AN ESSAY ON TOPICS IN APPLIED MICROECONOMICS

A Dissertation

by

MERADEE TANGVATCHARAPONG

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	Mark Hoekstra
Committee Members,	Jonathan Meer
	Andrew Barr
	Kalena Cortes
Head of Department,	Steven Puller

May 2021

Major Subject: Economics

ABSTRACT

This dissertation examines three topics in the field of applied microeconomics using quasi-experimental methods.

In the section entitled "Does Being Tracked with Better Peers Matter?: Regression Discontinuity Evidence", I study a popular education policy of tracking students into classrooms. Although tracking is widely used around the world, relatively few papers directly identify the impact of being tracked into classrooms with higher- or lower-ability peers on student achievement. Using administrative data from Thai middle schools with a regression discontinuity design, I show that being tracked into classrooms with higher-ability students alone does not lead to significant increases in student achievement.

In the section entitled "The Impact of Misinformation: Evidence from the Anti-Vaccination Movement in the US", I study the effects of misinformation about vaccination, which is an important public health tool. Specifically, I examine whether the dissemination of false information about the MMR vaccine changed people's immunization behavior. Using a difference-in-differences approach with individual-level vaccination records from the National Immunization Surveys, I estimate that the rise in misinformation about the MMR vaccine caused the MMR vaccine take-up rates at 15 and 29 months old to decrease by 3.3 and 4.1 percentage points, respectively.

In the section entitled "An Empirical Test of Anti-Muslim Bias: Evidence from Property Values", we propose a test for anti-Muslim bias by asking whether the introduction of a mosque in a neighborhood reduces property values. To do so, we link administrative data on property sales prices in Michigan to the opening dates of new mosques. We then compare sales prices over time for properties closer and farther away from newly-opened mosques. Results indicate no evidence of anti-Muslim bias. Estimates from repeat-sales specifications indicate the introduction of mosques to a neighborhood does not depress property prices.

DEDICATION

To my parents, Ubol and Somchai Tangvatcharapong, for their love, support, and encouragement.

ACKNOWLEDGMENTS

I am incredibly grateful to my adviser, Mark Hoekstra, for his continuous support and guidance throughout my time at Texas A&M. Thank you for all the countless hours you spent mentoring me and helping me to become the researcher that I am today. I also would like to thank the members of my dissertation committee: Jonathan Meer, Andrew Barr, and Kalena Cortes, for their valuable research advice and continuous support and encouragement.

I am deeply grateful to the Economics Department Heads, Tim Gronberg and Steve Puller, for their unparalleled support, especially in 2020 after my post-graduation plan was affected by COVID-19. I am also grateful to the Department of Economics staff, especially Chelsi Bass and Teri Tenalio. Thank you for your support, patience, and kindness. I have always felt extremely well taken care of at Texas A&M.

Special thanks to my friends and colleagues in the program especially, Abigail Peralta, Adam Bestenbostel, Andrea Kelly, Brittany Street, CarlyWill Sloan, Josh Witter, Manuel Hoffmann, and Moffii Odunowo. You guys have inspired me daily and have always been there for me, both professionally and personally. Thank you. Thanks also to my friends and former colleagues from back home. Thank you for all your Facetime calls and text messages. On my toughest days, I can always count on you all to make me feel better.

Lastly, I am grateful to my family. To all my aunts, uncles, and cousins, thank you for all the free food and your words of encouragement. Mom and dad, thank you so much for equipping me with all the tools that I need to be successful, and for your unconditional love, support, and encouragement. I was able to spend six years in graduate school without having to worry about anything because I always knew that no matter what happened, you would always be there for me.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supervised by a dissertation committee consisting of Professor Mark Hoekstra and Professors Jonathan Meer and Andrew Barr of the Department of Economics and Professor Kalena Cortes of the Bush School of Government and Public Service.

Section 4 is joint work with Abigail Peralta. All other work conducted for the dissertation was completed by the student independently.

Funding Sources

Graduate study was supported by assistantships from Texas A&M University and a summer fellowship from the Private Enterprise Research Center at Texas A&M University.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE Page

LIST OF TABLES

1. INTRODUCTION

The aim of this dissertation is to answer policy-relevant questions empirically using quasi-experimental methods. This is because in order to design effective and efficient public policy it is important to identify the right problem and the effectiveness of existing policies. The topics examined in this dissertation include student tracking, misinformation about vaccines, and religious discrimination.

Section 2 studies tracking, which is a popular education policy of sorting students into classrooms. However, despite the popularity of tracking, there is still debate over the merits of tracking students by ability. Of particular concern is whether tracking could harm students who were tracked into lower ability classrooms through increased exposure to lower-ability peers. Unfortunately, there is little direct evidence on whether these differences in peer exposure induced by tracking cause subsequent differences in student achievement. This section answers this question by estimating the effects of being tracked into higher-ability classrooms in a setting where only peer quality changes, thereby separating the effects of higher-ability peers from other confounding factors. Using a regression discontinuity approach with administrative data from public middle schools in Thailand where students are sorted into classrooms based on ability, I show that being tracked into classrooms with higher-ability peers does not improve student GPA.

Section 3 studies whether the surge in misinformation about vaccines changed immunization behavior. In recent years, we have observed increasing number of misinformation. The increasing amount of misinformation and its potential consequences has thus generated significant debate about the government's role in regulating misinformation and social media platforms' responsibility to fight it. However, little is known about the actual impact of misinformation on behavior. This section answers this question by asking whether the dissemination of misinformation about vaccines leads to meaningful changes in immunization behavior. To do so, I examine how vaccination rates responded to the unexpected surge in media coverage in 2007 of the verifiably false claim that the Measles-Mumps-Rubella (MMR) vaccine caused autism.

Specifically, I use a difference-in-difference approach to compare the MMR vaccination rates of children whose parents were most and least likely to be affected by the news over time. I determine parents' susceptibility using three predetermined characteristics: whether their child is a firstborn, the child's gender, and the parents' age. Results show that susceptible parents were 3.3 percentage points less likely to vaccinate their children with an MMR shot by the recommended age of 15 months and 4.1 percentage points less likely to do so by 29 months. This indicates that at a minimum, misinformation about vaccine safety caused parents to delay vaccinating their children by over a year, and at most prevented them from ever immunizing their children.

Section 4 proposes an empirical test for anti-Muslim bias by asking whether the introduction of a mosque in a neighborhood reduces property values. The Muslim population in the US is growing at a fast pace. Unfortunately, the growth of the Muslim population has also been associated with a perceived rise in anti-Muslim sentiment. While all this suggests that anti-Muslim bias is increasing over time, as with all forms of bias it is difficult to provide definitive evidence. We offer a new test of anti-Muslim bias by asking whether the introduction of a Mosque in a neighborhood leads to a reduction in housing prices. In the presence of substantial anti-Muslim bias, we would observe a reduction in sales prices. We link administrative data on property transactions in Detroit and Hamtramck, Michigan, to the opening dates of all places of worship in those two cities. We address concerns about the potential endogeneity in the location of mosques by exploiting the granular nature of our data, which allows for a spatial difference-in-differences research design using repeat-sales of the same properties. In particular, we compare prices for properties that are very close to a new mosque to those that are slightly farther away. Our results show that new mosque openings do not results in a decrease in property prices. Additionally, the effects of mosque openings and church openings on property prices are also similar. The finding here suggests that even if residents or prospective buyers choose to avoid the neighborhood as a result of anti-muslim bias against a new mosque, it appears that this response is not strong enough to offset demand for these properties.

2. DOES BEING TRACKED WITH BETTER PEERS MATTER?: REGRESSION DISCONTINUITY EVIDENCE

2.1 Introduction

Across the world, a common educational practice is to track students into different classrooms based on ability. For example, the OECD reports that 95 percent of students in the United Kingdom, Ireland, New Zealand, Australia, Israel, Albania, Kazakhstan, Singapore, Russia, and Malaysia attended schools where students were grouped by ability across classrooms (OECD, 2013). While student tracking is generally less common in the U.S., it is still widely used in some areas. Per the National Assessment of Education Process, 75 percent of U.S. schools track students by ability for 8th-grade mathematics (Loveless, 2013). However, despite the popularity of tracking, there is still debate over the merits of tracking students by ability. Of particular concern is whether tracking could harm students who were tracked into lower ability classrooms through increased exposure to lower-ability peers. Unfortunately, there is little direct evidence on whether these differences in peer exposure induced by tracking cause subsequent differences in student achievement. This is largely due to the fact that many programs that assign higher-ability students to different classrooms, such as gifted and talented programs, also often expose those students to different classroom settings, such as more intensive curriculum and higher-quality teachers. This makes it difficult to identify whether and how much the most common change induced by tracking, namely changes in peer quality, affects student achievement. The purpose of this paper is to estimate the effects of being tracked into higher-ability classrooms in a setting where only peer quality changes, thereby separating the effects of higher-ability peers from other confounding factors. In doing so, I also speak to the underlying reasons for the mixed evidence in the literature. To emphasize, this paper does not directly speak to the effects of the tracking system itself, but rather the effects of being tracked into classrooms with higher-ability peers.

To this end, I apply a regression discontinuity design using administrative data from public middle schools in Thailand where many schools regularly sort students into classrooms based on ability. To measure student ability, these schools have students sit for a preliminary exam before the start of seventh grade. School officials then use the score from the preliminary exam as a proxy for student ability and sort students into classrooms based on this preliminary exam score. This allows me to exploit the resulting cutoffs between classrooms to employ a regression discontinuity approach that compares the academic outcomes of students just above and below the cutoffs to identify the effects of peer quality in the tracking system. There are several reasons why the institutional setting here is ideal for a regression discontinuity design. First, these cutoffs between classrooms are not known to the students until after the exam is taken and classrooms are assigned, making it difficult, if not impossible, to manipulate one's position relative to a classroom cutoff. Second, classrooms above and below the cutoffs are required to follow the same curriculum, take the same exam, and have very similar sets of teachers. As a result, this setting supports the identifying assumption that all determinants of achievement other than peer quality vary smoothly across the classroom cutoff. I provide empirical evidence supporting this assumption in the Empirical Strategy section.

For the analysis, I implement the aforementioned design using administrative student data from four public middle schools in Thailand. My data set contains the preliminary score, classroom assignment, GPA, classroom timetable, teacher assignment, and characteristics of 1,602 seventh-grade students. The main outcome of interest is the seventh-grade cumulative GPA. Importantly, GPA in Thailand is based primarily on student performance in multiple-choice exams for which there is no grade curving. Students in the same school also take the same exams regardless of their assigned classroom and teachers. As a result, there is little scope for teacher bias or subjectivity to affect GPA in this context.

Consistent with the classroom allocation mechanism, I first show that scoring just above the cutoff does increase the likelihood of being assigned to the higher-ability classroom by 80 percentage points. However, results indicate there is no discontinuity in student performance at the cutoff. This indicates that being tracked into classrooms with significantly higher-ability peers does not lead to higher achievement. Specifically, my 2SLS estimates show that assignment to a higher-ability classroom is associated with a 0.94 standard deviation increase in peer quality, as measured by performance on the seventh-grade preliminary exam. However, this exposure is associated with a statistically insignificant 0.08 standard deviation reduction in performance, as measured by GPA. Importantly, this finding is robust to the bandwidth size as well as the inclusion of student characteristics. Additionally, the result is also robust to the inclusion of teacher fixed effects. This is consistent with the fact that students have the same or similar teachers across classroom cutoffs, as well as the fact that the teachers in my data set do not systematically choose to only teach classrooms above the cutoffs, e.g. higher-quality teachers do not only teach classrooms just above the cutoffs.

In addressing the effects of being tracked into classrooms with higher-ability peers, this paper is most closely related to a paper by Vardardottir (2013). Vardardottir (2013) uses student data from Iceland to identify the impact of being tracked into classrooms with higher-ability peers. The major difference between this paper and Vardardottir (2013) is that I observe a clear, visually compelling discontinuity in the likelihood of being placed in a classroom with higher-achieving peers at the cutoff. In contrast, there is no such discontinuity in the data underlying Vardardottir (2013).[1] As a result, a major contribution of my paper is to identify the effects of being tracked into a higher-ability classroom using a clean regression discontinuity framework. This enables me to give estimates a causal interpretation under a reasonable identifying assumption. In addition, this paper is the first to identify the effect of peer quality shifts due to tracking in Asia, where tracking is very common.

In addition to providing estimates in a clean regression discontinuity framework, another advantage of my study is that I am able to rule out positive effects of modest size. For example, Duflo, Dupas, and Kremer (2011) performed a field experiment in Kenya that enabled a regression discontinuity study of the effects of being tracked into higher-ability classrooms on student achievement. Similar to this study, they reported no statistically significant effects from

[1]This is evident in the local averages shown in Figure 1 on page 115

an increase of one standard deviation in peer quality and ruled out effects larger than 0.21 standard deviations. By comparison, estimates in this study enable me to rule out effects of only 0.08 standard deviations.

This study is also directly related to the literature on the general effectiveness of tracked classrooms and gifted and talented programs (Bui, Craig, and Imberman, 2014; Card and Giuliano, 2016; Booij, Haan, and Plug, 2016, 2017; Cohodes, 2020). The results in this paper support the finding in Bui, Craig, and Imberman (2014) which used a regression discontinuity design to estimate the impact of gifted and talented programs from a large school district in the US and found that achievement does not improve for students placed in gifted and talented programs. The results here also speak to the finding in Card and Giuliano (2016) which used a regression discontinuity design to study the impact of gifted classrooms in a large school district in the US. They found that the impact of the gifted classrooms was minimal for white students but large for minority students. Since white and minority students experience the same curriculum, teachers, and peers, they concluded that the large effects on minority students were likely from the removal of low teacher expectations and negative peer pressure. As Thailand is a relatively homogeneous country, it is not surprising that the result in this paper is in line with the results of white students in Card and Giuliano (2016). Additionally, estimates in this study also enable me to rule out the effects of the magnitude found by some studies on gifted and talented programs. For example, Booij, Haan, and Plug (2017) reports that a gifted and talented program in the Netherlands increased student GPA by 0.2 standard deviations. My result here suggests that a large portion of the positive effects found in Booij, Haan, and Plug (2017) are likely due to specific features of the gifted and talented program, such as the curriculum and teacher quality, rather than the change in peer quality.

As attending higher-quality schools is often associated with an increase in peer quality, this paper also speaks to the literature on returns to school quality. The results in this paper again suggest that the positive effects found by some studies (Shi, 2019; Pop-Eleches and Urquiola, 2013; Park, Shi, Hsieh, and An, 2009) could perhaps be attributed more to features of the schools other than higher peer quality in the classroom. It also provides a plausible explanation as to why many papers (Allensworth, Moore, Sartain, and de la Torre, 2017; Dee and Lan, 2015; Lucas and Mbiti, 2014; Dobbie and Fryer Jr, 2014; Abdulkadiroğlu, Angrist, and Pathak, 2014; Clark, 2010) have found selective schools to have little effect on student achievement. Indeed, the findings of this paper would predict that same result, unless selective schools also offered better teachers or other input into education production. In this way, this paper also complements the finding in Hoekstra, Mouganie, and Wang (2018) and Jackson (2013) that the returns to high school quality are likely the result of features of the schools other than peer quality. Additionally, this paper also complements the large literature on peer effects more generally.[2]

The results of this paper also have important implications for both parents and policymakers. Recent evidence suggests that in choosing schools, parents put much weight on peer quality (Abdulkadiroğlu, Pathak, Schellenberg, and Walters, 2020). However, results shown in this paper suggest that parents would be better off making decisions on factors other than peer quality, such as teacher quality or curriculum. Similarly, they also suggest that educators and policymakers should put more emphasis on other factors believed to improve student performance and less emphasis on the role of peer composition in the classroom. More importantly, these results have direct implications for school tracking. Specifically, they suggest that an evaluation of tracking should focus more on the effects it has on teaching and curriculum, and less on whether some students are left disproportionately exposed to higher- or lower-ability students.

2.2 School Tracking in Thailand

Thailand has a 6-3-3 school system where students attend elementary school for 6 years, middle schools for 3 years, and then high school for 3 years. Typically, public schools only provide either primary education or secondary education. This means that the majority of Thai students have to start at a new school when they transition from primary education to secondary education in seventh grade.

[2]For example, see Hoxby (2000), Lefgren (2004), Lavy and Schlosser (2011), Ohinata and Van Ours (2013), Sacerdote (2001), Zimmerman (2003), Carrell, Fullerton, and West (2009).

At public middle schools, the practice of sorting students into classrooms based on ability, or 'tracking', is common. Many schools ask the newly-enrolled seventh-grade students to sit for a preliminary exam before the school year starts. Schools then use the results of this preliminary exam as a proxy for student ability and then sort students into classrooms based on the preliminary exam score. For example, in a school where there are 120 seventh-grade students, the 40 students who scored the highest in the preliminary exam (rank 1-40) are normally sorted into class 1, the next 40 students (rank 41-80) are sorted into class 2, and the bottom 40 students (rank 81-120) are grouped together into class 3. Students who are assigned to the same class stay together in that class for at least a whole school year. Specifically, this means that they sit in the same classroom, follow the same timetable, and take all the same courses from the same teachers.

Importantly, one unique feature of the Thai schools in this paper is the fact that students in classrooms above and below the cutoff also take the same courses, follow the same curriculum, and take the same exams. And while the sets of teachers assigned to classrooms above and below the cutoffs might not always be completely identical, they are quite similar. This is because teachers teach more than one classroom within a grade, which means that most classrooms above and below the cutoff have the same teacher for that subject.[3] As a result, the only thing changing at the classroom cutoff here is essentially the level of student ability in the classroom. This setting thus allows me to apply a regression discontinuity approach to identify the impact of solely peer quality in the tracking system. This is because comparable students who are just above and below the threshold are assigned to different classrooms that are the same in all aspects except for peer quality.

2.3 Data

The analyses in this paper use administrative data of students who were enrolled in the seventh grade in four public middle schools in Bangkok between 2013-2014 and 2016-2017. The data set consists of students' preliminary exam scores, class assignment, timetable, teachers assigned, GPA, and student characteristics, which include gender, height, weight, class size, birth order, and

[3] I also account for any difference in teacher assignment by including teacher fixed effects.

parents' marital status.

These four schools are all public secondary schools from the suburban area of Bangkok and were not chosen with any ex-ante presumptions. Rather, they were the only schools that kept complete records of the class sorting criteria and also allowed me access to the administrative data. For all schools, I checked each classroom's timetable to see if all the classrooms in the same grade, especially the classrooms just above and below the same cutoff, follow the same curriculum. I found that while it is true that students in all classrooms take the same core courses, the curriculum is more flexible for non-core subjects, such as physical education (PE). For example, there are instances where students in all classrooms take PE, but different classrooms take different PE courses. In other words, although all classes take PE, some classrooms have basketball, while some classrooms have volleyball. This difference is likely due to the fact that schools do not have enough equipment and teachers to allow all students to take the same non-core courses. Therefore, while students in all classrooms still follow the same curriculum and take the same required number of non-core courses in each semester, the non-core courses they take are sometimes different. I, therefore, limit my sample to only the cutoffs where the classrooms above and below follow the same identical courses, so that the only thing changing at the cutoff is student quality. As a result, my analysis sample consists of 10 cutoffs and 1,602 students.[4]

The main outcome of interest in this paper is seventh-grade cumulative GPA. Importantly, GPA in Thai middle schools is based primarily on performance on multiple-choice exams for which there are no curves. Students who take the same course in the same school also take the same exact exams even when they are in different classrooms and are taught by different teachers. As a result, in contrast to other contexts, there is little scope for teacher subjectivity to affect student grades and GPA. Anecdotally, 10 percent of the final grades could be subject to teacher discretion, which is usually based on student attendance and attentiveness in class. Many students

[4]For one of the ten cutoffs, there are five classrooms below the cutoff and students who score below the cutoff are randomly assigned into one of the five classrooms. Out of these five classrooms, four classrooms follow the exact same curriculum as the classrooms above the cutoff, while one classroom does not have the same non-core courses as the classroom above the cutoff. There are 51 students in this particular classroom and I drop them from my sample. Since students are randomly assigned to this classroom, this should not affect results.

receive full points and very few students receive less than 5 percent out of this 10 percent portion. Since different schools in different school years could have different standards for GPA, I use standardized cumulative GPA instead of raw cumulative GPA. I standardize cumulative GPA by rescaling within each school and school year so that the mean of the standardized cumulative GPA is zero and the standard deviation is one.

Table A.1 summarizes the characteristics of students in my data set. First, Column 1 reports the descriptive statistics for all the students in the data set used for the analysis. Columns 2-4 report the same statistics, but limit the sample to only students closer to the cutoff. Specifically, Column 2 reports the descriptive statistics of students whose preliminary exam score is within 20 points from the cutoff. Column 3 and Column 4 report the same statistics of students whose preliminary exam score is within 10 points and 5 points from the cutoff, respectively. Based on Table A.1, the average preliminary exam score of the student in the sample is 48 percent. This is not surprising as the schools in the dataset are not selective schools. The average seventh-grade cumulative GPA is 2.9. Columns 2-4 also show that 96 percent of students in the sample have a preliminary exam score within 20 points of the cutoff, while 65 percent are within 10 points and 40 percent are within 5 points.

2.4 Empirical Strategy

To disentangle the effects of peers from confounding factors, I apply a regression discontinuity design (RDD) that compares students just above and below the cutoff. The key assumption is that all other determinants of the outcomes except peer quality vary smoothly across the threshold.

Since I have multiple cutoffs, each with different cutoff scores, I follow the method used in Pop-Eleches and Urquiola (2013) and employ the stacked RDD method. Specifically, I first normalize the cutoffs using equation 2.1. Then I pool all normalized cutoffs together for the regression discontinuity analysis.

$$r_{ic} = prelim_i - cutoff\ score_c \tag{2.1}$$

In equation 2.1 , the normalized preliminary exam score of student i from cutoff c is denoted by r_{ic}. $Prelim_i$ is student i's raw preliminary exam score and $cutoffscore_c$ denotes the cutoff score of cutoff c. Using equation 2.1, all the cutoff scores are recentered to zero. The normalized preliminary exam score (r_{ic}) indicates how far each student is from their associate cutoff as well as whether they are above or below the cutoff. The number is positive for those above the cutoff and negative for those below the cutoff. The formal regression discontinuity analysis in this paper then use the following standard regression discontinuity model:

$$Y_{ic} = \gamma_1 r_{ic} + \beta I[r_{ic} \geq 0] + \gamma_2 r_{ic} I[r_{ic} \geq 0] + \delta_c + \beta_x X_i + u_{ic} \tag{2.2}$$

Where Y_{ic} is the outcome variable of student i at cutoff c. r_{ic} is the model's running variable, which is student i's normalized preliminary exam score. $I[r_{ic} \geq 0]$ is a binary variable indicating whether student i is above the cutoff. δ_c represents a full set of cutoff dummies. X_i is a matrix containing student i's characteristics including class size, gender, height, weight, birth order, parents' relationship status. Importantly, the coefficient of interest here is β which indicates whether there is a discontinuity in the outcome (Y_{ic}) at the cutoff.

One important thing that should be noted here is that since each school could have multiple cutoffs in a school year, it is possible that some students are associated with two cutoffs at the same time. For example, from the example earlier where there are three classrooms in the seventh grade, students who are in class 2 are associated with two cutoffs: the one separating class 1 and class 2, and the one separating class 2 and class 3. The students who are associated with two cutoffs thus appear in the data set twice and have two different normalized preliminary scores calculated based on each of their two different cutoffs. Due to these repeated observations, I cluster standard errors at the individual level.

2.4.1 Test of Identification

As with any RD design, the key identification assumption here is that students just above and below the cutoff are comparable in the absence of treatment. I will be able to accurately estimate the impact of peer quality only if students just above and below the cutoff are comparable and the only things changing at the cutoff are their class assignment and the resulting peer quality. Under this assumption, any discontinuity in student achievement at the cutoff can be properly attributed to the increase in peer quality. In this section, I provide support for this approach by providing empirical evidence consistent with the identifying assumption.

To this end, I start by checking that students could not manipulate the cutoff. This is important because if some students could strategically place themselves just above the cutoff then it would mean that students just above and below the cutoff are fundamentally different. For example, one might worry if particularly motivated students were able to obtain scores just above the cutoff. Institutionally, there is no reason to believe that students would be able to manipulate their position relative to the cutoff. First and foremost, since the cutoff score was not known to the students before the preliminary exam, it would be difficult, if not impossible, for students to precisely predict where the cutoff will be and put in just the right amount of effort as to place themselves just above the cutoff. Moreover, there is also no retake of the preliminary exam. To provide further support for this institutional claim, I examine the distribution of students' normalized preliminary exam scores. If students could precisely manipulate their position relative to the cutoff, we would see a jump in the density of students at the cutoff. Figure A.1 shows that the distribution of students' normalized preliminary score is smooth across the cutoff. The data are therefore consistent with my understanding of how students are assigned to classrooms and suggest no evidence of manipulation around the cutoff.

In addition, I also test whether the observable characteristics of students are smooth across the threshold. If the identifying assumption holds, all characteristics should be smooth across the cutoffs. If students could manipulate the threshold, we might observe a discontinuity at the cutoff for some characteristics. Here, I look at all the observable characteristics available in the data set including gender, weight, height, birth order, parents' marital status, and class size. In Figure A.2, I show graphically that all characteristics are smooth across the cutoff. I then formally estimate the discontinuity of each covariate at the cutoff using the model described in the last section. The regression discontinuity estimates are reported in Table A.2 and again confirm that student characteristics are smooth across the cutoff.

Additionally, rather than focusing on each of the characteristics individually, I also use these observable characteristics to predict seventh-grade cumulative GPA for each student. I then look at whether these predicted GPAs are smooth at the cutoff. The benefit of this method is that it allows me to attribute appropriate weight to each characteristic according to how much it contributes to student GPA. Figure A.3 shows visually that there is no discontinuity in the predicted GPA at the cutoff. This again suggests that students just above and below the cutoff are comparable and that there is no manipulation of the threshold.

One limitation of the data is that I do not observe seventh-grade cumulative GPA for roughly 14 percent of the students in my data set. This is because the schools did not provide me with the records of students who had transferred to another school or dropped out. This could potentially bias my estimates if there is selective attrition across the cutoff. To assess this, I test for a discontinuity in the probability of being observed with seventh-grade cumulative GPA across the cutoff and show that there is no such discontinuity. Results are shown in Figure A.10 and Table A.6 which follow the tables and figures of the main analysis. In addition, I also check for discontinuities in student characteristics and predicted GPAs again using only the students for whom I observe the main outcome, i.e. seventh-grade cumulative GPA. The results hold and confirm that there is no discontinuity in student characteristics at the cutoff. These results are shown in Figure A.11, Figure A.12, and Table A.7, again following the tables and figures of the results from the main analysis.

Based on all the evidence shown in this section, I conclude that students observed in the sample on either side of the cutoff are comparable. This is consistent with the identifying assumption, and with the institutional background that suggests manipulation would be difficult, if not impossible, in this context. As a result, there is little reason to expect that student outcomes would be different on either side of the cutoff, absent the effect of being tracked with higher-ability peers.

2.5 Results

2.5.1 The Discontinuity in Classroom Assignment

First, I examine the first-stage relationship between students' normalized preliminary exam scores and their class assignments. Specifically, I examine how crossing the classroom cutoff affects students' probability of being in the higher-ability classroom. Figure A.4 shows visually that the probability of students being in the higher-ability classroom jumps from approximately 0 to 80 percent when they cross the cutoff. The reason why the compliance rate is not jumping from precisely 0 to 100 percent at the cutoff is that there are students who received special treatment and students who opted out of the assigned classrooms.[5]

I formally estimate and report the discontinuities in the probability of being tracked into higher-ability classrooms at the cutoff in Panel 1 of Table A.3. The odd-numbered columns show the estimates from the regression without any controls, while the even-numbered columns show the estimates from the regression with controls for student characteristics. Columns 1-2 show the estimates from the regression using the full sample, while Columns 3-8 report the estimates from when the sample is only limited to students closer to the cutoff. The estimates reported in this panel range from 0.74-0.86 and all are statistically significant at conventional levels. In addition, across all bandwidths, the estimates change little as controls are added, consistent with the identifying assumption.

2.5.2 The Discontinuity in Peer Quality

Next, I turn my attention to peer quality. In this section, I examine whether crossing the cutoff and therefore having a higher chance of being in the higher-ability classroom is associated with higher quality peers. I measure each student's peer quality by calculating the average of their

[5] In one of the schools, students could choose to opt-out of their assigned classroom and enroll in the 'gifted' classroom if they could pay the higher tuition of the 'gifted' classroom. I leave the 51 students who were enrolled in the 'gifted' classroom in this school in the sample in order to avoid selection bias due to their exclusion, as the decision to switch could depend on which side of the threshold they were on. In Table A.8, which is shown after the main results, I show that the decision to control or not control for these gifted classrooms in the regression does not affect my results.

classmates' standardized preliminary exam scores.[6]

Figure A.5 shows graphically that peer quality jumps by approximately 0.7 standard deviations at the classroom cutoff. The formal estimates are reported in Panel 2 of Table A.3. They show that corresponding to the increase in the probability of being in the higher-ability classroom, peer quality jumps by 0.70-0.82 standard deviations at the cutoff. Again, my estimates are stable across bandwidth sizes and robust to the inclusion of student characteristic controls.

