

Borgen, Nicolai T.; Kirkebøen, Lars J.; Kotsadam, Andreas; Raaum, Oddbjørn

Working Paper

Do funds for more teachers improve student outcomes?

Discussion Papers, No. 982

Provided in Cooperation with:

Research Department, Statistics Norway, Oslo

Suggested Citation: Borgen, Nicolai T.; Kirkebøen, Lars J.; Kotsadam, Andreas; Raaum, Oddbjørn (2022) : Do funds for more teachers improve student outcomes?, Discussion Papers, No. 982, Statistics Norway, Research Department, Oslo

This Version is available at:

<https://hdl.handle.net/10419/268058>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Do funds for more teachers improve student outcomes?

TALL

SOM FORTELLER

DISCUSSION PAPERS

982

Nicolai T. Borgen, Lars J. Kirkebøen, Andreas Kotsadam and
Oddbjørn Raaum

*Nicolai T. Borgen, Lars J. Kirkebøen,
Andreas Kotsadam and Oddbjørn Raaum*

Do funds for more teachers improve student outcomes?

Abstract:

We investigate the effects of a large-scale Norwegian reform that provided extra teachers to 166 lower secondary schools with relatively high student-teacher ratios and low average grades. We exploit these two margins using a regression discontinuity setup and find that the reform reduced the student-teacher ratio by around 10% (from a base level of 22 students per teacher), with no crowding out of other school resources or parental support. However, the reform did not improve test scores and longer-term academic outcomes, and we can reject even small positive effects. We do find that the reform improved the school environment from the students' perspective, but with the largest impact on aspects most weakly associated with better academic outcomes.

Keywords: Student-teacher-ratio, class size, test scores, non-cognitive skills, RDD

JEL classification: J24, I2

Acknowledgements: We thank seminar participants at several universities and conferences, Martin Andresen, Colin Green and Astrid Sandsør for useful comments and suggestions. This research is funded by the Norwegian Directorate of Education and Training and the European Research Council (grant number 818425).

Address: Nicolai T. Borgen: University of Oslo; e-mail: n.t.borgen@isp.uio.no
Lars J. Kirkebøen: Statistics Norway, Research department; e-mail: kir@ssb.no
Andreas Kotsadam: Ragnar Frisch Centre for Economic Research;
e-mail: andreas.kotsadam@frisch.uio.no
Oddbjørn Raaum: Ragnar Frisch Centre for Economic Research;
e-mail: oddbjorn.raaum@frisch.uio.no

Discussion Papers

comprise research papers intended for international journals or books. A preprint of a Discussion Paper may be longer and more elaborate than a standard journal article, as it may include intermediate calculations and background material etc.

The Discussion Papers series presents results from ongoing research projects and other research and analysis by SSB staff. The views and conclusions in this document are those of the authors

Sammendrag

I denne artikkelen studerer vi effekter av ekstra lærere til 166 ungdomsskoler i skoleårene 2013/2014 til 2016/2017. Skolene var valgt ut fordi de i 2011/2012 hadde lave gjennomsnittskarakterer (grunnskolepoeng under snittet) og mange elever per lærer (relativ stor gruppestørrelse 2, dvs. forholdet mellom elevtimer og lærertimer i ordinær undervisning, altså unntatt bl.a. spesialundervisning). Vi bruker disse grensene for ekstra lærere til å undersøke hvordan ekstra finansiering påvirket ressursinnsats og resultater.

Vi finner at gruppestørrelse 2 ble redusert med omtrent 10 prosent (fra omtrent 22 til 20 elevtimer per lærertime). Dette er i tråd med hva vi skulle forvente utfra hvor mange nye lærerårsverk skolene fikk. De nye årsverkene ble fylt av kvalifiserte lærere, og brukt til ordinær undervisning. Andre mål på ressursinnsats ved skolene (f.eks. spesialundervisning, særskilt språkopplæring og bruk av assistenter) ble ikke påvirket. Det var heller ingen effekter på lærernes sykefravær eller foreldrenes involvering i skolearbeidet (målt ved Elevundersøkelsen).

Vi finner ingen effekter på resultater på Nasjonale prøver, skriftlig eksamen eller tidlige mål på fullføring av videregående opplæring, og vi kan langt på vei utelukke selv små positive effekter. Vi finner at imidlertid at de ekstra lærerne forbedret skolemiljøet som målt ved Elevundersøkelsen, men med størst innvirkning på aspekter ved skolemiljøet som er svakest knyttet til faglige resultater.

1 Introduction

There is a long-standing debate as to whether providing more financial resources to schools in rich countries improves student outcomes or not. Two decades ago, Hanushek (2003) argued that international evidence showed that resources do not affect performance, while Krueger (2003) concluded that resources are systematically related to student achievement. Building on recent studies with credible identification strategies, Jackson (2020) and Jackson and Mackevicius (2021) show that, on average, increased general spending improves educational outcomes. However, their survey, which is restricted to US studies, shows a wide range of individual effect estimates. Furthermore, spending type matters and some studies suggest that extra teachers may have less impact than other kinds of operational spending.

The effects of reductions of class size, or the student-teacher ratio¹, are challenging to identify because class size is potentially endogenous with respect to student characteristics (Lazear, 2001), and more teachers per student may influence decisions by parents, schools, and education authorities regarding other inputs (Todd and Wolpin, 2003). Since large RCTs such as project STAR (Student-Teacher Achievement Ratio) are rare, the literature is dominated by natural experiments such as rule-induced class size, close elections, or other exogenous sources of differential funding of school districts. Not only have these studies provided mixed findings (see below), but the variation in student-teacher ratios used for identification is typically far from what is relevant for large-scale policy. To our knowledge, there is no direct evidence based on policy changes that explicitly provide funds to hire more teachers and whether such funding improves students' learning.²

We provide direct evidence of the short- and long-term effects of extra school funding earmarked for hiring more teachers. We investigate the effects of a large Norwegian intervention that provided a four-year funding of 600 extra teachers to 166 lower secondary schools with a student-teacher ratio higher than 20 and a student grade point average below the national mean. We exploit these two margins in a regression discontinuity (RD) setup and find that the reform reduced the student-teacher ratio by around 10% (from 22 to 20), without crowding out other school resources and reducing teacher qualifications and parental support. However, the funding did not raise test scores after one year, externally graded exam scores after three years, or later completion of upper secondary education. We do find that the increased number of teachers improved the school environment from the students' perspective, but with the largest positive effects on aspects of the environment that are relatively weakly associated with better academic outcomes.

The intervention provides an excellent opportunity to identify the effects of funding for extra teachers. First, the extra funding policy provided a substantial increase in the number of teachers and was implemented on a large scale. The policy increased the annual costs per student by USD 1 400 over four years in the targeted schools, from an already high level of spending.³ Thus, although the policy was economically and politically feasible, it may constitute an upper bound for likely large-scale spending increases when spending is already high.

Second, the policy was implemented in a manner that permits credible identification. The

¹While much of the empirical literature focuses on class size, extra teachers may also be used, for example, as teacher aides in the classroom or for specific interventions (Andersen et al., 2018; Kirkeb en et al., 2021).

²However, there are evaluations of specific in-school interventions, for example showing positive effects of intense tutoring (Guryan et al., 2021; Fryer Jr, 2017; Bonesr nning et al., 2021).

³OECD (2021) reports that per-pupil spending on lower secondary education in Norway is USD 15 410 (PPP adjusted), which is similar to the US and higher than the OECD average of USD 11 091.

two sharp margins of the eligibility criteria enable us to use a RD framework, which can credibly address concerns about omitted variable bias. We study placebo dimensions to test our identifying assumptions empirically.

Third, the policy effects can be investigated using rich register data, allowing us to study both short- and long-term academic outcomes, with minimal attrition and non-response problems. The total size of the intervention (166 schools over four years, with a total of about 15,000 students per cohort) gives us sufficient precision to rule out even small effects. The data also enable us to examine compensatory adjustments by parents or schools induced by the policy, check whether the policy impacted recruitment of qualified teachers. Moreover, by supplementing register data with survey data, we can also study the effects on students' school environment.

Our main contribution is the identification of a policy-relevant parameter using marginal educational investments. While natural experiments credibly identify class size effects, they often study variation in student-teacher ratios at margins less relevant for marginal educational investments. In addition, class size reductions triggered by maximum class size rules could be subject to various input substitutions, such as providing fewer teacher hours or less qualified teachers in smaller classes. From a policy perspective, such endogenous inputs are an intrinsic part of the policy effect (Todd and Wolpin, 2003) and should clearly not be held constant (Leuven et al., 2008). However, the specific input substitutions (or absence thereof) may be closely related to the specific class size margins that are studied. Moreover, the effects of a given class size reduction may differ depending on whether the initial class size is 40, 30, or 20. Thus, studies exploiting class size cutoffs indicate the effects of maximum class size rather than of marginal educational investments. In contrast, we contribute to this literature by providing evidence on effects of increasing funds to hire more teachers.

We add to the literature on the effects of school resources on educational outcomes. Recent US studies using school financing reforms or changes in the components of the school financing formula show positive effects (Jackson, 2020). Jackson and Mackevicius (2021) conduct a meta-analysis of 31 US studies with credible identification and conclude that school spending affects test scores and, to an even greater extent, longer-run educational attainment. Outside of the US, there are fewer studies on the effects of non-targeted resources with credible research designs. One exception is Hægeland et al. (2012), which find that adding resources stemming from hydropower revenues to schooling in Norway positively affected learning outcomes. Another exception is Gibbons et al. (2018), who use boundary discontinuities in England and find that higher spending leads to better educational attainment. Our estimates add to the small number of studies credibly investigating the effects of spending in countries other than the US.

Our estimated effects of extra funding are small compared to those summarized in Jackson and Mackevicius (2021). Three years of exposure to an annual per-student increase of USD 1,400 approximately corresponds to the four-year annual USD 1,000 standardized treatment of Jackson and Mackevicius (2021). After three years of treatment, we find an insignificant effect of 0.5% of a standard deviation on exam scores. This estimate is considerably smaller than the corresponding weighted average of 3.5% in Jackson and Mackevicius (2021). Although we cannot reject this 3.5% estimate in our study, our point estimate is only one-third of the 25th percentile in the distribution of effect sizes from the 31 studies included in Jackson and Mackevicius (2021). We are not yet able to observe completion of upper secondary school, but our estimate of the effect of

completing the second year of upper secondary school can be compared to the effect on high school completion of 1.9 percentage points in Jackson and Mackevicius (2021). At -0.3 percentage point, our point estimate is negative (although not significant), unlike all estimated effects on completion in Jackson and Mackevicius (2021). Furthermore, with a confidence interval of (-2.2, 1.6) we can reject equality with the weighted average of 1.9.

Most of the studies reported in Jackson and Mackevicius (2021) use general spending changes, but four recent studies were able to investigate what type of spending matters or in which context the effects are largest. Baron (2021) cleverly exploits elections in Wisconsin, where there is a legal requirement to have referenda before changing school spending. In particular, there are different referenda for operational and capital expenditures. Using a regression discontinuity framework, he finds that operational expenditures positively affect test scores and post-secondary enrollment, but there are no effects attributable to increased capital expenditures. Abott et al. (2020) also use close elections, but in seven US states where the compliers increased support staff and teacher wages but not the number of teachers. They find that the increased operational spending led to higher test scores and graduation rates. Brunner et al. (2021) use wind energy installations in the US, which increase district incomes. They find that these increased incomes led to higher capital school spending but not much higher current spending. In particular, the increased school funding did not increase teacher wages or reduce class sizes, and they find no effects on academic outcomes. Finally, Brunner et al. (2020) instead exploit district variation in union strength and find larger effects of school spending on students' performance in areas with stronger unions, where the funds were spent on teacher compensation (rather than hiring new teachers, as in districts with weak unions). These studies indicate that the type of spending seems to be important and our study provides more direct evidence of the (null) effect on student learning of hiring more teachers.

Our study also adds to the inconclusive literature on class size effects. In a landmark study, Krueger (1999) investigated the effects of Project STAR, which randomly assigned students to smaller classes from kindergarten through third grade, and found a large improvement in performance due to class size reductions. Later work investigated the longer-term effects of the STAR experiment and found positive effects on college attendance (Chetty et al., 2011). Conversely, Hoxby (2000) exploits "as if random" variation across cohorts in Connecticut schools and can rule out even small effects of class size on performance in math, reading, and writing in fourth and sixth grade. Woessmann and West (2006) also use variation across grades within schools to identify class size effects in 11 countries using TIMSS data. They only find effects in 2 of the 11 countries, Greece and Iceland.⁴

Most studies investigating the effects of class size use rule-induced reductions, where classes need to split if they reach a certain threshold. A classic example of this method is Angrist and Lavy (1999), which has been interpreted as finding positive effects due to smaller classes in Israel, although the results were actually mixed. They found consistent effects for fifth graders, mixed results for fourth graders, and zero effects for a sample of third graders. Angrist et al. (2019) find no effects using later cohorts in a follow-up study. Results from other contexts are also mixed. Positive effects due to class size reductions have been found in Sweden (Fredriksson et al., 2013), Denmark (Browning and Heinesen, 2007), and Bolivia (Urquiola, 2006). In contrast, Argaw and Puhani (2018) find no effect of class size in elementary school on choosing a more academic track

⁴Neither the US nor Norway was included in their sample.

in Germany, and Angrist et al. (2017) find no effects on gains in learning in Italy. Likewise, earlier Norwegian studies found mixed results. An early study by Bonesrønning (2003) found some weak evidence that larger classes in lower secondary school lead to less favorable short-run outcomes, but no subsequent Norwegian study has found the same. Leuven et al. (2008) find no effect on short-run test scores in lower secondary schools. Leuven and Løkken (2020) reject small effects on long-run outcomes from class-size reductions in primary and lower secondary schools in Norway, and Falch et al. (2017) find no long-run effects from size reductions in Norwegian lower secondary schools.

Finally, our study contributes to the burgeoning literature on the marginal effects of resources on non-test score outcomes and the school environment. Test scores do not capture all that students learn in school. This is most clearly evinced by the fact that successful interventions often have even larger effects on longer-run outcomes, such as the probability of attaining higher education, than they do on immediate test scores (e.g., Jackson 2018; Jackson et al. 2020). Therefore, an increasing number of recent studies aim to identify effects on non-test-score outcomes, which could be caused by a wide variety of skills, including what are often labeled non-cognitive skills. Schools can affect non-cognitive skills, which often have long-term impacts on educational and labor market outcomes (Chetty et al., 2011; Cornelissen and Dustmann, 2019; Heckman et al., 2013; Fredriksson et al., 2013). US evidence suggests that schools that improve ninth-grade socioemotional development (interpersonal skills, school connectedness, academic engagement, grit, and academic effort) also increase upper secondary completion and college enrollment. Jackson (2018) finds that teachers affect non-cognitive skills, as proxied by absences, suspensions, course grades, and normative progression in ninth grade. These teacher-induced improvements in non-cognitive skills are better predictors of longer run academic outcomes than teacher-induced changes in test scores. Our study finds intervention effects on the school environment, but none on academic outcomes. In auxiliary analyses, we show that the policy mostly affected school environment aspects that are weakly associated with learning and that the effects are too small for any educational outcome effects to materialize.

The remainder of the paper proceeds as follows: We describe the institutional context and data in Section 2 and present the empirical approach in Section 3. We present the estimated effects on student teacher ratios in Section 4, academic outcomes in Section 5., and school environment in Section 6. Section 7 discusses the magnitudes of the effects and different explanations for the limited effects on academic outcomes, while Section 8 concludes.

2 Context, reform, and data

Up to tenth grade (age 16), municipalities administer the public schools in Norway. The municipalities acquire revenue from income tax, user charges on services, and transfers from the central government. They spend money on schools as well as on child care, health care, and other services like water supply and renovation. Schools are free of charge and compulsory through grade 10, and the share of lower secondary private schools is small (about 3-4%⁵).

