The District-Wide Impact of Charter School Implementation

Austin Sell

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2022

Reading Committee:

Mark Long, Chair

Marieka Klawitter

David Layton

Program Authorized to Offer Degree:

Public Policy and Governance

University of Washington

Abstract

The District-Wide Impact of Charter School Implementation

Austin Sell

Chair of the Supervisory Committee:
Mark Long
Public Policy and Governance

In this dissertation I examine the impact of charter school implementation for public school districts which did not previously have any operating charter schools. In contrast with most of the previous literature, I examine the impact on district-wide outcomes, rather than focusing specifically on the impact for students of either charter schools or traditional public schools. I use an event study difference-in-differences methodology applied to U.S. national data to estimate the causal impact on these outcomes. In Chapter 2, I find small and statistically insignificant effects of charter school introduction on mean test scores (*ATT* = 0.0169, s.e. = 0.013) and increasing and statistically significant effects on achievement inequality (*ATT* = 0.0599, s.e. = 0.017). In Chapter 3, I find substantively and statistically significant effects of charter school introduction on both median household income (*ATT* = $1,377, s.e.=$555) and the proportion of district residents with at least a bachelor's degree (*ATT* = 0.022, s.e.=0.006). In Chapter 4, I find substantial and statistically significant effects of charter school introduction on district Average Freshmen Graduation Rates (ATT = 0.0597, s.e.= 0.0274). In each chapter, I discuss how policymakers should consider these findings in the context of weighing charter schools as a policy option.

ACKNOWLEDGEMENTS

I would like to extend my thanks to my dissertation committee Mark Long, Marieka Klawitter, and David Layton. I am deeply indebted in particular to Mark Long who served as the committee chair, and without whose guidance and support this dissertation would not exist. I am sincerely grateful to both Marieka Klawitter and David Layton for their support during my long duration at the Evans School.

I am also extremely appreciative of Meredith Honig, who went above and beyond as a GSR. Her substantive feedback during the prospectus defense helped me think about my research in context.

Comments received from Celeste Carruthers and others at APPAM 2021 were immensely helpful in formulating the second chapter.

I would also like to give a big thank you to the Evans School community who have helped and supported me through my time here, including Rebeca de Buen Kalman, Emmi Obara, Nicole Kovski, Ben Glasner, Kate Crosman, Adam Hayes, Vance Larsen. They and others never failed to lend an ear when I needed to bounce ideas off of somebody else.

Finally, thank you to Maura Rowell, who was the first person to model what finishing a dissertation looked like. Without her unwavering support, I am certain I would not have finished.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

DEDICATION

To my mom, whose love and support got me through the darkest days.

And to my dad, who more than anyone, pushed me to finish this dissertation.

I wish he could be here to see it.

EXECUTIVE SUMMARY

The nature of charter schools is that of change. One of the reasons they exist is to attempt to provide education in a way that is different than what has been done or what is possible in traditional public schools. That does not make them better or worse, but it means when a charter school comes along disruption to the status quo is sure to follow it.

That disruption is what is captured in the analysis presented here. Most of the research done on charter schools has focused on either how charter schools affect their own students or how they affect students enrolled in traditional public schools. From a policy perspective, I do not think that distinction is often very useful. Charter schools receive public funds just as traditional public schools do, making their existence a tradeoff of resources that could have been spent in other places. To really understand charter schools as a policy option, I believe that policymakers should think about their impact in the context of all public-school students. That framework is what drives the research in this dissertation.

Taking advantage of the variation in timing in which charter schools are opened in public school districts across the U.S., I employ an event study difference-in-differences methodology to answer questions about the impact that charter schools have on districts as a whole. This analysis was only recently made possible due to advances in the difference-in-differences literature by Callaway and Sant'Anna (2021). Unlike the standard difference-in-differences research design, which relies on observing treatment and comparison groups pre- and post-treatment occurring at a specific point in time, this more advanced method allows the treatment introduction to vary across time in multiple time periods. Consequentially, my methods are robust to large systemic shocks that could

have occurred simultaneously with treatment in a particular year and been confused with the treatment effect. Now such spurious effects would have to be timed exactly with charter introduction across geographic and temporal dimensions to become entangled with the treatment effect.

For example, the early 2000's brought sweeping accountability-centric changes to education. The No Child Left Behind Act created a focus on testing and assessment that continues to have a lasting impact on education today. This same era is where a significant portion of my data comes from. As a result, I do my best to use methods that do not capture widespread, multi-year changes that occurred in the same timeframe that many charter schools sprang into existence. Because the opening of a charter school is such a singular, definable event (i.e., I can identify precisely the years during which districts have or do not have charter schools), more gradual changes to school districts that occur over many years would be downplayed in the estimated results. Further, charter schools are introduced in my data at varying points in time. National changes in norms and practices would happen at a particular point in time; that could cause the impact of charters to be disentangled for one implementation cohort, but not for all of them. The true threat to internal validity would be if charter schools are consistently implemented as a part of a set of policy options in the same year across many different districts in the country.

In Chapter 1, I review the literature surrounding charter school and explain where my research fits in to our existing knowledge. In Chapter 2, I analyze the impact on district-wide average test scores and the inequality of those scores. I find that charter schools do not affect a district's average test scores, but that the inequality of those scores increases as a result of the introduction of charter schools. An average district increases the inequality of test scores by 0.06 standard deviations, implying that high-achieving students perform better, low-achieving students perform worse, or both. From a policy perspective, increasing inequality while simultaneously doing nothing to increase mean achievement is likely undesirable to most policymakers.

In Chapter 3, I examine the impact on median household income and the attainment of bachelor’s degrees for residents located in the geographic school district. I do not expect that charter schools can directly affect these measures in the short timeframe in which I examine them. However, a change in the composition of families in a district could potentially yield changes in these variables as well. If charter schools impact a school district by influencing the families that choose to live there, then that would be a strong statement to policymakers about what types of people are responding to the availability of charter schools. Indeed, I find that upon the introduction of a charter school to a school district, the average median household income increases by $1,377 and the percentage of residents with at least a bachelor’s degree increases by 2.2 percentage points. Both results are statistically significant.

In Chapter 4, I investigate the impact on a district’s graduation rates. Again, rather than looking at the graduation rate of specifically charter students or traditional public-school students, I use the graduation rates for the school district as a whole. The hypothetical mechanisms for increasing graduation rates are similar to how one might expect charter schools to increase a district’s average test scores. Competitive pressure from a well-performing charter school could also increase the quality of other nearby schools. It is interesting then that while I find no significant increase in mean test scores in Chapter 2, I do find substantively and statistically significant increases in graduation rates as a result of charter school introduction. Once a charter school begins operating in a school district, the average district sees an increase in graduate rates by nearly 6 percentage points.

Chapter 1. INTRODUCTION

1.1 HISTORICAL BACKGROUND

A charter school is a particular type of school in which some actor proposes the creation of a school that will be privately run, and if authorized, receives public money to do so. Minnesota passed the first charter school legislation in 1991, which was followed by the first charter school opening in the state in 1992 (Wohlstetter, Smith, and Farrell 2013). Since then, 45 states and the District of Columbia have decided to allow charter schools; Montana, Nebraska, North Dakota, South Dakota, and Vermont, do not currently allow charter schools (Wilkins 2020). While the process is different for every state, in general the state legislation sets up some entity to act as an authorizer of charter schools. This authorizer could be the state itself or could be some private, nonprofit organization. In many cases, local school boards are capable of authorizing charter schools. No state currently allows the authorizer to be a for-profit organization. The authorizer is responsible for evaluating proposals to open charter schools. They review the applications and, after some specified length of time, the authorizer will review the charter school's performance and decide if the school should continue to remain open.

As an example, the state of Washington currently has two charter authorities, the Washington State Charter School Commission and Spokane Public Schools. The first has the authority to authorize charter schools anywhere in the state, while the latter can only authorize charters inside the Spokane school district. Any school district can apply to be an approved charter authorizer, but to date only Spokane Public Schools has been approved by the State Board of Education. Additionally, Washington State has a complicated history with charter schools. The state rejected initiatives to legalize charter schools in 1996, 2000, and 2004. In 2012, Washington voters narrowly supported introducing charter schools into the state, and the first charter school opened a couple years later. Legal challenges ensued and the Washington Supreme Court declared that the use of general funds for charter schools was in violation of the state constitution. The following year, legislators passed a new law that defined the revenue stream for charter schools so that money would not come from the general fund. Legal challenges arose once more, but this time the Washington Supreme Court upheld the constitutionality of the law. As of the autumn of 2022, charter schools are still legal in Washington.

The number of charter schools in the United States has grown substantially, as well as the number of students that are being educated in them. There are now over 6,000 charter schools that educate over 2.5 million students (Berends 2015). In the last decade, there has been a trend of increasing charter schools that are a part of nonprofit Charter Management Organizations (CMOs). CMOs and their for-profit counterparts, Education Management Organizations (EMOs), are collections of charter schools. This is opposed to a singular, independent charter school. The vast majority of charter schools are not associated with a CMO or EMO, but the percentage of CMOs as a function of total number of charter schools is increasing (Berends 2015). Additionally, when scholars and policymakers consider the context of organizations in the charter school situation, it is often these CMOs that come to mind. Popular CMOs, such as KIPP (Knowledge is Power Program), are well described in both the academic literature and the media. While KIPP has certainly produced impressive results (Angrist et al. 2012; Henig 2008), it is important to remember that CMOs represent only about a quarter of all charter schools.

The idea of school choice originated with market theorists who believed that competition among schools was vital to increasing the quality of those schools. While the initial premise was that the government should issue vouchers to parents and empower them to make the decision on where their child is schooled (Friedman 1955), school choice advocates have also used these arguments in favor of charter schools. The autonomous nature of charter schools also creates the potential benefit of allowing new, improved approaches to organization and curriculum. Combined with competitive forces, this autonomy produces an incentive to innovate (Chubb and Moe 1990).

There are other theoretical lenses with which to view the charter school phenomena, such as sociological institutionalism. Viewed from such a framework, charter schools are another manifestation of the myths, norms, and rituals that bind organizations that operate in the educational sphere (Berends 2015). Despite the different formal constraints faced by charter schools, there exist the same informal constraints that apply to traditional public schools. As a result, charter schools are subject to isomorphic pressure that severely limits their ability to operate substantially differently than traditional public schools (DiMaggio and Powell 1983). Any innovation that too severely departs from the norm creates the impression that this is no longer a "real school" (Bidwell and Kasarda 1980; Tyack and Cuban 1995). Thus, charter schools do not significantly differ from the operation of traditional public schools, and we would not expect to observe any differences in the educational achievement of charter students.

However, the arguments that charter school advocates and opponents make often rely on market theory. Unlike sociological institutionalism, market theory focuses explicitly on the costs, benefits, and incentives for the various actors involved. Despite the norms that may shape the behavior of charter schools, the critical policy question is how charter schools may differ from traditional public schools in terms of the resources they require.

What makes charter schools unique in the context of other school choice alternatives, is the investment of resources it requires, and the necessary interplay among existing infrastructure. Unlike policies such as vouchers or open enrollment, which seek to increase choice among existing options, charter schools require the substantial investment of public resources to create a new educational option for parents. As a result, charter schools create a scenario in which traditional public and privately-run charter schools must compete for the same set of public resources.

1.2 Literature: Conceptualizing Outcomes and Costs

Much of the discussion surrounding charter schools as a policy option tends to focus on specific benefits or drawbacks of the implementation of such a policy. Numerous studies have examined how charter schools affect student outcomes, stratification, and choice. Likewise, studies have examined the necessary tradeoffs that come with implementing charter schools, both in terms of enrollment and financial resources. I will be examining these studies in more depth later in this paper.

It is rarer for the charter literature to take a more holistic approach in examining the benefits and costs of permitting charter schools. In order to make such a comparison, I conceptualize the important policy tradeoffs inherent in charter school implementation as described by the Education Production Function and the Education Cost Function. In many ways, these functions examine the same phenomenon from different angles. While policymakers often like to focus on the improvements to educational outcomes, the policy implementation to achieve those outcomes has a cost in terms of the public resources that were expended. Similarly, a given level of public investment is only able to improve student outcomes by so much. The decision to introduce or maintain charter schools requires looking at both the educational output and the cost it takes to reach that output.

1.2.1 *Education Production Function*

Like any production function, the Education Production Function models an output as the function of different inputs. This conceptualization of education thinks of schooling as a way to generate some level of a particular educational outcome, where policymakers can manipulate different inputs in order to affect the outcomes that they or their constituents care about. The Education Production Function is flexible in the sense that it can be applied to many different outputs. Common outputs that are measured in the literature include years of schooling (Psacharopoulos and Patrions 2004), college attendance (Place and Gleason 2019), graduation rates, and test scores.

Depending on the specific output, the inputs of a particular Education Production Function may vary, but the vast majority share some common characteristics. Educational outputs are typically modelled as cumulative. That is, a student's outputs in grade 4 necessarily builds upon the output from grade 3. As such, outputs at any point in time are more than simple snapshots, but rather a reflection of all the inputs for a student up to that point. Most inputs can be thought of as belonging to one of three categories: the student's individual characteristics, the student's home environment, and the factors that a student encounters while at school. Policymakers often direct their attention to the latter because it is most directly influenced by education policy. Inputs of this variety include teacher characteristics, class sizes, peer effects, and per pupil expenditures.

The Education Production Function informs policymakers how they should expect outputs to be affected as a result of different policy levers. It is through this lens that I examine the literature for evidence of how the implementation of charter schools affects student and district outputs and outcomes. Charter schools affect more than a single input and can be implemented in different ways. As such, it is important to parse out the specific relationships between different inputs affected by charter schools and the outputs presented in the literature.

1.2.2 *Education Cost Function*

Equally as important as the Education Production Function is the Education Cost Function, which returns the cost associated with a certain level of educational achievement given the input prices, such as teacher salary. In practice, this allows policymakers and researchers to model the expenditures required to achieve some adequate level of educational outcome, given the input prices (Imazeki and Reschovsky 2006).

Unlike the Education Production Function, the Education Cost Function does not have a wide set of possible outputs. Cost is almost always measured in the currency of the economy of the schooling environment (e.g., US Dollars). However, there is a notable conceptual difference between observable expenditures and the concept of cost. As Costrell et al. remark, the assumption behind the conceptualization of cost is that cost is at a minimum efficient level (2008). For example, to say that it costs $15 million per year to operate a particular school is to imply that $15 million is the lowest possible amount spent that yields a given level of educational achievement. On the other hand, expenditures refer to the money and resources that are actually used. Continuing the example, perhaps that school’s expenditures are actually $17 million, meaning that the actual amount of money spent was higher than the minimum amount needed to return the given level of educational achievement. While related, these are distinct concepts that are connected through efficiency. Observed expenditures can be used to estimate the relationship between cost and efficiency (Buerger and Bifulco 2019). In this case, the school has some level of inefficiency as it has spent $17 million to achieve a result that should only take $15 million. That is, the school has not actually reached a minimum efficient cost. As the Education Cost Function assumes the minimum efficient cost as its output, it is necessary to use an understanding of efficiency to relate the model to observable data.

A requirement for the Education Cost Function is a particular level of educational achievement (e.g., a certain test score average, graduation rate, etc.) Additionally, the cost function must include the input prices for that output. Faculty salaries, building costs, and material expenses are all typical input prices that are considered. Lastly, it is common to consider variables that may affect the way financial resources impact the school, for example enrollment or demographic information about the composition of the school or district (Duncombe and Yinger 2011).

As a lens for examining the literature, the Education Cost Function serves a valuable purpose in helping to understand how policymakers view educational interventions. As an investment of public resources, charter schools are necessarily constrained in the expenditures that are available to them. The decision to provide money to charter schools is one that is likely informed by the educational goals of policymakers, while simultaneously balancing the potential financial impact on other parts of the schooling system. To fully understand how charter schools operate as a policy option, the literature must be examined with a specific financial and cost-centric focus, especially as it pertains to the tradeoffs between charter schools and public schools. Additionally, reviewing the literature through this lens forces the scope of the discussion to a higher level. Schools do not typically make financial decisions independently, so the focus on the components of the Education

Cost Function will help to elucidate the ways in which resource-driven decisions at multiple levels of governance are interconnected.

1.3 LITERATURE: DIRECT IMPACT ON STUDENT OUTCOMES

1.3.1 *Impact on Charter Students*

The fundamental problem of any causal estimation is discovering how to find a counterfactual. This process is complicated by the heavy amount of self-selection and sorting that occurs in regard to school choice. If a student chooses to enroll in a charter school, it seems problematic to compare them to a student who did not choose to enroll in a charter school. Even controlling for every observable characteristic may miss some unobservable characteristic that was different between those two students. If that is the case, any causal estimate we produce by comparing those two students is likely to be biased. Fortunately, the empirical literature has discovered ways around this problem, although they all have specific drawbacks.

1.3.1.1 Random Assignment Lottery

One specific methodology exploits the fact that when a charter school has too many applicants for its capacity, it must choose which students to admit based on a lottery. The introduction of randomization allows researchers to create a counterfactual and establish a valid causal inference about the effect of charter schools (Angrist, Pathak, and Walters 2013; Angrist et al. 2016; Abdulkadiroğlu et al. 2011). Scholars examine the data of students who chose to apply for admission into a charter school, comparing those who were randomly selected for enrollment to those who were not. In theory, the only difference between these two groups of students is whether they attended the charter school, which was randomly assigned to them. Thus any differences observed in the dependent variable can be attributed to the charter school.

There are some drawbacks to this approach for estimating causal effects. First and foremost, it limits the scope of students to which the resulting estimates could be generalized. In order to exploit the randomization of the lottery, researchers only have the ability to study students who self-selected into applying for charter school admission. This identification strategy becomes problematic if the students who apply to charter schools are fundamentally different from the students who do not apply to charter schools. For example, if students who apply to charter schools are more highly motivated than students who do not, then the causal impacts estimated by exploiting the lottery mechanism are only valid for that subset of highly motivated students. We would be unable to generalize to the group of less motivated students (i.e., those who did not apply to the charter school). Limiting analysis only to this group means that using research of this type to advocate for broadly expanding the availability of charter schools has the potential to be misguided.

Another limitation is that this method is only applicable to a specific, likely unrepresentative, subset of charter schools. Schools are only required to hold an admission lottery if the number of applications exceeds the capacity of the school. It is plausible that these schools are different from the entire population of charter schools for a couple of reasons. First, by definition these are high demand schools. For whatever reason, parents have decided that these charter schools are a good fit for their children. It seems entirely reasonable that parents would be drawn to higher performing schools, which would mean that researchers are inadvertently selecting the best schools to study. Second, the range of schools that are potentially oversubscribed is not equal across different geographies. Schools in dense, urban areas simply have more potential students, making it easier to hit capacity. Rural schools may struggle to attract enough students to necessitate holding a lottery. The idea that school choice policies are more difficult to implement in rural areas is supported by the literature as the majority of studies utilizing the randomized lottery approach have looked at schools in urban areas (Betts and Tang 2016). If there is a fundamental difference in rural and urban students, then the results of the literature cannot be extrapolated to rural areas.

Finally, this method does not necessarily allow us to develop a deeper understanding of the mediators involved. Any observed increase in achievement is not because the student was at a charter school *per se*, but rather the specific things those charter schools did differently than public schools. These mediators can include more instructional time, different hiring practices, focused leadership, or something else entirely. The way that scholars tend to estimate the impact of charter schools ignores these potential mediators. Instead, the interpretation of the results largely rests upon the coefficient for the variable indicating whether the student attended the charter school or not. While this does give us information about charter school effectiveness, it provides a limited understanding of what is truly happening to produce those results. It should be noted that this is not a necessary component of the randomized lottery methodology; scholars could include mediators into their regression frameworks, although including such mediators may introduce bias into the model, as those mediators are not typically randomized. In practice the emphasis does not appear to be about developing this closer understanding of what is happening inside the black box (Berends 2015).

1.3.1.2 Panel Models

Another popular method used in the literature to deal with the selection bias problem is the use of student-level panel models (Booker et al. 2008; Bifulco and Ladd 2006). Using panel data to track the same students over time is a different way of constructing a counterfactual to estimate a causal impact. Instead of comparing two groups of students, each student is essentially compared to her past self. If other time-varying factors are controlled, this estimation technique should result in valid causal inferences. Panel models deal with the selection bias problem by using student fixed effects, which captures the unobserved variables for any given student. Thus, it does not matter if there are unobserved differences in motivation, as this will all be accounted for with the fixed effects. Notably, the fixed effects do not allow us to look at any specific effects of motivation or other unobserved time-invariant variables; it simply allows the estimation of the effect of other variables to remain unbiased.

Of course this technique does have limitations. As alluded to previously, collecting panel data is very resource intensive. By its very nature, panel data takes years to collect as researchers have to actually wait for the next stage of the process to happen. If it is decided that a group of students are going to be tracked for five years, then collecting that data has to take at least five years. In practice, researchers often use panel data that some other entity decided to collect at some point in the past. While that means the researcher does not have to wait for the data to develop, it does of course require that someone out there has already decided to collect the relevant data. This may introduce additional complications since the scholars conducting the analysis do not have input on how the data is collected and measured. In fact, it is often the case that education administrators or state agencies are the actors that required the data to be recorded. These actors have different motivations than education scholars. Scholars are explicitly motivated by uncovering causal impacts, while this is not necessarily the case for those collecting the data. As a result, researchers may end up with a dataset that suffers from various problems. For example, limiting attrition from the sample may not be a top priority for administrators collecting the data, but it is of vital importance for scholars looking to produce unbiased estimates. Additionally, such data is rarely collected for the purpose of scholarly analysis, making it likely that key variables may have not been collected.

Panel models also make some assumptions that may or may not be realistic. First, in order for the estimated effect to be interpreted as causal, researchers have to incorporate all relevant variables into the model. This problem, omitted variable bias, is not unique to panel models, but it does present some challenges unique to this context. Since this technique essentially allows for the comparison of a specific individual with her past self, it is important to include any potential variables that could influence a change in achievement from the past to the present. More importantly, fixed effects models assume that unobserved characteristics remain constant over time. For example, we would assume that a specific student's motivation remains unchanged over the time that the data is collected. If that assumption holds, the model is successful at eliminating the selection bias problem previously discussed. Unfortunately, there is no way to test whether this assumption is true or not given that it is an assumption about unobservable traits. In fact, the reality that the student has made the decision to enter a charter school may be correlated with other events in their life that may have prompted the decision, while also having a direct effect on their outcomes.

There are specific kinds of questions that panel models are more apt to answer than others. Because of the structure of the data, it does not make sense to use panel data to estimate the impact on outcomes that occur only once (Booker et al. 2011). Rather, it necessitates that outcomes are regularly measured. In the context of education and charter schools, this means that panel models are well equipped to estimate the impact of charter schools on outcomes such as test scores and attendance, but are unhelpful when trying to look at outcomes that occur only once, such as graduation or college enrollment.

Finally, the main problem with the fixed effects approach is that it relies on the fact that some individuals will attend both charter schools and public schools over the course of the data collection period. This is a very specific subset of students that may or may not be representative of the entire student population. Additionally, we know from the literature on school closures that there is the potential for substantial harm to a student's achievement when she switches schools (Engberg et al. 2012).

1.3.1.3 Matching Techniques

The final methodological approach to dealing with the selection bias issue that I will discuss is the use of matching. There are a variety of different matching techniques, but the central idea of all of them is that students in charter schools are statistically matched with students in traditional public schools based on a long list of observable characteristics. There are various ways to try to increase the quality of the match, including using variables such as test scores prior to the choice to enroll in charter schools (Booker et al. 2011), but all matching methods are only able to compare the observable traits a student has displayed. For this method to yield an accurate estimate of the treatment effect, it must be the case that there are two students who look identical in every way, with the only difference being whether they chose to attend a charter school or not. Then any difference in achievement should be attributed to the attendance at the charter school.

There are several complications involved with matching methods. First, the researcher has to decide which matching strategy to employ. Between exact matching, propensity score matching, nearest neighbor matching, or a unique matching algorithm, the researcher has to carefully consider which choice is right for their study knowing that the different options have implications for the ability to make causal inferences (Dehejia and Wahba 2002). The unfortunate reality is that these methods rely on untestable assumptions that may not be true of the data that is available. There is also the problem of acquiring enough data to make high quality matches. In order for this technique to work, there needs to be a large number of students that did not attend charter schools and have sufficient data available to be matched to charter school students. Acquiring these data is certainly not a trivial task. The most severe problem with matching is the potential for exacerbating bias related to unobservable characteristics. Consider again the case of two nearly identical students; the only observable difference being that one chose to attend a charter school while the other did not. The assumption that the matching technique makes is that the choice to attend the charter school after accounting for all of the observable traits is essentially random. However, it seems far more likely that there is some unobservable difference between these students, such as motivation, that is driving the choice to enroll in a charter school. If that is the case, then any matching technique will magnify the relative importance of the difference between the unobservable characteristics leading to biased estimates.

All of the methodologies discussed here have significant drawbacks in how they allow us to think about the findings they produce, but they are currently the best practices in the field. In the aggregate, this variety of techniques can be beneficial at addressing the limitations of another method. For example, one of the key weaknesses of lottery assignment is the restriction about which charter schools can be examined. Matching techniques have no such limitation so they can be more broadly applied to all charter schools.

