The Returns to Public Library Investment

Gregory Gilpin   Ezra Karger   Peter Nencka *

May, 2023

Abstract

Local governments spend over 12 billion dollars annually funding the operation of 15,427 public libraries in the United States, yet we know little about their effects. We use data describing the near-universe of public libraries to show that public library capital investment increases library visits, children's attendance at library events, and children's circulation by an average of 5–15% in the years following investment. Increases in library use translate into improved test scores in nearby school districts: a $200 or greater per student capital investment in local public libraries increases reading test scores by 0.01-0.04 standard deviations in subsequent years.

*Gregory Gilpin, gregory.gilpin@montana.edu, Montana State University. Ezra Karger, ezra.karger@chi.frb.org, Federal Reserve Bank of Chicago. Peter Nencka, nenckap@miamioh.edu, Miami University. Thanks to Riley Acton, Matthew Birnbaum, Lauren Jones, Kurt Lavetti, Dean Lillard, Bruce Weinberg, and 3 referees and the editor for helpful comments on prior versions of this work. We thank Zöe Arnaut for valuable research assistance. We also thank seminar participants at the Association for Public Policy and Management, Association for Education Finance and Policy, Case Western Reserve University, Institute for Museum and Library Science, NBER Summer Institute, Miami University, and the Southern Economic Association for their feedback. The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

Research shows that neighborhoods play a critical role in child development. Where a child grows up affects their academic performance, long-term educational attainment, and labor market outcomes (e.g., Chetty and Hendren, 2018a; Chetty and Hendren, 2018b; Chyn, 2018). However, despite a recognition that location matters for child development, there is much to learn about which local policies and institutions create beneficial environments. In particular, while an extensive literature studies whether and how local public school spending translates into student achievement, less is known about the effects of other policy levers available to local governments.

In this paper, we analyze a ubiquitous yet understudied local institution: the public library. In 2018, there were 9,261 library systems across the United States, with 15,427 branches. Libraries spent 12 billion dollars on operating costs, and patrons checked out over 2 billion items.[1] Beyond their collections, libraries also provide instructional programs on topics ranging from literacy to computer usage, job search, and tax preparation; and they serve as one of the few non-commercial indoor spaces available to the public.[2] A particularly important focus of library services is child development. In 2018, children checked out more than 750 million library items and attended library events more than 80 million times nationwide.

Despite the popularity and abundance of libraries, it is difficult to study their effects due to two major challenges. First, the vast majority of public library operations funding is determined locally; in 2018, city and county governments provided 86 percent of library funding, with only 6.8 and 0.3 percent coming from state and Federal sources, respectively.[3] This local operating spending is relatively stable from year-to-year, limiting researchers' ability to exploit policy-generated changes in annual expenditures. In addition, there is limited research causally linking changes in library expenditures to changes in library resources and patron usage. Without estimating the public library production function, it is difficult for researchers and policymakers to understand how library investment generates community benefits.

[1] All statistics are taken from the 2018 version of the Institute of Museum and Library Services' (IMLS) Public Library Survey. We discuss the construction of these data in Section II.

[2] Building on Andrew Carnegie's observation that libraries are "palaces for the people," Klinenberg (2018) argues that libraries improve community health outcomes by providing vulnerable residents with a refuge of last resort.

[3] The remainder of funding is classified as coming from "other" sources in the IMLS data, and includes private foundations and philanthropists.

We use new data and methods to address these challenges. We estimate how investment in public libraries affects library operations, patron usage, and local communities, focusing on student achievement. We compile library spending, revenue, and usage data collected annually in the Institute of Museum and Library Services' Public Library Survey. We link this data to district-level test scores from the Stanford Education Data Archive (SEDA) and zip code-level house price indices from Zillow and the Federal Housing Finance Agency (FHFA). Using this data, we examine a specific type of library expenditure: capital spending on major renovations and new library buildings. Unlike day-to-day expenses, capital investments are lumpy, allowing us to identify discrete changes in public library investment. We show that the timing of these large-scale investments is orthogonal to changes in other local government spending and local demographics. Using these sharp changes in investment, we estimate the causal effects of library capital spending with dynamic difference-in-difference models. We use a recently developed two-stage method presented in Gardner (2021) and Borusyak, Jaravel, and Spiess (2021) that avoids biases common in staggered difference-in-difference designs.

We first estimate the effect of capital spending on library resources and usage. We find that library capital spending events larger than $200 per local student sharply increase library visits, children's circulation, and children's attendance at library events by an average of 5-15 percent in the years after library investment. Capital investment also increases library employees, spending on salaries, and annual operating expenditures, and these increases exist on both an absolute and per user basis. In other words, library capital investment increases both the quality and usage of libraries.

Next, we test whether library capital investment affects children's academic achievement. We find that library capital spending increases nearby children's reading test scores by 0.01-0.04 standard deviations in the seven years following the library investment shock, with the largest effects coming five to seven years after the increased capital investment. We find mixed evidence that library investment affects math test scores, consistent with the reading-focused mission of libraries. For both test scores and library-level outcomes, we see sharp changes in outcomes in the years after library investment but no differential pre-trends in the years leading up to those capital spending events.

We test for heterogeneous effects of libraries across community and student characteristics. Effects are largest in smaller districts and in districts that spent the most per student on schools during our sample period. Together, these results suggest that the effects of libraries are larger when libraries are more salient in the local community and that libraries may play a complementary role with school spending. Across student characteristics, we see no evidence of differential effects by grade or socioeconomic status. We find some evidence that our effect sizes are smaller for Black and Asian students and larger for white and Hispanic students, but we cannot reject sizable effects for all subgroups.

Our results are robust to measurement and specification choices, paying particular attention to the possibility that results could be driven by selective migration of residents to areas with new and improved public libraries. While we cannot observe individual student data, we explore whether our results are explained by other time-varying local characteristics that correlate with library expansions. We estimate our baseline model with these characteristics as the outcome variables and show that library capital spending does not predict changes in student demographics, local adult characteristics, or local school spending, addressing concerns about selective migration. In addition, we show that our results are robust to controlling flexibly for these local characteristics and state-year fixed effects using the Gardner (2021) method, standard two-way fixed effects estimators, and the Callaway and Sant'Anna (2021) method. Lastly, we show that our results satisfy two placebo tests: (1) small library capital investment shocks have no effect on test scores; and (2) large library capital investment shocks near a school district, but outside of commuting distance (30–50 miles from the district), have no impact on student achievement.

Prior work

Our results build on a small literature that tries to estimate the demand factors that drive public library usage. For example, James (1985) studies how changes in local economic conditions impact public library use. Palmer (1981) and Ottensmann (1997) estimate the impact of patrons' distance to public libraries on the demand for services. Bekkerman and Gilpin (2013) estimate the impact of access to high-speed residential Internet access on the demand for library services. We focus on how library demand changes after large capital investments.

Several related papers focus on the community effects of libraries, particularly their potential effects on children. Goldhor and McCrossan (1966) and Guryan, Kim, and Quinn (2014) study the impact of summer reading clubs on reading skills. Bhatt (2010) instruments for library usage with the distance to the nearest public library and finds that libraries increase the amount of time children spend reading, reduce children's television consumption, and increase homework completion rates. Rodríguez-Lesmes, Trujillo, and Valderrama (2014) use a difference-in-difference approach to show that the construction of two public libraries in Bogota, Colombia, did not affect nearby high school students' test scores. Porter (2015) finds that when libraries stayed open for longer hours in Los Angeles, crime rates declined. Similarly, Neto, Nowicki, and Shakya (2021) find that crime fell after a Kansas City, MO library branch opened. Neto (2022) finds no effect of library programming on local labor market conditions using a lagged funding instrument.

Another set of papers examines the historical impact of public library investment in the early 1900s. Kevane and Sundstrom (2016) find that public library investment did not affect political participation; Berkes and Nencka (2023) show that Carnegie libraries increased patenting rates, and Karger (2021) finds that Carnegie libraries increased the long-run educational attainment and earnings of exposed children. These papers use identification strategies that exploit the lack of commonplace, high-quality libraries in the early 1900s. To our knowledge, this is the first paper to show that library spending directly affects student achievement in the modern period.

Our results also relate to a literature that discusses and structurally models the public financing of libraries. Early work recognized that libraries may be a public good, and their characteristics could justify the government provision of library services (Tiebout and Willis, 1965). In public finance, a literature analyzes the efficiency of public library spending relative to theoretical benchmarks and other public institutions (e.g. Feldstein, 1977; Getz, 1980; DeBoer, 1992; Vitaliano, 1997; Vitaliano, 1998; Worthington, 1999; Hemmeter, 2006; Neto and Hall, 2019). Xie and Waldfogel (2022) use the Public Library Survey to calibrate a structural model of the tradeoff between physical and electronic holdings in public libraries. We show that discrete investments in public library infrastructure have long-lasting effects on library visits, programming, and circulation patterns.

To benchmark our test score findings, we use a back-of-the-envelope calculation to compare the magnitude of our results to the effects of public school capital spending. Our work is methodologically similar to an extensive literature on the causal effects of school spending. Most relevantly, a subset of this literature focuses on the effects of large, discrete capital investments in schools.[4] In a recent meta-analysis, Jackson and Mackevicius (2023) argue that the balance of credible work suggests that both capital and operational school spending boosts test scores. Using Jackson and Mackevicius's estimates, we find that the typical large-scale library capital investment increases test scores by a larger amount per dollar relative to spending on schools. This suggests that library investment can be an important complement to public school investment. However, we caution that precise dollar-to-dollar cost/benefit comparisons are complicated by two facts: (1) the typical library capital investment is smaller than the typical school capital investment, making dollar-to-dollar comparisons difficult, and (2) test scores are only one facet of the social return to public library and school investment. Each may have broader and distinct effects on local communities.

I Library policy background

Public libraries have always been community centers, but their mission and programming have evolved over time. The modern idea of a public library in the United States emerged after the 1893 world fair in Chicago, where Melvil Dewey organized a library exhibit that inspired new libraries across the country (Wiegand and Wiegand, 2018). Children's sections in libraries became the norm starting in the early 1900s, and a long-standing mission of libraries is increasing child literacy. Following the development of the internet, libraries pivoted to providing more digital resources, including e-books. At the same time, libraries increased in-person programming for children. These events include story-time and early literacy services for newborns, reading tutoring for older children, and topic-area programming for teens.

[4] Methodologically, this school spending literature and our paper are similar to other papers that focus on continuous treatment variables that are analyzed as discrete treatment events (e.g., Neilson, 2020; Thompson, 2022; Baron, 2022).

This renewed focus on in-library programming is evident in data on library use. From 1992–2018, there was a 130 percent per capita increase in the number of children attending events at libraries. This growth has accelerated in recent years. Children's circulation increased 42 percent over 1992-2018.[5] Library usage remains widespread, with 53 percent of families with children in 2015 reporting that they used a library in the past 12 months (Pew, 2015).

Public library services are funded by a mix of local, state, and national governments, as well as private donations. Per capital library funding has remained relatively stable over the last 30 years, but the sources of that funding have changed over time. Most library revenue comes from local taxes, and the relative importance of local funding has grown over time. According to the IMLS Public Library Survey, in 1992, 79 percent of library income came from local sources; by 2018, this proportion had increased to 86 percent (IMLS, 2019). Increased local support made up for declines in Federal (from 1 percent to 0.3 percent) and state (from 12 percent to 7 percent) funding over this time period. Local funding models differ from jurisdiction to jurisdiction, but the most common approach involves property taxes that fund ongoing operations. In this paper, we focus on large, one-off capital expenditures. While these can be partially funded from operational savings or Federal/state grants, most are funded through one-time, city-issued bonds.

Why do libraries improve their capital stock, and how might we expect it to affect community and child outcomes? This question is central to this paper, and we aim to estimate these effects. Discussions in the library policy community about the value of new and improved buildings indicate beliefs about potential benefits. Many new buildings involve creating space that can house both collections and events, consistent with the evolving missions of libraries. For example, in a recent proposal, Westborough, Massachusetts, proposed to renovate their library to "continue evolving as the clear 'hub' of the community", adding spaces for meetings and an interactive childrens' section.[6] Westborough argued that their current building "prohibits [them] from fulfilling the needs of [their] current and future users" and that constructing a new library would attract more patrons and users. In this paper, we directly test for these possible impacts.

[5]Calculated from 1992-2018 IMLS data.

[6]For more details, see the proejct summary here: https://www.westboroughlibrary.org/about/building-project/

II Data

II.A Library data

Our data on libraries come from the Public Library Survey (PLS), a census of public library systems collected annually by the Institute for Museum and Library Services (IMLS) since 1988 (IMLS, 2020). The survey unit in the PLS is a library system, which can contain multiple library branches. Survey response rates are extremely high. In 2017, 9,042 of 9,216 eligible public libraries responded to the survey, a 98 percent response rate corresponding to over 17,000 library branches (IMLS, 2019). The PLS contains a rich set of information about library finances and usage. IMLS (2019) provides a detailed overview and summary of the variables that they collect. Most information about spending and usage is reported at the library system-level and is not available for individual branches.

We are particularly interested in capital expenditures, so we use the PLS's measure of capital spending as our primary treatment variable. IMLS defines capital expenditures as spending on major one-time projects that add to fixed assets, including purchasing land, building new structures, and renovating existing structures. We also analyze operating expenditures, which the PLS separates into several components, including spending on staff salaries, staff benefits, and collections. We convert all dollar values to 2019 dollars using the all-items CPI accessed from the Federal Reserve Economic Data system (Organization for Economic Co-operation and Development, 2022).

In addition to financial library statistics, the PLS contains comprehensive information on library usage and holdings. As measures of library usage and resources, we analyze the number of children who attend library-based events, the total circulation of a library's collections, the circulation of children's material, the number of visits, and the total stock of books and print serial subscriptions (e.g., magazines and newspapers).[7] We also examine library spending on operations, salaries, and book collections.

[7] There is no individual-level library usage information in these data. This means that we do not observe the number of unique users of library services, for example. Instead, we observe the number of user-by-visit events.

II.B Test score data

We measure the effect of public library capital investment on test scores using a standardized dataset of district-level test scores from version 4 of the Stanford Educational Data Archive (SEDA), compiled by The Educational Opportunity Project at Stanford University (Reardon et al., 2020; Fahle et al., 2021). The basis for these data is annual standardized tests that public school students were required to take following the passage of the Federal No Child Left Behind Act in 2002. In our main analyses, we use reading and math test scores from 2009 to 2018 for 3rd–8th graders in over 5,000 school districts across the United States. The SEDA panel reports test scores in standard deviation units on a national scale. The team that compiled the SEDA data constructed the mean and standard deviation of the test score measure for each grade, subject, and year using a heterogeneous ordered probit model and restricted-access data from the U.S. Department of Education (Fahle et al., 2021). We describe additional details of the construction of this dataset in Appendix A.A.

In robustness exercises, we show results that augment SEDA with two additional test score databases, allowing us to extend our panel before 2009 for some states and school districts. First, we use the National Longitudinal School-Level State Assessment Score Database (NLSLSASD), a predecessor to the SEDA panel from a different group of researchers that also collected test score information from states. These data include standardized test score information before 2009, though their coverage is unbalanced across states and grades (McLaughlin et al. 2002). As described in Appendix A.A, we transform binned proficiency percentiles from the NLSLSASD into district-level means and standard deviations that are comparable to the SEDA panel using the same methodology as SEDA researchers (Fahle et al., 2021). Second, we use test score data from restricted versions of the National Assessment of Educational Progress (NAEP). These data are available for a subset of school districts starting in 2000 for 4th and 8th graders (U.S Department of Education, 2022). We follow Brunner et al. (2022) and construct district-year-grade-subject aggregate scores from the NAEP data that are comparable to the test scores in SEDA.

II.C Housing price data

We use house price indices from two sources: Zillow and the Federal Housing Finance Agency (FHFA). Both datasets are available at the five-digit zip code-level. Zillow constructs its house price index by forecasting the sales price of all houses in its national database of more than 100 million properties (Zillow Group, 2020). Zillow calls these forecasts 'Zestimates.' Zillow then calculates the index for a zip code as the value-weighted average Zestimate in the area, excluding houses that undergo significant construction or renovation.[8] We also show housing price effects using the FHFA index of house prices at the five-digit zip code-level (Federal Housing Finance Agency, 2019). The FHFA constructs this index using a repeat sales measure (Bogin, Doerner, and Larson, 2019). While the Zillow and FHFA measures are constructed to measure two different quantities: changes in the value of all homes (Zillow) vs. changes in the value of newly sold homes (FHFA), we show in Section VI that the two indices are interchangeable in our context, and lead to virtually identical results.

II.D Local covariates and characteristics

To understand other factors that might have been changing at the same time as library expansions, we use data on time-varying local characteristics available at a school district-level. We use SEDA-constructed variables for the proportion of adults in a school district with a bachelor's degree or higher, the share that is SNAP eligible, the share unemployed, and the share of single-mother households. These measures derive from the American Community Survey, and their construction is discussed in Fahle et al. (2021) and Appendix A.A.

In addition, we incorporate time-varying school district covariates. We use demographic information on the share of Black, Hispanic, Asian, and Native American students, as well as the share of students who qualify for free or reduced-price lunch. We also compile information on total school district spending, capital spending, salary spending, and spending on salaries related to the instruction of students. We access this data through the Urban Institute's Education Data

[8]For more details about Zillow's methodology, see https://www.zillow.com/research/zhvi-methodology-2019-deep-26226/.

Portal API, which compiles information from the Common Core of Data (Common Core of Data, 2022).

II.E Sample construction and summary statistics

We construct a consistent sample at the school district-level for all the analyses we present in this paper. Because the survey unit of the library data is a library system and the unit of analysis in our test score data is a grade within a school district, we match the two datasets together using a flexible distance approach. Our matching of libraries to school districts balances the precision with which we can associate a given school district to a nearby library with our aim of including as many library systems in the analysis as possible.

To achieve this goal, we locate all library systems within two miles of the modal zip code in a school district. To the extent that multiple library systems are within a given 2-mile radius, we aggregate all library-level variables to the school district-level. We choose the 2-mile radius as our baseline to prioritize a tight match between library systems and schools. Given the small radius, it is likely that every library system in our sample matched to a school district is geographically relevant to residents of that school district. We show results where we flexibly relax this matching threshold to 1, 5, 10, 15, and 20 miles. In addition, we conduct a placebo analysis where we match school districts to libraries within 30-50 miles. This is a placebo since it is unlikely that improvements to these libraries would affect test scores, given their distance.

To increase the precision of our match between libraries and school districts, we limit our main analysis to districts that have at most five library buildings (including branches) within two miles of the modal zip code of the district. This avoids dense urban areas where the mapping between a library improvement project and its effects on local residences is likely to be attenuated. This sample restriction affects 2.8 percent of the sample of district-grade-year observations and 4.8 percent of the sample by district enrollment. We show results where we relax this sample restriction in Section V. For housing prices, we follow a similar procedure. Our housing price data is at the zip code-level, so we locate all zip codes within two miles of the modal school district zip code. We then take the average housing price values for those zip codes. After these steps, our analysis sample is unique at a grade-school district-year-level. It contains information on nearby library systems and housing markets within two miles of the school district.

In Table 1, we show statistics describing the library systems within two miles of each school district and student enrollment at the schools in that district. For example, the average district in the sample has 13 percent Hispanic students, and the nearby adult unemployment rate is 7 percent. The number of students enrolled in each district is much smaller than the number of people served by the libraries within two miles of that school district for three reasons: (1) elementary and middle school students make up only a small fraction of the population that could use a given library; (2) library systems often serve a larger area than is covered by one school district; and (3) our geographic rule for assigning libraries to school districts is broad: a school district may have one library system two miles to its east and one library system two miles to its west. In these cases, those two library systems likely serve many people who are not represented by the geographic bounds of the school district.

III Methodology and identification

III.A Event study framework

A standard approach for estimating dynamic treatment effects is to use an event study regression with unit and time fixed effects. In our context, the estimating equation to evaluate the effects of library capital investment shocks on place-level outcomes is:

$$Y_{gdst} = \sum_{e=-7}^{7} \beta_e 1(CapitalShock)_{dte} + \delta_{g,d} + \gamma_{s,t} + \epsilon_{dgst} \tag{1}$$

where Y_{gdst} is an outcome for grade g, in school district d, state s, and year t. The indicator variable $1(CapitalShock)_{dte}$ tracks the year surrounding a capital investment shock to the libraries near district d, and $\delta_{g,d}$ and $\gamma_{s,t}$ are grade-district ("unit") and state-year fixed effects respectively. The coefficients of interest are the vector β_e. Typically, researchers interpret β_e as the causal effect of the treatment on the outcome of interest e periods from the shock. However, this interpretation can be incorrect if there are heterogeneous treatment effects across cohorts or time. In particular, since standard event studies use early-treated units as controls for later-treated units, treatment effect dynamics can introduce bias.[9]

For these reasons, we use a two-stage event study procedure developed independently by Gardner (2021) and Borusyak, Jaravel, and Spiess (2021). In the first stage, we estimate our fixed effects (and later, other covariates) using a sample of never-treated and not-yet-treated school districts:

$$Y_{gdst} = \delta_{g,d} + \gamma_{s,t} + \nu_{dgst}. \tag{2}$$

In the second stage, we estimate:

$$Y_{gdst} - \hat{Y}_{gdst} = \sum_{e=-7}^{7} \beta_e 1(CapitalShock)_{dte} + u_{gdst}. \tag{3}$$

where $Y_{gdst} - \hat{Y}_{gdst}$ are the residual outcomes calculated for *all* units in our sample generated from the coefficients in Equation 2. In other words, in the first stage we use the untreated and not-yet-treated observations in our sample to calculate the relationship between the covariates and our outcome of interest. In the second stage, Equation 3, we use the residual from the first stage and calculate the sample average of the outcome in event time, e. As Gardner (2021) shows, this two-step process avoids contaminating event study coefficients with estimates from treated units in the sample population, avoiding standard issues with two-way fixed effects estimation.

Throughout, we calculate clustered standard errors at the school district-level. Given the two-stage procedure, standard analytical clustered standard errors incorrectly ignore the uncertainty of coefficients in the first stage. To address this, we estimate bootstrapped standard errors.[10] Our preferred bootstrap procedure is a version of the Bayesian clustered bootstrap (Rubin, 1981). Instead of drawing clusters with replacement, as in a standard bootstrap, we draw independent

[9] For a discussion of this bias and evidence of how this bias can change the magnitude and statistical significance of estimated causal parameters in published research, see Baker, Larcker, and Wang (2022).

[10] Gardner (2021) shows how to produce corrected analytical standard errors, though they are computationally infeasible to calculate given the size of our data.

exponential weights for each cluster (school district) in our panel in each iteration.[11] We use those weights in weighted regressions to produce the two-stage regression results for each bootstrap iteration (Cheng, Yu, and Huang, 2013). This produces an empirical distribution of coefficients across bootstrap iterations that we use to generate confidence intervals and standard errors for each event study coefficient.

We rely on the standard difference-in-difference parallel trend assumption for identification: treatment and control units would have followed parallel outcome trends after the capital investment if not for the existence of the treatment. In our setting, this implies that outcomes in school districts that did and did not invest in library systems would have followed parallel outcome paths in the absence of any treatment. This assumption is untestable, but we gauge its plausibility by seeing: (1) whether outcome trends in the years leading up to the treatment year are parallel across treatment and control units, (2) whether other observable characteristics between treatment and control units were parallel before and after the treatment, and (3) whether our results are sensitive to the inclusion of a battery of time-varying district and area covariates as controls. The second and third analyses help evaluate the concern that other community characteristics that could affect outcomes may have changed at the same time as library investments. We find that our results are robust to all three tests, bolstering the plausibility of our estimates.

We also explore additional estimation methodologies. In Section V, we show that we estimate similar treatment effects using alternative methodologies, including two-way fixed effects models that control with and without time-varying covariates. We also estimate models using Callaway and Sant'Anna (2021) robust difference-in-difference approach, which allows for matching based on pre-treatment covariates. We estimate similar effects using the Callaway and Sant'Anna approach, with or without this covariate matching.

[11] This procedure is computationally faster than standard wild cluster bootstrapping, and allows us to maintain the same set of estimated fixed effects across all specifications. We use 1,000 iterations for all results that we show in the paper.

III.B Defining capital spending events

Our baseline method described in the previous subsection uses a discrete treatment to generate clearly defined treatment events. This strategy helps us avoid interpretation issues that emerge when one uses continuous treatment variables in a difference-in-difference framework (Callaway, Goodman-Bacon, and Sant'Anna, 2021).

In our baseline model, we define a capital spending shock as the first year in the sample with at least $200 per student of capital library spending. We calculate the number of students using SEDA test score data at the school district-level. Because 3rd–8th graders make up about 10% of the U.S. population, a $200 per student shock can be interpreted as a $2,000 or larger per person increase in capital spending. As described above, we calculate capital spending after aggregating library data to the district-level by locating all library systems within two miles of a school district's modal school zip code.

We focus on library systems with capital investment shocks between 2010 and 2017, during which we can measure test score data. We also require that in year $t-1$, the library system has less than $200 per person of capital spending to ensure that our capital spending shocks represent sudden changes from prior levels. Using this shock definition, 12% of all school districts in our sample experienced a shock between 2010–2017. While our preferred methodology requires a discrete treatment, we also present results with alternative shock sizes in Section V. These results help trace out the intensive margin of the response to library spending.

If a district experiences multiple capital spending shocks, we use the earliest one to define treatment timing and ignore subsequent shocks. The potential for multiple shocks affects the interpretation of our results, since the probability of observing a shock in a given year depends on past shock observations. For example, a school district that experiences a large library construction project in year t is less likely than control units to see a new project started in year $t+1$. Our results, therefore, come from employing what Cellini, Ferreira, and Rothstein (2010) call a "reduced-form" method. This produces a policy-relevant parameter, since it reflects the expected outcome for a real-world investment in library resources. We account for this treatment dependence in our cost-benefit analysis by tracing out the impacts of our observed "events" on the actual path of capital investment (and other library spending) in subsequent years. We then compare this outlay to our outcomes of interest.[12]

IV Results

IV.A Effects of capital investment on library use and quality

Figure 1 shows the results of our baseline analysis for our main library outcomes. All figures show the coefficients estimated from Equation 3 with bootstrapped 95 percent confidence intervals. We present the logged outcomes as $log(x + 1)$ to account for zeroes.[13]

Panel A of Figure 1 shows how our discrete capital spending shock treatment correlates with logged library capital spending in the time periods around a treatment event. In the year of the first large capital expenditure for a library system, there is a sharp increase in capital spending that decreases to baseline levels after 3-4 years. Panels B and C show the effects of capital spending shocks on circulation of children's material and children's attendance at library events. We observe a sharp and persistent 5–15% increase in child library use that lasts for seven years after the capital investment. Interestingly, we also observe a suggestive drop in children's event attendance in the year of the shock. This is consistent with a brief decline in library capacity in the year when a library is under construction. Panel D shows the effect of capital spending on log total visits to the library. There is a small decline in visits during the year of capital investment followed by a sharp and persistent 8-10% increase in visits after investment. Across all three library usage outcomes, there are no pre-trends in the years leading up to the library capital spending event.

To summarize, Figure 1 shows that library capital spending has sharp, positive, and persistent effects on library use. In Figure 2, we show the effect of capital spending shocks on additional

[12] Appendix Figure A26 shows that conditional on having any library construction event, 86 percent of our sample only has one such event. In Section V, we show that our results are similar if we limit our sample to single-event districts.