Between-school differences are relatively modest in Norway, including differences in resources. The allocation of input in terms of teaching personnel is highly compensatory. To reduce achievement

⁵Statistics Norway, table 05232 (<https://www.ssb.no/statbank/table/05232>).

gaps, school administrations in some large municipalities allocate more resources to schools with disadvantaged student populations. As a result, students from low-income families with less-educated parents typically attend smaller classes (Leuven et al., 2008). We illustrate this compensatory policy by means of entry test scores in grade 8 (lower secondary school). As shown in panel A of Appendix Table A.1, a difference in entry test scores of 10% of a standard deviation is associated with 0.35 fewer students per teacher in lower secondary schools, relative to a mean of 15.9.


In the fall of 2012, the Norwegian Parliament decided to provide four-year funding for an extra 600 teachers for grades 8 to 10 per year to municipal school administrations (NOK 1.5 billion, or USD 258 million using the 2012 exchange rate), starting with the 2013/2014 cohort. In line with the tradition for compensatory resource allocation, the extra resources were channeled to 166 lower secondary schools with an average student-teacher ratio for regular instruction⁶ above 20 *and* average grades at the end of grade 10 below the national mean in the previous school year (2011/2012). Thus, the schools could not manipulate their distance to the cutoffs. For schools that met both conditions, the number of extra teachers varied according to school size. Schools received funds for one, two, three, four, or five additional teachers depending on whether the number of students was 0-99, 100-199, 200-299, 300-399, or more than 400. The schools with extra funding were located in 98 different municipalities, and all counties in Norway were represented.

The purpose of the intervention was to enable schools to better tailor the teaching to the individual students and thereby raise basic skills, improve the learning environment, and reduce special needs education. The extra resources were intended to increase the number of qualified teachers in regular teaching and were not to be used for special needs education. Apart from these requirements, schools were free to organize the regular teaching as they wanted (two teachers in regular classes, divide classes into smaller groups, etc.). A survey administered to the principals at all treatment schools (with a 65% response rate) provides information on how schools used the extra funds (Kirkebøen et al., 2017). About three of four principals (77%) report that the extra funds were used to have two teachers available for the class, while 66% report that the funds were used to divide the class into smaller groups. Very few principals report that the extra resources were always used to divide groups by students' skills levels. However, nearly 60% report that groups were sometimes divided in this way (with 40% reporting that it was never divided according to skills level). The principals report that the resources were primarily used in math (95%), Norwegian (90%), and English (78%). Finally, 70% report that resources were shared across all lower secondary grades. Given our precise null findings documented below, it is also interesting to note that 95% of principals believed that the funding improved students' learning outcomes, suggesting that the perceived effectiveness of interventions may differ considerably from the evidence based on a credible treatment effect analysis.

The starting point for our data set is the population of eligible lower secondary schools. The extra funding was allocated on the basis of 2011/2012 regular student-teacher ratios and grade point average (GPAs) based on teacher-assessed subject-specific performance and externally-graded exams scores by the end of grade 10. Teacher grades dominate the GPA, but the written and oral exams also count (weight of about 0.1). Student-teacher ratios and GPAs are well defined for

⁶The student-teacher ratio is measured as the number of regular-instruction student hours divided by the number of regular-instruction teacher hours. Thus, it disregards student and teacher hours spent on for example special needs teaching and special services for Norwegian language learners, and measures of the average group size in regular instruction settings.

Figure 1. Pre-reform GPAs and student-teacher ratios of treated and control schools



Note: The figure shows the schools by student-teacher ratio and GPA in the pre-reform year 2011/2012. These forcing variables are constructed from student and school registers. The red dotted lines mark the two cut-offs. The markers indicate whether the schools received extra teachers (from a separate data source).

859 out of 1 089 public lower secondary schools, containing 97.5% of the students. The remaining 230 schools are small, and many of these schools cater for special needs students.

Figure 1 shows how schools are distributed along the student-teacher ratio (x-axis) and the GPA (y-axis) dimensions in the pre-reform year (2011/2012). The treated schools, marked with blue x, have below-average GPAs and group size above the average. The two sharp margins enable us to use a regression discontinuity framework, as discussed below.⁷

We merge our school treatment status dataset with four different data sources that together allow us to study a variety of outcomes: (i) the compulsory school register (“GSI”), with school-level data enabling us to study resource use in schools, including student-teacher ratios, the share of qualified teachers, and how teacher resources are used (e.g., on regular teaching and special needs teaching); (ii) individual teacher data from matched employer-employee data, which enable us to study teacher characteristics; (iii) student-level outcome data from administrative registers, including standardized tests, end-of-compulsory-school exam scores and teacher grades, and an early measure of progression in upper secondary school; and finally, (iv) school-level responses from an annual national student survey, which allow us to study school environment attributes.

Table 1 shows key descriptive statistics for our main student-level estimation sample, separately for the pre-funding (2009-2012) and the funding (2013-2016) periods. Here year refers to the first

⁷We define treatment from the forcing variables in our data, which may not correspond exactly to the data used by the ministry when they allocated extra funds. For all but one school in our sample, the expected and recorded numbers of teachers match. This school, with 257 students, did not receive extra teachers, despite being just below the GPA cutoff (and well above the student-teacher ratio cutoff) in our data.

semester of the school year, and students are indexed by their first year in lower secondary (e.g., 2009 represents the school year 2009/2010 and the students in grade 8 this school year). We include the 2009-2019 cohorts in the analysis, in total about 618,000 students (about 56,000 students per cohort). 169,000 students, or 27%, attended treatment schools that received extra funding during 2013-2016. The extra resources were channeled to schools with below-average GPAs and above-average student-teacher ratios, and the schools that later received extra funding have GPAs about one standard deviation lower than other schools and about one standard deviation more students per teacher in the pre-treatment year of 2011, as shown by panel A in Table 1. This corresponds to a student-teacher ratio in regular teaching of 22 in the treated schools and 19 in the non-treated schools. Furthermore, the treated schools are on average significantly larger, with about 341 students compared to 279 students for the non-treated schools. GPA levels, student-teacher ratios, and school size all vary more among the control schools than among the treatment schools.

In panel B of Table 1, we report student characteristics. While the sex composition is balanced, the treated schools have fewer students with at least one college-educated parent, more fathers with below-median earnings, and a larger fraction of students with two immigrant parents. These differences are expected since the funding targeted schools that have below-average GPAs. However, the differences are mostly modest, with 6 percentage points more students having at least one college-educated parent and the same difference for above-median earnings in the untreated schools. Concerning students' outcomes in Panel C, the average 8th grade entry test score of students in the treated schools is 0.13 SD lower in the pre-treatment years and the differential drops slightly in the post-treatment years. For all outcomes, there is a modest improvement in outcomes for students in the treatment school from pre- to post-treatment years. This is to be expected from mean reversion, as the schools were selected on the basis of poor performance in 2011 (Chay et al., 2005). In the untreated schools, there is a mix of positive and negative changes from pre- to post-treatment cohorts.

Table 1. Descriptive statistics

	Untreated schools		Treated schools	
	Pre-years	Post-year	Pre-years	Post-years
<i>A. School characteristics</i>	2011		2011	
GPA (de-meaned and standardized)	0.202		-0.743	
	(1.013)		(0.544)	
Student-teacher ratio (de-meaned and standardized)	-0.245		0.734	
	(0.999)		(0.570)	
Student-teacher ratio	19.0		22.4	
	(3.56)		(2.21)	
School size (# of students)	279		341	
	(139)		(102)	
<i>B. Student characteristics</i>	2009-2012	2013-2016	2009-2012	2013-2016
Female students (share)	0.488	0.491	0.487	0.487
	(0.500)	(0.500)	(0.500)	(0.500)
No parent with higher education (share)	0.482	0.429	0.543	0.488
	(0.500)	(0.495)	(0.498)	(0.500)
Father has below median income (share)	0.488	0.485	0.533	0.540
	(0.500)	(0.500)	(0.499)	(0.498)
Two foreign-born parents (share)	0.084	0.109	0.147	0.177
	(0.278)	(0.312)	(0.354)	(0.382)
<i>C. Student outcomes</i>				
8th grade entry test score	0.022	0.016	-0.114	-0.082
	(0.920)	(0.924)	(0.920)	(0.919)
9th grade test score	0.024	0.021	-0.100	-0.074
	(0.917)	(0.922)	(0.928)	(0.928)
Written exam score 10th grade	0.053	0.071	-0.050	-0.009
	(0.989)	(0.973)	(0.985)	(0.972)
Teacher assessment grade 10th grade	0.044	0.031	-0.100	-0.067
	(0.985)	(0.984)	(1.012)	(1.011)
Completed 2.year of upper secondary school in time	0.781	0.816	0.759	0.797
	(0.413)	(0.388)	(0.428)	(0.402)
Number of students	168 305	160 132	63 135	60 712

Note: Years refer to year of school register data and 8th grade test score, and the first semester of the school year. Standard deviation in parentheses, schools weighted by number of students. The pre-treatment school characteristics are those used to assign schools to treatment, and are based solely on 2011 data. High school completion is only observed for 2015 and earlier cohorts, other outcomes are observed for 93-99% of students (most missing for written exam score and teacher grades). Test scores, exam scores and teacher grades are standardized within year in the total student population.

3 Empirical strategy

Schools were treated by increased funding if they had a student-teacher ratio (*STR*) above 20 and grade point average (*GPA*) below the national mean in 2011 (i.e., the school year 2011/2012, following the notation from the previous section). These two necessary conditions (cutoffs) place each school in one of the four quadrants of Figure 1. The funding assignment offers two different

margins for evaluating treatment effects

$$\begin{aligned} |STR - 20| \rightarrow 0, \quad GPA < Mean & \quad (Student - teacher ratio margin) \\ |GPA - Mean| \rightarrow 0, \quad STR > 20 & \quad (GPA margin) \end{aligned} \quad (1)$$

and two placebo margins that can be used to validate the design

$$\begin{aligned} |STR - 20| \rightarrow 0, \quad GPA > Mean \\ |GPA - Mean| \rightarrow 0, \quad STR < 20 \end{aligned} \quad (2)$$

The reform lends itself to an RD analysis, and we start by estimating the specifications above, using local regressions and parametric RD estimation. We use the `rdrobust` and `rdplot` packages (Calonico et al., 2014) to estimate local linear regressions and bias-corrected confidence intervals, which are based on higher-order polynomials, and to plot the results. Effects are estimated separately for the two different cutoffs, and the sample is restricted to ± 0.5 standard deviation along the specific dimension in these RD analyses. A crucial prerequisite for identification by means of RD estimation is that there is no strategic sorting around the cutoffs. Since the extra funding was allocated based on historical data, we do not expect that schools were able to manipulate the GPA or student-teacher ratios that determined eligibility. In Figure A.1 in the Appendix, we also demonstrate that there is no indication of bunching around any of the the cutoffs.

In RD estimation, there is a trade-off between unbiasedness (improved by a small bandwidth) and precision (improved by a wider bandwidth and more data). With two forcing variables, there are several choices to be made in how to summarize the effects and analyze the data that add to the questions of optimal bandwidths and the functional form of the running variable(s). For instance, separate analyses can be run for the two different margins, and this can be done either parametrically or non-parametrically. Alternatively, both cutoffs can be used simultaneously (Cattaneo et al., 2020).⁸

Our preferred model is a parametric difference-in-RD specification in which both cutoffs and all schools, irrespective of their distances to the cutoffs, are used to identify the effects of the policy. We include quadratic controls for the distances to the cutoffs. We document that our results are robust with respect to the choice of functional form and whether to include all observations, or only those close to a cutoff. Our choice of specification is guided by transparency (as it is easy to understand), credibility (by having “built-in” placebo checks), and precision gains (more of the data is exploited). To simplify the presentation of the model, let z_S be the running variable on the student-teacher ratio margin measuring the distance to the cutoff, and z_G be the running variable on the GPA margin, such that treated schools have $z_S > 0$ and $z_G > 0$ (all measured in the pre-treatment year 2011). We estimate

$$y_t = \delta_t(Z_G Z_S) + \gamma_{1t} Z_G + \gamma_{2t} Z_S + \eta_{1t} z_G + \eta_{2t} z_S + \eta_{11t} z_G^2 + \eta_{22t} z_S^2 + \eta_{12t} z_G z_S + x_t \beta + \epsilon_t, \quad (3)$$

⁸Rather than using the RMSE optimal bandwidth selection of the `rdrobust` and `rdplot` packages, we will use a fixed bandwidth of 0.5 SD for the RD estimation. We found that the RMSE optimal bandwidths (typically 0.2-0.3 SD) produced implausibly strong gradients around the cut-offs and correspondingly implausibly large, although imprecise, estimates, cf. Figures 2 and A.2.

where y_t denotes the outcome of interest for cohort t (defined by entry year in grade 8), z_G and z_S are our running variables, and Z_G and Z_S are dummies for $Z_G := z_G > 0$ and $Z_S := z_S > 0$.

The main coefficient of interest is δ_t , which shows the effects of being a student in cohort t in a school above both cutoffs. This difference-in-RD design, and the interpretation of δ_t , is similar to a two-by-two difference-in-difference setup, where the Z_G and Z_S variables control for potential (cohort-specific) differences in outcomes around the two cutoffs in untreated schools. However, unlike in a simple two-by-two setup, we also observe distances to the two cutoffs (the z variables), and include quadratic distance controls, as shown in (3). Under standard RD regularity assumptions and absent treatment, sufficiently flexible distance controls will control for differences in potential outcomes around the cutoffs, making the controls for above/below the two cutoffs superfluous. In this case, $\gamma_1 = \gamma_2 = 0$ and the inclusion of Z_G and Z_S variables is not necessary for consistent estimation of δ . By way of contrast, suppose there are discontinuities around the cutoffs in the absence of treatment. In that case, the inclusion of Z_G and Z_S presumably accounts for such differences and is necessary to identify δ . Additionally, the Z_G and Z_S variables serve as a kind of placebo test. Evidence that $\gamma_1 = \gamma_2 = 0$ indicates that the controls for the distance to the cutoffs are sufficient to control for differences between treated and untreated schools and further strengthens the credibility of the effect estimate δ . This placebo test builds upon the reasoning that, in the absence of the treatment, crossing both cutoffs should not affect the outcome if crossing the separate cutoffs has no effect.