1.3.2 *Key Findings of the Empirical Literature*

For such a large literature with numerous rigorous studies, it is surprising that there is not a strong consensus on what actually is the effect of charter schools. Some studies have found extremely positive results, while others find mixed or even negative impacts. One of the most consistently positive areas of charter schools is KIPP and other No Excuses charter schools (Angrist et al. 2012; Henig 2008; Nichols-Barrer et al. 2016; Cheng et al. 2015). A meta-analysis of No Excuses charter schools found that on average, students who attended such schools saw significant increases of 0.16 standard deviations in English and 0.25 standard deviations in math test scores (Cheng et al. 2015). The meta-analysis did include lottery admissions studies which potentially limits the scope to which we can generalize these findings; the individual studies in the meta-analysis may have been biased by selecting only oversubscribed schools. Additionally, the authors failed to explicitly address the problems of attrition in No Excuses charter schools, which is particularly problematic as these types of schools typically suffer from the highest levels of attrition (Powers n.d.). As the name implies, No Excuses charter schools have a low tolerance for students that do not meet high academic and disciplinary standards. As a result, the perceived achievement levels of the school may be inflated by the attrition of lower-performing students. It should be noted that the evidence suggests that the large upwards bias from attrition is unlikely to overwhelm all of the gains made by the school, based on how many students actually left (Henig 2008). KIPP schools specifically have been the subject of many studies and they have been very positive. Angrist et al. (2012) examine a particular KIPP school and find that disadvantaged students are likely to see the most gains. They conclude that the KIPP approach could be used to help mitigate the achievement gap. Other studies have also found positive results for KIPP schools, especially in terms of mathematics achievement (Nichols-Barrer et al. 2016; Henig 2008).

More generally, studies that use the lottery assignment design tend to find positive gains in test scores (Abdulkadiroğlu et al. 2011; Angrist, Pathak, and Walters 2013; Angrist et al. 2016; Hoxby and Murarka 2009; W. Dobbie and Fryer 2013) and long term job outcomes (W. S. Dobbie and Fryer 2016). However, a national study that looked at charter schools in several states found mixed results in different contexts (Gleason et al. 2010). Non-experimental methods, such as panel designs and matching, have discovered more mixed findings (Booker et al. 2008; Bifulco and Ladd 2006; Hanushek et al. 2007; Sass 2006; Buddin and Zimmer 2005), with some estimating positive effects while others remain inconclusive. A meta-analysis of studies that examine the effects of charters schools on reading and math scores found that overall, there are small gains for math scores in elementary and middle school (Betts and Tang 2016). However, high school math scores and reading scores were not statistically significant.

In general, the type of methodology used to study charter schools is very important in how we interpret the results because of the specific limitations of each. We can see in the literature that the likelihood of a strong positive effect is at least correlated with the research design employed, specifically lottery assignment. This understanding is important for the way that we should interpret the literature as a whole. While simply aggregating effect sizes may not show particularly impressive results, the consistent positive findings of these studies with high internal validity but low external validity indicate that instead of showing biased results, they are likely demonstrating the effects of the top tier of charter schools.

1.3.3 *Impact on Traditional Students*

While identifying the causal impact of charter schools on their own students is difficult, there at least exists a substantial portion of the literature that uses quasi-experimental methods. As discussed above, this is largely due to the frequent randomness that is baked into the selection process via enrollment lotteries. That randomness can be exploited to generate a causal interpretation. Unfortunately, such randomness is generally not found when it comes to estimating the impact of charter schools on district students. Charter schools do not locate randomly. As such, the evidence for and against competition effects is far more varied in methodology; scholars are unable to consistently use the same statistical methods to uncover causal estimates.

The best studies that examine the impact on traditional public school students have done so using student-level longitudinal data, which has many of the same drawbacks of the same type of data when estimating the effect of charter schools on their own students. Models that rely on this type of data need variation in students' attendance of charter versus traditional public schools to properly estimate effect sizes. One source of this variation could be following students over time who move in to or out of districts where charter schools are located nearby. While these types of students are included in some of the data that is used by the literature, I have not read any study that made this source of variation the primary focus. Instead, the longitudinal studies that are in the literature find variation in the exposure to charter schools over time, typically by examining data in districts before and after charter schools were introduced. These types of models work best when unmeasured characteristics of students, schools, and districts are unchanged over time, thus making the source of potentially changing outcomes the exposure to charter schools or some other measurable variable.

Compared to the multitude of studies examining the direct effect of charter schools on students, there is far less research done on the indirect effects (Gill and Booker 2015). The methodological challenges discussed pose some serious obstacles for quality research looking to examine the indirect effects. However, this has done little to slow both advocates and opponents of charter schools from making strong claims about how charter schools impact the achievement of traditional public-school students.

Proponents of charter schools argue that the presence of such schools indirectly benefits public school students in two ways. First, charter schools compete with public schools for the same students (or rather the same geographic set of students). Because public schools have funding that is directly tied to the number of students enrolled, they have a strong incentive to keep as many students enrolled as possible. Students that decide to enroll in a charter school represent a loss of funding revenue. Thus, public schools will do everything they reasonably can to improve their quality in order to attract and retain students. Notably, this effect may be heterogeneous across students as schools attempt to specifically attract those that are the least costly to educate.

The second way that charter schools allegedly benefit public school students is by allowing space for educators to implement innovative pedagogical practices. Charter schools are not restricted to the same curriculum and practices to which public schools abide. As such, they offer an opportunity for teachers and administrators to push the boundaries of best practices in education. Most of the literature examines the competitive effects of charter schools rather than the implementation of innovative techniques that is posited by charter school advocates. It is possible that this mechanism is in fact occurring, but even evaluating specific techniques for effectiveness presents substantial hurdles, much less examining the magnitude of their spread to public school students (Gleason 2019). While there are certainly some charter schools that have had consistent success in promoting high levels of student achievement (e.g., KIPP), it is exceedingly difficult to ascertain which of the policies and techniques these schools employ is responsible for their success and whether these lessons would actually translate to a traditional public school classroom.

Opponents of charter schools argue that traditional public-school students are harmed when charter schools open nearby. As explained above, traditional public schools and charter schools are competing for the same set of students to acquire access to the same source of revenue. This competition makes it likely that public schools will receive less funding, thus creating a financial drain that ultimately harms students. The argument that competition will increase the efficiency of public schools relies on the assumption that there already exists some level of inefficiency that needs to be mitigated; charter school opponents are typically unwilling to make this assumption, leading them to believe that competing over scarce funding resources inevitably harms public schools and their students.

The evidence in the literature examining the impact and direction of these alleged competitive effects is mixed, ranging from small and positive (competition with charter schools improves traditional public student outcomes) to slightly negative, although negative effects are rare (Ladd 2019). Studies have used data at the district, city, state, and national level in an attempt to answer this question, with some studies having mixed results even in the same geographic area. For example, in North Carolina, Bifulco & Ladd (2005) are unable to find evidence of positive competitive effects, whereas Jinnai (2013) did find consistent positive effects. A similar pattern is evident in New York City, with one study finding positive effects (Cordes 2018) in contrast to an earlier study where the small positive effects were dependent on the model specification (Winters 2012). Across the literature is seems difficult to back up the claim that there are strong, consistent, positive effects on traditional public school students as charter school proponents often argue. With so few studies spread across so many states and time periods, the detection of a treatment effect seems dependent on other factors such as the inclusion of particular variables in the model and how long the charter schools in the data have been operational.

One of the only studies that demonstrates negative competitive effects on public school students is also one of the only studies to explicitly address the problem of nonrandom charter school location (Imberman 2011), using an instrumental variable strategy to account for the introduced endogeneity. The analysis finds significant drops in math and language test scores, in contrast to much of the literature that uses longitudinal models. I am unaware of any study that attempts to replicate this technique. It is also worth noting that this particular study is one of the few that uses a single school district as its focus, rather than a city or state. In many ways this study is an outlier in the extant literature.

A sizeable portion of the literature finds no evidence of a competitive effect whatsoever. A geographically comprehensive study, which separately examined competitive effects for charter schools in Chicago, Denver, Milwaukee, Philadelphia, San Diego, Ohio, and Texas, found precisely estimated null effects for competitive effects on both the reading and math scores of traditional public-school students in every region except for Texas (Zimmer et al. 2009). Notably, the positive results found in Texas echo the results found a year earlier in a study that focused specifically on Texas (Booker et al. 2008). The only study that examines this question with a national sample also yields inconclusive results, although they are not estimated with the degree of precision as some of the other null results (Davis 2013).

1.4 Literature: Impact on Schools and Districts

Most of the research done on the introduction and implementation of charter schools into a district has focused on the direct impact on students. However, charter schools also create systemic impacts necessarily because they use the same resources as public schools, specifically government funding. Because there is only so much funding that the government will put toward K-12 education, charter schools disrupt the status quo by taking some of that funding which alters the funding streams for traditional public schools. The extent to which this happens depends on the funding arrangement for charter schools in a particular state. All states allocate funding to charter schools based on the number of students they enroll, but how much money is allocated per-pupil and the formula by which that is assessed varies by state (Shen and Berger 2011).

Charter schools and traditional public schools compete over the same set of students once the charter schools are open and operating. On some level, this competition is a zero-sum game where every student that a charter school enrolls is a source of funding that is no longer available to a traditional public school. However, the characterization of conflict between traditional public schools and charter schools misses the larger picture; traditional public schools and charter schools are often both accountable to the same school district. In most states, public school districts are a charter authority. That is, the charter school is opened by and ultimately accountable to the public school district. Thus, the decision to allow charter schools to operate is not potentially lost funding for public school districts because both traditional and charter schools are getting their funding through the same district. The key question is whether charter schools change the allocation of resources in a way that is less efficient for the entire system. To fully understand the systemic impact that charter schools have, it is necessary to first look at how charter schools themselves use that funding and then to examine how the presence of charter schools changes the financial situation for traditional public schools.

1.4.1 *Expenditures and Costs of Charter Schools*

Once approved for operation, charter schools receive funding in much the same way as traditional public schools. The amount received is tied to student enrollment, but the specific amount of per-pupil funding is often less for charter schools (Baker 2014). These policies are created at the state level so naturally there is variation across the country, but most charter schools receive less funding for an identical number of enrolled students because lawmakers see charter schools as having different expenditure responsibilities than traditional public schools.

Many state and federal regulations of traditional public schools require expenditures, some of which are explicitly funded while others are unfunded mandates. It is common for charter schools to be exempt from many of these state regulations. For example, traditional public schools are required to accommodate students with different types of disabilities pursuant to federal statute. Charter schools are also required to accommodate any enrolled students with disabilities or special needs, and are not allowed to turn away such students for reasons other than enrollment capacity. However, the structure of a particular charter school is important to how this policy is enforced. In some situations, the onus of holding the charter school accountable is on the school district with which the charter school has partnered. This can open up resources for the charter school that does not take away from their typical expenditures. In the event that a charter school is operating independently, the charter school itself is responsible for ensuring that all students have appropriate accommodations. The charter school's structural relationships can drastically change its responsibilities in terms of financial expenditures.

It is less important how much money charter schools are given and more important how efficiently they spend that money. Proponents of charter schools often claim that these schools can use resources more efficiently and in innovative ways because they are not constrained is many of the ways that traditional public schools are. If this were the case, we should be able to observe two potential outcomes. First, per the Education Cost Function, for charter schools and traditional public schools with similar input prices, we should observe a higher level of student output for the same level of expenditure. Second, again for similar input prices, we should observe an equivalent level of student output for a lower level of charter school expenditure. The difference in the output per expenditure level would be attributed to an efficiency factor, demonstrating that charter schools are better able to maximize the student output for a given dollar.

One important aspect to note is that this efficiency comparison relies on similar input pricing, which may not necessarily be true for charter schools. Looking at one example input of teachers' salaries, charter schools have significantly more flexibility in how they compensate teachers than do traditional public schools. Charter school teachers are not always unionized as is typically the case for public schools. These factors make it unlikely that charter schools see input pricing in the same way. To my knowledge, no analysis has been published that specifically looks to quantify the impact of this phenomenon.

There have been studies that seek to understand if there is an efficiency advantage for charter schools compared to traditional public schools. In Texas, charter schools were examined over a 5-year period and found that charter schools do in fact operate more efficiently (Gronberg et al. 2012). Specifically, they found that charter schools achieved similar levels of educational output at a lower cost. This performance was largely attributed to the charter schools' freedom in facing fewer regulations.

Another study in Texas used a different methodology to estimate the efficiency difference between charter schools and traditional public schools. By measuring student output in terms of value-added metrics and looking at the input distance function as a measure of technical efficiency, the authors found that charter schools were able to make more efficient use of their financial resources and improve student outcomes (Grosskopf et al. 2009).

1.4.2 *Expenditures and Costs of Traditional Public Schools*

Examining whether charter schools are more efficient is only one part of the systemic impact that charter schools have. Because both charters and traditional public schools are competing over the same pot of government funding, opponents of charter schools warn of the danger of charters draining resources from public schools (Arsen et al. 1999). Even if charter schools use the money more efficiently, that could still come at a cost of further constrained resources and lower educational output for the students in traditional public schools. Advocates of charter schools argue that this competition increases the efficiency of traditional public schools as well (Hoxby 2003), forcing them to try to do more with less.

These competing forces could both be in effect, but the real question is what the net effect of charter schools is on the financial state and performance of nearby public schools. Scholars have found a consistent negative financial impact (Arsen and Ni 2012, Bifulco and Reback 2014, Buerger and Bifulco 2019, Ladd and Singleton 2020). These findings are especially true for schools and districts that are already facing budgetary pressures. (Arsen and DeLuca 2016).

One of the biggest reasons that traditional public schools see a change in their costs is due to the responding change in student composition. The types of students that are enrolled are a crucial input to the Education Cost Function because every student does not have identical costs associated with their education. As discussed above, special education students often have accommodations that are required which necessitates additional expenditures. Likewise, low-income students and English Language Learners are associated with higher costs of education on average (Jimenez-Castellanos and Topper 2012).

When a charter school opens nearby, the composition of these types of students changes at the public school. There are mixed results in the literature for the change in enrollment of low-income students, with some studies finding increased percentages staying at traditional public schools (Abdulkadiroglu et al. 2009, Carnoy et al. 2005, Epple et al. 2016), while other studies have found lower percentages stay enrolled (Booker et al. 2008, Finnigan et al. 2004, Hoxby and Murarka 2009). However, the literature shows consistently that percentages of both special education students (Bifulco and Buerger 2015, Abdulkadiroglu et al. 2009, Mommandi and Welner 2018) and English Language Learners (Bifulco and Buerger 2015, Chingos and West 2015) increase after charter schools begin operating nearby.

These studies show that charter schools are not randomly enrolling students from existing public schools. That is to say, there is a selection process for which students decide to apply for and enroll in a charter school. Some types of students are less likely to transition to a charter school, meaning that traditional public schools will see a disproportionate share of their enrollments coming from those students. Arguably, public schools are faced with higher financial costs from educating these students, creating a ceiling on how efficient they can be with the limited resources they have. There is some evidence that efficiency gains from charters could offset financial losses to publics in the long run (Buerger and Bifulco 2019) but some are skeptical that non-transitory costs can be offset (Ladd 2019).

1.5 Questions This Dissertation Answers

Despite nearly three decades of research on charter schools, the literature still lacks a definitive answer on how charter schools should be applied as a policy tool. It is evident that there are clear benefits that charter schools offer in at least some situations; there is plenty of analysis showing improvements in educational achievement for students who attend charter schools, even if the results are not ubiquitous. Even students in traditional public schools seem to receive some residual benefits from competition.

What is less clear is whether this policy tool is an efficient use of public resources. That question can only be answered with a fuller understanding of how the introduction of charter schools changes the circumstances surrounding school districts. While some parts of the literature have examined the financial burden that charter schools can impose on traditional public schools via the siphoning of students (Buerger and Bifulco 2019), there is very little work that explicitly examines the district-wide changes that occur when a charter school is introduced. A school district wondering whether it should authorize a charter school would benefit from that type of analysis.

In the next three chapters, I conduct such an analysis to better understand the net impact of charter school introduction. In Chapter 2, I analyze the impact on district-wide average test scores and the inequality of those scores. The introduction of charter schools could theoretically increase competitive pressure on traditional public schools, leading to increased average achievement.

Conversely, charter schools could create a financial drain on charter schools that decrease their performance. Either mechanism could allow for the gap between the highest and lowest achieving students to increase or decrease. In Chapter 3, I examine the impact on median household income and the attainment of bachelor's degrees for residents located in the geographic school district. I do not expect that charter schools can directly affect these measures in the short timeframe in which I examine them. However, a change in the composition of families in a district could potentially yield changes in these variables as well. If charter schools impact a school district by influencing the families that choose to live there, then that would be a strong statement to policymakers about what types of people are responding to the availability of charter schools. In Chapter 4, I investigate the impact on a district's graduation rates. Again, rather than looking at the graduation rate of specifically charter students or traditional public-school students, I use the graduation rates for the school district as a whole. The hypothetical mechanisms for increasing graduation rates are similar to how one might expect charter schools to increase a district's average test scores. Competitive pressure from a well-performing charter school could also increase the quality of other nearby school

Chapter 2. DOES THE INTRODUCTION OF CHARTER SCHOOLS LEAD TO WIDENING INEQUALITY IN ACHIEVEMENT WITHIN SCHOOL DISTRICTS?

2.1 ABSTRACT

Using a matched event study difference-in-differences analysis applied to U.S. national data from 2008-18 on district-level means and standard deviations of reading and math test scores in grades 4-8, I find small and statistically insignificant effects of charter school introduction on mean test scores, with point estimates for averaged effects generally in the range of 0.0 to 0.05 s.d. I find that test score inequality at the district level rises for the first six years after authorization of charter schools within the district. These increasing effects on achievement inequality are statistically significant in the 4th, 5th, and 6th years after authorization and peak at a 0.15 s.d. increase in the district-wide standard deviation of test scores. I discuss how policymakers should consider the tradeoff between (perhaps) modestly increased mean achievement with a concurrent increase in achievement inequality.

2.2 BACKGROUND AND THEORY

Prior research on charter schools has focused on evaluating the effects of charter schools on their own enrolled students (e.g., Angrist et al. 2016; Betts and Tang 2016; Epple, Romano, and Zimmer 2016) or on nearby traditional public-school students (e.g., Gill 2016; Bifulco and Ladd 2006; Imberman 2011). It is rarer to take an approach that examines the impact on the district as a whole, and such studies are typically focused on school finance (Buerger and Bifulco 2019; Ladd 2019). Yet, in considering whether to authorize a charter school, which is often a decision allocated by the states to school districts, the district needs to consider the costs and benefits of introducing charter schools for both future charter school students and their traditional public-school students simultaneously. That is, the school district needs to answer questions about how charter schools will change the nature of the district itself.

Equity has been a long-standing goal of policy that is often cited by policymakers, particularly in the realm of education and school choice. Education has a rich history as an avenue for increasing equality in society going back as far as *Brown v. Board of Education*, and equity continues to be a critical goal today. Research has investigated how charter schools and other school choice policies affect degrees of segregation among students (Garcia 2008). It is equally important to understand how equality in achievement is affected by these policies. While inequality, variation, and heterogeneity are largely empirical descriptors of how outcomes could be distributed, they are distinct but related to the equity conversation. Equity requires that all students are given the resources to succeed respective of their specific needs. Equality, in reference to outcomes, is examining whether students are actually achieving at the same level. A policy intervention that has an unequal impact on outcomes could be indicative that it creates inequitable opportunities for students. Many education studies focus on the impact on the average student. However, without understanding the underlying distributional impact, researchers may be missing an important part of the picture. Examining the level of inequality in achievement is a crucial piece of the puzzle in terms of understanding how charter schools work as policy tools and for whom they work.

I seek to answer the following question: Does the introduction of charter schools lead to widening inequality in achievement within school districts that authorize these charter schools? I take advantage of the variation in timing of charter school openings to estimate the causal impact of these openings. I estimate difference-in-differences models applied to data from the Stanford Education Data Archive, capturing the within-district standard deviation of test scores as my measure of inequality. I take care in this analysis to consider the existence of pre-policy disparate trends and take advantage of recent innovations in difference-in-differences modeling approaches in cases of multiple treatments implemented at different times, focusing specifically on the approach of Callaway and Sant'Anna (2021).

There are reasons to expect that charter schools might cause an impact on achievement inequality. Charter school education methods may yield more (or less) heterogeneity in student learning than traditional public schools. The very nature of charter schools allow them to explore alternative educational strategies that traditional public schools may find difficult to employ (Dressler 2001). In theory, charter schools are held accountable for their results, thus it follows that they would focus on results-driven approaches. For example, many charter schools utilize a "no excuses" methodology to education that would be unlikely to be possible in traditional public schools (Angrist, Pathak, and Walters 2013; Cheng et al. 2015). If charter schools create a higher achievement floor across a school district, but do not significantly raise the achievement ceiling, then inequality would be measurably reduced. In contrast, if charter schools raise the achievement ceiling, while not impacting the achievement floor, the inequality would be measurably increased. Notably, both of these options would increase the mean achievement in the district, but they would have markedly different implications for policymakers wishing to consider equality of outcomes. Further, student selection into charters and traditional public schools might generate peer effects in these schools that amplify pre-existing differences in achievement, widening existing inequality.

I use U.S. national data from 2008-18 on district-level means and standard deviations of reading and math test scores in grades 4-8, as described in the next section. Using a matched event study difference-in-differences analysis, I find small and statistically insignificant effects of charter school introduction on mean test scores, with point estimates for averaged effects generally in the range of 0.0 to 0.05 s.d. I find that test score inequality at the district level rises for the first six years after authorization of charter schools within the district. These increases in achievement inequality are statistically significant in the 4th, 5th, and 6th years after authorization and peak at a 0.15 s.d. increase in the district-wide standard deviation of test scores. In years 7-9, the positive effects on inequality drop off and are statistically insignificant in each year and, in fact, slightly negative in the 9th year (yet, imprecisely estimated). In the conclusion, I discuss how policymakers should consider the tradeoff between (perhaps) modestly increased mean achievement with a concurrent increase in achievement inequality.

2.3 METHODS

There have been recent innovations in estimating difference-in-differences models in the context where the treatment has been introduced in multiple different areas at different times and with heterogenous and dynamic treatment effects (de Chaisemartin and D'Haultfœuille 2020; Goodman-Bacon 2018; Callaway and Sant'Anna 2021). I use the approach of Callaway and Sant'Anna (CS, 2021) to estimate the causal effect of charter school authorization on subsequent levels of district-mean achievement and district-standard deviation of achievement. The CS model can be estimated assuming the following conditions hold:

a) Irreversibility of treatment (i.e., "once a unit becomes treated, that unit will remain treated" p. 203). In my context, this is almost true. From the data discussed in the next section, I drop 193 districts that had charters in operation in a year between 2009 and 2017 and then had no charters in operation in a subsequent year, through 2018 (1.8% of the full sample of districts).

b) Access to panel data. The data used form a nearly balanced panel. I implement the Callaway and Sant'Anna method by using their "did" package for R, available at: https://bcallaway11.github.io/did/. This command allows treatment effects to be estimated with unbalanced panel data.

c) Limited treatment anticipation. This would mean no changes in district outcomes in advance of charter school operation that is due to the anticipation of charter schools that are soon to be available. While not directly testable, examination of the periods before charter schools begin operating in a school district can shed light on the plausibility of this assumption. Without treatment anticipation, there should be no detectable effect on the eventual treatment in periods before the treatment has actually occurred. Evidence of pre-intervention trends would raise doubts that this assumption holds. As discussed in the results section, there is little evidence of any treatment anticipation.

d) Conditional parallel trends based on a "never-treated" group. That is, "conditional on covariates, the average outcomes for the group first treated in period g and for the 'never-treated' group would have followed parallel paths in the absence of treatment" (p. 204). CS note that this assumption may be favored by practitioners when "there is a sizeable group of units that do not participate in the treatment in any period, and, at the same time, these units are similar enough to the 'eventually treated' units" (p. 205). In the context here, I have a very large sample of never-treated districts and, as described below, rather than condition on covariates, I pare the list of never-treated districts down to those that "match" treated districts on pre-policy outcome levels.

e) Overlap, i.e., "a positive fraction of the population starts treatment in period g, and that, for all g and t, the generalized propensity score is uniformly bounded away from one" (p. 205). Again, since I opt here to match on pre-policy outcomes instead of conditioning on covariates, the necessity of this assumption is limited. Further, the sample is large and diverse enough that conditioning on additional covariates does in fact lead to substantial overlap.

To create the best chance of satisfying the parallel trends assumption, I create a matched sample for this analysis. Each district that has a charter school open in year g is matched with a single, never-treated district based on their pre-policy trends. Specifically, for each treated district d^* that becomes treated starting in year g, I identify their "nearest neighbor" as the never-treated district that has the smallest value of S_d as given in the following equation:

$$S_d = \sum_{t=2008}^{g-1} (Y_{dt} - Y_{d^*t})^2 \qquad (2.1)$$

Using this sample of districts that begin treatment in year g and their matched control district counterparts, assuming no anticipation effect, I estimate the average treatment effect on the treated, as given by CS "when pre-treatment covariates play no role in identification" (p. 206):

$$ATT_{gt} = \mathrm{E}\left[Y_t - Y_{g-1} | G_g = 1\right] - \mathrm{E}\left[Y_t - Y_{g-1} | C = 1\right] \qquad (2.2)$$

The first term in Equation (2.2) is the change in the outcome between year t and the year before the treatment was introduced, $g - 1$, among districts that are observed to implement the treatment in year g (i.e., districts for which $G_g = 1$). The second term in Equation (2.2) is the change in the outcome between year t and the year before the treatment was introduced, $g - 1$, among control districts that are observed to never implement the treatment (i.e., districts for which $C = 1$). Thus, Equation (2.2) is a restatement of the familiar difference-in-differences framework, but allowing for implementation in different years, g, and assessment of outcomes at different number of years after implementation, t.

