

IZA DP No. 3474

Quasi-Experimental Estimates of the Effect of Class Size on Achievement in Norway

Edwin Leuven
Hessel Oosterbeek
Marte Rønning

April 2008

Quasi-Experimental Estimates of the Effect of Class Size on Achievement in Norway

Edwin Leuven

*University of Amsterdam, Tinbergen Institute,
CEPR, CREST-INSEE and IZA*

Hessel Oosterbeek

University of Amsterdam and Tinbergen Institute

Marte Rønning

*Statistics Norway and
Norwegian University of Science and Technology*

Discussion Paper No. 3474

April 2008

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Quasi-Experimental Estimates of the Effect of Class Size on Achievement in Norway

Using a comprehensive administrative database we exploit independent quasi-experimental methods to estimate the effect of class size on student achievement in Norway. The first method is based on a maximum class size rule in the spirit Angrist and Lavy (1999). The second method exploits population variation as first proposed by Hoxby (2000). The results of both methods (and of variations on these methods) are very similar and cannot reject that the class size effect is equal to zero. The estimates are very precise; we can rule out effects as small as 1.5 percent of a standard deviation for a one student change in class size during three years in a row.

JEL Classification: I2

Keywords: class size, educational production

Corresponding author:

Edwin Leuven
Amsterdam School of Economics
University of Amsterdam
Roetersstraat 11
Room 10.06, Building E
1018 WB Amsterdam
The Netherlands
E-mail: e.leuven@uva.nl

1 Introduction

One of the still unresolved issues in education research concerns the effects of class size on students' achievement. It is by now well-understood that endogeneity problems may severely bias naive OLS-estimates of the class size effect, and that exogenous sources of variation in class size are key for a credible identification of the class size effect. Various recent studies acknowledge this and apply convincing identification methods. This has, however, not led to a definite conclusion about the magnitude or even the sign of the class size effect.

Most of the (quasi-) experimental studies report that a reduction in class size boosts achievement (Angrist and Lavy, 1999; Boozer and Rouse, 2001; Browning and Heinesen, 2007; Krueger, 1999; Urquiola, 2006). Some of the studies that attempt to correct for endogeneity do, however, report opposite results (e.g. Hoxby 2000). Moreover, the studies reporting negative effects of bigger classes vary considerably in the size of this effect thereby limiting the relevance of these results for policy conclusions.

Of course there need not be a universal effect of class size reduction on achievement. Effects may vary with characteristics of the students affected by the policy, or by contextual factors such as remedial instruction for low performing students or the quality of teachers' education (cf. Wößmann and West, 2006). This would imply that for policy purposes, studies have to be conducted for separate levels of education and for separate countries (or perhaps for groups of very similar countries).

This paper provides evidence about the effect of class size in Norwegian lower secondary schools on achievement. This evidence is obtained from two different and independent approaches. The first approach uses exogenous variation due to maximum class size rules in Norwegian lower secondary education. This approach was first used by Angrist and Lavy (1999). The second approach exploits variation in actual class size that is attributable to demographic variation. This approach was applied previously by Hoxby (2000).

Some features of our study are worth emphasizing. First, for this study we have access to an extraordinary rich dataset. The dataset covers two entire cohorts of students participating in nationwide tests in the schoolyears 2001/02 and 2002/03. The data are administrative thereby giving rise to no or only little measurement error in actual class size and enrollment. Achievement is measured as high stake test scores, thereby also reducing measurement error in the dependent variables. Together, these characteristics of the dataset enable us to produce very precise estimates of class size effects. Secondly, for all students in our sample we know actual class sizes during the three years they spent in lower secondary school. This allows a clear-cut interpretation of the effects that we estimate.

Furthermore, no previous study has presented precise estimates of the effects of class size on achievement for any of the Nordic countries (Denmark, Finland, Norway and Sweden). Previous studies have either considered other outcomes than achievement such as years of education (Bingley et al., 2005; Browning and Heinesen, 2007), or used a very small sample therefore lacking precision (Bonesrønning, 2003; Lindahl, 2005). Finally, most of the evidence on class size effects pertains to primary education, our study is among the few dealing with class size effects in lower secondary education.^{1,2}

The results reported in this paper consistently point to a lack of any impact of class size on achievement. Effects as small as 1.5 percent of a standard deviation for a one student change in class size during three consecutive years, can be ruled out. This holds irrespective of identification method (maximum class size rule or demographic variation), subject tested (math, languages) or the control variables included in the regressions. The finding also holds across various subgroups of the population and is also independent of teacher characteristics.