[13] In subsequent sections, we show that these results are robust to alternative transformations, including the inverse hyperbolic sine and logged per capita values.

library system-level outcomes. We see only delayed changes in the total number of books in the library system after a capital shock, but approximately 2-5% increases in the number of employees, payroll, and total operating expenditures after a capital spending event. These results suggest that the collection size of libraries takes longer to adjust after a large capital investment, but that other measures of library services quickly increase after investment, consistent with the increases in library use illustrated in Figure 1. Figures A2 and A3 show those same figures but with a standard two-way fixed effect regression framework, without applying any correction for contamination in using early-treated units as controls for later-treated units. Here, we see similar sharp changes around capital spending events. However, we also see evidence of pre-trends for some outcomes, consistent with the strong effect of capital spending shocks on outcomes creating contamination bias due to early-treated units serving as controls for later-treated units.

Recent work by Chen and Roth (2023) shows that using log or inverse hyperbolic sign transformations can produce difficult-to-interpret treatment effects when some outcome observations are zero-valued. In Figure A4, we show results for four library-level outcomes along the intensive margin, removing observations from our panel with zero-valued outcomes and instead focusing only on logged outcomes for observations with non-zero outcomes. We see similar effects on library resources and use to our baseline result, demonstrating that our main results on library use are robust to transformation concerns.

As another robustness check to our main results, we show results using logged per user library use (Figure A5) and library resources (Figure A6). To calculate these measures, we rely on each library's reported service population in the PLS, which can vary over time. This analysis identifies intensive margin effects by taking into account each library system's potential expansion of its service area in the years following the capital investment shock. We find similar results in the per capita and baseline specifications. This implies that libraries are not attracting additional users merely by expanding their service area. In Figures A7 and A8, we show very similar results if we use an inverse hyperbolic sine transformation for our library outcome variables instead of a $log(x + 1)$ transformation. Finally, in Figure A9 we show that we also observe increases in two sub-categories of library quality: the number of log subscriptions to print serial subscriptions that libraries hold and their annual log spending on collections.

IV.B Effects of capital investment on student test scores

Figure 3 shows results for reading test scores. As with the library outcome analysis, we define capital shocks as a $200 or greater increase in per student library capital spending within two miles of each school district. Our main outcome is district-grade-year measures of average reading test scores normalized to be mean zero and standard deviation one at the subject-grade-year-level. A one-unit increase in these test scores corresponds to a one standard deviation increase in test scores. We observe no pre-trend in test scores in the years leading up to a capital library shock, suggesting that districts are not positively selected on test score dynamics in the years preceding major library investments. After capital investment shocks, we observe gradual increases in reading test scores. On average, reading scores increase by approximately 0.01 standard deviations in the short-run, with larger effects (0.02–0.04 standard deviations) emerging in the later years of the estimation window. The timing of these results corresponds to the increases in library use and quality following capital expenditures, as observed in Figures 1 and 2.

A potential concern is that these results are biased by time-varying changes in local communities or school districts that (1) correlate with the timing of library investments and (2) affect test scores. We find no evidence of pre-trends in our event studies, implying that test scores and library usage in communities that invest in libraries are not changing differentially in the years leading up to library capital investment. However, if local policies or demographics that affect test scores change at the same time as library spending, our results could be affected.

We address this concern in multiple steps. First, Panel B of Figure 3 shows results conditional on an extensive battery of time-varying covariates, including local demographic shares, the share of free and reduced-price lunch students, local school district revenue by source, and local school outlays on capital spending, construction, salaries, and instructional services. We estimate nearly identical event study coefficients conditional on these covariates, suggesting that our findings are not driven by time-varying differences in districts that had nearby expansions in public libraries.

As an additional check on our identifying assumptions, we vary the distance from school districts to libraries that we used to calculate shock exposure. Our baseline estimate uses all capital spending within two miles of a school district. Libraries are local institutions, and it would be difficult for far-away libraries to affect students' scores. Consistent with this intuition, our results in Figure 4 (Panel A) show that capital spending shocks between 30–50 miles from a given district do not affect students' reading test scores. This is a placebo test that suggests that we are not merely picking up a general trend towards higher test scores in broader regions that choose to invest in libraries. Similarly, Panel B of Figure 4 shows that small capital investment — between 1 and 30 dollars per student — has no effects on reading test scores. Our finding that public library investment causes an increase in reading test scores is confined to local, large public library investments. We confirm this point through additional robustness testing described in Section V.

In summary, we find that library capital investment increases measures of library use, library quality, and nearby reading test scores in the seven years after library investment. We summarize these results together in Table 2, which contains estimated post-period coefficients for 1, 3, 5, and 7 years after a library spending event for our main outcome. This table also contains an overall pre-post coefficient, estimated by replacing our relative time indicators in Equation 3 with a single dummy variable equal to one for treated units after their treatment date and otherwise equal to zero.[14] This summary highlights the consistent pattern of results that we observe across library spending, library use, and test score outcomes.

IV.C Reading results heterogeneity

To help understand our main test score results, we investigate whether library investment has heterogeneous effects across student or community characteristics.

[14] We also estimate this aggregate effect using the Gardner (2021) method. These overall effects are not the simple average of our event study coefficients. Instead, the overall effects are relative to the entire pre-period (not just $t-1$), and the calculation is weighted by sample sizes within relative time bins. For example, more weight is placed on comparisons between $t-2$ and $t+2$ than comparisons between $t-2$ and $t+6$ because we have more observations for $t+2$. For that reason, we focus most of our analysis on the event study coefficients, which let us observe treatment effect dynamics.

First, we explore how reading effects differ by race. In Figure 5, we show results from the baseline reading event study estimated separately for the test scores of (A) white, (B) Black, (C) Hispanic, and (D) Asian students. These results are less precisely estimated than our baseline results; SEDA does not report test scores in districts when there is an insufficient number of observations to construct their test score measures, and many districts have only a small number of students of a given race, leading to more year-to-year variability in the test scores of small groups. The results in Figure 5 suggest that our findings are mainly driven by positive effects on the reading test scores of white and Asian students. However, the standard errors shown in these figures are large enough that we cannot rule out sizable, positive effects for all groups.

Unfortunately, the PLS library data does not measure library use by race, so we cannot compare these results to similar measures of library usage. However, these results suggest that programs aimed at boosting library engagement for specific racial groups may be needed to realize the full benefits of library investment. [15] In particular, residents in the south tend to live further away from libraries than those in the north (Donnelly, 2015), and more effort may be needed to encourage library attendance in those circumstances. We also know that librarians are disproportionately white (American Library Association, 2012), and that a history of discrimination in public libraries (e.g., Wiegand and Wiegand, 2018) may still have long-term effects. Future work on additional dimensions of library access by race will be valuable.

Next, we estimate heterogeneous effects by student economic status. Figure 6 shows results estimated separately by students whom SEDA classifies as "economically disadvantaged" (Panel A) and "not economically disadvantaged" (Panel B), respectively. Higher-income and more educated families are more likely to use public libraries (Pew, 2015). However, the marginal value of library materials and programming may be higher for families with fewer resources. On balance, it is therefore unclear which group of students would benefit more from library investment. Figure 6 shows similar patterns of results in both groups of students. As with race, we do not observe information on economic status directly in our data. Future work on library use and benefits

[15] These results may be complicated by the higher concentration of white students in smaller school districts, where we see larger effects, as discussed below.

by socioeconomic status would be valuable to understand how library investment differentially affects educational outcomes.

We also investigate how effect sizes differ by grade. Students across all age groups may benefit from library services, though the types of books consumed and events attended vary by age. Figure 7 shows the baseline results estimated separately for student test scores reported in grades 3-4 (Panel A), 5-6 (Panel B), and 7-8 (Panel C). We observe similar results in all grades, though positive effects emerge more quickly for the youngest students (grades 3-4) in the years immediately following library investment. Interpreting our results in the context of a dynamic education production function, it is plausible that libraries have the most impact at earlier ages, when it can be easier to affect the trajectory of human capital development with targeted investments (e.g., Heckman, 2007). However, more analysis is needed on this point, since the precision of our heterogeneity estimates limits our ability to rule out alternative hypotheses.

Next, we explore heterogeneous effects across district sizes. Students in smaller school districts are likely to be more affected by a given per capita library shock, since they are more likely to live near only one library system. This implies that they are "fully treated" by any library spending shock. By contrast, students in larger districts are more likely to live near multiple library branches. For these students, when one nearby library improves, it does not necessarily improve the local library they regularly visit. In addition, students in large districts may have more local educational amenities, like parks and museums, nearby. In Figure 8, we show results separately after splitting the sample into school district size terciles. Panel A of Figure 8 shows that our reading test score effects are largest in the smallest school districts, consistent with the increased salience of libraries in those communities and less measurement error in our ability to link relevant library spending to nearby school districts. Consistent with this heterogeneity, in Figure A18 we estimate our baseline specification on the full sample of school districts but weight observations by the number of test-takers in a given district-grade-year cell, and we observe attenuated effects. This is because, as shown in Figure 8, our results are driven by smaller school districts.

Finally, we estimate how our library results differ across levels of investment in *school* capital spending using Common Core data on annual district capital expenditures. Per student school capital spending has only a weakly positive correlation with school district size ($r = 0.132$), so this exercise is distinct from estimating heterogeneity by district size. In Figure 9, we show our baseline reading test score results after splitting the sample into district capital spending terciles. We find that our results are driven by districts that have higher amounts of school capital spending. These results suggest a possible complementary role of library investment in the education production function.

V Robustness of test score results

In this subsection, we demonstrate the robustness of our test score results. First, we show additional evidence that our findings are unlikely to be driven by time-varying, district-level confounders given a lack of observed dynamic selection in the timing of library investment. Second, we show that our findings follow predictable patterns when we vary the shock size and distance thresholds that we use to construct our sample and define our treatment events. Third, we show that our results are robust to alternative specifications and to using additional datasets to extend the length of our panel. Finally, we show that our results are robust to two alternative estimation strategies: the standard two-way fixed effects ordinary least squares estimator and the Callaway and Sant'Anna (2021) staggered treatment timing estimator.

V.A Dynamic selection

A potential concern is that our results could be biased by time-varying changes in local communities or school districts that (1) correlate with the timing of library investments and (2) affect test scores. In Section IV and Figure 3, we showed that our test score results were robust to controlling for a battery of time-varying covariates. To explore this result further, we directly test whether the timing of library capital investment is correlated with changes in local characteristics that could affect student test scores. We estimate identical versions of our baseline model with these local characteristics—instead of test scores—as outcome variables. Figures A10-A13 illustrate that point estimates are very small and statistically indistinguishable from zero for a wide variety of these local characteristics—including the share of Black and Hispanic students, the share of students eligible for reduced-price lunch, the share of adults with a bachelor's degree, and local unemployment—in years following a sudden increase in library capital spending, consistent with regression results shown with and without controls in Figure 3. Lastly, Figure A13 shows that local school district expenditures, including spending on capital projects and instructional staff, do not change in the years after library investment. This provides evidence that the timing of library capital investment is plausibly orthogonal to local changes that might have affected reading test scores.

V.B Sample and treatment event construction

Our baseline specification matches libraries to school districts to libraries within two miles and defines treatment events as the first time that a district experiences a per student library capital expenditure of $200 or more. In this subsection, we show robustness to both of these choices and demonstrate that our results follow intuitive patterns as we vary both the distance-match and shock-size parameters.

First, we hold the $200 dollar threshold fixed but vary the radius that we use to assign nearby libraries to a given school district. Figure A14 shows versions of the analysis with 1-, 2- (our baseline), 5-, 10-, 15-, and 20-mile matching distances. As we change the thresholds, the results behave intuitively. We observe larger effects at closer distance matches, though standard errors are larger because we have fewer capital spending treatment events. As we increase the distance radius, newly added library events are less salient to a given school district. Consistent with this intuition, we see smaller but more precisely estimated effects as we increase the shock size.

Second, we hold the sample construction distance fixed and vary the size of the library capital spending that defines our treatment events. This allows us to explore the intensive margin of the response. Our baseline estimates define a per student capital shock as $200 dollars or more of nearby library capital spending. Figure A15 shows reading results for this baseline shock size (Panel D) and alternative lower-bound thresholds of $50, $100, $150,..., $300 per student capital shocks in Panels A–F respectively. For a given distance threshold, varying the treatment event threshold produces intuitive results: Smaller (larger) shock thresholds lead to smaller (larger) effects on reading test scores.

Next, we show results with our baseline sample construction and events size where we exclude school districts with multiple capital spending events within our sample period. As discussed in Section III, multiple spending events can affect the interpretation of our results because there may be a dynamic change in the probability of future capital spending events as a function of prior treatment status. Figure A16 shows our baseline reading test results (Panel A) and results where we exclude multi-event districts (Panel B). We estimate very similar results in both cases. This is unsurprising, given that only 16 percent of districts with any event have multiple events. Finally, we show a similar exercise in Figure A17 where we relax our sample requirement that school districts have no more than five branches within two miles of their modal zip code. This affects only a small share of observations, and Figure A17 illustrates that this sample restriction does not affect our conclusions.

V.C Longer test score panel and heterogeneity across cohorts

In this subsection, we use additional datasets to extend the length of the test score panel for a subset of districts. The SEDA data gives us a panel of school district-level reading test scores from 2009-2018, but in Figure 3 we report event-time coefficients over a 14-year span. This means that no school district exists in our panel with both t-7 and t+7 observations. While we observe positive test score effects within 1-2 years of library construction, we cannot estimate a long pre-period for some of the units that generate the larger results that we see in the later portions of the post-event window. In prior sections, we showed little evidence of dynamic selection by both putting potential confounders on the left-hand side of our baseline specification and by controlling for a wide range of covariates in our main regression. The lack of dynamic selection in observable characteristics suggests that panel imbalance at the end of our sample period is unlikely to be driving our results. In this subsection, we describe additional exercises that establish that panel imbalance is unlikely to cause identification or external validity concerns.

First, we extend the test score panel back to before 2009 using two datasets. We combine SEDA test scores with restricted-access data from schools that are part of the National Assessment of Educational Progress (NAEP)—a widely used dataset that gives us 4th- and 8th-grade reading and math test score data for additional districts between 2002 and 2007. We also gather and compile test scores from the National Longitudinal School-Level State Assessment Score Database (NLSLSASD). The NLSLSASD is a precursor to SEDA that attempted to collect standardized math and reading test score information from schools and school districts in states that collected this data. In some cases, the NLSLSASD allows us to observe test scores going as far back as the late 1990s. These datasets are described in more detail in Section II. While both sources are less comprehensive than SEDA, they allow us to observe more pre-period observations for units treated earlier in our sample window. For example, the number of observations we use to estimate the $t = -7$ pre-period coefficient in our baseline model increases by 45 percent when we use both datasets.

Figure A19 shows the results of our baseline analysis when we incorporate these additional datasets. Panel A shows the results after we add NAEP test scores; Panel B shows results where we incorporate both the NAEP and the NLSLSASD. Our results are essentially unchanged when we include both of these new sources of data: we observe flat pre-trends in the years leading up to a library capital spending event and similar post-period effects.

Second, we estimate our baseline model on our 2009-2018 sample after partitioning treated units into two groups based on treatment timing. Figure A20 shows our baseline result estimated for the full sample (black line), units that had a capital spending event before 2014 (short-dotted grey line), and units treated after 2014 (long-dotted grey line). Splitting the sample in this manner increases confidence intervals, but we cannot reject equal estimates across these cohort groups. We see flat pre-trends for both the early- and later-treated groups, and we observe a similar pattern of treatment effect coefficients in the four years after treatment for coefficients where we see estimates from both sets of cohort timing groups.

V.D Additional specifications, measures, and estimation approaches

In this subsection, we show robustness to our results to additional Gardner (2021) specification choices and test score measures. We then show that our results are robust to two alternative estimators in addition to the Gardner (2021) approach.

First, Figure A21 shows reading test score results without conditioning on state-year fixed effects. Panel A shows our baseline model with no state-year covariates while Panel B shows a version that conditions on the set of continuous, time-varying covariates described in Sections II and III. Both panels show similar increases in reading test scores as the main test score results shown in Figure 3.

We next show results for math test scores as an outcome variable using models with and without state-year fixed effects in Figure A22. It is unclear whether to expect increases in math test scores after a library capital construction event. On the one hand, most libraries hold wide-ranging tutoring services for youth that cut across disciplines and hold an extensive range of STEM-focused books and materials. On the other hand, libraries focus on reading, and it is most plausible that we would observe the largest impact on test scores in that subject. Figure A22 shows results consistent with that hypothesis: we see less consistently estimated positive impacts on math than reading scores. There is some evidence of a positive increase in math scores in models that condition on state-year fixed effects, but point estimates are, on average, smaller than the comparable reading test score effects and are less precisely estimated. When not conditioning on state-year fixed effects (Panel A), we see no evidence of an increase in test scores, by contrast to the positive results on reading scores estimated in Figure A21 with or without conditioning on state-year fixed effects.

Next, we estimate our baseline reading results using a two-way fixed effect event estimator equivalent to Equation 1 in Section III. As discussed in Section III, the standard two-way fixed effects estimator may produce biased estimates of the underlying causal effect depending on the nature of treatment effect dynamics. Figure A23 (Panel A) shows a baseline version of a standard event study controlling for district-grade and state-year fixed effects. We see results consistent with our main specifications—public library capital investment leads to a 0.02 standard deviation increase in reading test scores 4–7 years after the investment. In Panel B of Figure A23, we show results from the same methodology but add the time-varying, district-level covariates that we include in our baseline Gardner (2021) analysis. As in the Gardner analysis, the inclusion of these covariates does not change our results.

Lastly, in Figure A24 we show reading results using the Callaway and Sant'Anna method (2021, CS) to calculate treatment effects. CS provide an alternative estimation procedure that, similar to Gardner (2021), avoids issues with comparing later to early-treated units in staggered difference-in-difference settings. Unlike Gardner, the CS method does not easily extend to controlling for time-varying fixed effects and other covariates. However, it can be extended to incorporate pre-treatment covariates to create propensity-score-based matches between treatment and control units. This adjustment is needed if one believes that the parallel trend assumption only holds conditional on covariates. In Panel A and B, we show the CS estimation results respectively without and with this covariate adjustment. In both cases, we observe results that are similar to our baseline Gardner results and the two-way fixed effect results described in the previous paragraph.

VI Housing prices

We next measure the effect of public library investment on local housing prices. Housing prices represent homebuyers' willingness to pay for local amenities. We estimate how the value of homes changes in each zip code within two miles of a school district that experienced a capital spending shock. In recent empirical work, researchers examine the relationship between public school investment, resulting increases in children's test scores, and changes in local housing prices (e.g., Cellini, Ferreira, and Rothstein, 2010; Nielson and Zimmerman, 2014; Conlin and Thompson, 2017; and Bayer, Blair, and Whaley, 2020). These papers discuss and cite an underlying theoretical literature on hedonic regressions and the value of local amenities. If housing prices increase after local school spending increases, homebuyers reveal a preference for an equilibrium increase in school spending: They are willing to pay more for these amenities than the cost they incur as taxpayers.[16]

To identify homebuyers' willingness to pay for local library investment, we use our baseline two-stage Gardner (2021) event study method and our sample of school districts linked to nearby libraries. We calculate two measures of nearby housing prices: Zillow's zip code-level indices (in log dollars) and the FHFA's repeat-price zip code sales index. For each school district, we find all zip codes within two miles of the modal school zip code and calculate the average of each measure of price within those zip codes. We then estimate equations 2 and 3 for each outcome, testing whether housing prices change in the years after library investment.

Figure 10 (Panel A) shows that capital spending shocks have no measurable effect on house prices using Zillow's logged house price index as the outcome. Using 95 percent confidence intervals, we can reject effects larger than a 0.3% decrease or a 1% increase in the year following the capital spending shock. Looking at the dynamic estimates, we can also reject a one percentage point increase or decrease in housing prices up to four years after the shock. There is no apparent pre-trend in housing prices in the years leading up to the capital spending event. Table 2 shows similar results from a pre-/post- difference-in-difference approach, finding null effects with similar precision if we compare the average post-investment house price index value to the pre-investment average. Panel B of Figure 10 and Table 2 show similar null effects for the FHFA measure of housing prices.

The capital spending shocks that we investigate are small enough that any plausible negative effect on housing prices due to increased local taxes cannot be rejected by our results. However, if library spending is valued at more than its taxpayer cost, we might expect to see housing prices rise. For example, Nielson and Zimmerman (2014) find that a school construction program that increased reading test scores by 0.15 standard deviations raised nearby housing prices by 10 percent. Using our 0.02 standard deviation reading test score estimate from five years after a capital

[16]The intuition behind this logic is that (1) local taxes are capitalized into housing prices—if a local community charges a $1,000 tax for each home-owner and does nothing with the money (disposes of it), then housing prices should decrease in proportion to the net present cost of the tax; and (2) if local amenities improve, then the increased willingness to pay for those amenities will be reflected in equilibrium housing prices.

spending shock, a similar level of under-provision would predict a (0.02/0.15)*10% or 1.3 percent increase in local housing prices. Taking these point estimates at face value, our results provide suggestive evidence that households internalize the benefits of the library capital investments and that library spending is not undervalued at a similarly high level as Neilson and Zimmermann found for schools: We can generally rule out a comparable increase in housing prices based on the precision of our housing price results. That said, our confidence intervals are wide enough that we cannot rule out smaller increases in housing prices that might be predicted if households value library investment at more than their cost but less than the under-provision implied by Neilson and Zimmermann's work on schools.

Moreover, we view the lack of pre- and post-trends in housing prices as key evidence that aggregate neighborhood characteristics are not changing in the years surrounding major library spending events, even as children's reading test scores and library usage increase. This finding complements our dynamic selection analysis of Section V, where we show balance in local covariates before and after a library capital spending event. Housing prices take into account any characteristic that could plausibly affect local willingness to pay to live in a community, including characteristics that are difficult to measure. Consistent with our finding that other community characteristics are not changing before or after library investment, we observe no significant change in housing prices before and after a library spending event.

VII Comparisons to public school capital investment

In our baseline analysis, we define a capital spending treatment event as the first year that a school district has a $200 per student (or greater) local public library spending event (subject to not having more than $200 per student in capital spending in the prior year). To convert this number into something that is amenable to cost-benefit analyses, we make a number of adjustments. First, the average size of capital investments that meet this threshold or are above it is $810 per student in year zero. Second, as shown in Figure 1, our definition of a capital spending event identifies the first year of on-average five years of elevated capital spending. Adding these years of elevated spending to the cost of our event, our capital spending implies that our capital spending shocks are equivalent to $2,028 per student in capital spending over five years.

Focusing on our overall pre-post effect, a $2,028 per student one-time library capital spending shock increases test scores by 0.01 standard deviations. In a recent meta-analysis, Jackson and Mackevicius (2023) find that a one-time $1,000,000 renovation project affecting a school with 600 students—equivalent to a $\frac{\$1{,}000{,}000}{600} = \$1,667$ one-time capital expenditure per affected student—causes a 0.00225 standard deviation increase in average test scores across the papers that they review. Comparing our results to that paper, we find that our library capital spending shocks generate average test score increases that are roughly five times as large as the results from Jackson and Mackevicius at a $(1\text{-}\frac{2{,}000}{1{,}667}) = 20\%$ higher initial capital cost.

This does not provide evidence that library capital projects have higher benefit-cost ratios, and we caution against this direct comparison for several reasons. First, school capital spending is orders of magnitude more costly than library spending, and both are likely subject to diminishing marginal returns. Jackson and Mackevicius (2023) report that the typical construction of a new elementary school costs $44,000 per student. It is difficult to know if school capital spending effects scale linearly in investment amounts — if the first dollars of school spending have larger effects, our comparison will overstate the relative benefit of libraries. Second, while we focus on capital spending to define our events, Figure 2 shows that operations spending at libraries increases by 2-7 percent in the years after a capital spending investment. In their meta-analysis, Jackson and Mackevicius find point estimates for the marginal effect of per dollar school operations spending that are larger than the effects of capital spending. An implicit assumption in our comparison is that school and library capital spending leads to similar proportional increases in operating spending in the years following the capital spending. Although, given the small dollar increases in library operating spending in the years following the capital spending shock, we expect this mechanism to be relatively unimportant, adding only marginally to the per student cost of our typical library spending shock.

Next, the 95% confidence interval around our estimate contains test score effects that are both significantly larger and smaller than the meta-analyzed overall effect from Jackson and

Mackevicius (2023). This is a natural result of our paper being the first to observe a causal effect of public library investment on children's test scores; we cannot improve precision by pooling our estimate with other independent causal estimates. More work on the effects of public libraries would help to pin down these effects.

Finally, our grade 3–8 test score effects likely underestimate the importance of libraries: a single library expansion could affect others in surrounding areas, students in grades K–2 and 9–12, achievement in subjects where test scores are not measured in grades 3–8 (e.g., science), and adults and non-school-aged children (e.g., pre-K programs).[17]

For these reasons, we do not think that our dollar-to-dollar cost/benefit comparisons of library and school capital investments can be directly used to determine optimal policy. In addition, the heterogeneity in Figure 9 shows that the marginal effects of library spending are increasing as a function of local school capital spending, suggesting a complementary relationship between the two types of spending. But our results do suggest that under reasonable assumptions: (1) library capital spending, like school capital spending, can have positive effects on local test scores, and (2) smaller-scale public investments in local amenities — like libraries — can positively affect children's academic performance.

VIII Conclusion

A growing literature suggests that neighborhoods can have long-lasting and important effects on child development. However, knowing that neighborhoods matter is not enough: for place-based policies to be effective, we need to know which characteristics of local communities cause changes in childhood experiences and which are simply correlated with desirable community characteristics. In this paper, we study the causal effects of investments in one of the most commonly used and lauded functions of local governments: the public library.

Every year, library administrators in local communities petition voters to approve additional

[17] School construction also may have broader public impacts beyond direct effects on children, since school facilities are often used by other members of the community. We do not attempt to compare the possible spillover effects across libraries and school districts.

funds to expand and improve local public libraries. To date, little is known about the impact of these investments. We use an event study methodology to show that library capital spending has at least two effects. First, capital library investments cause patrons to use the library more. More residents visit the library, the stock of library materials increases, and more children attend library events. Second, library capital investments cause students to perform better on standardized tests. In particular, we observe persistent improvement in reading test scores in the seven years after library capital spending. These improvements in reading test scores coincide with increases in library use and are not driven by changing demographics in local communities.

We study public library capital spending in the post-smartphone era, so our results highlight the importance of public libraries to children, even in a world with widespread access to the internet and other new technologies that compete for children's attention. We hope more work will investigate the many other ways that public libraries can affect local communities and investigate how these effects differ for children of different races and socioeconomic characteristics.

References

[1] American Library Association, Diversity Counts [database], 2012, accessed at https://www.ala.org/aboutala/offices/diversity/diversitycounts/divcounts

[2] Baker, Andrew C., David F. Larcker, and Charles CY Wang. 2022. "How much should we trust staggered difference-in-differences estimates?" *Journal of Financial Economics* 144 (2): 370-395.

[3] Baron, E. Jason. 2022. "School spending and student outcomes: Evidence from revenue limit elections in Wisconsin." *American Economic Journal: Economic Policy* 14 (1): 1-39.

[4] Bayer, Patrick, Peter Q. Blair, and Kenneth Whaley. 2020. "A National Study of School Spending and House Prices."