I estimate ATT_{gt} numerous times: with g set equal to each year 2010 to 2018 and with t set equal to each year from g to 2018 (producing 45 estimates of ATT_{gt}); with ATT_{gt} estimated for reading and math test scores in grades 4th, 5th, 6th, 7th, and 8th (10 combinations) and these two subjects and five grades combined (1 combination); and with ATT_{gt} computed for Y set equal to the district's mean test score or the standard deviation of the district's test scores (2 combinations). In total, I compute 990 estimates of ATT_{gt} (i.e., 45×(10+1)×2). For the purposes of producing sound policy analysis, having 990 estimates is not particularly helpful. In response, I combine these estimates by computing the weighted average estimated effects found for each post-implementation year (i.e., for each year between g and 2018). This combination allows us to see how the impact of charter schools evolves in subsequent years. Note that I have nine cohorts for which I estimate ATT_{gt} for the first year after implementation, but I only have one cohort (i.e., districts that began operating charter schools in 2010) that is followed for nine years after implementation. Correspondingly, the standard errors on the combined estimates will be smaller when I am estimating effects for the immediate years after implementation and will be larger when I am estimating effects for many years after implementation. These standard errors are calculated using a simple multiplier bootstrap procedure, which has the advantage of simultaneously computing valid confidence bands in g and t without creating multiple-testing problems (Callaway and Sant'Anna 2021, p. 214). I additionally provide an estimate of the overall treatment effect in the observed post-policy years (i.e., collapsing post-policy years to one estimate) for the effects on both the mean test score and the standard deviation of test scores.

2.4 DATA

The district level data on the outcome, Y_{dt}, comes from the Stanford Education Data Archive (SEDA 2021). There are two primary outcomes of interest that come from the SEDA data. First, I use the mean test scores in the district as estimated in SEDA. These achievement estimates combine scores from multiple tests and are scaled relative to national test score distributions. These scores are standardized such that a value of 1 represents a standard deviation improvement in performance over the national baseline. The second outcome is the standard deviation of those scores, as an indication of the variability of achievement in school districts. I compute the standard deviation as function of the SEDA-reported standard error of the estimated mean and the number of enrolled students. SEDA includes information for both of these metrics for the years 2009-2018.

The SEDA data is uniquely suited for these research purposes for three reasons. First, the data includes estimates for every public school and public school district across the country and does so in a way that is explicitly made to make those estimates comparable. Previous literature on the effects of charter schools are often limited to single state analysis due to the nature of education data. Much of the education data used in this research area is collected at the state level, primarily because of the differences in how different states organize and assess their students. The work that SEDA has done in making comparisons possible across state lines allows for a fuller, national analysis. The few studies that use a national scope have been limited to relatively small samples and national assessments that do not span many schools and grade levels (Betts and Tang 2016). For example, Gleason et al. (2010) conducted a national study of charter middle schools. While their methodology is rigorous, they note the limited external validity of a sample size of just 36 schools. Even with one of the few studies with a large national sample, the focus is on the direct effect of charter school education on students enrolled in those charter schools, rather than a broader investigation to the impact on district-wide outcomes (Center for Research on Education Outcomes (CREDO) 2013).

Second, SEDA breaks down data, even at the district level, by subject and grade level, allowing detailed analysis of the potential heterogeneous effects of charter schools. Because the data includes every school, these subcategorizations are possible without losing a substantial amount of statistical power in the analysis.

Third, the inclusion of variation in measured achievement in the SEDA data is central to testing our hypothesis about the impact on achievement inequality. I investigate how the introduction of charter schools affects not just the average achievement in a district, but also how variable that achievement is. The focus on the distributional nature of educational achievement is a key contribution to how I understand the impact of charter schools and how they operate as policy options.

I merge into the SEDA data district-level information from the National Center for Education Statistic's Common Core of Data (CCD), from which I generate indicators for the presence of a charter school in operation in the district.

I begin with 10,478 districts. I drop 193 that had a charter school in operation in a year between 2009 to 2017 and then had no charter schools operating a subsequent year (through 2018). Then, I drop 787 districts that always had charter schools operating in the years 2009-2017. This leaves 9,498 districts, of which 371 districts began operating charters in the period 2010-17. My analytical dataset includes these 371 treated districts matched with their nearest neighbor control districts. Descriptive statistics for the dependent variables are located in Table A-1.

2.5 RESULTS

Figure A-1 shows the estimates of the treatment effects for the mean of test scores (i.e., averaged across math and reading and across grades 4^{th} – 8^{th}). The figure is centered at Year 0, which is the year the treatment went into effect. At either end of the figure, there are fewer data points because there is only one cohort that I am able to observe for the maximum pre-policy and post-policy gap in time. The blue circles give the point estimates of the effects in the post-policy years and the bars show 95% confidence intervals on these point estimates. The right-most estimate that is shown is the estimated effect for the 9^{th} year after implementation of districts that opened charter schools in 2010. As shown, this estimated effect, while positive, is not statistically significant. Moving to the left, I show two estimates of the effect for the 8^{th} year after implementation. This shows the effects for districts that opened charters in 2010 (slightly to the left) and 2011 (slightly to the right). Again, these estimated effects are slightly positive, but not statistically significant. None of the estimated post-policy effects are statistically significant.

Figure A-2 collapses the estimates shown in Figure A-1 by year of implementation. The combined estimates are slightly positive, peak at 0.05 s.d., and appear to be trending upwards given more years after the introduction of charter schools. But none of the estimates are statistically significant. Importantly, none of the pre-policy estimates (shown in red) are statistically significant and there is no visual evidence of a pre-policy trend. That is, control districts are trending similarly to districts that implement charter schools. This fact helps bolster confidence in the assumption of counterfactual parallel trends.

I further collapse the nine post-intervention estimates to form a single estimate of the effect on mean test scores; this estimate suggests charter school introduction causes an improvement of the district-wide mean tests scores by 0.0169 s.d., but this estimate is not statistically significant (s.e.= 0.0134). This result could be the consequence of two possibilities. First, the introduction of charter schools in a district fails to raise achievement for both students enrolled in charter schools and traditional public schools. The second possibility is that the gains made by one set of students (e.g., those enrolled in charter schools) are offset by the losses incurred by the other. In either scenario, the net effect of charter schools as a public policy is no benefit to students with respect to test scores.

This overall insignificant effect on mean tests scores might mask heterogeneity in results by grade level and subject. To examine this possibility, Figure A-3 shows the estimated effects disaggregated by both grade (four through eight) and subject (math and reading). While none of the estimates are statistically significant, if there is any longer-term increase in mean achievement, it seems that it may be driven by gains in reading scores rather than gains in math scores. Figure A-4 collapses these estimates by year of implementation.

Figure A-5, Figure A-6, and Figure A-7, which repeat this analysis using the standard deviation of the district's test scores as the dependent variable, show the major results of this chapter. In Figure A-5, I show the estimates of ATT_{gt} for inequality in test scores. I find insignificant effects on inequality across all implementation cohorts for the first three years after implementation. In years 4-6 after implementation, the results are noisy, but the point estimates suggest an increase in inequality. Figure A-6 reduces this noise in these estimates by combining them across implementation cohorts. I indeed find that the effects on inequality are positive and statistically significant in years 4-6, and peak at 0.15 s.d. That is, the "standard deviation of the district's test scores" are increased by 0.15 standard deviations in the 6^{th} year after implementation of charter schools. The effect of the introduction of charter schools on test score inequality appears to be a short burst and is not statistically significant in the 7^{th}-9^{th} years after introduction, and is, in fact, estimated to be negative in the 9^{th} year (although noisily estimated as just a single cohort of districts is followed this long). When I collapse the nine post-intervention estimates, I find that charter school introduction causes inequality in test scores to significantly widen, with the cross-district standard deviation of test scores increasing by 0.0599 s.d. (s.e. = 0.0169). To reiterate, while I find no effect of charter introduction on mean test scores district-wide, I find that the introduction of charter schools causes widening inequality in test scores within the district.

Figure A-7 shows these estimates disaggregated by grade and subject. The trend of widening inequality is consistent across all grades and subject, but the trend is particularly observed for math scores. Figure A-8 collapses these estimates by year of implementation.

As mentioned in the Methods section above, this statistical technique produces a plethora of estimates to examine. The simplest, "one-number", *ATT* estimate for each combination of grade and subject is reported in Table A-2. The overall takeaway from the results is that the introduction of charter schools does not have a detectable impact on mean test scores within a school district, but does significantly increase the inequality of those scores.

Figure A-9 illustrates the combined effects of the two models for a hypothetical school district whose test scores follow a standard normal distribution. I assume that the policy has a symmetric effect on the distribution of scores, as follows. The pre-intervention distribution, shown in blue, has a mean of 0 and a standard deviation of 1. The post-intervention distribution, shown in red, has a mean of 0.0169 and a standard deviation of 1.0599 to simulate the effects for the average district. The lower peak at the mean and fatter tails indicate that a post-intervention district has more students moving away from the average test score. In context, this means that less students are performing at an average level on assessments, with an increase in both higher- and lower-achieving students. This is a particularly important finding for policymakers who view public education as a mechanism for providing students with a baseline of opportunities (i.e., the sentiment behind the No Child Left Behind legislation). Studies that demonstrate null effects on achievement as a result of charter schools do not necessarily imply that no students were made worse off.

2.6 Conclusion

When charter schools are introduced into a public school district, there appear to be modest effects on the distribution of district achievement. First, I find null to slightly positive effects on mean achievement in the school district. These effects vary both by the cohort in which charter schools were introduced, as well as the number of years after the intervention occurred, but these effects are never estimated to be statistically significant. Second, I find null to slightly positive effects on the standard deviation of achievement in the school district. These effects also vary by cohort and years after intervention and are positive and statistically significant in the 4^{th}, 5^{th}, and 6^{th} years after charter schools begin operating in the district.

Taken together, these two effects tell an interesting story about how charter schools shape the distribution of outcomes in a school district. While raising mean achievement is likely to be viewed by nearly everyone as a desirable outcome, increasing inequality is not as clear cut. Due to limitations in the source data, our approach does not indicate the manner in which the standard deviation increased. That is, I am unable to differentiate between increasing inequality that is symmetric (i.e., no change in the skewness of the distribution of outcomes) or asymmetric.

One explanation that synthesizes these outcomes is that students who are at the high end of the achievement distribution improve as a result of charter school introduction, particularly as time goes on. However, students at the lower end of the distribution are largely unaffected. This increases the average achievement of the district, while asymmetrically increasing the standard deviation. On the other hand, students at the lower end of the distribution could be negatively impacted, but in a way that is outweighed by the gains in achievement seen by other students. In either case, it is unlikely that the combination of increasing mean achievement and widening achievement inequality is being driven by positive outcomes for lower-achieving students.

If I assume that the estimated null effects on mean achievement are, in fact, indicative of true zero effects and if I assume that the estimated positive effect on inequality in test scores is indeed accurate, then social welfare is likely reduced by the introduction of charter schools. To support this assertion, I would need to make another set of assumptions. For example, suppose that (a) individual utilities are a linear and increasing function of test scores and (b) social welfare is a concave function of individual utilities (i.e., that the shape of the social welfare function would suggest a social preference for less inequality in utilities holding constant mean utility). With these assumptions, the introduction of charter schools would lower social welfare as it would have no effect on mean utility while raising inequality in utility. Alternatively, if we assumed that (a) individual utilities are a concave and increasing function of test scores and (b) social welfare is utilitarian and equals the sum of individual utilities, then I would get the same result. Under these assumptions, social welfare would be reduced by the introduction of charters as the utility gains to the winners would be less than the utility losses to the losers.

Of course, I do not want to overemphasize the suggestion that social welfare is reduced by the introduction of charter schools for two reasons. First, despite having the full sample of U.S. school districts evaluated over a decade, the estimated effects are somewhat noisy. The statistically significant increase in inequality of test scores is only observed for three years and seems to evaporate in later years. Second, it would be unwise to assume that the effects of charter schools on welfare should only be measured by test scores in grades 4-8. A fuller analysis would account for effects on other outcomes of interest (e.g., discipline, engagement with learning, income) and would evaluate the longer-run effects.

Chapter 3. DOES THE INTRODUCTION OF CHARTER SCHOOLS LEAD TO SHORT-TERM CHANGES IN HOUSEHOLD INCOME AND EDUCATIONAL ATTAINMENT?

3.1 ABSTRACT

Using an event study difference-in-differences analysis based on the conditional parallel trends assumption, I find substantively and statistically significant effects of charter school introduction on both median household income and the proportion of district residents with at least a bachelor's degree. The average treatment effect of the former is $1,377, or 2.7% when using a log transformation. The average treatment effect of the latter is 2.2 percentage points. In the conclusion, I discuss how policymakers should consider these findings and what they say about the changing composition of a school district in response to the introduction of charter schools.

3.2 BACKGROUND AND THEORY

This chapter also takes a district-wide approach to the impact of charter schools. Unlike most of the literature surrounding charter schools, this analysis is not interested in the effects on public school students directly. Instead, I examine how the characteristics of a geographic school district change in response to the introduction of charter schools. I intentionally use variables that I suspect charter schools would be unable to impact directly within the timeframe of the analysis. The mechanism through which the outcomes would change is a change in the composition of families living within the boundaries of the district.

I seek to answer the following question: Does the introduction of charter schools to a public school district lead to short-term changes in the income and educational attainment of families living within these districts? By "short-term" changes, I am referring to impacts that happen within 10 years of charter school introduction. For most of the districts I analyze, the timeframe will be even shorter than that. The reason the short timeframe is important to note is because it reinforces the most probable causal mechanism. I contend that a charter school is unlikely to raise incomes in a district at any detectable level by educating students, giving them marketable skills, and boosting their wages within the district the charter school is located. The same is true for the attainment of bachelor's degrees. At a minimum, charter schools would need 5 years to have any theoretical direct effect on students, and that would only be for students who were high school seniors at the time the charter school beings operating. For younger students, the theoretical minimum is even longer.

The mechanism that would drive these potential changes would thus be a change in the composition of families within the district (i.e., families entering/leaving the district boundaries). Evidence in the literature suggests that some families are drawn to charter schools, even when they are not the highest academically performing school in the area (Kleitz et al. 2000; Villavicencio 2013; Bulkley and Fisler 2003). If families seek out districts where charter schools are opening, that could be one potential mechanism for the composition of the district to change.

I use U.S. national data form 2008-2018 on median household incomes and educational attainment rates within public school districts. Using an event study difference-in-differences analysis based on the conditional parallel trends assumption, I find substantively and statistically significant effects of charter school introduction on both median household income and the proportion of district residents with at least a bachelor's degree. The average treatment effect of the former is $1,377, or 2.7% when using a log transformation. The average treatment effect of the latter is 2.2 percentage points. In the conclusion, I discuss how policymakers should consider these findings and what they say about the changing composition of a school district in response to the introduction of charter schools.

3.3 METHODS

I use a difference-in-differences methodology for time-varying treatments created by Callaway and Sant'Anna (2021). This approach allows me to estimate the causal impact of charter school operation in a public school district on subsequent levels of median household income in those districts. The interpretation of a causal impact relies on the parallel trends assumption; in the absence of treatment, both treated and untreated districts would have continued to trend parallel with each other. While this is an untestable assumption, we can look at the evidence from pre-intervention periods to see if there is evidence of any trend that would suggest that the treatment and comparison group are fundamentally different.

Beyond the typical difference-in-differences assumption, the CS method relies on the irreversibility of treatment. That is, once a school district has a charter school operating, the assumption is that a charter school will always be operating in that district. This assumption is not strictly true in the data. There are 193 districts of the initial 10,478 that went from having a charter school to not having a charter school. In this case, I had two ways to deal with the violated assumption. First, I could leave the data as it is. This would subtly change the interpretation of the causal estimate from *actually* having a charter school present and operating in the school district to *ever* having had a charter school present. Second, I could remove these districts from the sample to keep the causal interpretation the same, with the knowledge that the analysis is based on a sample that didn't allow for reversible treatments. Because of the small proportion of districts affected, I opted for the second option and removed these districts from the sample.

Instead of using a matched sample in this analysis, I opt to condition on covariates due to the increased analytical sample size. Unlike in Chapter 2, which relied on test score data that was suppressed for the smallest of districts, the use of 5-year ACS data on income allows for a broader range of examined districts. Further, the data used in this chapter includes useful demographic information to condition on. For robustness, I present the results of both the unconditioned and conditioned models.

Using this sample, I estimate the average treatment effect on the treated, given by:

$$ATT_{gt} = \mathrm{E}\left[Y_t - Y_{g-1} | G_g = 1\right] - \mathrm{E}\left[Y_t - Y_{g-1} | C = 1\right] \tag{3.1}$$

I estimate ATT_{gt} multiple times in the analysis, for each group of districts treated in a particular year. For the years of data I have in the sample, this produces 45 separate estimates. To better understand the actual dynamics of the policy intervention, I then consolidate these estimates and their confidence intervals into groupings by year after the policy intervention. Because the greatest number of cohorts exist nearest to the point of intervention, the aggregated confidence intervals will be narrowest there. For the estimates that are several years removed from the intervention, the confidence intervals will be wider.

3.4 Data

The data for the outcome of interest comes from the American Community Survey (ACS). The ACS samples many different geographic regions, including elementary, secondary, and unified school districts. I use the median income in a district as the dependent variable in the analysis. For each school district, the ACS reports a median income based on an average of respondents in that geography over the last five years. Because of the nature of this rolling average, it can be difficult to observe any small, rapid changes to an area. In the context of my analysis, the data construction limitation has the effect of attenuating estimated treatment effects. However, collecting and aggregating the data in this way allows for more information about much smaller geographic areas, such as the school districts that are the subject of this analysis.

In addition to median income, I examine two other dependent variables. First, I transform the district's median household income using the natural log. This effectively changes the model estimation from one of being a marginal level increase to a marginal percent increase. I also estimate the impact of charter schools on the proportion of residents obtaining at least a bachelor's degree.

The demographic covariates that I use are the percent of a district's students that are Black, Hispanic, Asian, and Native American. The percentage of White students is omitted as a reference category to avoid perfect multicollinearity. In reality, these five race variables do not actually create perfect multicollinearity in the data, due to their sum adding up to slightly less than 100%. These categories are not mutually exclusive and some students may be recorded in a way that does not fit in to any of the five categories. Nonetheless, I exclude White students as a reference category but it should be interpreted as the combination of White students and an unreported "other" category. In addition to these variables, I also condition on the percent of students who are English Language Learners.

I use this data for the years 2008-2018 and merge in data from the National Center for Education Statistics Common Core of Data (CCD). The CCD reports whether a district has any operating charter schools on an annual basis. Thus, I am able to construct a panel data set for school districts from the years 2008-2018 that includes whether the district included a charter school for a given year as well as the median income of households and rate of bachelor's degree attainment located in the district for that year. Descriptive statistics for the dependent variables are located in Table B-1.

3.5 Results

I present here the results of both the unconditioned and conditioned models. Figures B-1 through B-6 refer to the unconditioned model. Figures B-7 through B-12 refer to the conditioned model. The aggregated *ATT* for each dependent variable and model are shown in Table B-2.

Figure B-1 shows the estimates of the treatment effects for median household income in the school district. It is centered at Year 0, the year that a charter school began operating for that particular cohort. Each point represents one of the 45 estimates of ATT_{gt} generated by the model. The red points indicate the years before charter schools were operating in the district. The expectation for the red points is thus that they are equal to 0. Prior to the treatment being introduced, we should observe no differences in the trends between districts that are never treated and districts that are eventually treated. While there is some noise in the estimates before charter schools were introduced, none of these estimates are statistically significant. Further, there doesn't appear to be any strong trends in the pre-policy period.

The blue points show the estimated treatment effect after the policy has been introduced. While 13 of the 45 ATT_{gt} estimates are less than a magnitude of $500, there are multiple estimates that are both positive and significant. These indicate a positive effect on median household income as a result of the introduction of charter schools in a school district. Across all the estimates, this change in income ranges from slightly negative (and one large negative estimate) to a nearly $10,000 increase.

Figure B-2 collapses the estimates shown in the previous figure by the year in which charter schools began operating in the district. This aggregation allows for the computation of smaller standard errors, yielding statistically significant results for some of the post-policy periods. Immediately following treatment, the average district saw median household income rise by nearly $1,800. The estimate is also positive and significant in the third year of policy implementation (year 2). The remaining early years are bordering on significance, with relatively high estimates. Eventually, the estimates trend toward zero, with the last two periods seeing slightly negative and statistically insignificant effects. It is worth reiterating that these last periods are estimates based on comparatively less data than the estimates that are close to the actual period of charter school implementation.

Collapsing all 45 estimates into a single number gives a treatment effect of $1,218 (s.e. = $537). That is, on average, a school district begins operating charter schools sees their median household income increase by $1,218.

Figures B-3 and B-4 display the ATT_{gt} and ATT_t estimates for logged median household income. As one might expect, the overall pattern of the estimates presented is similar to the that of Figures B-1 and B-2, which use the untransformed median income. However, the scale is now reported in percent increases to median income. For example, the estimated ATT_t for t=0 in Figure B-4 shows an effect size of approximately 0.03, or a 3% increase in median income.

Collapsing all 45 estimates into a single number gives a treatment effect of 0.0203 (s.e. = 0.009), which represents an average 2% increase in median household incomes in school districts that begin operating charter schools.

Figures B-5 and B-6 show the ATT_{gt} and ATT_t estimates for the proportion of district residents that have attained at least a bachelor's degree. The disaggregated estimates in Figure B-5 shows the range of impacts is slightly negative to significantly positive. The aggregated estimates in Figure B-6 make it clear that there is a significant, immediate increase in the proportion of residents who have obtained at least a bachelor's degree. The largest estimates come between 3 and 6 years after treatment, with the largest effect being slightly more than 0.02, an increase of 2 percentage points in the proportion of district residents. Once aggregated, the estimated *ATT* is 0.0158 (s.e. = 0.004), which indicates a 1.6 percentage point increase the proportion of district residents who have an educational attainment of at least at bachelor's degree.

Figures B-7 though B-12 repeat the analysis, with the addition that the model is now conditioned on the percentage of students who are Black, Hispanic, Asian, Native American, and English Language Learners. The results are substantively similar to the unconditioned models. The estimated *ATT* on household median income is $1,377 (s.e. = $555). For logged median income the estimate is 0.0267 (s.e. = 0.010). For the educational attainment of bachelor's degrees, the estimate is 0.022 (s.e. = 0.006). Based on the models conditioned on the demographic variables, the results show that a school district that begins operating charter schools would see median household incomes rise by $1,377 or 2.7%. Likewise, the percentage of residents who have obtained at least a bachelor's degree would increase by 2.2 percentage points.

3.6 Conclusion

While I believe there is reason to believe that the estimates I have provided here are causal, I would not go as far as to say they are direct. That is, I do not believe that the introduction of a charter school to a public school district somehow directly transfers additional income to the households located in that district through a mechanism such as salary. Similarly, charter schools do not confer bachelor's degrees. Instead, the evidence suggests that the presence of charter schools may impact the composition of households within a school district, potentially attracting households with higher incomes and higher educational attainment to relocate. If this were the case, it seems unlikely that we would observe a diminishing effect over time. Such an effect would be possible if there was some sort of delayed negative impact to household incomes in the area (e.g., the presence of charter schools created an issue with the local economy that took multiple to years to develop). However, given the nature of the imprecise estimation of the periods furthest from the point of intervention, it is possible that the true effect is not diminishing at all.

Part of the school choice debate revolves around the use of public funds that could otherwise be given to public schools. The result presented here is an interesting addition to that discussion because most public schools receive a substantial portion of their revenue through the collection of local taxes. A nearby charter school may divert some funds that could have otherwise gone to traditional public schools, but based on these findings there is evidence that it also bring more money into the local economy, through higher household incomes and more educated workers, indirectly increasing the pool of potential tax revenue.

Chapter 4. DOES THE INTRODUCTION OF CHARTER SCHOOLS IMPACT DISTRICT GRADUATION RATES?

4.1 ABSTRACT

Using a matched event study difference-in-differences analysis, I find substantial and statistically significant effects of charter school introduction on district Average Freshmen Graduation Rates (AFGR), with the average treatment effect increasing AFGR by 6 percentage points. I also explore the possibility of heterogenous treatments effects across time and find evidence that charter schools do not create such an impact immediately upon implementation. In the conclusion, I discuss how policymakers should weigh this evidence when considering charter schools as a policy option.

4.2 BACKGROUND AND THEORY

In this chapter, I continue to take a district-wide approach to the impact of charter schools. Instead of examining the impact that charter schools have on their own students or the students of nearby traditional public schools, I investigate the impact that charter schools create on all public-school students in a district, including their own.

I seek to answer the question: To what extent does the introduction of charter schools in a public school district impact graduation rates in that district? I once again take advantage of the variation in timing of charter school openings to estimate a causal impact using the Callaway and Sant'Anna (2021) difference-in-differences methodology.

There is a substantial literature on whether charter schools impact graduation rates or other longer term outcomes. Much of that literature has focused on the direct impacts that charter schools have on their own students (e.g., Sass et al. 2016; Dobbie and Fryer 2016; Berends 2015; Booker et al.

2011). These studies find that charter schools, on average, have null to slightly positive effects on high school graduation rates. Other areas in the literature have focused on how charter schools affect the achievement outcomes of nearby traditional public schools (e.g., Booker et al. 2008; Bohte 2004). There is some evidence of modest gains for those students, but evidence suggests they can be simultaneously harmed by student turnover and drained resources (Arsen, Plank, and Sykes 1999). In the course of potentially affecting a district's outcomes, both of these avenues could come into play.