The remainder of this paper is organized as follows. The next section gives a brief summary of related studies. Section 3 describes the relevant institutional features of the Norwegian educational system. Section 4 gives a description of the data employed in the empirical analysis. Section 5 continues with an exposition of the empirical approaches applied in this paper and their limitations. Section 6 presents and discusses the main findings. Section 7 investigates the possibility of heterogeneous class size effects, and Section 8 concludes.

2 Related literature

During more than a decade, the common wisdom among economists was that extra resources for education - measured as the teacher-pupil ratio or as expenditures per pupil - have no systematic relation with students' achievement. This view was mainly based on Hanushek's (1986) influential review of the literature. Only recently has this received wisdom been challenged by a series of studies that use experimental and quasi-experimental approaches to identify the causal impact of class size on achievement.

Krueger (1999) analyzes data from a large-scale field experiment conducted in Tennessee. Students and their teachers were randomly assigned to a group of regular size (22-25 students), to a group of regular size including a teaching assistant, or to a small group (13-17), during their first four years in school. Krueger's findings are in line with

¹For a subsample of our schools we can also conduct analyses using class size in primary school as a class size measure. We have chosen not to present these results since they are very similar to the results for lower secondary schools.

²Häkkinen et al. (2003) is a study at the lower secondary level for Finland, but they look at a broader measure of school resources.

what other people have reported about this project namely that students in smaller classes perform better on standardized achievement tests. Scores increase by four percentile points for the first year that a student is exposed to a small class and by one percentile point for each subsequent year. In a follow-up study, Krueger and Whitmore (2001) demonstrate that reduced class sizes in early school years can have long-lasting effects. Students who attended small classes in this experiment are more likely to take a college-entrance exam and have somewhat higher test scores. The effect on exam taking are mainly concentrated among minority students.

Angrist and Lavy (1999) were the first to exploit the exogenous variation generated by maximum class size rules to obtain quasi-experimental estimates of the class size effect on achievement. They exploit the fact that according to official guidelines for Israeli public schools, maximum class size equals 40. If the size of an enrollment cohort in a school exceeds (a multiple of) 40, an extra class should be created. This rule creates discontinuities in the relation between cohort enrollment size and class size, which Angrist and Lavy then use in a regression discontinuity framework to identify the effect of class size on achievement.³ When they do not correct for endogeneity bias, their estimates point to a positive relation between class size and achievement. In contrast to that, estimates based on the discontinuities in grade enrollment point to a negative effect of class size on achievement.

Other papers exploiting maximum class size rules include Bonesrønning (2003) for Norway, Urquiola (2006) for Bolivia, Piketty (2004) and Gary-Bobo and Mahjoub (2006) for France, Browning and Heinesen (2007) and Bingley et al. (2005) for Denmark, and Wößmann (2005) for Denmark, France, Germany, Greece, Iceland, Ireland, Norway, Spain, Sweden and Switzerland. Especially the papers by Bonesrønning and Wößmann are of interest for the analysis presented in this paper because they also deal with class size effects in Norway. Bonesrønning uses data from a small-scale self-collected dataset. Class size effects are estimated using the discontinuity at 30 as an instrument for actual class size. The reported effects depend on the exact specification but vary between 0.13 and 0.26 of a standard deviation for a 10-students reduction in class size, and are significantly different from zero but not very precisely estimated. The main differences between this study and ours are the following: (i) we apply two different methods rather than only one; (ii) we employ a much larger dataset, so that our estimates have much more precision; (iii) we use high-stake exam scores as our achievement measure (iv) in our regression discontinuity specifications we include controls for enrollment in a grade, something that Bonesrønning does not. Without controls for enrollment, the estimate of the class-size effect will also

³In section 5 we will discuss their identification strategy in more detail because we will apply the same method in this paper.

pick up effects of enrollment.

Wößmann uses data from 38 Norwegian schools (1,351 pupils) who participated in TIMSS. Without controls for enrollment a small but significantly negative estimate of the class size effect is found which vanishes once controls for enrollment are included in the specification. Interestingly, Wößmann uses a cutoff of 28 rather than on 30 arguing that gives the best fit to his data so that it seems that this implicit rule is actually used in most Norwegian schools. This “letting the data decide” approach is at odds with the basic philosophy underlying the regression discontinuity approach because assignment to treatment and control is no longer based on an exogenous rule but on schools’ choices.