[5] Bekkerman, Anton, and Gregory Gilpin. 2013. "High-speed Internet Growth and the Demand for Locally Accessible Information Content." *Journal of Urban Economics* 77: 1-10.

[6] Berkes, Enrico, and Peter Nencka. 2023. "Knowledge Access: The Effects of Carnegie Libraries on Innovation."

[7] Bhatt, Rachana. 2010. "The Impact of Public Library Use on Reading, Television, and Academic Outcomes." *Journal of Urban Economics* 68 (2): 148-166.

[8] Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2021. "Revisiting event study designs: Robust and efficient estimation." arXiv preprint arXiv:2108.12419.

[9] Bogin, Alexander N., William M. Doerner, and William D. Larson. 2019. "Local House Price Dynamics: New Indices and Stylized Facts." *Real Estate Economics* 47 (2): 365-398.

[10] Brunner, Eric, Ben Hoen, and Joshua Hyman. 2022. "School district revenue shocks, resource allocations, and student achievement: Evidence from the universe of US wind energy installations." *Journal of Public Economics* 206: 104586.

[11] Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna. 2021. "Difference-in-differences with a continuous treatment." arXiv preprint arXiv:2107.02637.

[12] Callaway, Brantly, and Pedro HC Sant'Anna. 2021. "Difference-in-differences with multiple time periods." *Journal of Econometrics* 225 (2): 200-230.

[13] Cellini, Stephanie Riegg, Fernando Ferreira, and Jesse Rothstein. 2010. "The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design." *Quarterly Journal of Economics* 125 (1).

[14] Chen, Jiafeng and Jonathan Roth. 2023. "Log-like? Identified ATEs defined with zero-valued outcomes are (arbitrarily) scale-dependent." arXiv preprint arXiv:2212.06080.

[15] Cheng, Guang, Zhuqing Yu, and Jianhua Z. Huang. 2013. "The cluster bootstrap consistency in generalized estimating equations." *Journal of Multivariate Analysis* 115: 33-47.

[16] Chetty, Raj, and Nathaniel Hendren. 2018. "The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects." *The Quarterly Journal of Economics* 133 (3): 1107-1162.

[17] Chetty, Raj, and Nathaniel Hendren. 2018. "The impacts of neighborhoods on intergenerational mobility II: County-level estimates." *The Quarterly Journal of Economics* 133 (3): 1163-1228.

[18] Chyn, Eric. 2018. "Moved to opportunity: The long-run effects of public housing demolition on children." *American Economic Review* 108 (10): 3028-56.

[19] Common Core of Data, Education Data Portal (Version 0.16.0), Urban Institute, accessed August, 2022, https://educationdata.urban.org/documentation/, made available under the ODC Attribution License.

[20] Conlin, Michael, and Paul N. Thompson. 2017. "Impacts of new school facility construction: An analysis of a state-financed capital subsidy program in Ohio." *Economics of Education Review* 59.

[21] DeBoer, Larry. 1992. "Economies of scale and input substitution in public libraries." *Journal of Urban Economics*, 32 (2): 257–268.

[22] Donnelly, Francis P. 2015. "Regional variations in average distance to public libraries in the United States." *Library Information Science Research*, 37 (4): 280-289.

[23] Fahle, Erin M., Benjamin R. Shear, Demetra Kalogrides, Sean F. Reardon, Belen Chavez, and Andrew D. Ho. 2021. "Stanford Education Data Archive Technical Documentation Version 4.0 February 2021."

[24] Federal Housing Finance Agency. 2019. "House Price Index." Accessed 2019. https://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index.aspx.

[25] Feldstein, Kathleen Foley. 1977. "The Economics of Public Libraries." PhD dissertation, Massachusetts Institute of Technology.

[26] Gardner, John. 2021. "Two-stage differences in differences." arXiv preprint arXiv:2207.05943.

[27] Getz, Malcolm. 1980. *Public Libraries: An Economic View*. Baltimore, MD: Johns Hopkins University Press.

[28] Goldhor, Herbert, and John McCrossan. 1966. "An exploratory study of the effect of a public library summer reading club on reading skills." *Library Quarterly*, 36 (1): 14-24.

[29] Guryan, Jonathan, James Kim, and David Quinn. 2014. "Does reading during the summer build reading skills? Evidence from a randomized experiment in 463 classrooms." *NBER Working Paper 20689*.

[30] Heckman, James J. 2007. "The economics, technology, and neuroscience of human capability formation." *Proceedings of the National Academy of Sciences*, 104 (33): 13250-13255.

[31] Hemmeter, Andrew. 2006. "Estimating Public Library Efficiency Using Stochastic Frontiers." *Public Finance Review*, 34 (3): 328-348.

[32] IMLS. 2019. Public Library Survey Fiscal Year 2019 Data File Documentation and User's Guide [technical documentation].

[33] IMLS. "Public Library Survey." 2020. [dataset] Accessed 2020. https://www.imls.gov/research-evaluation/data-collection/public-libraries-survey.

[34] Jackson, C. Kirabo and Claire Mackevicius. 2023. "What Impacts Can We Expect from School Spending Policy? Evidence from Evaluations in the U.S." Forthcoming in *American Economic Journal: Applied Economics.*

[35] James, Stephen. 1985. "The Relationship between Local Economic Conditions and the Use of Public Libraries." *Library Quarterly*, 55 (3): 255-272.

[36] Karger, Ezra. 2021. "The Long-Run Effect of Public Libraries on Children."

[37] Kevane, Michael J., and William A. Sundstrom. 2016. "Public Libraries and Political Participation, 1870-1940."

[38] Klinenberg, Eric. 2018. *Palaces for the People: How Social Infrastructure can Help Fight Inequality, Polarization, and the Decline of Civic Life.* Broadway Books.

[39] McLaughlin, D.H., V. Bandeira de Mello, S. Cole, C. Blankenship, H. Hikawa, K. Farr, and R. Gonzalez. 2002. *National Longitudinal School-Level State Assessment Score Database: Analyses of 2000/2001 school-year scores.* Report submitted to the U.S. Department of Education by the American Institutes for Research.

[40] Neto, Amir B. Ferreira, and Joshua C. Hall. 2019. "Economies of Scale and Governance of Library Systems: Evidence from West Virginia." *Economics of Governance*, 20 (3): 237-253.

[41] Neto, Amir B. Ferreira, Jennifer Nowicki, and Shishir Shakya. 2021. "Do Public Libraries Help Mitigate Crime? Evidence from Kansas City, MO."

[42] Neto, Amir B. Ferreira. 2022. "Do public libraries impact local labour markets? Evidence from Appalachia." *Spatial Economic Analysis*, 1-23.

[43] Neilson, Elijah. 2020. *Essays on the Determinants of Human Capital Investment.* PhD diss., Clemson University.

[44] Nielson, Christopher A., and Seth D. Zimmerman. 2014. "The effect of school construction on test scores, school enrollment, and home prices." *Journal of Public Economics*, 120: 18-31.

[45] Ottensmann, J. 1997. "Using Geographic Information Systems to Analyze Library Utilization." *Library Quarterly*, 67 (1): 24-49.

[46] Organization for Economic Co-operation and Development, Consumer Price Index: All Items for the United States [USACPIALLAINMEI], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/USACPIALLAINMEI, accessed May 24, 2022.

[47] Palmer, E. 1981. "The Effect of Distance on Public Library Use: A Literature Survey." *Library Research*, 3 (4): 315-354.

[48] Pew Research Center. 2015. "Chapter 1: Who Uses Libraries and What They do at Their Libraries." Retrieved from https://www.pewresearch.org/internet/2015/09/15/who-uses-libraries-and-what-they-do-at-their-libraries/.

[49] Porter, Andrew Joseph. 2015. *Essays in Crime and Behavior*. PhD diss., University of California, Irvine.

[50] Reardon, S. F., Ho, A. D., Shear, B. R., Fahle, E. M., Kalogrides, D., Jang, H., and Chavez, B. 2020. Stanford Education Data Archive (Version 4.0). Retrieved from https://edopportunity.org/get-the-data/.

[51] Rubin, Donald B. 1981. "The Bayesian Bootstrap." *The Annals of Statistics*: 130-134.

[52] Rodríguez-Lesmes, Paul, Josè D. Trujillo, and Daniel Valderrama. 2014. "Are public libraries improving quality of education? When the provision of public goods is not enough." *Revista Desarrollo y Sociedad* 74: 225-274.

[53] Thompson, Owen. 2022. "School desegregation and Black teacher employment." *Review of Economics and Statistics* 104 (5): 962-980.

[54] Tiebout, Charles M., and Robert J. Willis. 1965. "The Public Nature of Libraries." In *The Public Library and the City*: 94-101.

[55] U.S. Department of Education. Institute of Education Sciences, National Center for Education Statistics, National Assessment of Educational Progress (NAEP), Assessments (accessed 2022) [dataset]

[56] Vitaliano, Donald F. 1997. "X-Inefficiency in the Public Sector: the Case of Libraries" *Public Finance Review* 25 (6): 629-643.

[57] Vitaliano, Donald F. 1998. "Assessing Public Library Efficiency using Data Envelopment Analysis." *Annals of Public and Cooperative Economics* 69 (1): 107-122.

[58] Wiegand, Shirley A., and Wayne A. Wiegand. 2018. *The Desegregation of Public Libraries in the Jim Crow South: Civil rights and local activism*. LSU Press.

[59] Worthington, Andrew. 1999. "Performance Indicators and Efficiency Measurement in Public Libraries." *Australian Economic Review* 32 (1): 31-42.

[60] Xie, Claire, and Joel Waldfogel. 2022. "The First Sale Doctrine and the Digital Challenge to Public Libraries."

[61] Zillow Group. 2020. "Housing Price Data." Accessed 2020. https://www.zillow.com/research/data/.

Figure 1: Impact of capital expenditure shock on library use

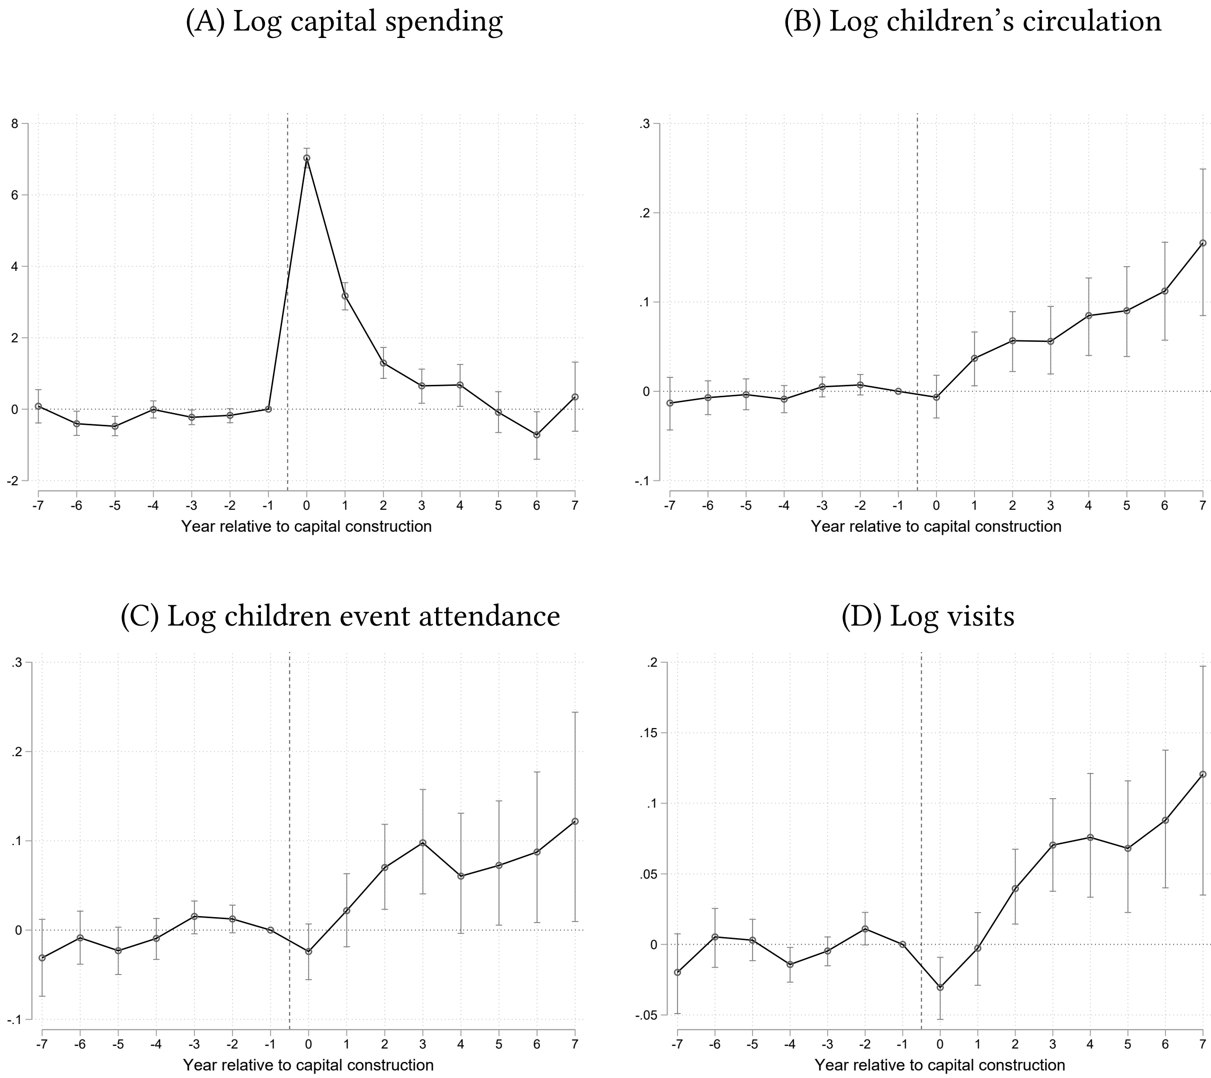

Event study estimates generated using the Gardner estimation procedure described in Section III. Outcomes are aggregated to the school district-level. The outcome variables are (A) Log capital spending, (B) log children's circulation, (D) log children's event attendance, and (D) log visits. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Results are conditional on state-year and district-grade fixed effects.

Figure 2: Impact of capital expenditure shock on library resources

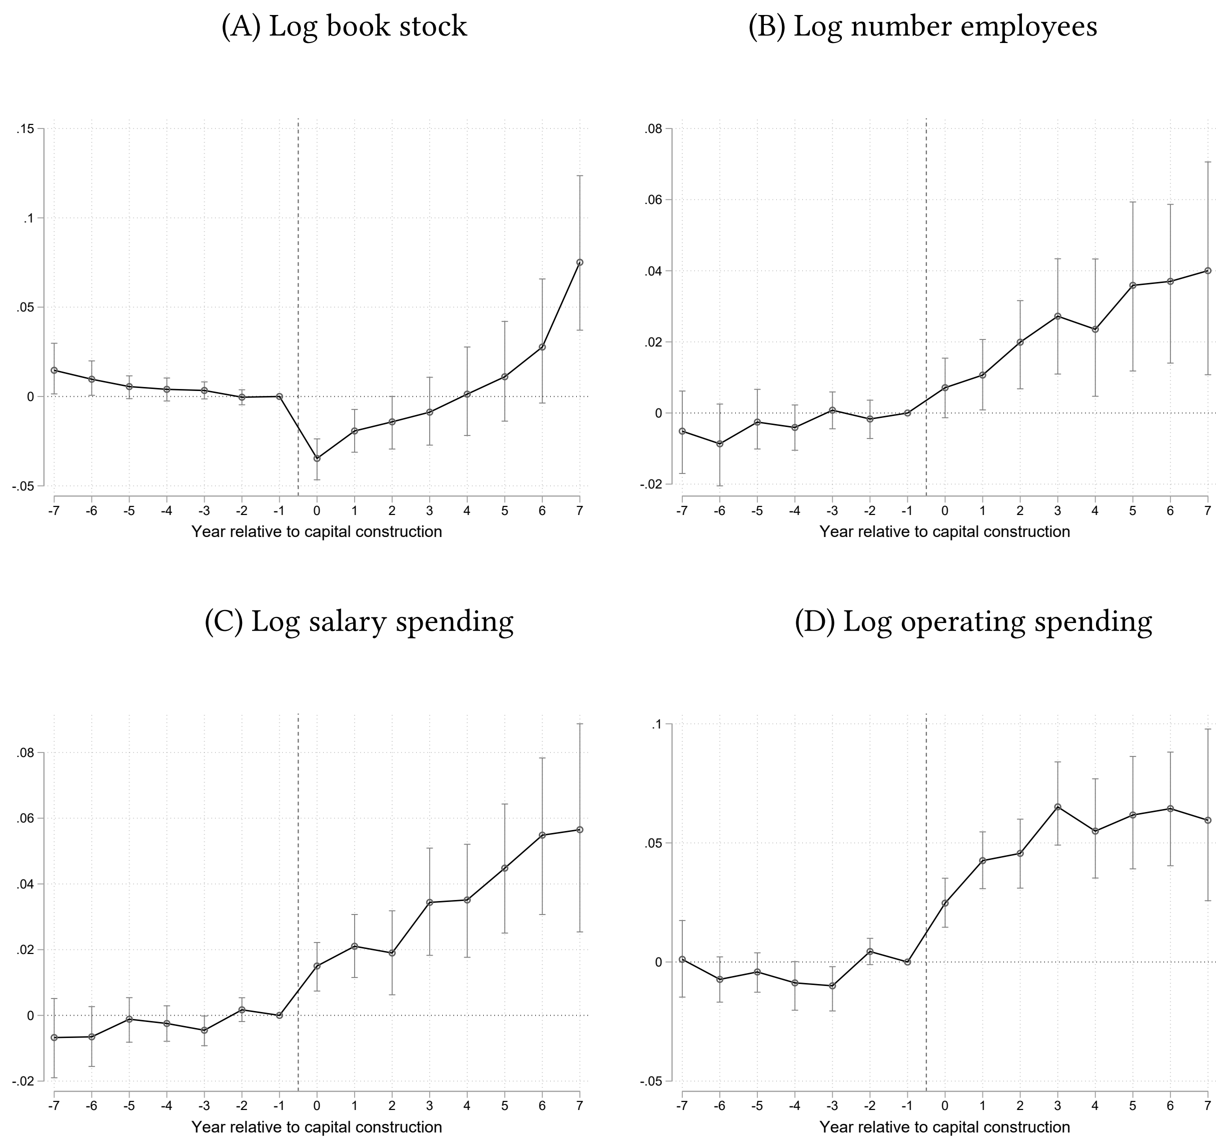

Event study estimates generated using the Gardner estimation procedure described in Section III. Outcomes are aggregated to the school district-level. The outcome variables are (A) log books, (B) log number employees, (C) log expenditures on salaries, and (D) log operating expenditures. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Results are conditional on state-year and district-grade fixed effects.

Figure 3: Impact of library capital spending shocks on reading test scores

(a) Reading test scores

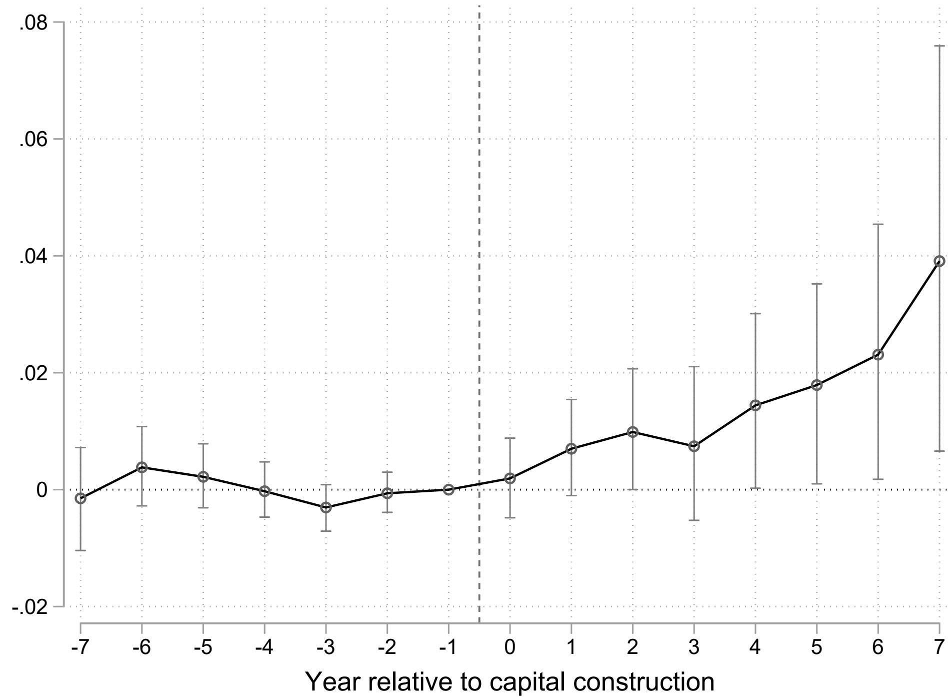

(b) Reading test scores, conditional on time-varying local covariates

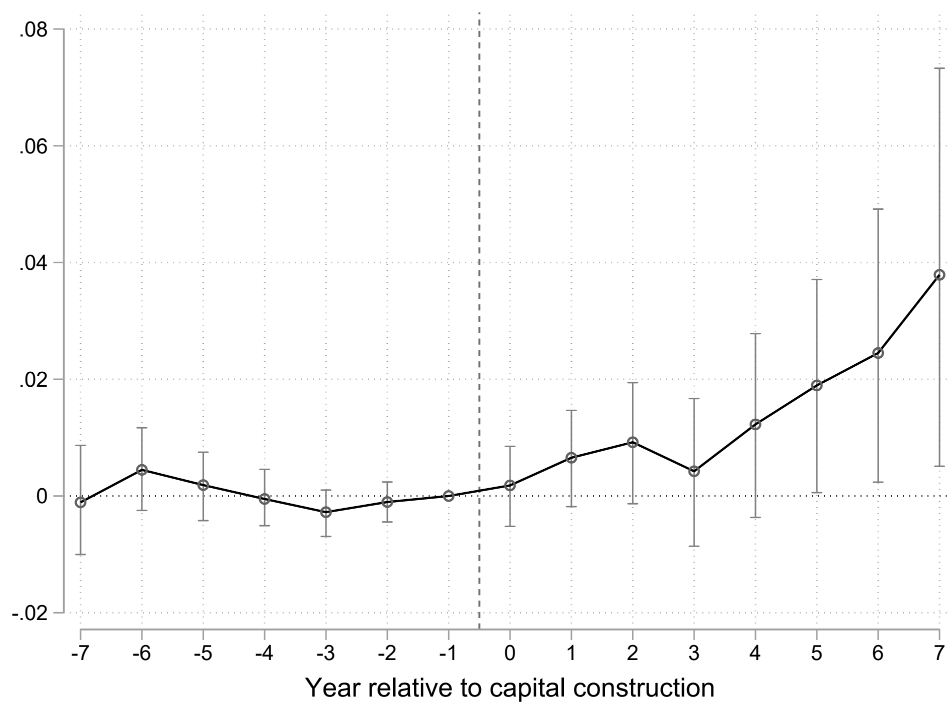

Event study estimates generated using the Gardner estimation procedure described in Section III. Outcomes are measured at the grade-school district-level. All figures show bootstrapped 95 percent confidence intervals, calculated by drawing bootstrap samples that account for clustering within school districts. All results are conditional on state-year and district-grade fixed effects. The results in Panel B are additionally conditional on a set of time-varying, district-level covariates described in Section III.

Figure 4: Impact of library capital spending shocks on reading test scores, distance and shock size placebo tests

(a) Capital shocks 30-50 miles away

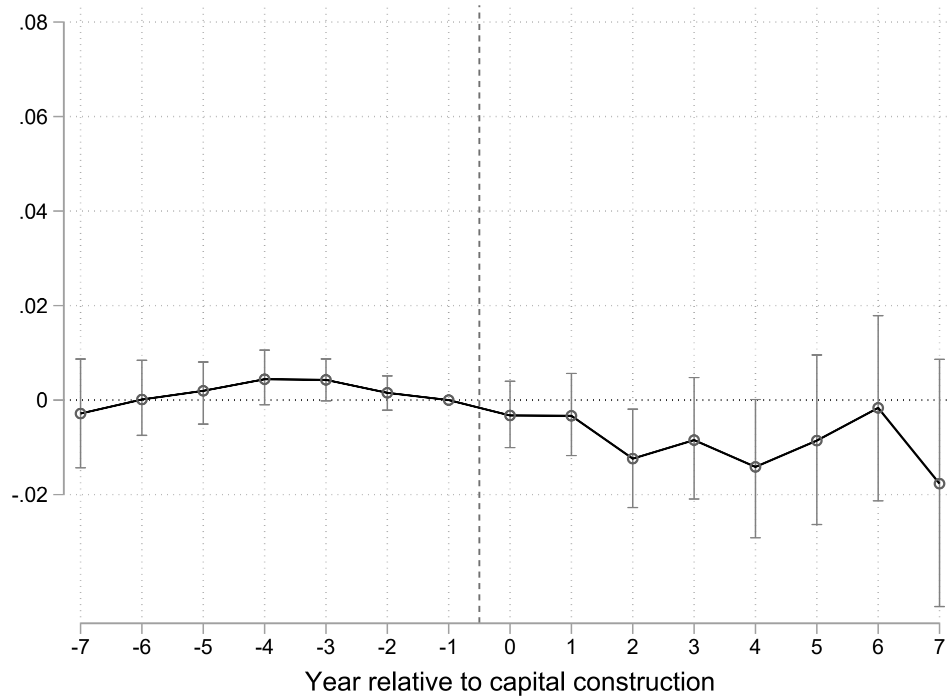

(b) Small capital shocks

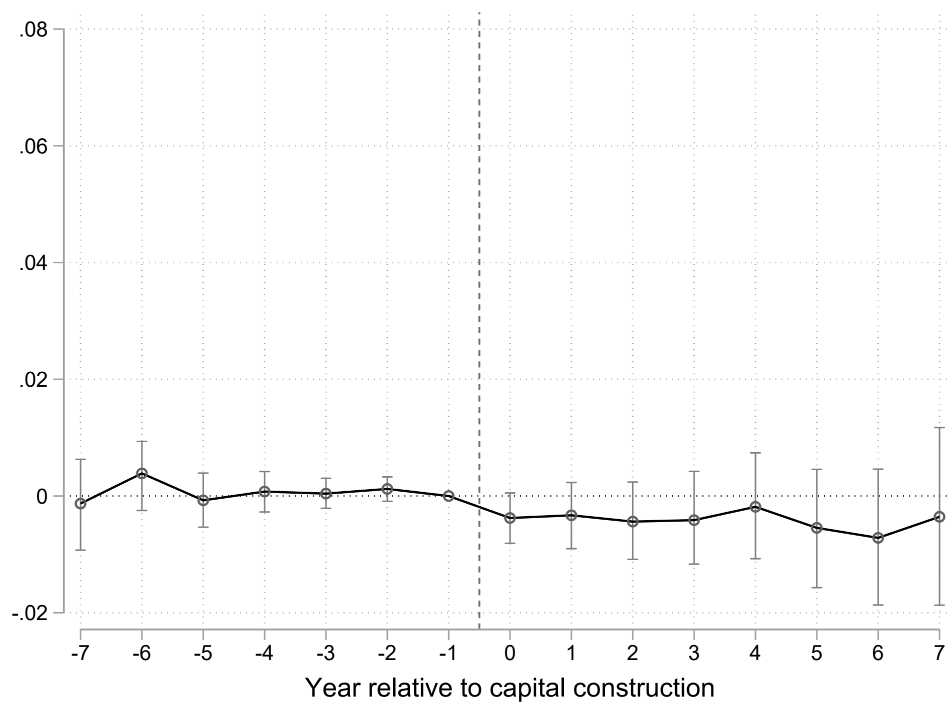

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals, calculated by drawing bootstrap samples that account for clustering within school districts. Regressions include state-year and district-grade fixed effects. Small capital shocks are defined as spending events between 1 and 30 dollars per student.