Charter schools operate differently than traditional public schools, thus creating the opportunity for advances to be made and for traditional public schools to change in response. Studies conducted on the difference of these types of schools have found that the work culture is significantly different, with tradeoffs for teachers based on their professional values (Bomotti, Ginsberg, and Cobb 1999; Malloy and Wohlstetter 2003). These differences can impact the educational experience for students through the support that charter staff receives (Ni 2012) and the level of professional development offered to teachers (Wei, Patel, and Young 2014).

To assess the impact of the introduction of charter schools, I use U.S. national data from 1998-2009 on district-level graduation rates. Using a matched event study difference-in-differences analysis, I find substantial and statistically significant effects of charter school introduction on district Average Freshmen Graduation Rates, with the average treatment effect increasing AFGR by 6 percentage points. I also explore the possibility of heterogenous treatments effects across time and find evidence that charter schools do not create such an impact immediately upon implementation.

4.3 METHODS

In this chapter, I continue to employ the Callaway and Sant'Anna (2001) difference-in-differences methodology to estimate the causal effect of charter school introduction on subsequent levels of district graduation rates. The key assumption that must hold for a causal interpretation is the parallel trends assumption. That is, in the absence of intervention the outcome for the treated and untreated observations would have trended in a parallel fashion. Of course, the assumption cannot be tested directly but as shown below, the matching pre-policy trends provide good evidence that the assumption is plausible.

To create the best possible comparison between treated districts and untreated districts, I created a matched sample based on the pre-treatment trends of each treated district. Specifically, for each treated district d^* that becomes treated starting in year g, I identify their "nearest neighbor" as the never-treated district that has the smallest value of S_d as given in the following equation:

$$S_d = \sum_{t=1998}^{g-1} (Y_{dt} - Y_{d^*t})^2 \tag{4.1}$$

Using the CS methodology, I estimate the average treatment effect on the treated given by:

$$ATT_{gt} = \mathrm{E}\left[Y_t - Y_{g-1} | G_g = 1\right] - \mathrm{E}\left[Y_t - Y_{g-1} | C = 1\right] \tag{4.2}$$

The ATT_{gt} will be estimated 66 times in this analysis, with g starting with the group of districts that first had a charter school introduced in the 1999-2000 school year and ending with those that had a charter school introduced in 2009-10. The first cohort will only have a single pre-policy period to examine (1998-99), while having 11 post-policy periods. The opposite is true for the last cohort. As a result, the standard errors on aggregated estimates will be smallest when looking at the periods immediately before or after the policy implementation. Moving further into pre- or post-policy territory leaves fewer eligible cohorts for analysis and thus larger standard errors.

4.4 DATA

The district-level data on the outcome of interest, Y_{dt}, comes from the National Center for Education Statistics Common Core of Data. This outcome is constructed using collected data on the number of diplomas granted in a district and the number of students enrolled in grade 8-10 in the previous years for that district. These measures are recorded annually between 1998 and 2010. The outcome variable, what the NCES refers to as the Average Freshman Graduation Rate, is calculated using the following equation:

$$Y_{dt} = \frac{Diplomas_{dt}}{AFGR\_Base_{dt}} \tag{4.3}$$

The numerator is the total number of diplomas granted within a district, *d*, in year, *t*. The denominator is a constructed measure meant to create a baseline of potential graduating students for a given district. The AFGR base is constructed by:

$$AFGR\_Base_{dt} = \frac{\sum Gr8_{d,t-4} + Gr9_{d,t-3} + Gr10_{d,t-2}}{3} \tag{4.4}$$

The constructed baseline of potential graduates is calculated by averaging together the prorated enrollment of a district's eighth graders 4 years prior, the prorated enrollment of a district's ninth graders 3 years prior, and the prorated enrollment of a district's tenth graders 2 years prior. The prorated enrollment can be calculated with:

$$Gr\theta_{dt} = Gr\alpha_{dt} + \frac{Gr\alpha_{dt}}{TE_{dt}} \times UG_{dt} \tag{4.5}$$

The prorated enrollment of a grade, θ, in a district, *d*, in year, *t*, begins with the non-prorated enrollment for that grade, $Gr\alpha_{dt}$. The district's ungraded students, UG_{dt}, are apportioned to each grade consistent with the ratio of the non-prorated enrollment to the district's total enrollment, $\frac{Gr\alpha_{dt}}{TE_{dt}}$.

The CCD reports the AFGR for a small selection of years, but not enough to do any meaningful analysis. Fortunately, all of the measure components are reported from 1998 to 2010, allowing me to construct the AFGR as outlined in NCES documentation.

I begin with 13,231 districts. I drop 172 that had a charter school in operation in a year between 1998 to 2009 and then had no charter schools operating a subsequent year (through 2018). Then, I drop 901 districts that always had charter schools operating in the years 1998-2009. This leaves 12,158 districts, of which 417 districts began operating charters in the period 2010-17. My analytical dataset includes these 417 treated districts matched with their nearest neighbor control districts. Table C-1 shows descriptive statistics for the AFGR across the Treatment, Potential Control, and Matched Groups.

4.5 Results

Figure C-1 shows the estimates of the treatment effects for graduation rates. It is centered at Year 0, which is the first year of observation with the charter school operating within the school district. Each point represents one of the 66 estimates of ATT_{gt} generated by the model. The red points indicate the years before charter schools were operating in the district. The expectation for the red points is thus that they are equal to 0. Prior to the treatment being introduced, we should observe no differences in the trends between districts that are never treated and districts that are eventually treated. The fact that the estimates are consistently at or near 0 in the pre-policy periods is good evidence that the parallel trends assumption is plausible.

Initially, most of the post-treatment estimates are centered around 0. However, as years pass after treatment the estimates steadily become more positive. Toward the end of the analysis period, when there are fewer cohorts to examine, the estimates start to trend back toward 0.

Figure C-2 collapses the estimates shown in the previous figure by the year in which charter schools began operating in the district. This aggregation allows for the computation of smaller standard errors, yielding statistically significant results for many of the post-policy periods. Immediately following treatment, the average district saw virtually no effect until the third year after implementation (year 2). Starting at this point, the ATT_t is positive every year, although no individual year is statistically significant.

Collapsing all 66 estimates into a single number gives a treatment effect of 0.0597 (s.e. = 0.0274). That is, on average a school district that has an operating charter school sees graduation rates increase by 6 percentage points.

4.6 CONCLUSION

The analysis conducted here highlights the utility of the Callaway and Sant'Anna method and the advantage of leveraging policy implementation staggered across time. A simplistic understanding of the difference-in-differences methodology assumes a treatment effect that occurs immediately following the policy introduction. However, that is not the case as observed in this chapter. Charter schools are introduced into a school district and the average treatment effect is to raise Average Freshmen Graduation Rates by 6 percentage points. This effect appears to be heterogeneous over time, consistent with the causal mechanisms that charter schools would theoretically use to create such an effect. First, the charter schools themselves may have higher graduation rates, raising the average graduation rate of the district. Second, the charter schools may create competitive pressure on nearby traditional public schools, increasing the quality of education and thus graduation rates. In both cases, it would be expected that there may be some amount of time before any true effects start to transpire.

Increased graduation rates are an important finding for policymakers to consider when weighing charter schools as a policy option. Because of the district-wide lens this analysis takes, the resulting increase in graduation rates can only come from one of two places. First, students are moving into districts with charter schools in them and graduating at higher rates than the previous average. Second, students within a district are graduating after charter schools are introduced that otherwise would not have graduated. From a policymaker's perspective, the latter seems to be unambiguously desirable. If the former is happening, it seems to say something about students' and families' preferences for charter schools, which is also an important consideration for policymakers.

FINAL THOUGHTS

The story presented here is one that I think is different than most of the existing charter school literature, although not one that is inconsistent with its findings. The focus on district-wide measures to examine charter schools as a policy tool creates different conclusions than the typical discussion of tradeoffs that happens in other discussions in the literature. Further, this analysis has brought insight into one of the more interesting mechanisms on how charter schools could create policy impacts—changing the composition of the district.

While the methods employed here do not allow me to examine this directly, the evidence presented does create a compelling case that such changes are occurring in school districts. The finding in Chapter 2 that average test scores do not increase, yet inequality does, implies that some students are doing better while other students are doing worse. The existence of students that would do better in the face of exposure to charter schools, either directly or indirectly, creates an incentive for families to seek out districts with those schools. It is unclear if the students who do better in districts with charter schools are those from families with higher income and educational attainment. It is more likely to me that those families are simply the ones with the resources to move and seek out districts with charter schools. The implication that such families exist and are potentially seeking out charter schools is important policy information. As such, future research should examine this possibility more explicitly.

The findings of Chapters 3 and 4 provide additional evidence that families with resources could be seeking out districts with charter schools. While there does exists a direct causal mechanism for charter school introduction to impact graduation rates, the effect size is larger than I might have expected for that to be the end of the story. However, if an additional mechanism, such as families of students who are already likely to graduate, were added to the mix, the findings begin to seem more plausible.

What is clear is that introducing charter schools into a district has an impact that is much wider in scope than simply the effect on the students enrolled in that specific school. There are clear changes to student achievement in terms of widening inequality (i.e., some students do better while others do worse) and overall higher average graduation rates. This result seems to be driven, at least in part, by the changing composition of the households in the district. Households with higher income and educational attainment are coming into the district, potentially driving out lower income and less educated households. If charter schools are being sought out by these types of households, which are associated with higher achieving students in the literature, then charter schools may only be impacting district outcomes by changing the students encompassed in the district. Further, those students must come from somewhere. If households are relocating to districts that have opened charter schools, they must necessarily be leaving from districts without them (or are entering public education from the private sphere). In addition to households and students, teachers have to be hired to work at these charter schools. To the extent that these teachers are looking to work at charter schools specifically, school districts may be able to attract different teachers than they were previously able to hire.

These findings are complicated by the fact that many of the results seem to not persist over time. While there are issues with having enough statistical power to be confident about the long-term impact of charter schools, many of the outcomes that I have analyzed seem to revert back to pre-charter school levels. That presents an interesting story of short-term effects of charter schools as a disrupting force, with a decaying effect as time continues. It may well be that the true district-wide impact of charter schools is that they create opportunities for things to be shuffled around.

REFERENCES

Abdulkadiroğlu, Atila, Joshua D. Angrist, Susan M. Dynarski, Thomas J. Kane, and Parag A. Pathak. 2011. “Accountability and Flexibility in Public Schools: Evidence from Boston’s Charters And Pilots.” *The Quarterly Journal of Economics* 126 (2): 699–748. https://doi.org/10.1093/qje/qjr017.

Angrist, Joshua D., Sarah R. Cohodes, Susan M. Dynarski, Parag A. Pathak, and Christopher R. Walters. 2016. “Stand and Deliver: Effects of Boston’s Charter High Schools on College Preparation, Entry, and Choice.” *Journal of Labor Economics* 34 (2): 275–318. https://doi.org/10.1086/683665.

Angrist, Joshua D., Susan M. Dynarski, Thomas J. Kane, Parag A. Pathak, and Christopher R. Walters. 2012. “Who Benefits from KIPP?” *Journal of Policy Analysis and Management* 31 (4): 837–60. https://doi.org/10.1002/pam.21647.

Angrist, Joshua D., Parag A. Pathak, and Christopher R. Walters. 2013. “Explaining Charter School Effectiveness.” *American Economic Journal: Applied Economics* 5 (4): 1–27. https://doi.org/10.1257/app.5.4.1.

Arsen, David, David Plank, and Gary Sykes. 1999. “School Choice Policies in Michigan: The Rules Matter.” For full text: http://edtech. https://eric.ed.gov/?id=ED439492.

Berends, Mark. 2015. “Sociology and School Choice: What We Know After Two Decades of Charter Schools.” *Annual Review of Sociology* 41 (1): 159–80. https://doi.org/10.1146/annurev-soc-073014-112340.

Betts, Julian R., and Y. Emily Tang. 2016. *A Meta-Analysis of the Literature on the Effect of Charter Schools on Student Achievement*. Society for Research on Educational Effectiveness. https://eric.ed.gov/?id=ED566972.

Bifulco, Robert, and Helen F. Ladd. 2006. “The Impacts of Charter Schools on Student Achievement: Evidence from North Carolina.” *Education Finance and Policy* 1 (1): 50–90. https://doi.org/10.1162/edfp.2006.1.1.50.

Bohte, John. 2004. “Examining the Impact of Charter Schools on Performance in Traditional Public Schools.” *Policy Studies Journal* 32 (4): 501–20. https://doi.org/10.1111/j.1541-0072.2004.00078.x.

Bomotti, Sally, Rick Ginsberg, and Brian Cobb. 1999. “Teachers in Charter and Traditional Schools.” *Education Policy Analysis Archives* 7 (0): 22. https://doi.org/10.14507/epaa.v7n22.1999.

Booker, Kevin, Scott M. Gilpatric, Timothy Gronberg, and Dennis Jansen. 2008. “The Effect of Charter Schools on Traditional Public School Students in Texas: Are Children Who Stay behind Left Behind?” *Journal of Urban Economics* 64 (1): 123–45. https://doi.org/10.1016/j.jue.2007.10.003.

Booker, Kevin, Tim R. Sass, Brian Gill, and Ron Zimmer. 2011. “The Effects of Charter High Schools on Educational Attainment.” *Journal of Labor Economics* 29 (2): 377–415. https://doi.org/10.1086/658089.

Buddin, Richard, and Ron Zimmer. 2005. “Student Achievement in Charter Schools: A Complex Picture.” *Journal of Policy Analysis and Management* 24 (2): 351–71. https://doi.org/10.1002/pam.20093.

Buerger, Christian, and Robert Bifulco. 2019. “The Effect of Charter Schools on Districts’ Student Composition, Costs, and Efficiency_ The Case of New York State.” *Economics of Education Review*, no. 69: 61–72. https://doi.org/10.1016/j.econedurev.2019.01.003.

Bulkley, Katrina, and Jennifer Fisler. 2003. “A Decade of Charter Schools: From Theory to Practice.” *Educational Policy* 17 (3): 317–42. https://doi.org/10.1177/0895904803017003002.

Callaway, Brantly, and Pedro H. C. Sant’Anna. 2021. “Difference-in-Differences with Multiple Time Periods.” *Journal of Econometrics*, Themed Issue: Treatment Effect 1, 225 (2): 200–230. https://doi.org/10.1016/j.jeconom.2020.12.001.

Center for Research on Education Outcomes (CREDO). 2013. “National Charter School Study.” *CREDO, Stanford University*, June, 104.

Chaisemartin, Clément de, and Xavier D’Haultfœuille. 2020. “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects.” *ArXiv:1803.08807 [Econ]*, March. http://arxiv.org/abs/1803.08807.

Cheng, Albert, Collin Hitt, Brian Kisida, and Jonathan N. Mills. 2015. “No Excuses Charter Schools: A Meta-Analysis of the Experimental Evidence on Student Achievement.” SSRN Scholarly Paper ID 2652401. Rochester, NY: Social Science Research Network. https://papers.ssrn.com/abstract=2652401.

Chubb, John, and Terry M. Moe. 1990. *Politics, Markets, and America's Schools*. The Brookings Institution Press.

Dehejia, Rajeev H., and Sadek Wahba. 2002. "Propensity Score-Matching Methods for Nonexperimental Causal Studies." *Review of Economics and Statistics* 84 (1): 151–61. https://doi.org/10.1162/003465302317331982.

Dobbie, Will, and Roland G. Fryer. 2013. "Getting Beneath the Veil of Effective Schools: Evidence From New York City." *American Economic Journal: Applied Economics* 5 (4): 28–60. https://doi.org/10.1257/app.5.4.28.

Dobbie, Will S., and Roland G. Fryer. 2016. "Charter Schools and Labor Market Outcomes." Working Paper 22502. National Bureau of Economic Research. https://doi.org/10.3386/w22502.

Dressler, Boyd. 2001. "Charter School Leadership." *Education and Urban Society* 33 (2): 170–85.

Engberg, John, Brian Gill, Gema Zamarro, and Ron Zimmer. 2012. "Closing Schools in a Shrinking District: Do Student Outcomes Depend on Which Schools Are Closed?" *Journal of Urban Economics* 71 (2): 189–203. https://doi.org/10.1016/j.jue.2011.10.001.

Epple, Dennis, Richard Romano, and Ron Zimmer. 2016. "Charter Schools: A Survey of Research on Their Characteristics and Effectiveness." In *Handbook of the Economics of Education*, 5:139–208. Elsevier. https://doi.org/10.1016/B978-0-444-63459-7.00003-8.

Friedman, Milton. 1955. *The Role of Government in Education.* Rutgers University Press.

Garcia, David R. 2008. "The Impact of School Choice on Racial Segregation in Charter Schools." *Educational Policy* 22 (6): 805–29. https://doi.org/10.1177/0895904807310043.

Gill, Brian. 2016. "The Effect of Charter Schools on Students in Traditional Public Schools: A Review of the Evidence." Education Next. November 2, 2016. https://www.educationnext.org/the-effect-of-charter-schools-on-students-in-traditional-public-schools-a-review-of-the-evidence/.

Gleason, Philip, Melissa Clark, Christina Clark Tuttle, and Emily Dwoyer. 2010. *The Evaluation of Charter School Impacts: Final Report. NCEE 2010-4029*. National Center for Education Evaluation and Regional Assistance. https://eric.ed.gov/?id=ED510573.

Goodman-Bacon, Andrew. 2018. “Difference-in-Differences with Variation in Treatment Timing.” w25018. Cambridge, MA: National Bureau of Economic Research. https://doi.org/10.3386/w25018.

Hanushek, Eric A., John F. Kain, Steven G. Rivkin, and Gregory F. Branch. 2007. “Charter School Quality and Parental Decision Making with School Choice.” *Journal of Public Economics* 91 (5–6): 823–48. https://doi.org/10.1016/j.jpubeco.2006.09.014.

Henig, Jeffrey R. 2008. “What Do We Know About the Outcomes of KIPP Schools?” *National Education Policy Center*, November. http://nepc.colorado.edu/publication/outcomes-of-kipp-schools.

Hoxby, Caroline M., and Sonali Murarka. 2009. “Charter Schools in New York City: Who Enrolls and How They Affect Their Students’ Achievement.” Working Paper 14852. National Bureau of Economic Research. https://doi.org/10.3386/w14852.

Imberman, Scott A. 2011. “The Effect of Charter Schools on Achievement and Behavior of Public School Students.” *Journal of Public Economics* 95 (7–8): 850–63. https://doi.org/10.1016/j.jpubeco.2011.02.003.

Kleitz, Bretten, Gregory R. Weiher, Kent Tedin, and Richard Matland. 2000. “Choice, Charter Schools, and Household Preferences.” *Social Science Quarterly* 81 (3): 846–54.

Ladd, Helen F. 2019. “HOW CHARTER SCHOOLS THREATEN THE PUBLIC INTEREST.” *Journal of Policy Analysis and Management* 38 (4): 1063–71. https://doi.org/10.1002/pam.22163.

Malloy, Courtney L., and Priscilla Wohlstetter. 2003. “Working Conditions In Charter Schools: What’s the Appeal for Teachers?” *Education and Urban Society* 35 (2): 219–41. https://doi.org/10.1177/0013124502239393.

Ni, Yongmei. 2012. “Teacher Working Conditions in Charter Schools and Traditional Public Schools: A Comparative Study.” *Teachers College Record* 114 (3).

Nichols-Barrer, Ira, Phil Gleason, Christina Clark Tuttle, Thomas Coen, and Virginia Knechtel. 2016. *Do Charter School Networks Deflate as They Expand? Trends in the Impacts of KIPP Schools during a Period of Rapid Growth in the KIPP Network*. Society for Research on Educational Effectiveness. https://eric.ed.gov/?id=ED567594.

Powers, Jeanne M. n.d. “Review of No Excuses Charter Schools.” Accessed February 1, 2017. http://www.greatlakescenter.org/docs/Think_Twice/TT_Powers_CharterSchools.pdf.

Reardon, S. F., A. D. Ho, B. R. Shear, E.M. Fahle, D. Kalogrides, H. Jang, and B. Chavez. 2021. “Stanford Education Data Archive (Version 4.1).” The Educational Opportunity Project at Stanford University. 2021. http://purl.stanford.edu/db586ns4974.

Sass, Tim R. 2006. “Charter Schools and Student Achievement in Florida.” *Education Finance and Policy* 1 (1): 91–122. https://doi.org/10.1162/edfp.2006.1.1.91.

Sass, Tim R., Ron W. Zimmer, Brian P. Gill, and T. Kevin Booker. 2016. “Charter High Schools’ Effects on Long-Term Attainment and Earnings.” *Journal of Policy Analysis and Management* 35 (3): 683–706. https://doi.org/10.1002/pam.21913.

Villavicencio, Adriana. 2013. “‘It’s Our Best Choice Right Now’: Exploring How Charter School Parents Choose.” *Education Policy Analysis Archives* 21 (81): 23.

Wei, Xin, Deepa Patel, and Viki M Young. 2014. “Opening the ‘Black Box’: Organizational Differences between Charter Schools and Traditional Public Schools.” *Education Policy Analysis Archives* 22 (3). https://eric.ed.gov/?id=EJ1018832.

Wilkins, JoAnne. 2020. “50-State Comparison: Charter School Policies.” Education Commission of the States. January 28, 2020. https://www.ecs.org/charter-school-policies/.

Wohlstetter, Priscilla, Joanna Smith, and Caitlin C. Farrell. 2013. *Choices and Challenges: Charter School Performance in Perspective*. Harvard Education Press.

A. APPENDIX FOR CHAPTER 2 TABLES AND FIGURES

Table A-1: Descriptive Statistics for Dependent Variables

	Mean	**SD**	**Median**	**Min**	**Max**	**N**
Mean Test Score						
Treatment	-0.10	0.37	-0.11	-1.51	2.87	371
Potential Control	0.05	0.37	0.06	-3.74	1.62	9127
Matched	-0.09	0.36	-0.10	-1.22	0.92	371
SD Test Score						
Treatment	1.27	0.29	1.20	0.81	4.13	371
Potential Control	1.09	0.18	1.05	0.63	5.85	9127
Matched	1.21	0.22	1.15	0.87	2.72	371

Table A-2: Disaggregated Treatment Effect Estimates

	Mean			SD		
	Math	Reading	Overall	Math	Reading	Overall
All Grades	0.0089	0.0249	0.0169	0.0715*	0.0534*	0.0599*
	(0.015)	(0.014)	(0.013)	(0.015)	(0.014)	(0.017)
Grade 4	0.003	0.015	-0.008	0.0607*	0.051*	0.0634*
	(0.018)	(0.017)	(0.014)	(0.017)	(0.018)	(0.012)
Grade 5	-0.005	0.011	0.006	0.0426*	0.019	0.0401*
	(0.017)	(0.014)	(0.013)	(0.012)	(0.011)	(0.010)
Grade 6	-0.012	0.021	0.013	0.0413*	0.0256*	0.0385*
	(0.015)	(0.013)	(0.012)	(0.013)	(0.009)	(0.009)
Grade 7	-0.005	0.019	0.023	0.045	0.0295*	0.0492*
	(0.020)	(0.013)	(0.014)	(0.026)	(0.012)	(0.012)
Grade 8	-0.028	0.019	0.003	0.0438*	0.014	0.043*
	(0.021)	(0.014)	(0.013)	(0.022)	(0.015)	(0.013)

*p<.05

The treatment effects reported here are the aggregated estimates of the Callaway and Sant'Anna model. That is, for each combination of grade and subject, the 45 estimated ATT_{gt} for the pairings of district cohorts (g) and year after implementation (t) are aggregated into a single average treatment effect.