2.5.3 Reduced-Form Estimation: Effects on Seventh-Grade GPA

In the previous section, results indicate that crossing the cutoff is associated with an increase of approximately 0.70-0.82 standard deviations in peer quality. In this section, I examine whether this could, in turn, lead to an increase in academic performance, as measured by GPA. If it does, because peer quality is the only thing changing at the cutoff, it would suggest that crossing the cutoff increases student academic achievement through improvement in peer quality.

Figure A.6, which plots the relationship between students' normalized preliminary score and standardized GPA, graphically shows this reduced-form relationship. From Figure A.6, it is clear that there is no discontinuity in student GPA at the cutoff. This suggests that crossing the cutoff, and therefore having higher-quality peers, does not lead to better student outcomes.

I formally estimate the discontinuity in student GPA at the cutoff by estimating the model in equation 2.2 with standardized seventh-grade cumulative GPA as the outcome variable. The estimates are shown in Table A.4. They are all statistically insignificant at conventional levels and range from -0.09 to -0.11 (Columns 1, 4, 7, 10). When I also include characteristic controls in my specification, across bandwidth sizes, the estimates change little. They are still statistically

[6]The standardized preliminary score of student j who is a 7th-grade student in school s in school year y is calculated using

$$standardized\, prelim_j = \frac{prelim_j - mean\, prelim_{sy}}{s.d.\, prelim_{sy}}$$

and i's peer quality is calculated using

$$peer\ quality_{ic} = peer\ quality_i = \frac{1}{n_{class(i)} - 1} \sum_{j \neq i, j \in class(i)} standardized\, prelim_j$$

insignificant at the conventional levels and range from -0.08 to -0.12 (Columns 2, 5, 8, 11). This suggests that the finding of no positive effects are robust to the inclusion of controls and bandwidth sizes.

However, one might be concerned about teacher quality across classrooms. For instance, if the classrooms above the cutoff always get worse teachers, then my estimates of peer effects could be biased downward. As mentioned before, institutionally, this should not be an issue as most of the teachers in the data set teach both the classrooms above and below the cutoff. Nevertheless, I address this issue empirically by adding teacher fixed effects to my specification. The estimates from this specification with teacher fixed effects become a little more negative and range from -0.11 to -0.19 (Columns 3, 6, 9, 12), but are still statistically insignificant at conventional levels. This suggests that if anything, students with higher ability peers may have access to higher-quality teachers, causing my unconditional estimates to be an upper bound. I emphasize, however, that the estimates without and with teacher fixed effects are not statistically different from each other.

In any case, since the estimates across specifications and bandwidth sizes are negative and statistically insignificant, the important thing we could take from the results is that being tracked into classrooms with higher-ability peers does not lead to significantly higher achievement for students. Importantly, the majority of the upper bounds of the 95 percent confidence intervals, which are also shown graphically in Figure A.7, indicate that the effect of crossing the cutoff and therefore having higher-ability peers is not greater than 0.07 standard deviations.

2.5.4 2SLS Estimates

Next, in Table A.5, I report local average treatment effect (LATE) estimates of being tracked into higher-ability classrooms using 2SLS. Intuitively, these estimates are the reduced-form estimates divided by the increase in the likelihood of attending the higher-ability classroom at the cutoff as shown in Panel 1 of Table A.3. Estimates from Panel 1 of Table A.5 indicate that peer quality increases by approximately 0.94 standard deviations when students are tracked into higher-ability classrooms. At the same time, Panel 2 of Table A.5 reports that being tracked into higher-ability classrooms and therefore having peers that are 0.94 standard deviations better is associated with a statistically insignificant 0.10-0.16 standard deviation decrease in student GPA. Additionally, Panel 3 of Table A.5 rescales the estimates and shows that an increase of one standard deviation in classroom peer quality results in a statistically insignificant decrease in student GPA of 0.10-0.18 standard deviations.

In addition, Figure A.8 plots the LATE estimates of being in higher-ability classrooms on student achievement (seventh-grade cumulative GPA) along with their 95 percent confidence intervals across bandwidth sizes. We can see that the estimates are all negative, statistically insignificant, and relatively stable across bandwidth sizes. Importantly, more than 80 percent of the upper bound estimates across bandwidth sizes are smaller than 0.08 standard deviations. This enables me to rule out any positive effects bigger than 0.08 standard deviations.

While Figure A.8 plots the LATE estimates of being in the high-ability classrooms which are associated with an increase of 0.94 standard deviations in peer quality, Figure A.9 shows the LATE estimates of an increase of one standard deviation in peer quality across bandwidth sizes. Because of the large first-stage discontinuity, the estimates in Figure A.9 are very similar to those in Figure A.8. They are all negative and statistically insignificant and the upper bounds suggest that an increase of one standard deviation in classroom peer quality could not lead to an increase in student achievement that is larger than 0.08 standard deviations. To summarize, results from 2SLS estimations indicate that being tracked into better classrooms is associated with an increase of 0.94 standard deviations in peer quality. However, that increase in peer quality does not lead to positive effects on GPA, as point estimates are negative and I am able to rule out positive effects larger than 0.08 standard deviations.

2.6 Discussion

The absence of a positive effect for students tracked into classrooms with significantly higher-achieving peers seems puzzling. There are multiple possible interpretations of this finding. One is that perhaps there are positive effects for one group that are offset by negative or null effects for another group. To investigate this possibility, at least as it relates to perhaps the most salient difference—student gender—I look at peer effects on male and female students separately. The results, shown visually in Figure A.13, suggest that the impacts of peer quality are similar for both genders and that there are no positive peer effects for either male or female students. And while some of the formal estimates, shown in Table A.9, are statistically significant, they are marginally significant and not robust. Importantly, the underlying plots shown in Figure A.13 show little evidence that there is any discontinuity in student GPA at the cutoff for either gender. Therefore, it seems highly unlikely that I do not detect effects because of this reason.

While the most obvious interpretation of these findings is that exposure to higher-achieving peers does not benefit students, it is also possible that any benefits from that exposure are offset by other differences. For example, when a student is tracked into a classroom with higher-ability peers, they also automatically become a small fish in a large pond and has a lower rank in the classroom. This means that the impact of being tracked into a higher-ability classroom captures the net effect of increased exposure to high-ability peers, but also lower relative rank. Murphy and Weinhardt (2020) have looked into the effects of ordinal rank on student achievement and found large effects. These effects could potentially offset any positive effects from exposure to higher-achieving peers. I note that my setting is not unique in this sense; any policy that increases one's exposure to higher-achieving peers will also mechanically lower rank. In theory, I could untangle the two effects by looking at the heterogeneous effects across cutoffs. For example, I could compare the estimates at cutoffs with big increases in peer quality to the estimates from cutoffs with a small increase in peer quality, but both of which include similar effects on rank. Unfortunately, I do not have enough data and heterogeneity across cutoffs to do so in a constructive way. As a result, in this paper, I do not attempt to separate the two effects, but instead identify the reduced-form policy-relevant effect of being tracked into a classroom with higher- ability peers.

Additionally, the change in peer quality could also affect students through the change in teacher behavior. Specifically, teachers might tailor their instruction to the quality of the students in each classroom and therefore teach students in the high- and low- ability classrooms differently. If that is the case, it is possible that the change in teacher instruction affects students in the way that it offsets any positive effects from exposure to high-ability peers. In this paper, since I do not observe or have information on teacher instruction, I do not attempt to separate these indirect peer effects from the direct effects of having high-achieving peers. Instead, as stated earlier I focus on identifying the reduced-form policy-relevant effect of being tracked into a classroom with higher-ability peers.

2.7 Conclusion

This paper estimates the impacts of being tracked into classrooms with higher-achieving peers on student achievement using administrative data from public middle schools in Thailand. Using an RDD approach, reduced-form results show that crossing the classroom cutoff is associated with a large increase of approximately 80 percentage points in the likelihood of being assigned to the higher-ability classroom. This, in turn, translates to an increase of 0.7-0.8 standard deviations in peer quality at the cutoff. However, the increase in peer quality at the cutoff does not lead to an increase in student achievement as the seventh-grade cumulative GPAs remain smooth across the cutoff. Two-stage least squares estimates indicate that being in a higher-ability classroom is associated with a 0.94 standard deviation increase in peer quality, and results in a statistically insignificant 0.10-0.16 standard deviation reduction in student GPA.

Importantly, my finding of no positive effects is robust to bandwidth size and the inclusion of student characteristic controls and teacher fixed effects. In addition, upper bound estimates also allow me to rule out positive peer effects of modest sizes. Specifically, my upper-bound estimates indicate that the effects of a significant increase of one standard deviation in peer quality on student GPA, at least in Asian contexts similar to this, could not be larger than 0.08 standard deviations.

These results are in line with the findings in Duflo, Dupas, and Kremer (2011), which also found that an increase of one standard deviation in peer quality in classrooms in tracking schools leads to no statistically significant increase in student achievement. However, the strength of this paper is that I am able to rule out effects larger than 0.08 standard deviations, while Duflo, Dupas, and Kremer (2011) were only able to rule out effects larger than 0.21 standard deviations. My findings also enable me to speak to the literature on the effectiveness of tracked classrooms, such as gifted and talented programs, in general. My result complement the findings in Bui, Craig, and

Imberman (2014); Card and Giuliano (2016); Cohodes (2020) which estimated few and insignificant test score effects of special classrooms. And while Booij, Haan, and Plug (2017) found effects of 0.2 standard deviations of a gifted and talented program in the Netherlands, my results suggest that exposure to significantly higher-quality peers could not increase student achievement more than 0.08 standard deviations. Thus, it seems likely that a large portion of the positive effects found in Booij, Haan, and Plug (2017) are the results of specific features of the program.

In addition, my estimates also rule out effects that are small relative to previous papers on the benefits of attending higher-quality schools, which also have better peers. For example, Pop-Eleches and Urquiola (2013) studied Romanian secondary schools and found that attending higher-quality schools where peers are on average 0.1 standard deviations higher in quality results in an increase of 0.02-0.10 standard deviation increase in high school exit exam. Given the results in this paper indicate that a one standard deviation increase in peer quality could not lead to an increase of more than 0.08 standard deviations in student achievement, results here suggest that less than half of the effects found in Pop-Eleches and Urquiola (2013) could be attributed to the increase in peer quality in better schools. In particular, my results complement the findings of Hoekstra, Mouganie, and Wang (2018) and Jackson (2013) that the returns to high school quality are likely the result of other features of the schools other than peer quality and that peer quality explains very little of those returns.

More generally, the results of this study suggest that parents and policymakers should perhaps focus less on peer quality when making decisions as to how to best improve educational outcomes for children. Additionally, an equivalent way of interpreting the results is that being tracked into lower-ability classrooms and therefore being exposed to lower-ability peers does not result in lower student achievement. The results suggest that at least in this context, concerns that tracking systems might disproportionately harm students tracked into lower-ability classrooms seem overemphasized. Rather, future work on tracking should focus more on the effects it has on teaching and curriculum, and less on whether some students are left disproportionately exposed to higher- or lower-ability students.

3. THE IMPACT OF MISINFORMATION: EVIDENCE FROM THE ANTI-VACCINATION MOVEMENT IN THE US

3.1 Introduction

Recent advancements in technology have enabled information to travel faster and reach far more people than before. Unfortunately, this also means that it has become easy to spread misinformation and false stories. Additionally, while inconsequential false stories such as the flat-earth conspiracy have always existed, many of the current false stories are more likely to affect important outcomes. For example, Allcott and Gentzkow (2017) report that during the 2016 presidential election cycle fake news stories regarding presidential candidates were shared at least 37.6 million times on Facebook. In addition, they also estimate that average American adults likely saw and remembered at least one fake news stories in the months before the election. The increasing amount of misinformation and its potential consequences has thus generated significant debate about the government's role in regulating misinformation and social media platforms' responsibility to fight it. However, little is known about the actual impact of misinformation on behavior (Lazer et al., 2018). In theory, false stories can be seen as a distorted signal uncorrelated with the truth. This distorted signal could then lead consumers to make different decisions than they otherwise would have (Allcott and Gentzkow, 2017). Nevertheless, there has been little empirical evidence to confirm this theory. The purpose of this paper is to ask whether the dissemination of completely misinformation leads to meaningful changes in behavior.

I estimate the effect of misinformation by studying how vaccination rates responded to the unexpected surge in media coverage in 2007 of the claim, which was shown to be false in the early 2000s, that the MMR (Measles-Mumps-Rubella) vaccine causes autism. This exogenous shock in misinformation, along with the fact that some parents are *ex ante* more likely to be sensitive to this misinformation than others, allows me to identify the effects of misinformation about vaccine safety on parents' vaccination decisions.

There are several reasons why the surge in media coverage on the alleged link between vaccines and autism is the ideal setting in which to study the impact of false stories. First, the claim that the MMR vaccine, or any vaccine, causes autism is false and could be easily verified by 2007. The claim that the MMR vaccine causes autism stems from a now-retracted paper by Wakefield et al., which was published in 1998 in The Lancet, a major British medical journal. However, major medical and scientific bodies have since conducted further studies and refuted the claim as false; the Institute of Medicine (IOM) in May 2004, the Food and Drug Administration (FDA) in September 2006, and the Centers for Disease Control and Prevention (CDC) in July 2007. Therefore, whenever the media covered the stories or gave the platform to anti-vaccination activists to propel the claim without explicitly refuting it, especially after 2007, they were broadcasting false information. Second, in contrast to some other misinformation, false information about vaccine safety can affect important health outcomes. Parents who do not vaccinate their children not only expose their own children to the risk of serious diseases but also makes it harder for the community to retain herd immunity as well.[1] Third, the surge in misinformation about vaccine safety by the media in 2007 was unexpected to parents, because it was largely driven by high-profile court cases alleging that vaccines cause autism and celebrities' decisions to speak out on the issue. One notable instance of this was Jenny McCarthy making multiple appearances on talk shows, including The Oprah Winfrey Show.

I begin my analysis by looking at the media coverage on the alleged link between vaccines and autism to confirm that there is a surge in false news stories. Specifically, I collect news transcripts from six major television networks in the US (ABS, CBS, NBC, CNN, MSNBC, and Fox News) from 2001 to 2012 via LexisNexis. I use coverage on major television networks as a proxy for media coverage because although many people get their news through other sources, 44% of Americans still prefer television as the platform they most prefer for news (Mitchell, 2018). I classify a new story as false if it only reported on the alleged link between vaccines and

[1] Herd immunity is defined as the resistance to the spread of contagious disease within a population that results if a sufficiently high proportion of individuals are immune to the disease, especially through vaccination. For example, the vaccination rate required to achieve herd immunity is 83-94% for measles (Fine, 1993).

autism without refuting it as false. I then show that the number of false news stories about vaccine safety reported on these six networks rose dramatically from an average of 7.5 stories per year between 2001 and 2006 to 33 stories in 2007 and then 79.5 stories in 2008.[2]

To identify the effects of false news stories, I exploit this shock in false news about the MMR vaccine along with its differential impact on parents. Specifically, I expect misinformation should have larger effects on parents who are *ex ante* more likely to be sensitive and receptive to the false news about the MMR vaccine. Therefore, I identify the effects by comparing the vaccination rates of children whose parents are *ex ante* most sensitive and least sensitive to the news over time. While this approach will likely result in an underestimation of effects given all parents were likely somewhat affected by the false news, it enables me to use a difference-in-differences approach to distinguish effects from other time-varying factors. Specifically, I do so using individual-level vaccination data obtained from healthcare providers of 19-35-month-old American children surveyed in the 2002-2012 National Immunization Surveys (NIS). I determine parents' susceptibility to the news using three predetermined characteristics: whether the child is a firstborn, a boy, and the mother is over 30 years old. Parents are classified as most sensitive if they have all these three characteristics present and least sensitive if they have none. I use these three characteristics to determine parents' sensitivity for the following reasons. First, experienced parents were likely already exposed to information about vaccines prior to the surge in media coverage of the false claim in 2007 because of their past experience with their older children. As a result, the false news stories after 2007 likely only accounted for a small fraction of their information. Second, the child's gender and parental age are predictors of parents' sensitivity to stories involving autism risks because boys and children of older parents are known to be at a much higher risk of autism than their counterparts.[3] Importantly, the identifying assumption behind this approach is that the least sensitive parents and the most sensitive parents would have

[2]As explained later in the Data section, a news story is counted as one false story if both research assistants classified it as reporting on the false claim but not explicitly refuting it as false and 0.5 false story if only one research assistant did so.

[3]Autism is four times more common among boys than girls (CDC,2007). Children of older parents could be as much as five times as likely to be on the autism spectrum than children of younger parents (Reichenberg et al., 2006; Durkin et al., 2008).

changed their vaccination behavior in the same way in the absence of the surge in misinformation about vaccines.

Results indicate that the surge in false news stories about the MMR vaccine caused susceptible parents to become 3.3 percentage points less likely to vaccinate their children with an MMR shot by 15 months old. Importantly, this is the maximum age at which the CDC recommends the first MMR shot be administered. To assess whether parents were delaying the MMR shot or completely forgoing it, I examine the effects on take-up at 29 months old, which is the oldest age at which vaccination rates are consistently recorded in the survey. Results indicate that the resistance to the vaccine persisted. I estimate a 4.1 percentage point reduction in MMR shot take-up at 29 months old. This indicates that at a minimum, misinformation caused parents to delay vaccinating their children by over a year, and at most prevented them from ever immunizing their children. These results are robust to including time-varying controls and allowing family and state characteristics to have different effects on the MMR vaccine take-up rates each year. In addition, I also test whether my results are dependent on how I define treatment and control groups and find that the results are qualitatively similar when using more loosely defined treatment and control groups. Finally, subgroup analysis results suggest that college-educated mothers are more likely to be affected by this misinformation than non-college-educated mothers. This is consistent with Chang (2018) which finds that college-educated mothers were more likely to be affected by the initial vaccine controversy in 1998.

Indeed, the estimated reduction in vaccine take-up of 3 to 4 percentage points, which is likely an underestimate given the approach, is economically meaningful. A 3.27 percentage point (4.2 percent) drop in the MMR vaccine take-up at 15 months is equivalent to an increase of 15 percent in unvaccinated 15-month-olds. And a decrease of 4.13 percentage points (4.4 percent) in the MMR vaccine take-up at 29 months translates to an increase of 59 percent in unvaccinated 29-month-olds.[4] Lo and Hotez (2017) also, through model calibration, predict that a similar-sized decline of 5 percent in the MMR vaccine coverage of children 2-11 years old in the US would result in a

[4]Based on Table B.2, 22 percent of 15-month-olds and 7 percent of 29-month-olds were unvaccinated with an MMR shot.

three-fold increase in annual measles outbreaks.

In providing evidence that an increase in misinformation can lead to meaningful changes in behavior, this paper contributes to two bodies of literature. First, it complements the literature studying vaccine controversies. Smith, Ellenberg, Bell, and Rubin (2008), Anderberg, Chevalier, and Wadsworth (2011), and Chang (2018) study the impact of the vaccine controversy in 1998 when the MMR vaccine was first linked to autism and found that the MMR vaccine take-up rate decreased after 1998. My study differs from these studies in that while the claim in 1998 was believed to be true given it was published in a prestigious medical journal, by 2007 this claim had been clearly refuted. In this way, while these studies estimated the effect of new information that was expected to be reliable, my study identifies the effect of verifiably false information. In addition, this paper also complements Carrieri, Madio, and Principe (2019) which studies the impact of misinformation about the MMR vaccines in Italy in 2012 and finds a decrease in child immunization for all types of vaccines. The main difference between this paper and Carrieri, Madio, and Principe (2019) is in the nature of the events that triggered the surge in media coverage. The surge in the coverage of the misinformation about the MMR vaccine in Italy was due to a regional court officially recognizing a causal link between the MMR vaccine and autism in 2012. On the other hand, in my setting, the false claim was not endorsed by any government body or authority figure. Rather, the surge in media coverage of this false claim in the US was mainly driven by famous people speaking out on the issue and court hearings of a case *alleging* that vaccines cause autism. Despite the hearings, it is important to note that the US court never officially endorsed this false claim and eventually ruled against it in 2009 and 2010. Therefore, the main difference between the two papers is that the misinformation in Carrieri, Madio, and Principe (2019) was endorsed by an authority figure and could possibly be deemed reliable whereas the misinformation in my paper was not. All in all, combined with these previous findings, my paper shows that misinformation about vaccines reported by the media can affect people's decisions as much as or even more than perceived reliable information.

Second, this paper complements research on misinformation and media bias. Consistent with the theoretical framework of fake news provided by Allcott and Gentzkow (2017), results here suggest that misinformation can have important consequences. Importantly, these reductions in immunizations affect not only people's own welfare but also the welfare of those around them. Furthermore, the results also indicate that the general population does not easily detect misinformation, especially when it is reported by major media outlets. This is in line with the finding that consumers do not accurately determine the reliability of health content on the internet documented in Allam, Schulz, and Nakamoto (2014), Knapp, Madden, Wang, Sloyer, and Shenkman (2011), and Kutner, Greenburg, Jin, and Paulsen (2006). Lastly, this paper also speaks to related literature on the effects of media bias (DellaVigna and Kaplan, 2007; Gentzkow and Shapiro, 2006, 2010; Gerber, Karlan, and Bergan, 2009; Chiang and Knight, 2011; Enikolopov, Petrova, and Zhuravskaya, 2011; Prat, 2018; Martin and Yurukoglu, 2017). These studies built theoretical frameworks and provided empirical evidence that media slant can change individual beliefs and behavior. The results of this paper show that, in addition to media slant, completely false information reported by the media can also change behavior, even when it is easy for both the media and consumers to verify that the information is wrong.

3.2 Background: Media Coverage of the Anti-Vaccination Claim in the US

Although vaccines are regarded as one of the most successful medical interventions of the 20th century (CDC, 1999), some opposition to vaccines has always existed (Hussain et al., 2018). In 1998, however, the claim that vaccines are dangerous was propelled into the mainstream by the media when an article by Wakefield et al. (1998) suggested a causal link between the MMR vaccine and autism. The article was published in the Lancet, a major British medical journal. Anderberg, Chevalier, and Wadsworth (2011) studied the effects of this 1998 vaccine controversy and found that the MMR vaccine take-up rate declined sharply in the immediate years following the controversy. While the controversy did not garner as much media attention in the US as in the UK, Smith, Ellenberg, Bell, and Rubin (2008) and Chang (2018) also observed that the MMR vaccine take-up rates in children 19-35 months old in the US dropped by approximately 1-2 percentage points immediately following the Wakefield publication, but returned to pre-controversy levels by

2003. Importantly, the Wakefield et al. article was eventually retracted by the Lancet in 2010 after several subsequent studies disproved its results. While this retraction process took some time, I note that 10 of the 12 coauthors of the paper have retracted the paper in 2004 and issued a statement stating that they no longer interpret the results of their study as suggesting a causal link between the MMR vaccine and autism.

In the US, the topic of vaccine safety gained popularity again in 2007 when the media coverage on vaccine safety increased dramatically. This rise in the coverage was due in part to several vaccine court hearings of a case alleging that vaccines cause autism [5], and in part to the increasing number of celebrities publicly claiming that vaccines cause autism. Notably, Jenny McCarthy, an actress and TV host, famously went on talk shows including the Oprah Winfrey Show to talk about her belief that the MMR vaccine causes autism and how her son got diagnosed with autism after the MMR shot. For example, during the interview with Winfrey, McCarthy talked about her experience:

"Right before his MMR shot, I said to the doctor, I have a very bad feeling about this shot. This is the autism shot, isn't it? And he said, 'No, that is ridiculous. It is a mother's desperate attempt to blame something on autism.' And he swore at me.... And not soon thereafter, I noticed that change in the pictures: Boom! Soul, gone from his eyes."

Mnookin (2011) estimated McCarthy's message to have reached at least 15-20 million viewers based on her appearance on The Oprah Winfrey Show, Larry King Live, and Good Morning America alone.

Figure B.1 shows the number of news coverage on six major television networks (ABC, CBS, NBC, CNN, MSNBC, and FNC) of the false claim that vaccines cause autism from 2001 to 2012. As stated earlier, the coverage was few and far between from 2001 to 2006 before rising dramatically in 2007.

A critical aspect of the surge in media coverage on vaccine safety in 2007 is that at that point prominent medical bodies had already refuted the claim of any link between vaccines and autism.

[5] Despite the hearings, the US court never officially endorsed this false claim and eventually ruled against it in 2009 and 2010.

This includes the Institute of Medicine (IOM) in May 2004, the Food and Drug Administration (FDA) in September 2006, and the Centers for Disease Control and Prevention (CDC) in July 2007. In addition, as alluded earlier, the Wakefield et al. paper that had initially proposed the link had been disproved by multiple papers. Despite all that, Figure B.1 shows an increase in the number of news stories reporting on the alleged link between vaccines and autism without explicitly refuting it as false in 2007. This means that although the alleged link between vaccines and autism had been thoroughly debunked by that time, the public was exposed to a dramatic increase in misinformation alleging the link between vaccines and autism in 2007. I leverage this unanticipated increase in misinformation to estimate the causal impact of misinformation.

3.3 Data

To analyze the exposure to false news stories, I look at the number of television news stories on the alleged link between vaccines and autism over time. I use coverage on major television networks as a proxy for media coverage because although many people also access news through other sources, 44% of Americans still report television as the platform the most preferred for news (Mitchell, 2018). I obtained the news transcripts of six major television networks in the US from January of 2001 to December of 2012 from LexisNexis. The six networks were ABC, CBS, NBC, CNN, MSNBC, and FNC. To determine the number of false news stories, I first identified new stories that mentioned vaccines (or vaccination) and autism in the same section. I then hired two research assistants to read these news transcripts. I classify a new story as a false story if both research assistants flagged the story as 'reporting on the alleged link between vaccines and autism without explicitly refuting the claim'. If only one research assistant did so, I classify the story as 0.5 false story.[6] Table B.1 reports the number of false stories over time and matches the visual representation in Figure B.1. The number of false stories about vaccine safety rose dramatically from an average of 7.5 stories per year between 2001 and 2006 to 33 stories in 2007 and then 79.5 stories in 2008.

[6]The research assistants were instructed to sort and read the news transcripts in random order, rather than chronologically. This is to avoid any bias that could occur if they associate a certain time period with news of certain types.

To identify the impact of misinformation about vaccines on individual behavior, I look at parents' decisions regarding vaccination. In particular, since the MMR vaccine is the vaccine at the center of the vaccine-autism claim, I look at the MMR vaccine take-up rate as my main outcome. Individual-level data on vaccination decisions used in this paper comes from the 2002-2012 National Immunization Survey (NIS), which is conducted yearly by the Centers for Disease Control and Prevention (CDC). For each survey, the CDC surveys parents of 19-35 month-old children about their children's vaccination history. In addition, the CDC also asks for consent to obtain the vaccination records from their medical providers. Approximately 70% of the parents consent to the CDC acquiring vaccination records from their healthcare providers. Since healthcare provider records offer much more accurate information than parents' memory or a shot card, I only include children whose provider data is available in my analysis. For the analysis in this paper, I only include the data starting from 2002 to avoid the confounding effects from the first MMR vaccine controversy in 1998 when the Wakefield et al. paper first published. I only include the data up until 2012 because I only have media data up until 2012. I show in the Robustness section that the results are robust to alternative starting and ending years.

The National Immunization Surveys classify children into three age groups: 19-23 month olds, 24-29 month olds, and 30-35 month olds. I use the vaccination information of children from all age groups, i.e. all 19-35 month olds whose provider data is available, to look at the MMR vaccine take-up rate at 15 months old. Since the CDC recommends that the first MMR shot is given to a child at 12-15 months old, looking at the MMR vaccine take-up rate at 15 months old allows me to see if parents follow the CDC's recommendation. In addition, it is also important to see if parents only delay vaccinating their children or decline to vaccinate altogether. To address this question, I examine the MMR vaccine take-up rate of older children. The oldest children in my data set are 30-35 months old. This means that I have complete vaccination information up to when these children were 29 months old. I thus use the vaccination information of children 30-35 months old to look at the MMR vaccine take-up rate at 29 months old to see if parents have caught up to the vaccination schedule.