Hoxby uses demographic variation to identify the class size effect. She exploits that - after correcting for a trend - cohort sizes within school districts can be larger or smaller in some years than in others.⁴ Using data of elementary school pupils in the state of Connecticut, she does not find any statistically significant effect of class size on student achievement, and her estimates are precise enough to rule out even modest effects.

Two recent papers apply identification strategies in the same spirit as Hoxby’s approach, although both studies do not control for a trend in cohort size. Urquiola (2006) identifies the class size effect using variation in population size between school in rural areas that are so small that fewer than 30 students (a number that fits in one class) are enrolled. These schools are more likely to be in small communities where class size is mainly determined by cohort size. Urquiola uses data from third graders in Bolivia and finds significantly negative effects of class size on achievement.

Wößmann and West (2006) exploit within-school differences in average class size between adjacent grades. An attractive feature of their study is that it uses data from students in 11 different countries, although this comes at the price that numbers of (identifying) observations are in some cases rather small and precision therefore low. Sizeable positive effects of smaller classes are reported for Greece and Iceland, for four countries even small class size effects can be ruled out, and for another four countries large beneficial effects can be ruled out. As an explanation for the differences in class size effects across countries, Wößmann and West (2006) advance the hypothesis that smaller classes are only beneficial where/when the average capability of the teaching force appears to be low.

Finally, Lindahl (2005) implements a value added approach to estimate the effect of class size using Swedish data. He fails to find statistically significant class size effects using standard value added methods. When he identifies the class size effect by taking the difference between school- and summer-period test score changes, he does find that class size matters. We are not able to implement this approach because we do not have repeated achievement measures for the students in our sample.

⁴This method will also be discussed in more detail in section 5.

The variation in the findings reviewed here is reflected in the controversy between two of the leading education economists, see: Krueger (2003) versus Hanushek (2003). In an attempt to reconcile both views, Todd and Wolpin (2003) stress the difference between policy effects and parameters of the education production function. According to these authors, estimates of the class size effect obtained from experimental and quasi-experimental research designs should be interpreted as policy effects, whereas estimates obtained from non-experimental research designs are aimed at the identification of the education production function. To learn about the production technology, one needs exogenous variation in class size *holding other inputs constant*. To learn about the policy impact, one needs exogenous variation in class size *not holding other inputs constant*. The estimates presented in this paper are therefore probably best understood as policy effects. An important other input not controlled in experimental and quasi-experimental studies is the input of parents. Parents may for instance respond to a reduction in class size by spending less time teaching their children at home. In that case school and parental inputs are substitutes and the policy effect will be smaller than the technological effect. In principle, school and parental inputs can also work as complements in which case the policy effect would exceed the technological effect. A similar line of reasoning holds with respect to for example teacher effort. Hægeland et al. (2007) argue that the maximum class size rule is accompanied by a similar input substitution in terms of school resources. We will show that this is not a concern for our findings. When we replace class size by a commonly used measure for school resources in Norway we find significant differences around the discontinuities and our results for achievement are unchanged.

This brief review of related studies only includes studies by economists. For a recent review of the class size literature from a non-economic perspective (although references to most of the studies cited above are included) see Hattie (2005). His reading of the literature is that class size effects are often very small, and he proposes as a candidate explanation that teachers tend to use the same teaching methods independent of class size.

3 Institutional Settings

Compulsory schools in Norway are run and owned by the 435 municipalities.^{5,6} Municipalities receive funding to run their various activities (including schools) through a combination of a local income tax, property taxes and transfers from the central government (see f.e. Hægeland et al. (2007) for more details).

⁵In Norway the names local school district, local government and municipality are interchangeable.

⁶In addition to compulsory public schooling local governments are responsible for elderly care, preschool education, and infrastructure. Spending on education consists of about 30% of total spending of the available budget.

For the students in our sample compulsory schooling consisted of nine years; grades 1 to 6 in primary school; and grades 7 to 9 in lower secondary school.⁷ Less than three percent of the students are enrolled in private schools. Thus private schools do not provide a realistic alternative to public schools and are therefore dropped from the analysis. Schools have catchment areas, implying that parental school choice between schools for given residence is not allowed. Most students go to separate primary and lower secondary schools, but because of the rural settlement pattern in Norway about 23 percent of the students are enrolled in so-called “combined” schools. These are schools that offer both primary and lower secondary education. Students who start taking their primary education in a combined school typically continue to take their lower secondary education there because of lack of school choice for given residence and moving costs. Combined schools are often situated in relatively rural areas. If total enrollment is less than 40 schools are allowed to mix grades in classrooms. We exclude these schools from the analysis (these schools account for 3.2 percent of the student population).