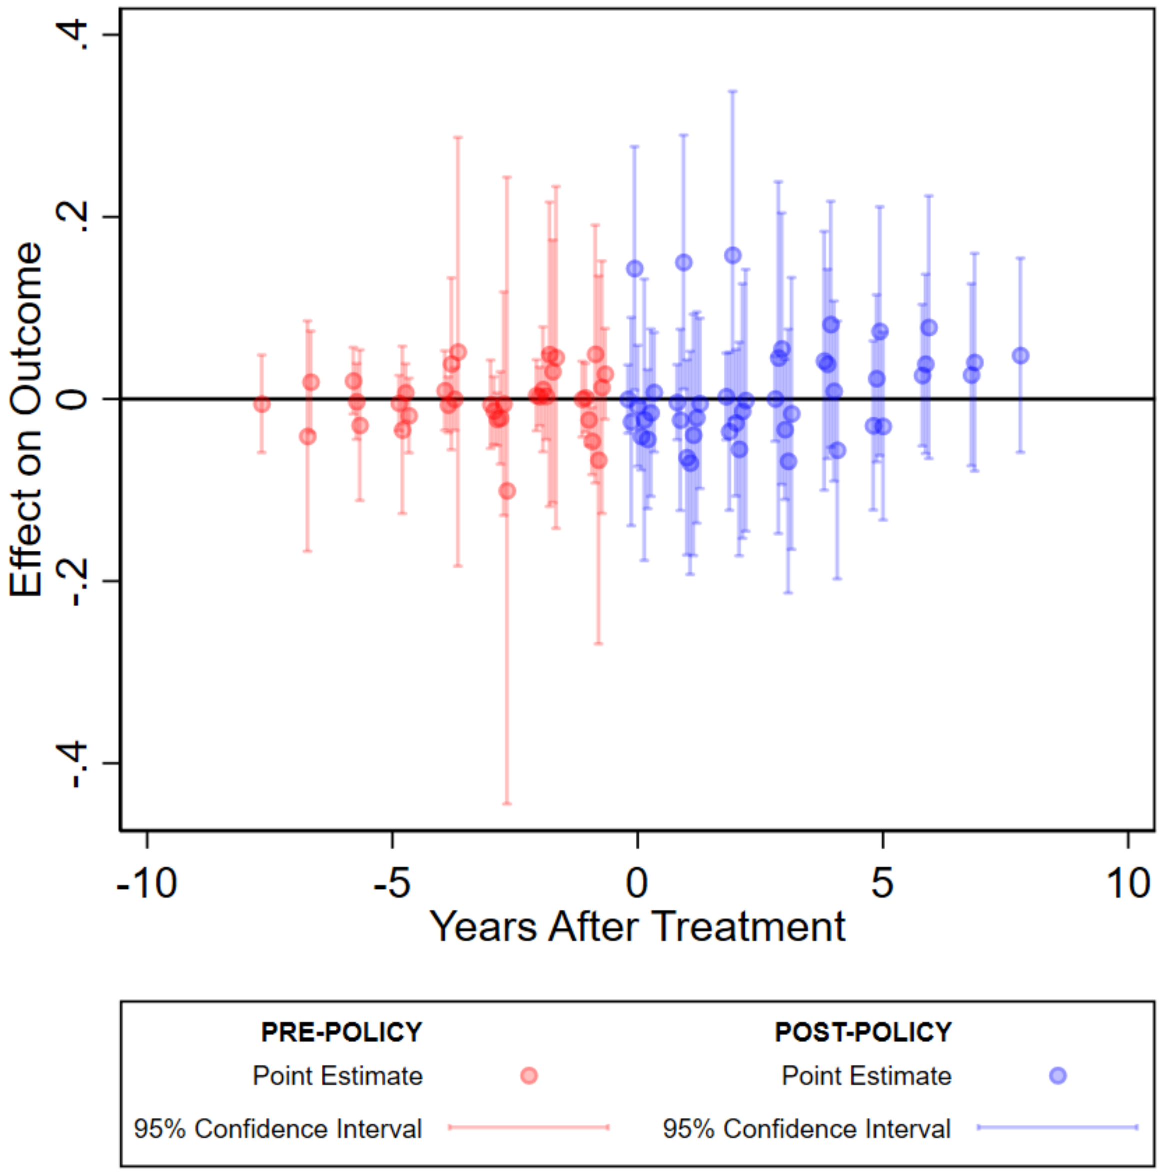

Figure A-1: Estimated ATT_{gt} for the District's Mean Test Score by Years After Treatment (t) and Implementation Cohort (g)

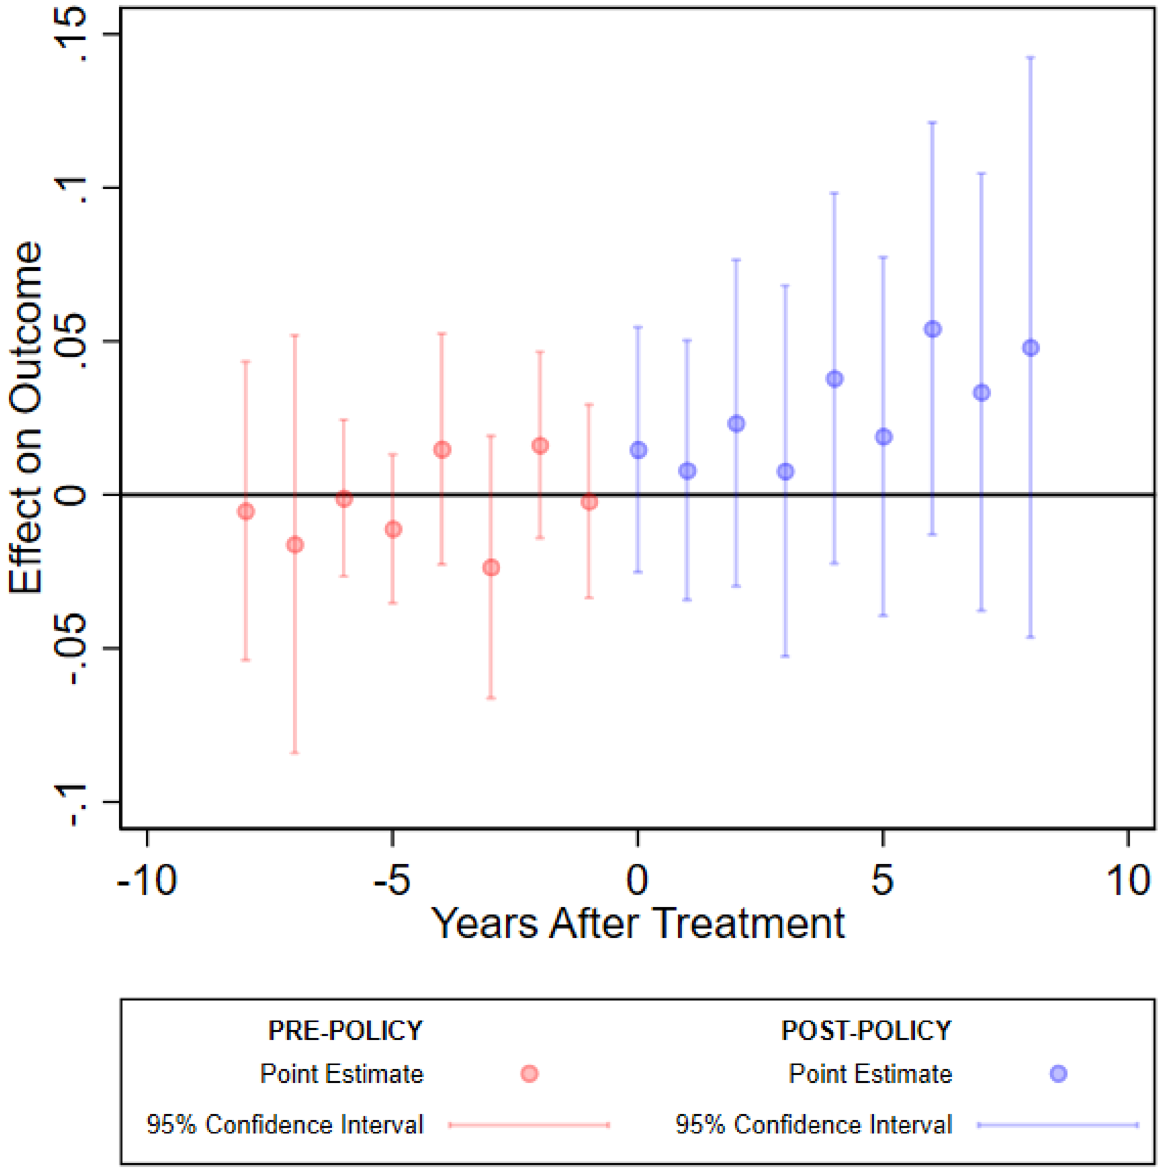

Figure A-2: Estimated ATT_t for the District's Mean Test Score by Years After Treatment (t)

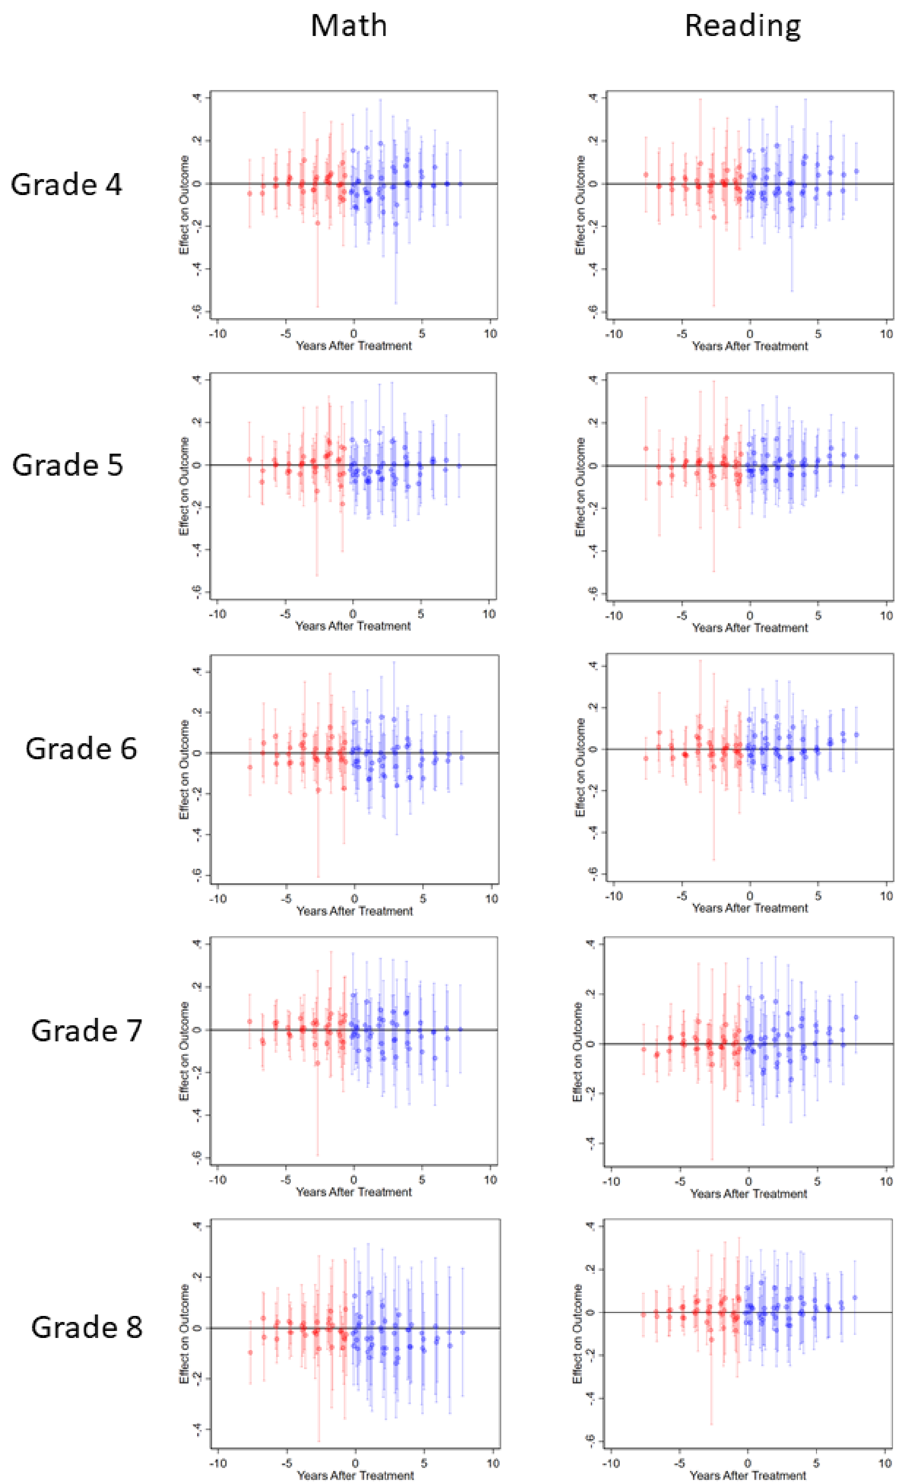

Figure A-3: Estimated ATT_{gt} for the District's Mean Test Score by Years After Treatment (t) and Implementation Cohort (g), Disaggregated by Subject and Grade

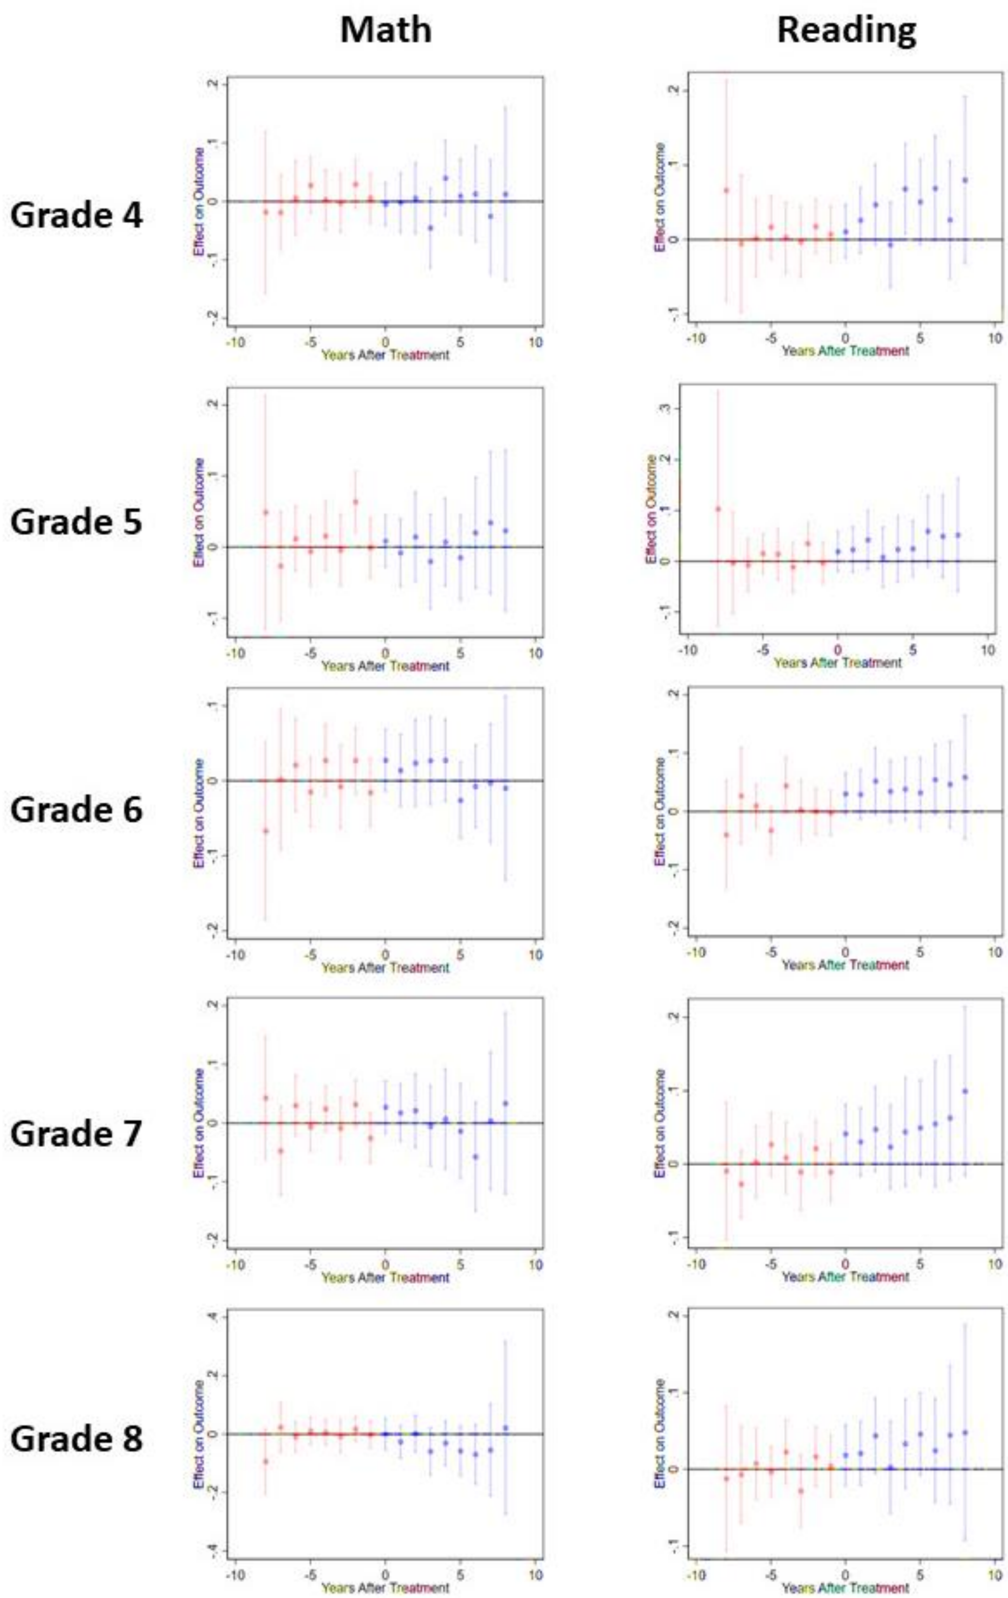

Figure A-4: Estimated ATT_t for the District's Mean Test Score by Years After Treatment (t), Disaggregated by Subject and Grade

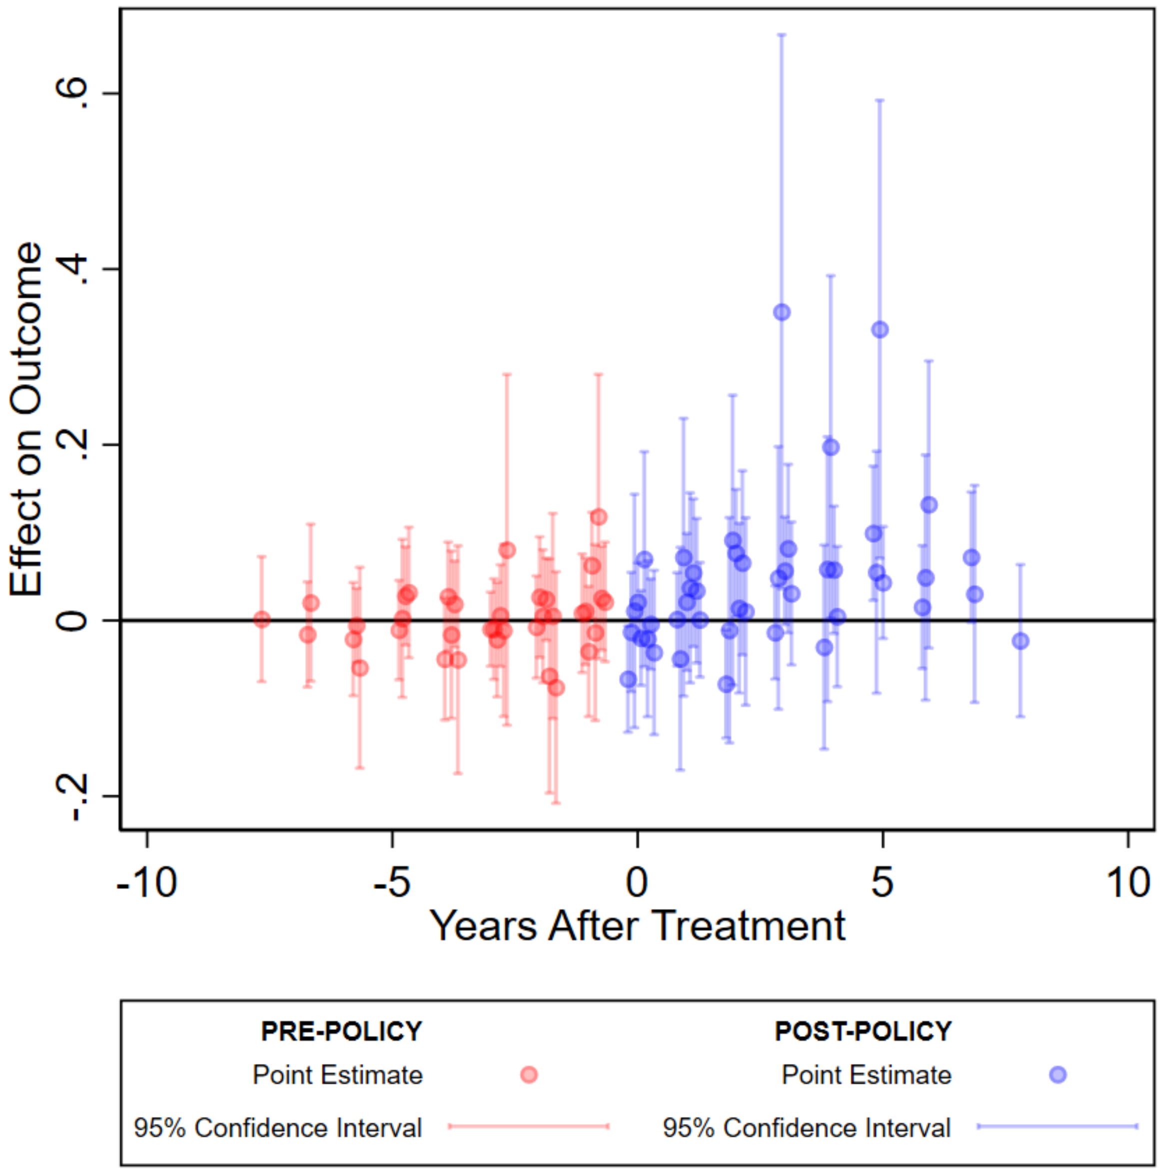

Figure A-5: Estimated ATT_{gt} for the District's Standard Deviation of Test Scores by Years After Treatment (t) and Implementation Cohort (g)

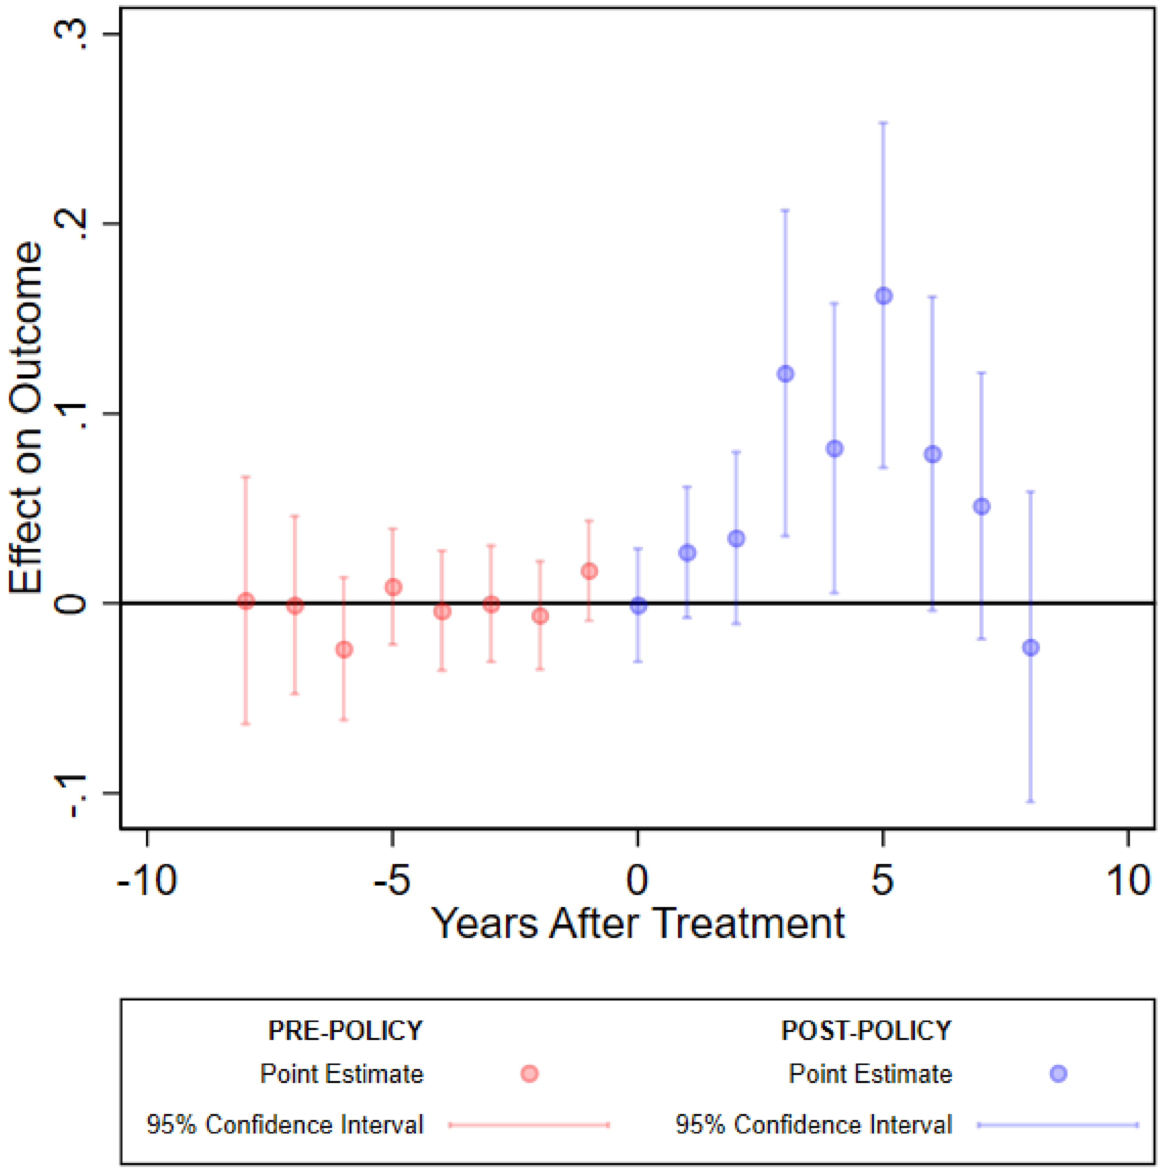

Figure A-6: Estimated ATT_t for the District's Standard Deviation of Test Scores by Years After Treatment (t)