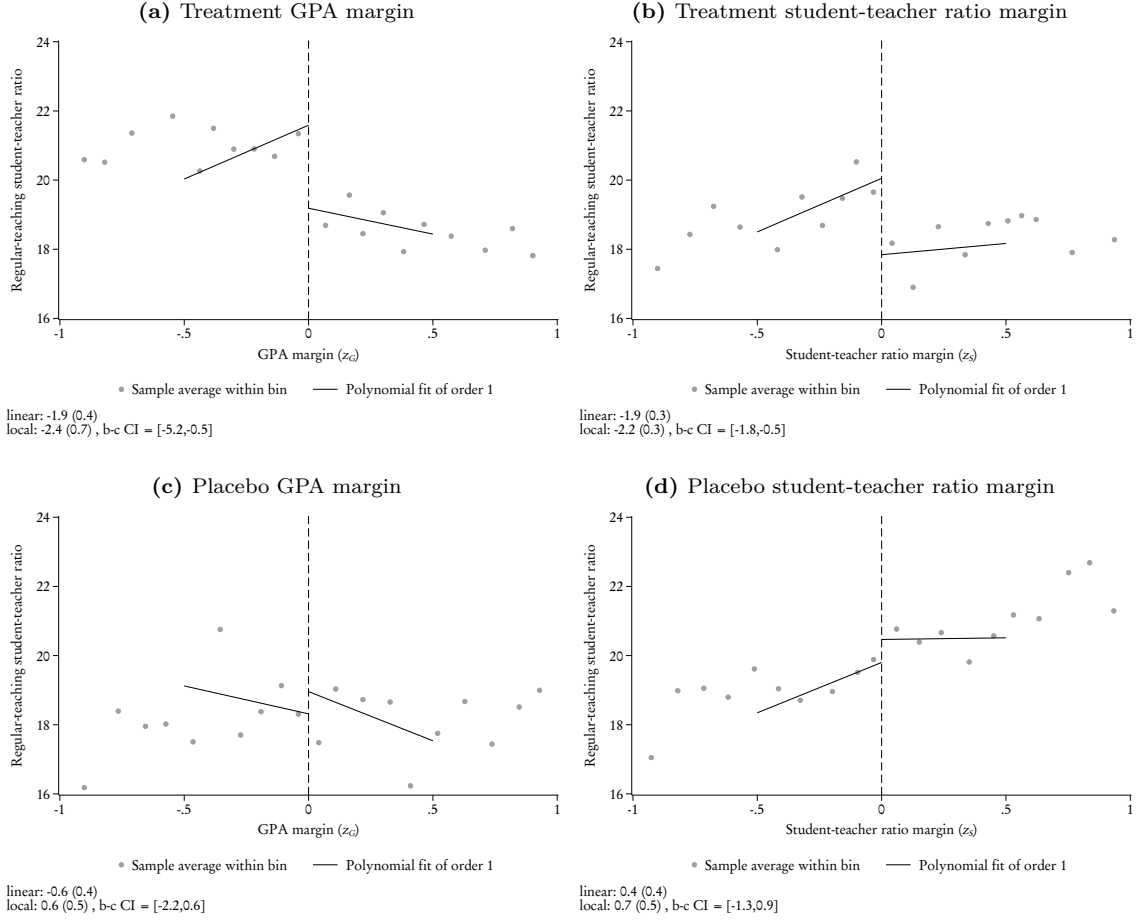
As indicated by the cohort subscripts on the δ , γ , and η coefficients in (3), all inference is based on variation within the cohort, comparing outcomes of treated and untreated schools, conditional on the forcing variables. For treated cohorts, δ provides our main policy effect estimates. For later cohorts who attended grades 8-10 after the extra funding was terminated, δ provides estimates of lasting effects on the schools, potentially including local continuation of the extra teachers using other funding. For cohorts in grades 8 to 10 before the reform, δ shows whether there were pre-existing differences between treatment and control schools, providing us with another placebo check. In addition, we run placebo regressions where we test for discontinuities in background characteristics such as entry test scores and parental education.

The control variables in x include gender, age, year fixed effects, a cubic in the 8th grade test score, and parental education. Control variables are included to increase precision, and results without controls are very similar (as shown in the Appendix). Because treatment is at the cohort-school level, standard errors will be clustered at school level when we study individual-level outcomes. The results will be summarized in coefficient plots with confidence intervals.

4 School resources

The policy funded new teaching positions in targeted schools over a period of four years (2013-2016). We start by examining how it affected the student-teacher ratio. We first present results for the two treatment and placebo margins using the RD design (Figure 2) before turning to our preferred difference-in-RD specification (Figure 3). Panel (a) of Figure 2 shows the RD estimate for the treatment GPA margin (with $z_S > 0$). There is a clear drop in the student-teacher ratio around the cutoff from about 21.5 to 19. The linear, local linear, and bias-corrected local linear estimates are all similar and statistically significant with a reduction of about two students per

Figure 2. RD estimates of the effect on student-teacher ratio



Note: The graphs show RD estimates. Data are from the years 2013-2016 and the outcome is the student-teacher ratio for regular teaching. Figure notes show coefficients and standard errors from linear and local linear regressions. All analyses use school-level data and student weights. The lines show the local linear regressions and are estimated using Calonico et al. (2014), with triangular weights and a fixed bandwidth of 0.5. Bias-corrected confidence intervals (estimated using higher-order polynomial) in brackets. Bins are quantile-based.

teacher. Panel (b) shows similar results for the treatment student-teacher ratio margin (with $z_G > 0$). As in the case of the GPA margin, there is a drop of about two students per teacher, from 20 to 18, and all estimates are of the same magnitude and statistically significant.

In sub-figures (c) and (d), we show the corresponding results for the two placebo margins (i.e., with $z_S < 0$ and $z_G < 0$, respectively). Consistent with our expectations, we find zero effects on the student-teacher ratio from crossing the two separate cutoffs. All placebo estimates are small and statistically insignificant.

Our main estimates of Equation (3) are presented in Figure 3. The figure displays the point estimates and 95% confidence intervals for separate school cohorts in treatment schools (δ ; blue circle symbol), GPA margin placebo schools (i.e., $z_G > 0$ and $z_S < 0$, γ_1 ; green diamond symbol), and placebo schools at the student-teacher ratio margin (i.e., $z_S > 0$ and $z_G < 0$, γ_2 ; red triangle symbol). All estimates are within-year, relative to the group of schools with high GPA and a

low student-teacher ratio, and conditional on the parametric specification in Equation (3). Note that the 2011 estimates will be zero as the outcome is the same as one of the forcing variables in this year. The treatment coefficients during the shaded period in the graph (2013-2016), can be interpreted as effects of the funding.

Figure 3 shows that the policy reduced the class size during the treatment years (blue circles in the shaded area of the figure). The average effect over the four treatment years is a reduction of 2.3 in the student-teacher ratio, which is substantial compared to the treated school average of 22.4 before the reform. In Figure A.3 in the Appendix, we find the same result using logs (11%). Moreover, in Appendix Figure A.4, we show that the policy increased the number of teachers by about 3.9 teachers (17%) in treated schools, closely corresponding to the number of teachers funded. Thus, there is no evidence of municipalities' funding being crowded out or substitution to other schools (Section 7 discusses substitution effects in more detail). The evidence clearly shows that the funding policy had the intended effect on teacher input.⁹ While there are no significant effects on student-teacher ratios after the termination of the reform, there are indications of a gradual reversion to control school levels. This is consistent with reports that some municipalities continued with a reduced student-teacher ratio funded by other sources.


We check the validity of our design using three distinct placebo dimensions as well as the pre-treatment balancing tests discussed in Section 5, all of which support our identification strategy. The first placebo dimension is treatment school coefficients before the funding period (2009-2012). Changes in the student-teacher ratio (or other outcomes) before the funding period would suggest trends within treatment schools that may bias our effect estimates. We see no indications of significant differences in schools that later become treatment schools.

The second and third placebo dimensions are the non-funded schools with low GPA and high student-teacher ratio, respectively. Since these placebo schools did not receive funding during the period, these schools can be used to rule out other changes concurrent with the funding policy that presumably would have affected all schools with either low GPAs or high student-teacher ratios. These coefficients seem to trend over the reform period and are close to significant in 2016. However, as discussed in Section 3, these coefficients do not have to be zero for effect estimates to be valid. Over time, the forcing variables may become less predictive of later outcomes, and controls for low initial GPA and high student-teacher ratios may become more important for correct inference. Nevertheless, there are no significant differences on either placebo margin in any year or in total over both margins and all reform years (p-value of 0.135), and the placebo estimates are small compared to the effect estimate. These results indicate that the parametric specification of the forcing variables is sufficient to control for between-school differences around the cutoffs and strengthens the credibility of the estimated treatment effects.

In Appendix Table A.3, we investigate the robustness and compare estimates from alternative specifications, including Equation (3) with global linear controls, as well as local specifications with either linear or quadratic controls. The local specifications are similar to those in Figure 2, based on observations within 0.5 SD of either cutoff using triangular weights based on distance to the closest cutoff. Although the estimated effect on student-teacher ratio is slightly larger with the

⁹ This was expected, but in light of another Norwegian education policy of resources to primary schools not leading to lower student-teacher ratios (Reiling et al., 2021), not obvious.

Figure 3. Effects on the student-teacher ratio



Note: The graph shows estimates and confidence intervals from estimating eq. (3). The outcome is the student-teacher ratio. The different series correspond to treatment effects, δ in eq. (3), and placebo effects from the non-treatment margins, γ in eq. (3). The treatment period is shaded. The text below the figure reports the sample mean and standard deviation of the outcome, estimated effect and standard errors for a pooled analysis of the treatment years, and the p -value of a joint test of all placebo effects (γ) for all treatment years. The regression uses school-level data with student weights and robust standard errors.

global linear specification, the effect is not significantly different across specifications. In the global linear specification, the coefficient for the STR placebo margin is also significant, suggesting that the quadratic terms are needed to fully capture school heterogeneity. The estimated coefficients from the two local specifications are very similar to those of our preferred specification, but the standard errors are larger.

In sum, we conclude that there is strong and robust evidence of a substantial reduction in the student-teacher ratio, from about 22 to 20 students per teacher in regular instruction.

5 Academic outcomes

The main objective of the intervention was to improve student outcomes. The number of exposure years differs across cohorts and outcomes. We start by studying the effects on the standardized test scores in grade 9, when students had been exposed to the treatment for about one school year. Figure 4 shows results from school-level RD analyses of the two margins, similar to what we did for school resources in Section 4. The notes to the graphs report results from student-level RD analyses with and without additional student-level control variables (with standard errors clustered at school level). If the extra teachers improved student outcomes, we should see a jump in test scores at the two margins. There is no indication of any test score discontinuity in neither panels (a) and (b) (treatment margins) nor in panels (c) and (d) (placebo margins). The estimated effects are small in all sub-figures and never statistically significant. However, the RD estimates

are not very precise, and in several cases, we cannot reject intervention effects of 5% of a standard deviation.

Figure 5(a) shows the effects of our main specification in Equation (3) on students' 9th-grade standardized test scores. One year with a lower student-teacher ratio did not improve the treated students scores on the national standardized test. The average estimated effect across the four treated cohorts is 0.6% of a standard deviation, with a confidence interval from -1.6% to 2.8%. The effect of improved funding to hire more teachers on 9th-grade test scores is minor at best. The results are similar without control variables (see Appendix Figure A.8), but the precision is much lower. All placebo test estimates are small and insignificant. The pre-reform placebo effects are similar in size to the reform effects. Moreover, the within-year placebo estimates are close to zero, thus providing additional evidence of the validity of our design.


Compositional changes within treatment schools represent a potential concern. We examine this by studying placebo effects on students' entry test scores in lower secondary schools as well as on parental characteristics. Since funding in lower secondary schools cannot influence test scores at entry, any significant placebo effects would indicate compositional changes correlated with the treatment. Although we control for students' entry test scores when estimating the effect on 9th grade test scores, which accounts for this exact source of bias, any within-school changes in students' academic abilities correlated with treatment would be worrying since it would suggest that other unobserved factors were co-occurring with our treatment. Reassuringly, there are no discontinuities in entry test scores nor in parental earnings, parental education, or immigrant background in treatment schools before, during, or after the funding period, as seen in Appendix Figure A.9. Nor are there any systematic differences in any of the placebo schools. All in all, the placebo and balance tests make us confident that composition bias is unlikely.

The precisely estimated zero average effect on 9th grade test scores strongly restricts potential significant effects for subgroups. Any effect of the extra funding is either small, limited to a small group of students, or counteracted by a negative effect for other students. Previous research has found that it is primarily disadvantaged students who benefit from extra funding (Jackson, 2020). In Appendix Figure A.10, we investigate heterogeneous effects by individual (gender, earlier test scores) and parental characteristics (immigrant background, income and education). We do not find significant effects for any group. Indeed, we can reject effects larger than 4% of a standard deviation for most groups. The single exception is children of immigrants, for whom we can only reject effects larger than 6%. This mostly reflects lower precision for this group (which is smaller than the others studied), since the point estimate of less than 2% hardly indicates any large effect. In Appendix Figure A.11 we investigate heterogeneity by school characteristics and find no significant effect for most types of schools, with the exception of marginally significant effects in schools with lower average 8th grade test scores. In all, we conclude that there is very limited heterogeneity in the effects.

Treatment school students had been exposed to a lower student-teacher ratio for just over one school year when they took the 9th grade test. Since a longer exposure may have a more substantial impact, we also estimate the effects at grade 10 and beyond. By the end of compulsory schooling (grade 10), the 2011 and 2016 cohorts had been exposed to one year of extra funding, the 2012 and 2015 cohorts to two years, and the 2013 and 2014 cohorts to three years.

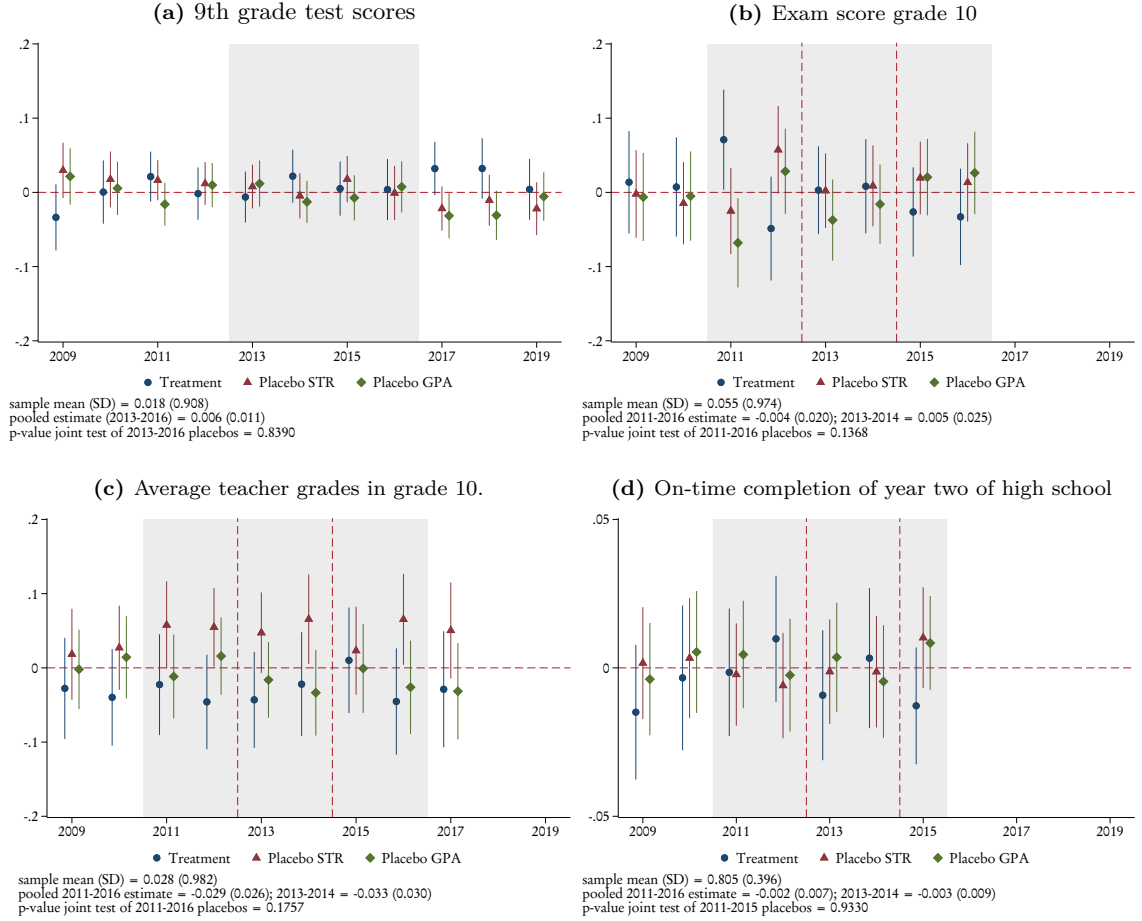
Panel (b) of Figure 5 shows the effects on externally graded exam scores at the end of

Figure 4. RD estimates of effects on 9th grade test scores



Note: The graphs show RD estimates for the average score from the 9th grade test. The data are provided by students sitting the grade 9 test in 2013-2016. The graph is based on school-level data and student weights. The lines show the local linear regressions and are estimated using Calonico et al. (2014), with no additional control variables, triangular weights and a fixed bandwidth of 0.5. Bins are quantile-based. Figure notes show coefficients and standard errors from student-level linear and local regressions. Student controls in the linear and local regressions include gender, year dummies, a cubic in the 8th grade test score, and parental education. Bias-corrected confidence interval (estimated using higher-order polynomial) in brackets.