Although we are measuring class size effects in lower secondary school (instead of primary school as most studies do), it is important to note that students in lower secondary schools in Norway are assigned to the same class during the whole day and year. Each class is taught by different teachers for different subjects; the teachers rotate among classes.

Another feature of the Norwegian compulsory school system is that grade retention almost never happens. Strøm (2004) attributes this to “the strong integration and equalizing policy that all students within a cohort should be treated equal, and be given education in their ordinary classes.” As a result, at the end of compulsory schooling all Norwegian students have identical length of schooling. This is important for our analysis since we have data at the student level of the nationwide tests conducted in 2001/2 and 2002/3 and of class size from these years as well as from previous years. Because students do not repeat grades we know the class size history of individual students during their stay in lower secondary school, provided that they did not change school. We are not aware of any study documenting the extent of students’ school mobility in Norway. Hægeland et al. (2007), who use the same student data as we do, report that 95.3 percent of the students have lived in their graduation municipality throughout all 3 years in lower secondary schools. While this does not prove low mobility across schools, it is certainly not inconsistent with it. Notice furthermore that even selective movements are not a problem when we instrument average class during 3 years in lower secondary school by predicted class size in 7th grade.

⁷From the school year 1997/98 onwards it became compulsory to start schooling at the age of 6 (instead of 7), and from then on 10 years of schooling was implemented in Norway. The reform was implemented in such a way that the length of primary school was extended with one year.

4 Data

4.1 Test scores

We use administrative enrollment data from Statistics Norway that cover all students who were in the final grade of lower secondary school (9th grade) for the school years 2001/2002 and 2002/2003. We merged this dataset with test score data from centralized exit exams (also from Statistics Norway). The students in Norway have to sit such exams at the end of their final school year in lower secondary.⁸ Their results on this exam are important for further schooling possibilities (which is upper secondary education), and the exam is considered to be a high stake test by all parties involved; students, their parents, teachers and school administrators.⁹ Although the curriculum includes many different subjects a written centralized exam is only undertaken in three subjects: mathematics, English and Norwegian.¹⁰ To reduce the administrative burden each student takes the exam in one subject only. In which subject a student will take the exam, is determined by a random device shortly before the exams take place so that students have minimal scope to prepare for that specific subject only. Students are told only three school days in advance in which subject they are going to make the exam. All students in the same class do their exam in the same subject, but students in different classes of the same school may sit their exam in different subjects.

Each exam is graded by two independent external examiners, and a third one in case of disagreement between the first two. Examiners receive detailed guidelines for the grading of exams from the Ministry of Education. Each examiner grades 100-120 exam papers, in most cases from different schools. Examiners do not know the name of the pupils, but do know the name of the school(s). The grading of the exams is on a scale running from 1 to 6, where 1 is fail, 2 the lowest pass and 6 is the top score. It is explicitly mentioned in the grading instructions that examiners should not normalize exam grades, neither within nor between schools. The distributions of test scores have a bell curve shape, and there are no signs of floor or ceiling effects (low frequencies for 1 and 6). Average scores for each of the four subjects are around 3.5 with standard deviations almost equal to 1.

⁸Although this exit exam existed already for many years, its results became available for research purposes for the first time in the schoolyear 2001/02.

⁹Although all students have the right to continue at upper secondary schools and above 95 percent do so, their choice set among different schools and different study tracks depends on their achievement in lower secondary schools.

¹⁰In Norway there are two official written languages, main Norwegian (Hovedmål) and a second Norwegian language (Sidemål), and students who are examined in Norwegian get a score on both languages.

4.2 *Class size*

Class size information comes from the Norwegian Ministry of Education (Grunnskolen Informasjonssystem) which registers, for all schools in Norway, the number of classes per grade. We combine this information with the enrollment data mentioned above to calculate for each student average class size at the grade level as

$$\text{class size} = \text{enrollment} / \# \text{ classes.}$$

Note that we have data on average class size per grade and not actual class size (except when schools have exactly one class in the grade). It is important to note that this eliminates biases resulting from within school sorting, while the associated measurement error is removed by our 2SLS approach.