Figure 5: Heterogeneous reading test score effects by student race and ethnicity

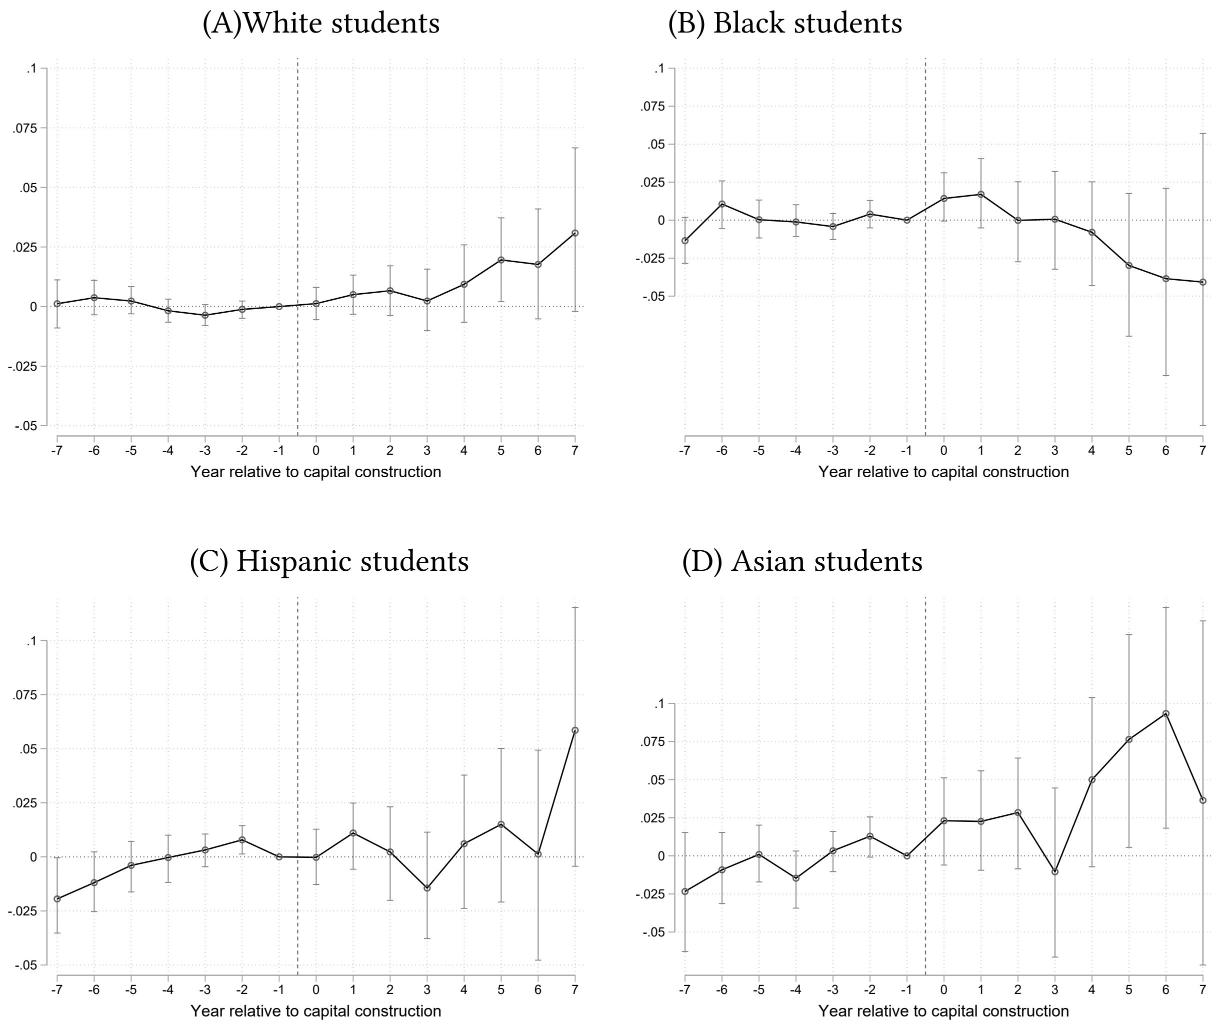

Event study estimates generated using the Gardner estimation procedure described in Section III. Effects are estimated separately for each indicated subgroup. All figures show bootstrapped 95 percent confidence intervals, calculated by drawing bootstrap samples that account for clustering within school districts. Regressions include state-year and district-grade fixed effects.

Figure 6: Heterogeneous reading test score effects by student economic status

(a) Economically disadvantaged

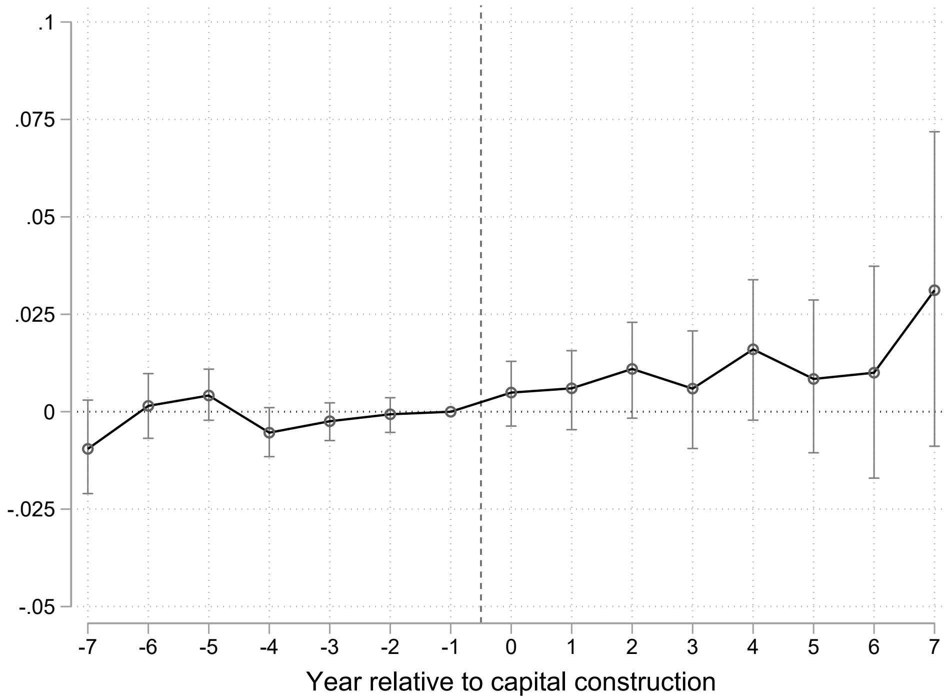

(b) Not economically disadvantaged

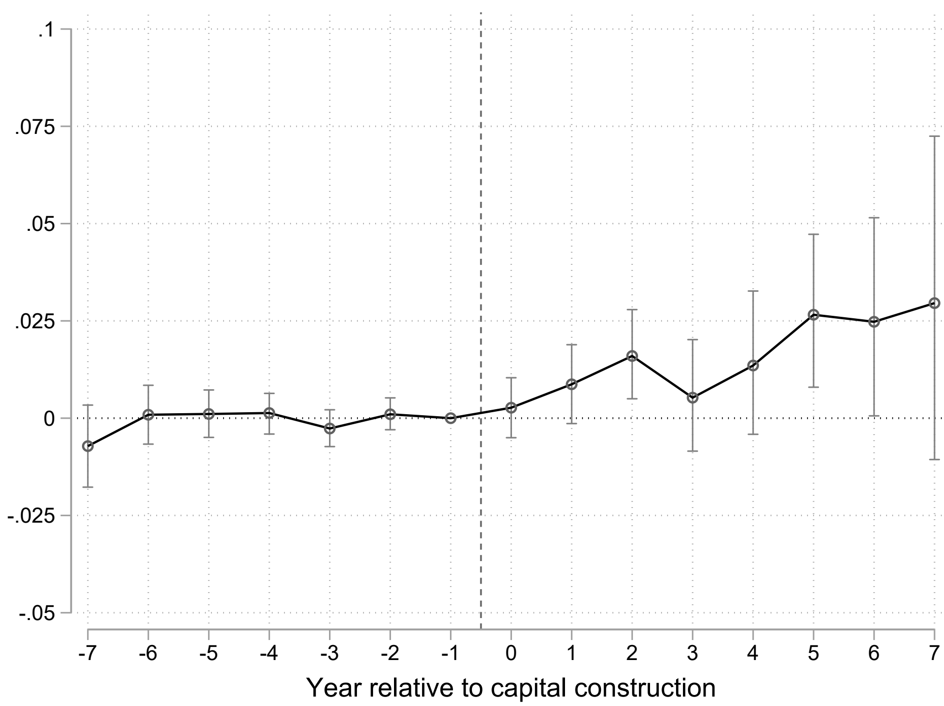

Event study estimates generated using the Gardner estimation procedure described in Section III. Effects are estimated separately for each indicated subgroup. All figures show bootstrapped 95 percent confidence intervals, calculated by drawing bootstrap samples that account for clustering within school districts. Regressions include state-year and district-grade fixed effects.

Figure 7: Heterogeneous reading test score effects by student grades

(a) Grades 3-4

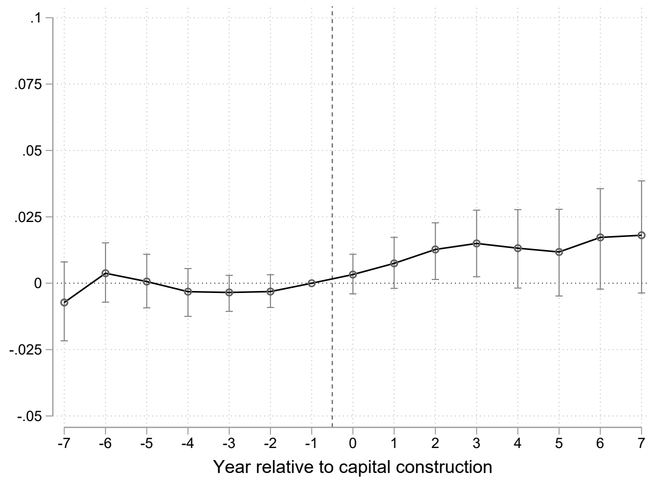

(b) Grades 5-6

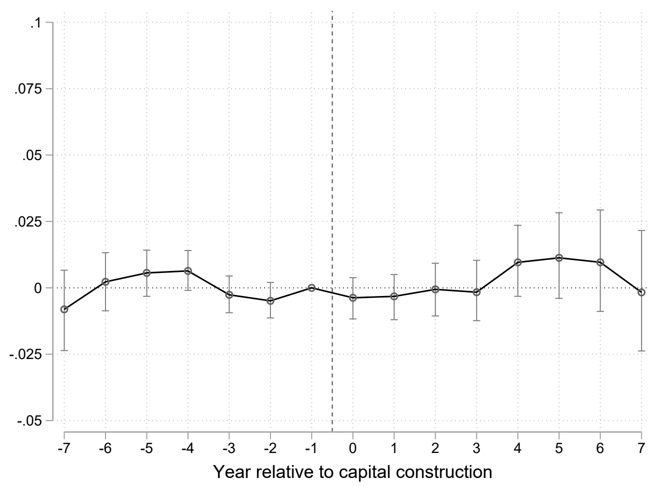

(c) Grades 7-8

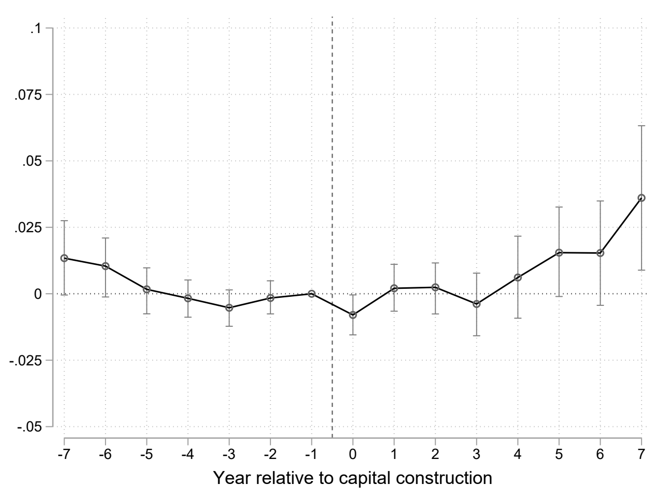

Event study estimates generated using the Gardner estimation procedure described in Section III. Effects are estimated separately for each indicated subgroup. All figures show bootstrapped 95 percent confidence intervals, calculated by drawing bootstrap samples that account for clustering within school districts. Regressions include state-year and district-grade fixed effects.

Figure 8: Heterogeneous reading test score effects by school district size

(a) Smallest size tercile

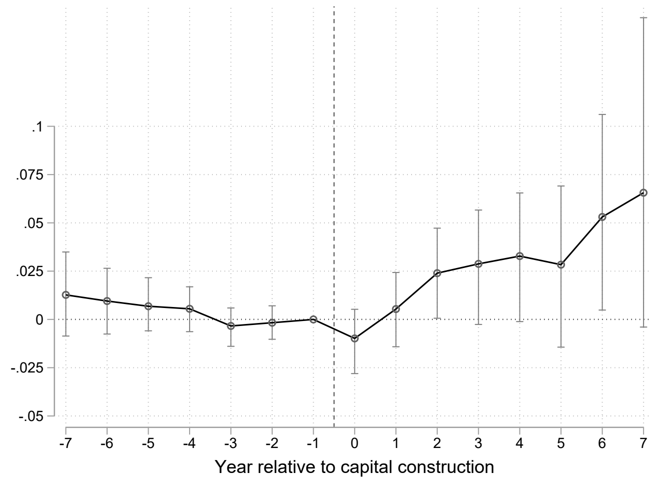

(b) Middle size tercile

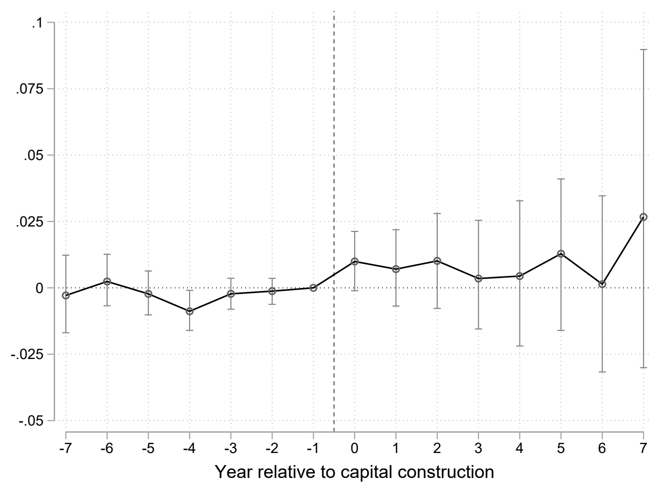

(c) Largest size tercile

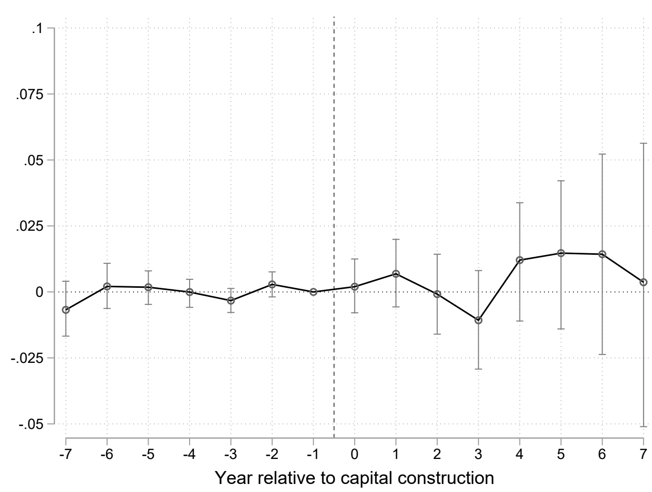

Event study estimates generated using the Gardner estimation procedure described in Section III. Effects are estimated separately for each indicated subgroup. All figures show bootstrapped 95 percent confidence intervals, calculated by drawing bootstrap samples that account for clustering within school districts. Regressions include state-year and district-grade fixed effects.

Figure 9: Heterogeneous reading test score effects by school district capital spending

(a) Smallest spending tercile

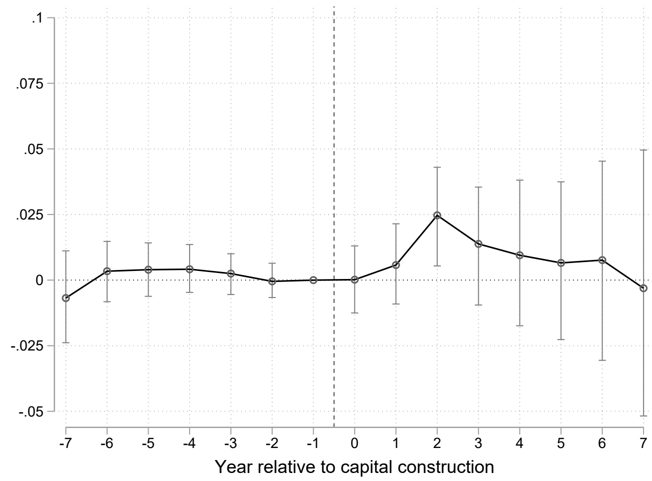

(b) Middle spending tercile

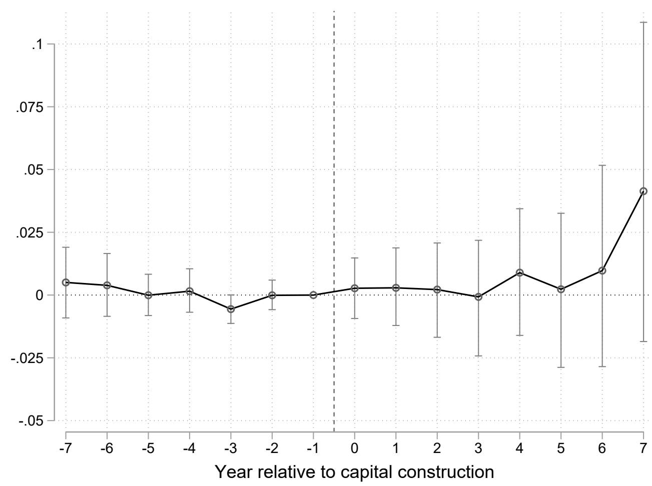

(c) Largest spending tercile

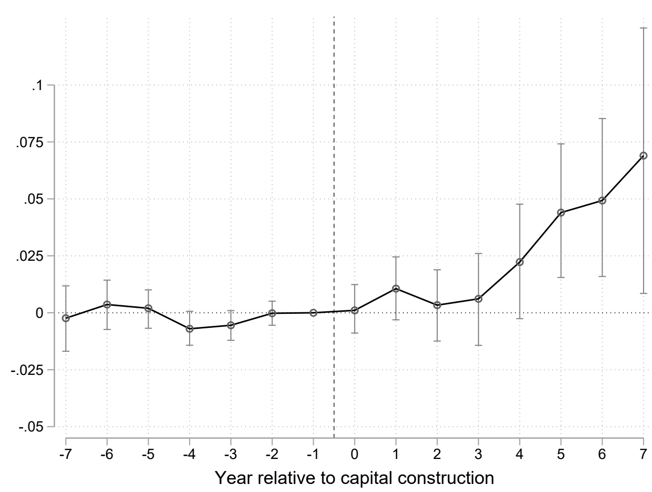

Event study estimates generated using the Gardner estimation procedure described in Section III. Effects are estimated separately for each indicated subgroup. All figures show bootstrapped 95 percent confidence intervals, calculated by drawing bootstrap samples that account for clustering within school districts. Regressions include state-year and district-grade fixed effects.

Figure 10: Impact of library capital spending on local housing prices

(a) Zillow house price index (log dollars)

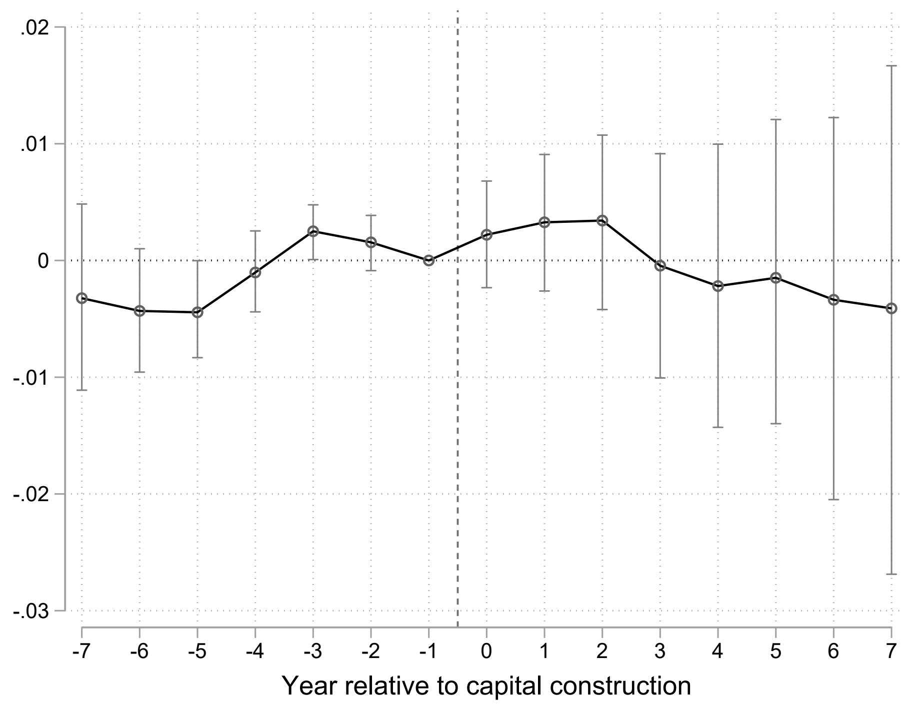

(b) FHFA housing index (percentage points)

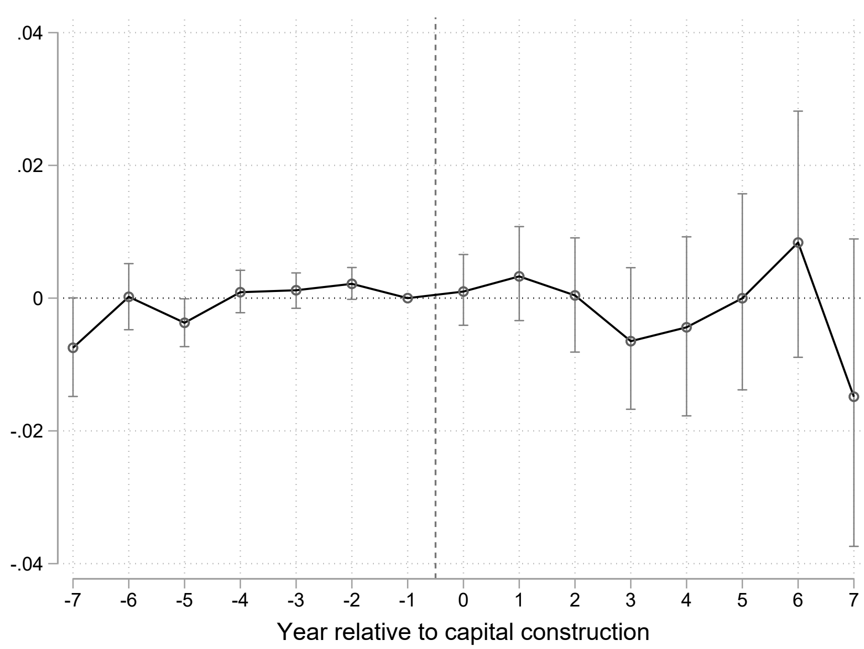

Event study estimates generated using the Gardner estimation procedure described in Section III. Outcomes are logged housing price data from Zillow (Panel A) and the FHFA housing index in percentage points (Panel B) after aggregating to the school district-level using a 2 mile distance threshold. Bootstrapped standard errors are clustered by school district.

Table 1: Summary statistics: school district-by-year panel

	Obs.	Mean	10th %	Median	90th %	Std. Dev.
Per student library capital spending ($s)	517,789.00	43.60	0.00	0.00	36.70	421.82
Child materials circulation (thousands)	279,137.00	59.35	2.63	20.56	153.76	111.53
Child event attendees (thousands)	279,191.00	5.83	0.30	2.68	14.72	9.71
Visits (thousands)	274,237.00	111.15	7.85	50.84	279.95	164.70
Books (thousands)	281,984.00	64.71	13.30	38.15	147.49	76.38
Employees	281,557.00	11.20	1.13	5.31	27.53	16.87
Salary spending (thousands)	217,933.00	561.13	72.86	245.00	1,396.89	873.83
Operating spending (thousands)	280,064.00	856.42	53.58	306.11	2,174.35	1,549.42
Library systems	517,789.00	0.62	0.00	1.00	1.00	0.69
Branches	517,789.00	0.15	0.00	0.00	0.00	0.57
Reading test scores	517,789.00	0.03	−0.43	0.04	0.49	0.37
Math test scores	488,470.00	0.04	−0.47	0.04	0.55	0.41
Reading test-takers	517,789.00	330.55	32.00	123.00	682.00	1,009.68
Math test-takers	488,685.00	322.35	32.00	122.00	661.00	1,009.40
Share Native American students	517,595.00	0.02	0.00	0.00	0.03	0.10
Share Asian students	517,595.00	0.02	0.00	0.01	0.05	0.05
Share Hispanic students	517,595.00	0.13	0.00	0.05	0.42	0.20
Share Black students	517,595.00	0.08	0.00	0.02	0.24	0.17
Share free lunch elig. students	517,595.00	0.40	0.13	0.38	0.69	0.21
Share reduced lunch elig. students	517,595.00	0.08	0.03	0.08	0.13	0.05
Share econ. disadvantaged students	513,040.00	0.49	0.18	0.49	0.79	0.23
Share BA+ adults	514,327.00	0.23	0.11	0.19	0.41	0.13
Share single mom adults	514,327.00	0.15	0.09	0.14	0.23	0.06
Share SNAP-eligible adults	514,327.00	0.11	0.03	0.10	0.19	0.06
Share unemployed adults	514,327.00	0.07	0.04	0.07	0.11	0.03
Federal revenue (per student, $1,000s)	516,432.00	1,091.44	405.11	899.44	1,827.88	1,043.51
State revenue (per student, $1,000s)	516,432.00	6,198.13	3,143.41	5,865.38	9,260.04	2,958.23
Local revenue (per student, $1,000s)	516,432.00	6,278.83	2,189.82	4,894.67	12,360.64	5,007.19
Total expenses (per student, $1,000s)	516,423.00	13,468.77	9,024.17	12,032.26	19,622.88	5,516.71
Capital expenses (per student, $1,000s)	516,426.00	1,161.29	105.72	504.50	2,608.77	2,347.11
Capital construction expenses (per student, $1,000s)	516,411.00	849.92	0.00	194.77	2,102.42	2,227.31
Salary expenses (per student, $1,000s)	516,355.00	6,544.41	4,677.52	6,009.20	9,288.75	2,176.58
Instructional salary expenses (per student, $1,000s)	516,355.00	4,463.57	3,175.04	4,083.64	6,368.96	1,520.88

Notes: This table shows summary statistics describing our school district-year panel described in Section II. All summary statistics are at the school district-by-year-level for school districts with fewer than 20 library buildings within two miles of the main zip code of the school district. Data cover the years 2009—2018 because those are the years for which we have test score data. Data come from (1) the IMLS's PLS census of almost all libraries in the United States; (2) the SEDA test score dataset (see Fahle et al., 2021 for more details); and (3) the Common Core of Data—a database produced by NCES describing school-level and district-level financing information. Test-taker counts, student demographics, and the number of students on a free or reduced price lunch program come from the SEDA data.