Figure 5. Effects on academic outcomes



Note: The graphs show estimates and confidence intervals from estimating eq. (3). Outcomes are a) 9th grade test scores, b) exam score, c) teacher grades and d) completion of year two of high school. Control variables are gender, age, year fixed effects, a cubic in the 8th grade test score, and parental education. The different series correspond to treatment effects, δ in eq. (3), and placebo effects from the non-treatment margins, γ in eq. (3). The X-axis is the year of the 8th grade test, treated cohorts are shaded. In sub-figures b-d the dashed vertical lines indicate cohorts treated for three years. The figure notes show the sample mean and standard deviation of the outcome, estimated effect and standard errors for a pooled analysis of the treatment years and the p -value of a joint test of all placebo effects (γ) for all treatment years. The regression uses student-level data and clusters standard errors at school level.

compulsory schooling.¹⁰ The average effect across all treated cohorts is -0.4% of a standard deviation. According to a dose-response rationale, being exposed to more resources over a longer period should have a larger impact (see Jackson and Mackevicius, 2021). However, the average impact for school cohorts exposed for three years to extra funding, delimited with dashed vertical lines, is only 0.5% of a standard deviation and insignificant.¹¹

We can also investigate the effects on a broader set of academic outcomes as well as non-cognitive outcomes such as the school environment. These outcomes reflect potentially important consequences of more teachers, since test scores do not fully capture all aspects of students' learning. To the extent that non-cognitive skills impact longer-term outcomes, we would expect that such effects show up in teacher-graded tests and school dropout. Specifically, since the funding did not improve the 9th grade standardized test scores nor 10th grade exam scores, any effects on 10th grade teacher-assessed grades and high school completion would likely have been caused by non-cognitive skills (or effects on teachers' grading practices). We find no significant effect on teacher-assessed grades. The estimated average effect across all treated cohorts is -2.8% of a standard deviation for teacher-assessed grades, and the impact for those exposed for three years is -3.3% of a standard deviation. Finally, there is no significant effect for on-time completion of the second year of upper secondary school, with the 95% confidence interval ranging, in terms of percentage points, from a reduction of 1.6 to an increase of 1.2.¹²

Appendix Figure A.8 shows that the results are similar without controls. As for the student-teacher ratio, the estimated effects on academic outcomes are robust with respect to the choice of specification, see Appendix Tables A.5-A.7. As for the student-teacher ratio there are some indications that the global linear specification is insufficiently flexible to control for all differences not related to the reform. The other specifications all give similar results, and among these the global quadratic gives the highest precision.

6 School environment

This section examines whether the funding affected students in ways not captured by academic outcomes. In a nationwide survey of Norwegian 10th graders, students respond anonymously on subjects such as well-being in school, teacher support, and bullying. We construct an index based

¹⁰RD graphs similar to Figure 4 for the outcomes in panels (b)-(d) of Figure 5 are included in the Appendix, as Appendix Figures A.5-A.7.

¹¹While the point estimate for 2011 is positive and significant, there is no clear pattern for the whole set of estimates, and a joint test of all the effects for the period 2011-2016 does not reject the null of zero effects (p -value = 0.136). Furthermore, the positive effect estimate in 2011 is also accompanied by negative estimates at both placebo margins, reducing the total difference between treated schools and schools with high GPAs and low STRs. Finally, we did not find any effect on this cohort's 9th grade test scores (discussed above), concurrent teacher grades, or on later school dropout (both discussed below). Thus, we interpret the 2011 effect as spuriously significant and not indicating any real effect on this cohort. Given the number of hypotheses tested, one or more spuriously significant result is not unexpected, and the 2011 estimate is only just significant (p -value = 0.039).

¹²We are not able to observe the treated cohorts complete upper secondary school, which nominally takes from 3-4.5 years (depending on the track) and which is customarily measured five years after completion of lower secondary (or enrolling in upper secondary school). However, based on earlier cohorts, completion of the second year strongly predicts eventual completion of upper secondary school. At about 80%, the share of students completing the second year on time is similar to the share completing upper secondary school within five years. Students completing the second year on time have a 50 percentage points higher probability of completing upper secondary school in Norway. Conditioning on results from lower secondary school, gender, and parental education reduces this difference to 34 percentage points. In Figure A.12 we investigate on-time enrollment in the third year, and find insignificant negative effects.

on aggregate student responses from 11 sub-indices to test for any effect on the school environment. The students score each item on a scale from one to five, and the sub-indices represent school-level averages of students' average scores across a small number of related questions. Except for bullying, higher values imply a better school environment.¹³ To construct the index, we scale each sub-index by dividing by the student-level standard deviation and average the scaled sub-indices (bullying is rescaled as 1 - sub-index before averaging, such that higher values imply better outcomes for all sub-indices).¹⁴

Unlike in analyses of academic outcomes, we find evidence suggesting that the extra teachers improved the school environment, as measured by our index, by about 5% of a student-level standard deviation, as shown by Figure 6. While all single-year effect estimates are insignificant, they are all positive, and the average effects over the 2011-2016 cohorts (exposed for one, two, or three years to extra teachers) and 2013-2014 cohorts (exposed for three years to extra teachers) are significant at the 5% level.¹⁵

The validity of the design and the credibility of the positive treatment effects are supported by the small and insignificant placebo effects in Figure 6. Since the student survey was redesigned in 2012, we do not have data for all sub-indices in all pre-treatment years. However, there are no indications of pre-treatment differences for the sub-indices that are covered throughout (Figures A.14 and A.15). Moreover, the school environment improvements among treated schools disappear once the funding ends, as shown by the precisely estimated zero treatment coefficients for 2017 and 2018. Notably, the effects are already reduced for the 2016 cohort. When the 2016 cohort answered the survey in the fall of 2018, the funding had terminated more than a year before, and there was little difference remaining in student-teacher ratios (see Figure 3). In contrast, the 2015 cohort answered the survey a few months into the first school year after the end of the reform, when differences in student-teacher ratios were still large, and when some schools may still have had extra teachers (see Figure 3). While the estimates are not sufficiently precise to reject equality of the yearly effects, this cross-cohort pattern suggests that the school environment effects are relatively short-lived.


In light of the absent effects on academic outcomes, a skeptic might ask whether the school environment effects are spurious; that is, merely random artifacts of testing several outcomes and specifications. While no definite conclusion can be drawn, a multiple testing adjustment alleviates such concerns. Although we did not file a pre-analysis plan with a pre-determined number of tests, the paper is based on a policy evaluation commissioned by the Norwegian Directorate for

¹³For example, the students are asked "Do you enjoy school?" with possible answers "Not at all", "Not very much", "Somewhat", "Quite well", and "Very well", and asked "Do you feel that your teachers care about you?" with possible answers "None", "Only one", "A few", "Most of them", and "All".

¹⁴ The student-level standard deviation ranges from 0.7-1.0 in the original one to five units, meaning effects in original units are slightly larger than the standardized effects we report.

¹⁵ Corresponding estimates for each of the 11 different sub-indices are presented in Figures A.14 and A.15. These constitute all published sub-indices with two exceptions: share bullied, which is a transformation of the bullying sub-index, and which is measured as a share of the students, and support from parents, which does not reflect the classroom environment, and which we will return to in Section 7. RD graphs investigating effects on the overall index on all margins are presented in Appendix Figure A.13.

Figure 6. Effects on the school environment index



sample mean (SD) = 4.572 (0.215)
pooled 2011-2016 estimate = 0.048** (0.015); 2013-2014 = 0.061** (0.026)
p-value joint test of 2011-2016 placebos = 0.9948

Note: The graphs show estimates and confidence intervals from estimating (3) by cohort. Outcome is an index summarizing the sub-indices from the student survey presented in Figures A.14 and A.15. The different series correspond to treatment effects, δ in eq. (3), and placebo effects from the non-treatment margins, γ in eq. (3). The treated cohorts are shaded. The dashed vertical lines delimit the cohorts in their third year of treatment when they answer the survey. The figure notes show the sample mean and standard deviation of the outcome, estimated effect and standard errors for a pooled analysis of the treatment years, and the p -value of a joint test of all placebo effects (γ) for all treatment years. The regression uses school-level data with student weights.

Education and Training. Their call for proposals and the funded protocol entail some pre-registered decisions. Consistent with the overall aim of the reform (see Section 2) and the call, our funded protocol specified the outcomes of interest as school resources (with special needs teaching being an outcome and student-teacher ratios essentially a first stage), learning outcomes (including 9th grade test scores and exam scores), school environment, and longer-term outcomes (early measures of upper secondary school completion). Thus, the school environment is one out of four pre-specified domains and five pre-specified outcome measures. Our design also involves a choice of whether to include partly treated cohorts or not when averaging the effects for the longer-term outcomes, which was not specified in the proposal. Still, the number of main hypotheses tested is arguably smaller than or equal to 10. The p -values of the average effects on the school environment index are 0.018 for the 2013 and 2014 cohorts (fully treated) average and 0.001 for all the 2011-2016 cohorts. The latter is thus significant at the 5% level even with Bonferroni-adjustment for testing 10 hypotheses, which is known to be conservative, especially with non-independent tests. The average effects for the 2013-2014 cohorts is not, but this is due to lower precision, as the estimated effect on the school environment is largest for the fully treated. Thus, the school environment effects are robust to multiple testing adjustment and unlikely to be spurious.¹⁶

In order to understand our divergent results for academic outcomes and the school environment, we look for any pattern across areas or aspects of the school environment. In Panel (a) of Figure 7, we summarize the treatment effects on the pooled cohorts both for the index from Figure 6 and for the separate sub-indices presented in Appendix Figures A.14 and A.15. All outcomes are measured in terms of comparable student-level standard deviations. Although all estimates are positive, not all are statistically significant, and the magnitude varies across dimensions. The effects are strongest for assessments that support learning (formative assessment), student democracy, guidance on educational choices, support from teachers, culture for learning, and school well-being, with smaller effects for bullying, sense of mastery, common rules, academic challenge, and school motivation.


Funding effects on academic outcomes that operate via a slightly better school environment obviously depend on the impact of a slight improvement in the school environment on such outcomes. Thus, the next question is, what is the impact of the school environment on academic outcomes? In Appendix Table A.2, we investigate the association between the school environment index and student outcomes. Unsurprisingly, cohorts of students that report a better environment also have more favorable outcomes. Moreover, this association is significantly positive even within school and conditional on student characteristics. While these associations may not be credible causal estimates,¹⁷ they provide a relevant baseline for comparing estimated effects on the school environment and student outcomes.

The funding effect on the school environment is modest. To illustrate, an effect of extra funding on the school environment of 5% of a standard deviation (Figure 7) implies that one in twenty students answers that “most teachers” are supportive rather than just “a few teachers”, or that one in twenty “enjoy school very much” rather than “enjoy school”. Even if a better school environment

¹⁶The significant effect on the school environment is also robust with respect to the choice of global vs local specification, see Table A.8).

¹⁷There may still be unmeasured student characteristics contributing to the school environment and later outcomes. Also, by looking at within-school associations, we disregard how time-invariant school quality may impact both student outcomes and the school environment.

Figure 7. Effects on separate sub-indices and associations between sub-indices and longer-term outcomes



Note: Sub-graph (a) shows pooled effect estimates and confidence intervals of δ resulting from estimating (3) for 2011-2016 and 2013-2014 students, respectively. Sub-graph (b) shows point estimates and confidence intervals for within-school associations between the school environment indices and longer-term outcomes. Upper secondary school completion is measured five years after completing compulsory schooling. The index is an average of all the other sub-indices' outcomes. Outcomes are sorted by the 2011-2016 effect in sub-graph (a). All regressions use school-level data with student weights.

predicts improved educational outcomes, the within-school associations are also modest as shown in column (4) of Appendix Table A.2. If we multiply these effects of a better school environment on academic outcomes by the 5% improvement in the school environment (Figure 6), the funding of extra teachers is expected to improve students' exam scores by 0.3% of a standard deviation and completion of the second year of upper secondary school by 0.2 percentage point as a result of the better school environment.

Furthermore, the funding appears to have the strongest impact on the school environment indicators that matter least for student outcomes. While Panel (a) of Figure 7 shows the effects of the extra teachers on the different sub-indices, Panel (b) shows within-school associations with student outcomes by sub-index. With the exception of motivation and completion of upper secondary school, all associations are positive, and mostly significantly so. However, while neither the effect estimates nor the within-school associations are sufficiently precise for equality to be rejected across most sub-indices, there is a clear negative correlation between the two. The estimated effects of extra teachers are greater on sub-indices that are more weakly related to educational outcomes.

7 Discussion

The intervention studied in this paper represented a significant investment of public funds that reduced the student-teacher ratio by about 10% in the schools and cohorts affected. More teachers led to an improved school environment as experienced by the students. This section first discusses effect magnitudes and cost effectiveness, and we then offer possible explanations for why it did not improve academic outcomes.

7.1 Magnitudes and cost effectiveness

We find no effect on any of the academic outcomes. Standardized test scores, exam scores, and teacher-assessed grades are all measured in terms of student-level standard deviations, and thus have comparable scales. The estimated effects on test scores are the most precisely estimated, allowing us to rule out effects larger than 3% of a standard deviation. However, these effects are estimated after just over one year of exposure, while one may want to know whether sustained exposure to lower student-teacher ratios is more effective. Additionally, to compare the results with the latest summary of US studies, focusing on longer-term exposure (Jackson and Mackevicius, 2021), we concentrate on intervention impacts on exam scores after three years of exposure. The estimated intervention effect on exam scores was 0.5% of a standard deviation, with an upper bound of 5.5%.¹⁸

Cost-effectiveness is a key test for any intervention. Recall that the per-student costs in our intervention amount to USD 1,400 per year. To put a value on the potential effects on learning, we use the estimated association between school value-added and future earnings from Kirkebøen (2021), who finds that one standard deviation higher exam score increases future labor earnings by 1.5%. Using 2014 data, we estimate the average present value of labor earnings in Norway at

¹⁸This three-year exam scores estimate is more precise than a linear extrapolation of the effects on test scores.

age 13 at USD 1.03 million.¹⁹ Thus, the upper bound of the effect on exam scores (5.5% of an SD) is valued at USD 854 per student, or 61% of the cost, while the value of an effect corresponding to the point estimate (0.5% of an SD) is only 6% of the cost. Therefore, more funds to lower the student-teacher ratio are not justified in terms of lifetime earnings.

Another question is whether the intervention provides sufficient power to rule out effect sizes found in previous studies. In their meta-analysis of recent, well-identified US studies, Jackson and Mackevicius (2021) find that a USD 1,000 increase in per-pupil public school spending for four years leads to an average increase in test scores of 3.5% of a standard deviation and in high school graduation of 1.92 percentage points. Our intervention is similar in magnitude to their standardized treatment, as the extra funding amounts to about USD 1,400 per student year for three years. The point estimates for effects on exam scores and completing (the second year of) upper secondary school are considerably smaller than the corresponding US results, with an estimated increase in exam scores of 0.5% of a standard deviation and a 0.3 percentage point reduction in upper secondary school completion.