Table B.2 provides summary statistics of children included in my analysis. Panel 1 reports on all children in the 2002- 2012 National Immunization Surveys whose provider data is available, i.e. all 19-35 month olds, while Panel 2 reports the statistics of only 30-35 month-old children. Overall, 78% of children are vaccinated with an MMR shot by 15 months old and 93% are vaccinated by 29 months old. This suggests that at least approximately 15% of parents do not strictly follow the CDC's recommendation, but eventually vaccinate their children. In addition, the vaccination rates at both ages are in general higher among the children most likely to be affected by misinformation about vaccines (boy/firstborn/mother/$\geq$30) than those least likely to be affected (girl/not firstborn/mother<30).

3.4 Empirical Method

3.4.1 Measuring False News Exposure and Identifying the Post Period

I begin my analysis by identifying first which cohorts of children were affected by the increase in false news stories. I do so by looking at the number of false news stories to which parents are exposed. I first define the period when parents are most likely to pay attention to information about vaccine recommendations and vaccine safety as the 'exposure period'. For each child, I consider the exposure period to start in the month that the child was born and end in the month that I measure the child's MMR vaccine take-up. If I had information on each child's birthdate, I would identify each child's exposure period and then count the number of false news stories reported on television in this exposure period and use this number as a measure of parents' false news exposure. However, although the National Immunization Survey (NIS) data is rich in many ways, it does not provide information on the date of birth, the date of the interview, or age at the time of the interview. Therefore, I cannot directly back out the birth month and calculate parents' false news exposure for each child in my dataset individually. The NIS data does, however, provide information on which age group the child falls into at the time of the interview (19-23, 24-29, 30-35 months old). I thus calculate for the average news exposure for children in each age group in each interview year using this age group information along with two hypotheses. First, I assume that children of all ages are as equally likely to appear in the survey. Second, I assume that the probability of getting interviewed in each month is uniformly distributed throughout the year.

Figures B.2 and B.3 show the average false news exposure of parents interviewed in each survey year. Figure B.2 shows the average false news exposure up until when the child was 15 months old. Panel A shows that for parents whose child was 19-23 months old at the time of the interview, the first cohort that experienced the surge in false news was those interviewed in 2008. Panels B and C show the average false news exposure of parents whose child was 24-29 months old and 30-35 months old at the time of the interview, respectively. Both panels show that for both groups of parents, the first cohort that experienced the surge in false news was the one interviewed in 2009. Figure B.3 shows the average false news exposure up until when the child was 29 months old. I only look at the average false news exposure for parents whose child was 30-35 months old at the time of the interview here, because they are the only group with relevant information of children at 29 months old. We can see the average false news exposure rose dramatically for the cohort interviewed in 2008.

3.4.2 Classifying Treatment and Control Groups

To identify the effects of misinformation, we would ideally compare a group that was randomly exposed to misinformation to a group that was not exposed to misinformation. However, this is difficult for several reasons. First, people usually choose what they watch on television. For example, it could be the case that people who are less likely to vaccinate are the ones more likely to watch false news reports about vaccines on television. Second, more than 95% of US homes have television service (EIA, 2005) and therefore almost everyone was exposed to television and thus misinformation about vaccines to some degree. This makes it hard to identify a control group. In this paper, I overcome these issues by using a difference-in-differences approach that compares the groups that are *ex ante* most and least sensitive to misinformation about vaccines over time. Using this approach, the least sensitive group serves as the control group. The advantage of this approach is that I am able to distinguish the effect of false news exposure from other common time-varying factors, as well as group-specific factors. The disadvantage is because all parents are to some extent treated, this approach will underestimate the effect of misinformation on immunization behavior.

To identify which group of parents is the most sensitive and which group is the least sensitive to the misinformation about vaccines, it is important to consider which factors would make some parents more sensitive to the false news stories than others. Here, I propose that parents' sensitivity to false news stories about vaccines is based on both their parenting experience and their child's risk of being on the autism spectrum. There are two major reasons why misinformation about vaccines should be less impactful on experienced parents. First, because experienced parents would have started paying attention to information about vaccines earlier than first-time parents, the false news stories after 2007 would account for a smaller percentage of information for experienced parents. Therefore, the false news about vaccines, which increased dramatically in 2007, should be less impactful to experienced parents than first-time parents. Second, experienced parents are also more likely to have already formed their opinion on the issue from past experience and therefore less likely to be receptive to the new information than new parents. Therefore, among parents of same-age children in the data, experienced parents would likely be less sensitive to new information and thereby less affected by the increase in misinformation about vaccines.

Next, since the false news stories link vaccines to autism risk, parents whose child is at higher risk of being on the autism spectrum would likely be more sensitive to the news. In terms of autism risk, two characteristics—parental age and gender—have been consistently reported by both the CDC and media outlets to be associated with higher autism risk. For example, the CDC reported in February of 2007 that the autism spectrum disorder is 3-5 times more common among boys than girls (CDC, 2007). Similarly, several news networks reported on a study by Reichenberg et al. (2006) that found that children of men over 40 years old were 5.75 times more likely to have autism spectrum disorder compared with children of men under 30 years old.[7] A large study by Durkin et al. (2008) also found that firstborn children of two older parents were three times more

[7]McNamara, M. (2006) 'Men's Biological Clocks Are Ticking, Too', CBS, 15 November (https://www.cbsnews.com/news/mens-biological-clocks-are-ticking-too/)
Robin, R. (2007) 'It Seems the Fertility Clock Ticks for Men, Too', The New York Times, 27 Feb (https://www.nytimes.com/2007/02/27/health/27sper.html)

likely to develop autism than were third- or later-born offspring of 20-34 years old mothers and fathers under 40 years old.

I, therefore, determine parents' sensitivity to the news using three predetermined characteristics: whether the child is a firstborn, a boy, and the mother is over 30 years old. Mother's age is used as a proxy for parental age as it is the only consistent information about parental age available from the survey and the majority of couples are not more than 5 years apart in age.[8] Parents are classified as most sensitive to the false news stories if they have all three characteristics present and least sensitive if they have none. As a result, within my sample, I define the group that is the most sensitive to the misinformation about vaccines as boys who are a firstborn and whose mother is over 30 years old, and the group that is the least sensitive as girls who are not a firstborn and whose mother is younger than 30 years old.

Using these treatment and control groups, I implement a generalized difference-in-differences approach to identify the impact of misinformation about vaccines. Specifically, I compare the MMR vaccine take-up rate of boys who are a firstborn and whose mother is over 30 years old to the take-up rate of girls who are not a firstborn and whose mother is younger than 30 years old before and after the surge in misinformation. Formally, I estimate the impact of the increase in misinformation on parents' decision to vaccinate their child using the following model:

$$MMR_{it} = \alpha_t + \theta MostSensitive_i + \beta_x X_{it} + \beta MostSensitiveXPost_{it} + u_{it} \tag{3.1}$$

Where the outcome, MMR_{it}, is a binary variable indicating whether child i whose parent was interviewed in year t has been given at least one shot of MMR vaccine. In this paper, I focus on looking at this outcome at two points in time: when child i was 15 months old and 29 months old. I look at whether child i has been given any MMR shot at 15 months old because the CDC recommends that parents vaccinate their children with a dose of MMR vaccine at 12-15 months old, and therefore this will show whether parents stop following the CDC's recommendation.

[8] Based on the 2013 Current Population Survey, for 76.7% of heterosexual married couples, the husband and wife are less than 5 years apart in age.

Additionally, it is also important to assess whether misinformation has long-run effects on vaccination take-up, or if it only delays it. This is why I test for effects on children in the oldest age group in my sample, who are 30-35 months old at the time of the interview, to look at MMR vaccine take-up rate at 29 months old.

α_t is survey year fixed effects. $MostSensitive_i$ is an indicator variable for whether child i's parents are classified as the most sensitive, i.e. whether child i is a boy, a firstborn, and has a mother who is over 30 years old. X_i is a matrix containing child i's characteristics including state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group at the time of the interview, whether they live in the state they were born in, and whether their state allows personal belief exemption from vaccination. $MostSensitiveXPost_{it}$ is an indicator variable for whether child i is in the most sensitive group in the post period. The post-period starts in the year when we first see the dramatic increase in false news exposure as discussed in the Empirical Method section. Importantly, the coefficient of interest here is β which measures the effects of misinformation on parents' decision to vaccinate. Specifically, it measures whether parents most sensitive to the surge in misinformation vaccinate their children differently than parents who are the least sensitive.

In all specifications, survey weights are used and robust standard errors and their corresponding p-values are reported. In addition, accounting for within-cluster dependence in estimating standard errors of regression estimates is important. Ideally, we want to cluster at the level of treatment or higher. However, since I only have two clusters, I follow the wild bootstrap method proposed in Cameron, Gelbach, and Miller (2008) which clusters at the year level. These wild-bootstrap p-values are reported for all specifications. Furthermore, I also perform a randomization inference exercise. Specifically, I randomly reassign child gender, mother's age, and firstborn status based on the true distribution of each variable in each year, and then estimate the effect (β) based on the reassignment. I do this for 1,000 replications and plot the distribution of the 1,000 coefficients estimated. I then compute the proportion of these 1,000 coefficients that have larger absolute value than the actual estimate and interpret this number as the two-tailed empirical p-value.

As with any difference-in-differences design, the underlying assumption for this approach is that MMR vaccine take-up rates of children in the control group and treatment group would have changed similarly over time in the absence of the increase in misinformation. I provide support for this assumption by first showing the visual representation of the raw data that shows the MMR vaccine take-up rates for control and treatment groups track each other prior to the post period. Second, I also formally test for the divergence in outcomes between the treatment and control groups in the pre periods using a dynamic difference-in-differences approach.

One potential concern with this approach is that perhaps results would differ for alternative definitions of treatment and control groups. To provide further support for my identification strategy, I also perform multiple analyses using more loosely defined treatment and control groups. Specifically, I do this in three different ways. First, I include more children in my control group. Namely, instead of excluding children who have one or two predetermined characteristics[9], I include them in my control group. Second, I include more children in my treatment group, i.e. instead of excluding children who have one or two predetermined characteristics, I include them in my treatment group. And lastly, I use two instead of three predetermined characteristics to determine treatment and control groups. With more loosely defined treatment and control groups, we would expect the effects to be weaker, but of the same sign.

Additionally, another potential concern is that exposure to misinformation might have caused some parents to become less (or more) likely to allow the CDC to obtain their official vaccination record from their healthcare providers. If this is the case, the estimate might simply just reflect the change in the consent rates and not parents changing their vaccination behavior. For example, a lower consent rate from parents who did not vaccinate their children would result in a lower number of unvaccinated children being included in the data, even when the parents did not change their vaccination behavior. This, in turn, would affect the vaccination rates of the treatment and

[9]In the main specification, the control group is children with zero of the three predetermined characteristics and the treatment group is children with all three predetermined characteristics. The three predetermined characteristics are: whether the child is a firstborn, a boy, and the mother is over 30 years old.

control groups and then result in treatment effects, even when there is no actual change in the vaccination behavior. To provide supporting evidence that this is likely not the case, I look at the consent rates of the treatment and control groups over time. Figure B.4 and Table B.3 both indicate that there is no significant effect of false news exposure on the consent rate.

3.5 Results

3.5.1 Main Results

I begin by looking at the raw data of the MMR vaccine take-up rates over time. Figures B.5 and B.6 show the MMR vaccine take-up rates at 15 months old and 29 months old, respectively. Time is re-centered so that year=+1 is the first year parents experienced the surge in false news exposure. For both the MMR vaccine take-up rates at 15 and 29 months old, Figures B.5 and B.6 show that prior to the surge in false news exposure, the take-up rates among children in the treatment group (boys who are a firstborn and whose mother is over 30 years old) and control group (girls who are not a firstborn and whose mother is younger than 30 years old) track each other well over the years. This is important since the validity of a difference-in-differences approach hinges on the parallel trend assumption. Additionally, the figures also show that before the increase in false news exposure, children in the treatment group are consistently more likely to be vaccinated than children in the control group both at 15 months old and 29 months old. However, after the increase in false news stories, the gap in vaccination rates between the two groups closes. The gap closes by about half for the MMR vaccination rate at 15 months old and closes completely for the MMR vaccination rate at 29 months old.

To assess the parallel trends assumption more rigorously, I estimate a dynamic difference-in-differences model, controlling for year fixed effects, group fixed effects, and observable characteristics, to check if the treatment group diverges from the control group in any year before the increase in false news exposure. Figures B.7 and B.8 plot the dynamic difference-in-differences estimates for MMR vaccine take-up at 15 months old and 29 months old respectively. Both figures reaffirm that for both outcomes, there is no evidence of divergence in trends before the increase in false news exposure. In addition, both figures also show that after the increase in false news exposure, both the MMR vaccine take-up rates at 15 months and 29 months of children in the treatment group fall. This suggests that increased exposure to misinformation about vaccine safety does not only lead parents to deviate from the CDC's recommended schedule, but also delays vaccination by a minimum of a year, and possibly longer.

Next, I formally estimate the average treatment effect of the increase in exposure to false news stories and report the results in Table B.4. Column 1 shows the average treatment effect of misinformation on the MMR vaccination rates using the simplest difference-in-differences model, without any controls. Based on this specification, the rise in misinformation about vaccine safety causes the MMR vaccine take-up rate at 15 months old to drop by 4.57 percentage points and the MMR vaccine take-up rate at 29 months old to drop by 4.53 percentage points. Column 2 reports the estimates from the preferred specification, shown in equation 3.1, which also includes controls for observable characteristics, state fixed effects, and state vaccination exemption law. If my results are driven by the change in the characteristics of children in my control or treatment groups and not by the increased exposure to misinformation, then these controls should absorb my treatment effects. The estimates from this specification are only slightly smaller than those reported in column one but are still in the same direction and statistically significant. Based on these estimates, the increased exposure to misinformation about vaccines causes the MMR vaccine take-up rates at 15 and 29 months old to decrease by 3.27 percentage points and 4.13 percentage points, respectively.

Finally, families with different characteristics, such as income, parents' education level, and race, may respond differently to year-to-year shock. For example, richer parents might have better access to vaccines in the year where there is a vaccine shortage. Since my treatment and control groups are different in terms of family income and mother's education (as shown in Table B.2), in Column 3 of Table B.4, I allow observable characteristics to affect the MMR vaccine take-up rate differently each year. The estimate reported from this specification for the MMR vaccine take-up rate at 15 months old is no longer statistically significant at the conventional level but the magnitude still remains at a similar level of -2.31 percentage points. The estimate for the effect on the MMR vaccine take-up rate at 29 months old is robust and remains stable at a statistically-significant 4.16 percentage points reduction. This shows that the effects were not driven by the differences in characteristics between the two groups. In this table, wild-bootstrap p-values, which allow the correlation between take-up rates within the same year, are also reported alongside with the robust p-values. As shown in the table, wild-bootstrap p-values and robust p-values are very similar, and using the wild-bootstrap approach does not change my results. Finally, Table B.4 also reports randomization inference p-values for estimates from the preferred specification. The randomization inference p-values for both the effects at 15 months and 29 months are similar to the wild-bootstrap p-values and robust p-values. The effect at 15 months is significant at the five percent level (randomization inference p-values=0.019) and the effect at 29 months is significant at the one percent level (randomization inference p-values=0.003).[10]

Overall, these results suggest that misinformation about vaccines' link to autism caused both the MMR vaccination rates at 15 months old and 29 months old to drop by at least 3-4 percentage points. This indicates that at a minimum, misinformation caused parents to delay vaccinating their children by over a year, and at most prevented them from ever immunizing their children.

3.5.2 Subgroup Analysis by Mother's Education

Next, I examine whether the impact of this false claim about the MMR vaccine varies across parents' education. This is because it is possible that highly-educated parents process and apply health information differently than other parents (Grossman, 1972).[11] Specifically, I test whether mothers with a college degree are more likely to be affected by the false claim and changed their vaccination behavior more than mothers without a college degree.[12]

Table B.5 shows the effects on the vaccine take-up rate of children with college-educated mothers and non-college-educated mothers separately. Column 1 shows the results for the whole

[10]The distributions of coefficients from this randomization exercise are shown in Figures B.9.

[11]In fact, there is a large body of literature devoted to studying the link between education and health decisions and health outcomes. For example, Brunello, Fort, Schneeweis, and Winter-Ebmer (2016), Lange (2011), and Kenkel (1991). For an extensive literature review on this topic, see Brunello, Fort, Schneeweis, and Winter-Ebmer (2016).

[12]The only information about parents' education available in the dataset is mother's education.

sample. Column 2 shows the results for children whose mother has a college degree, while Column 3 shows the results for children whose mother does not have a college degree.

For the MMR vaccine take-up at 15 months old, the results indicate that the reduction in take-up rate is almost entirely driven by college-educated mothers. The effect is a statistically insignificant reduction of 4.46 percentage points among children of college-educated mothers, whereas, it is only a statistically insignificant 0.1 percentage point reduction among children of non-college-educated mothers.

For the MMR vaccine take-up at 29 months old, the results indicate that the difference in effects across subgroups is small. The effect is a statistically insignificant reduction of 4.25 percentage points among children of college-educated mothers and a statistically insignificant 3.35 percentage point reduction among children of non-college-educated mothers.

Overall, the results here suggest that the effect of misinformation about vaccines may be larger for college-educated mothers, though this difference is starker at 15 months than at 29 months.

3.6 Robustness

3.6.1 Effects on MMR Vaccine Take-Up at Other Ages

In addition to the main results discussed in the Main Results section, I also look at the effects of misinformation on the MMR vaccine take-up rates at other ages besides 15 and 29 months old. I estimate the average treatment effects on the MMR vaccine take-up rate at each age from 15-29 months old using the preferred specification shown in equation 3.1. For the estimate at each age, I only include children who at the time of the interview are older than the age at which I measure the MMR vaccine take-up.[13] In addition, since the age at which I measure the MMR vaccine take-up changes the exposure period,[14] I also re-examine the exposure period, the exposure to false news, and revise the first post year for each estimation.

[13]This is because we only have information on the vaccination history of each child up until the time of the interview. For example, when the outcome is the MMR vaccine take-up at 20 months old, I only include children who were older than 20 months old at the time of the interview in the analysis. Since there are three age groups of children in my dataset: 19-23 months, 24-29 months, 30-35 months, this means that only children in age groups 24-29 and 30-35 months old are included in the analysis of the MMR vaccine take-up rate at 20 months old.

[14]Specifically, the exposure period relevant for when the outcome is MMR vaccine take-up at z months would be from when the child was born until when the child was z months old and not after.

The results, reported in Table B.6 and visually in Figure B.10, show that the estimates are relatively similar across ages. They are all negative and range from -1.3 to -4.6 percentage points with 80% of them being statistically different from zero at the 10% level. This indicates that the negative effects of misinformation observed in the earlier section are not driven by the selection of the 15 and 29 month ages.

3.6.2 Using More Loosely Defined Control and Treatment Groups

In the main analysis, I compare children who are most and least likely to be affected by the treatment. I classify children into these two groups using three characteristics that are associated with susceptible parents: whether the child is a firstborn, the child is a boy, and the mother is over 30 years old. Children with all of these three characteristics present are classified as most likely to be affected whereas children with none of these characteristics are classified as least likely to be affected. These two groups are then used as my treatment and control groups. In this section, I perform multiple analyses using more loosely defined treatment and control groups to test the robustness of my findings to alternative classifications. As explained in the earlier section, I redefine my control and treatment groups in three major ways: 1. expanding the definition of the control group, 2. Expanding the definition of the treatment group, and 3. defining treatment and control groups using only two characteristics. Using the more loosely defined treatment and control groups, we would expect to see the treatment effects become smaller in magnitude, but not completely disappear.

The results of this exercise for the MMR vaccine take-up rate at 15 months old are reported in Table B.7 and the same results for the MMR vaccine take-up at 29 months old are reported in Table B.8. Column 1 shows the results of the main identification strategy. Columns 2-3 show the estimates when I add more children into my control groups by including children with only one or two of the three characteristics associated with susceptible parents in the control group as well. Columns 4-5 show the estimates when I increase my treatment group by including children with only one or two of the three characteristics associated with susceptible parents in my treatment group. Columns 6-8 show the estimates when I only use two characteristics in defining my control and treatment groups. For any two characteristics I use, my treatment group is the children with both 2 characteristics present and the control group is the children with neither of the 2 characteristics present. All the estimates reported are, as expected, smaller in magnitude than the estimates from the main identification. And although some estimates are no longer significant at conventional levels, all of them are still negative, and all but one of them still report a relatively low p-value. In particular, the estimates for the MMR vaccine take-up at 29 months old are very robust to alternative definitions of treatment and control groups.

3.6.3 Other Robustness Checks

In addition, since the dependent variable is binary, I also use logistic regression to estimate my main results. The results are shown in Table B.9. Similar to the linear regression results, the logistic regression results show reductions in the MMR vaccine take-ups. Furthermore, I also test the robustness of the results to changing the starting or ending year of my sample. Results are shown in Tables B.10 through B.14, and indicate that changing the start or end years does not change the results. For the estimates of the effect at 15 months, although some of them are no longer significant at conventional levels, they are all still negative. Specifically, estimates of the effect at 15 months range from -0.0237 to -0.0375. The smaller estimates are those from data sets that extend the end year, consistent with Figures B.1 and B.2 that show exposure to misinformation has been decreasing since 2010-2011. Estimates of the effect at 29 months are all statistically significant at the five percent level and range from -0.0346 to -0.0443.

It is also worth considering what would have to be true for a confounding factor to drive the results estimated in this paper. The confounder must i) have caused a coincidental divergence in the MMR vaccine take-up rates between the most and least sensitive group in the post-period, but not in the years before; ii) be orthogonal to any of the observable characteristics; iii) affect boys, firstborns, and children of older parents more than girls, later-borns, and children of younger parents. This seems unlikely. In addition, if the surge in tv coverage of the misinformation was not exogenous and was actually a result of growing concern about vaccine safety in the population, it seems unlikely that the divergence in the MMR vaccine take-ups between the most and least sensitive group would only start in the post-period but not in the years before. For these reasons, I therefore interpret estimates as the causal impact of misinformation about vaccines.

3.7 Discussion and Conclusion

This paper studies the effect of misinformation on individuals using the unanticipated rise in television coverage of the alleged link between vaccines and autism in 2007 as an exogenous shock in misinformation to parents. Using vaccination data obtained from healthcare providers of 19-35-month-old children surveyed in the 2002-2012 National Immunization Surveys (NIS), I find that misinformation about vaccines resulted in a drop of at least 3.3 percentage points in the MMR vaccine take-up rate at 15 months old which is the CDC's recommended age. In addition, misinformation also led to a drop of at least 4.1 percentage points in the MMR vaccine take-up rate at 29 months old. This indicates that at a minimum, misinformation caused parents to delay vaccinating their children by over a year, and at most prevented them from ever immunizing their children.

The estimates here are economically meaningful, especially considering that my identification strategy of comparing the most and least sensitive groups likely results in the underestimation of effects. The estimated drop in the MMR vaccine take-up at 15 months old is equivalent to an increase of 15 percent in unvaccinated 15-month-olds while the estimated decrease in the MMR vaccine take-up at 29 months old is equivalent to an increase of 59 percent in unvaccinated 29-month-olds. In addition, Lo and Hotez (2017) predict that a similar-sized reduction in the MMR vaccine coverage of children 2-11 years old in the US would result in a three-fold increase in annual measles outbreaks. Importantly, results here suggest that people can change behavior in important ways that not only affect their own welfare but also the welfare of those around them. Additionally, these estimates are comparable or even bigger than the reported effects of new and reliable information found in prior literature. For example, Smith, Ellenberg, Bell, and Rubin (2008) reports that the number of American children who received all childhood immunizations except for the MMR vaccine rose from 0.8 percent to 2.1 percent after the publication of Wakefield et al. (1998) which first suggested a link between the MMR vaccine and autism. Chang

(2018) also examines the effects of the 1998 vaccine controversy in the US and reports that the overall MMR vaccine take-up declined by 1.1 to 1.5 percentage points in the immediate year following the Wakefield et al. (1998) publication. Combined with these findings, my results suggest that misinformation reported by the media can change individual behavior as much as reliable information and that the general public is not able to discern false information even when it is easy to verify.

The subgroup analysis results in this paper also suggest that college-educated mothers are more affected by this misinformation about the MMR vaccine. This is in line with Chang (2018) which uses non-college educated mothers as a counterfactual in a difference-in-difference framework and finds that an increase of 10 news stories about the vaccine controversy in 1998 led college-educated mothers to be 0.4 percent less likely to vaccinate their children with an MMR shot.

These results also have clear relevance for public policy regarding fake news and misinformation. Much of the debate over the responsibility of social media companies and the government in combating misinformation depends on whether misinformation actually matters. Results presented here provide clear evidence that misinformation can change a behavior that not only affects those individuals but also potentially imposes negative externalities on those around them. This suggests that there are potentially large social benefits from preventing the dissemination of misinformation.

4. AN EMPIRICAL TEST OF ANTI-MUSLIM BIAS: EVIDENCE FROM PROPERTY VALUES

4.1 Introduction

The Muslim population in the US is growing at a fast pace. It is projected that Islam will overtake Judaism as the second-largest religion in the United States by 2040, in large part due to immigration from Muslim-majority countries (Mohamed, 2018). The growth of the Muslim population has been associated with a perceived rise in anti-Muslim sentiment. In 2017, the Pew Research Center analyzed data published by the FBI and found that hate crimes against Muslims are rising, with the number of assaults on Muslims, higher relative to the months immediately following the September 11 attacks. Likewise, a Pew Research Center survey revealed that a majority of American Muslims feel that being Muslim is becoming more difficult in the US (Kishi, 2017). Even the FBI data likely do not fully capture the situation. The Bureau of Justice Statistics surveys suggest that hate crimes are severely under-reported to the FBI, with the real numbers about 25 times higher (Sandholtz et al., 2013).

While all this suggests that anti-Muslim bias is increasing over time, as with all forms of bias it is difficult to provide definitive evidence. We offer a new test of anti-Muslim bias by asking whether the introduction of a Mosque in a neighborhood leads to a reduction in housing prices. We note that Mosques are a particularly good point of emphasis for this test, given that anti-Mosque incidents have increased by a factor of eight in the last 15 years, as tracked by the American Civil Liberties Union (ACLU).[1]

The purpose of this study is to provide an empirical test of whether there is anti-Muslim bias in the United States. To do so, we ask whether new mosque openings are capitalized into housing prices. In the presence of substantial anti-Muslim bias, we would observe a reduction in sales prices. We link administrative data on property transactions from 2002 to 2018 in Detroit and

[1] The ACLU regularly updates the following page: https://www.aclu.org/issues/national-security/discriminatory-profiling/nationwide-anti-mosque-activity

Hamtramck, Michigan, to the opening dates of all places of worship in those two cities. We address concerns about the potential endogeneity in the location of mosques by exploiting the granular nature of our data, which allows for a spatial difference-in-differences research design using repeat-sales of the same properties. In particular, we compare outcomes for properties that are very close to a new mosque to those that are slightly farther away. The identifying assumption of our approach is that the changes in property prices across these two areas would have been similar in the absence of the new mosque. Importantly, we show that outcomes trended similarly in the years prior to the entry of the new mosques, which is consistent with the identifying assumption.

Our results provide evidence against the hypothesis that anti-Muslim sentiment is capitalized into property prices following the entry of a new mosque. In our baseline specification, we rule out negative effects greater than 2 percent for properties within walking distance (around 0.2 mile) from a new mosque. Using a wild cluster bootstrap, we are able to rule out even smaller negative effects. While our main results focus on repeat-sales data that enables us to include individual property fixed effects, we also obtain similar results when using the universe of property sales and directly controlling for observed property-level characteristics. We also do not see any evidence of differential changes in the turnover rate of neighboring properties in the time period around mosque openings.

We also assess whether the absence of a negative effect on property values is because other positive aspects of a nearby house of worship offset the effect of anti-Muslim bias. To do so, we test for an effect of new churches in the neighborhood, and show that the null effect for mosques is mirrored by a null effect of churches. This provides further evidence against the hypothesis that anti-Muslim bias is capitalized into housing prices, as both churches and mosques have the same (null) impact on housing prices. We note that one potential limitation of our approach is that there may be offsetting effects of demand for housing near mosques. While non-Muslim residents might want to move away due to anti-Muslim bias, they may be offset by Muslims wanting to move to those areas. Instead, our approach captures the net impact of both potential effects. We interpret our null results to mean that if religious discrimination against Muslims exists, it is not strong enough to offset increased demand for properties near mosques. We also analyze the predicted ethnic origin based on names of buyers and find little evidence of an influx of Muslim buyers.

To our knowledge, this is the first paper to use data on repeat sales of properties to test whether the entry of new mosques cause anti-Muslim bias to be capitalized into housing prices. We are closely related to research on the effects of terror attacks on anti-Muslim sentiment (Lepage, 2020 and Ratcliffe and von Hinke Kessler Scholder, 2015). Unlike these studies, we seek to assess religious discrimination by focusing solely on increased salience from new mosques, separately from the effects of terrorist attacks, as terrorist attacks may also cause a shock to perceived safety.