An attractive feature of the data is that we do not only have information on contemporaneous class size (the school year of the exam) but also on class size during the previous years in lower secondary school. Unfortunately class sizes across years within the same school are too highly correlated to examine their separate impacts (the correlations are always higher than 0.9). We therefore have chosen to estimate the effect of class size as the average class size during the three years of lower secondary school. In this way we avoid confounding the impact of class size in grade 9 with the impact of class size in earlier grades. Moreover, impact estimates of class size defined this way are relevant from a policy point of view. Hoxby (2000) also focuses on the average class size a cohort has experienced up until the time it takes the test.

Figure 1 shows the distribution of average class size in lower secondary schools. It shows that the majority of average class sizes per grade is between 20 and 30 students. A substantial fraction (25%) of the classes does, however, have 21 students or less. The (unweighted) average class size is equal to 23.3 (s.d. 4.1). The graph also suggest that the maximum class size of 30 is enforced since no class has a size exceeding the threshold of 30. At the grade level average class size exceeds 30 in only 22 out of 5,032 cases.

4.3 *Control variables*

From other administrative databases we merged information on students' age, gender, ethnic minority background, household income, whether parents live together or not, and years of education of both parents. This information pertains to the same year as the exam scores (2001/2002 and 2002/2003). Comparable information is available and exploited in most (but not all) other studies examining the effect of class size on achievement.

We also control for a number of teacher characteristics. We use employer register data on the teachers from the Ministry of Labor and Government Administration. Since we are

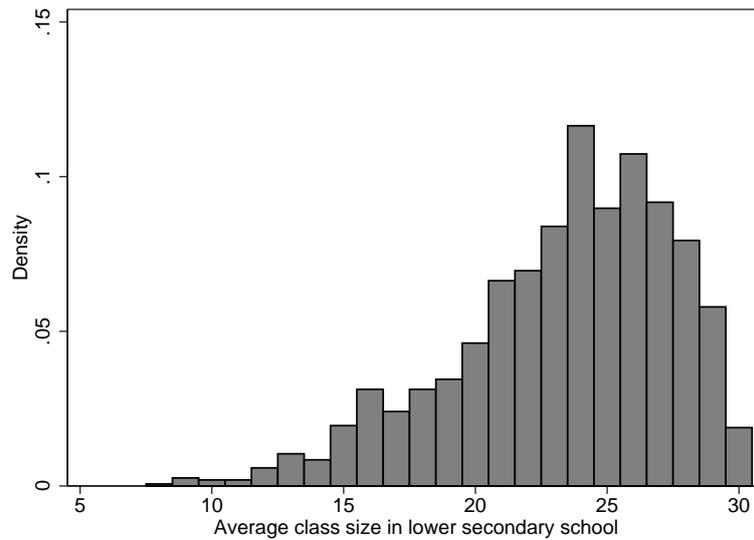


Figure 1: Distribution of class size

not able to link teachers and students we aggregate the teacher data up to the school level where we weigh with the "workload per teacher". Included teacher controls are experience, gender, temporary contract and years of schooling. Finally we include the (log) size of the school district measured in terms of inhabitants and the number of people in the school district who live in rural areas. We also control for whether schools are combined schools.

Table 1 presents descriptive statistics of the control variables. Half of the students is female, the average age is 14.5 years, the average levels of fathers' and mothers' education are almost equal, only 5 percent of the sample consists of students with an immigrant background, and 31 percent of the students do not live in a household with both parents present. Almost a quarter of the students attends a school combining primary and lower secondary education, (weighted by numbers of students) average class size equals 24.8 and enrollment in the first grade of lower secondary school (7th grade) equals 87.4. Teachers have on average 4.6 years of teacher training, 18.9 years of work experience as a teacher, 56 percent of the teachers is women and 17 percent of the teachers is on a temporary contract (again all weighted by numbers of students).¹¹

Class size is not distributed randomly in the population because of sorting of students and teachers, and the targeting of educational resources. This is illustrated by Table 2 which reports the results of regressions of actual average class size in lower secondary school on individual, teacher and school characteristics. The first column reports the results of a regression that only includes students' characteristics. This shows that actual class size increases with family income, parents' levels of education, being immigrants, and with

¹¹ Student characteristics are measured at the moment of testing, whereas school and teacher information relates to October 1st of the school year.