For exam scores, our point estimate is about one-seventh of the US average. Moreover, it is only one-third of the 25th percentile in the distribution of effect sizes from the 31 studies included in Jackson and Mackevicius (2021)’s meta-regressions. However, since the 95% confidence band includes effects as large as 5.5% of a standard deviation, we cannot reject equality with the average US effect of 3.5% (p-value = 0.244). If we include all the 2011-2016 cohorts (which have been treated for one to three years), we nearly reject an effect equal to the US average (p-value = 0.050). For completion rates, our point estimate is negative (although not significant), unlike all estimated effects on completion in Jackson and Mackevicius (2021). The upper limit of our confidence band is 1.6 percentage points, allowing us to reject equality with the US average of 1.92 percentage points (p-value = 0.020). Thus, although our test scores and completion effects are markedly smaller than the US results, only the confidence interval of the completion estimate fails to cover the US estimate.

7.2 Possible explanations for the limited effects on academic outcomes

Teacher funding effects may depend on the institutional context, which raises the question of whether our results are policy-relevant outside of Norway. In particular, if there are diminishing marginal returns to school spending, we may expect additional funding to have small effects in a country like Norway, where spending is already high (Schleicher, 2018). Indeed, although the ratio of students to teaching staff is similar to many other European countries, OECD statistics show that it is slightly below the EU23 mean and the US mean (OECD, 2021).²⁰ However, this diminishing returns explanation is challenged by the fact that we identify marginal effects

¹⁹This estimate is based on average earnings by age for the entire Norwegian population, discounted to age 13 of the student, i.e. at the start of lower secondary school, and summed over age: $\sum_{age=16}^{67} (1+discount)^{-(age-13)} \bar{y}_{age}$. The real discount rate is 4%, in line with what the Norwegian Ministry of Finance recommends for public investment.

²⁰ The ratio is 9 in Norway, compared to 11 in the EU23. The average class size is 23 for the EU23, not far from the ratio between regular teacher hours and student hours in this paper. Unfortunately, OECD does not include class size information for Norway.

for students in schools with an above-average student-teacher ratio and below-average results. Moreover, there is no consensus on diminishing marginal returns to school spending; for example, Jackson and Mackevicius (2021) find little evidence of this in the US.

Furthermore, by finding no effects of extra teachers on academic outcomes, our study fits with US studies on differential effects across types of school spending. For example, recent studies suggest that while additional operational expenditures (e.g., teacher pay) impact student outcomes (Baron, 2021; Abbott et al., 2020; Brunner et al., 2021), spending increases in areas where they lead to the hire of more teachers seem to have no effects (Brunner et al., 2020). Moreover, although the findings on class size in the literature are mixed, our study fits with several studies that find no effects as a result of reducing class size in various countries (Angrist et al., 2019, 2017; Argaw and Puhani, 2018), including Norway (Falch et al., 2017; Leuven et al., 2008; Leuven and Løkken, 2020). Thus, our results appear to be relevant for policymakers outside Norway.

While we cannot explain why the extra school funding did not improve academic outcomes, or how the intervention could have been designed to raise student performance, the data allow us to investigate some potential mechanisms. If school and parental inputs are substitutes in educational production, parents may reduce the effort and time devoted to their children’s learning activities in response to a reduction in the student-teacher ratio. The evidence on parental reactions to changes in school inputs is far from conclusive, however (Rabe, 2019). For example, while a Swedish study points to lower parental effort when classes are small (Fredriksson et al., 2016), the opposite has been found in Norway (Bonesrønning, 2004). To test the potential part played by parental adjustment, we use a question in the student survey. Students are asked about to the extent to which parents show interest in school, help with homework, encourage them, and expect them to do their best. Using the same model as for the school environment effect(s), we find no indication that the presence of more teachers reduces parental support (Appendix Figure A.16). If anything, the (insignificant) positive coefficient suggests a complementary response rather than a substitution effect, as students perceive that parents increase their support when the student-teacher ratio is reduced. Thus, it is unlikely that the lack of effects on academic outcomes is explained by a reduction in parental input that counters an effect due to more teachers.

School principals and municipalities may potentially reallocate resources within or between schools as a response to the extra funding (see discussion in Hoxby, 2000). Since the effect on total teachers corresponds to the intended increase, there is no indication of between-school reallocation. Within-school reallocation can happen along several dimensions, studied in Appendix Figure A.17.²¹ The funding increased regular teacher hours very similarly to total hours (including for example special needs education), and we find no change in teacher hours used for special needs teaching or special services for Norwegian language learners. Another possible margin of adjustment is teaching assistant hours, typically used to support special needs teaching. We do not find any effect on the ratio of assistant hours to student hours, nor any effect on assistant hours used for regular teaching.

Furthermore, if the treatment schools had difficulties recruiting competent teachers, we would expect to see reduced teacher qualifications in these schools (Gilraine, 2020). However, the increase

²¹In Figures A.17 and A.18, we study teacher-student ratios, rather than student-teacher ratios, in order to have a fixed denominator and be able to decompose the teacher hours in the nominator. The denominator is total student hours, including special needs teaching, and thus slightly different from the nominator of our measure of student-teacher ratio. However, this difference is minor for student hours.

in teacher hours taught by qualified teachers is very similar to the total effect on teacher hours. Additionally, there is no effect of the funding on the number of hours taught by teachers without formal qualifications (Appendix Figure A.18). Moreover, from matched employer-employee data (Appendix Figure A.19),²² we find neither statistically significant nor quantitatively substantial effects on average time since teachers completed their education, tenure at the school, the share of teachers with a teaching degree, or average sickness absence of the teachers (which may be a measure both of (un)available teaching resources and of teacher workload, if a heavy workload induces illness and absences).

The fact that we find no effects on teacher characteristics despite a substantial relative increase in the number of teachers reflects characteristics of the reform and the teacher labor market. While the extra funding substantially increased the number of teachers in the treated schools, these schools represent only about a quarter of total lower secondary school students (and a smaller share of schools and teachers, as the treatment schools are large schools with high student-teacher ratios). Furthermore, teachers may move between primary, lower secondary, and upper secondary schools. Also, there is substantial mobility in the teacher labor market, such that even in the treatment schools, new hires only increased moderately in response to the reform.

Overall, while crowding out of other school input and recruitment constraints are potential explanations for the null effects on student academic outcomes, the evidence provides no basis for this interpretation.

8 Conclusion

We have studied a large-scale intervention that funded 600 extra teachers per year distributed among 166 lower secondary schools in Norway for a four-year period, comparable to an increase in cost per student of USD 1,400 in each of the three years of lower secondary education. The funding assignment was based on two sharp conditions, allowing us to credibly identify the effects of additional resources using a regression discontinuity framework. We find that the extra funding reduced the student-teacher ratio by around 10%. There is no evidence of crowding out of other school inputs and the treatment schools did not seem to meet any restrictions when recruiting additional teachers. Nor did the policy induce other compensatory adjustments, such as changes in parental effort, teacher composition, teacher sickness absence, or special needs education. The reduced student-teacher ratio did not improve academic outcomes, as measured by 9th grade test scores, end-of compulsory schooling grades, or upper secondary school progression. Inspired by recent evidence of the effects of school inputs on non-cognitive outcomes, we also tested whether more teachers changed how students perceived the school environment. Although the funding improved the school environment, this effect was too small to induce a sizable effect on academic outcomes.

A major advantage of this study is that its findings have clear policy implications. In the quest for causal identification, studies often need to exploit natural experiments at margins far from relevant policies. In our study, however, credible identification and policy relevance go hand in hand, since we directly identify the parameter of interest for funding policies by the national

²²This data source is available only for 2015-2018 cohorts.

government: will increasing funding to provide more teachers improve student outcomes? For countries like Norway, we conclude that such funding may slightly improve the school environment, but will likely have no impact on the students' academic achievement or upper secondary school dropout.

References

- Abott, C., V. Kogan, S. Lavertu, and Z. Peskowitz (2020). School district operational spending and student outcomes: Evidence from tax elections in seven states. *Journal of Public Economics* 183, 104–142.
- Andersen, S. C., L. Beuchert, H. S. Nielsen, and M. K. Thomsen (2018). The Effect of Teacher’s Aides in the Classroom: Evidence from a Randomized Trial. *Journal of the European Economic Association* 18(1), 469–505.
- Angrist, J. D., E. Battistin, and D. Vuri (2017). In a small moment: Class size and moral hazard in the italian mezzogiorno. *American Economic Journal: Applied Economics* 9(4), 216–49.
- Angrist, J. D. and V. Lavy (1999). Using maimonides’ rule to estimate the effect of class size on scholastic achievement. *The Quarterly Journal of Economics* 114(2), 533–575.
- Angrist, J. D., V. Lavy, J. Leder-Luis, and A. Shany (2019). Maimonides’ rule redux. *American Economic Review: Insights* 1(3), 309–24.
- Argaw, B. A. and P. A. Puhani (2018). Does class size matter for school tracking outcomes after elementary school? quasi-experimental evidence using administrative panel data from germany. *Economics of Education Review* 65, 48–57.
- Baron, E. J. (2021). School spending and student outcomes: Evidence from revenue limit elections in wisconsin. *American Economic Journal: Economic Policy*.
- Bonesrønning, H. (2003). Class size effects on student achievement in norway: Patterns and explanations. *Southern Economic Journal*, 952–965.
- Bonesrønning, H. (2004). The determinants of parental effort in education production: do parents respond to changes in class size? *Economics of Education Review* 23(1), 1–9.
- Bonesrønning, H., H. Finseraas, I. Hardoy, J. M. Vaag Iversen, O. H. Nyhus, V. Opheim, K. V. Salvanes, A. M. J. Sandsør, and P. Schone (2021). Small group instruction to improve student performance in mathematics in early grades: Results from a randomized field experiment. *Available at SSRN 3979348*.
- Browning, M. and E. Heinesen (2007). Class size, teacher hours and educational attainment. *Scandinavian Journal of Economics* 109(2), 415–438.
- Brunner, E., B. Hoen, and J. Hyman (2021). School district revenue shocks, resource allocations, and student achievement: Evidence from the universe of us wind energy installations.
- Brunner, E., J. Hyman, and A. Ju (2020). School finance reforms, teachers’ unions, and the allocation of school resources. *Review of Economics and Statistics* 102(3), 473–489.
- Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust data-driven inference in the regression-discontinuity design. *The Stata Journal* 14(4), 909–946.
- Cattaneo, M. D., R. Titiunik, and G. Vazquez-Bare (2020). Analysis of regression-discontinuity designs with multiple cutoffs or multiple scores. *The Stata Journal* 20(4), 866–891.
- Chay, K. Y., P. J. McEwan, and M. Urquiola (2005). The central role of noise in evaluating interventions that use test scores to rank schools. *The American Economic Review* 95(4), 1237–1258.
- Chetty, R., J. N. Friedman, N. Hilger, E. Saez, D. W. Schanzenbach, and D. Yagan (2011). How does your kindergarten classroom affect your earnings? evidence from project star. *The Quarterly Journal of Economics* 126(4), 1593–1660.

- Cornelissen, T. and C. Dustmann (2019). Early school exposure, test scores, and noncognitive outcomes. *American Economic Journal: Economic Policy* 11(2), 35–63.
- Falch, T., A. M. J. Sandsør, and B. Strøm (2017). Do smaller classes always improve students’ long-run outcomes? *Oxford Bulletin of Economics and Statistics* 79(5), 654–688.
- Fredriksson, P., B. Öckert, and H. Oosterbeek (2013). Long-term effects of class size. *The Quarterly Journal of Economics* 128(1), 249–285.
- Fredriksson, P., B. Öckert, and H. Oosterbeek (2016). Parental responses to public investments in children: Evidence from a maximum class size rule. *Journal of Human Resources* 51(4), 832–868.
- Fryer Jr, R. G. (2017). The production of human capital in developed countries: Evidence from 196 randomized field experiments. In *Handbook of economic field experiments*, Volume 2, pp. 95–322. Elsevier.
- Gibbons, S., S. McNally, and M. Viarengo (2018). Does additional spending help urban schools? an evaluation using boundary discontinuities. *Journal of the European Economic Association* 16(5), 1618–1668.
- Gilraine, M. (2020). A method for disentangling multiple treatments from a regression discontinuity design. *Journal of Labor Economics* 38(4), 1267–1311.
- Guryan, J., J. Ludwig, M. P. Bhatt, P. J. Cook, J. M. Davis, K. Dodge, G. Farkas, R. G. Fryer Jr, S. Mayer, H. Pollack, et al. (2021). Not too late: Improving academic outcomes among adolescents. *NBER WP* (Number 28531).
- Hægeland, T., O. Raaum, and K. G. Salvanes (2012). Pennies from heaven? using exogenous tax variation to identify effects of school resources on pupil achievement. *Economics of Education Review* 31(5), 601–614.
- Hanushek, E. A. (2003). The failure of input-based schooling policies. *The Economic Journal* 113(485), 64–98.
- Heckman, J., R. Pinto, and P. Savelyev (2013). Understanding the mechanisms through which an influential early childhood program boosted adult outcomes. *American Economic Review* 103(6), 2052–86.
- Hoxby, C. M. (2000). The effects of class size on student achievement: New evidence from population variation. *The Quarterly Journal of Economics* 115(4), 1239–1285.
- Jackson, C. K. (2018). What do test scores miss? the importance of teacher effects on non-test score outcomes. *Journal of Political Economy* 126(5), 2072–2107.
- Jackson, C. K. (2020). *Does school spending matter? The new literature on an old question*. American Psychological Association.
- Jackson, C. K. and C. Mackevicius (2021). The distribution of school spending impacts. *NBER WP* (Number 28517).
- Jackson, C. K., S. C. Porter, J. Q. Easton, A. Blanchard, and S. Kiguel (2020). School effects on socioemotional development, school-based arrests, and educational attainment. *American Economic Review: Insights* 2(4), 491–508.
- Kirkebøen, L. (2021). School value-added and long-term student outcomes. Discussion Papers 970, Statistics Norway, Research Department.

- Kirkebøen, L. J., T. Gunnes, L. Lindenskov, and M. Rønning (2021). Didactic methods and small-group instruction for low-performing adolescents in mathematics. results from a randomized controlled trial. Discussion Papers 957, Statistics Norway, Research Department.
- Kirkebøen, L. J., A. Kotsadam, O. Raaum, S. Andresen, and J. Rogstad (2017). Effekter av satsing på økt lærertetthet. *Mimeo SSB*.
- Krueger, A. B. (1999). Experimental estimates of education production functions. *The Quarterly Journal of Economics* 114(2), 497–532.
- Krueger, A. B. (2003, 02). Economic Considerations and Class Size. *The Economic Journal* 113(485), F34–F63.
- Lazear, E. P. (2001). Educational production. *The Quarterly Journal of Economics* 116(3), 777–803.
- Leuven, E. and S. A. Løkken (2020). Long-term impacts of class size in compulsory school. *Journal of Human Resources* 55(1), 309–348.
- Leuven, E., H. Oosterbeek, and M. Rønning (2008). Quasi-experimental estimates of the effect of class size on achievement in norway. *Scandinavian Journal of Economics* 110(4), 663–693.
- OECD (2021). *Education at a Glance 2021*.
- Rabe, B. (2019). Do school inputs crowd out parents investments in their children? *IZA World of Labor*.
- Reiling, R. B., K. V. Salvanes, A. M. J. Sandsør, and B. Strøm (2021). The effect of central government grants on local educational policy. *European Journal of Political Economy* 69, 102006.
- Schleicher, A. (2018). Insights and interpretations. *Pisa 2018* 10.
- Todd, P. E. and K. I. Wolpin (2003). On the specification and estimation of the production function for cognitive achievement. *The Economic Journal* 113(485), F3–F33.
- Urquiola, M. (2006). Identifying class size effects in developing countries: Evidence from rural bolivia. *Review of Economics and Statistics* 88(1), 171–177.
- Woessmann, L. and M. West (2006). Class-size effects in school systems around the world: Evidence from between-grade variation in timss. *European Economic Review* 50(3), 695–736.