We also contribute to an existing body of research that use hedonic pricing models to quantify the contribution of places of worship to housing values (e.g., Do, Wilbur, and Short, 1994, Carroll, Clauretie, and Jensen, 1996; Ottensmann, Bielefeld, and Payton, 2006). In particular, our study is closely related to a working paper by Brandt, Maennig, and Richter (2013). Using a model that allows for spatial dependence, as well as a rich set of controls, they find that places of worship in Hamburg, Germany are associated with an increase in housing prices, and that the effect of mosques is not different from that of other churches. Our approach differs from theirs in that we leverage the opening of new mosques to identify causal effects.

In addressing the effects of mosque openings, this paper also speaks to the literature on the impact of potentially controversial facilities on housing prices. For example, Dröes and Koster (2016) examined the effects of wind turbines on housing prices and estimated a decrease of 1.4 percent in housing prices. Daams, Proietti, and Veneri (2019) studied the impact of asylum seekers' reception centers and found that the openings of these centers resulted in a drop of as much as 9.3 percent in house prices. Chirakijja (2021) studied the impact of prison openings and estimated that prisons decrease housing values by 2-4 percent. As our lower bound estimates allow us to rule out decreases larger than 2 percent in housing prices, it would seem that although some people may have concerns and prejudice against mosques, they are not nearly as large as those against prisons and asylum seekers' reception centers.

4.2 Muslims and Mosques in the US

4.2.1 Growth of the Muslim Population

Although Muslims still form only about 1% of the population, immigration from countries with significant Muslim populations makes Islam among the fastest growing religion in the US. Muslims tend to be more concentrated in some parts of the country. States such as New Jersey, New York, and Arkansas have a greater share of Muslims in the adult population than other states.[2] There is also variation within states, with metro areas like Detroit and its enclave Hamtramck having a significant share of Muslims among their adult populations.

The two cities we examine in this paper, Detroit and Hamtramck, are two of the most densely populated cities in Michigan. Detroit is the largest city in the state, while Hamtramck is essentially an enclave of Detroit. While still majority Christian, both cities have large Muslim populations. Hamtramck in particular, after decades of being comprised of mostly Catholic Polish immigrants, has a large Muslim population due to immigration from countries such as Bangladesh, Yemen, and Bosnia (Perkins, 2010). The Muslim population is ethnically diverse, with almost equal parts South Asian, Middle Eastern, and European. These factors makes Detroit and Hamtramck uniquely suited to an examination of the effects of mosques on housing prices.

4.2.2 Potential Effects of Mosques

The construction of mosques has proven to be a controversial issue. In Hamtramck, some residents have complained about the externality effects of mosques, especially the daily Adhan, or the call to prayer, which starts at sunrise and repeats several times until the evening (Bailey, 2015). The city of Hamtramck has allowed the broadcast since 2004. Some have also voiced fears over Islam becoming a state religion, as some Muslim immigrants come from countries where Islam is a state religion (Barro and McCleary, 2005). Without further analysis, it is difficult to determine whether these complaints are tainted by anti-Muslim sentiment rather than purely a general sentiment against places of worship in neighborhoods.

[2] https://www.pewforum.org/religious-landscape-study/religious-tradition/muslim/

We note that mosques and Islamic centers not only provide religious services, but also playgrounds, day care centers, and a place for social gatherings among people in the neighborhood. These are potentially positive contributions to neighborhood amenities, similar to what other places of worship provide.

Because mosques necessarily become a neighborhood feature, we focus on their effect on housing prices. There is established empirical literature that examines the extent to which neighborhood attributes are capitalized into housing prices. As with this literature, we rely on the framework introduced by Rosen (1974) and adapt it to study the effect of mosques.

4.3 Data and Empirical Strategy

4.3.1 Data

We rely on administrative data for our key variables. To determine the location, as well as the opening dates of the mosques and churches in Detroit and Hamtramck, we queried the Corporations Online Filing System on the Michigan Department of Licensing and Regulatory Affairs website.[3] This gave us the universe of mosque and church openings during the time period 2002-2018. Table C.1 summarizes the mosque and church openings that occur within the time span of our data. We observe 10 mosque openings and 13 church openings between 2002 and 2018.

For the analysis, we obtained the property sales data from the cities of Detroit and Hamtramck. We restrict our analysis to non-token sales.[4] The outcome data are summarized in Table C.2, which summarizes data on all property sales (Panel A) and repeat sales (Panel B). The full data contain almost 12,000 property sales. Repeat sales are properties that were sold more than once during our study period. They comprise 73 percent of all property sales we observe. We use repeat sales to account for time invariant property-level characteristics in our primary specification. Our primary sub-sample thus contains around 8,500 sales. The mean sales price is under $60,000, and most properties within the study zone are residential.

Because property sales data include the names of buyers and sellers, we are also able to conduct

[3] See https://cofs.lara.state.mi.us/SearchApi/Search/Search. We used the keywords masjid, mosque, Islamic Center, Muslim, Muslem, church, and synagogue.

[4] Token sales are sales of value under $5,000.

ethnic name analysis to determine the likelihood that the parties involved in a property transaction are Muslim. The tool we use to conduct this analysis is NamePrism (Ye et al., 2017; Ye and Skiena, 2019).[5]

4.3.2 Empirical Strategy

We base our empirical strategy on the standard hedonic pricing approach developed by Rosen (1974), and applied in the empirical literature in studies such as Pope and Pope (2015), McMillen (2004), Zabel (2008) among many others. We modify the standard model by incorporating a spatial difference-in-differences design, exploiting the plausibly exogenous location of the new mosques, shown in Figure C.1. The modified hedonic model is of the form :

$$log(P_{it}) = \alpha_t + \gamma \mathbf{X} + \beta \cdot post\ mosque\ opening \times treated\ zone_{it} + \epsilon_{it} \tag{4.1}$$

where $log(P_{iq})$ is the log of the sales price of property i that was sold in year-month t, α_t represents year-month fixed effects, which captures period specific shocks common to all properties, and $\mathbf{X}$ is a set of controls for property-level characteristics. $post\ mosque\ opening \times treated\ zone_{it}$ is a binary variable indicating that property i is in the treated zone, i.e. located within 0.2 mile, or approximately 500 steps of a new mosque, and that the transaction in year-month t occurred after the associated mosque's opening. The analysis neighborhood is defined to be twice the size of the treated zone, in this case, a 0.40 mile radius around the new mosque. The coefficient of interest here is β, which measures the effect of mosque openings on log property prices for properties within 0.20 mile of a new mosque, relative to the properties between 0.20 and 0.40 mile of that mosque.

While we estimate Equation 4.1 with all available controls for property-level characteristics, due to data limitations, our data might not capture all the important differences between properties that drive sales prices. Our preferred estimates are instead based on a repeat-sales framework using data on properties that were sold more than once between 2002 and 2018. Formally, we estimate

[5] We are grateful to Junting Ye and Steven Skiena for allowing us to access this web tool.

the effects of mosque openings using the following model:

$$log(P_{it}) = \alpha_i + \alpha_t + \beta \cdot post\ mosque\ opening \times treated\ zone_{it} + \epsilon_{it} \tag{4.2}$$

where α_i represents property fixed effects, which captures time-invariant property-level characteristics that contribute to the sales price. As in Equation 1, α_t is year-month fixed effects. $post\ mosque\ opening \times treated\ zone_{it}$ is a binary variable indicating that property i is in the treated zone and that the transaction in year-month t occurred after the associate mosque's opening. The coefficient of interest is still β, which measures the effect of mosque openings on log property prices, accounting for the time-invariant property-level characteristics.

The inclusion of property fixed effects helps us account for differences we cannot observe in the data, such as the number of bedrooms and bathrooms, the general condition of the house, among others. The identifying assumption is that the change in the sales price experienced by the properties slightly farther away from the new mosque provides a valid counterfactual for the change in sales prices that the treated properties would have experienced had the new mosque not opened. Robust standard errors are clustered at the level of the neighborhood surrounding the newly opened mosque. We also use a wild cluster bootstrap to account for the possibility that we may have too few treated clusters (Roodman et al., 2019).

To assess the validity of our identifying assumption, as well as to visualize the dynamics of the response of house prices, we also estimate a dynamic version of Equation 4.2:

$$log(P_{it}) = \sum_l \delta_l \cdot I[\frac{t - mosque\ opening\ date_i}{90} = l] \times treated\ zone_i + \alpha_i + \alpha_t + \varepsilon_{it} \tag{4.3}$$

where t is the sale date of property i. $mosque\ opening\ date_i$ is the opening date of property i's associated new mosque. l signifies the lead or lag quarter, of property's i transaction date relative to the associated mosque's opening date. $treated\ zone_i$ again indicates whether property i is in the treated zone. α_i is property fixed effects and α_t is year-month fixed effects. Here, Equation 4.3 estimates the dynamic effects (δ_l) at each point in time of mosque openings on log property prices.

The validity of the research design would be supported by estimating null effects for the leading terms.

4.4 Results

4.4.1 The Effect of Mosque Openings

First, we examine the validity of the parallel trends assumption required for a difference-in-differences method. To do so, we estimate a dynamic difference-in-differences model as shown in Equation 4.3. Panel A of Figure C.2 shows the dynamic difference-in-differences estimates when the treated properties are defined to be those within 0.20 mile distance from a newly opened mosque. The comparison properties are those that are up to twice the distance away. This cutoff was chosen to approximate a mosque being walking distance to a property.[6] Figure C.2 shows that there is little evidence of divergence in property prices before a new mosque opens in the neighborhood, providing some evidence in favor of our identifying assumption. Moreover, there does not appear to be a significant effect of new mosques on property prices. The event studies for other definitions of the treated zone give similar results.

This is borne out by the formal estimates presented in Panel A of Table C.3, which reports the average treatment effects of new mosque openings across different specifications. Each column presents either the estimate of the coefficient on our variable of interest, $post\ mosque\ opening \times treated\, zone_{it}$ from the repeat-sales model of Equation 4.2, or from the modified hedonic model of Equation 4.1 when using all sales with controls. There appears to be no significant negative effect of mosque openings, in contrast to what detractors of mosques suggest. Column 1 reports the estimates from the preferred specification, i.e., Equation 4.2, for all repeat sales when the treated zone is 0.20 mile from a new mosque and the comparison group are properties between 0.20 and 0.40 mile away from the new mosque. We estimate an insignificant increase of approximately 6 percent in property prices. Importantly, the lower bound estimates here allow us to rule out negative effects larger than 2 percent, while our 95% confidence interval based on a wild cluster bootstrap rules out any meaningful negative effects.

[6]The 0.20 mile distance is roughly 500 steps.

Columns 2 and 3 report the estimates from the modified-hedonic model shown in Equation 4.1. We use two sets of property characteristics, grouped by data availability, since data on some property characteristics is sometimes available only for a subset of property sales. The lower bounds of these estimates also allow us to rule out negative effects larger than 2 percent.

Although our results show insignificant effects of mosque openings on property prices, we cannot immediately conclude that this is evidence of no religious discrimination, especially if church openings result in significant increases in property prices. We assess the possibility that church openings may affect property prices differently from mosque openings by estimating the impact of church opening on property prices. The estimated effects of church openings are reported in Panel B of Table C.3 and are similar, though less precise, to the estimates of mosque openings in Panel A. The similarity in magnitude of these estimates is more easily seen in Figure C.3, which graphs the estimated coefficients for our variables of interest, $post mosque opening \times treated zone_{it}$ and $post\ church\ opening \times treated\ zone_{it}$, across different possible definitions of the treated and comparison zones.

Our discussion of the potential effects of new mosque openings in the Muslims and Mosques in the US section suggests that mosques may provide amenities to Muslims that might drive an increase in demand for housing near mosques large enough to offset anti-Muslim bias. Using names of buyers, we conduct an analysis that asks whether new mosque openings lead to an influx of Muslims moving closer to mosques. Specifically, we estimate the impact of mosque openings on the likelihood of the property buyer being Muslim, using models similar to Equations 4.1 and 4.2. The results of this analysis are presented in Table C.4. Our estimates do not indicate a large influx of Muslim buyers into areas near new mosques.

4.4.2 Robustness Checks

As our main estimates rely on a definition of the treated zone of properties being within 0.2 mile of a new mosque, we examine the sensitivity of our estimates to changing this cutoff. We start by defining the treated zone as being within 0.15 mile, and then increasing by 0.05 mile increments until 0.3 mile. With each definition of the treated zone, we define the comparison zone as being twice the distance to the mosque (i.e., comparison zone for the 0.15 mile definition is the area greater than 0.15 mile but less than or equal to 0.30 mile to the new mosque). We also do this exercise for church openings. The results are presented in Columns 4-12 of Panel A in Table C.3 for mosque openings, and Panel B for church openings.

These tables report the effects of mosque and church openings, i.e. the coefficient of variables $post\ mosque\ opening \times treated\ zone_{it}$ and $post\ church\ opening \times treated\ zone_{it}$, respectively, for different constructed treated zones. The estimated coefficients are not significantly different from our main estimate, indicating that our results are not sensitive to arbitrarily small changes in how we define the properties treated by the new mosque or church.

In Table C.5, we also estimate Equations 4.1 and Equation 4.2 using data on only residential property sales. Our estimates are qualitatively similar, as would be expected since more than 90 percent of property sales transactions in our sample are of residential properties. Finally, in Appendix Table C.6, we show that there is no change in the turnover rate of properties in Detroit during the time period around a new mosque opening, consistent with our main results.

4.5 Conclusion

We combine a spatial difference-in-differences framework with a standard hedonic price model to propose a test for religious discrimination based on the impact of new mosques on property values. Finding a negative effect on property values would have indicated that new mosques cause property values to fall as a result of religious discrimination.

Given the rhetoric surrounding new mosque openings, one would expect that new mosque openings would significantly reduce property values. However, the results presented in this paper provide little evidence of this. Using a spatial difference-in-differences research design on repeat-sales data of properties in Detroit and Hamtramck, we are able to rule out negative effects larger than 2 percent. These results are robust to differently constructed treated zones, and to alternative specifications that use all property sales, with observable property characteristics. We also find little evidence that our null results are driven by anti-Muslim bias being offset by Muslim buyers wanting to move closer to mosques.

Taken together, our findings of the effects of mosque and church openings suggest that even if residents or prospective buyers choose to avoid the neighborhood as a result of a new mosque opening, it appears that this response is not strong enough to offset demand for these properties. While our estimates rule out reductions in property prices of greater than two percent, Daams, Proietti, and Veneri (2019) and Chirakijja (2021) estimated much larger negative effects of asylum seekers reception center and prisons, respectively, facilities around which there also exist potentially discriminatory rhetoric. Because our estimates suggest that new mosques do not appear to have the same negative impact on property prices as these facilities, we conclude that there is little evidence of religious discrimination against Muslims, at least in a setting that abstracts from terrorist attacks and the ensuing effects on perceived safety.

5. SUMMARY AND CONCLUSIONS

This dissertation examines three topics in applied microeconomics. In Section 2, I study tracking and find that being tracked into classrooms with higher-ability peers, instead of lower-ability peers, does not lead to improvement in student GPA. This result suggests that concerns that tracking systems might disproportionately harm students tracked into lower-ability classrooms seem overemphasized. In Section 3, I show that at a minimum, misinformation about the MMR vaccine caused parents to delay vaccinating their children by over a year, and at most prevented them from ever immunizing their children. Results presented in this section provide clear evidence that misinformation can change a behavior that not only affects those individuals but also potentially imposes negative externalities on those around them. This suggests that there are potentially large social benefits from preventing the dissemination of misinformation. Finally, Section 4 shows that new mosque openings do not suppress property values and that the effects of church and mosque openings on property values are similar. This suggests that even if there is bias against Muslims, it is not strong enough to offset the demand and suppress the values of properties near mosques.

As both education and health are important foundations of human capital development, Sections 2 and 3 provide empirical evidence that would be useful in designing policies concerning children's outcomes. Although the finding in Section 4 indicates little evidence of anti-Muslim bias in Detroit and Hamtramck, the test proposed in this section could be used by policymakers to identify religious discrimination in other areas. This would also lead to opportunities for a more effective policy response.

REFERENCES

Abdulkadiroğlu, A., Angrist, J., and Pathak, P. (2014). "The elite illusion: Achievement effects at boston and new york exam schools." *Econometrica*, *82*(1), 137–196.

Abdulkadiroğlu, A., Pathak, P. A., Schellenberg, J., and Walters, C. R. (2020). "Do parents value school effectiveness?" *American Economic Review*, *110*(5), 1502–39.

Allam, A., Schulz, P. J., and Nakamoto, K. (2014). "The impact of search engine selection and sorting criteria on vaccination beliefs and attitudes: two experiments manipulating google output." *Journal of medical internet research*, *16*(4), e100.

Allcott, H., and Gentzkow, M. (2017). "Social media and fake news in the 2016 election." *Journal of economic perspectives*, *31*(2), 211–36.

Allensworth, E. M., Moore, P. T., Sartain, L., and de la Torre, M. (2017). "The educational benefits of attending higher performing schools: Evidence from chicago high schools." *Educational Evaluation and Policy Analysis*, *39*(2), 175–197.

Anderberg, D., Chevalier, A., and Wadsworth, J. (2011). "Anatomy of a health scare: education, income and the mmr controversy in the uk." *Journal of Health Economics*, *30*(3), 515–530.

Bailey, S. P. (2015). "In the first majority-muslim us city, residents tense about its future." *The Washington Post*, *21*.

Barro, R. J., and McCleary, R. M. (2005). "Which countries have state religions?" *The Quarterly Journal of Economics*, *120*(4), 1331–1370.

Booij, A. S., Haan, F., and Plug, E. (2016). "Enriching students pays off: Evidence from an individualized gifted and talented program in secondary education." *IZA discussion paper*.

Booij, A. S., Haan, F., and Plug, E. (2017). "Can gifted and talented education raise the academic achievement of all high-achieving students?" *IZA Discussion Paper*.

Brandt, S., Maennig, W., and Richter, F. (2013). "Do places of worship affect housing prices? Evidence from Germany." *Hamburg Contemporary Economic Discussions*, (48).

Brunello, G., Fort, M., Schneeweis, N., and Winter-Ebmer, R. (2016). "The causal effect of education on health: What is the role of health behaviors?" *Health economics*, *25*(3), 314–336.

Bui, S. A., Craig, S. G., and Imberman, S. A. (2014). "Is gifted education a bright idea? assessing the impact of gifted and talented programs on students." *American Economic Journal: Economic Policy*, *6*(3), 30–62.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). "Bootstrap-based improvements for inference with clustered errors." *The Review of Economics and Statistics*, *90*(3), 414–427.

Card, D., and Giuliano, L. (2016). "Can tracking raise the test scores of high-ability minority students?" *American Economic Review*, *106*(10), 2783–2816.

Carrell, S. E., Fullerton, R. L., and West, J. E. (2009). "Does your cohort matter? measuring peer effects in college achievement." *Journal of Labor Economics*, *27*(3), 439–464.

Carrieri, V., Madio, L., and Principe, F. (2019). "Vaccine hesitancy and (fake) news: Quasi-experimental evidence from italy." *Health economics*, *28*(11), 1377–1382.

Carroll, T. M., Clauretie, T. M., and Jensen, J. (1996). "Living next to godliness: Residential property values and churches." *The Journal of Real Estate Finance and Economics*, *12*(3), 319–330.

Centers for Disease Control and Prevention and others (2007). "Prevalence of the autism spectrum disorders (asds) in multiple areas of the united states, 2000 and 2002." *Community Report from the Autism Developmental Disabilities Monitoring (ADDM) Network*.

Centers for Disease Control and Prevention (CDC) and others (1999). "Ten great public health achievements–united states, 1900-1999." *MMWR. Morbidity and mortality weekly report*, *48*(12), 241.

Chang, L. V. (2018). "Information, education, and health behaviors: Evidence from the mmr vaccine autism controversy." *Health economics*, *27*(7), 1043–1062.

Chiang, C.-F., and Knight, B. (2011). "Media bias and influence: Evidence from newspaper endorsements." *The Review of Economic Studies*, *78*(3), 795–820.

Chirakijja, J. (2021). "The local economic impacts of prisons." *Available at SSRN 3794967*.

Clark, D. (2010). "Selective schools and academic achievement." *The BE Journal of Economic*

Analysis & Policy, *10*(1).

Cohodes, S. R. (2020). "The long-run impacts of specialized programming for high-achieving students." *American Economic Journal: Economic Policy*, *12*(1), 127–66.

Daams, M. N., Proietti, P., and Veneri, P. (2019). "The effect of asylum seeker reception centers on nearby house prices: Evidence from the netherlands." *Journal of Housing Economics*, *46*, 101658.

Dee, T., and Lan, X. (2015). "The achievement and course-taking effects of magnet schools: Regression-discontinuity evidence from urban china." *Economics of Education Review*, *47*, 128–142.

DellaVigna, S., and Kaplan, E. (2007). "The fox news effect: Media bias and voting." *The Quarterly Journal of Economics*, *122*(3), 1187–1234.

Do, A. Q., Wilbur, R. W., and Short, J. L. (1994). "An empirical examination of the externalities of neighborhood churches on housing values." *The Journal of Real Estate Finance and Economics*, *9*(2), 127–136.

Dobbie, W., and Fryer Jr, R. G. (2014). "The impact of attending a school with high-achieving peers: Evidence from the new york city exam schools." *American Economic Journal: Applied Economics*, *6*(3), 58–75.

Dröes, M. I., and Koster, H. R. (2016). "Renewable energy and negative externalities: The effect of wind turbines on house prices." *Journal of Urban Economics*, *96*, 121–141.

Duflo, E., Dupas, P., and Kremer, M. (2011). "Peer effects, teacher incentives, and the impact of tracking: Evidence from a randomized evaluation in kenya." *American Economic Review*, *101*(5), 1739–74.

Durkin, M. S., Maenner, M. J., Newschaffer, C. J., Lee, L.-C., Cunniff, C. M., Daniels, J. L., Kirby, R. S., Leavitt, L., Miller, L., Zahorodny, W., et al. (2008). "Advanced parental age and the risk of autism spectrum disorder." *American journal of epidemiology*, *168*(11), 1268–1276.

EIA, US (2005). "Residential energy consumption survey." *US Energy Information Administration 2009*.

Enikolopov, R., Petrova, M., and Zhuravskaya, E. (2011). “Media and political persuasion: Evidence from russia.” *American Economic Review*, *101*(7), 3253–85.

Fine, P. E. (1993). “Herd immunity: history, theory, practice.” *Epidemiologic reviews*, *15*(2), 265–302.

Gentzkow, M., and Shapiro, J. M. (2006). “Media bias and reputation.” *Journal of political Economy*, *114*(2), 280–316.

Gentzkow, M., and Shapiro, J. M. (2010). “What drives media slant? evidence from us daily newspapers.” *Econometrica*, *78*(1), 35–71.

Gerber, A. S., Karlan, D., and Bergan, D. (2009). “Does the media matter? a field experiment measuring the effect of newspapers on voting behavior and political opinions.” *American Economic Journal: Applied Economics*, *1*(2), 35–52.

Hoekstra, M., Mouganie, P., and Wang, Y. (2018). “Peer quality and the academic benefits to attending better schools.” *Journal of Labor Economics*, *36*(4), 841–884.

Hoxby, C. (2000). “Peer effects in the classroom: Learning from gender and race variation.” Tech. rep., National Bureau of Economic Research.

Hussain, A., Ali, S., Ahmed, M., and Hussain, S. (2018). “The anti-vaccination movement: A regression in modern medicine.” *Cureus*, *10*(7).

Jackson, C. K. (2013). “Can higher-achieving peers explain the benefits to attending selective schools? evidence from trinidad and tobago.” *Journal of Public Economics*, *108*, 63–77.

Kenkel, D. S. (1991). “Health behavior, health knowledge, and schooling.” *Journal of Political Economy*, *99*(2), 287–305.

Kishi, K. (2017). “Assaults against muslims in us surpass 2001 level.” *Pew Research Center*, *15*.

Knapp, C., Madden, V., Wang, H., Sloyer, P., and Shenkman, E. (2011). “Internet use and ehealth literacy of low-income parents whose children have special health care needs.” *Journal of medical Internet research*, *13*(3), e75.

Kutner, M., Greenburg, E., Jin, Y., and Paulsen, C. (2006). “The health literacy of america’s adults: Results from the 2003 national assessment of adult literacy. nces 2006-483.” *National Center for*

Education Statistics.

Lange, F. (2011). "The role of education in complex health decisions: evidence from cancer screening." *Journal of Health Economics*, *30*(1), 43–54.

Lavy, V., and Schlosser, A. (2011). "Mechanisms and impacts of gender peer effects at school." *American Economic Journal: Applied Economics*, *3*(2), 1–33.

Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., et al. (2018). "The science of fake news." *Science*, *359*(6380), 1094–1096.

Lefgren, L. (2004). "Educational peer effects and the chicago public schools." *Journal of urban Economics*, *56*(2), 169–191.

Lepage, L.-P. (2020). "Discrimination and segregation in the real estate market: Evidence from terrorist attacks and mosques."

Lo, N. C., and Hotez, P. J. (2017). "Public health and economic consequences of vaccine hesitancy for measles in the united states." *JAMA pediatrics*, *171*(9), 887–892.

Loveless, T. (2013). "The resurgence of ability grouping and persistence of tracking." *Part II of the 2013 Brown Center Report on American Education, the Brookings Institution.[2]*, *1*.

Lucas, A. M., and Mbiti, I. M. (2014). "Effects of school quality on student achievement: Discontinuity evidence from kenya." *American Economic Journal: Applied Economics*, *6*(3), 234–63.

Martin, G. J., and Yurukoglu, A. (2017). "Bias in cable news: Persuasion and polarization." *American Economic Review*, *107*(9), 2565–99.

McMillen, D. P. (2004). "Airport expansions and property values: the case of Chicago O'Hare Airport." *Journal of urban economics*, *55*(3), 627–640.

Mitchell, A. (2018). "Americans still prefer watching to reading the news-and mostly still through television." *Pew Research Center*.

Mnookin, S. (2011). *The panic virus: a true story of medicine, science, and fear*. Simon and Schuster.

Mohamed, B. (2018). “New estimates show us muslim population continues to grow.” *Pew Research Center*, *3*.

Murphy, R., and Weinhardt, F. (2020). “Top of the class: The importance of ordinal rank.” *The Review of Economic Studies*, *87*(6), 2777–2826.

OECD (2013). “Pisa 2012 results: What makes schools successful? resources, policies and practices (volume iv).”

Ohinata, A., and Van Ours, J. C. (2013). “How immigrant children affect the academic achievement of native dutch children.” *The Economic Journal*, *123*(570), F308–F331.

Ottensmann, J., Bielefeld, W., and Payton, S. (2006). “The location of nonprofit organizations influences residential housing prices.”

Park, A., Shi, X. S., Hsieh, C.-T., and An, X. (2009). “Does school quality matter?: Evidence from a natural experiment in rural china.”

Perkins, A. (2010). “Negotiating alliances: Muslims, gay rights and the christian right in a polish-american city (respond to this article at http://www. therai. org. uk/at/debate).” *Anthropology today*, *26*(2), 19–24.

Pop-Eleches, C., and Urquiola, M. (2013). “Going to a better school: Effects and behavioral responses.” *American Economic Review*, *103*(4), 1289–1324.

Pope, D. G., and Pope, J. C. (2015). “When walmart comes to town: Always low housing prices? always?” *Journal of Urban Economics*, *87*, 1–13.

Prat, A. (2018). “Media power.” *Journal of Political Economy*, *126*(4), 1747–1783.

Ratcliffe, A., and von Hinke Kessler Scholder, S. (2015). “The london bombings and racial prejudice: Evidence from the housing and labor market.” *Economic Inquiry*, *53*(1), 276–293.

Reichenberg, A., Gross, R., Weiser, M., Bresnahan, M., Silverman, J., Harlap, S., Rabinowitz, J., Shulman, C., Malaspina, D., Lubin, G., et al. (2006). “Advancing paternal age and autism.” *Archives of general psychiatry*, *63*(9), 1026–1032.

Roodman, D., Nielsen, M. Ø., MacKinnon, J. G., and Webb, M. D. (2019). “Fast and wild: Bootstrap inference in stata using boottest.” *The Stata Journal*, *19*(1), 4–60.

Rosen, S. (1974). “Hedonic prices and implicit markets: product differentiation in pure competition.” *Journal of political economy*, *82*(1), 34–55.

Sacerdote, B. (2001). “Peer effects with random assignment: Results for dartmouth roommates.” *The Quarterly journal of economics*, *116*(2), 681–704.