Table 2: Coefficient estimates for post-treatment period, by outcome and relative year

Relative time	Cap. spending	Children circ.	Children event attend	Visits	Reading scores	FHFA Housing	Zillow Housing
1	3.170	0.037	0.022	-0.003	0.007	0.003	0.003
	[2.776, 3.537]	[0.006, 0.066]	[-0.019, 0.063]	[-0.029, 0.022]	[-0.001, 0.015]	[-0.003, 0.011]	[-0.003, 0.009]
3	0.653	0.056	0.098	0.070	0.007	-0.006	0.000
	[0.166, 1.120]	[0.019, 0.095]	[0.040, 0.157]	[0.038, 0.103]	[-0.005, 0.021]	[-0.017, 0.005]	[-0.010, 0.009]
5	-0.089	0.090	0.072	0.068	0.018	0.000	-0.001
	[-0.657, 0.488]	[0.039, 0.139]	[0.005, 0.145]	[0.023, 0.116]	[0.001, 0.035]	[-0.014, 0.016]	[-0.014, 0.012]
7	0.343	0.166	0.122	0.121	0.039	-0.015	-0.004
	[-0.625, 1.310]	[0.085, 0.248]	[0.010, 0.244]	[0.035, 0.197]	[0.007, 0.076]	[-0.038, 0.009]	[-0.027, 0.017]
Pre-/Post-	2.511	0.052	0.046	0.032	0.010	-0.000	0.001
	[2.188, 2.803]	[0.024, 0.078]	[0.006, 0.084]	[0.007, 0.057]	[0.002, 0.020]	[-0.007, 0.007]	[-0.006, 0.008]

Notes: This table shows point estimates from the event studies measuring the effect of capital spending shocks on the outcome named in the column heading. Under each point estimate is a bootstrapped 95% confidence interval that takes into account clustering within school districts. The first four rows of point estimates indicate the causal effects 1, 3, 5, and 7 years after the capital spending shock. The last row of point estimates indicates the effect of the capital spending shock in the post-treatment periods for treated school districts relative to pre-period observations for treated school districts. Point estimates rely on the method of Gardner (2021).

Online Appendix: The Returns to Public Library Investment

Gregory Gilpin Ezra Karger Peter Nencka

A Data Appendix

This appendix contains additional information on the datasets that we use in this project.

A.A Stanford Educational Data Archive

We use version 4.00 of the Stanford Educational Data Archive (SEDA) to track test scores for youth across time, as well as the characteristics of nearby local communities.

SEDA reflects an effort to standardize data reported by the states about the 3-8th graders and their performance on mandated language arts and math exams. In particular, Federal law requires these annual tests, but does not prescribe the exact form that they must take. States choose the format of their test and report the results as the number of students in each school who score above performance-based thresholds. The number of thresholds can vary by states and ranges from 2 to 5 cutoffs in our sample.

The staff that constructed SEDA have access to restricted data with full information of the number of students who scored above each threshold and their characteristics from 2008-2018. This is an improvement relative to publicly released data on ED*Facts* which suppresses information for small schools and student subgroups. SEDA staff clean these test scores and then estimate Heteroskedastic Ordered Probits to estimate the state-grade-subject specific test score cutoffs. They then standardize these cutoffs by linking them to information from the Nation's Report Card (NAEP). Since the NAEP is a common test across states, this process creates test scores that can be compared across states. More information on the construction of the test score data is given in Fahle et al (2021). SEDA standardizes test scores at the district-grade-subject-level; we do not observe individual-level school information in the SEDA data.

In addition to test scores, SEDA provides estimates of socio-demographic and economic characteristics of school districts drawn from the American Community Survey (ACS). We use these ACS covariates in some of our analysis and robustness checks. The construction of these covariates and how they are mapped to school districts is described in greater detail in Fahle et al (2021).

Figure A1: Distribution of library shocks in the school district sample

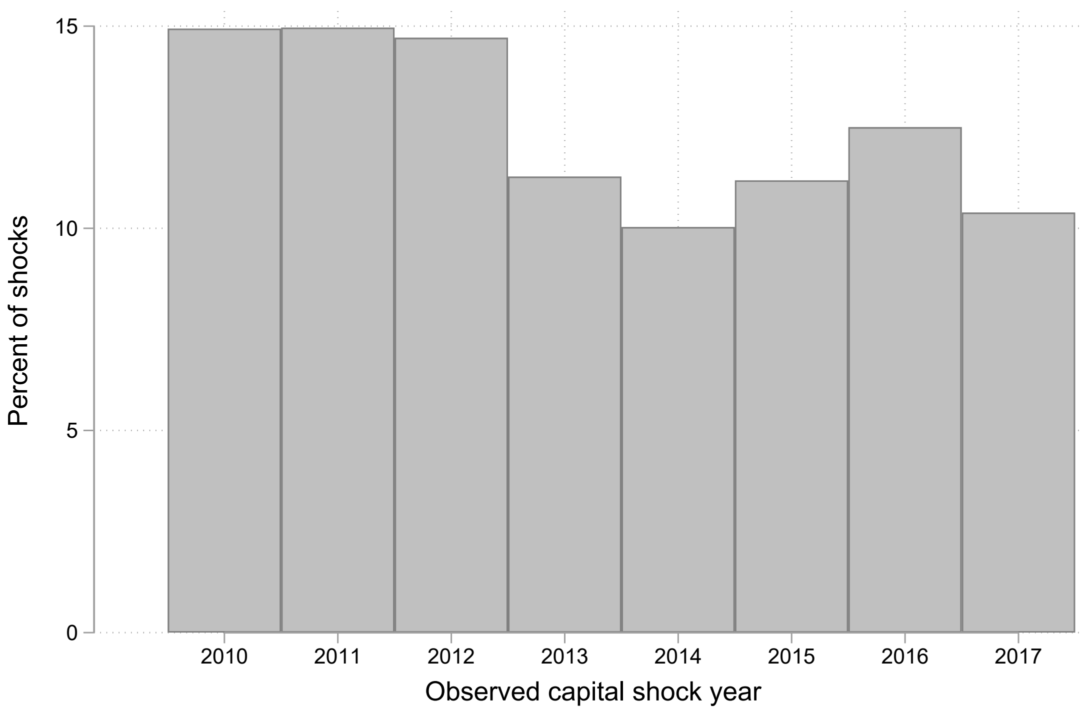

This figure shows the distribution of capital shock years in the school district sample for the sub-sample of districts with observed shocks. Approximately 7 percent of all districts have a shock between 2010 and 2017.

Figure A2: Impact of capital expenditure shock on library use, TWFE

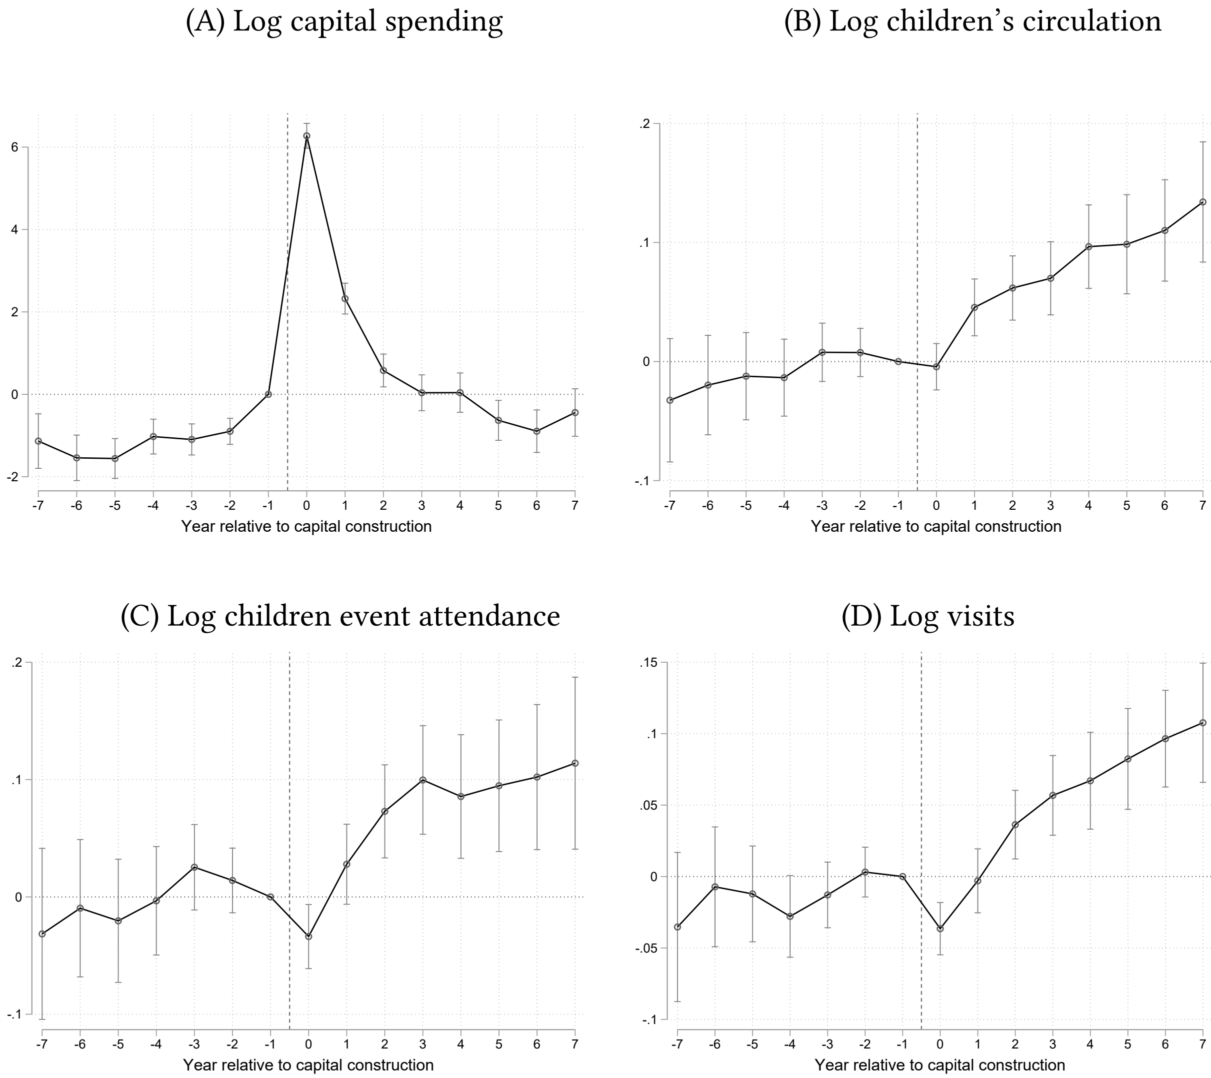

Event study estimates generated using the estimation procedure described in Section III with a standard two-way fixed effect approach instead of the Gardner approach used in our main figures. Outcomes are aggregated to the school district-level. The outcome variables are (A) Log capital spending, (B) log children's circulation, (D) log children's event attendance, and (D) log visits. All figures show 95 percent confidence intervals that account for within-school district clustering. Results are conditional on state-year fixed effects.

Figure A3: Impact of capital expenditure shock on library resources, TWFE

(A) Log book stock

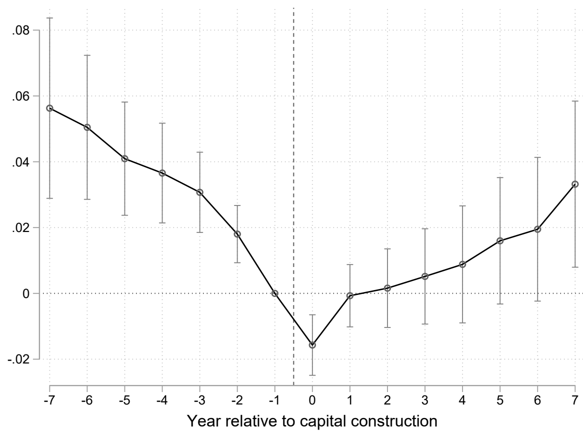

(B) Log number employees

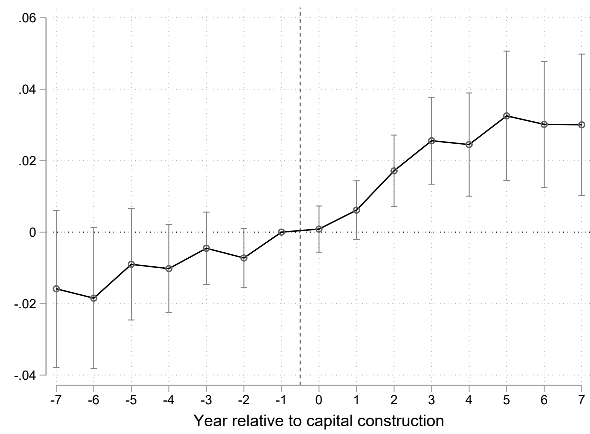

(C) Log salary spending

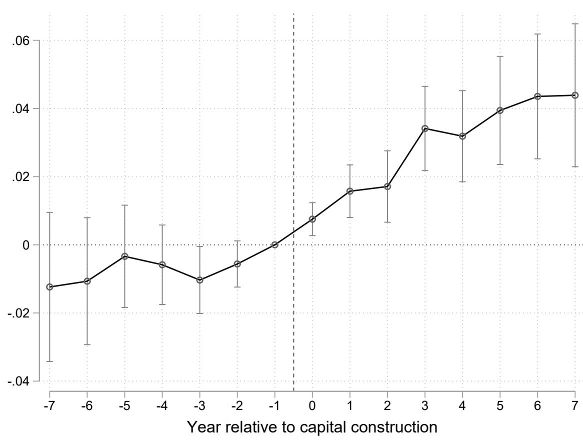

(D) Log operating spending

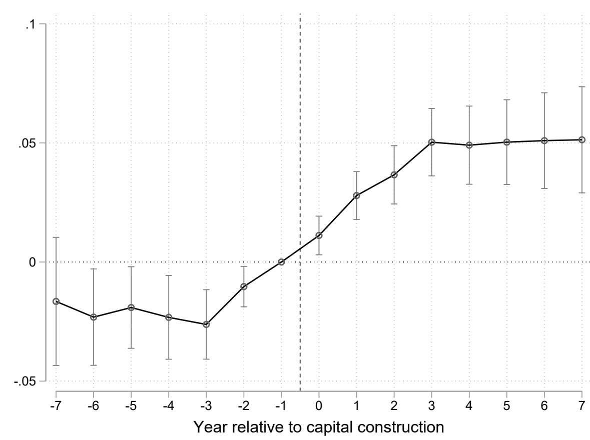

Event study estimates generated using the estimation procedure described in Section III with a standard two-way fixed effect approach instead of the Gardner approach used in our main figures. Outcomes are aggregated to the school district-level. The outcome variables are (A) log books, (B) log number employees, (C) log expenditures on salaries, and (D) log operating expenditures. All figures show 95 percent confidence intervals that account for within-school district clustering. Results are conditional on state-year fixed effects.

Figure A4: Impact of capital expenditure shock on library resources, intensive margin

(A) Log children event attendance

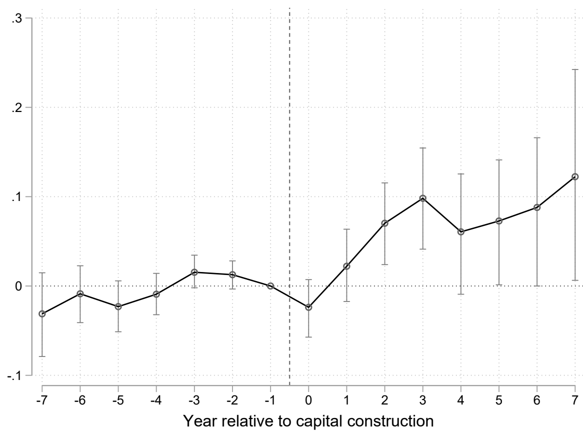

(B) Log visits

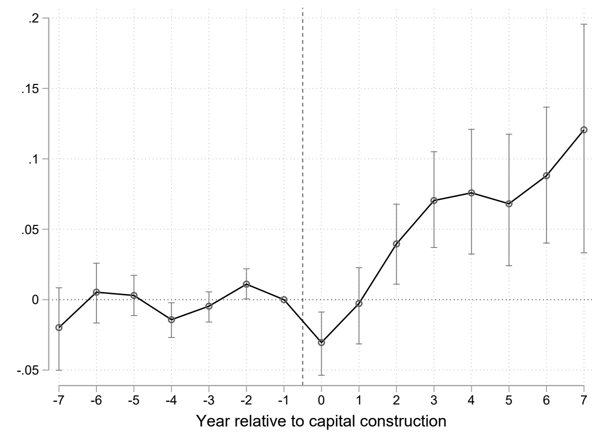

(C) Log operating spending

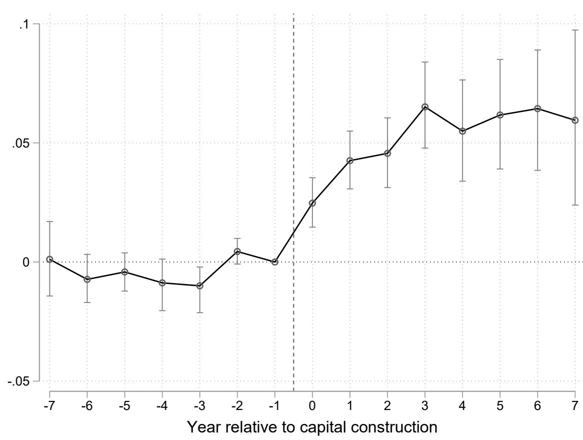

(D) Log children's circulation

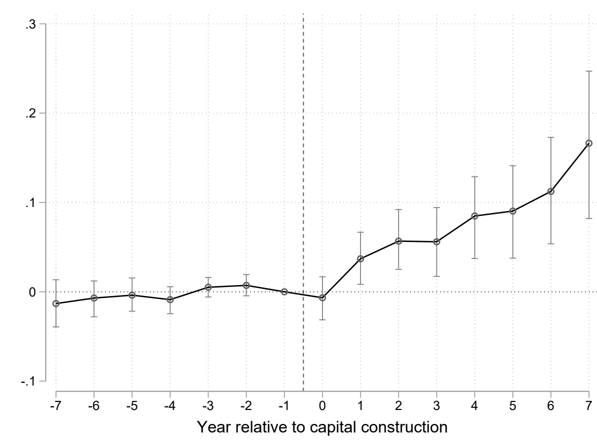

Event study estimates generated using the Gardner estimation procedure described in Section III. Outcomes are aggregated to the school district-level. The outcome variables are (A) log books, (B) log number employees, (C) log expenditures on salaries, and (D) log operating expenditures. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Results are conditional on state-year and district-grade fixed effects. Panel excludes observations where the outcome is zero (in levels).

Figure A5: Impact of capital expenditure shock on per user library use

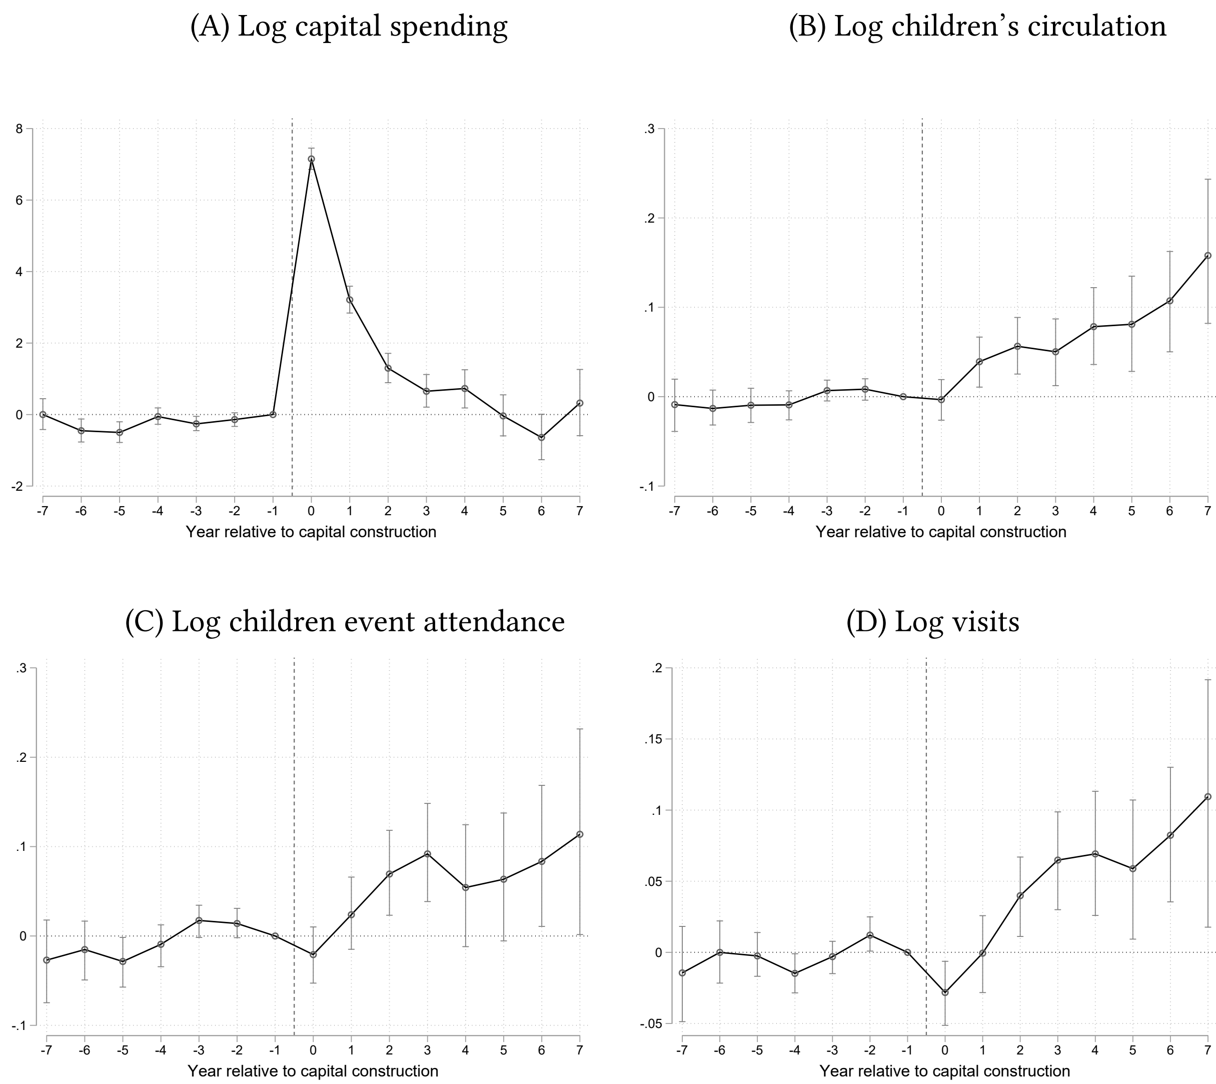

Event study estimates generated using the Gardner estimation procedure described in Section III. The outcome variables are (A) Log capital spending, (B) Log children's circulation, (D) log children's event attendance, and (D) log visits. All outcome variables are expressed as logged per-10,000-user figures. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects.

Figure A6: Impact of capital expenditure shock on per user library resources

(A) Log book stock

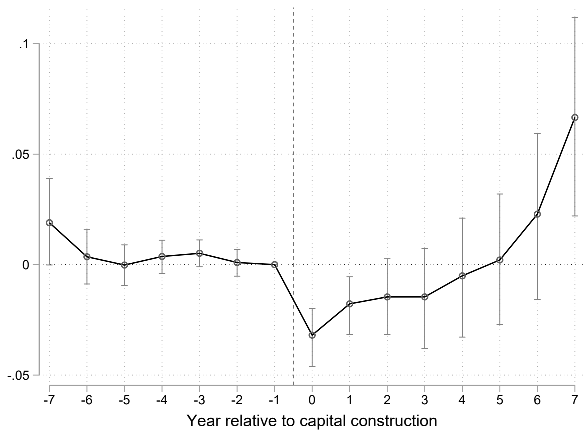

(B) Log number employees

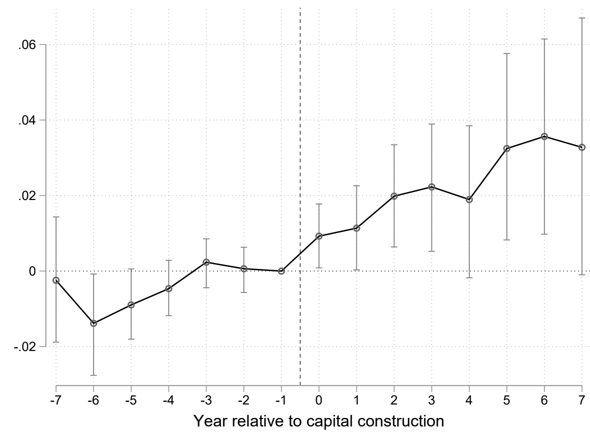

(C) Log salary spending

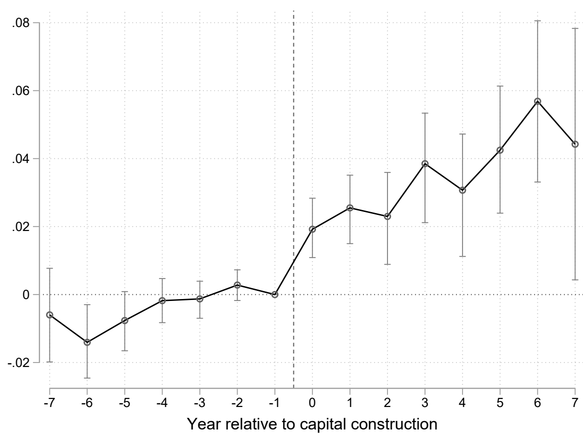

(D) Log operating spending

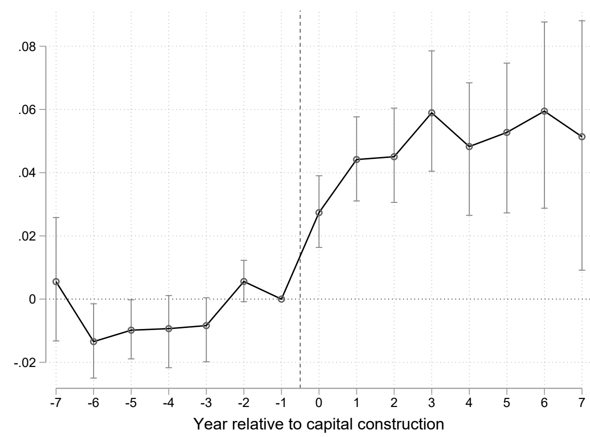

Event study estimates generated using the Gardner estimation procedure described in Section III. The outcome variables are (A) Log books, (B) Log number employees, (C) log expenditures on salaries, and (D) log operating expenditures. All outcome variables are expressed as logged per-10,000-user figures. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects.