A Results referred to in the text

Table A.1. Compensating school inputs. Association between entry test score and student teacher ratio in the pre-treatment school years (2011-2012)

	All schools		Large municipalities
	(1)	(2)	(3)
<i>A. Student-teacher ratio</i>			
Entry test score (8th grade)	3.490** (0.309)	2.430** (0.357)	3.235** (0.527)
Mean (std dev)		15.9 (2.8)	17.9 (2.2)
<i>B. Student-teacher ratio in regular teaching</i>			
Entry test score (8th grade)	2.568** (0.364)	0.884** (0.450)	1.461** (0.620)
Mean (std dev)		20.0 (3.2)	22.0 (2.5)
Fixed effects	Year	Year*Municipality	Year*Municipality
N schools	1 574	1 574	239
N students	336 556	336 556	75 665

Note: Each cell is an estimate from a separate model that regresses resource inputs in grades 8-10 on average entry test score in 8th grade (beginning of lower secondary) for different samples and specifications. Sample of pre-treatment school years 2011-2012. For a year t , test scores are the average 8th grade scores for the years $t, t-1, t-2$. The top row shows total student hours/total teacher hours (Panel A) and the bottom row shows regular student instruction hours/regular teaching hours (Panel B). Column (1) controls for year fixed effects, columns (2) and (3) control for Year*Municipality fixed effects. In column (3) we restrict the sample to large municipalities (defined as having more than 9 schools). Standard errors clustered at school level in parentheses. ** $p < .05$.

Table A.2. Associations between school environment index and academic outcomes

	(1)	(2)	(3)	(4)
Teacher-assigned grades (10th grade)	0.319** (0.027)	0.069** (0.019)	0.235** (0.014)	0.098** (0.013)
Exam scores (10th grade)	0.349** (0.030)	0.127** (0.017)	0.180** (0.017)	0.062** (0.015)
Completion of second year of upper secondary school	0.075** (0.007)	0.045** (0.006)	0.076** (0.006)	0.058** (0.006)
Completion of upper secondary (within 5 years)	0.088** (0.009)	0.043** (0.007)	0.066** (0.012)	0.042** (0.011)
Control variables	No	Yes	No	Yes
Fixed effects	No	No	Yes	Yes

Note: Each cell is an estimate from a separate model that regresses academic outcomes on the school environment index among cohorts from 2011 to the latest year for which figures are available (2016 for exam score, 2012 for completion of upper secondary)–XX. Standard errors clustered at school level in parentheses. ** $p < .05$.

Table A.3. Robustness - regular-teaching student teacher ratio

	(1)	(2)	(3)	(4)
Treatment	-2.779** (0.196)	-2.293** (0.250)	-2.348** (0.260)	-2.105** (0.345)
STR margin	0.533** (0.195)	0.320 (0.212)	0.236 (0.232)	0.062 (0.263)
GPA margin	0.192 (0.198)	-0.020 (0.212)	0.102 (0.229)	0.019 (0.253)
Specification	Global		Local	
Forcing variables	linear	quadratic	linear	quadratic
N	3089	3089	1687	1687

Note: Data are schools 2013-2016. Specifications (1)-(2) use all schools. Specifications (3)-(4) use school within .5 SD of either threshold and triangular weights (similar to the RD analyses). Specifications (1) and (2) use linear controls for the forcing variables, while (2) and (4) use quadratic controls. All analyses are at the school level, weighted with the number of students. ** p<.05, * p<.10.

Table A.4. Robustness - 9th grade test score

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-0.073* (0.029)	0.044 (0.034)	0.004 (0.010)	0.037 (0.039)	0.092 (0.052)	0.018 (0.016)
STR margin	0.062* (0.030)	-0.006 (0.032)	0.007 (0.009)	-0.017 (0.037)	-0.052 (0.041)	0.001 (0.012)
GPA margin	0.070* (0.029)	0.013 (0.027)	0.001 (0.009)	-0.001 (0.033)	-0.025 (0.035)	-0.014 (0.012)
Specification	Global		Local			
Forcing variables	linear	quadratic	quadratic	linear	quadratic	quadratic
Background chars.			yes			yes
N	259280	259280	250230	156248	156248	151023

Note: Data are students 2013-2016. Specifications (1)-(3) use all schools. Specifications (4)-(6) use school within .5 SD of either threshold and triangular weights (similar to the RD analyses). Specifications (1) and (4) use linear controls for the forcing variables, while (2), (3), (5) and (6) use quadratic controls. Specifications (3) and (6) include controls for student characteristics (pretest, both parents' education). All analyses are at the student level, standard errors are adjusted for clustering at the school level. ** p<.05, * p<.10.

Table A.5. Robustness - exam score

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-0.077** (0.028)	0.010 (0.035)	-0.020 (0.021)	0.000 (0.040)	0.034 (0.053)	-0.032 (0.030)
STR margin	0.055 (0.030)	0.005 (0.032)	0.021 (0.018)	0.005 (0.037)	-0.019 (0.040)	0.037 (0.024)
GPA margin	0.057* (0.029)	0.015 (0.030)	0.002 (0.018)	0.015 (0.035)	0.002 (0.037)	0.010 (0.022)
Specification		Global			Local	
Forcing variables	linear	quadratic	quadratic	linear	quadratic	quadratic
Background chars.			yes			yes
N	250189	250189	241221	150892	150892	145684

Note: Data are students 2013-2016. Specifications (1)-(3) use all schools. Specifications (4)-(6) use school within .5 SD of either threshold and triangular weights (similar to the RD analyses). Specifications (1) and (4) use linear controls for the forcing variables, while (2), (3), (5) and (6) use quadratic controls. Specifications (3) and (6) include controls for student characteristics (pretest, both parents' education). All analyses are at the student level, standard errors are adjusted for clustering at the school level. ** p<.05, * p<.10.

Table A.6. Robustness - teacher grades

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-0.048* (0.022)	0.006 (0.028)	-0.031 (0.026)	0.015 (0.031)	0.072 (0.042)	-0.000 (0.039)
STR margin	0.072** (0.023)	0.040 (0.024)	0.053* (0.023)	0.016 (0.028)	-0.017 (0.033)	0.040 (0.029)
GPA margin	0.025 (0.021)	-0.002 (0.021)	-0.015 (0.022)	0.012 (0.024)	-0.014 (0.027)	-0.008 (0.027)
Specification		Global			Local	
Forcing variables	linear	quadratic	quadratic	linear	quadratic	quadratic
Background chars.			yes			yes
N	254221	254221	245048	153334	153334	148000

Note: Data are students 2013-2016. Specifications (1)-(3) use all schools. Specifications (4)-(6) use school within .5 SD of either threshold and triangular weights (similar to the RD analyses). Specifications (1) and (4) use linear controls for the forcing variables, while (2), (3), (5) and (6) use quadratic controls. Specifications (3) and (6) include controls for student characteristics (pretest, both parents' education). All analyses are at the student level, standard errors are adjusted for clustering at the school level. ** p<.05, * p<.10.

Table A.7. Robustness - on-time completion of HS year 2

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	-0.012 (0.007)	-0.002 (0.010)	-0.006 (0.008)	-0.002 (0.012)	0.007 (0.018)	0.003 (0.014)
STR margin	0.006 (0.007)	0.002 (0.008)	0.002 (0.007)	-0.003 (0.011)	-0.008 (0.013)	-0.005 (0.010)
GPA margin	0.007 (0.008)	0.003 (0.008)	0.003 (0.007)	-0.004 (0.010)	-0.008 (0.012)	-0.007 (0.010)
Specification	Global			Local		
Forcing variables	linear	quadratic	quadratic	linear	quadratic	quadratic
Background chars.			yes			yes
N	157584	157584	149891	94982	94982	90419

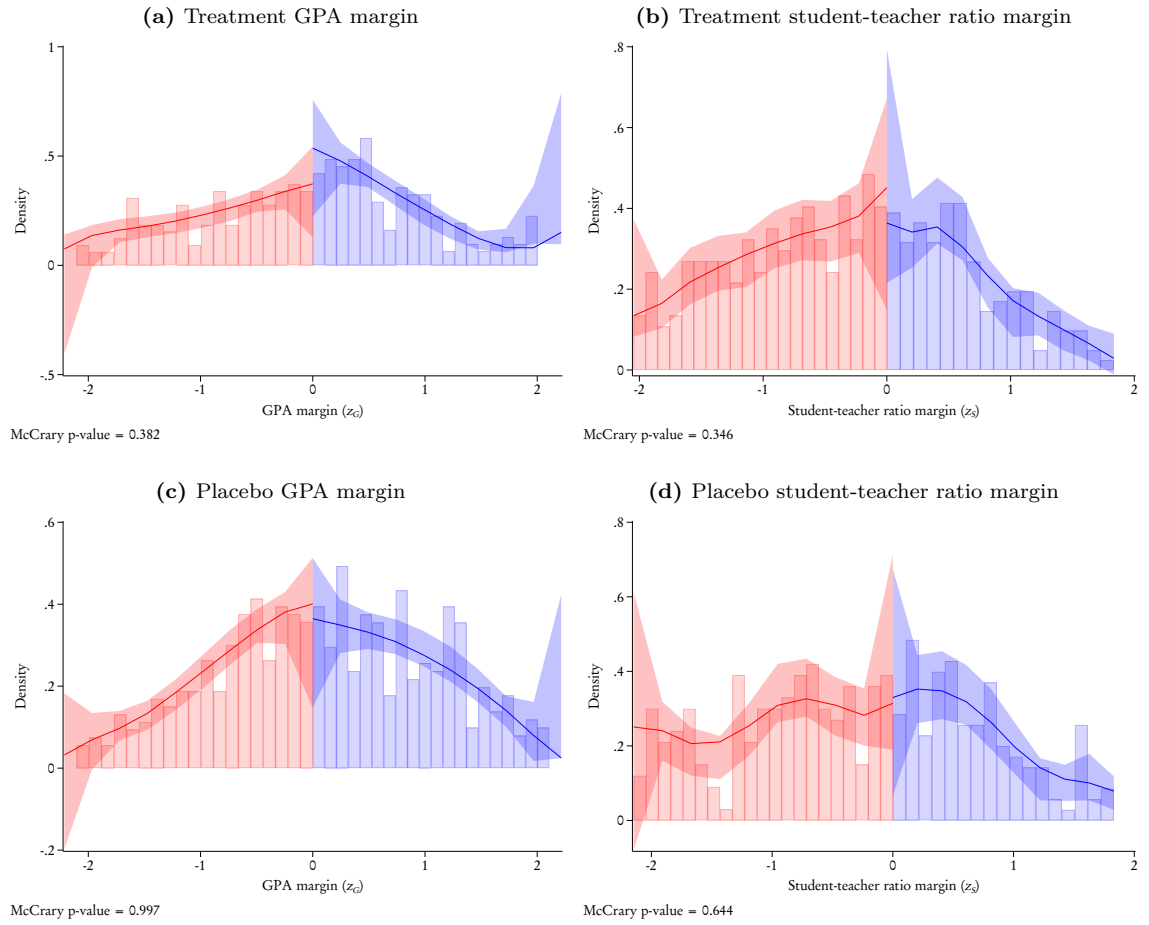
Note: Data are students 2013-2016. Specifications (1)-(3) use all schools. Specifications (4)-(6) use school within .5 SD of either threshold and triangular weights (similar to the RD analyses). Specifications (1) and (4) use linear controls for the forcing variables, while (2), (3), (5) and (6) use quadratic controls. Specifications (3) and (6) include controls for student characteristics (pretest, both parents' education). All analyses are at the student level, standard errors are adjusted for clustering at the school level. ** p<.05, * p<.10.

Table A.8. Robustness - school environment index

	(1)	(2)	(3)	(4)
Treatment	0.043** (0.014)	0.052** (0.018)	0.072** (0.018)	0.086** (0.024)
STR margin	0.006 (0.014)	-0.000 (0.015)	-0.013 (0.016)	-0.021 (0.018)
GPA margin	-0.006 (0.014)	-0.010 (0.015)	-0.015 (0.016)	-0.022 (0.017)
Specification	Global		Local	
Forcing variables	linear	quadratic	linear	quadratic
N	2831	2831	1588	1588


Note: Data are schools 2013-2016. Specifications (1)-(2) use all schools. Specifications (3)-(4) use school within .5 SD of either threshold and triangular weights (similar to the RD analyses). Specifications (1) and (2) use linear controls for the forcing variables, while (2) and (4) use quadratic controls. All analyses are at the school level, weighted with the number of students. ** p<.05, * p<.10.

Figure A.1. Balancing




Note: The graphs show RD density estimates. Data are from the 2011, the base year for assigning the treatment. Figure notes show p -values for a test of discontinuity at the threshold. Densities and p -values are estimated using `rddensity` Calonico et al. (2014), using school-level data.

Figure A.2. RD-estimates of the effect on student-teacher ratio, RMSE-optimal bandwidths




Note: The graphs show RD-estimates corresponding to (1). Data are years 2013-2016, outcome is (regular-teaching) student-teacher ratio. Figure notes show coefficients and standard errors from linear and local regressions. All analyses uses student weights. The lines show the local linear regressions and are estimated using Calonico et al. (2014), with triangular weights and RMSE-optimal bandwidth. Bias-corrected confidence interval (estimated using higher-order polynomial) in brackets. Bins are quantile-based.

Figure A.3. Effects on log student-teacher ratio



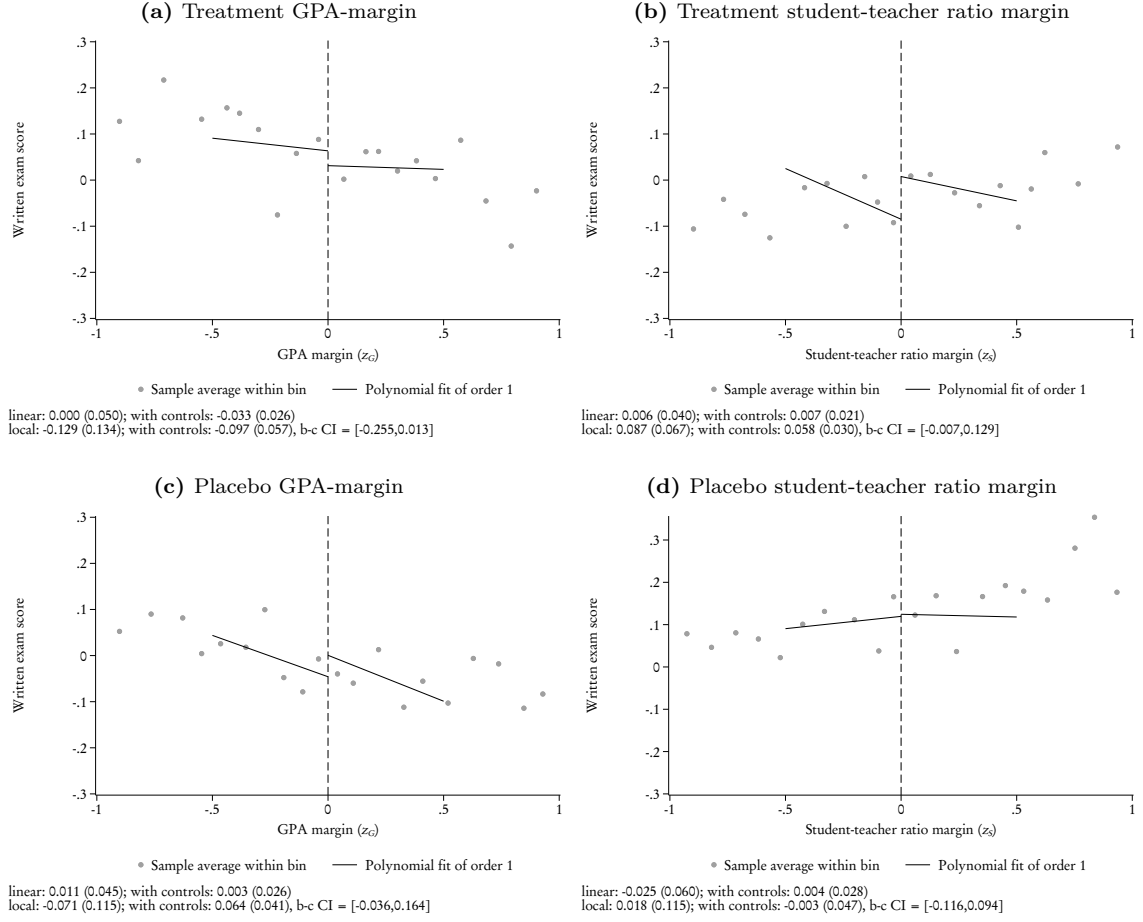
Note: The graph show estimates and confidence intervals from estimating (3). Outcome is log of (regular-teaching) student-teacher ratio. The regression uses student weights. See note to Figure 3 for further details.