Sandholtz, N., Langton, L., and Planty, M. (2013). *Hate crime victimization, 2003-2011*. US Department of Justice, Office of Justice Programs, Bureau of Justice Statistics.

Shi, Y. (2019). “Who benefits from selective education? evidence from elite boarding school admissions.” *Economics of Education Review*, 101907.

Smith, M. J., Ellenberg, S. S., Bell, L. M., and Rubin, D. M. (2008). “Media coverage of the measles-mumps-rubella vaccine and autism controversy and its relationship to mmr immunization rates in the united states.” *Pediatrics*, *121*(4), e836–e843.

Vardardottir, A. (2013). “Peer effects and academic achievement: a regression discontinuity approach.” *Economics of Education review*, *36*, 108–121.

Ye, J., Han, S., Hu, Y., Coskun, B., Liu, M., Qin, H., and Skiena, S. (2017). “Nationality classification using name embeddings.” In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 1897–1906.

Ye, J., and Skiena, S. (2019). “The secret lives of names? name embeddings from social media.” In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 3000–3008.

Zabel, J. E. (2008). “Using hedonic models to measure racial discrimination and prejudice in the us housing market.” In *Hedonic methods in housing markets*, 177–201, Springer.

Zimmerman, D. J. (2003). “Peer effects in academic outcomes: Evidence from a natural experiment.” *Review of Economics and statistics*, *85*(1), 9–23.

APPENDIX A

FIGURES AND TABLES FOR SECTION TWO

A.1 Figures

Figure A.1: Histogram of running variable

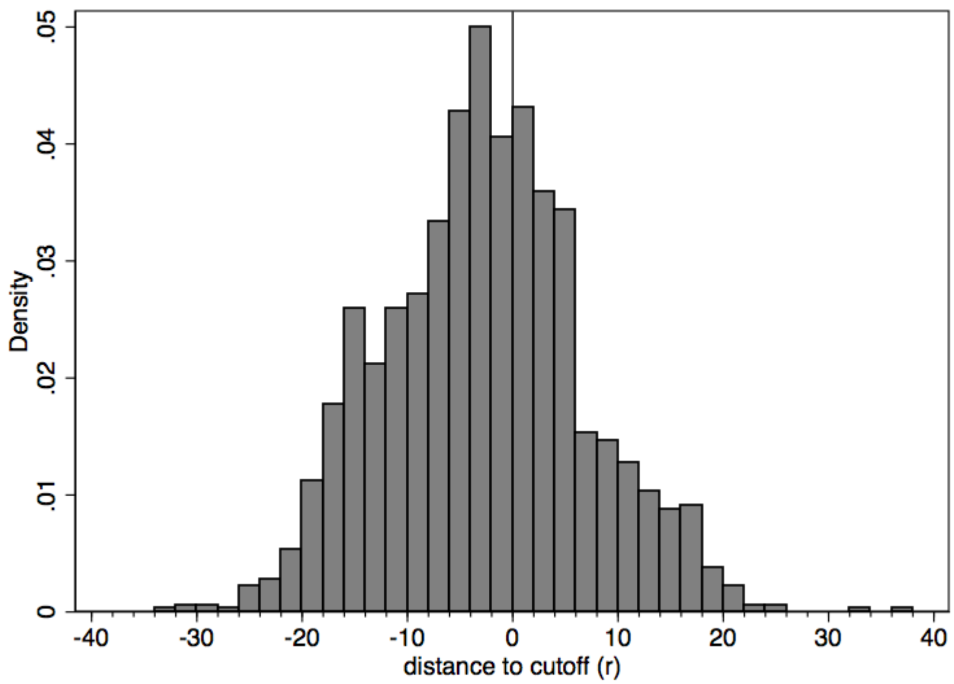

Notes: This figure shows the density of the running variable. The running variable used here is the distance to cutoff, i.e. how far each student's preliminary score is from the cutoff. The number is positive if they score above the cutoff, and negative if they score below the cutoff.

Figure A.2: Student characteristics across cutoff

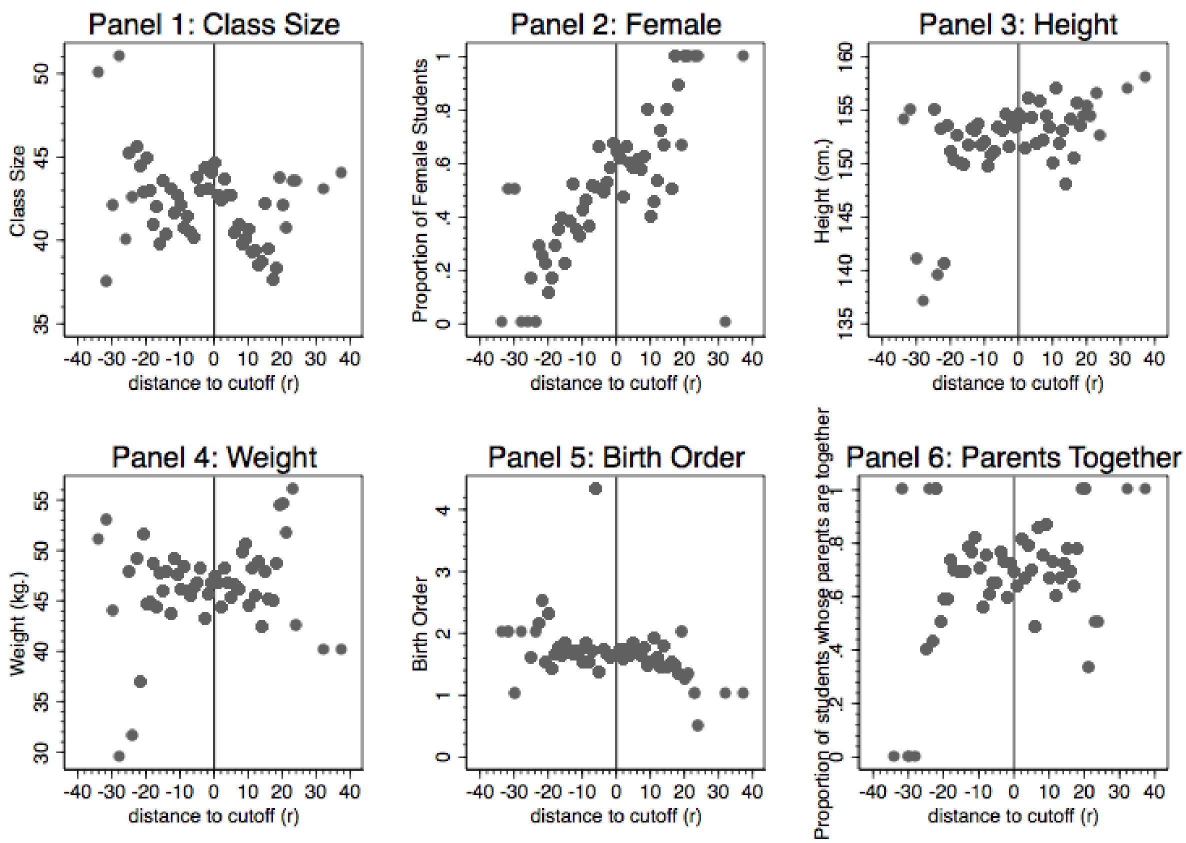

Notes: This figure plots student characteristics across the cutoff. The running variable is the distance to cutoff, i.e. how far each student's preliminary score is from the cutoff. The number is positive if they score above the cutoff, and negative if they score below the cutoff.

Figure A.3: Predicted 7th grade cumulative GPA based on student characteristics

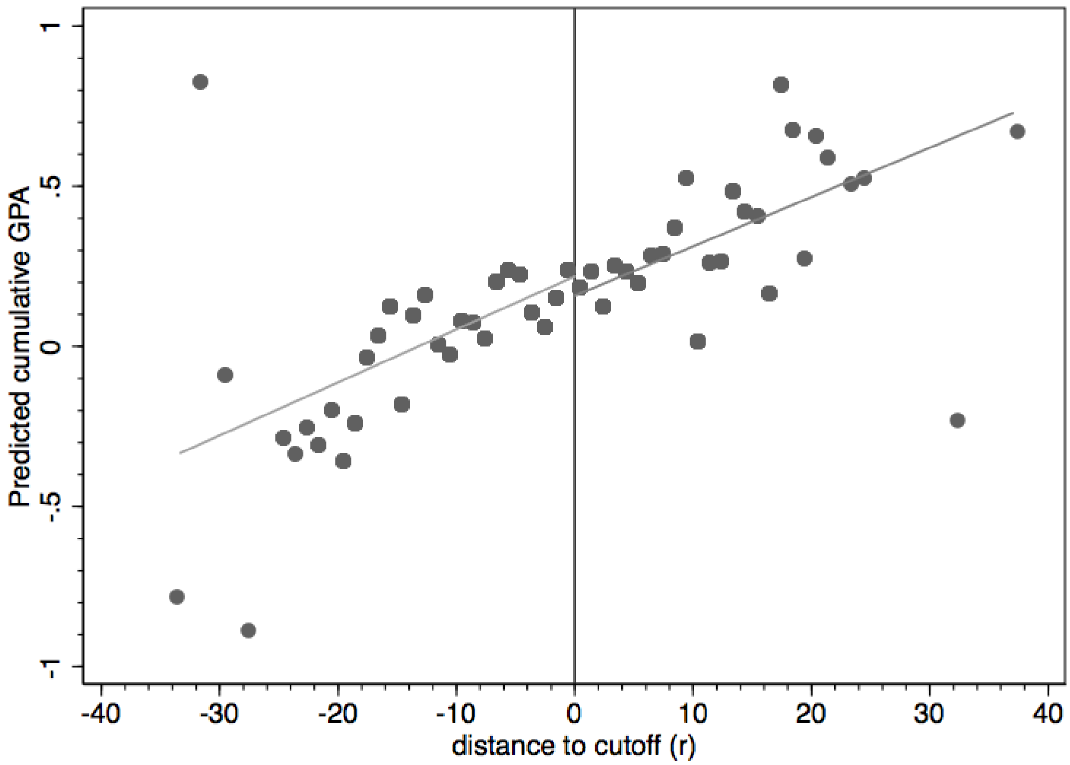

Notes: This figure plots predicted cumulative GPA across the cutoff. Predicted cumulative GPAs are based on the regression of (standardized) cumulative GPA on student characteristics.The running variable is the distance to cutoff, i.e. how far each student's preliminary score is from the cutoff. The number is positive if they score above the cutoff, and negative if they score below the cutoff.

Figure A.4: Likelihood of being in the higher-ability classrooms

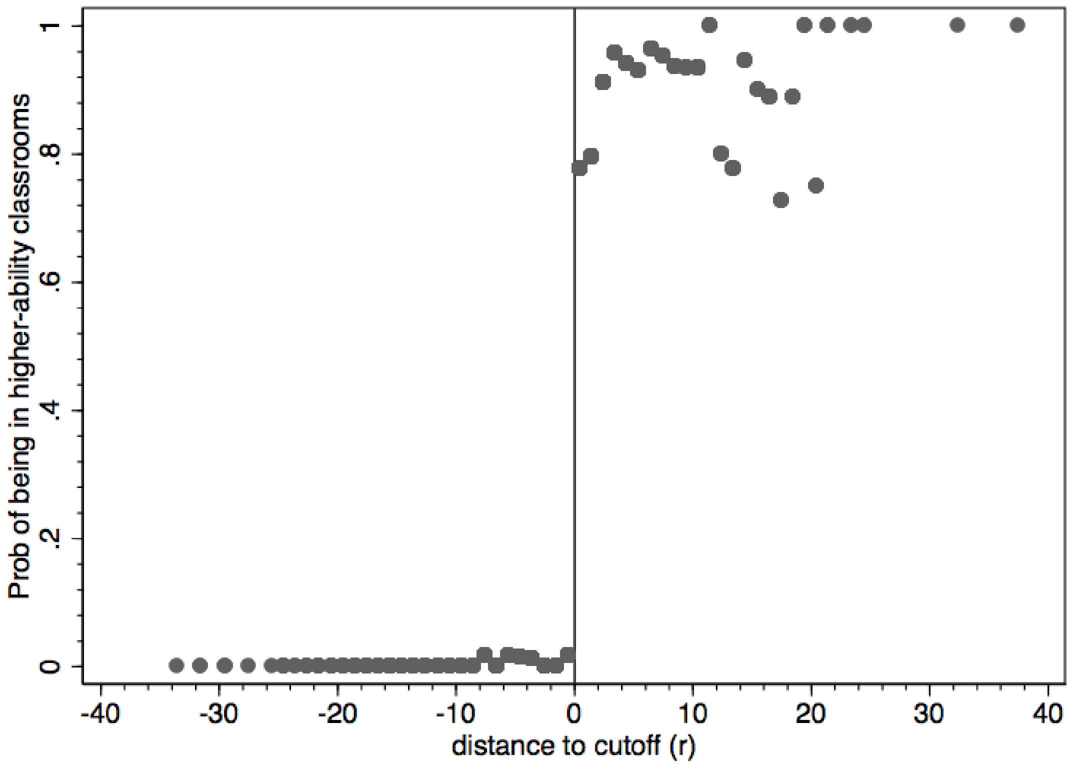

Notes:This figure plots the first-stage relationship between the likelihood of student being in the higher-ability classroom and the running variable. The running variable is the distance to cutoff, i.e. how far each student's preliminary score is from the cutoff. The number is positive if they score above the cutoff, and negative if they score below the cutoff.

Figure A.5: Peer quality

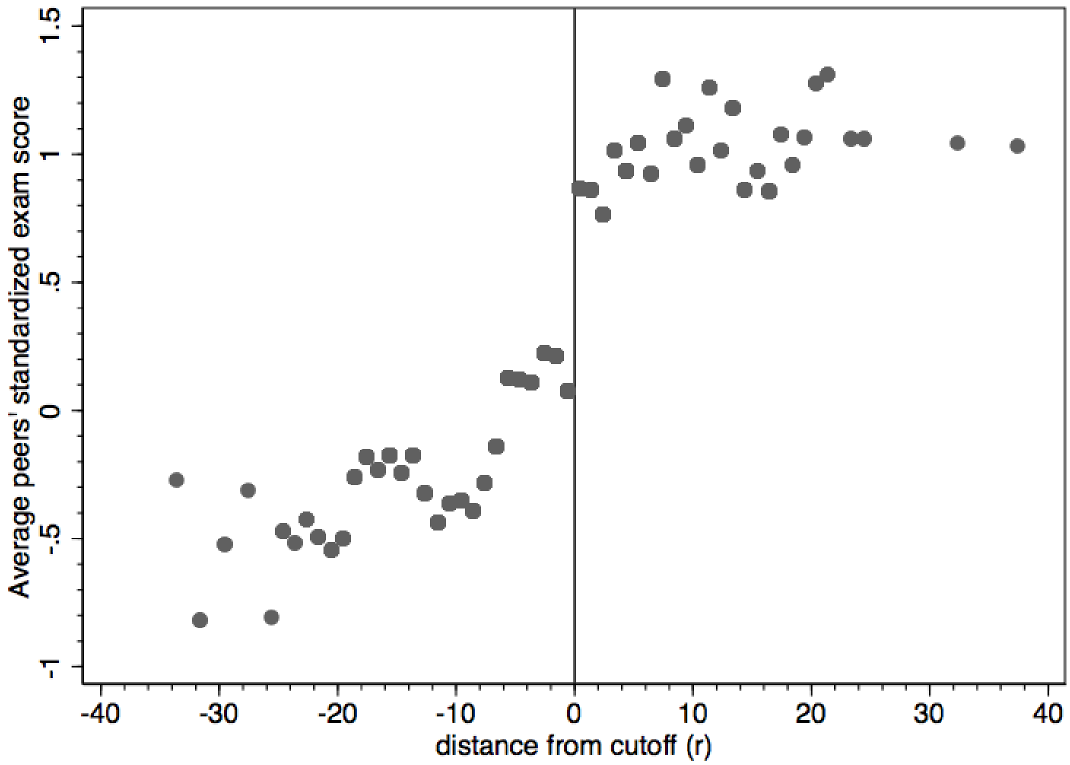

Notes: This figure plots the relationship between peer quality and the running variable. The running variable is the distance to cutoff, i.e. how far each student's preliminary score is from the cutoff. The number is positive if they score above the cutoff, and negative if they score below the cutoff. Each student's peer quality is proxied by the average of their classmates' standardized preliminary exam scores.

Figure A.6: Standardized seventh-grade cumulative GPA

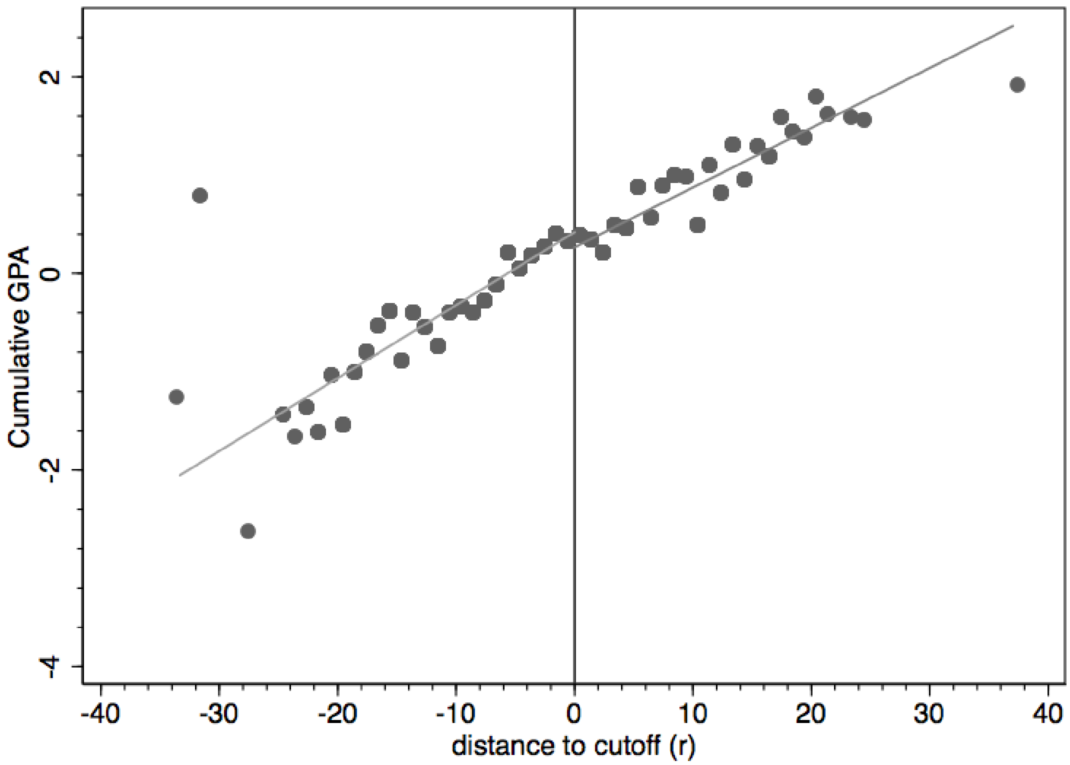

Notes: This figure shows the reduced-form relationship between students' cumulative seventh-grade GPA and the running variable.The running variable is the distance to cutoff, i.e. how far each student's preliminary score is from the cutoff. The number is positive if they score above the cutoff, and negative if they score below the cutoff.

Figure A.7: Reduced-form estimates using different bandwidth sizes

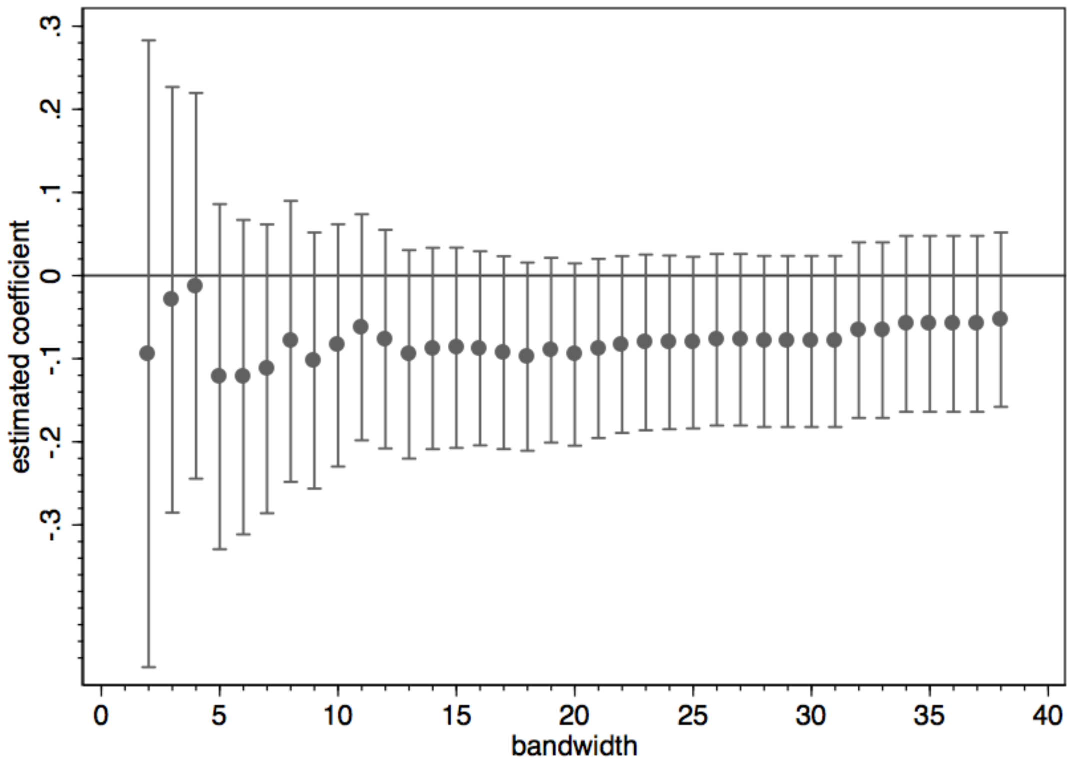

Notes: This figure shows the reduced-form estimates, i.e. the effects of crossing the cutoff on seventh-grade GPA, from regressions using different bandwidth sizes. Estimates are from the specification shown in Equation 2 with controls for student characteristics.

Figure A.8: LATE estimates of being tracked into higher-ability classrooms on student GPA across bandwidth

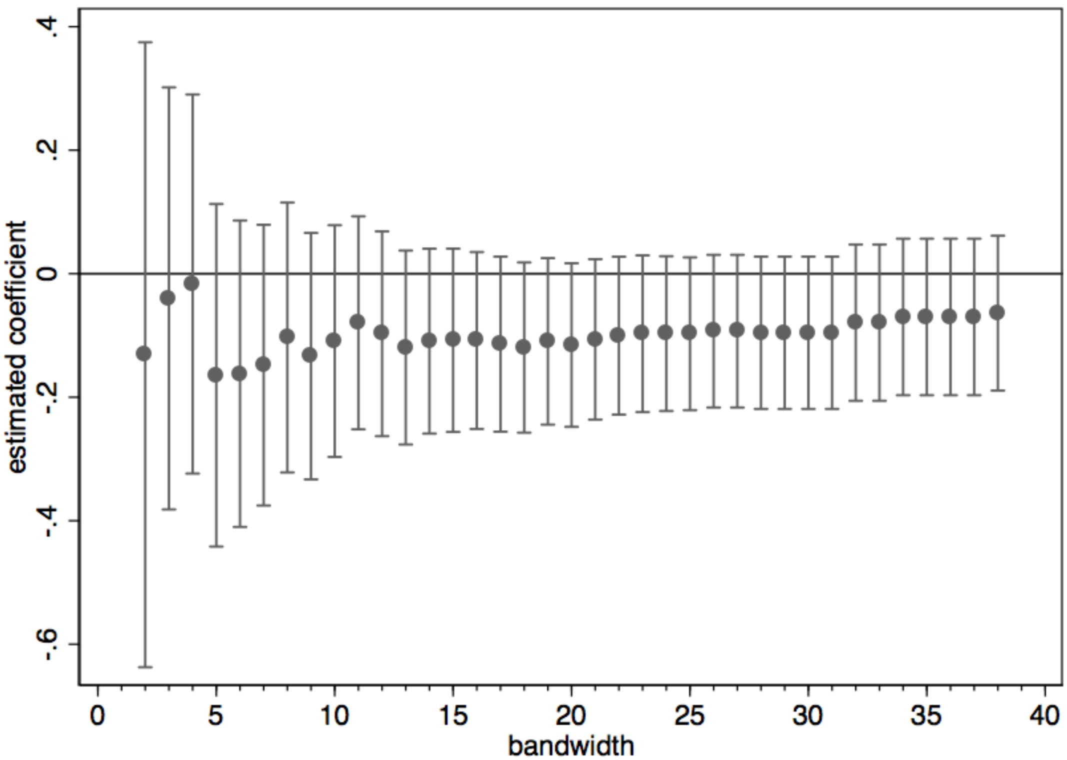

Notes: This figure shows the LATE estimates of being in the higher-ability classrooms on students seventh-grade GPA from regressions using different bandwidth sizes. Estimates are from the specification shown in Equation 2 with controls for student characteristics.

Figure A.9: LATE estimates of an increase of one s.d. in peer quality on student GPA across bandwidth

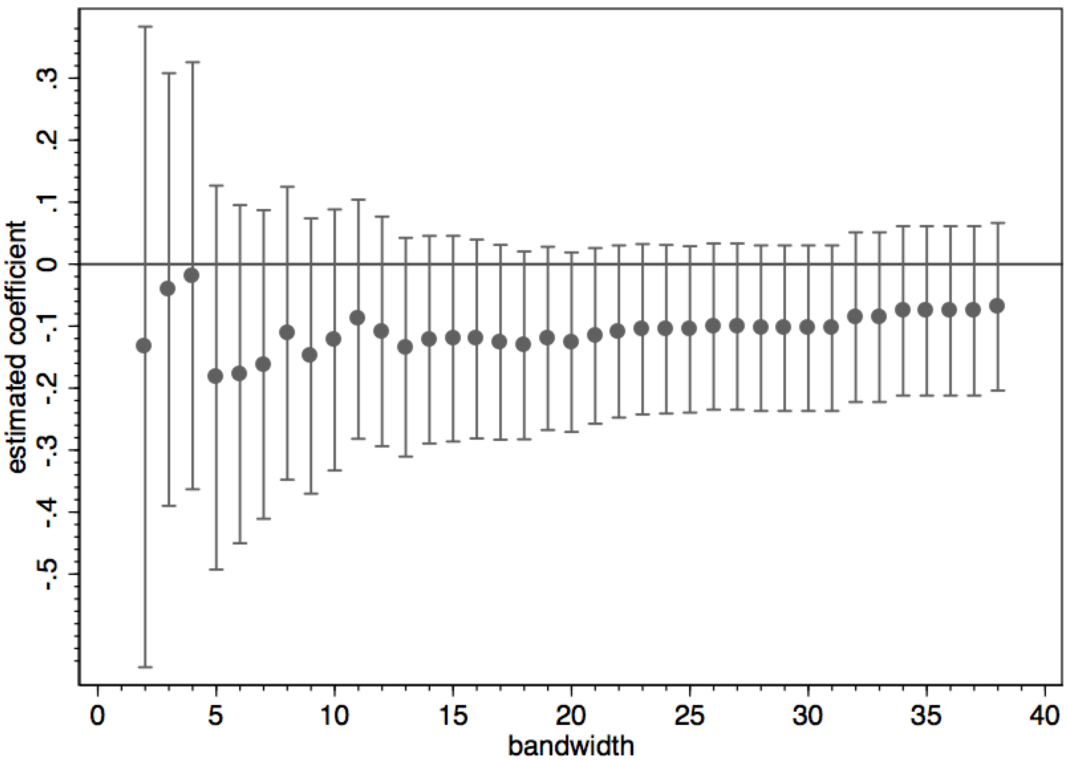

Notes: This figure shows the LATE estimates of being tracked into classrooms with 1 s.d. higher peer quality on student seventh-grade GPA from regressions using different bandwidth sizes. Estimates are from the specification shown in Equation 2 with controls for student characteristics.