Figure A7: Impact of capital expenditure shock on library use, inverse hyperbolic sine transformed outcomes

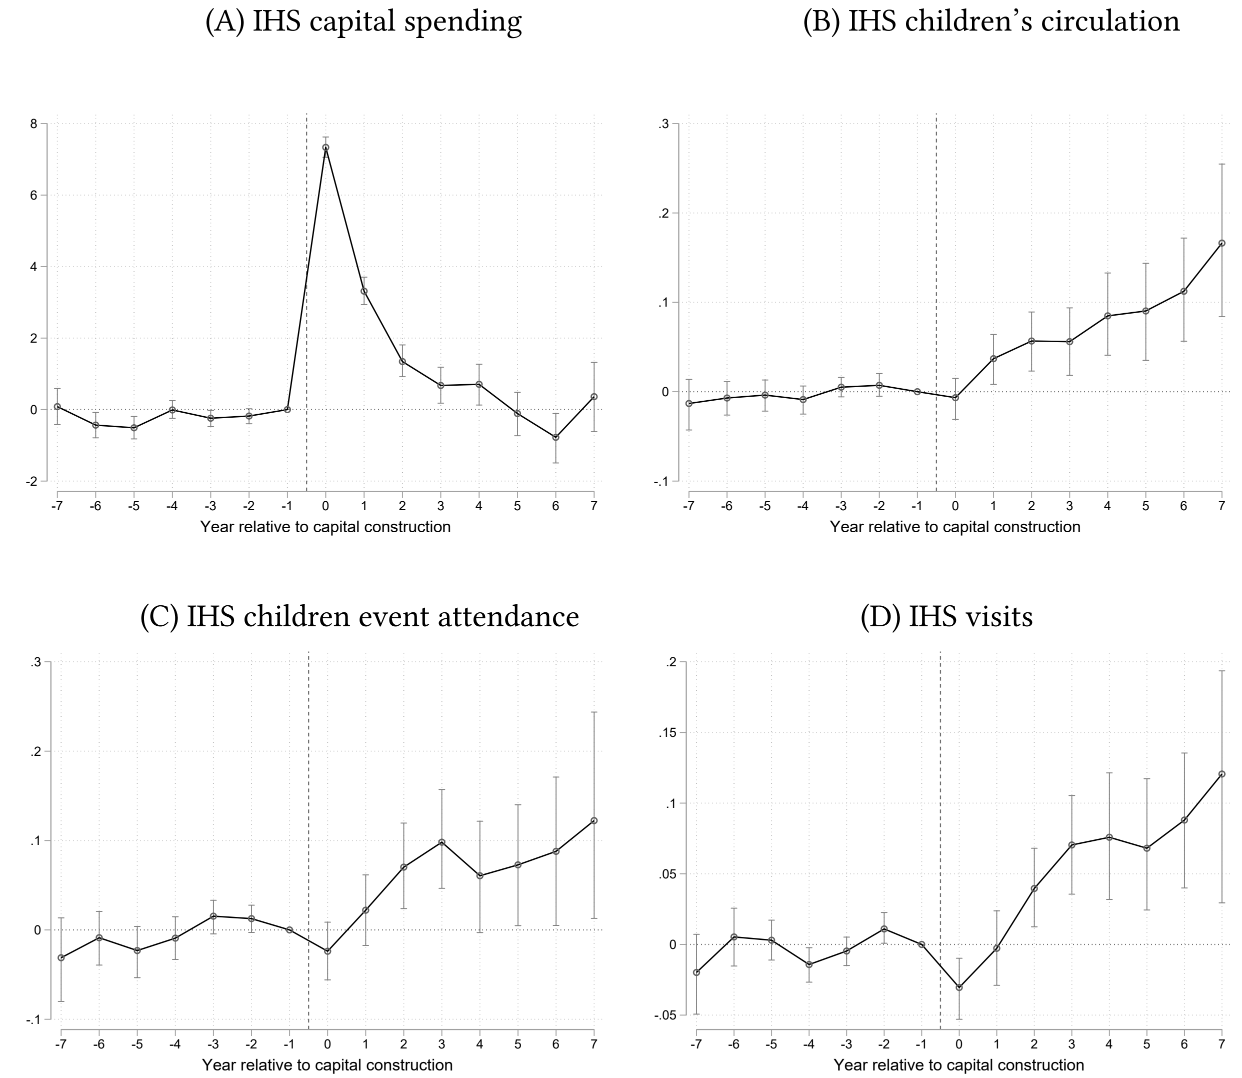

Event study estimates generated using the Gardner estimation procedure described in Section III. All outcome variables are transformed using the inverse hyberbolic sine method. The outcome variables are (A) IHS capital spending, (B) IHS children's circulation, (D) IHS children's event attendance, and (D) IHS visits. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects.

Figure A8: Impact of capital expenditure shock on library resources, inverse hyperbolic sine transformed outcomes

(A) IHS book stock

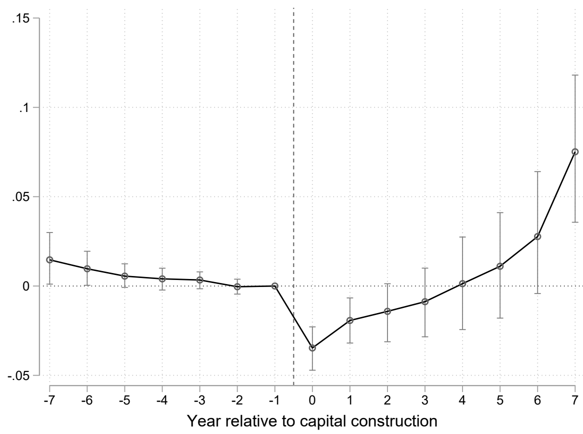

(B) IHS number employees

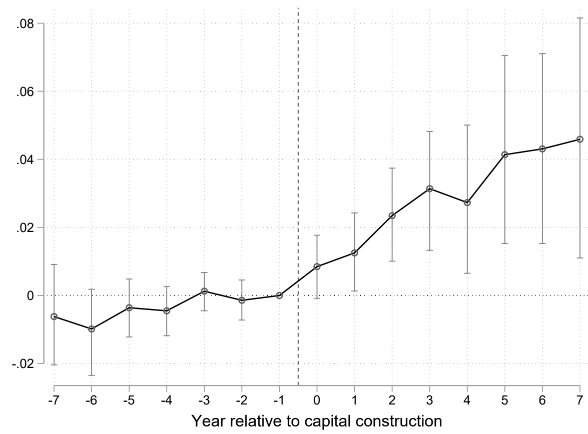

(C) IHS salary spending

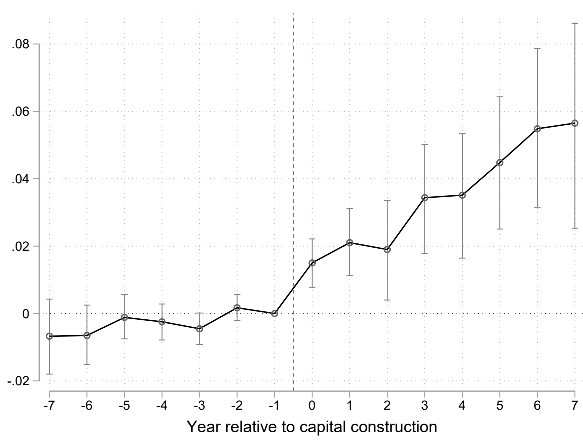

(D) IHS operating spending

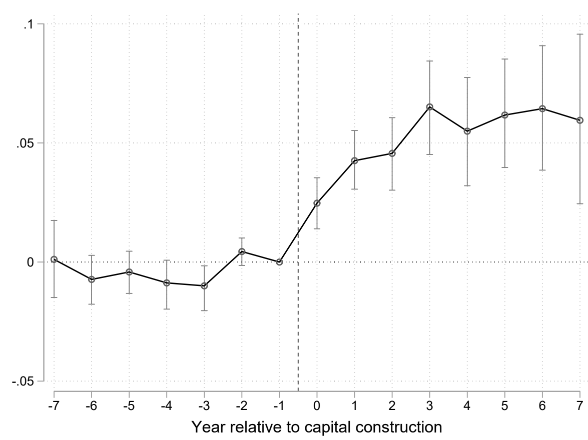

Event study estimates generated using the Gardner estimation procedure described in Section III. All outcome variables are transformed using the inverse hyberbolic sine method. The outcome variables are (A) IHS books, (B) IHS number employees, (C) IHS expenditures on salaries, and (D) IHS operating expenditures. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects.

Figure A9: Impact of capital expenditure shock on library resources, alternative quality variables

(a) Log library media subscriptions

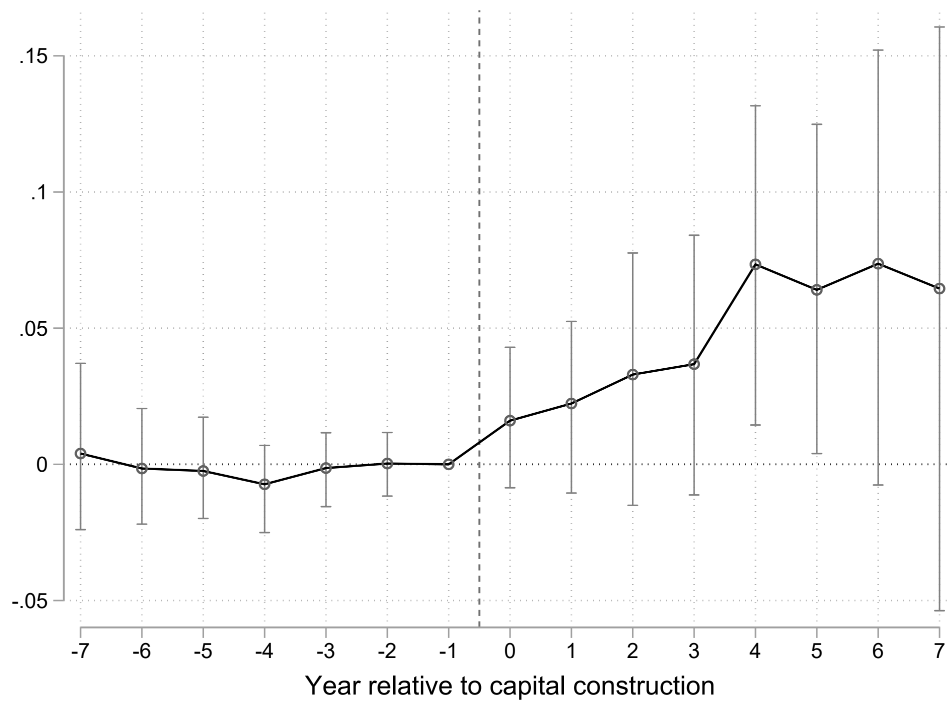

(b) Log spending on collections

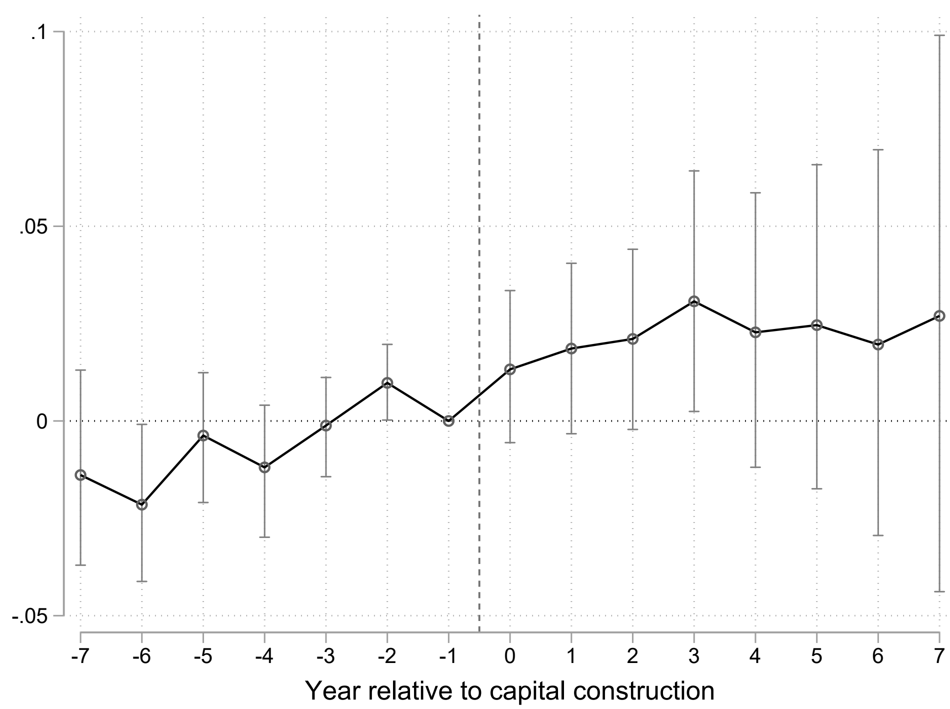

Event study estimates generated using the Gardner estimation procedure described in Section III. The outcome variables are (A) Log media subscriptions and (B) Log expenditures on collections. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects.

Figure A10: Dynamic correlation between capital expenditure shocks and school district characteristics

(A) Share Black students

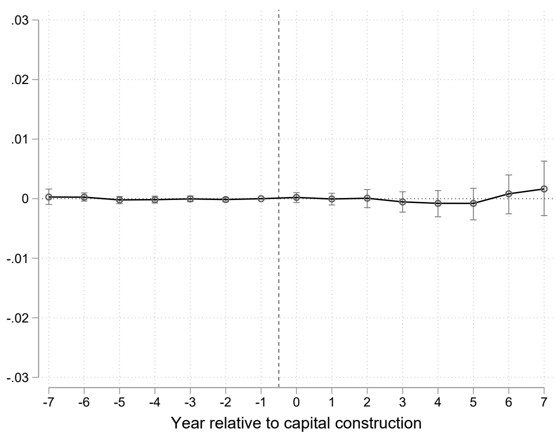

(B) Share Hispanic students

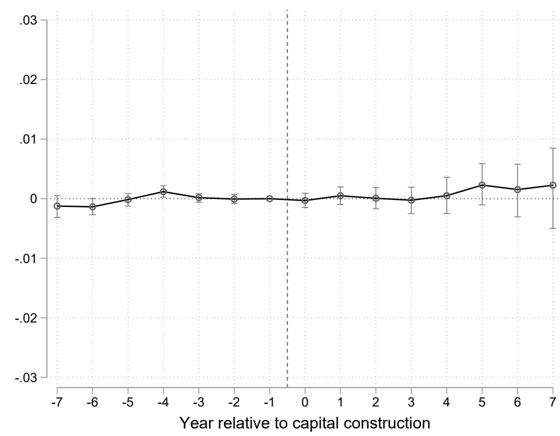

(C) Share Asian students

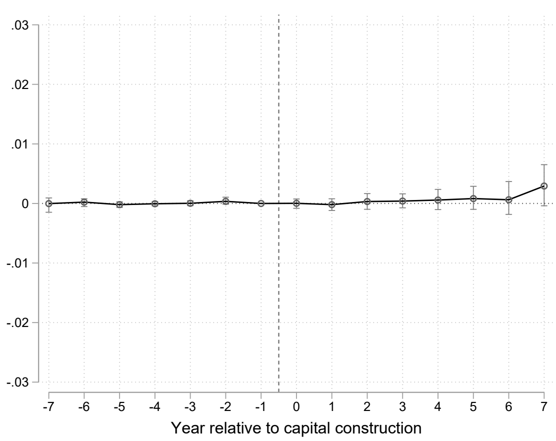

(D) Share Native American students

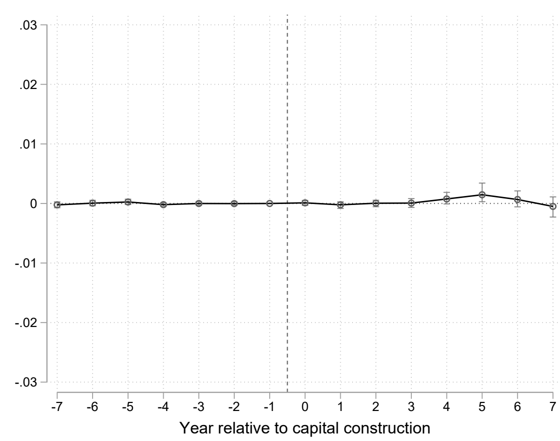

Event study estimates generated using the Gardner estimation procedure described in Section III. The outcome variables are (A) Share Black students, (B) Share Hispanic students, (C) Share Asian students, and (D) Share Native American students. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects.

Figure A11: Dynamic correlation between capital expenditure shocks and additional school district characteristics

(a) Share of students who qualify for free lunch

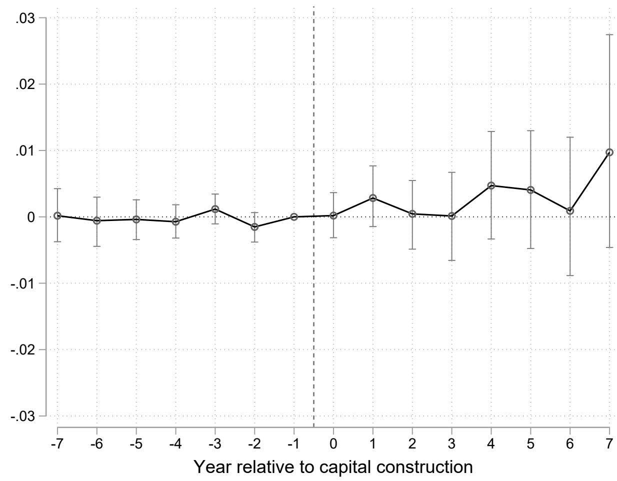

(b) Share of students who qualify for reduced price lunch

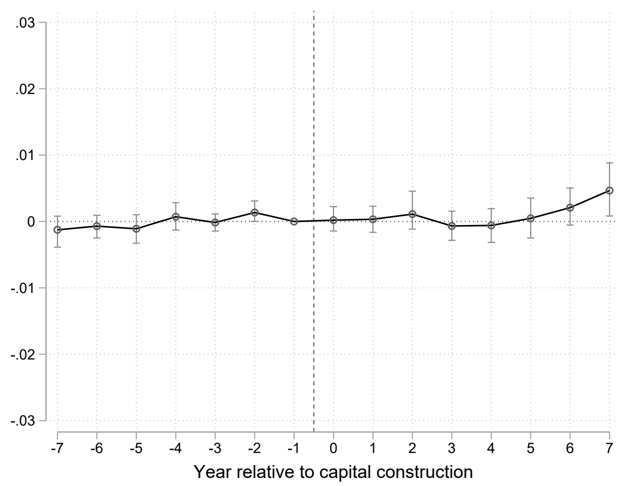

(c) Share of disadvantaged students

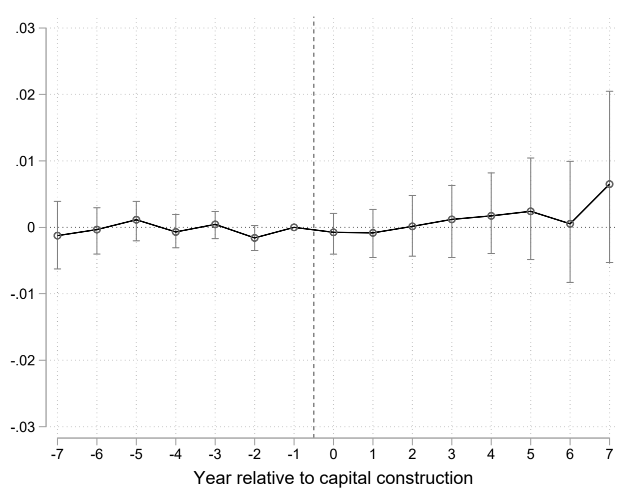

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects.

Figure A12: Dynamic correlation between capital expenditure shocks and community characteristics

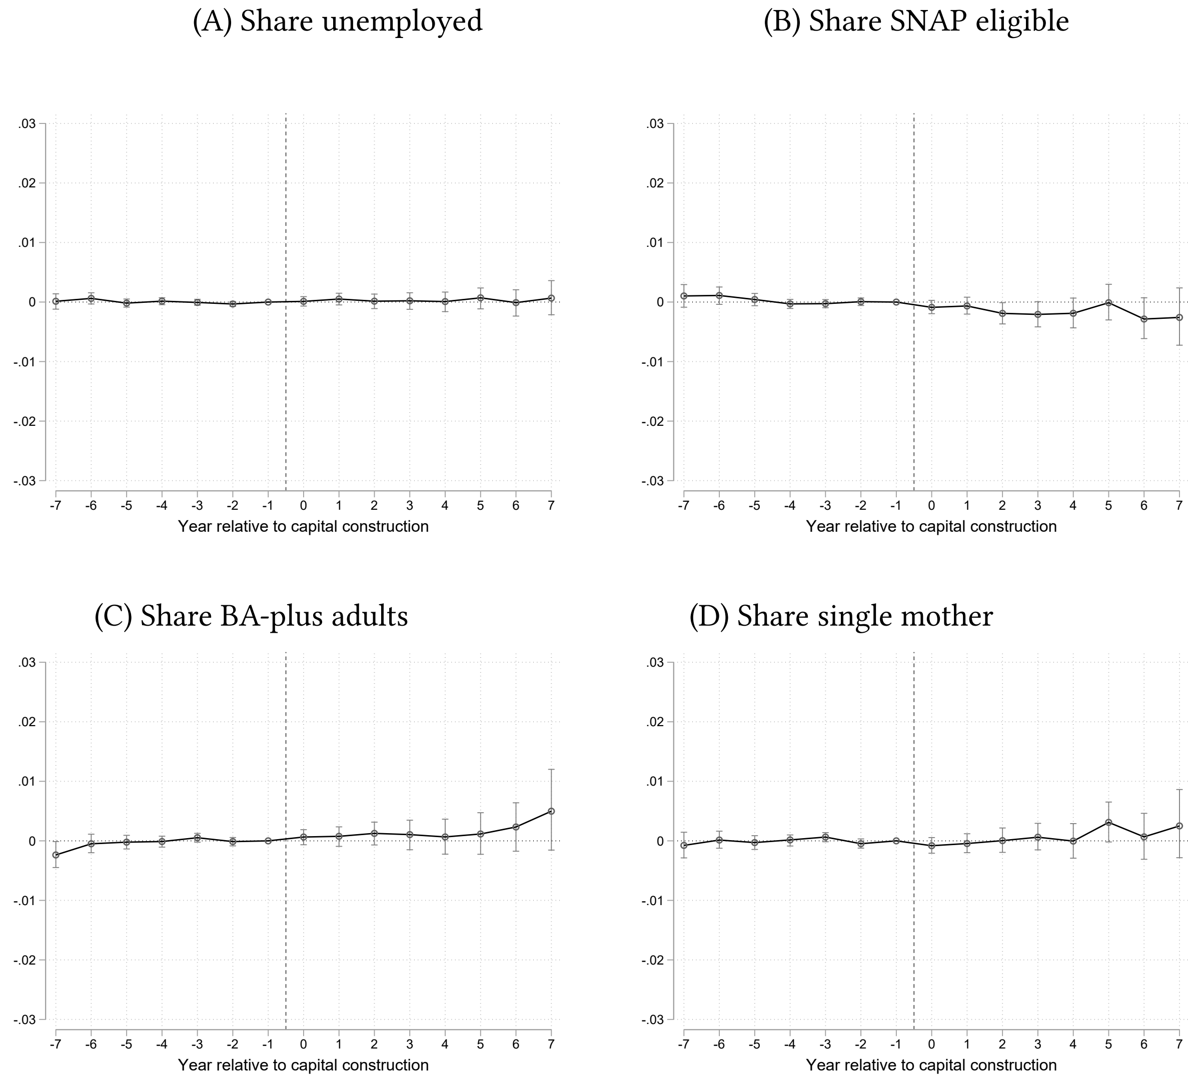

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects.

Figure A13: Dynamic correlation between capital expenditure shocks and school district finances

(A) Log total spending

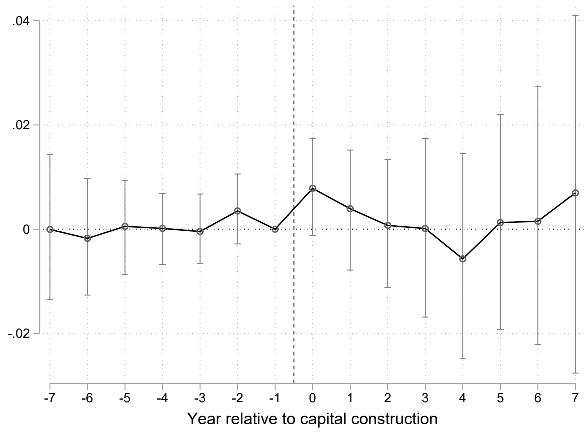

(B) Log capital spending

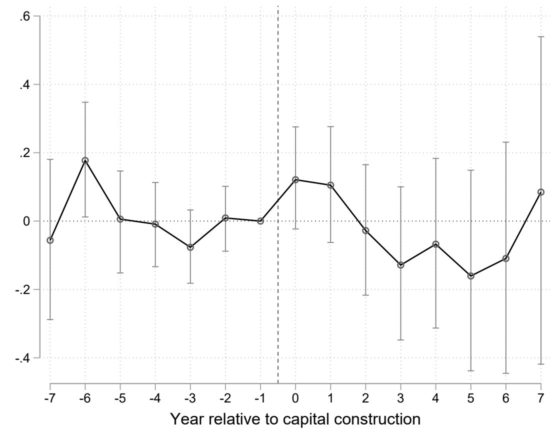

(C) Log salary spending

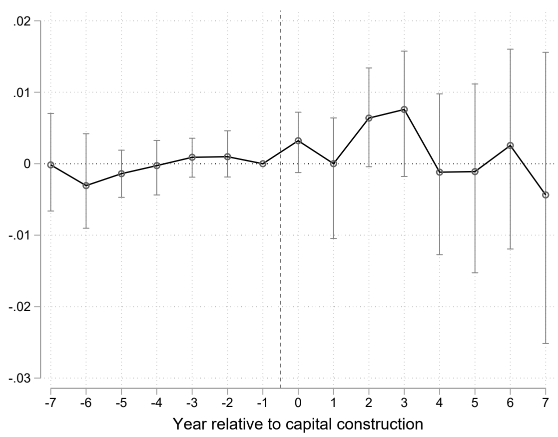

(D) Log instructional salary spending

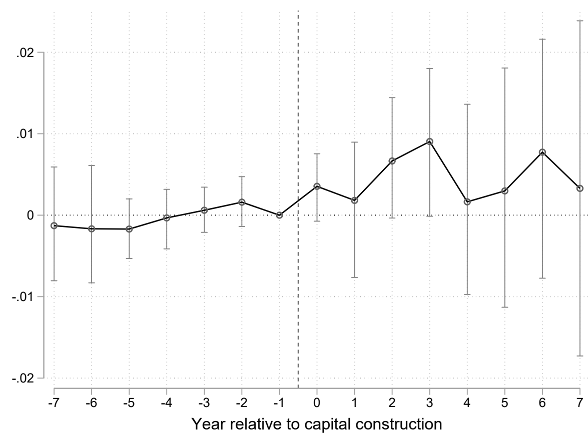

Event study estimates generated using the Gardner estimation procedure described in Section III. The outcome variables are (A) Log total school spending, (B) Log capital school spending, (C) Log salary school spending, and (D) Log instructional salary school spending. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects.