Figure A.4. Effects on number of teachers



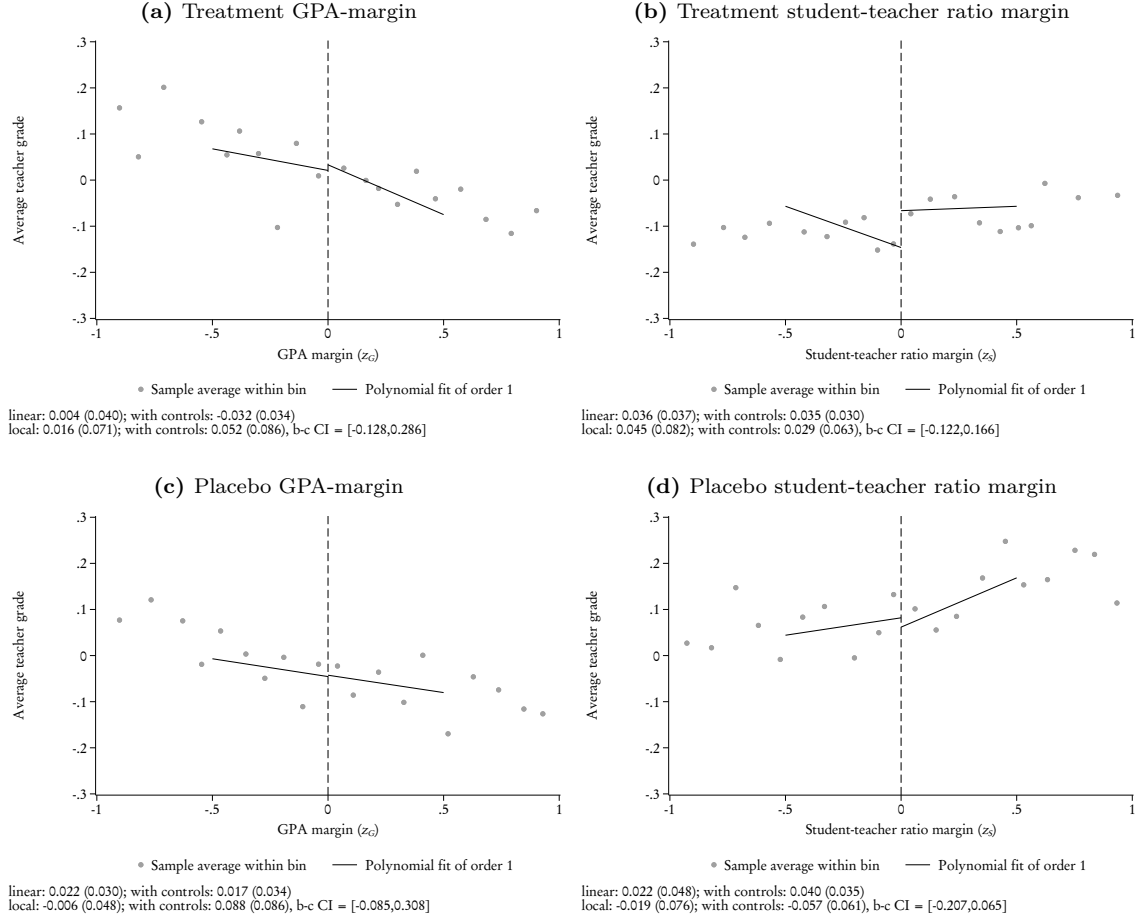
Note: The graph shows estimates and confidence intervals from estimating (3). Outcomes are number of teacher at the school and log number of teachers. The regression uses student weights. See note to Figure 3 for further details.

Figure A.5. RD estimates of effects on exam score




Note: The graphs show RD estimates for exam scores. The data are provided by students sitting the grade 9 test in 2013-2016. The graph is based on school-level data and student weights. The lines show the local linear regressions and are estimated using Calonico et al. (2014), with no additional control variables, triangular weights and a fixed bandwidth of 0.5. Bins are quantile-based. Figure notes show coefficients and standard errors from student-level linear and local regressions. Student controls in the linear and local regressions include gender, year dummies, a cubic in the 8th grade test score, and parental education. Bias-corrected confidence interval (estimated using higher-order polynomial) in brackets.

Figure A.6. RD estimates of effects on teacher grades



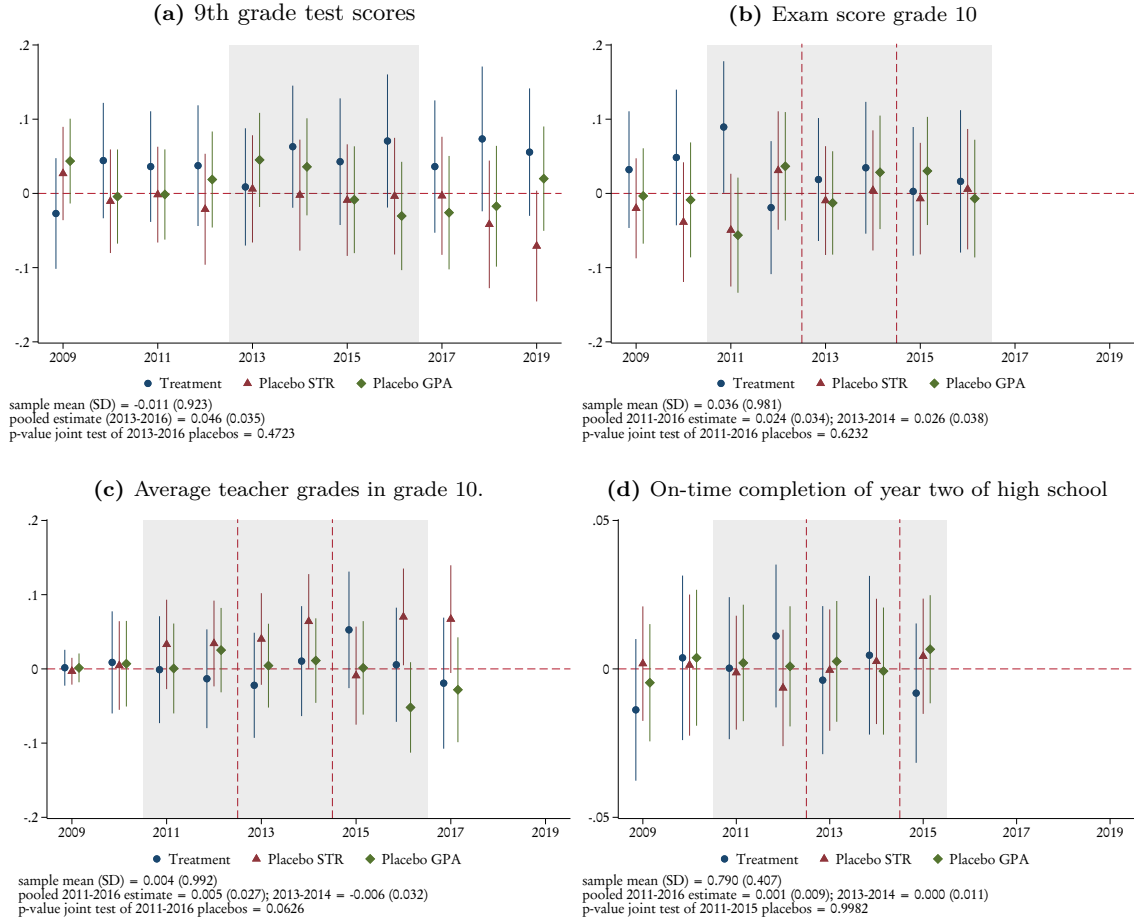
Note: The graphs show RD estimates for the average teacher grades. The data are provided by students sitting the grade 9 test in 2013-2016. The graph is based on school-level data and student weights. The lines show the local linear regressions and are estimated using Calonico et al. (2014), with no additional control variables, triangular weights and a fixed bandwidth of 0.5. Bins are quantile-based. Figure notes show coefficients and standard errors from student-level linear and local regressions. Student controls in the linear and local regressions include gender, year dummies, a cubic in the 8th grade test score, and parental education. Bias-corrected confidence interval (estimated using higher-order polynomial) in brackets.

Figure A.7. RD estimates of effects on on-time completion of year two of high school




Note: The graphs show RD estimates for on-time completion of the second year of high school. The data are provided by students sitting the grade 9 test in 2013-2016. The graph is based on school-level data and student weights. The lines show the local linear regressions and are estimated using Calonico et al. (2014), with no additional control variables, triangular weights and a fixed bandwidth of 0.5. Bins are quantile-based. Figure notes show coefficients and standard errors from student-level linear and local regressions. Student controls in the linear and local regressions include gender, year dummies, a cubic in the 8th grade test score, and parental education. Bias-corrected confidence interval (estimated using higher-order polynomial) in brackets.

Figure A.8. Effects on academic outcomes, no student level controls




Note: The graphs show estimates and confidence intervals from estimating eq. (3). Outcomes are a) 9th grade test scores, b) exam scores, c) teacher grades and d) completion of year two of upper secondary school. Control variables are year fixed effects. The different series correspond to treatment effects, δ in eq. (3), and placebo effects from the non-treatment margins, γ in eq. (3). The x-axis is the year of the 8th grade test, treated cohorts are shaded. In sub-figures b-d the dashed vertical lines indicate cohorts treated for three years. The figure notes show the sample mean and standard deviation of the outcome, estimated effect and standard errors for a pooled analysis of the treatment years and the p -value of a joint test of all placebo effects (γ) for all treatment years. The regression uses student-level data and clusters standard errors at school level.

Figure A.9. Placebo effects on pre-determined student characteristics




Note: The graphs show estimates and confidence intervals from estimating (3). Outcomes are pre-determined student characteristics. The different series correspond to treatment effects, δ in eq. (3), and placebo effects from the non-treatment margins, γ in eq. (3). The treated cohorts are shaded. The dashed vertical lines indicate the cohorts treated for three years. The figure note shows the sample mean and standard deviation of the outcome, estimated effect and standard errors for a pooled analysis of the treatment years and the p -value of a joint test of all placebo effects (γ) for all treatment years. The regression uses student-level data and clusters standard errors at school level.

Figure A.10. Heterogeneous effects on standardized test scores




Note: The graphs show estimates and confidence intervals from estimating (3) fully interacted with a binary variable. Outcome is 9th grade test scores. The different estimates correspond to treatment effects, δ in eq. (3), and placebo effects from the non-treatment margins, γ in eq. (3). The regression uses student-level data and clusters standard errors at school level.

Figure A.11. Heterogeneous effects on standardized test scores, by school characteristics




Note: The graphs show estimates and confidence intervals from estimating (3) fully interacted with a binary variable. Outcome is 9th grade test scores. The different estimates correspond to treatment effects, δ in eq. (3), and placebo effects from the non-treatment margins, γ in eq. (3). The regression uses student-level data and clusters standard errors at school level.

Figure A.12. Effects on enrollment in third year of upper secondary school



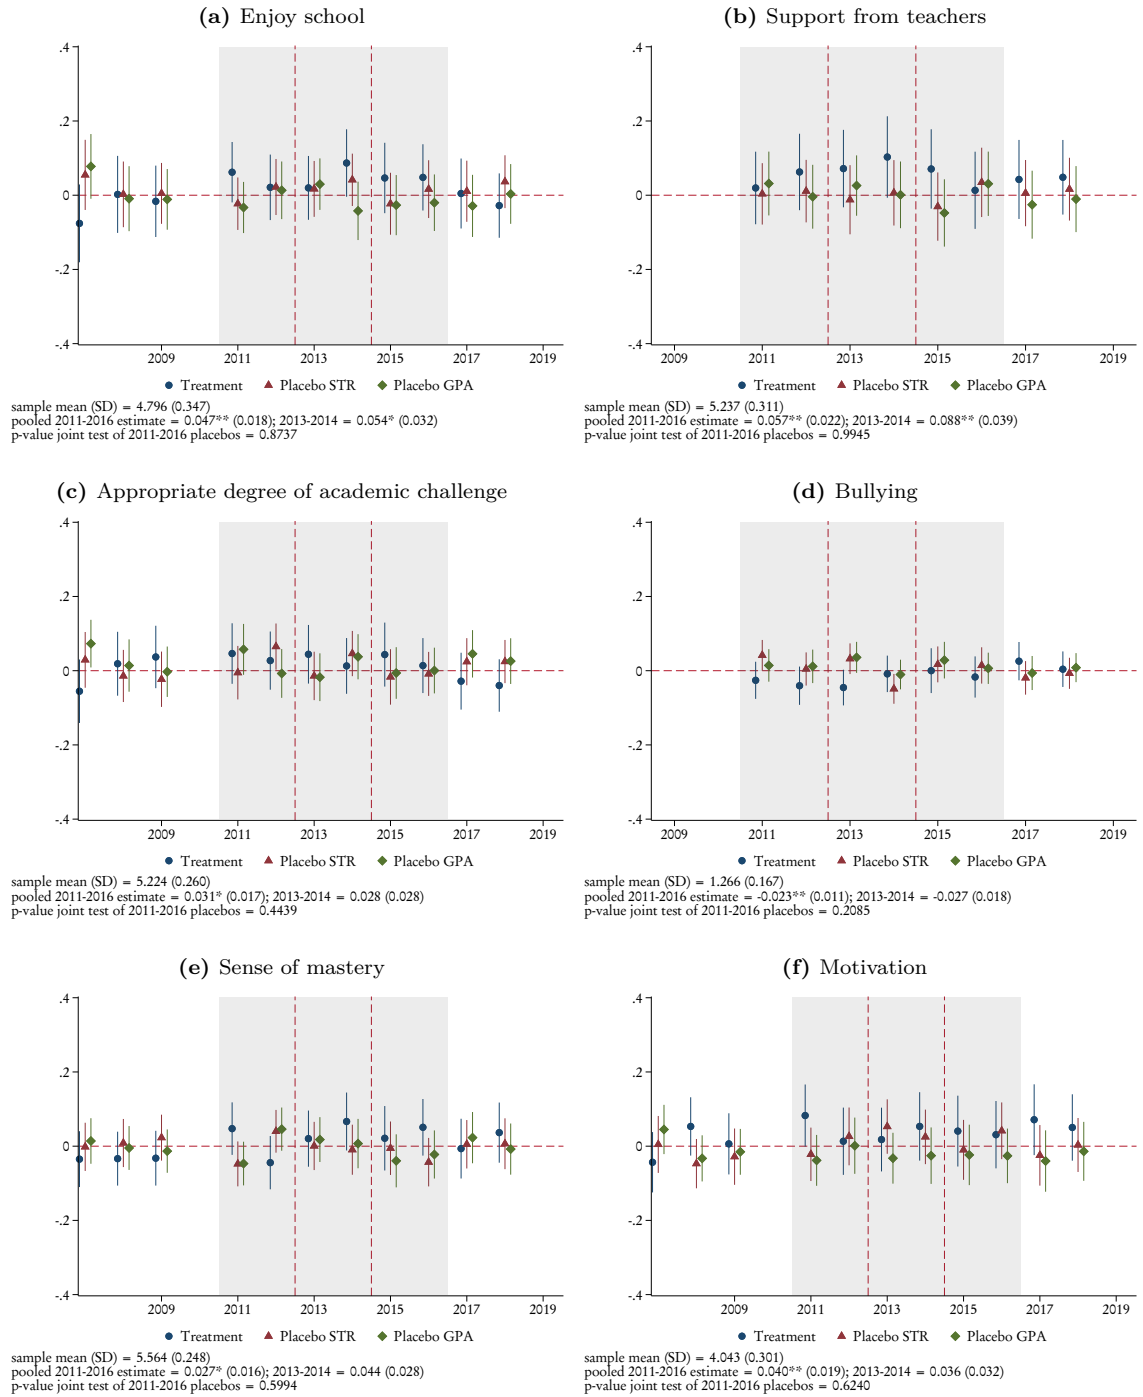
Note: The graph shows estimates and confidence intervals from estimating (3). The outcome is completion of the second year of upper secondary school and the control variables are gender, age, year fixed effects, a cubic in the 8th grade test score, and parental education. The different series correspond to treatment effects, δ in eq. (3), and placebo effects from the non-treatment margins, γ in eq. (3). The treated cohorts are shaded. The dashed vertical lines indicate the cohorts treated for three years. The figure note shows the sample mean and standard deviation of the outcome, estimated effect and standard errors for a pooled analysis of the treatment years and the p -value of a joint test of all placebo effects (γ) for all treatment years. The regression uses student-level data and clusters standard errors at school level.