Figure A.10: Observability of seventh-grade cumulative GPA across cutoff

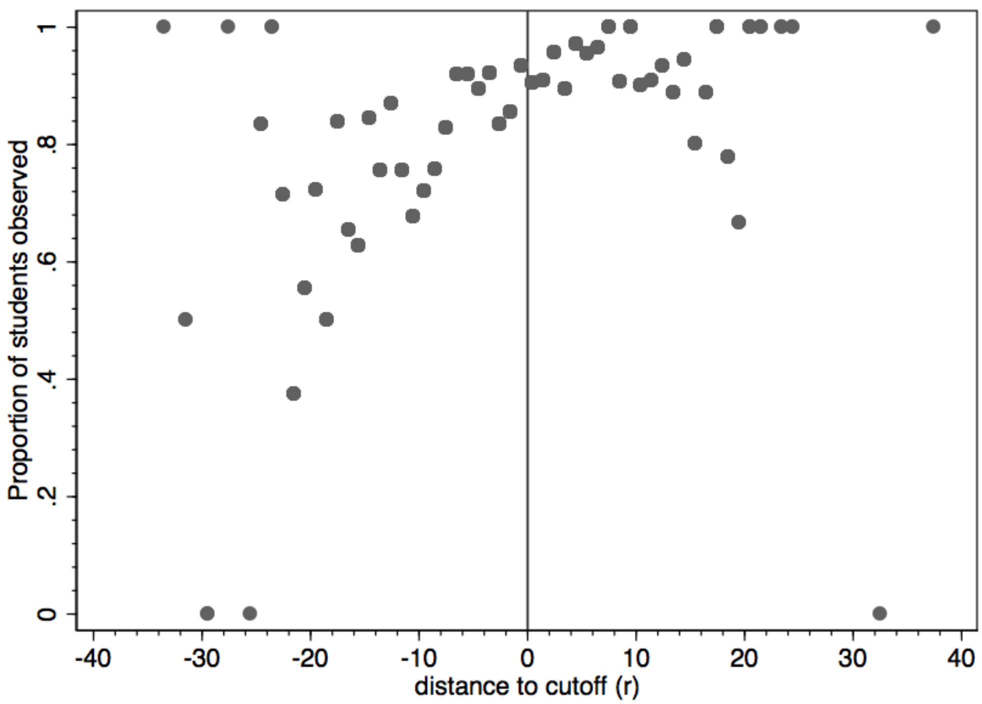

Notes: This figure plots the relationship between observability of seventh-grade cumulative GPA and the running variable. The running variable is the distance to cutoff, i.e. how far each student's preliminary score is from the cutoff. The number is positive if they score above the cutoff, and negative if they score below the cutoff.

Figure A.11: Student characteristics across cutoff (only students whose seventh-grade cumulative GPA is observed)

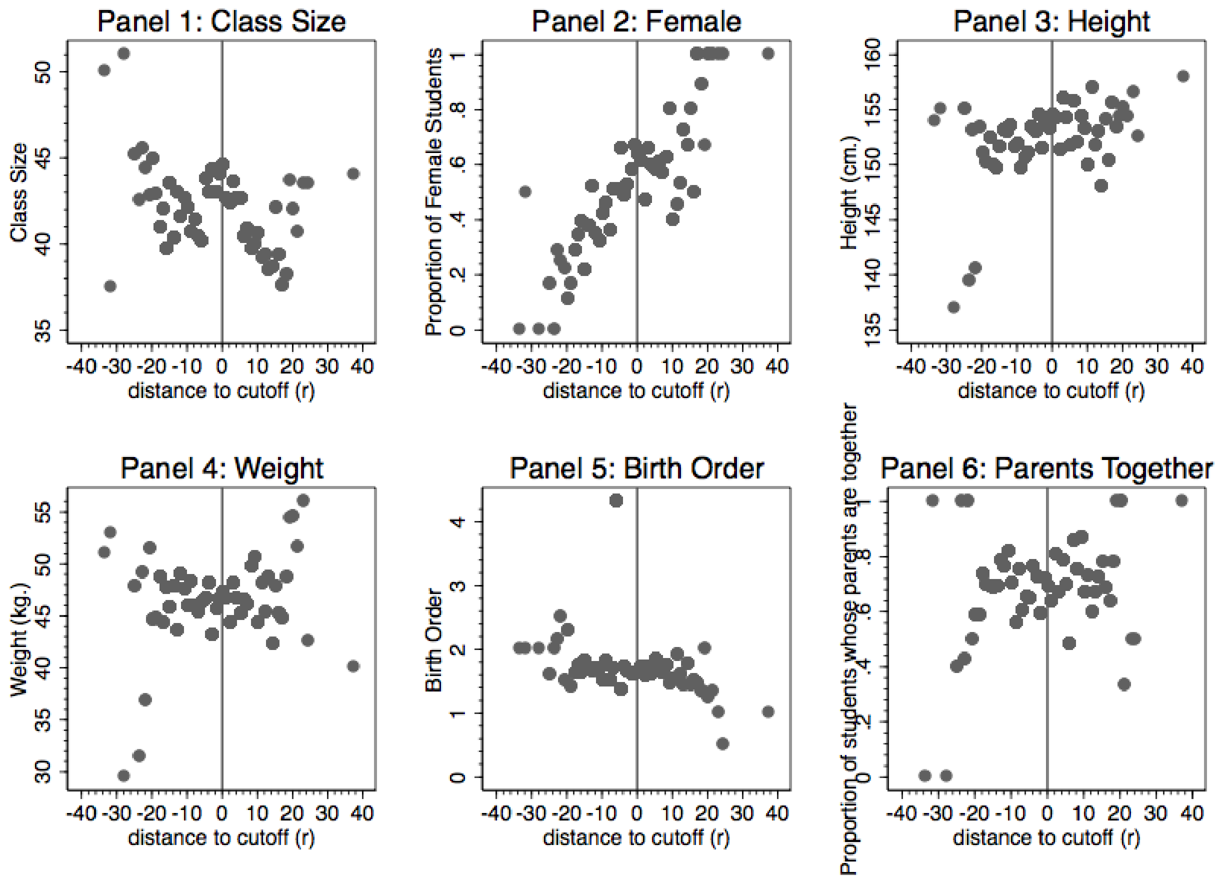

Notes: This figure plots student characteristics across the cutoff of students whose seventh-grade cumulative GPA is observed. The running variable is the distance to cutoff, i.e. how far each student's preliminary score is from the cutoff. The number is positive if they score above the cutoff, and negative if they score below the cutoff.

Figure A.12: Predicted seventh-grade cumulative GPA based on student characteristics (only students whose seventh-grade cumulative GPA is observed)

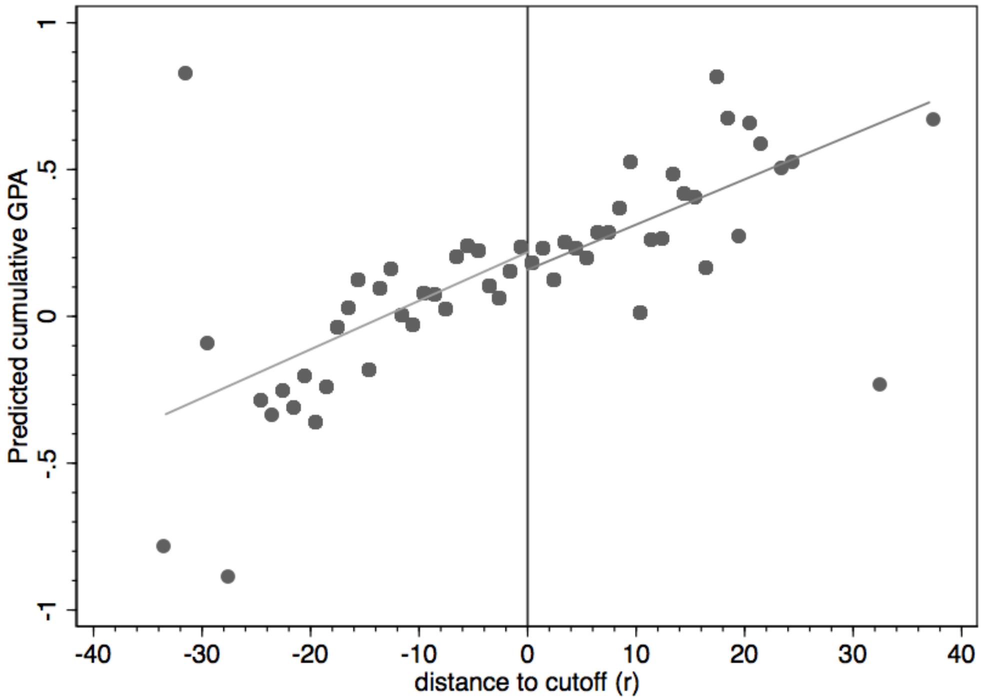

Notes: This figure plots predicted cumulative GPA across the cutoff of students whose seventh-grade cumulative GPA is observed. Predicted cumulative GPAs are based on the regression of (standardized) cumulative GPA on student characteristics.The running variable is the distance to cutoff, i.e. how far each student's preliminary score is from the cutoff. The number is positive if they score above the cutoff, and negative if they score below the cutoff.

Figure A.13: Seventh-grade cumulative GPA across cutoff by gender

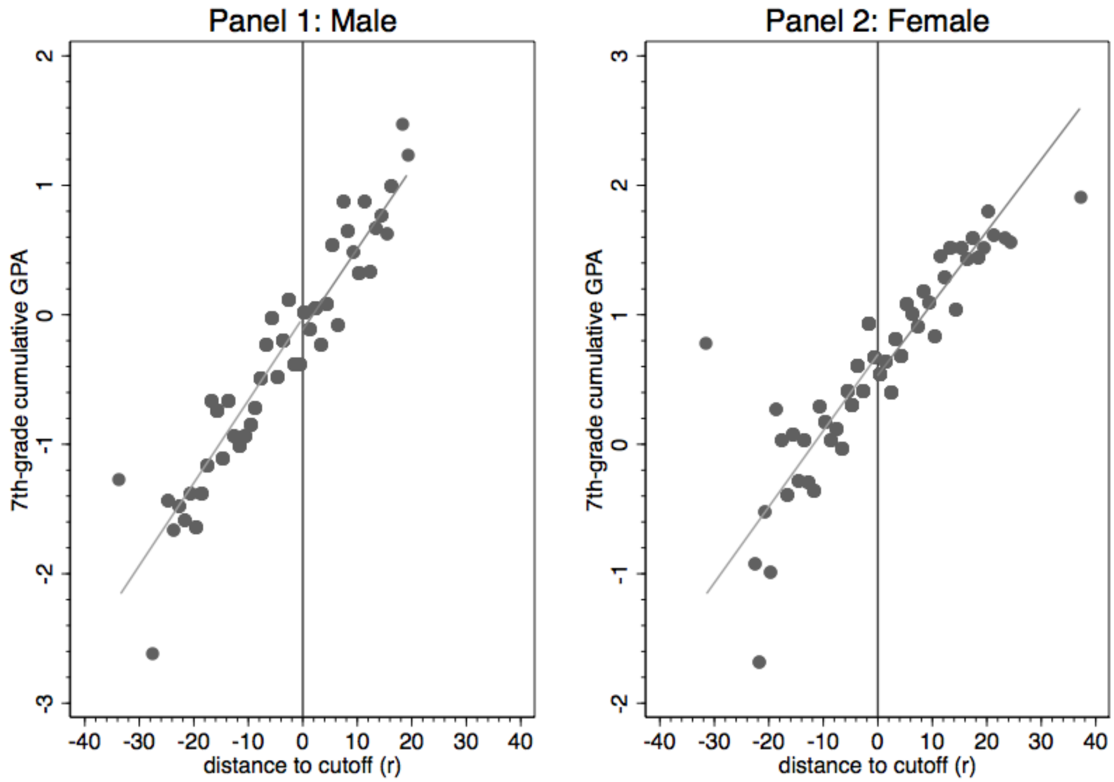

Notes: This figure plots the reduced-form relationship between seventh-grade cumulative GPA and the running variable for each gender separately. The running variable is the distance to cutoff, i.e. how far each student's preliminary score is from the cutoff. The number is positive if they score above the cutoff, and negative if they score below the cutoff.

A.2 Tables

Table A.1: Summary statistics

	(1) full sample	(2) $-20 < r < 20$	(3) $-10 < r < 10$	(4) $-5 < r < 5$
preliminary score	47.79	48.31	50.70	51.89
	(12.52)	(11.71)	(8.764)	(7.100)
distance to cutoff (r)	-2.971	-2.593	-0.952	-0.297
	(9.535)	(8.686)	(4.796)	(2.791)
class size	42.12	42.07	42.54	43.38
	(6.531)	(6.608)	(7.249)	(7.599)
female	0.511	0.517	0.563	0.585
weight (kg)	46.47	46.49	46.47	46.24
	(10.50)	(10.36)	(10.99)	(11.31)
height (cm)	152.6	152.7	153.1	153.6
	(9.045)	(8.974)	(8.775)	(8.788)
birth order	1.744	1.743	1.787	1.624
	(4.275)	(4.346)	(5.208)	(0.821)
Parents are together	0.697	0.701	0.697	0.711
7th-grade cumulative GPA	2.904	2.912	2.962	2.972
	(0.600)	(0.586)	(0.530)	(0.518)
standardized 7-th grade cumulative GPA	0.164	0.178	0.284	0.312
	(0.996)	(0.973)	(0.873)	(0.852)
Observations	1602	1543	1050	660

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table reports mean, standard deviation, and number of observations for each variable. Distance to cutoff, i.e. normalized preliminary score, is the running variable and is calculated based on Equation 1. r signifies the running variable, i.e. distance to cutoff. Therefore, each columns represent the estimates for each bandwidth choice. Standard deviations are in parentheses.

Table A.2: Regression discontinuity estimates of student characteristics

	(1) full sample	(2) $-30 < r < 30$	(3) $-20 < r < 20$	(4) $-10 < r < 10$	(5) $-5 < r < 5$
Class size	0.3427*	0.3958**	0.4903**	0.4193	0.3842
	(0.1871)	(0.1909)	(0.1978)	(0.2637)	(0.3737)
Female	-0.03921	-0.05341	-0.04857	-0.02265	0.004405
	(0.03833)	(0.03885)	(0.04147)	(0.05466)	(0.07627)
Weight (kg)	-0.09691	-0.2505	0.06910	0.04397	1.2388
	(0.8806)	(0.9108)	(0.9152)	(1.3042)	(1.8857)
Height (cm)	0.1375	0.2473	0.3475	-0.4088	0.8303
	(0.7188)	(0.7425)	(0.7836)	(0.9714)	(1.3944)
Birth Order	-0.1084	-0.1020	-0.1390	0.1620	-0.03075
	(0.2663)	(0.2618)	(0.2494)	(0.1822)	(0.1352)
Parents together	0.008322	0.02086	0.02942	-0.0001297	-0.004345
	(0.03786)	(0.03885)	(0.04094)	(0.05367)	(0.07503)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports the estimates of the discontinuity in each observed characteristic at the cutoff. The estimations use the regression discontinuity model shown in Equation 2 with each characteristics as the outcome. r signifies the running variable, i.e. distance to cutoff. Standard errors are in parentheses and are clustered at the student level. All regressions use rectangular kernel.

Table A.3: Regression discontinuity estimates for treatment (first stage)

	full sample		$-30 < r < 30$		$-20 < r < 20$		$-10 < r < 10$		$-5 < r < 5$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel 1. Probability of being in the higher-ability classroom										
Preliminary score	0.855***	0.834***	0.855***	0.831***	0.852***	0.823***	0.796***	0.770***	0.763***	0.739***
above or at cutoff	(0.0211)	(0.0190)	(0.0217)	(0.0195)	(0.0226)	(0.0203)	(0.0298)	(0.0275)	(0.0414)	(0.0387)
($r \geq 0$)										
N	1602	1542	1595	1536	1543	1489	1050	1023	660	643
Panel 2. Peer Quality: Average standardized preliminary exam score of peers										
Preliminary score	0.815***	0.794***	0.810***	0.789***	0.797***	0.774***	0.732***	0.712***	0.713***	0.698***
above or at cutoff	(0.0256)	(0.0259)	(0.0260)	(0.0264)	(0.0267)	(0.0272)	(0.0342)	(0.0346)	(0.0488)	(0.0489)
($r \geq 0$)										
N	1602	1542	1595	1536	1543	1489	1050	1023	660	643
Controls										
Cutoff fixed effects	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Student characteristics		Y		Y		Y		Y		Y

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table presents the first-stage estimations. Panel 1 reports the estimates of the discontinuity in the likelihood of being in the higher-ability classroom at the cutoff. Panel 2 reports the estimates of the discontinuity in peer quality at the cutoff. The estimations were conducted using the regression discontinuity model shown in Equation 2. Each student's peer quality is proxied by the average of their classmates' standardized preliminary exam scores. r signifies the running variable, i.e. distance to cutoff. Standard errors are in parentheses and are clustered at the student level. Student characteristics include class size, gender, height, weight, birth order, parents' relationship status. All regressions use rectangular kernel.

Table A.4: Reduced-form estimates

	$-30 < r < 30$			$-20 < r < 20$			$-10 < r < 10$			$-5 < r < 5$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Preliminary score above or at cutoff ($r \geq 0$)	-0.0908 (0.0577)	-0.0793 (0.0525)	-0.169 (0.111)	-0.102* (0.0613)	-0.0950* (0.0559)	-0.182 (0.112)	-0.0946 (0.0831)	-0.0841 (0.0744)	-0.189 (0.121)	-0.114 (0.121)	-0.122 (0.106)	-0.114 (0.132)
Controls												
Cutoff fixed effects	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Student characteristics		Y	Y		Y	Y		Y	Y		Y	Y
Teacher fixed effects			Y			Y			Y			Y
N	1366	1362	1362	1331	1328	1328	949	947	947	598	597	597

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table presents the reduced-form estimates, i.e. the discontinuity in seventh-grade cumulative GPA at the cutoff. The estimations were conducted using the regression discontinuity model shown in Equation 2. r signifies the running variable, i.e. distance to cutoff. Standard errors are in parentheses and are clustered at the student level. Student characteristics include class size, gender, height, weight, birth order, parents' relationship status. All regressions use rectangular kernel.

Table A.5: 2SLS estimates

	(1)	(2)	(3)	(4)
	$-30 < r < 30$	$-20 < r < 20$	$-10 < r < 10$	$-5 < r < 5$
Panel 1. Peer quality				
Being tracked into higher-ability classroom	0.9496***	0.9406***	0.9246***	0.9439***
	(0.02131)	(0.02199)	(0.02933)	(0.04030)
N	1536	1489	1023	643
Panel 2. Standardized 7th grade cumulative GPA				
Being tracked into higher-ability classroom	-0.09562	-0.1155*	-0.1091	-0.1645
	(0.06291)	(0.06756)	(0.09570)	(0.1415)
N	1362	1328	947	597
Panel 3. Standardized 7th grade cumulative GPA				
Peer quality increases by 1 s.d.	-0.1033	-0.1259*	-0.1223	-0.1830
	(0.06812)	(0.07386)	(0.1074)	(0.1580)
N	1362	1328	947	597
Controls				
Cutoff fixed effects	Y	Y	Y	Y
Student characteristics	Y	Y	Y	Y

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table presents the 2SLS estimates. Panel 1 reports the local average treatment effect (LATE) estimates of being tracked into the higher-ability classroom on peer quality. Panel 2 reports the LATE estimates of being tracked into the higher-ability classroom on seventh-grade cumulative GPA. And Panel 3 reports the LATE estimates of being tracked into classroom with 1-s.d.-higher peer quality on seventh-grade cumulative GPA. r signifies the running variable, i.e. distance to cutoff. Standard errors are in parentheses and are clustered at the student level. Student characteristics include class size, gender, height, weight, birth order, parents' relationship status. All regressions use rectangular kernel.

Table A.6: Regression discontinuity estimations of observability of seventh-grade cumulative GPA

	$-30 < r < 30$		$-20 < r < 20$		$-10 < r < 10$		$-5 < r < 5$	
	observed	observed	observed	observed	observed	observed	observed	observed
above or at cutoff ($r \geq 0$)	0.009454 (0.02570)	0.0007272 (0.02268)	0.01702 (0.02765)	0.005156 (0.02410)	0.01856 (0.03324)	0.009819 (0.02987)	-0.001476 (0.04326)	0.009649 (0.03929)
Controls								
Cutoff fixed effects	Y	Y	Y	Y	Y	Y	Y	Y
Student characteristics	Y	Y	Y	Y	Y	Y	Y	Y
Teacher fixed effects		Y		Y		Y		Y
N	1595	1536	1543	1489	1050	1023	660	643

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table presents the estimates of the discontinuity in observability of seventh-grade cumulative GPA at the cutoff. The estimations were conducted using the regression discontinuity model shown in Equation 2. Each student's peer quality is proxied by the average of their classmates' standardized preliminary exam scores. r signifies the running variable, i.e. distance to cutoff. Standard errors are in parentheses and are clustered at the student level. Student characteristics include class size, gender, height, weight, birth order, parents' relationship status. All regressions use rectangular kernel.

Table A.7: Regression discontinuity estimates of student characteristics (only students whose seventh-grade cumulative GPA is observed)

	(1) full sample	(2) $-30 < r < 30$	(3) $-20 < r < 20$	(4) $-10 < r < 10$	(5) $-5 < r < 5$
Class size	0.4230*	0.4534**	0.5730**	0.5428	0.4877
	(0.2057)	(0.2078)	(0.2158)	(0.2865)	(0.4065)
Female	-0.02082	-0.02440	-0.008761	0.004948	0.03803
	(0.04082)	(0.04138)	(0.04423)	(0.05771)	(0.08024)
Weight (kg)	-0.01964	-0.1493	0.07158	0.03222	1.3482
	(0.9425)	(0.9652)	(0.9685)	(1.3726)	(1.9759)
Height (cm)	0.4673	0.4777	0.5954	-0.2890	1.2295
	(0.7716)	(0.7897)	(0.8289)	(1.0105)	(1.4292)
Birth Order	-0.1468	-0.1479	-0.1559	0.1387	-0.03113
	(0.2982)	(0.2946)	(0.2800)	(0.1842)	(0.1398)
Parents together	-0.001599	0.003263	0.008496	-0.005200	-0.01253
	(0.03912)	(0.03984)	(0.04182)	(0.05457)	(0.07642)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports the estimates of the discontinuity in each observed characteristic at the cutoff using only the students whose seventh-grade cumulative GPA is observed. The estimations use the regression discontinuity model shown in Equation 2 with each characteristics as the outcome. r signifies the running variable, i.e. distance to cutoff. Standard errors are in parentheses and are clustered at the student level. All regressions use rectangular kernel.

Table A.8: Reduced-form estimates (including vs. not including gifted classrooms)

	$-30 < r < 30$			$-20 < r < 20$			$-10 < r < 10$			$-5 < r < 5$		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel 1: Without controls for gifted classrooms												
Preliminary score above or at cutoff ($r \geq 0$)	-0.0908	-0.0793	-0.169	-0.102*	-0.0950*	-0.182	-0.0946	-0.0841	-0.189	-0.114	-0.122	-0.114
	(0.0577)	(0.0525)	(0.111)	(0.0613)	(0.0559)	(0.112)	(0.0831)	(0.0744)	(0.121)	(0.121)	(0.106)	(0.132)
N	1366	1362	1362	1331	1328	1328	949	947	947	598	597	597
Panel 2: With control for gifted classrooms												
Preliminary score above or at cutoff ($r \geq 0$)	-0.0923	-0.0900	-0.169	-0.103*	-0.106*	-0.182	-0.0982	-0.0986	-0.189	-0.112	-0.125	-0.114
	(0.0576)	(0.0534)	(0.111)	(0.0612)	(0.0570)	(0.112)	(0.0831)	(0.0759)	(0.121)	(0.121)	(0.106)	(0.132)
N	1366	1362	1362	1331	1328	1328	949	947	947	598	597	597
Controls												
Cutoff fixed effects	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Student characteristics		Y	Y		Y	Y		Y	Y		Y	Y
Teacher fixed effects			Y			Y			Y			Y

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: As students in one of my schools could opt out of the assigned classroom and enroll in the gifted classroom after taking the preliminary exam, dropping students in the gifted classrooms from this school might incur selection issue. I, therefore, chose to keep them in my sample. This table shows that the decision to control or not control for the gifted classroom does not affect my results. Panel 1 reports the reduced-form estimates when I do not control for gifted classroom status. Panel 2 reports the reduced-form estimates when I control for whether a student is enrolled in the gifted classroom. The estimations were conducted using the regression discontinuity model shown in Equation 2. r signifies the running variable, i.e. distance to cutoff. Standard errors are in parentheses and are clustered at the student level. Student characteristics include class size, gender, height, weight, birth order, parents' relationship status. All regressions use rectangular kernel.

Table A.9: Reduced-form estimates by gender

	$-30 < r < 30$		$-20 < r < 20$		$-10 < r < 10$		$-5 < r < 5$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel 1. Male								
Preliminary score	-0.0861	-0.314*	-0.111	-0.344*	-0.0286	-0.312	0.0850	-0.0287
above or at cutoff	(0.0908)	(0.178)	(0.0936)	(0.180)	(0.136)	(0.203)	(0.209)	(0.256)
N	635	635	616	616	407	407	245	245
Panel 2. Female								
Preliminary score	-0.0578	-0.164	-0.102	-0.189	-0.143*	-0.188	-0.263**	-0.259
above or at cutoff	(0.0707)	(0.127)	(0.0703)	(0.124)	(0.0847)	(0.137)	(0.115)	(0.164)
N	730	730	712	712	540	540	352	352
Controls								
Cutoff fixed effects	Y	Y	Y	Y	Y	Y	Y	Y
Student characteristics	Y	Y	Y	Y	Y	Y	Y	Y
Teacher fixed effects		Y		Y		Y		Y

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

Notes: This table presents the reduced-form estimates, i.e. the discontinuity in seventh-grade cumulative GPA at the cutoff, for each gender separately. The estimations were conducted using the regression discontinuity model shown in Equation 2. r signifies the running variable, i.e. distance to cutoff. Standard errors are in parentheses and are clustered at the student level. Student characteristics include class size, height, weight, birth order, parents' relationship status.

APPENDIX B

FIGURES AND TABLES FOR SECTION THREE

B.1 Figures

Figure B.1: Number of television coverage on the topic of vaccines and its link to autism

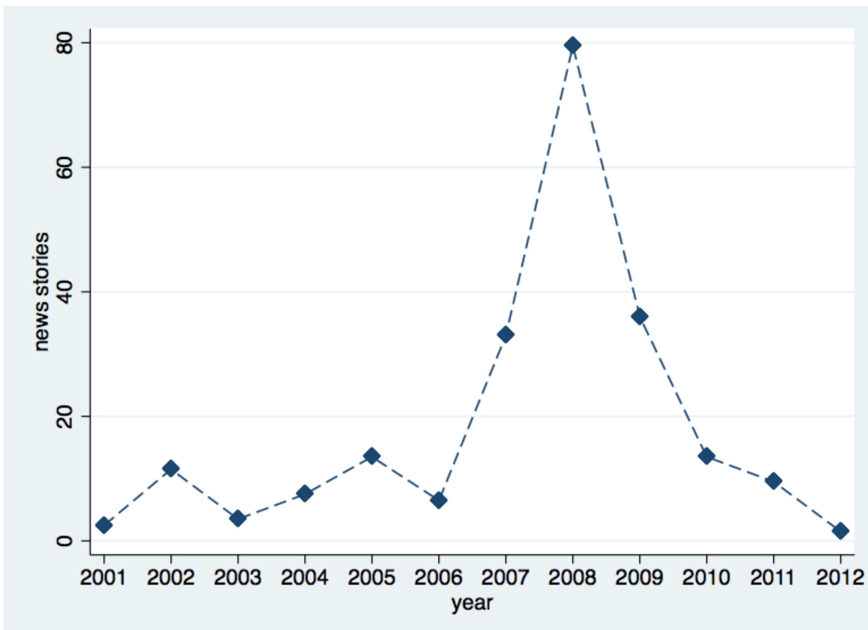

Notes: This figure demonstrates the number of news stories reporting on the alleged link between vaccines and autism without explicitly refuting it as false. The numbers are based on the coverage on 6 television networks: ABC, CBS, NBC, CNN, MSNBC, and Fox News. This figure is a visual representation of Table 1. Data source: LexisNexis.