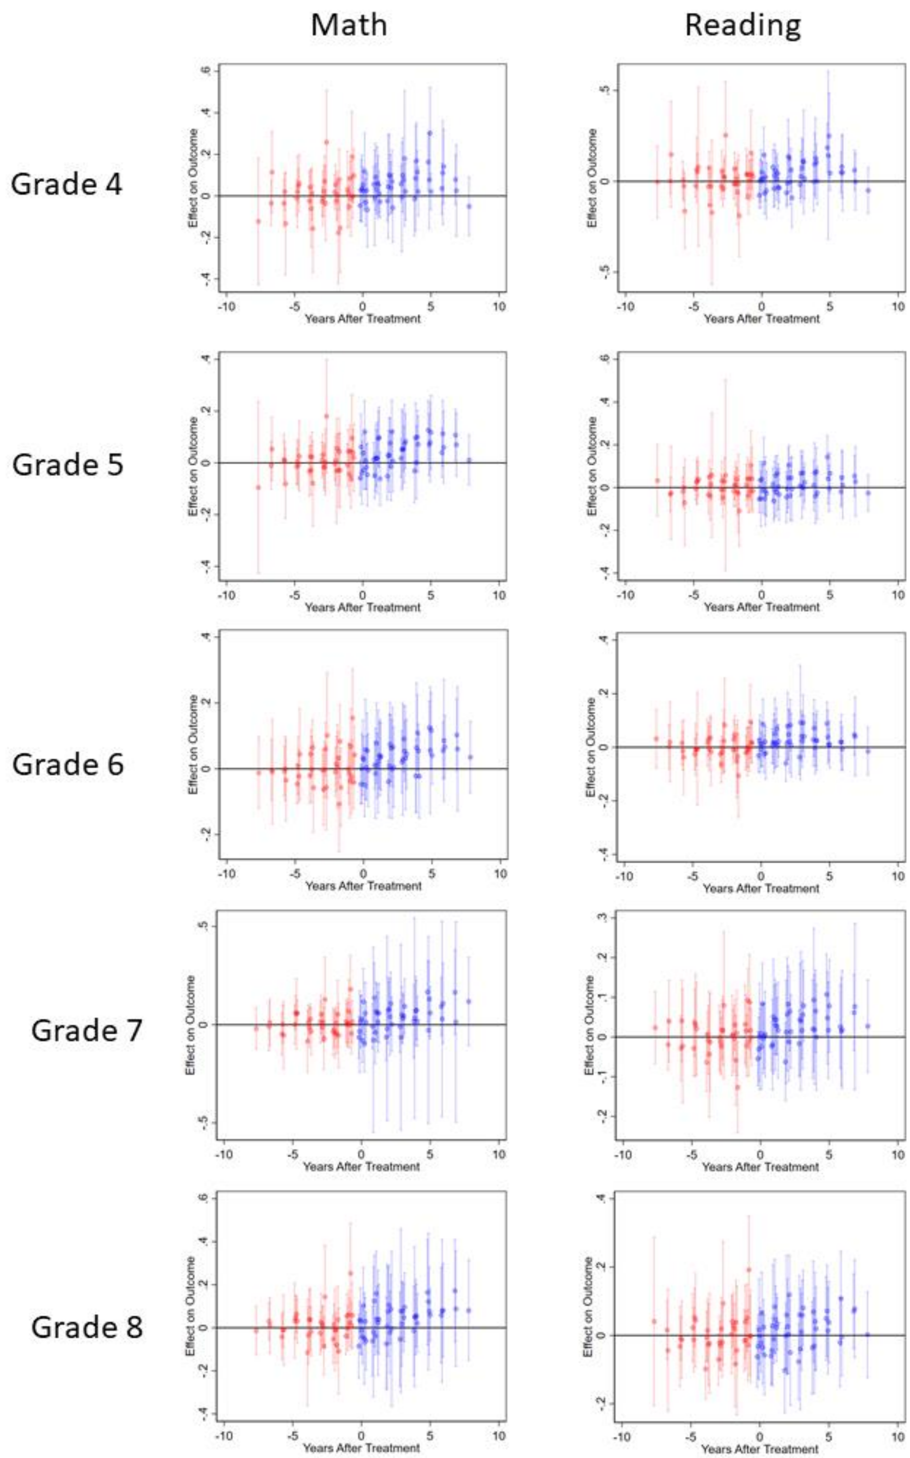

Figure A-7: Estimated ATT_{gt} for the District's Standard Deviation of Test Scores by Years After Treatment (t) and Implementation Cohort (g), Disaggregated by Subject and Grade

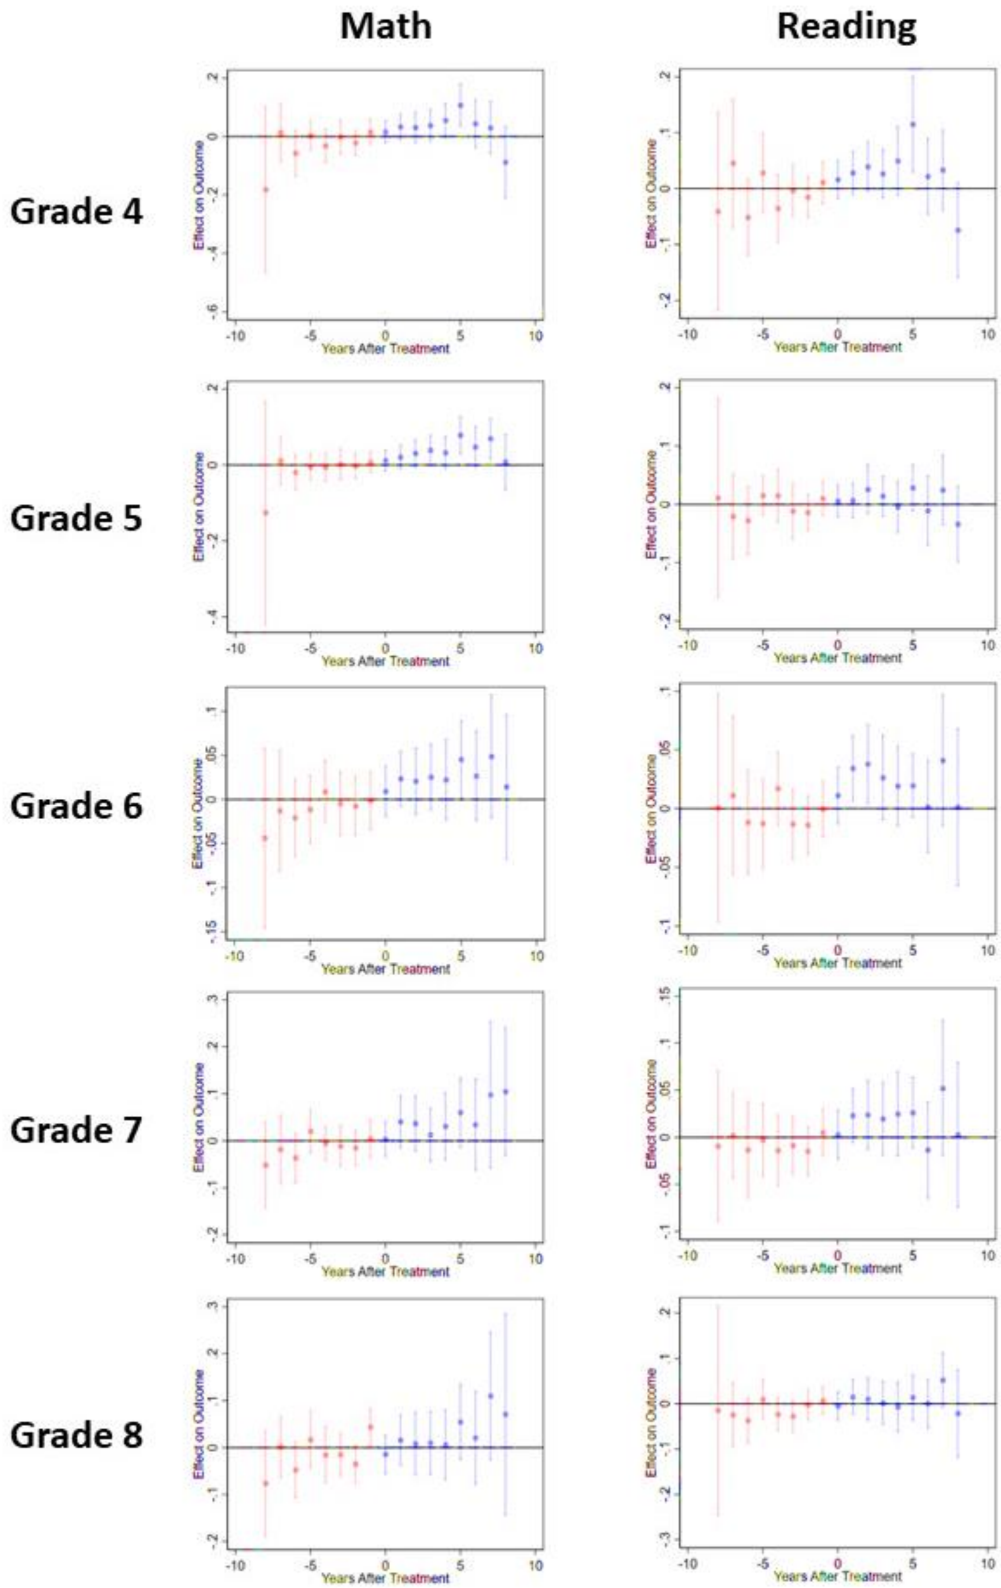

Figure A-8: Estimated ATT_t for the District's Standard Deviation of Test Scores by Years After Treatment (t), Disaggregated by Subject and Grade

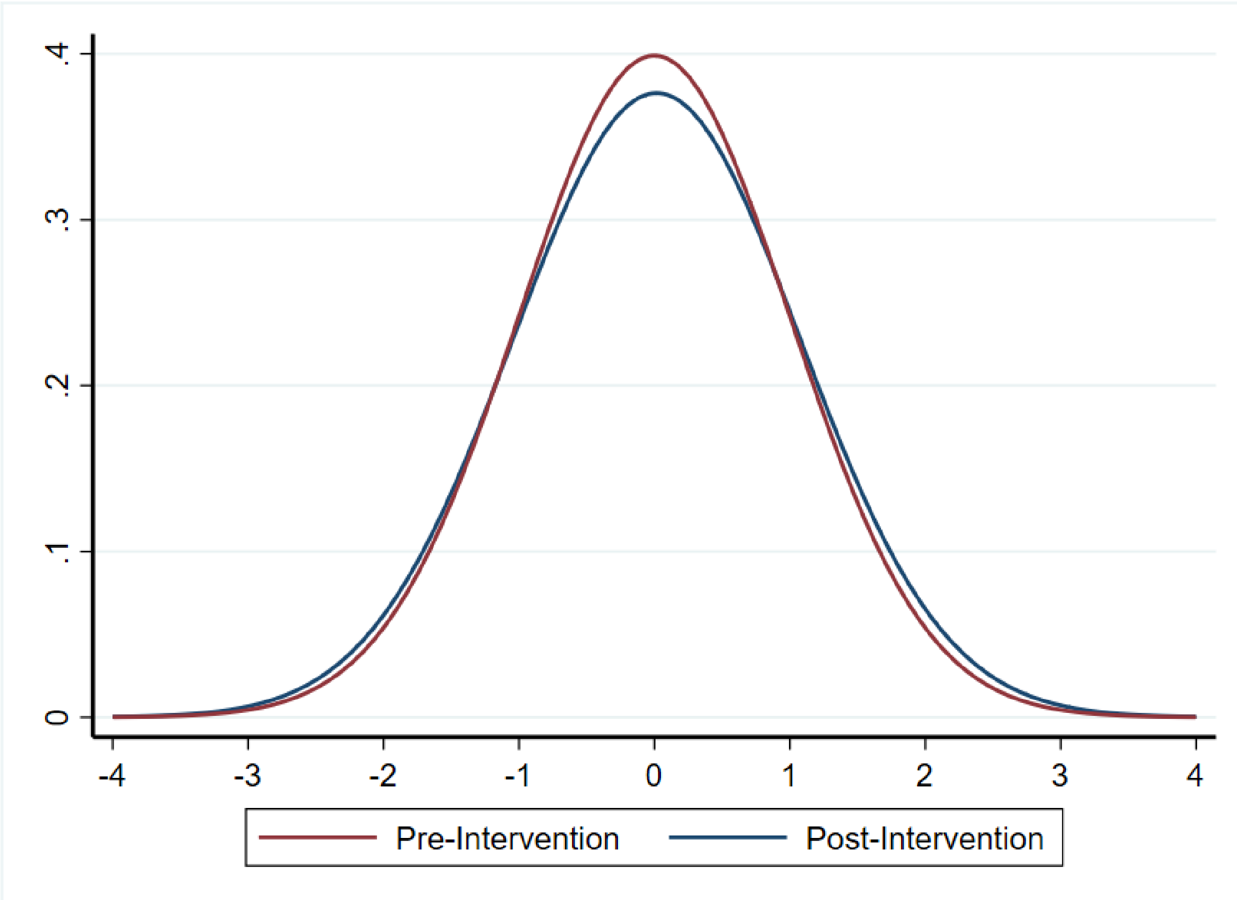

Figure A-9: Hypothetical Distribution of District Test Scores

B. APPENDIX FOR CHAPTER 3 TABLES AND FIGURES

Table B-1: Descriptive Statistics for Dependent Variables

	Mean	**SD**	**Median**	**Min**	**Max**	**N**
Median Income						
Treatment	$52,815	$17,125	$49,205	$14,217	$179,578	1,349
Comparison	$54,186	$20,052	$49,295	$12,340	$269,922	8,715
BA+						
Treatment	0.25	0.13	0.23	0.02	0.82	1,349
Comparison	0.23	0.13	0.19	0.001	0.89	8,715

Table B-2: Treatment Effect Estimates

Outcome	**Unconditioned**	**Conditioned**
Median Income	$1,218* ($537)	$1,377* ($555)
Median Income (logged)	0.0203* (0.009)	0.0267* (0.010)
Bachelor's Degree +	0.0158* (0.004)	0.022* (0.006)

*p<.05

The treatment effects reported here are the aggregated estimates of the Callaway and Sant'Anna model. That is, for each combination of grade and subject, the 45 estimated ATT_{gt} for the pairings of district cohorts (g) and year after implementation (t) are aggregated into a single average treatment effect.

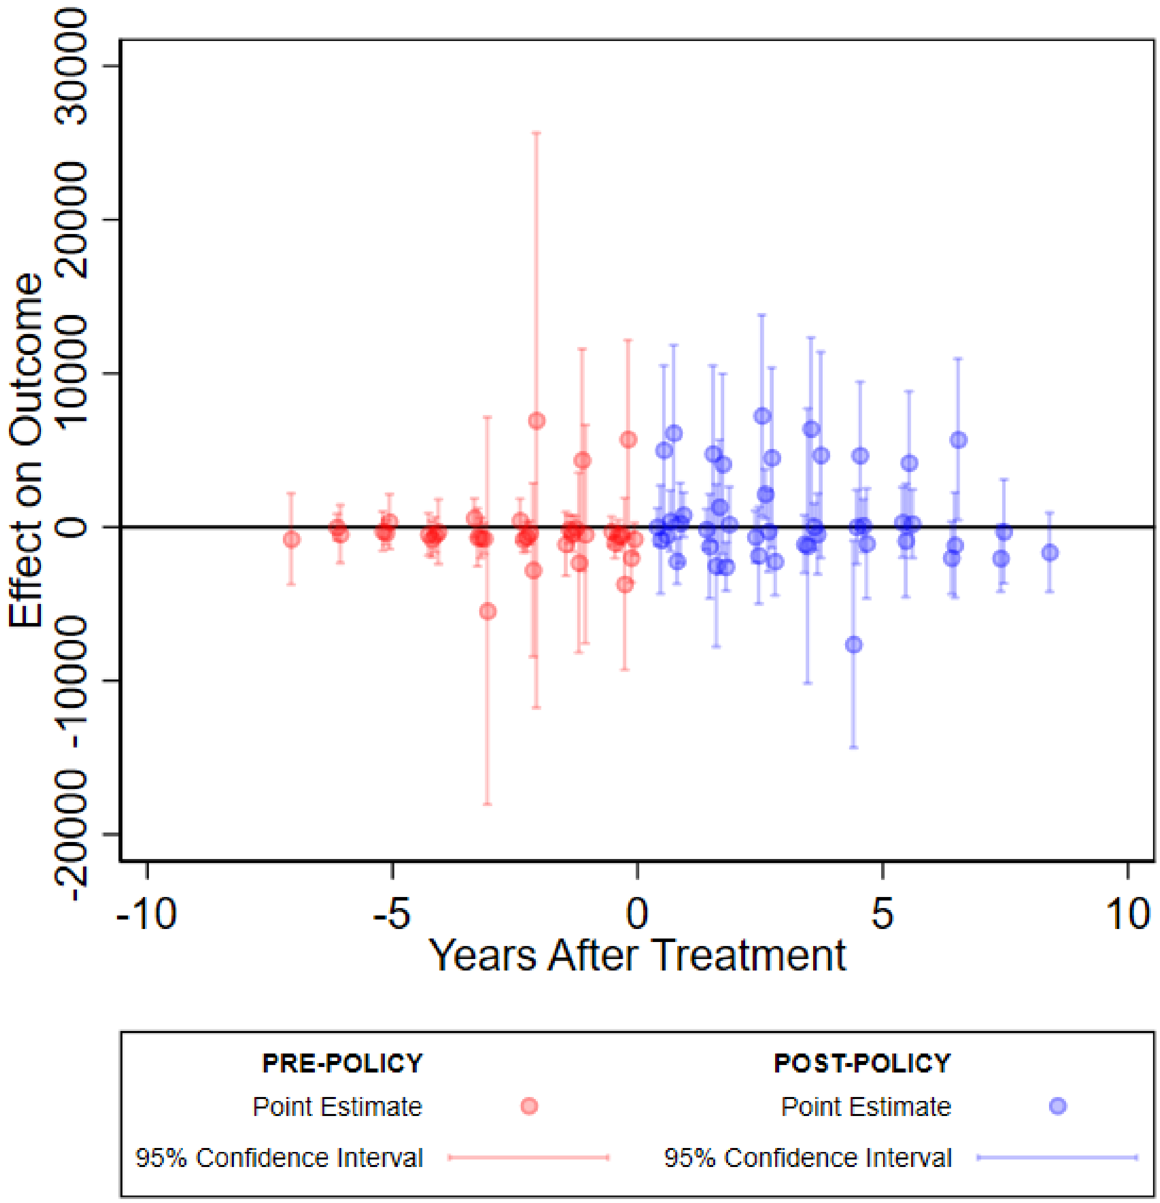

Figure B-1: Estimated ATT_{gt} for the District's Median Household Income by Years After Treatment (t) and Implementation Cohort (g), Unconditioned Model

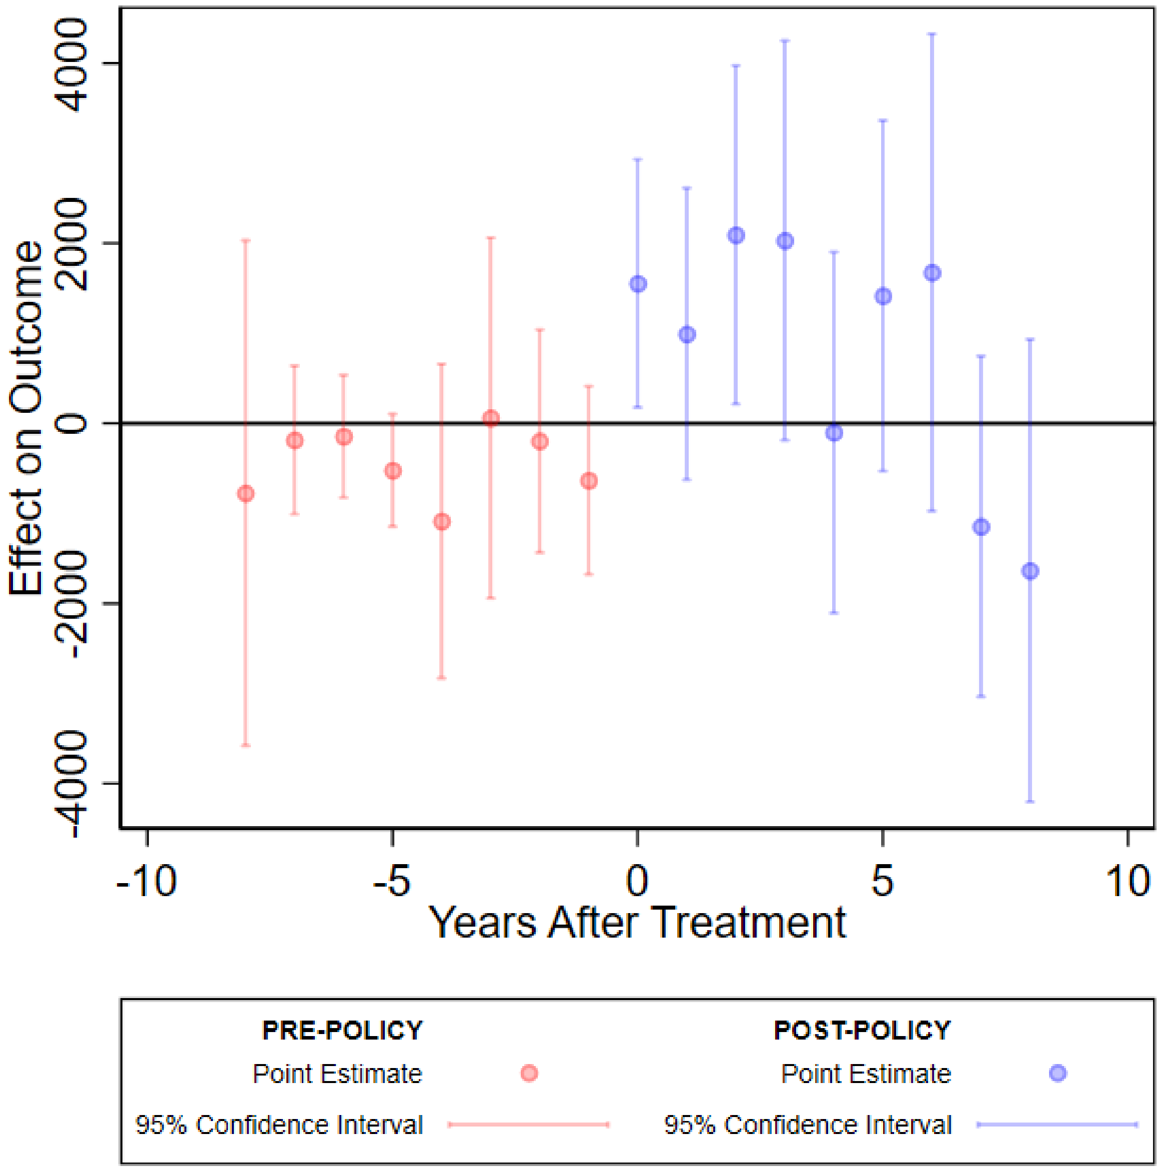

Figure B-2: Estimated ATT_t for the District's Median Household Income by Years After Treatment (t), Unconditioned Model

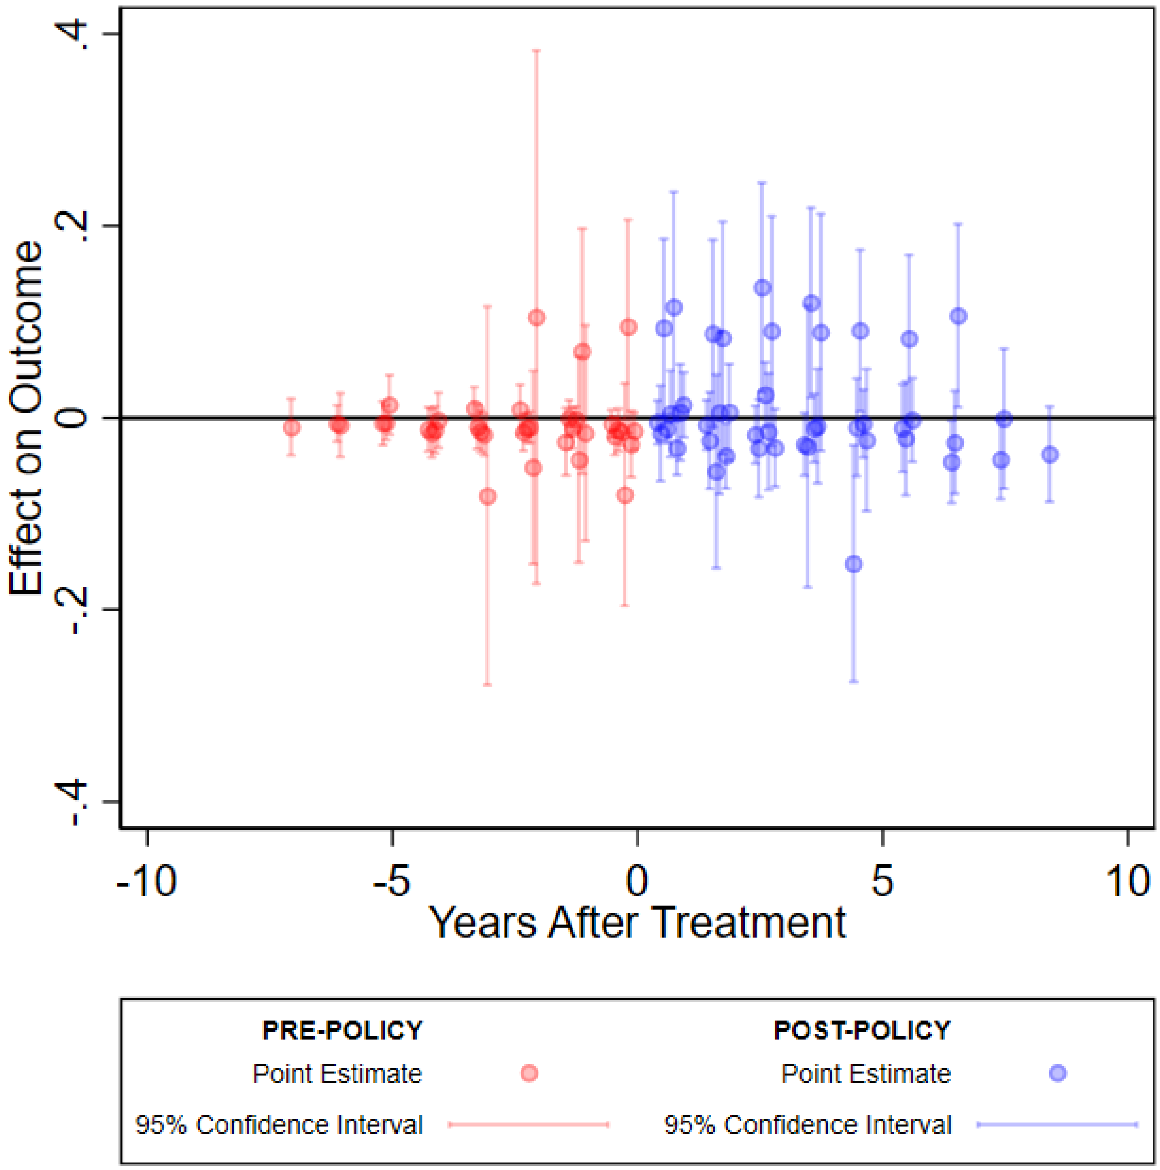

Figure B-3: Estimated ATT_{gt} for the District's Logged Median Household Income by Years After Treatment (t) and Implementation Cohort (g), Unconditioned Model