Table 1: Sample summary statistics

	Mean (1)	s.d. (2)
<i>Individual characteristics</i>		
Girl	0.49	(0.50)
Age	14.53	(0.31)
ln(Family income)	13.24	(0.76)
Education mother (years)	11.94	(2.90)
Education father (years)	12.23	(3.15)
1st or 2nd generation immigrants	0.05	(0.22)
Parents non-cohabiting	0.31	(0.46)
ln(pop. size school district)	10.11	(1.44)
ln(rural pop. size school district)	7.97	(0.74)
<i>School characteristics:</i>		
Combined school	0.18	(0.38)
Average class size grades 7-10	24.80	(3.24)
Enrollment grade 7	87.43	(40.13)
<i>Teacher characteristics (year t)</i>		
Average teacher education (years)	4.62	(0.19)
Average teacher experience	18.86	(3.24)
Fraction of female teachers	0.56	(0.11)
Fraction of teachers with a temp. contract	0.17	(0.11)
Year = 2002	0.52	0.50
N	111,463	
N schools	781	

Table 2: Regression of actual average class size in lower secondary school on observables

	(1)	(2)
Girl	0.001 (0.024)	0.004 (0.022)
Age	-0.039 (0.041)	0.065 (0.037)*
ln(family income)	0.296 (0.036)***	0.112 (0.022)***
Education mother (years)	0.039 (0.007)***	0.014 (0.006)**
Education father (years)	0.064 (0.008)***	0.018 (0.005)***
1st or 2nd generation immigrants	0.859 (0.145)***	-0.258 (0.122)**
Parents non-cohabiting	0.185 (0.037)***	-0.007 (0.032)
ln(pop. size school district)		0.893 (0.064)***
ln(rural pop. size school district)		0.202 (0.131)
Year=2002		0.105 (0.130)
adj. R-squared	0.015	0.173
N	111,463	111,463
N schools	781	781

Note: Standard errors are heteroscedasticity robust and corrected for school level clustering. ***/*** statistically significant at the 10/5/1 percent level.

parents not living together. When we add to this specification characteristics of the district, the positive effects of family income and parents' education remain significant but become smaller. The effect of having an immigrant background reverses and the effect of living in a household with both parents present is no longer significant (and also changed sign). These latter results indicate that immigrant families and separated parents are concentrated in more densely populated districts. The results in Table 2 make clear that class size is not randomly distributed across students. Students from more affluent families typically attend larger classes, suggesting that class size reduction is used as a compensatory policy. Ignoring selective placement in small and large classes would therefore most likely lead to an underestimation of the true effect of class size reduction on achievement. This calls for a strategy that addresses these problems. The next section discusses in detail the approaches that we will pursue.

5 Empirical approaches

We follow the literature and assume that achievement of student i (y_i) is generated by the following equation:

$$y_i = x_i' \beta + w_{s(i)}' \alpha + \delta \cdot \overline{cs}_{s(i)} + \eta_{s(i)} + \psi_{t(i)} + \varepsilon_i \quad (1)$$

x_i is a vector of observable attributes of the student and his parents, $w_{s(i)}$ a vector of observable school and teacher characteristics and $s(i)$ identifies the school of pupil i , $\overline{cs}_{s(i)}$ is the average class size that student i attended during her school career in school s , $\eta_{s(i)}$ a school effect, $\psi_{t(i)}$ is an effect for the year in which student i is in her final year of lower secondary school (2001/2 or 2002/3) and ε_i is all other determinants of achievement such as unobserved attributes of the student, parents and community.¹² The coefficient of interest is δ , the class size effect. Note that a value-added specification is not feasible because achievement is only measured at the end of lower secondary school.

5.1 Conditioning on observables

Table 3 shows the relation between student achievement and class size estimated using OLS. Results are presented from various specifications and separately for mathematics and languages, where we have pooled English and the two Norwegian subjects (and included subject dummies).¹³ Columns (1) and (4) are obtained from a specification without covariates; both estimates are positive and significant, indicating that pupils in larger classes perform better than pupils in smaller classes. The results in columns (2) and (5) are obtained from a specification that includes individual characteristics as controls, while columns (3) and (6) report the results from specifications that also include school and teacher characteristics. We control for school district characteristics in specifications (2), (3), (5) and (6). Including controls produces small negative but statistically insignificant estimates on class size in all specifications. If conditioning on a rather rich set of observables is sufficient to correct for biases of selective placement into (schools with) different class sizes, we can say that a one pupil reduction in class size improves test scores by no more than 0.8 percent of a standard deviation with 95% probability.

Not all selection into classes of different size needs, however, to be on observables. To address any remaining endogeneity of class size we need variation in actual average class size that is arguably not subject to the choices of parents and schools' principals or teachers. We exploit two sources of such exogenous variation, one induced by a maximum

¹²In the estimations we allow for clustering of this error term at the school level.