Figure B.2: False news exposure from when child was born to 15 months old

(A) 19-23 months old at time of interview

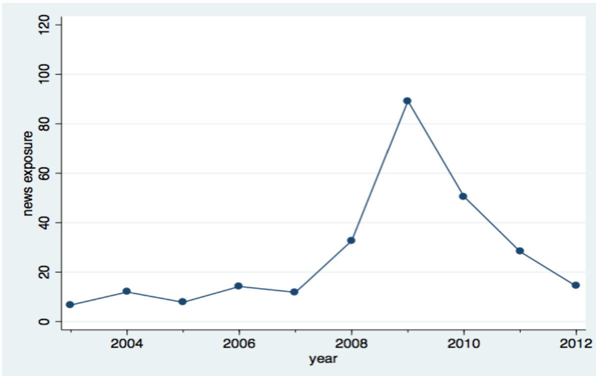

(B) 24-29 months old at time of interview

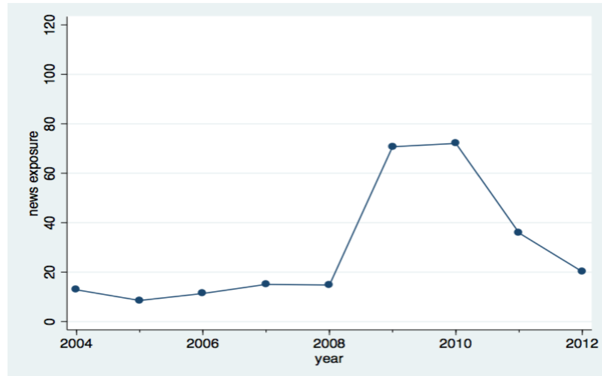

(C) 30-35 months old at time of interview

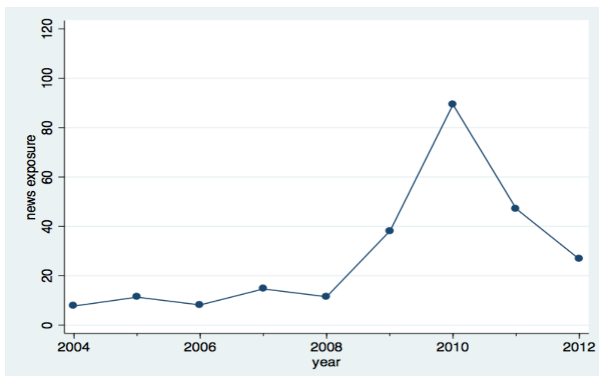

Notes: This figure shows the average false news exposure of parents in each interview cohort. Exposure is the number of false new stories that parents were exposed to from when the child was born to when the child was 15 months old. Data source: LexisNexis.

Figure B.3: False news exposure from when child was born to 29 months old

Child was 30-35 months old at time of interview

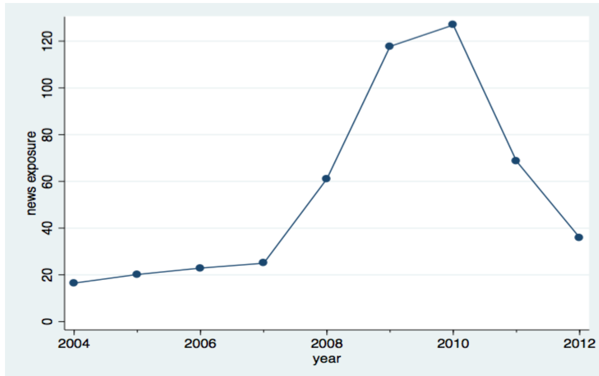

Notes: This figure shows the average false news exposure of parents in each interview cohort. Exposure is the number of false new stories that parents were exposed to from when the child was born to when the child was 29 months old. Data source: LexisNexis.

Figure B.4: Percent of parents who consent to the CDC obtaining vaccination record from healthcare providers

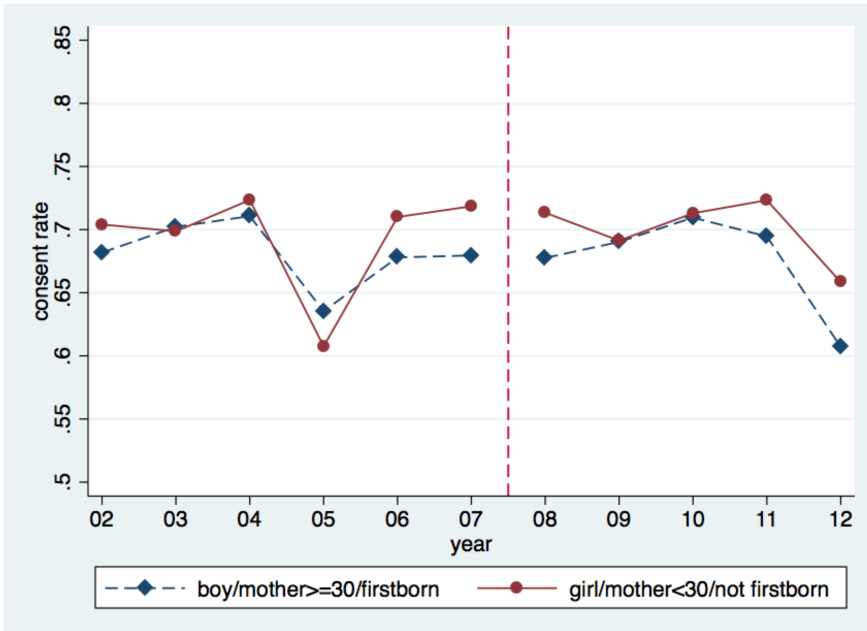

Notes: This figure shows the consent rate of parents in the control and treatment groups over time. The consent rate is the percent of the parents who were surveyed by the CDC who allowed the CDC to obtain vaccination data from healthcare providers. Data source: 2002-2012 National Immunization Surveys.

Figure B.5: MMR vaccine take-up rate at 15 months old

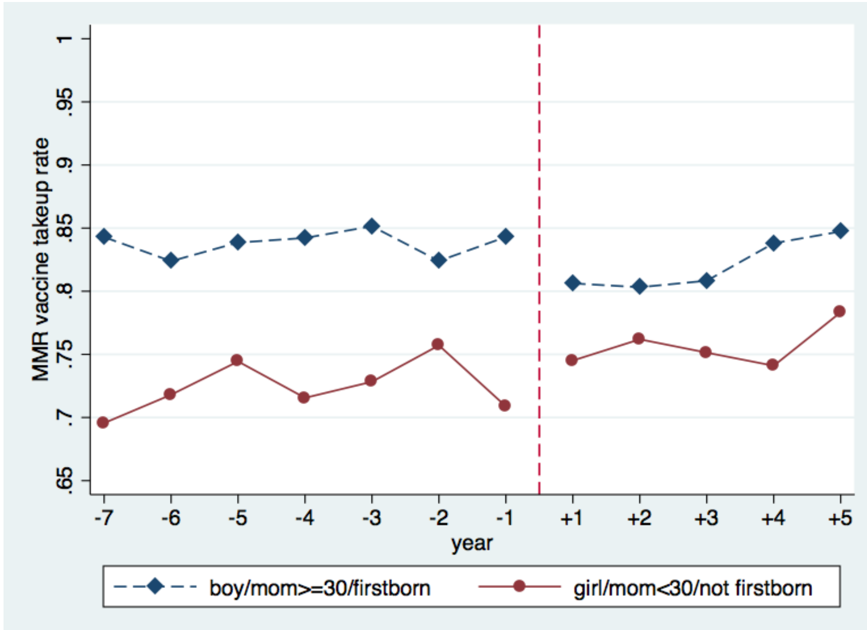

Notes: This figure shows the MMR vaccine take-up rate at 15 months old of children in the treatment and control groups. Treatment group is children with all 3 risk factors present, i.e. boys who are a firstborn and whose mom is older than or 30 years old. Control group is children with none of the risk factors present, i.e. girls who are not a firstborn and whose mother is under 30 years old. Data source: 2002-2012 National Immunization Surveys.

Figure B.6: MMR vaccine take-up rate at 29 months old

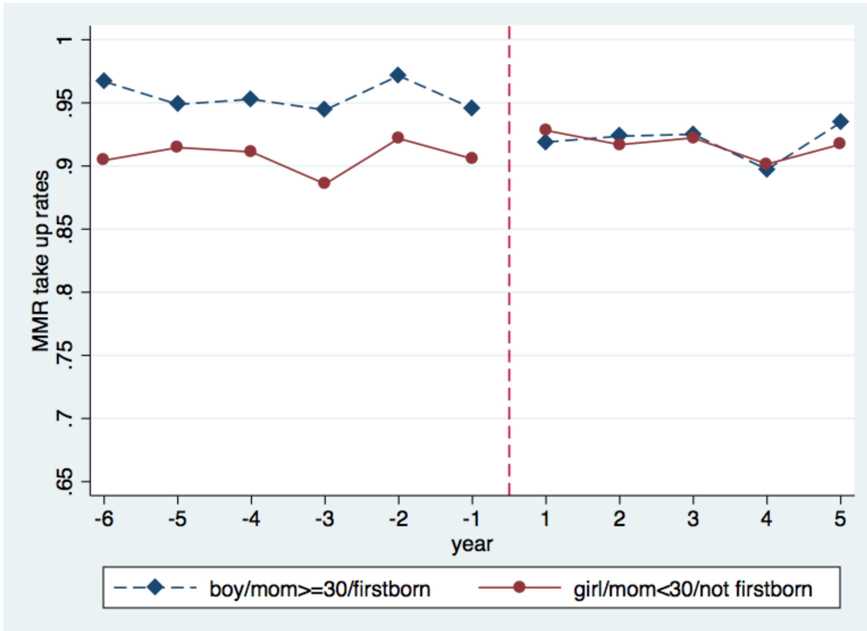

Notes: This figure shows the MMR vaccine take-up rate at 29 months old of children in the treatment and control groups. Treatment group is children with all 3 risk factors present, i.e. boys who are a firstborn and whose mom is older than or 30 years old. Control group is children with none of the risk factors present, i.e. girls who are not a firstborn and whose mother is under 30 years old. Data source: 2002-2012 National Immunization Surveys .

Figure B.7: Dynamic difference-in-differences estimates for MMR vaccine take-up rate at 15 months old

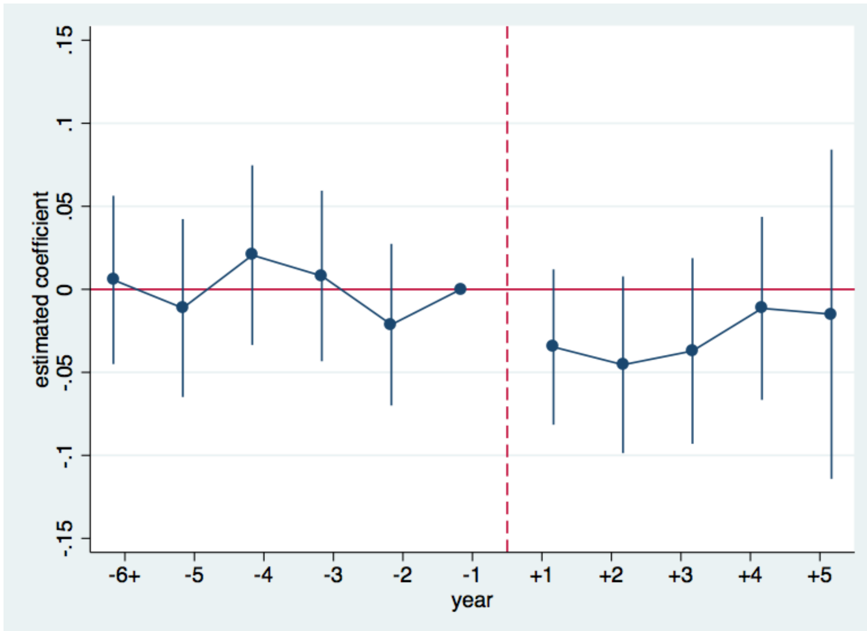

Notes: This figure shows the coefficients estimated from the dynamic difference-in-differences estimation for the MMR vaccine take-up at 15 months old. Treatment group is children with all 3 risk factors present, i.e. boys who are a firstborn and whose mom is older than or 30 years old. Control group is children with none of the risk factors present, i.e. girls who are not a firstborn and whose mother is under 30 years old. Data source: 2002-2012 National Immunization Surveys.

Figure B.8: Dynamic difference-in-differences estimates for MMR vaccine take-up rate at 29 months old

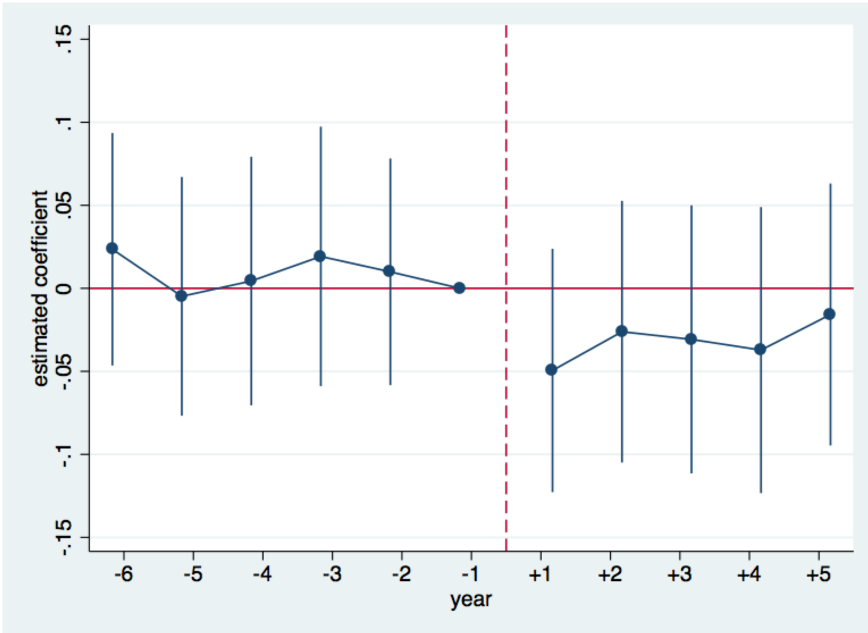

Notes: This figure shows the coefficients estimated from the dynamic difference-in-differences estimation for the MMR vaccine take-up at 29 months old. Treatment group is children with all 3 risk factors present, i.e. boys who are a firstborn and whose mom is older than or 30 years old. Control group is children with none of the risk factors present, i.e. girls who are not a firstborn and whose mother is under 30 years old. Data source: 2002-2012 National Immunization Surveys.

Figure B.9: Distribution of coefficients obtained from randomly reassigning treatment

(a) Take-up rate at 15 months old

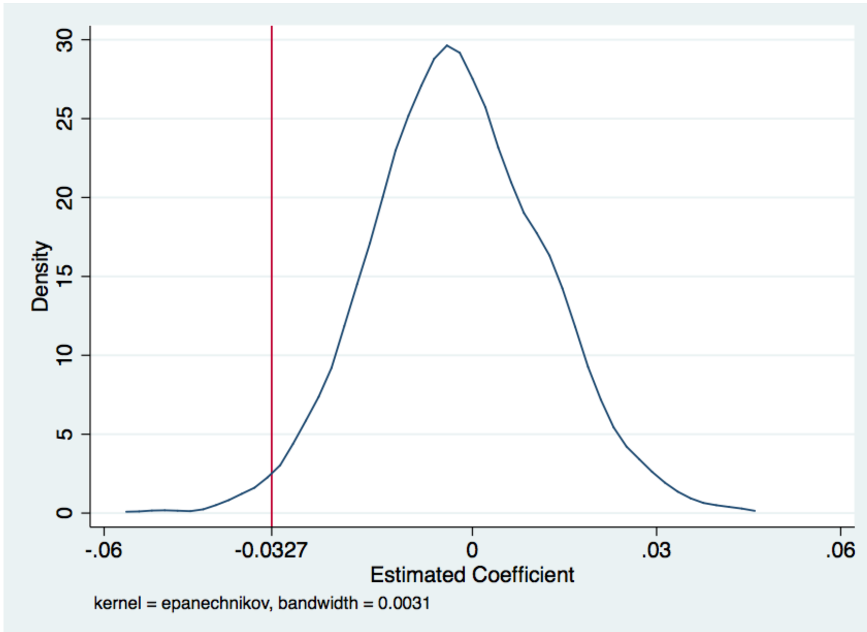

(b) Take-up rate at 29 months old

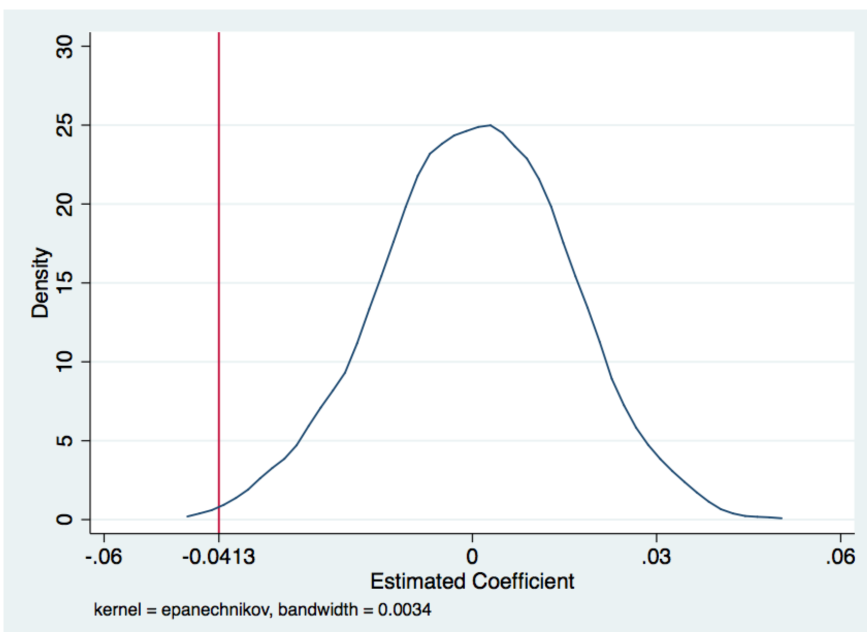

Notes: This figure shows the distribution of estimates obtained from a randomization exercise. Specifically, I randomly reassign child gender, mother's age, and firstborn status based on the true distribution of each variable in each year, and then estimate the effect (β) based on the reassignment. I do this for 1,000 replications and plot the distribution of the 1,000 coefficients estimated.Data source: 2002-2012 National Immunization Surveys.

Figure B.10: Estimated effects of misinformation on MMR vaccine take-up rates at age 15-29 months, using the main specification (3 risk factors vs. 0 risk factor present)

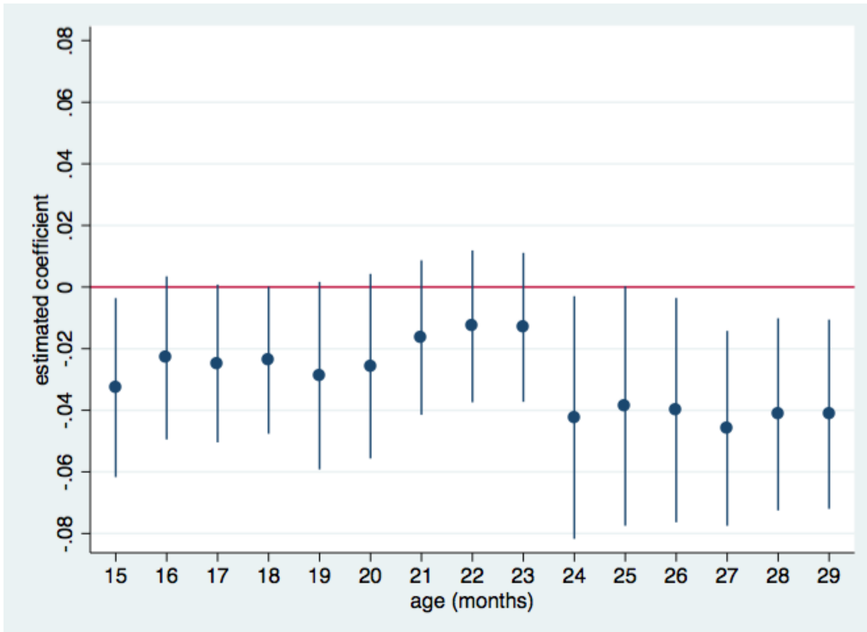

Notes: This figure shows the estimated coefficients for the effects at 15 month - 29 month. Data source: 2002-2012 National Immunization Surveys.

B.2 Table

Table B.1: Number of false news stories alleging the link between vaccines and autism

year	number of news stories
2001	2.5
2002	11.5
2003	3.5
2004	7.5
2005	13.5
2006	6.5
2007	**33**
2008	**79.5**
2009	36
2010	13.5
2011	9.5
2012	1.5

Notes: This table shows the number of news stories reporting on the alleged link between vaccines and autism without explicitly refuting it as false over the year. The numbers are based on the coverage on 6 television networks: ABC, CBS, NBC, CNN, MSNBC, and Fox News. Data source: LexisNexis.

Table B.2: Summary statistics

Panel 1: children 19-35 months old

	All		**Least Sensitive to Misinformation**		**Most Sensitive to Misinformation**	
	mean	sd	mean	sd	mean	sd
MMR shot at 15 months	0.78	(0.41)	0.74	(0.44)	0.83	(0.38)
Male	0.51	(0.50)	0		1	
Firstborn	0.43	(0.49)	0		1	
Mother $\geq$ 30	0.56	(0.50)	0		1	
White	0.73	(0.44)	0.68	(0.47)	0.76	(0.43)
Black	0.15	(0.36)	0.20	(0.40)	0.11	(0.31)
In poverty	0.31	(0.46)	0.50	(0.50)	0.13	(0.34)
Mother with college degree	0.31	(0.46)	0.09	(0.28)	0.56	(0.50)
Mother is married	0.68	(0.46)	0.52	(0.50)	0.83	(0.37)
19-23 months old	0.30	(0.46)	0.32	(0.47)	0.29	(0.46)
24-29 months old	0.34	(0.47)	0.34	(0.47)	0.34	(0.47)
30-35 months old	0.36	(0.48)	0.34	(0.47)	0.36	(0.48)
Moved state after birth	0.08	(0.27)	0.08	(0.27)	0.08	(0.28)
Observations	196684		16987		22239	

Panel 2: children 30-35 months old

	All		**Least Sensitive to Misinformation**		**Most Sensitive to Misinformation**	
	mean	sd	mean	sd	mean	sd
MMR shot at 29 months	0.93	(0.26)	0.91	(0.28)	0.94	(0.24)
Male	0.51	(0.50)	0		1	
Firstborn	0.42	(0.49)	0		1	
Mother $\geq$ 30	0.58	(0.49)	0		1	
White	0.73	(0.44)	0.68	(0.47)	0.76	(0.43)
Black	0.15	(0.36)	0.22	(0.41)	0.10	(0.30)
In poverty	0.30	(0.46)	0.49	(0.50)	0.13	(0.34)
Mother with college degree	0.31	(0.46)	0.09	(0.28)	0.54	(0.50)
Mother is married	0.69	(0.46)	0.51	(0.50)	0.82	(0.38)
Moved state after birth	0.09	(0.29)	0.09	(0.29)	0.09	(0.29)
Observations	70702		5655		8196	

Notes: All estimates obtained using sampling weights provided by the National Immunization Survey.The 'least sensitive to misinformation' group refers to girls who are not a firstborn and whose mother is <30 years old. The 'most sensitive to misinformation' group refers to boys who are a firstborn and whose mother is $\geq$ 30 years old. Data source: 2002-2012 National Immunization Surveys.

Table B.3: Effects of misinformation on parents consenting to the CDC acquiring vaccination record from healthcare provider

	(1) Consent
MostSensitive X Post	-0.0111
	(0.0138)
P-value	0.4207
N	56360

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

Notes: Robust standard errors in parentheses. Controls include state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group, mover status, and state's personal exemption law. All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table B.4: Effects of misinformation on MMR vaccine take-up rates

MMR vaccine take-up rate at 15 months old

	(1) MMR at 15 months	(2) MMR at 15 months	(3) MMR at 15 months
MostSensitive X Post	-0.0457***	-0.0327**	-0.0231
	(0.0144)	(0.0148)	(0.0183)
P-value	0.0015	0.0276	0.2054
Wild bootstrap p-value	0.0040	0.0260	0.0340
Randomization inference p-value		0.019	
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	39226	39226	39226

MMR vaccine take-up rate at 29 months old

	(1) MMR at 29 months	(2) MMR at 29 months	(3) MMR at 29 months
MostSensitive X Post	-0.0453***	-0.0413***	-0.0416**
	(0.0154)	(0.0157)	(0.0187)
P-value	0.0033	0.0084	0.0264
Wild bootstrap p-value	0.0030	0.0030	0.0190
Randomization inference p-value		0.003	
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	13851	13851	13851

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

Notes: Robust standard errors in parentheses. Wild bootstrap p-values are obtained using the method explained in Cameron, Gelbach and Miller (2008). Controls include state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group, mover status, and state's personal exemption law. All regressions are estimated using the sampling weights provided by the National Immunization Surveys. Data source: 2002-2012 National Immunization Surveys.

Table B.5: Subgroup analysis by mother's education

MMR vaccine take-up rate at 15 months old

	(1)	(2)	(3)
	All	College degree	No college degree
MostSensitive X Post	-0.0327**	-0.0446	-0.0010
	(0.0148)	(0.0285)	(0.0199)
P-value	0.0276	0.1183	0.9592
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Controls X Year			
N	39226	17671	21555

MMR vaccine take-up rate at 29 months old

	(1)	(2)	(3)
	All	College degree	No college degree
MostSensitive X Post	-0.0413***	-0.0425	-0.0335
	(0.0157)	(0.0349)	(0.0218)
P-value	0.0084	0.2239	0.1242
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Controls X Year			
N	13851	6300	7551

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

Notes: Robust standard errors in parentheses. Controls include state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group, mover status, and state's personal exemption law. All regressions are estimated using the sampling weights provided by the National Immunization Surveys. Data source: 2002-2012 National Immunization Surveys.

Table B.6: Effects of misinformation on MMR vaccine take-up rates at 15-29 months old

	15 months	16 months	17 months	18 months	19 months
MostSensitive X Post	-0.0327**	-0.0230*	-0.0248*	-0.0238*	-0.0288*
	(0.0148)	(0.0135)	(0.0131)	(0.0122)	(0.0155)
P-value	0.0276	0.0886	0.0575	0.0510	0.0641
Wild bootstrap p-value	0.0260	0.0260	0.0701	0.0491	0.1572
Outcome mean	0.78	0.82	0.84	0.87	0.88
Year FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
N	39226	39226	39226	39226	27504

	20 months	21 months	22 months	23 months	24 months
MostSensitive X Post	-0.0257*	-0.0164	-0.0128	-0.0131	-0.0424**
	(0.0153)	(0.0128)	(0.0126)	(0.0123)	(0.0201)
P-value	0.0928	0.1996	0.3096	0.2888	0.0350
Wild bootstrap p-value	0.1882	0.3333	0.4284	0.4495	0.0130
Outcome mean	0.89	0.89	0.90	0.90	0.91
Year FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
N	27504	27504	27504	27504	13851

	25 months	26 months	27 months	28 months	29 months
MostSensitive X Post	-0.0386*	-0.0400**	-0.0459***	-0.0413***	-0.0413***
	(0.0198)	(0.0186)	(0.0161)	(0.0159)	(0.0157)
P-value	0.0516	0.0314	0.0045	0.0094	0.0084
Wild bootstrap p-value	0.0220	0.0080	0.0020	0.0040	0.0030
Outcome mean	0.91	0.92	0.92	0.92	0.93
Year FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
N	13851	13851	13851	13851	13851

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

Notes: Robust standard errors in parentheses. Wild bootstrap p-values are obtained using the method explained in Cameron, Gelbach and Miller (2008). Controls include state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group, mover status, and state's personal exemption law. All regressions are estimated using the sampling weights provided by the National Immunization Surveys. All estimates are obtained using the main specification, i.e. difference-in-difference with year fixed effects, state fixed effects, observable controls. Data source: 2002-2012 National Immunization Surveys.

Table B.7: Effects of misinformation on MMR vaccine take-up rate at 15 months old with more loosely defined treatment and control groups

	Baseline	Increase control group		Increase treated group		Only using 2 characteristics to define treatment group		
	3 characteristics vs. 0 characteristic	3 characteristics vs. 0/1 characteristic	3 characteristics vs. 0/1/2 characteristics	3/2 characteristics vs. 0 characteristic	3/2/1 characteristics vs. 0 characteristic	boy & mother≥30 vs. girl & mother<30	boy & firstborn vs. girl & not firstborn	mother≥30 & firstborn vs. mother<30 & not firstborn
MostSensitive x Post	-0.0327**	-0.0128	-0.0127	-0.0271***	-0.0250***	-0.0154	-0.0039	-0.0284***
	(0.0148)	(0.0110)	(0.0103)	(0.0102)	(0.0089)	(0.0098)	(0.0093)	(0.0106)
P-value	0.0276	0.2455	0.2186	0.0078	0.0050	0.1178	0.6777	0.0071
Wild bootstrap p-value	0.0260	0.1291	0.1291	0.0110	0.0050	0.2322	0.6587	0.0220
Outcome Mean	0.78	0.78	0.78	0.78	0.78	0.78	0.78	0.78
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	39226	114537	196684	121373	196684	98830	98013	78293

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

Notes: Robust standard errors in parentheses. Wild bootstrap p-values are obtained using the method explained in Cameron, Gelbach and Miller (2008). Controls include state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group, mover status, and state's personal exemption law. All regressions are estimated using the sampling weights provided by the National Immunization Surveys. Data source: 2002-2012 National Immunization Surveys.