Figure A.13. RD estimates of effects on school environment index




Note: The graphs show RD estimates for the school environment index for the years 2013-2016. The graph is based on school-level data and student weights. The lines show the local linear regressions and are estimated using Calonico et al. (2014), with no additional control variables, triangular weights and a fixed bandwidth of 0.5. Bins are quantile-based. Figure notes show coefficients and standard errors from linear and local regressions. Bias-corrected confidence interval (estimated using higher-order polynomial) in brackets.

Figure A.14. School environment sub indices 1




Note: The graphs show estimates and confidence intervals from estimating (3). Outcomes indices from a student survey constructed by the Norwegian Directorate of Education and Training. The regression uses school level with student weights. See note to Figure 6 for further details.

Figure A.15. School environment sub indices 2



Note: The graphs show estimates and confidence intervals from estimating (3). Outcomes indices from student survey constructed by the Norwegian Directorate of Education and Training. The regression uses school-level with student weights. See note to Figure 6 for further details.

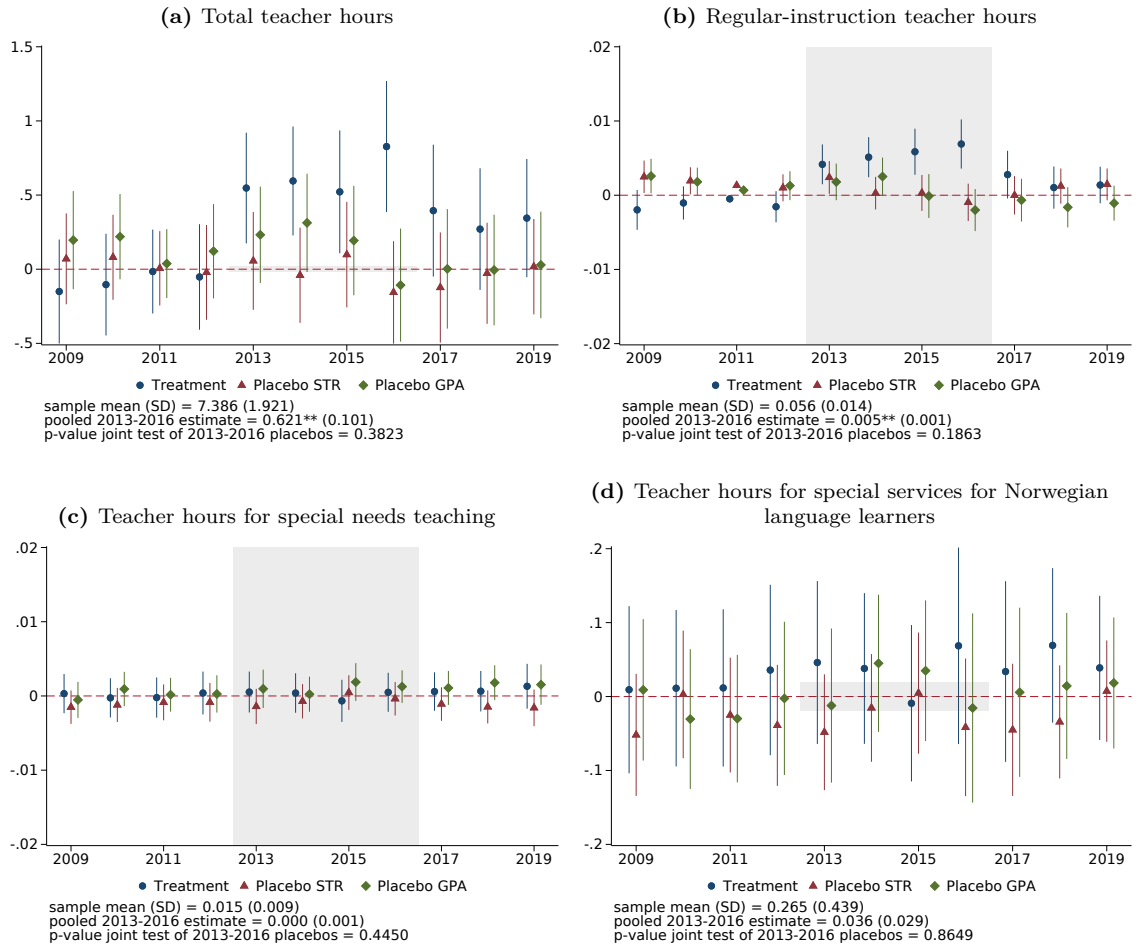
Figure A.16. Support from parents



sample mean (SD) = 4.450 (0.256)
 pooled 2011-2016 estimate = 0.012 (0.016); 2013-2014 = 0.044 (0.028)
 p-value joint test of 2011-2016 placebos = 0.9766

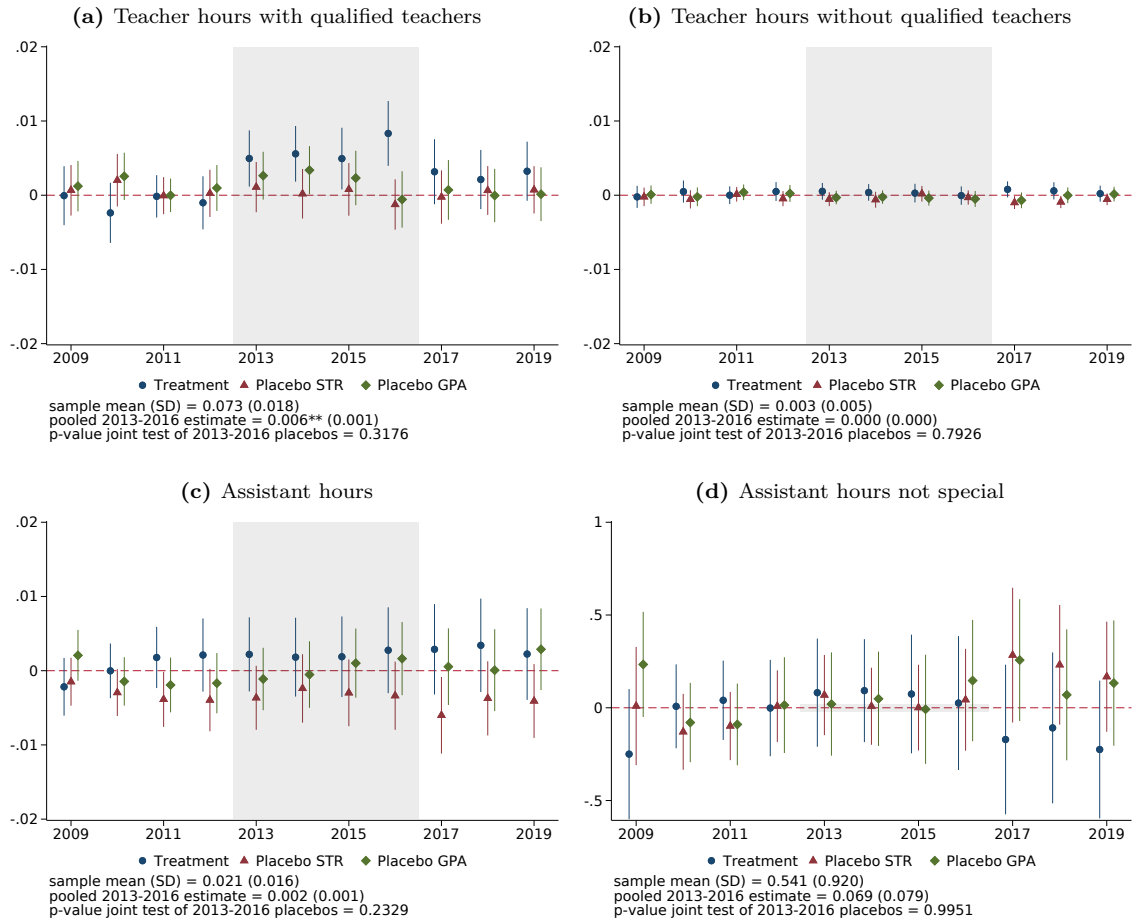
Note: The graphs show estimates and confidence intervals from estimating (3). Outcomes indices from student survey constructed by the Norwegian Directorate of Education and Training. The regression uses school level with student weights. See note to Figure 6 for further details.

Figure A.17. Effects on total teacher hours and different uses of teacher hours



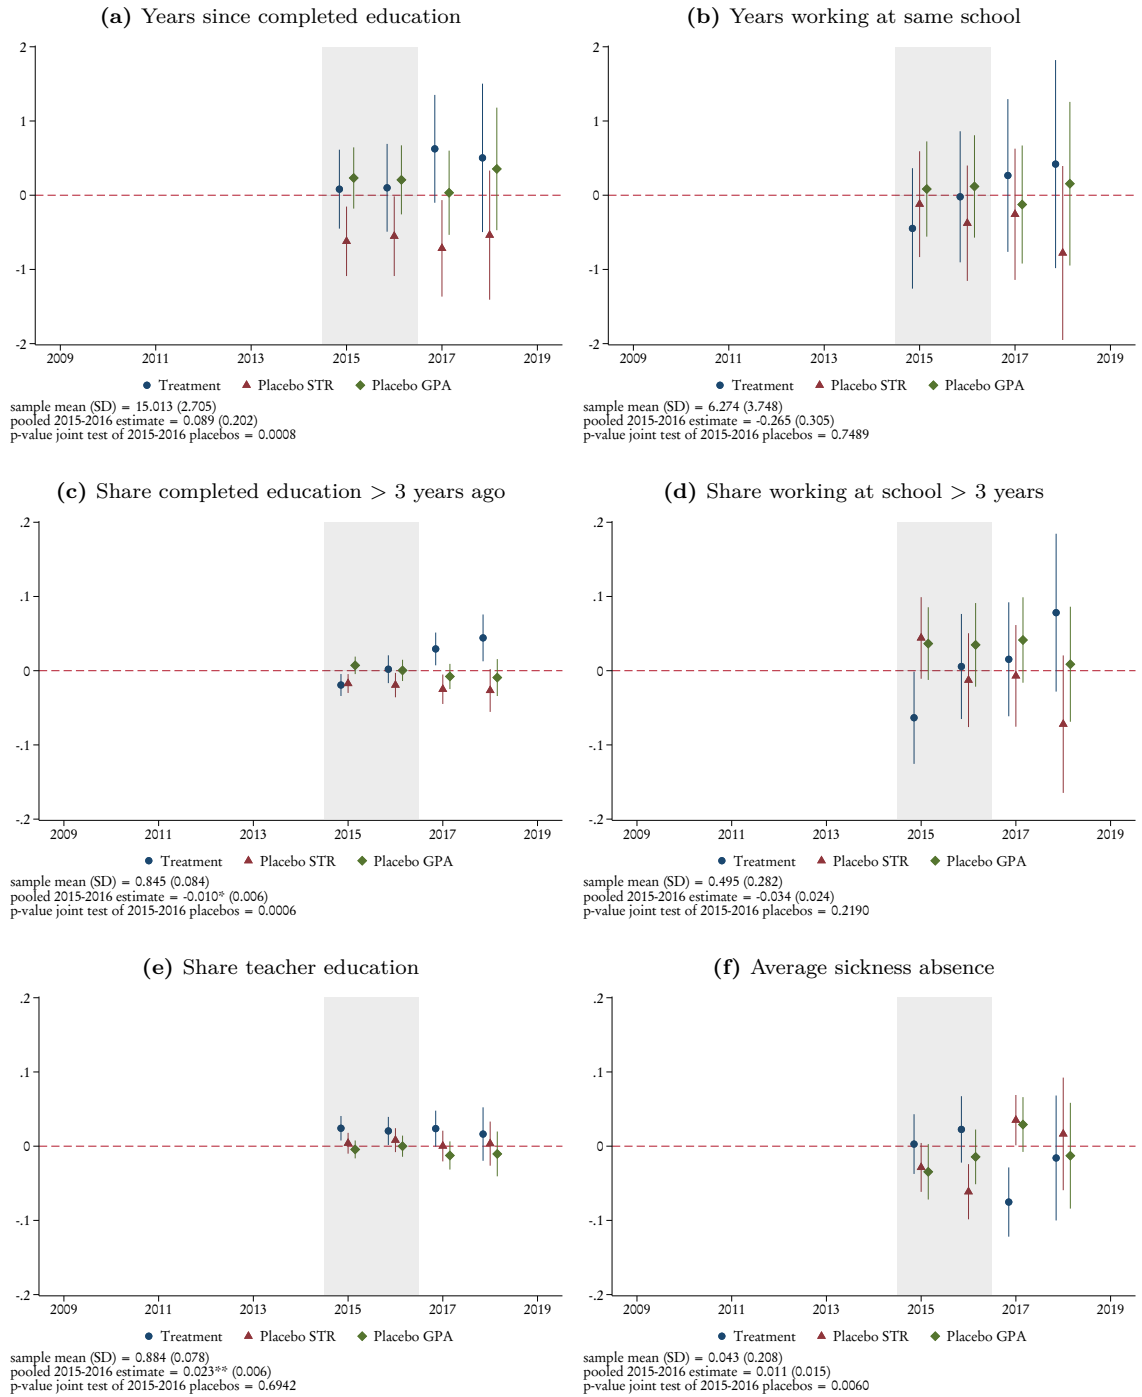
Note: The graph shows estimates and confidence intervals from estimating (3). Outcomes are total teacher hours and teacher hours decomposed by use, relative to total student hours. The regression uses student weights. See note to Figure 3 for further details.

Figure A.18. Effects on teacher and assistant hours by qualifications and use



Note: The graph shows estimates and confidence intervals from estimating (3). Outcomes are total teacher hours and assistant hours by qualifications and use, relative to total student hours. The regression uses student weights. See note to Figure 3 for further details.

Figure A.19. Effects on teachers



Note: The graph shows estimates and confidence intervals from estimating (3). Outcomes are average teacher characteristics derived from matched employer-employee data. The regression uses student weights. See note to Figure 3 for further details.