¹³We pooled the languages to reduce the number of outcomes to be presented and discussed. Findings are, however, very similar for the various languages separately.

Table 3: OLS. Dependent variable is exam results in 9th grade, unit of observations is the individual student's exam result

	Mathematics			Languages (pooled)		
	(1)	(2)	(3)	(4)	(5)	(6)
Average class size grade 7-9	0.010 (0.004)**	-0.002 (0.003)	-0.002 (0.003)	0.009 (0.003)***	-0.001 (0.002)	-0.001 (0.002)
<i>Individual characteristics</i>						
Girl		0.015 (0.012)	0.015 (0.012)		0.506 (0.010)***	0.505 (0.010)***
Age (years)		0.069 (0.020)***	0.070 (0.020)***		0.031 (0.015)**	0.030 (0.015)**
ln(Family income)		0.179 (0.009)***	0.180 (0.009)***		0.143 (0.008)***	0.143 (0.008)***
Education mother (years)		0.075 (0.002)***	0.075 (0.002)***		0.057 (0.002)***	0.057 (0.002)***
Education father (years)		0.066 (0.002)***	0.066 (0.002)***		0.049 (0.002)***	0.049 (0.002)***
1st or 2nd generation immigrants		-0.442 (0.032)***	-0.434 (0.032)***		-0.229 (0.025)***	-0.224 (0.024)***
Parents have different address		-0.347 (0.013)***	-0.345 (0.013)***		-0.186 (0.009)***	-0.185 (0.009)***
ln(pop. size school district)		0.015 (0.008)*	0.025 (0.008)***		0.009 (0.005)	0.013 (0.006)**
ln(rural pop. size school district)		-0.019 (0.016)	-0.025 (0.016)		-0.014 (0.010)	-0.017 (0.010)*
<i>School characteristics:</i>						
Average teacher education (years)		0.035 (0.057)	0.035 (0.057)		0.061 (0.043)	0.061 (0.043)
Average teacher experience		0.015 (0.003)***	0.015 (0.003)***		0.010 (0.002)***	0.010 (0.002)***
Fract. of female teachers		-0.095 (0.116)	-0.095 (0.116)		0.049 (0.080)	0.049 (0.080)
Fract. teachers with a temp. contract		0.072 (0.090)	0.072 (0.090)		0.022 (0.067)	0.022 (0.067)
Combined school		0.049 (0.028)*	0.049 (0.028)*		0.024 (0.023)	0.024 (0.023)
R-squared	0.001	0.165	0.167	0.012	0.173	0.174
N pupils	36,915	36,915	36,915	74,548	74,548	74,548
N schools	608	608	608	752	752	752

Note: Standard errors are heteroscedasticity robust and corrected for school level clustering. */**/**: statistically significant at the 10/5/1 percent level.

All regressions include a control for year of observation; the language regressions also include dummies for the language the student took the exam in.

class size rule and one based on population variation. The next subsections describe what these methods entail and how these can be implemented in the Norwegian context.

5.2 *Maximum class size rules*

Lower secondary schools in Norway are subject to maximum class size rules of 30 students per class. This rule creates a discontinuous relation between enrollment and class size. Just above multiples of 30 class size drops substantially. Following Angrist and Lavy (1999) we exploit this maximum class size rule in a regression discontinuity design. For this approach to work, schools need to be located randomly around the thresholds and no other discontinuities that may affect outcomes should exist.

Although identification in the regression discontinuity design is ultimately local (e.g. Hahn et al., 2001), Angrist and Lavy also proposed to instrument actual class size with predicted class size while conditioning on a smooth function of enrollment which is supposed to capture any direct effect of this variable on achievement. This identifying assumption essentially boils down to an exclusion restriction with respect to the discontinuities. In the analyses we will control for a cubic function of enrollment. An alternative for controlling for a smooth function of cohort enrollment is to restrict the sample to the regions around the kinks. We will present separate results from analyses for a sample which is restricted to schools with cohort enrollment in grade 7 at most 5 students away from the kinks and will refer to this regression discontinuity sample as ' $DS \pm 5$ '.

Our analysis differs from previous analyses because we use predicted class size in grade 7, the first year in lower secondary school, to instrument average class size during the three years in lower secondary school. The reason behind this is that cohort enrollment in grades 8 and 9 and thereby predicted class size in these grades may depend on actual class size in grade 7, and is therefore potentially endogenous. Such dependence could for instance result from parents' decisions to move from schools where they experienced large classes in 7th grade to schools where they observed small classes in 7th grade.¹⁴ Although we reported evidence above which suggests that student mobility during lower secondary education is limited, our estimates will not be affected by selective mobility if it is orthogonal to our instrument.