Table B.8: Effects of misinformation on MMR vaccine take-up rate at 29 months old with more loosely defined treatment and control groups

	Baseline	Increase control group		Increase treated group		Only using 2 characteristics to define treatment group		
	3 characteristics vs. 0 characteristic	3 characteristics vs. 0/1 characteristic	3 characteristics vs. 0/1/2 characteristics	3/2 characteristics vs. 0 characteristic	3/2/1 characteristics vs. 0 characteristic	boy & mother≥30 vs. girl & mother<30	boy & firstborn vs. girl & not firstborn	mother≥30 & firstborn vs. mother<30 & not firstborn
MostSensitive X Post	-0.0413***	-0.0275**	-0.0238**	-0.0253**	-0.0223*	-0.0226**	-0.0153	-0.0277**
	(0.0157)	(0.0114)	(0.0106)	(0.0128)	(0.0124)	(0.0106)	(0.0100)	(0.0109)
P-value	0.0084	0.0158	0.0249	0.0481	0.0726	0.0335	0.1281	0.0114
Wild bootstrap p-value	0.0030	0.0360	0.0220	0.0100	0.0240	0.1071	0.0500	0.0060
Outcome mean	0.93	0.93	0.93	0.93	0.93	0.93	0.93	0.93
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	13851	40628	70702	43925	70702	35649	35096	27659

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

Notes: Robust standard errors in parentheses. Wild bootstrap p-values are obtained using the method explained in Cameron, Gelbach and Miller (2008). Controls include state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group, mover status, and state's personal exemption law. All regressions are estimated using the sampling weights provided by the National Immunization Surveys. Data source: 2002-2012 National Immunization Surveys.

Table B.9: Effects of misinformation on MMR vaccine take-up rates using logistic regression

	MMR at 15 months	MMR at 15 months	MMR at 15 months
outcome			
MostSensitive X Post	-0.2814***	-0.2112**	-0.1371
	(0.0877)	(0.0910)	(0.1122)
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	39226	39226	39226

	MMR at 29 months	MMR at 29 months	MMR at 29 months
outcome			
MostSensitive X Post	-0.7405***	-0.7010***	-0.6761**
	(0.2243)	(0.2285)	(0.2693)
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	13851	13851	13851

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

Notes: Robust standard errors in parentheses. All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table B.10: Effects of misinformation on MMR vaccine take-up rates: data from 2001-2012

	(1) MMR at 15 months	(2) MMR at 15 months	(3) MMR at 15 months
MostSensitive X Post	-0.0499***	-0.0375***	-0.0301*
	(0.0141)	(0.0144)	(0.0178)
P-value	0.0004	0.0093	0.0915
Wild bootstrap p-value	0.0040	0.0260	0.0340
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	43548	43548	43548

	(1) MMR at 29 months	(2) MMR at 29 months	(3) MMR at 29 months
MostSensitive X Post	-0.0465***	-0.0433***	-0.0438**
	(0.0151)	(0.0153)	(0.0184)
P-value	0.0020	0.0047	0.0170
Wild bootstrap p-value	0.0030	0.0030	0.0190
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	15285	15285	15285

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

Notes: Robust standard errors in parentheses. Controls include state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group, mover status, and state's personal exemption law. All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table B.11: Effects of misinformation on MMR vaccine take-up rates: data from 2003-2012

	(1) MMR at 15 months	(2) MMR at 15 months	(3) MMR at 15 months
MostSensitive X Post	-0.0406***	-0.0279*	-0.0198
	(0.0148)	(0.0153)	(0.0188)
P-value	0.0062	0.0682	0.2930
Wild bootstrap p-value	0.0040	0.0260	0.0340
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	35243	35243	35243

	(1) MMR at 29 months	(2) MMR at 29 months	(3) MMR at 29 months
MostSensitive X Post	-0.0425***	-0.0401**	-0.0366*
	(0.0162)	(0.0165)	(0.0195)
P-value	0.0087	0.0154	0.0611
Wild bootstrap p-value	0.0030	0.0030	0.0190
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	12492	12492	12492

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

Notes: Robust standard errors in parentheses. Controls include state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group, mover status, and state's personal exemption law. All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table B.12: Effects of misinformation on MMR vaccine take-up rates: data from 2004-2012

	(1) MMR at 15 months	(2) MMR at 15 months	(3) MMR at 15 months
MostSensitive X Post	-0.0442***	-0.0339**	-0.0258
	(0.0154)	(0.0159)	(0.0195)
P-value	0.0042	0.0332	0.1869
Wild bootstrap p-value	0.0040	0.0260	0.0340
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	31241	31241	31241

	(1) MMR at 29 months	(2) MMR at 29 months	(3) MMR at 29 months
MostSensitive X Post	-0.0451***	-0.0443**	-0.0370*
	(0.0173)	(0.0175)	(0.0206)
P-value	0.0092	0.0115	0.0721
Wild bootstrap p-value	0.0030	0.0030	0.0190
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	11081	11081	11081

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

Notes: Robust standard errors in parentheses. Controls include state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group, mover status, and state's personal exemption law. All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table B.13: Effects of misinformation on MMR vaccine take-up rates: data from 2002-2013

	(1) MMR at 15 months	(2) MMR at 15 months	(3) MMR at 15 months
MostSensitive X Post	-0.0320**	-0.0237	-0.0216
	(0.0141)	(0.0145)	(0.0177)
P-value	0.0230	0.1027	0.2206
Wild bootstrap p-value	0.0040	0.0260	0.0340
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	41823	41823	41823

	(1) MMR at 29 months	(2) MMR at 29 months	(3) MMR at 29 months
MostSensitive X Post	-0.0381**	-0.0346**	-0.0378**
	(0.0152)	(0.0152)	(0.0175)
P-value	0.0119	0.0226	0.0306
Wild bootstrap p-value	0.0030	0.0030	0.0190
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	14890	14890	14890

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

Notes: Robust standard errors in parentheses. Controls include state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group, mover status, and state's personal exemption law. All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

Table B.14: Effects of misinformation on MMR vaccine take-up rates: data from 2002-2014

	(1)	(2)	(3)
	MMR at 15 months	MMR at 15 months	MMR at 15 months
MostSensitive X Post	-0.0315**	-0.0237*	-0.0232
	(0.0134)	(0.0139)	(0.0167)
P-value	0.0191	0.0884	0.1646
Wild bootstrap p-value	0.0040	0.0260	0.0340
Outcome mean	0.78	0.78	0.78
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	44542	44542	44542

	(1)	(2)	(3)
	MMR at 29 months	MMR at 29 months	MMR at 29 months
MostSensitive X Post	-0.0372**	-0.0351**	-0.0413**
	(0.0145)	(0.0147)	(0.0169)
P-value	0.0104	0.0167	0.0145
Wild bootstrap p-value	0.0030	0.0030	0.0190
Outcome mean	0.93	0.93	0.93
Year FE	Yes	Yes	Yes
Controls		Yes	Yes
Controls X Year			Yes
N	15994	15994	15994

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

Notes: Robust standard errors in parentheses. Controls include state fixed effects, race, poverty status, mother's education, mother's marital status, child's age group, mover status, and state's personal exemption law. All regressions are estimated using the sampling weights provided by the National Immunization Surveys.

APPENDIX C

FIGURES AND TABLES FOR SECTION FOUR

C.1 Figures

Figure C.1: Map of mosque openings in Detroit and Hamtramck

Notes: This figure shows the new mosques (black dots) we study in this paper, and illustrates the spatial difference-in-differences framework. For example, consider the mosque inside the blue and red circles. The properties inside the dark blue circle are in the “treated zone” while the properties inside the red circle but outside the blue circle are in the comparison group.

Figure C.2: Dynamic difference-in-difference estimates of mosque openings on log property prices

Panel A: Treated Radius = 0.20 mile

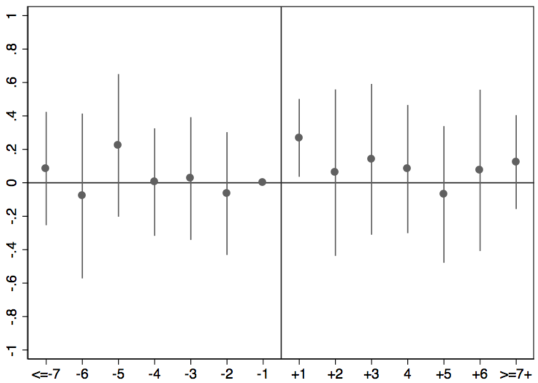

Panel B: Treated Radius = 0.15 mile

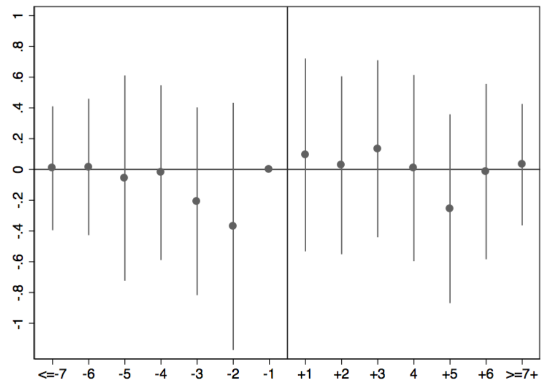

Panel C: Treated Radius = 0.25 mile

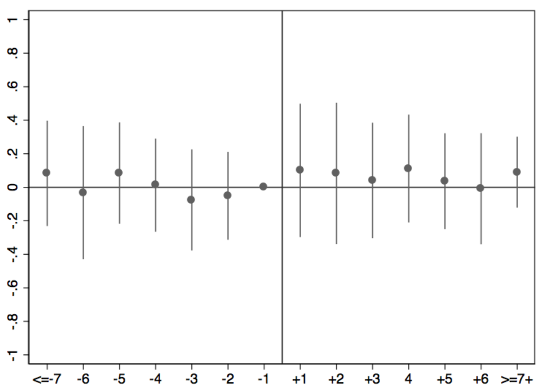

Panel D: Treated Radius = 0.30 mile

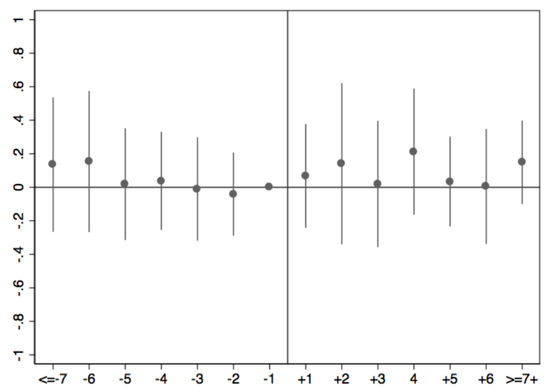

Notes: This figure presents the dynamic difference-in-differences estimates of mosque openings on the log of property prices. The X-axis shows quarter relative to mosque opening. Each dot plots δ_l, i.e. the coefficient of the treated zone indicator interacted with the relative quarter dummy during which the property sale took place, from Equation 4.3. The sample only includes sales transactions data of properties that were sold at least twice between 2002 and 2018 (*repeat sales*). The sample radius is two times the size of the treated radius. Standard errors in parentheses are clustered at the level of the neighborhood surrounding the newly opened church.

Figure C.3: Estimated effects of mosque openings and church openings, by definition of treated radius and type of property

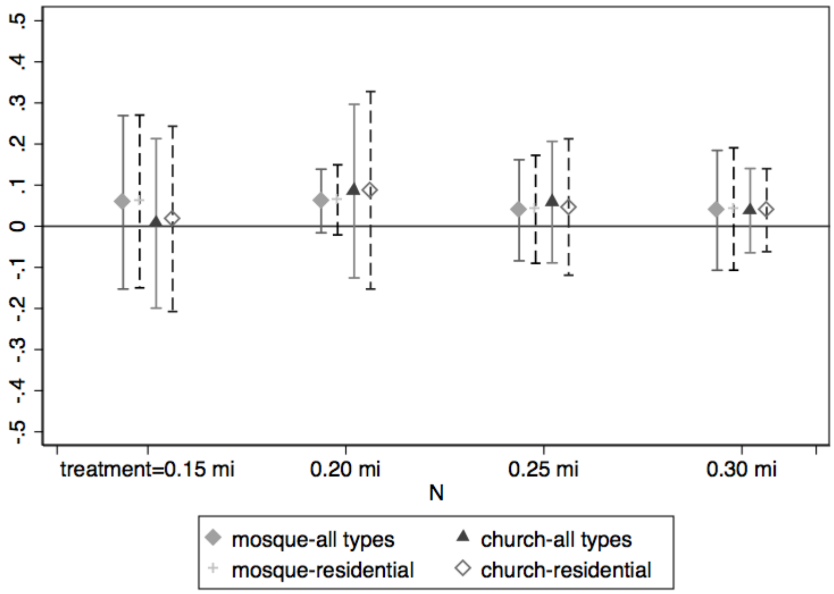

Notes: This figure presents the estimated effects of mosque openings and church openings on log property prices across property type and treatment and sample radius choice. All estimations use the *repeat-sales specification* in Equation 4.2 with only sales transactions data of properties that were sold at least twice between 2002 and 2018. Sample radius are twice the size of the treated radius. Standard errors are clustered at the level of the neighborhood surrounding the newly opened mosque (or church).

C.2 Tables

Table C.1: Mosque and church openings between 2002-2018

Mosque		Church	
Year	Openings	Year	Openings
2003	2	2002	1
2004	1	2003	2
2005	1	2004	2
2006	1	2005	2
2008	1	2006	1
2010	1	2007	1
2011	1	2008	1
2012	1	2011	1
2016	1	2013	1
		2014	1
Total	10	Total	13

Notes: The table shows the number of mosque and church openings each year during the study period (2002-2018).

Table C.2: Property sales data

Panel A: All sales

	All		Controls		Treated	
	mean	sd	mean	sd	mean	sd
sale price	58,660.09	(206872.89)	51,550.32	(149947.56)	68,855.76	(267878.54)
residential	0.97	(0.18)	0.97	(0.17)	0.96	(0.19)
acreage	0.13	(1.01)	0.10	(0.64)	0.18	(1.38)
front	32.00	(45.80)	29.79	(29.24)	35.18	(62.15)
depth	90.99	(52.86)	89.80	(47.86)	92.70	(59.26)
floor area	1,563.25	(5192.65)	1,527.27	(2853.93)	1,614.97	(7350.03)
Observations	11771		6935		4836	

Panel B: Only repeat sales

	All		Controls		Treated	
	mean	sd	mean	sd	mean	sd
sale price	60,278.95	(238829.22)	52,325.91	(171705.99)	71,716.25	(310590.37)
residential	0.98	(0.14)	0.98	(0.14)	0.98	(0.14)
acreage	0.13	(1.12)	0.10	(0.74)	0.19	(1.50)
front	31.33	(43.95)	29.35	(29.77)	34.19	(58.50)
depth	90.00	(51.43)	89.04	(47.22)	91.39	(56.91)
floor area	1,549.24	(5624.59)	1,510.11	(2470.09)	1,605.34	(8262.06)
Observations	8548		5042		3506	

Notes: This table shows the summary statistics of property sales transactions used for the analysis. Panel 1 shows the summary statistics of *all sales transactions* that happened between 2002 and 2018 of the properties that are within the 0.4 mile radius from the new mosques. Panel 2 shows the summary statistics of the sales transactions of the properties that are within 0.4 mile radius from the new mosques and were *sold at least twice* between 2002 and 2018. Treated properties are those within the treatment radius which is 0.2 miles, while the comparison properties are between 0.2 and 0.4 mile from the new mosques.

Table C.3: Effects of mosque and church openings on the log of property prices (all property types)

Panel A: Effects of mosque openings

	treated radius=0.20 mile			0.15 mile			0.25 mile			0.30 mile		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	repeat sales	all sales	all sales	repeat sales	all sales	all sales	repeat sales	all sales	all sales	repeat sales	all sales	all sales
treated x post	0.0614	0.103	0.0547	0.0583	0.159	0.0637	0.0388	0.00118	0.00101	0.0388	-0.0131	-0.0292
	(0.0395)	(0.0546)	(0.0373)	(0.108)	(0.0921)	(0.0711)	(0.0627)	(0.0386)	(0.0324)	(0.0744)	(0.0375)	(0.0345)
95% CI	[-0.020, 0.142]	[-0.009, 0.214]	[-0.022, 0.131]	[-0.165, 0.282]	[-0.032, 0.350]	[-0.084, 0.211]	[-0.089, 0.167]	[-0.077, 0.080]	[-0.065, 0.067]	[-0.111, 0.189]	[-0.089, 0.063]	[-0.099, 0.040]
Bootstrap 95% CI	[-0.008, 0.141]	[-0.026, 0.235]	[-0.038, 0.139]	[-0.167, 0.253]	[-0.035, 0.376]	[-0.077, 0.242]	[-0.065, 0.154]	[-0.082, 0.096]	[-0.079, 0.089]	[-0.098, 0.172]	[-0.105, 0.075]	[-0.117, 0.061]
Year-month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property fixed effects	Yes			Yes			Yes			Yes		
Area fixed effects		Yes	Yes		Yes	Yes		Yes	Yes		Yes	Yes
Controls 1		Yes			Yes			Yes			Yes	
Controls 2			Yes			Yes			Yes			Yes
N	8548	11771	7432	6373	8788	5558	10543	14560	9271	11879	16552	10736

Panel B: Effects of church openings

	treated radius=0.20 mile			0.15 mile			0.25 mile			0.30 mile		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	repeat sales	all sales	all sales	repeat sales	all sales	all sales	repeat sales	all sales	all sales	repeat sales	all sales	all sales
treated x post	0.0855	0.0553*	0.0318*	0.00697	0.00729	-0.00437	0.0587	0.0275	0.00477	0.0380	0.0267	0.0149
	(0.108)	(0.0273)	(0.0161)	(0.105)	(0.0305)	(0.0204)	(0.0754)	(0.0354)	(0.0309)	(0.0523)	(0.0345)	(0.0271)
95% CI	[-0.149, 0.320]	[-0.004, 0.115]	[-0.003, 0.067]	[-0.222, 0.236]	[-0.059, 0.074]	[-0.049, 0.040]	[-0.103, 0.220]	[-0.048, 0.103]	[-0.061, 0.071]	[-0.074, 0.150]	[-0.047, 0.101]	[-0.043, 0.073]
Bootstrap 95% CI	[-0.114, 0.372]	[-0.018, 0.114]	[-0.008, 0.063]	[-0.212, 0.261]	[-0.070, 0.089]	[-0.058, 0.046]	[-0.074, 0.231]	[-0.060, 0.102]	[-0.067, 0.072]	[-0.062, 0.148]	[-0.068, 0.095]	[-0.062, 0.070]
Year-month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property fixed effects	Yes			Yes			Yes			Yes		
Area fixed effects		Yes	Yes		Yes	Yes		Yes	Yes		Yes	Yes
Controls 1		Yes			Yes			Yes			Yes	
Controls 2			Yes			Yes			Yes			Yes
N	7969	12086	9358	4563	7030	5375	12397	18463	14430	18062	26509	21033

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

Notes: This table reports the estimated effects of mosque openings and church openings on the log of property prices in Panel 1 and Panel 2, respectively. Columns 1, 4, 7, and 10 report the estimated effects from the *repeat-sales model* shown in Equation 4.3 with the log of property prices as the outcome variable. The sales transaction data used in this specification only include the properties that were sold at least twice between 2002 and 2018. Columns 2, 5, 8, and 11 report the estimated effects from the *modified hedonic model* shown in Equation 4.1. In this specification, we use all sales transactions between 2002 and 2018 and include the controls for acreage, front, and depth of the property. Columns 3, 6, 9, and 12 also report the estimated effects from the *modified hedonic model* shown in Equation 4.1. However, in this specification, the controls include the property's floor area and year built in addition to acreage, front, and depth. The sample radius is two times the size of the treated radius. Standard errors in parentheses are clustered at the level of the neighborhood surrounding the newly opened mosque/church. 95% CI reports clustered-robust confidence intervals. Bootstrap 95% CI reports the confidence intervals from wild cluster bootstrap.

Table C.4: Effects of mosque openings on the likelihood of buyers being Muslim

	treated radius=0.20 mile			0.15 mile			0.25 mile			0.30 mile		
	(1) repeat sales	(2) all sales	(3) all sales	(4) repeat sales	(5) all sales	(6) all sales	(7) repeat sales	(8) all sales	(9) all sales	(10) repeat sales	(11) all sales	(12) all sales
treated x post	-0.0126	-0.00112	0.00725	-0.0234	-0.0109	0.00122	0.00364	-0.00183	0.000398	-0.0130	-0.00157	-0.00259
	(0.0332)	(0.00903)	(0.00655)	(0.0314)	(0.0123)	(0.00757)	(0.0230)	(0.00773)	(0.00855)	(0.0244)	(0.00659)	(0.00819)
95% CI	[-0.081, 0.056]	[-0.020, 0.017]	[-0.006, 0.021]	[-0.089, 0.042]	[-0.036, 0.015]	[-0.014, 0.017]	[-0.043, 0.050]	[-0.018, 0.014]	[-0.017, 0.018]	[-0.062, 0.036]	[-0.015, 0.012]	[-0.019, 0.014]
Bootstrab 95% CI	[-0.081, 0.042]	[-0.021, 0.018]	[-0.011, 0.020]	[-0.078, 0.028]	[-0.039, 0.018]	[-0.012, 0.018]	[-0.038, 0.042]	[-0.019, 0.015]	[-0.017, 0.024]	[-0.059, 0.029]	[-0.017, 0.012]	[-0.023, 0.016]
mean	0.267	0.265	0.265	0.276	0.275	0.275	0.260	0.258	0.258	0.253	0.251	0.251
Year-month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property fixed effects	Yes			Yes			Yes			Yes		
Area fixed effects		Yes	Yes		Yes	Yes		Yes	Yes		Yes	Yes
Controls 1		Yes			Yes			Yes			Yes	
Controls 2			Yes			Yes			Yes			Yes
N	8179	11272	7333	6090	8409	5479	10106	13972	9152	11404	15915	10597

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

Notes: This table reports the estimated effects of mosque openings on the likelihood of the property buyer being Muslim. Columns 1, 4, 7, and 10 report the estimated effects from the *repeat-sales model* shown in Equation 4.2 with the likelihood of the property buyer being Muslim as the outcome variable. The sales transaction data used in this specification only include the properties that were sold at least twice between 2002 and 2018. Columns 2, 5, 8, and 11 report the estimated effects from the *modified hedonic model* shown in Equation 4.1. In this specification, we use all sales transactions between 2002 and 2018 and include the controls for acreage, front, and depth of the property. Columns 3, 6, 9, and 12 also report the estimated effects from the *modified hedonic model* shown in Equation 4.1. However, in this specification, the controls include the property's floor area and year built in addition to acreage, front, and depth. The sample radius is two times the size of the treated radius. Standard errors in parentheses are clustered at the level of the neighborhood surrounding the newly opened mosque. 95% CI reports clustered-robust confidence intervals. Bootstrap 95% CI reports the confidence intervals from wild cluster bootstrap.

Table C.5: Effects of mosque and church openings on the log of property prices (residential properties)

Panel A: Effects of mosque openings

	treated radius=0.20 mile			0.15 mile			0.25 mile			0.30 mile		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	repeat sales	all sales	all sales	repeat sales	all sales	all sales	repeat sales	all sales	all sales	repeat sales	all sales	all sales
treated x post	0.0642	0.0903	0.0491	0.0604	0.155	0.0531	0.0412	0.0105	0.0152	0.0420	-0.00365	-0.00362
	(0.0435)	(0.0539)	(0.0359)	(0.107)	(0.0968)	(0.0716)	(0.0670)	(0.0362)	(0.0340)	(0.0760)	(0.0328)	(0.0353)
95% CI	[-0.025, 0.153]	[-0.020, 0.201]	[-0.025, 0.123]	[-0.162, 0.283]	[-0.046, 0.355]	[-0.095, 0.202]	[-0.095, 0.178]	[-0.063, 0.084]	[-0.054, 0.084]	[-0.112, 0.195]	[-0.070, 0.062]	[-0.075, 0.068]
Bootstrap 95% CI	[-0.012, 0.148]	[-0.029, 0.224]	[-0.040, 0.133]	[-0.162, 0.254]	[-0.039, 0.393]	[-0.085, 0.239]	[-0.073, 0.166]	[-0.074, 0.098]	[-0.070, 0.104]	[-0.092, 0.182]	[-0.079, 0.077]	[-0.092, 0.086]
Year-month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property fixed effects	Yes			Yes			Yes			Yes		
Area fixed effects		Yes	Yes		Yes	Yes		Yes	Yes		Yes	Yes
Controls 1		Yes			Yes			Yes			Yes	
Controls 2			Yes			Yes			Yes			Yes
N	8379	11382	7202	6257	8507	5406	10284	14012	8948	11537	15826	10302

Panel B: Effects of church openings

	treated radius=0.20 mile			0.15 mile			0.25 mile			0.30 mile		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	repeat sales	all sales	all sales	repeat sales	all sales	all sales	repeat sales	all sales	all sales	repeat sales	all sales	all sales
treated x post	0.0875	0.0358	0.0280	0.0179	0.00173	0.00247	0.0468	0.0203	0.000642	0.0391	0.0142	0.00411
	(0.123)	(0.0280)	(0.0165)	(0.115)	(0.0382)	(0.0350)	(0.0847)	(0.0302)	(0.0222)	(0.0515)	(0.0270)	(0.0217)
95% CI	[-0.180, 0.355]	[-0.025, 0.097]	[-0.008, 0.064]	[-0.233, 0.268]	[-0.081, 0.085]	[-0.074, 0.079]	[-0.135, 0.228]	[-0.045, 0.085]	[-0.047, 0.048]	[-0.071, 0.150]	[-0.044, 0.072]	[-0.042, 0.051]
Bootstrap 95% CI	[-0.159, 0.408]	[-0.060, 0.097]	[-0.040, 0.072]	[-0.211, 0.310]	[-0.104, 0.109]	[-0.090, 0.120]	[-0.103, 0.241]	[-0.074, 0.086]	[-0.063, 0.045]	[-0.065, 0.143]	[-0.074, 0.065]	[-0.062, 0.043]
Year-month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property fixed effects	Yes			Yes			Yes			Yes		
Area fixed effects		Yes	Yes		Yes	Yes		Yes	Yes		Yes	Yes
Controls 1		Yes			Yes			Yes			Yes	
Controls 2			Yes			Yes			Yes			Yes
N	7431	10807	8585	4238	6210	4897	11582	16659	13339	16909	24060	19514

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

Notes: Panel 1 of this table reports the estimated effects of mosque openings on the log of property prices of residential properties. Panel 2 of this table reports the estimated effects of church openings on the log of property prices of residential properties. Columns 1, 4, 7, and 10 report the estimated effects from the *repeat-sales model* shown in Equation 4.2 with the log of property prices as the outcome variable. The sales transaction data used in this specification only include the properties that were sold at least twice between 2002 and 2018. Columns 2, 5, 8, and 11 report the estimated effects from the *modified hedonic model* shown in Equation 4.1. In this specification, we use all sales transactions between 2002 and 2018 and include the controls for acreage, front, and depth of the property. Columns 3, 6, 9, and 12 also report the estimated effects from the *modified hedonic model* shown in Equation 4.1. However, in this specification, the controls include the property's floor area and year built in addition to acreage, front, and depth. The sample radius is two times the size of the treated radius. Standard errors in parentheses are clustered at the level of the neighborhood surrounding the newly opened mosque/church. 95% CI reports clustered-robust confidence intervals. Bootstrap 95% CI reports the confidence intervals from wild cluster bootstrap.

Table C.6: Turnover rates of residential and single-family properties in Detroit

	(1) treated radius=0.20 miles	(2) 0.15 miles	(3) 0.25 miles	(4) 0.30 miles
treated x post	0.000542	0.000588	0.000259	-0.000270
	(0.000575)	(0.000513)	(0.000213)	(0.000218)
95% CI	[-0.001, 0.002]	[-0.001, 0.002]	[-0.000, 0.001]	[-0.001, 0.000]
Bootstrap 95% CI	[-.0018, .0023]	[-.0011, .0026]	[-.0011, .0012]	[-.0010, .0009]
Year-month fixed effects	Yes	Yes	Yes	Yes
Property fixed effects	Yes	Yes	Yes	Yes
N	806208	564876	1073448	1357008

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

Notes: In this table, we examine the possibility that mosque openings increased the probability that properties were sold/turned over during the affected time period. Because our data for Hamtramck only includes properties that sold during the time period, we are restricted to Detroit for this turnover analysis. Each column presents estimates from a separate regression where we vary the width of the treated radius. Standard errors in parentheses are clustered at the level of the neighborhood surrounding the newly opened mosque/church. 95% CI reports clustered-robust confidence intervals. Bootstrap 95% CI reports the confidence intervals from wild cluster bootstrap.