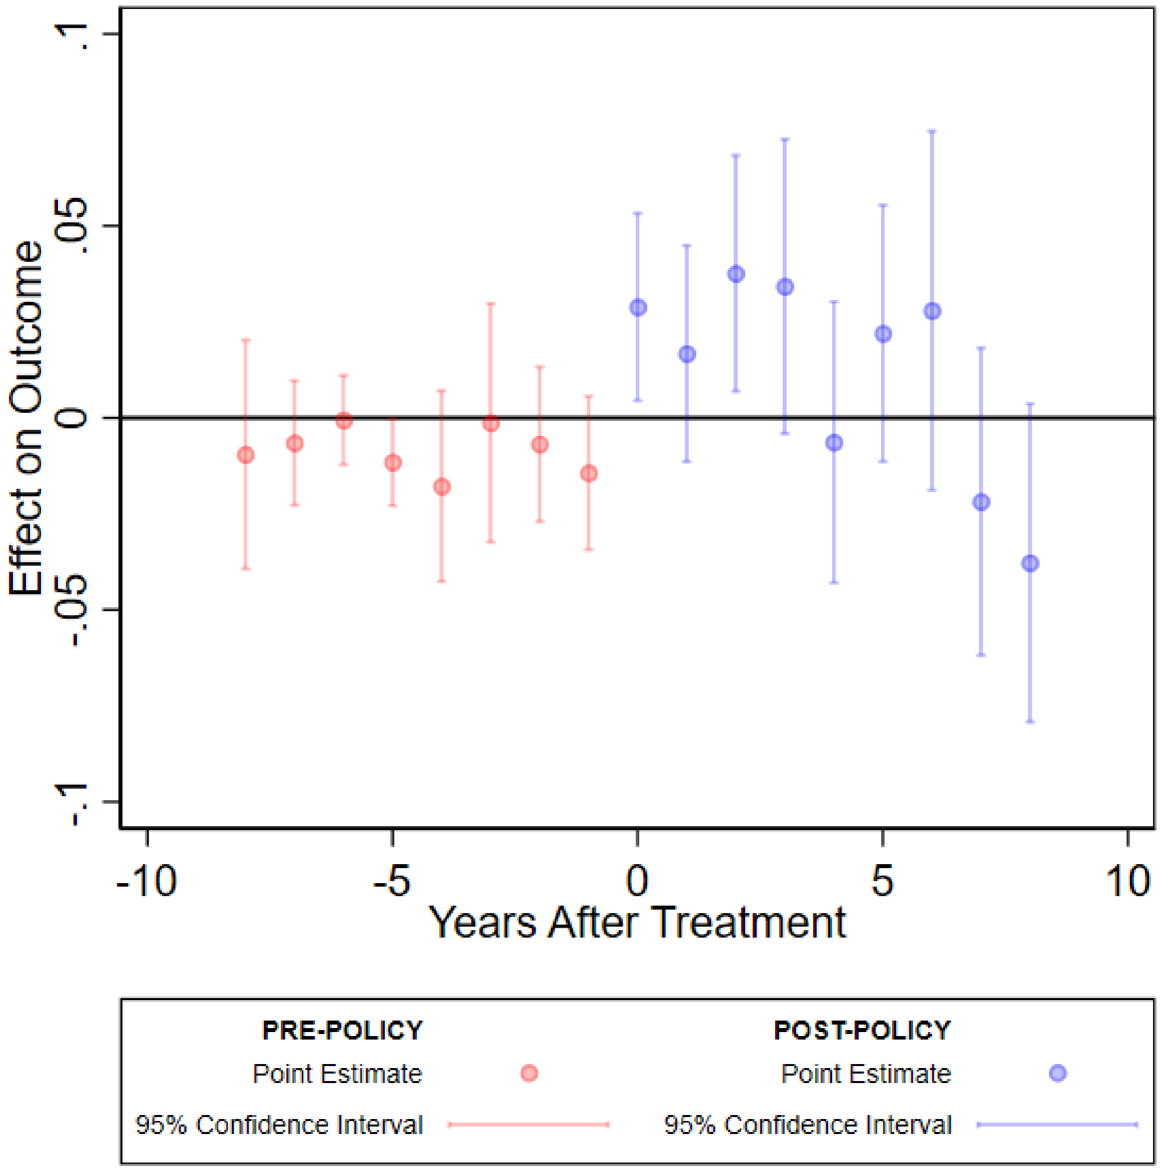

Figure B-4: Estimated ATT_t for the District's Logged Median Household Income by Years After Treatment (t), Unconditioned Model

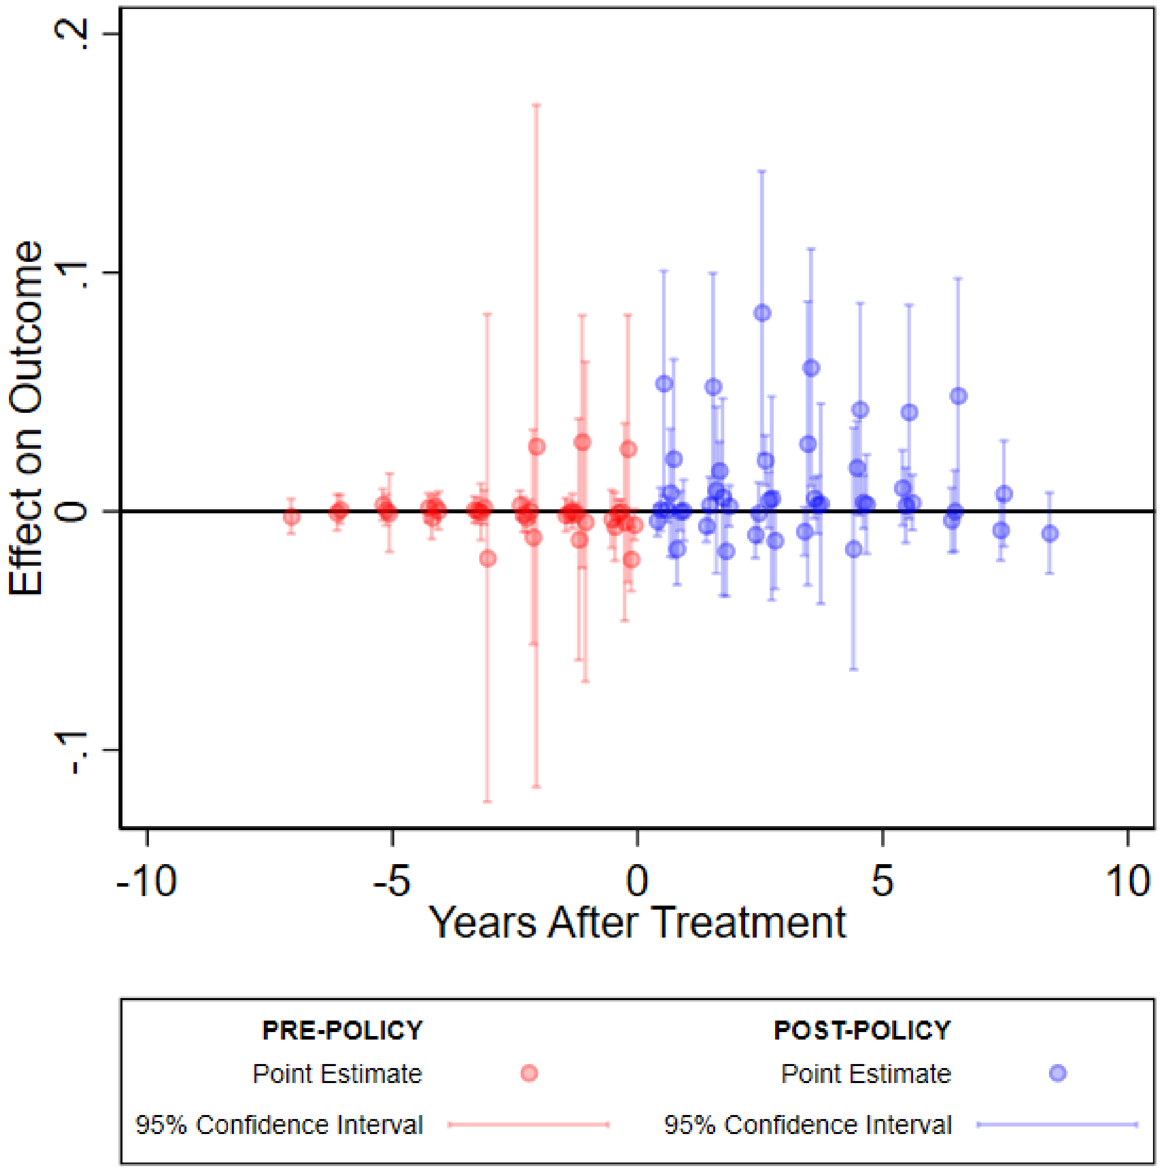

Figure B-5: Estimated ATT_{gt} for the District's Bachelor's Degree Attainment by Years After Treatment (t) and Implementation Cohort (g), Unconditioned Model

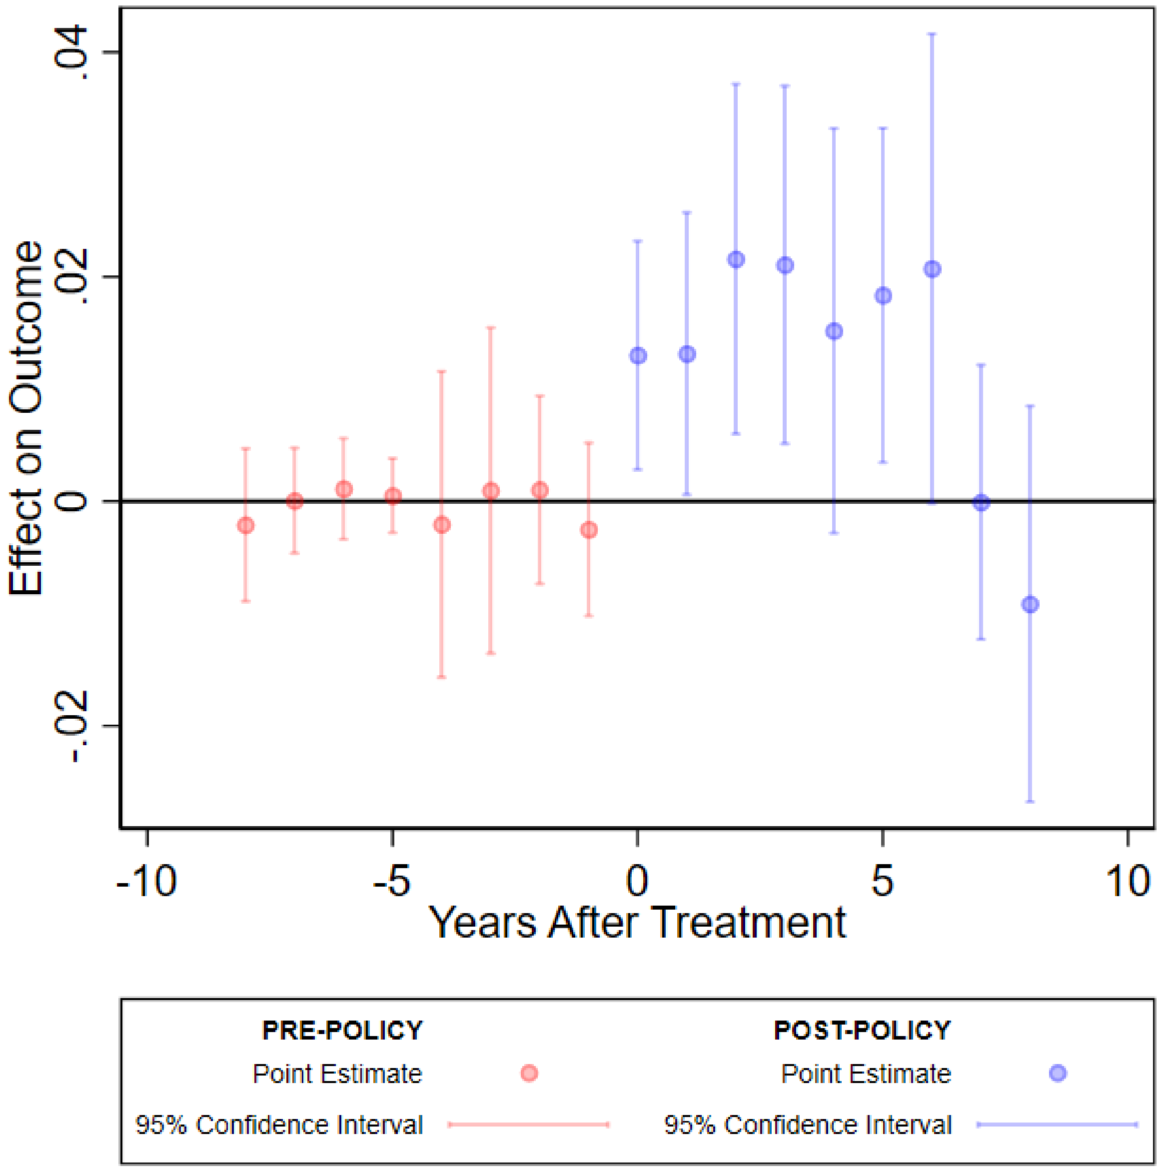

Figure B-6: Estimated ATT_t for the District's Bachelor's Degree Attainment by Years After Treatment (t), Unconditioned Model

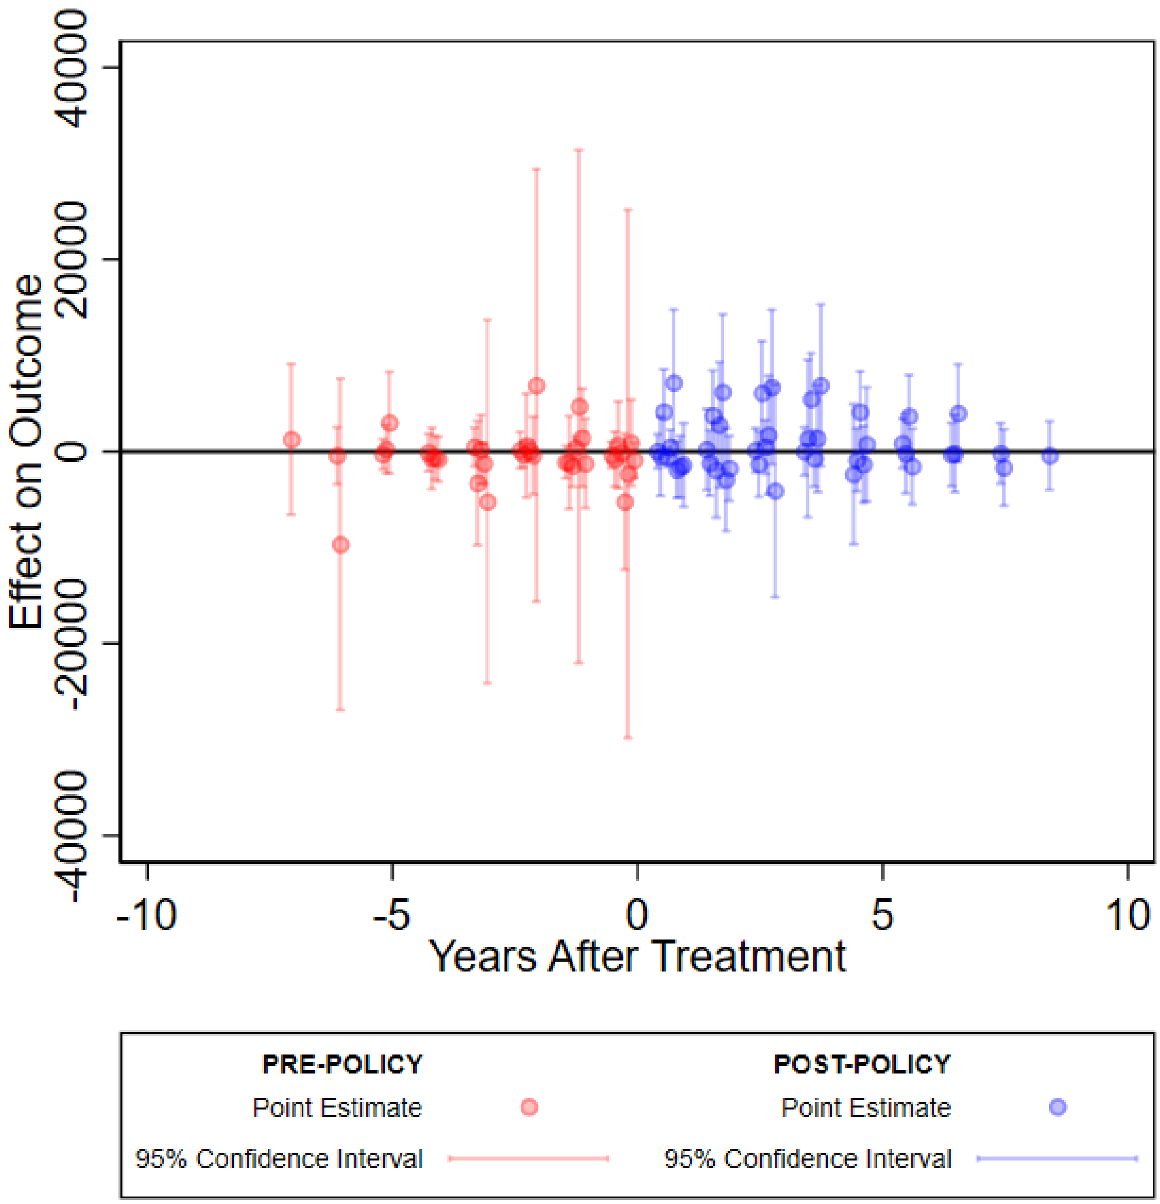

Figure B-7: Estimated ATT_{gt} for the District's Household Median Income by Years After Treatment (t) and Implementation Cohort (g), Conditioned Model

Model is conditioned on the district's percentage of students who are Black, Hispanic, Asian, Native American, and English Language Learners

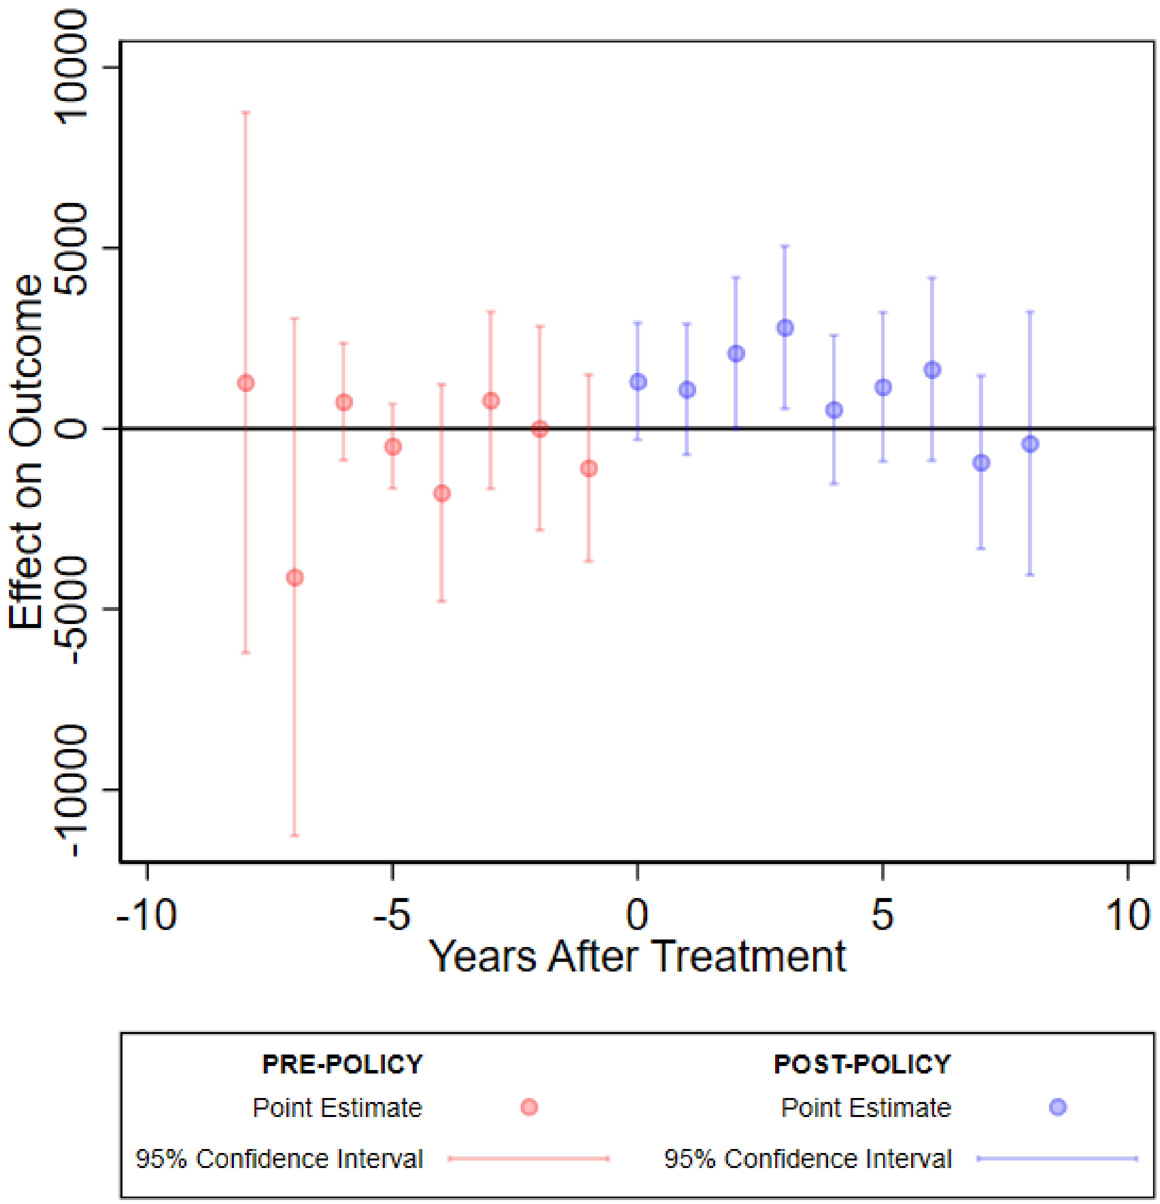

Figure B-8: Estimated ATT_t for the District's Household Median Income by Years After Treatment (t), Conditioned Model

Model is conditioned on the district's percentage of students who are Black, Hispanic, Asian, Native American, and English Language Learners