Figure A14: Impact of library capital spending shocks on reading test scores, alternative distance thresholds

(a) 1 mile

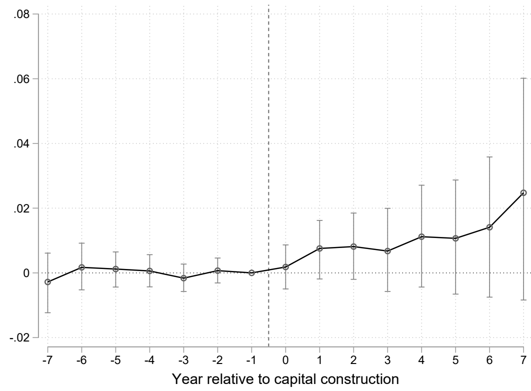

(b) 2 miles

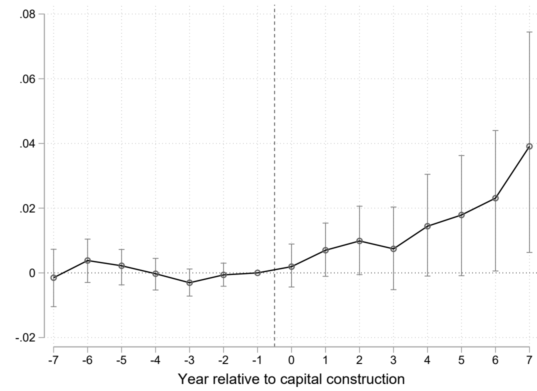

(c) 5 miles

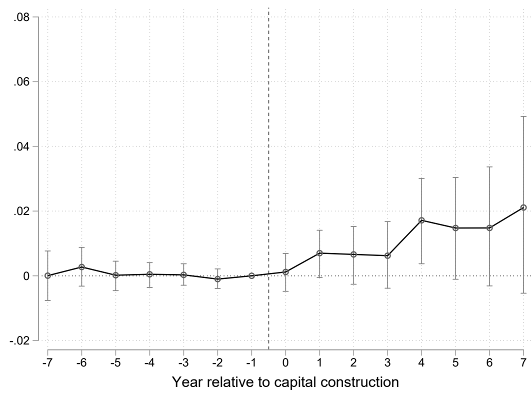

(d) 10 miles

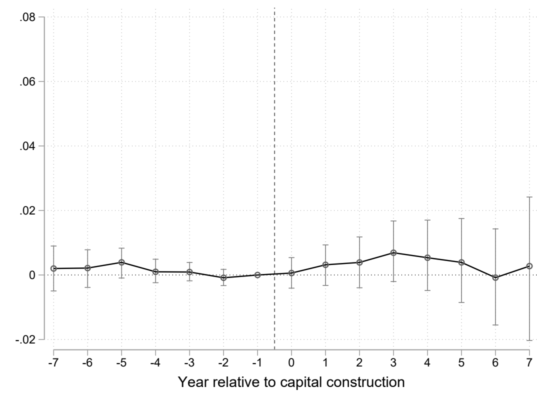

(e) 15 miles

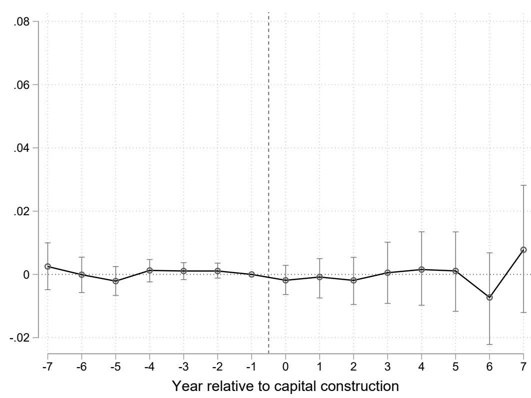

(f) 20 miles

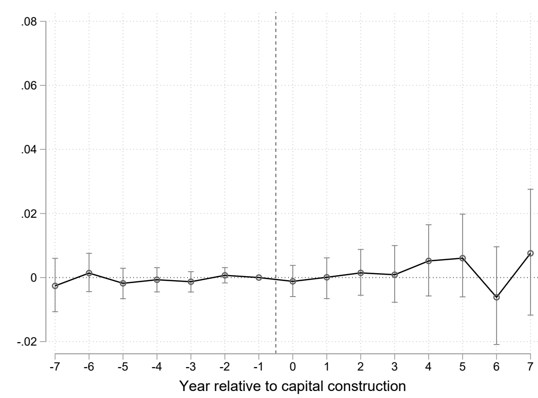

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects. Each figure shows results letting the capital spending events occurs within the indicated distance of a school district.

Figure A15: Impact of library capital spending shocks on reading test scores, alternative capital spending event thresholds

(a) $50 per student

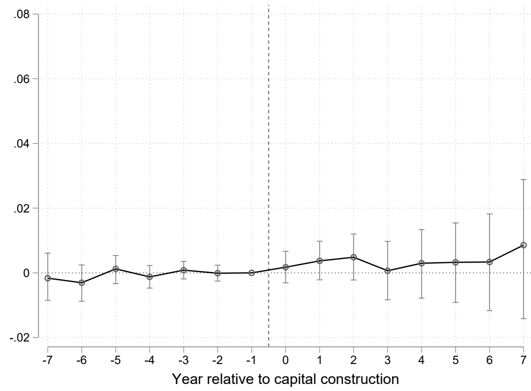

(b) $100 per student

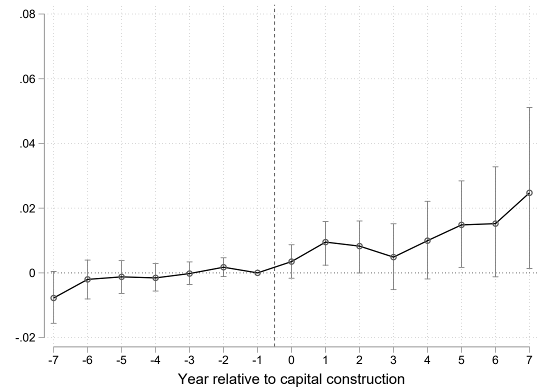

(c) $150 per student

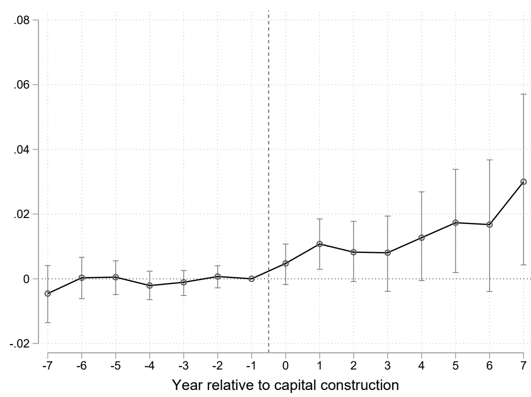

(d) $200 per student (baseline)

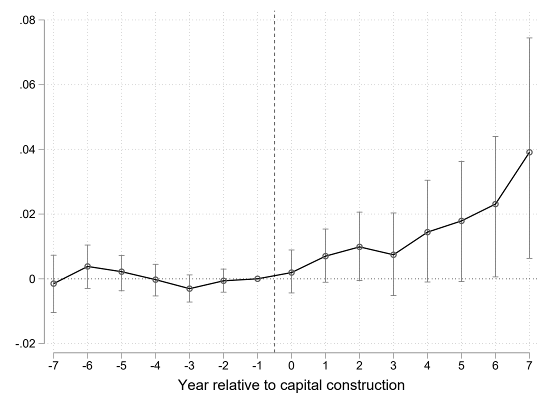

(e) $250 per student

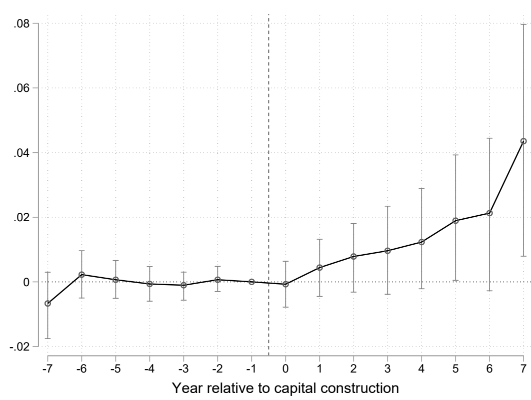

(f) $300 per student

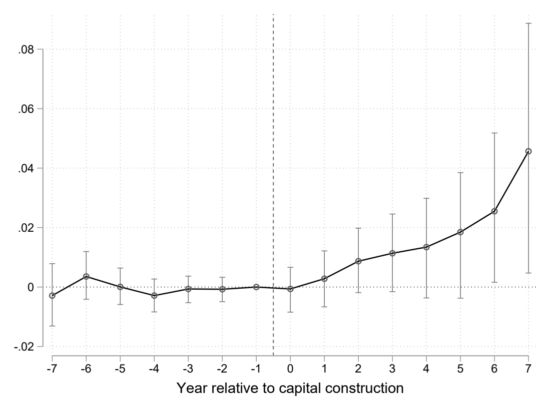

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects. Each figure shows results letting capital spending events be defined by the indicated per student dollar threshold.

Figure A16: Impact of library capital spending shocks on reading test scores, excluding districts with multiple library construction shocks

(a) Baseline

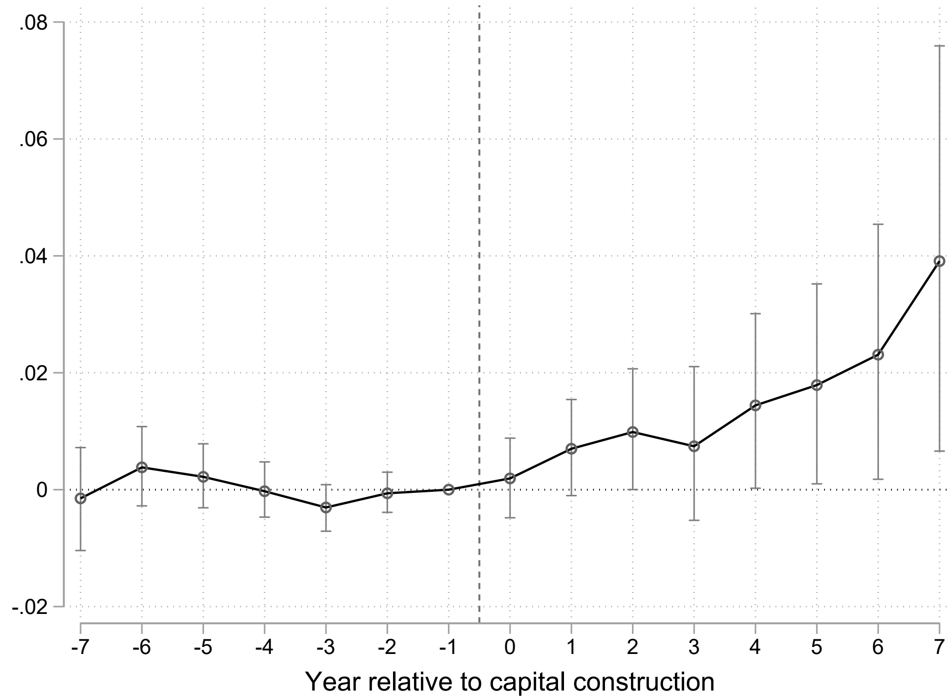

(b) Exclude multi-shock school districts

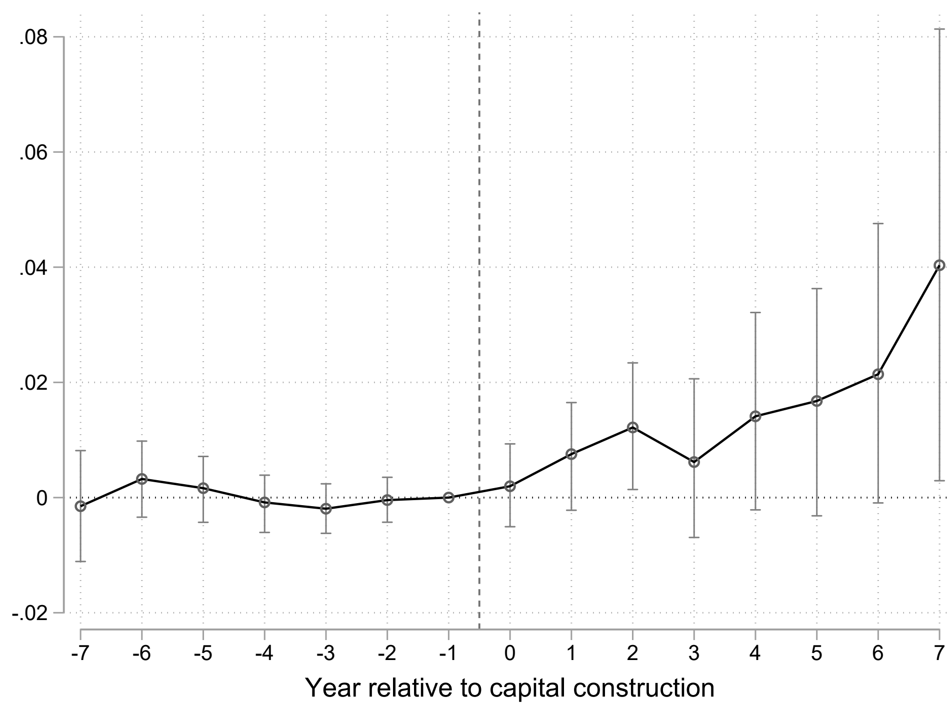

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects. Panel A shows our baseline results. Panel B shows results that exclude school districts that had more than one capital spending event during our sample period.

Figure A17: Impact of library capital spending shocks on reading test scores, including all school districts

(a) Baseline

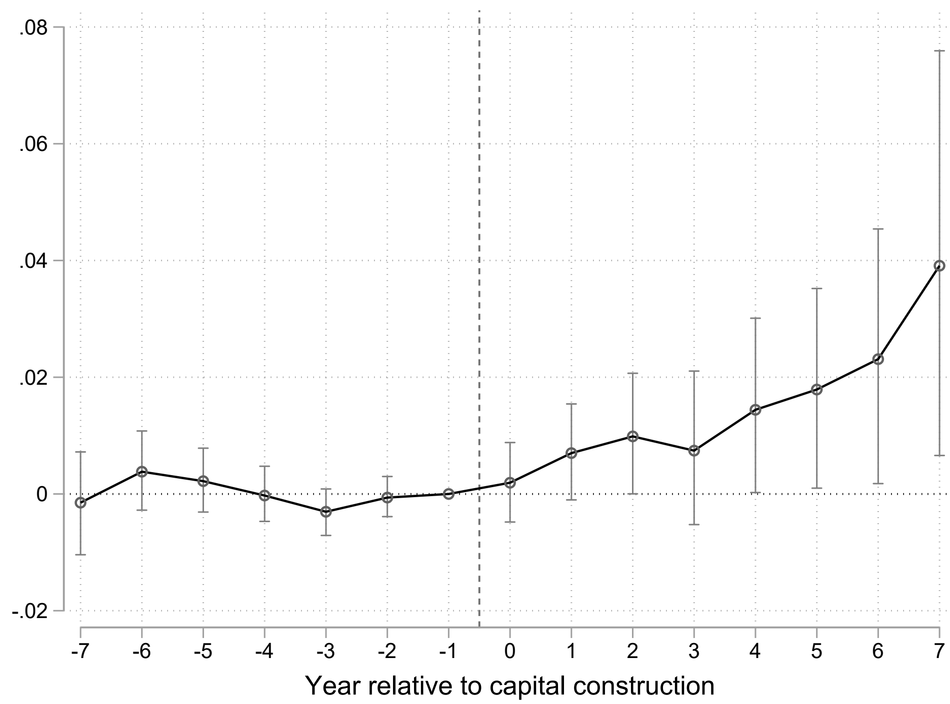

(b) Include many-branch districts

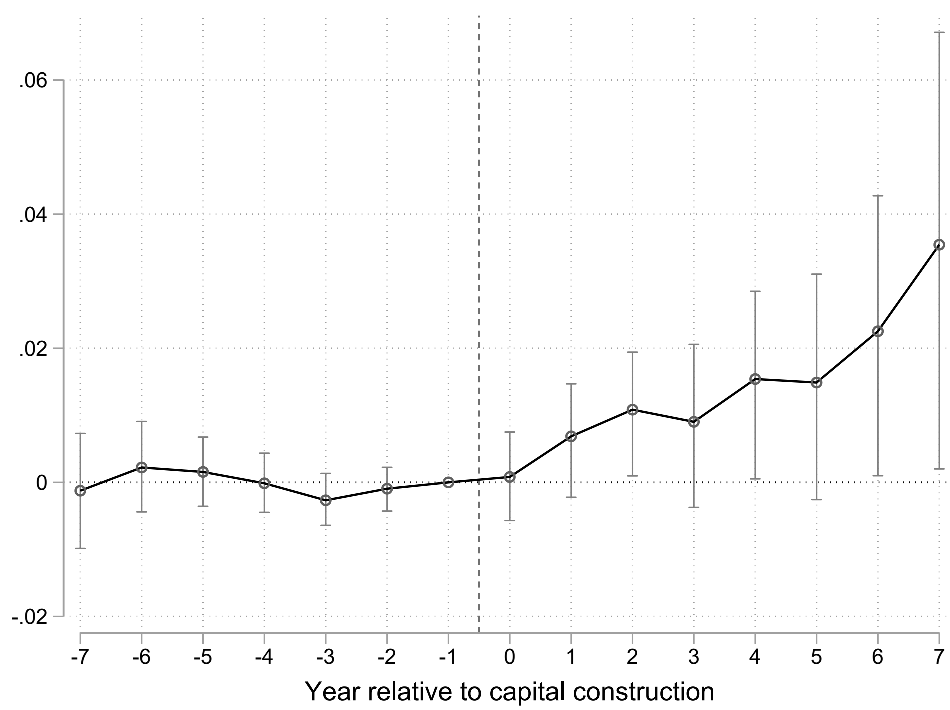

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects. Panel A shows our baseline results. Panel B shows results that do not limit the sample to districts with 5 or fewer nearby library branches.

Figure A18: Impact of library capital spending shocks on reading test scores, weighting by district size

(a) Baseline (no weighting)

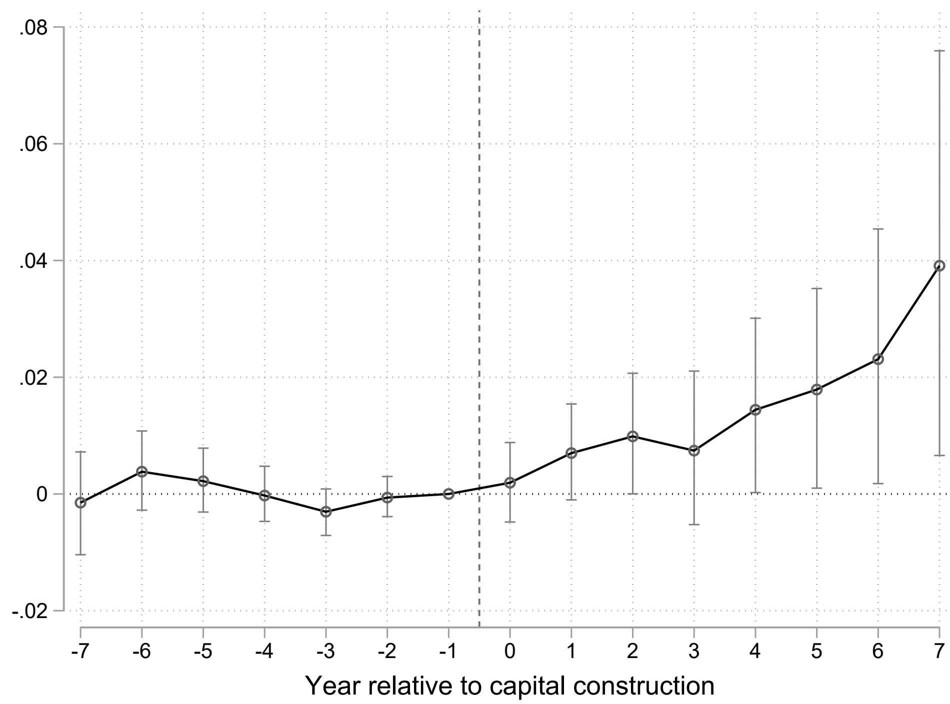

(b) Weighting by district size

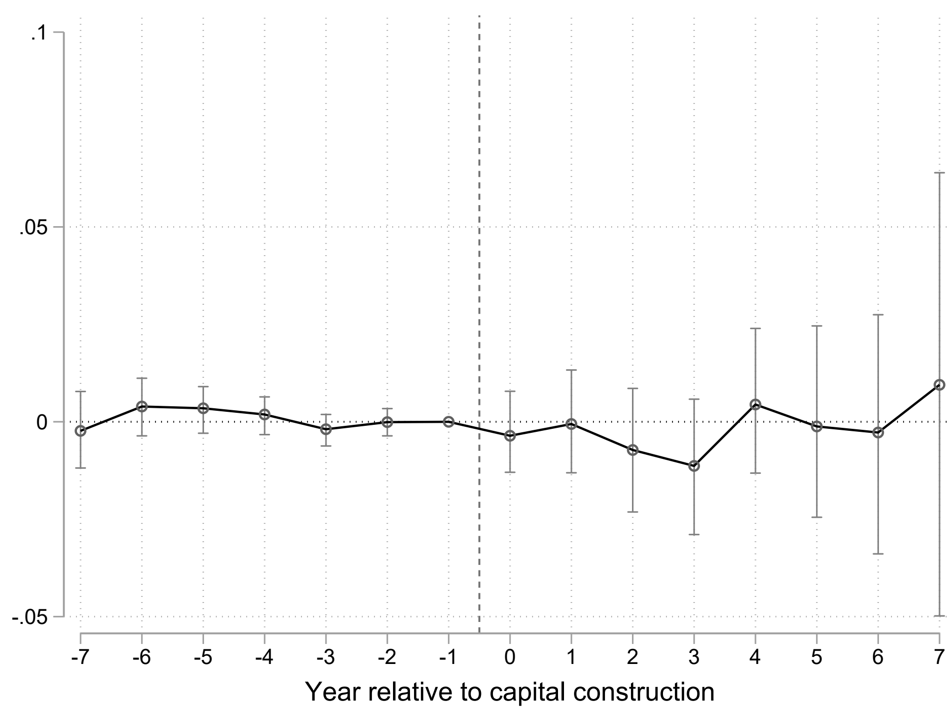

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects. Panel A shows our baseline results. Panel B shows results that weight observations by the number of test-takers in a given grade-district-year.

Figure A19: Impact of library capital spending shocks on reading test scores, extended panel

(a) Incorporating NAEP data

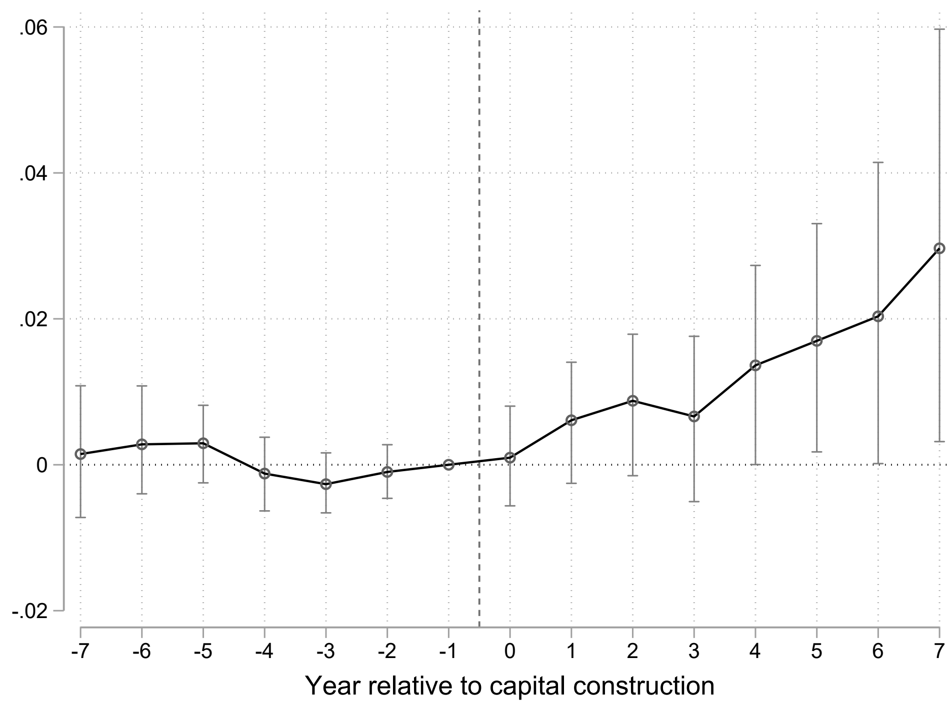

(b) Incorporating NAEP and NLSLSASD data

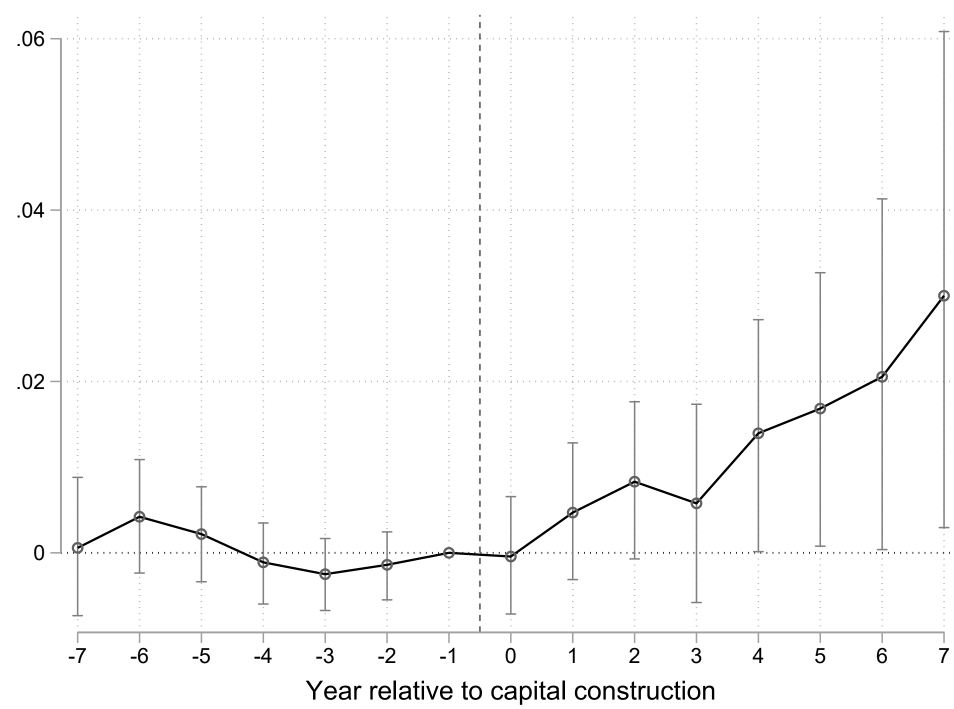

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include state-year and district-grade fixed effects. Panel A shows results that incorporate data from NAEP. Panel B shows results that incorporate data from both the NLSLSASD and NAEP.