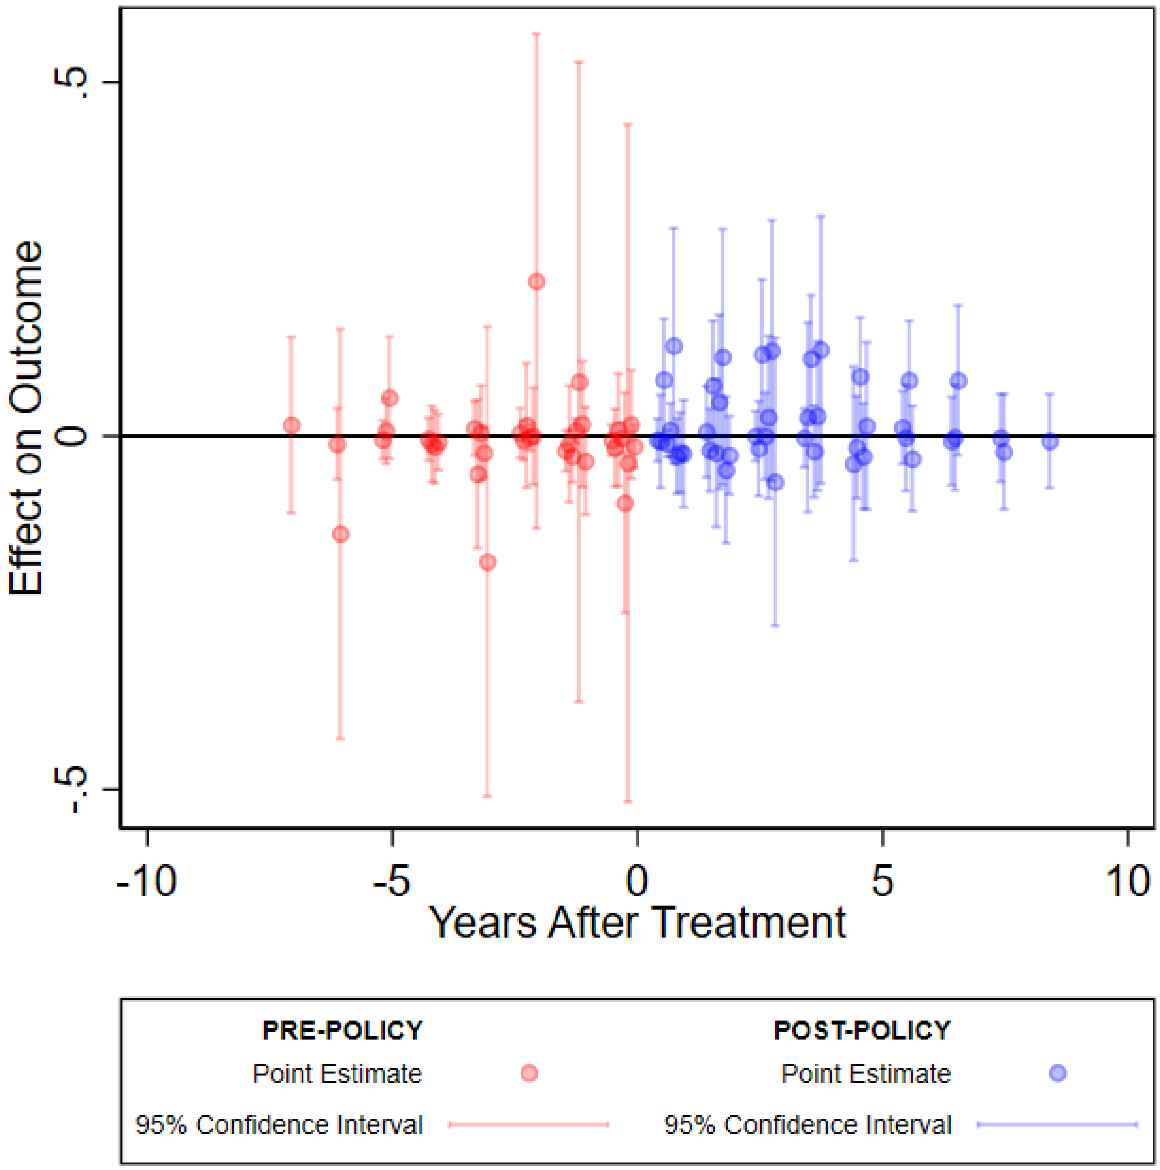

Figure B-9: Estimated ATT_{gt} for the District's Logged Household Median Income by Years After Treatment (t) and Implementation Cohort (g), Conditioned Model

Model is conditioned on the district's percentage of students who are Black, Hispanic, Asian, Native American, and English Language Learners

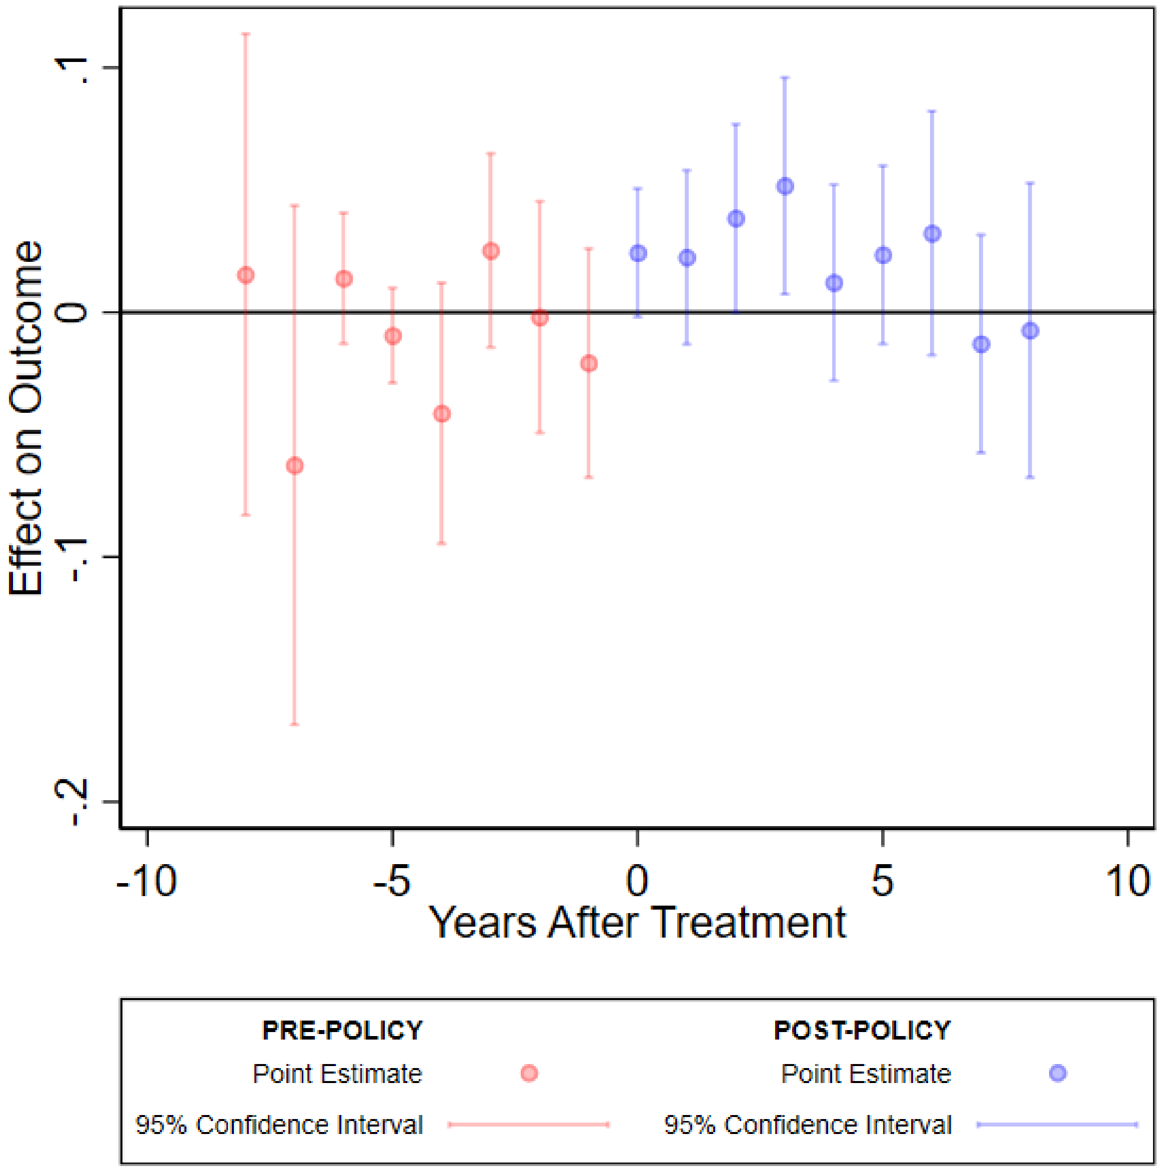

Figure B-10: Estimated ATT_t for the District's Logged Household Median Income by Years After Treatment (t), Conditioned Model

Model is conditioned on the district's percentage of students who are Black, Hispanic, Asian, Native American, and English Language Learners

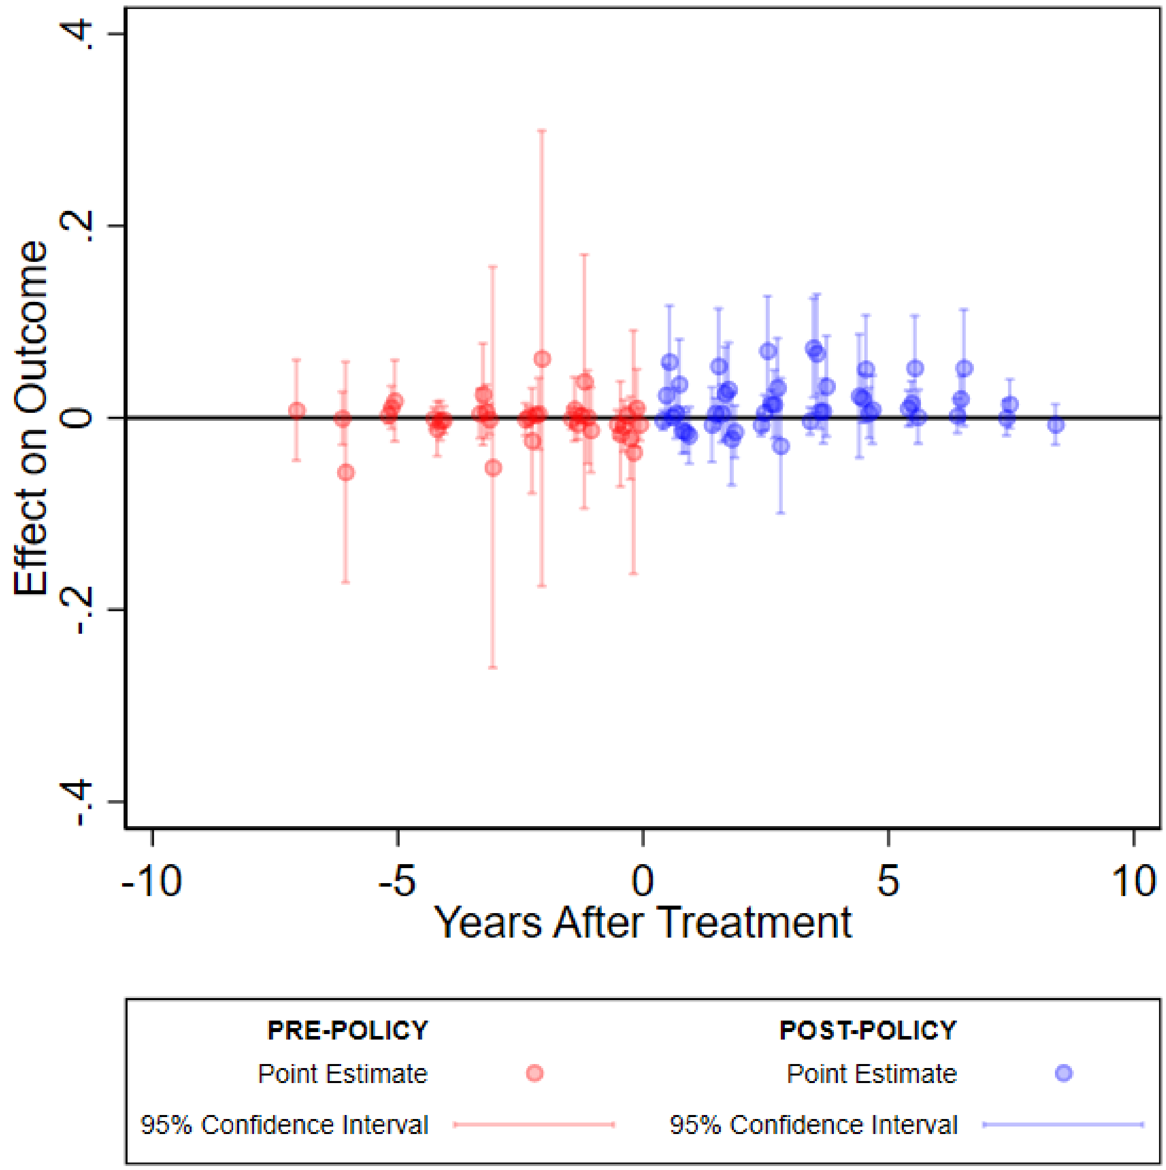

Figure B-11: Estimated ATT_{gt} for the District's Bachelor's Degree Attainment by Years After Treatment (t) and Implementation Cohort (g), Conditioned Model

Model is conditioned on the district's percentage of students who are Black, Hispanic, Asian, Native American, and English Language Learners

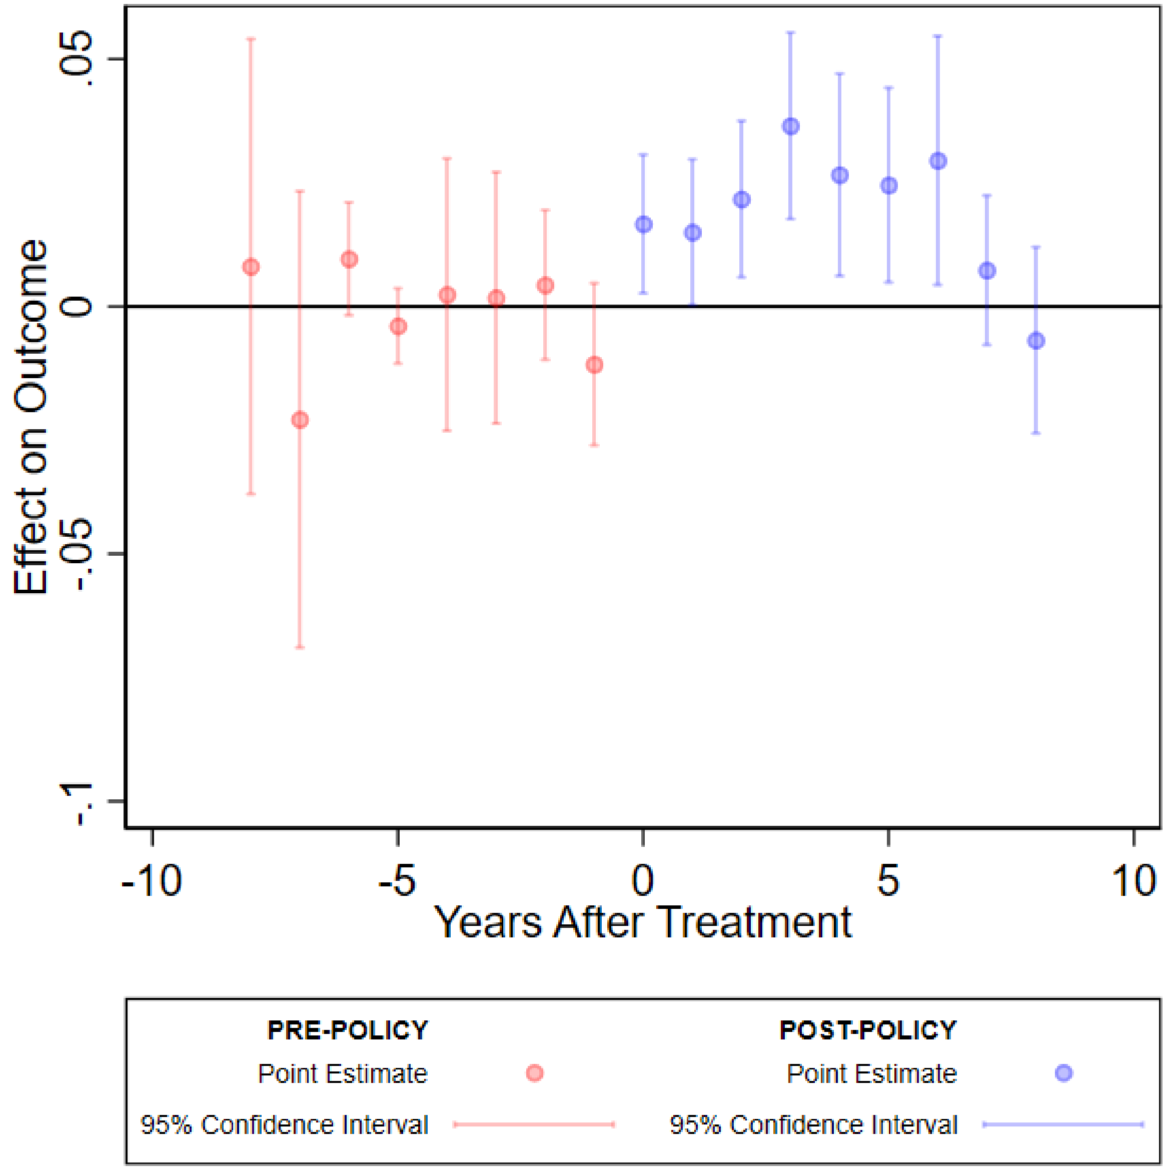

Figure B-12: Estimated ATT_t for the District's Bachelor's Degree Attainment by Years After Treatment (t), Conditioned Model

Model is conditioned on the district's percentage of students who are Black, Hispanic, Asian, Native American, and English Language Learners

C. APPENDIX FOR CHAPTER 4 TABLES AND FIGURES

Table C-1: Descriptive Statistics for Dependent Variable

	Mean	**SD**	**Median**	**Min**	**Max**	**N**
AFGR						
Treatment	0.80	0.27	0.76	0.00	1.00	417
Potential Control	0.82	0.18	0.82	0.00	1.00	12,158
Matched	0.75	0.17	0.77	0.00	1.00	417

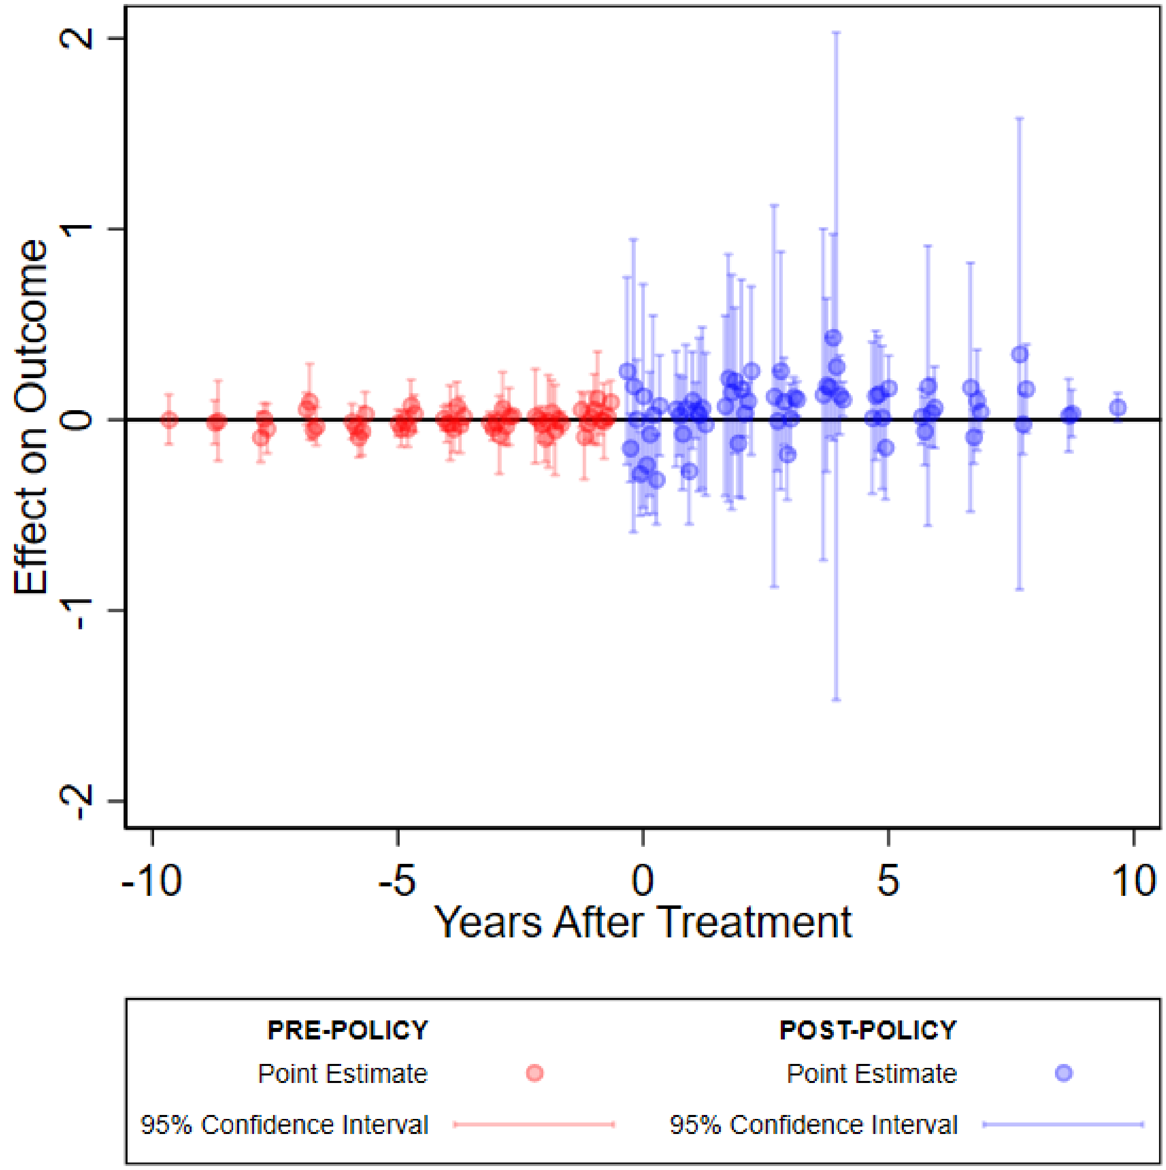

Figure C-1: Estimated ATT_{gt} for the District's Average Freshmen Graduation Rate by Years After Treatment (t) and Implementation Cohort (g)

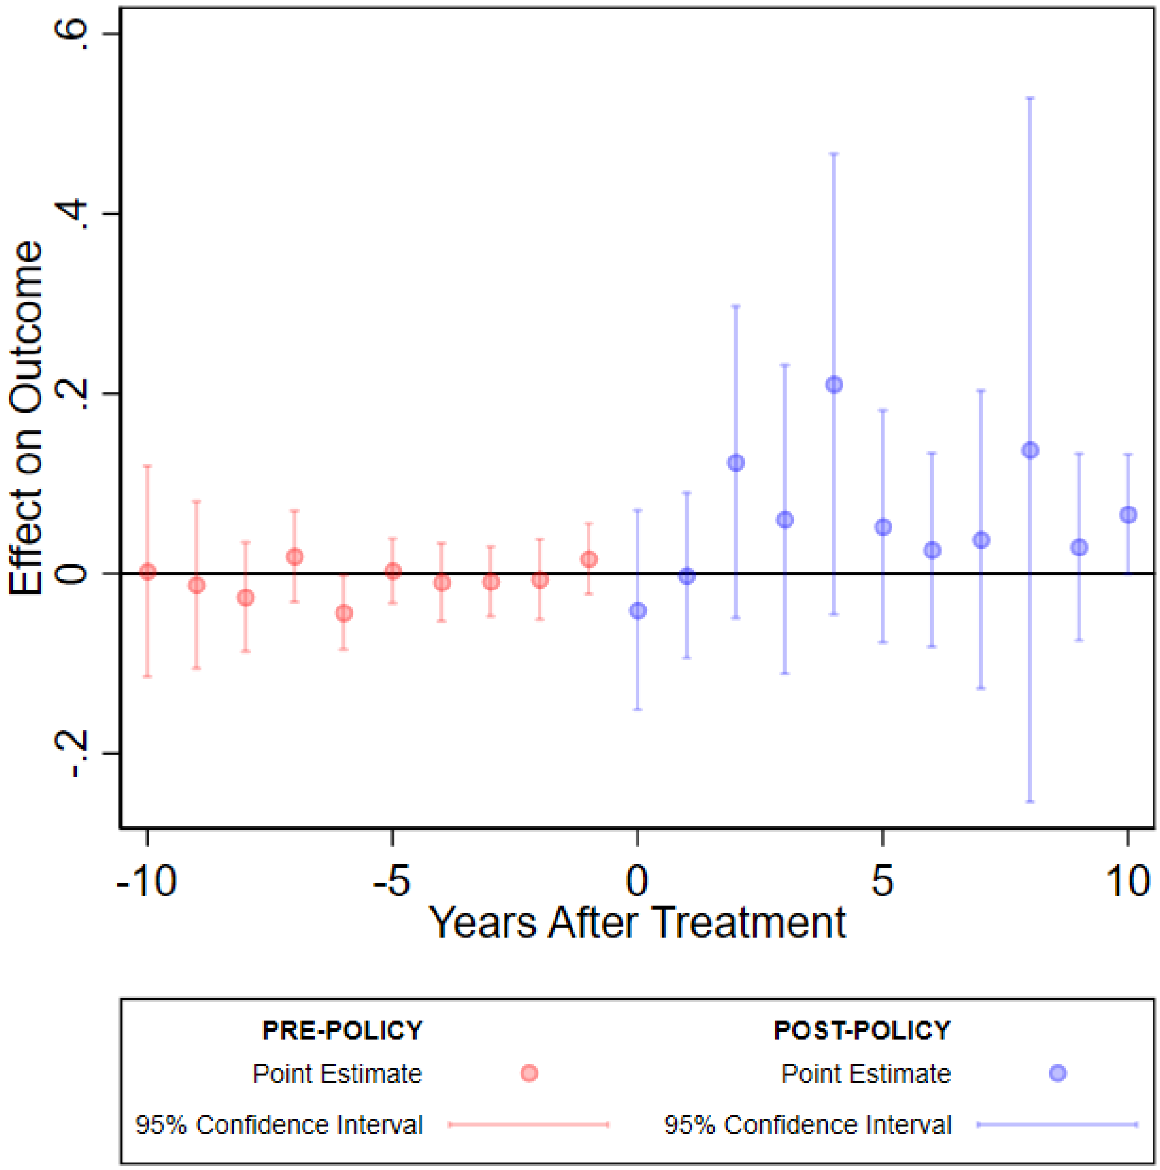

Figure C-2: Estimated ATT_t for the District's Average Freshmen Graduation Rate by Years After Treatment (t)