Figure A20: Impact of library capital spending shocks on reading test scores, by treatment timing

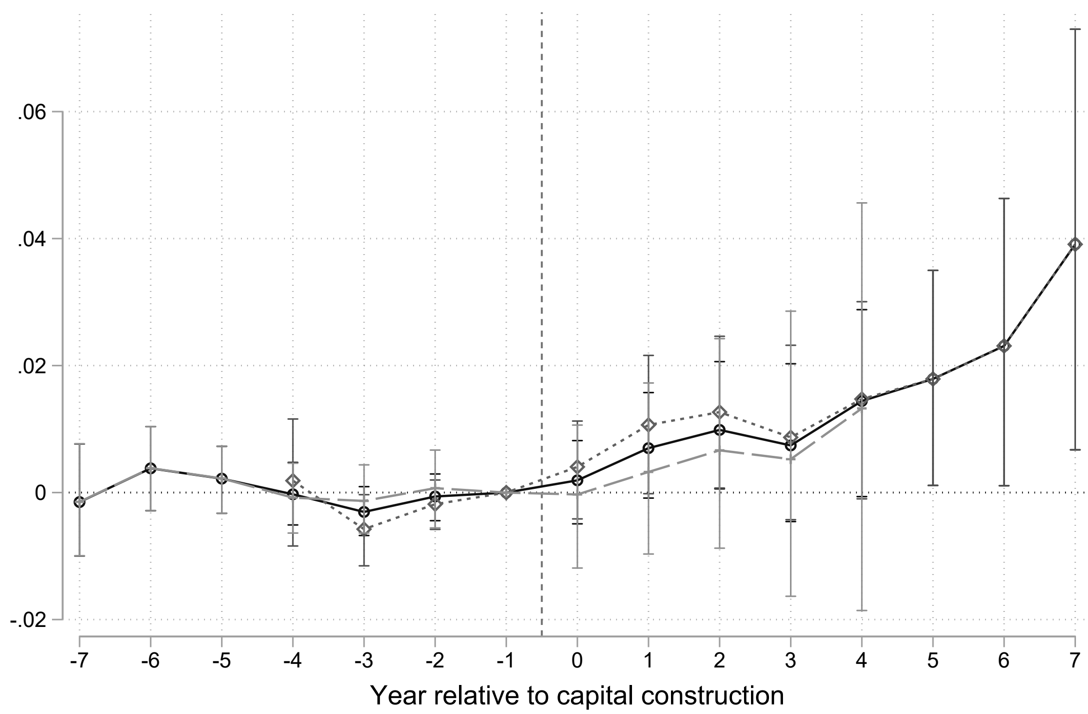

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Results are split by treatment timing: The black line corresponds to all treated units, the dashed line to units treated after or in 2014, and the dotted line to units treated before 2014.

Figure A21: Impact of library capital spending shocks on reading test scores, models without state-year fixed effects

(a) Baseline model

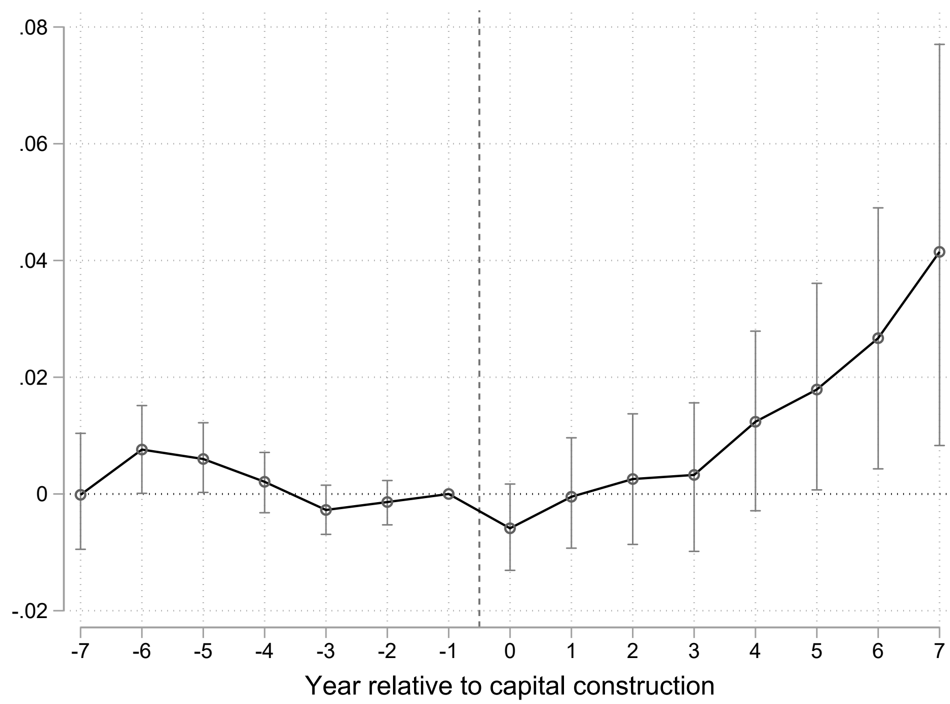

(b) With time-varying controls

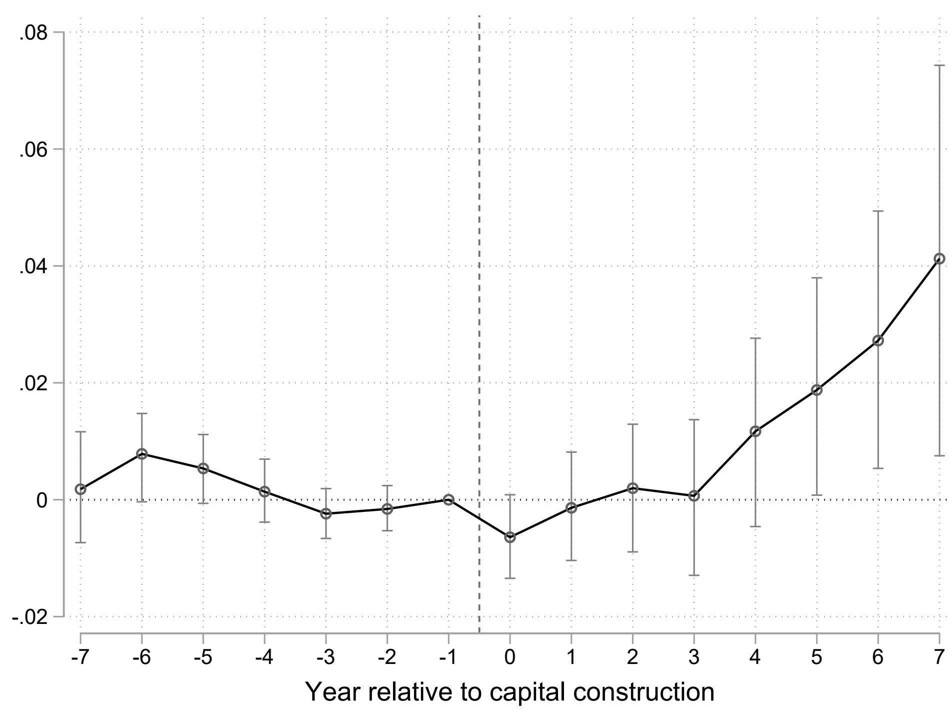

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. Regressions include district-grade fixed effects. Panel A shows results without time-varying controls. Panel B shows results with the set of time-varying controls described in Sections II and III.

Figure A22: Impact of library capital spending shocks on math test scores, with and without state-year fixed effects

(a) Without state-year fixed effects

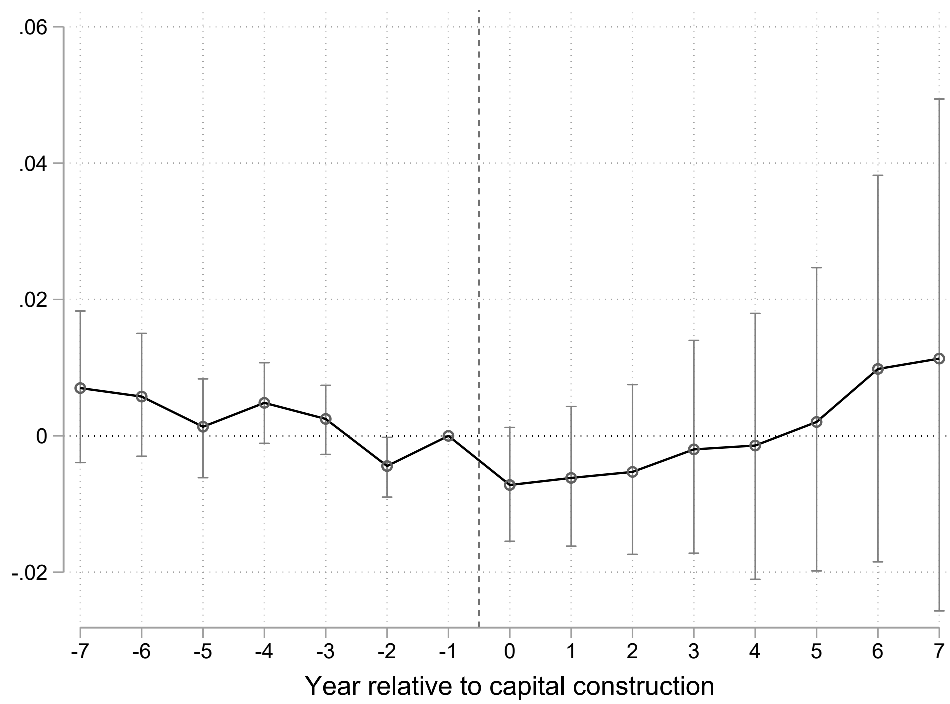

(b) With state-year fixed effects

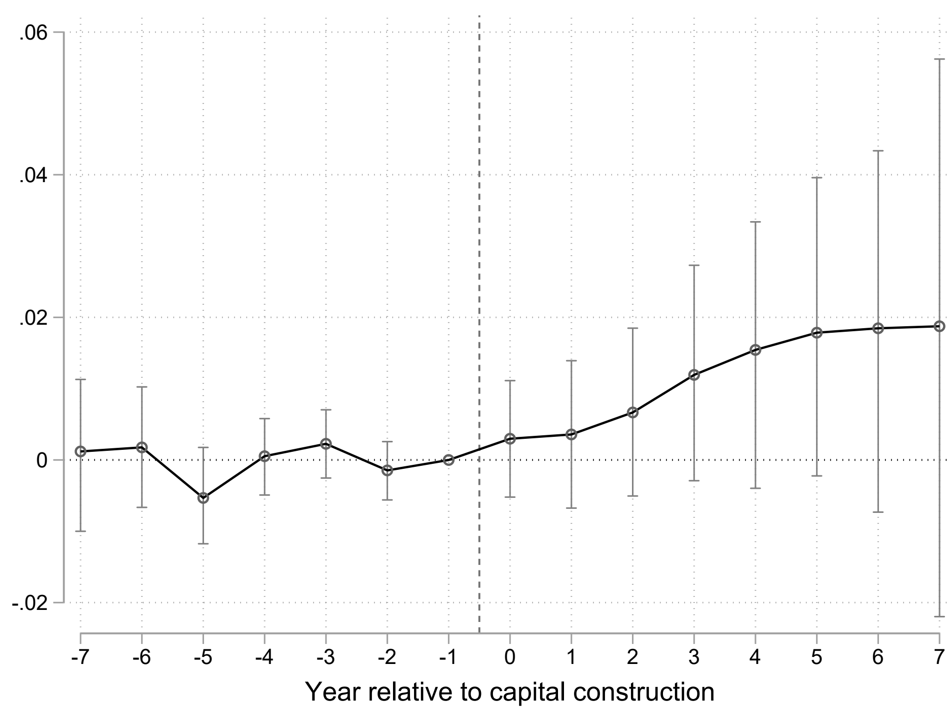

Event study estimates generated using the Gardner estimation procedure described in Section III. All figures show bootstrapped 95 percent confidence intervals that account for within-school district clustering. The outcome variable is mathematics test scores. Regressions include district-grade fixed effects. Panel A shows results without conditioning on state-year fixed effects. Panel B shows results with state-year fixed effects.

Figure A23: Impact of library capital spending shocks on reading test scores, TWFE method

(a) Without time-varying covariates

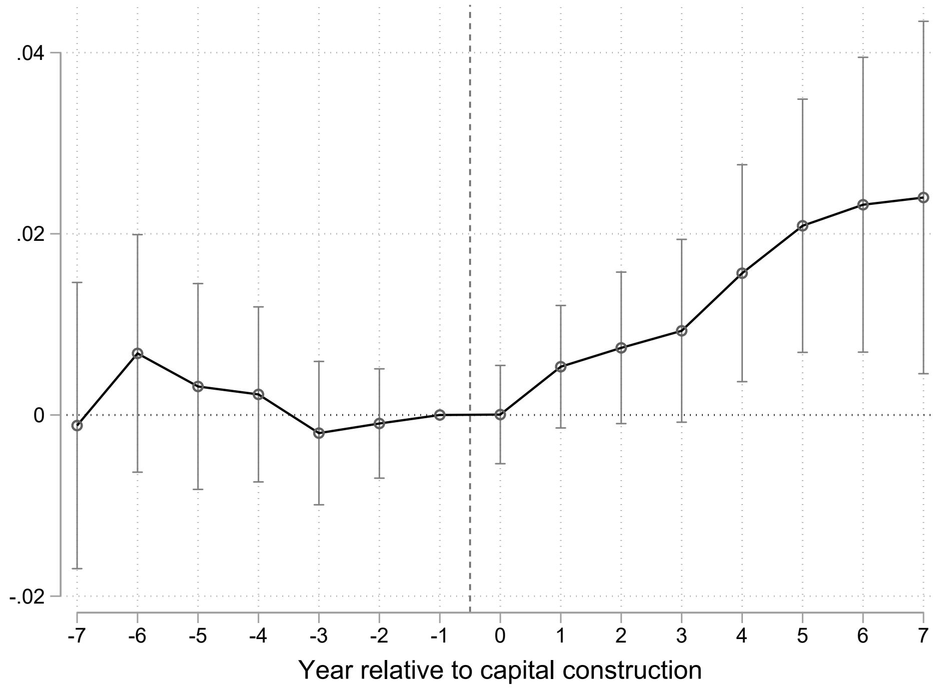

(b) With time-varying covariates

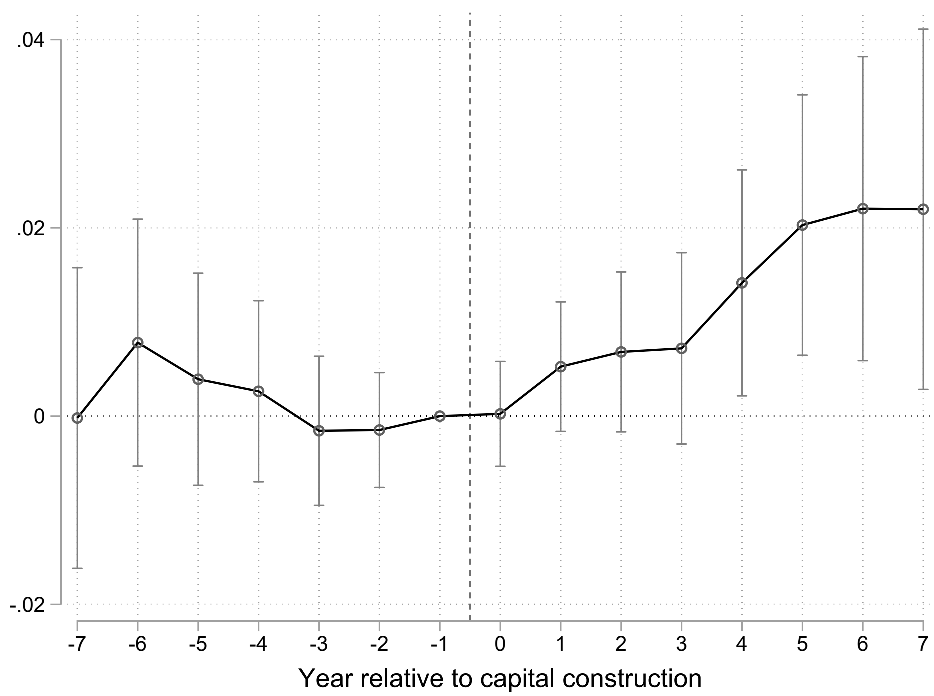

Event study estimates generated using two-way fixed effects estimators, as described in Section III. All figures show 95 percent confidence intervals that account for within-school district clustering. Panel A shows results after conditioning on state-year and district-grade fixed effects. Panel B shows results that additionally control for the time-varying, district-level covariates described in Sections II and III.

Figure A24: Impact of library capital spending shocks on reading test scores, Callaway and Sant'Anna (2021) method

(a) No covariates

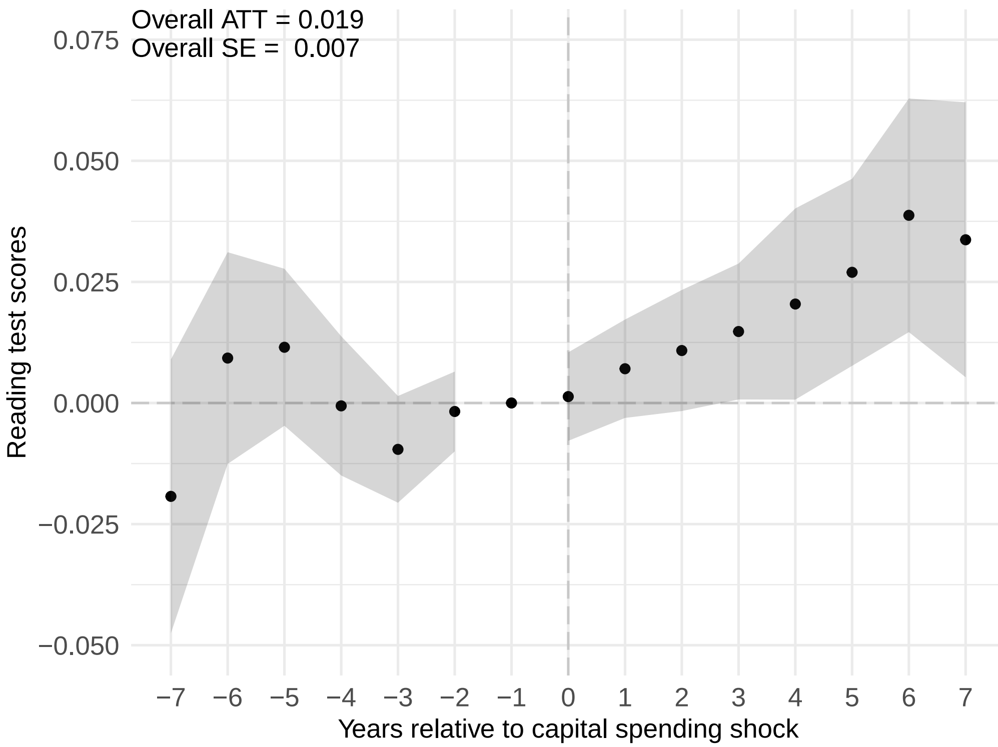

(b) With time-varying covariates

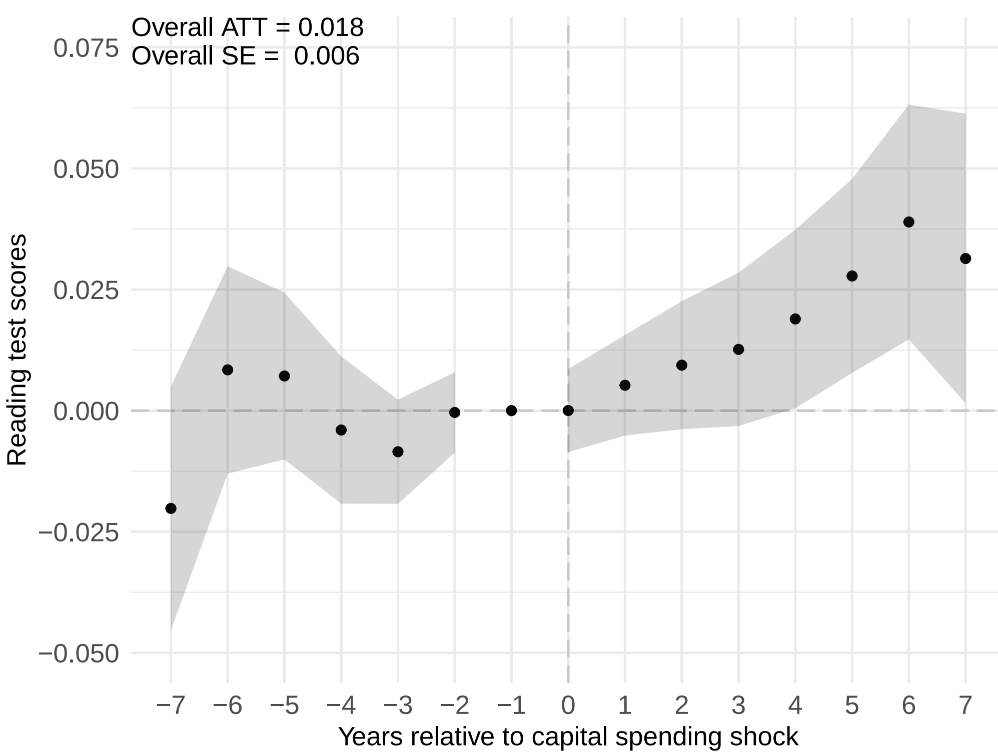

Event study estimates generated using the Callaway and Sant'Anna (2021) method. Standard errors are clustered by school district. Panel A shows baseline results, Panel B shows results that condition on pre-treatment covariates. The list of pre-treatment covariates is described in Sections II and III.

Figure A25: Distribution of log capital spending

(a) Including zero-spending observations

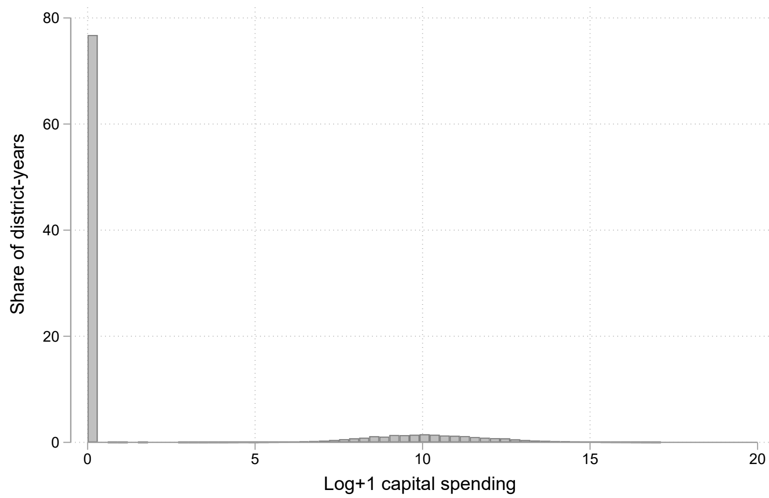

(b) Conditional on any capital spending

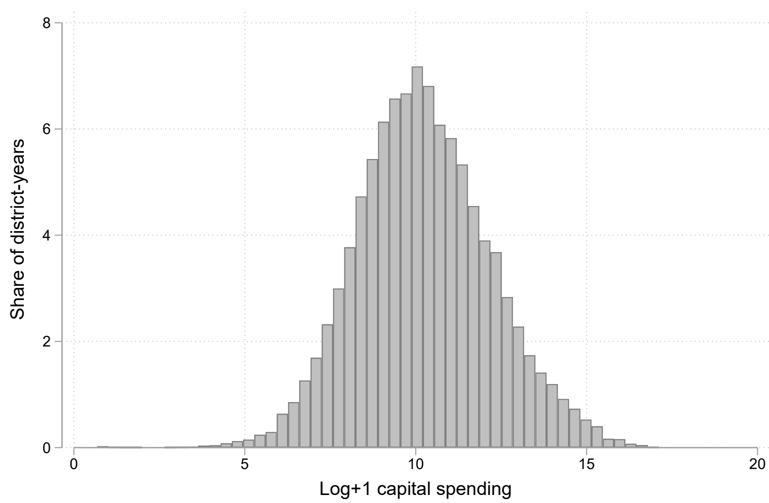

These figures show the distribution of log library capital + 1 in our sample after aggregating to the school district-level. Panel A shows this distribution unconditionally; Panel B shows it conditional on having non-zero capital spending.

Figure A26: Distribution of number of library capital construction shocks, conditional on having at least one

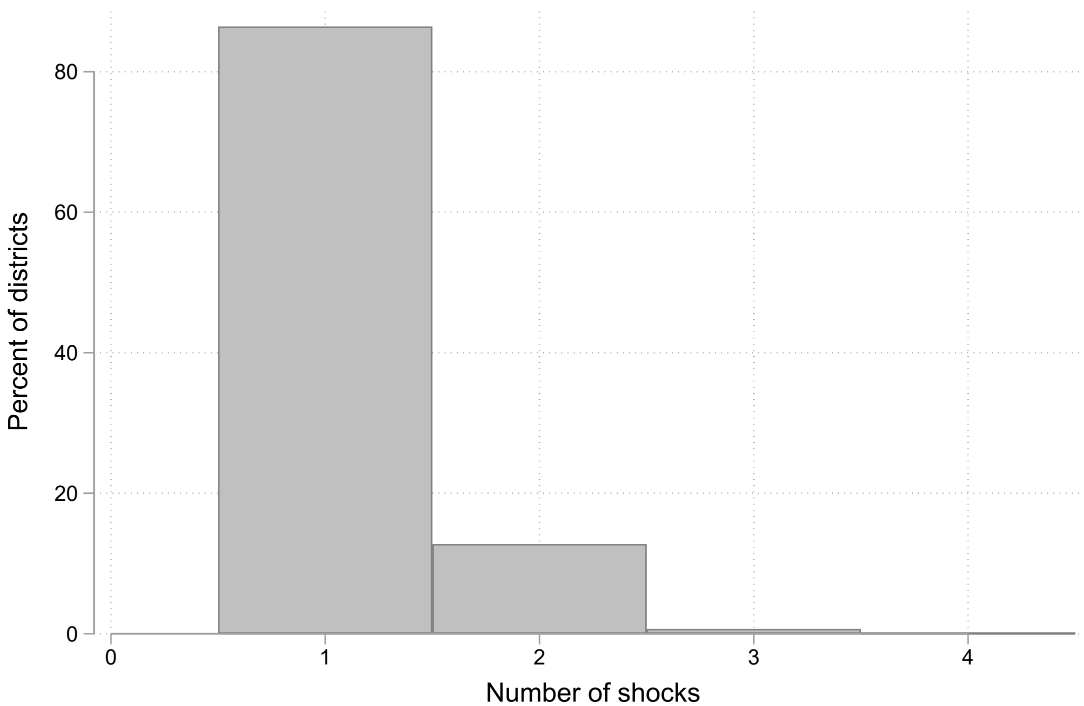

This figure shows the distribution of library capital construction shocks, conditional on having at least one in the school district sample for the sub-sample of districts with observed shocks.