Figure 2 plots predicted class size and average actual class size against cohort enrollment. Average actual class size closely tracks predicted class size especially around the first kink. Table 4 reports the results from the corresponding first stage regressions, for specifications with different sets of control variables, including polynomials of enrollment.

¹⁴For students in combined schools enrollment in grade 7, the grade when they enroll in lower secondary school, may depend on actual class size during the previous primary school period. But the fact that the maximum class size rule changes from 28 in grade 6 to 30 in grade 7 causes that also in the combined schools students in grade 7 are confronted with a new class size.

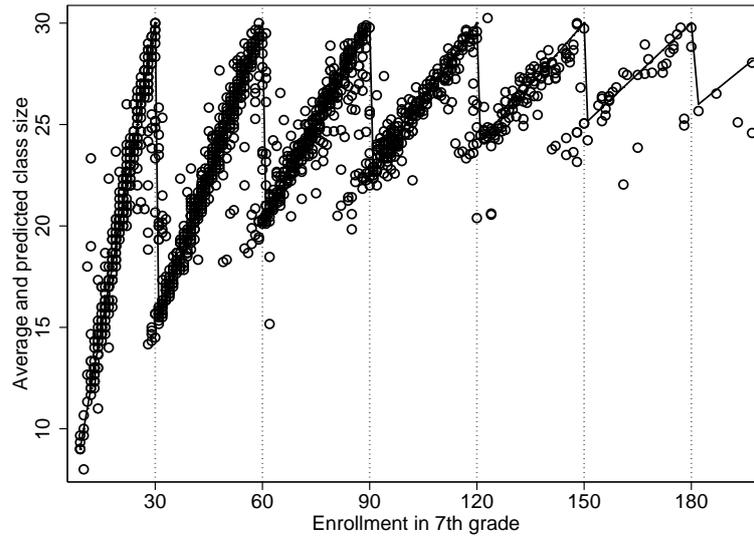


Figure 2: Average and predicted class size, lower secondary school

The first column presents these results for the entire sample of schools, the estimates in the second column are only based on the restricted sample of schools with enrollment levels at most 5 students away from a multiple of 30 (hence, between 25 and 35, between 55 and 65, etc.). These results show that the predicted class size is a strong instrument, the F-values are equal to 454 and 193. Notice though that the coefficients are between 0.52 and 0.63 and are thus smaller than 1, which would result if schools would perfectly stick to the rule. They are however substantially higher than the first stage estimates in Angrist and Lavy (1999) which are between 0.346 and 0.542, when they control linearly for enrollment.

Like in any regression discontinuity design one needs to ensure that the exclusion restriction is not violated. One way of testing this is to check that schools and/or parents do not sort around the cutoffs. We do not observe bunching after the cutoffs in Norway. To go one step further we also compared the characteristics of students and schools around the kinks. A proper regression discontinuity design is like a local randomized experiment and observed characteristics should therefore be balanced. Although this is a necessary and not a sufficient condition (which also requires balancing of unobserved characteristics), evidence for balancing is strong support since it seems difficult to imagine unobservables that matter for outcomes but that are orthogonal to observables that affect outcomes.

To test this, we restricted the sample to students in classes at most 5 students away from the cutoffs and regressed the indicator for being above the cutoff (versus being below it) on various sets of observable characteristics (including enrollment). In the specification that regresses the above/below indicator on student characteristics the p-value for joint significance of these characteristics equals 0.443. In the specification that regresses the above/below indicator on school characteristics the p-value for joint significance of these

Table 4: The impact of predicted class size according to the maximum class size rule on average actual class size

	Full Sample		DS±5	
	Coef..	s.e.	Coef.	s.e.
Predicted class size in grade 7	0.627	(0.029)***	0.532	(0.038)***
Enrollment/10	0.802	(0.181)***	0.416	(0.540)
(Enrollment/10) ²	-0.064	(0.020)***	-0.040	(0.056)
(Enrollment/10) ³	0.002	(0.001)**	0.001	(0.002)
<i>Individual characteristics</i>				
Girl	-0.032	(0.015)**	-0.079	(0.038)**
Age	0.031	(0.024)	0.073	(0.054)
ln(Family income)	0.021	(0.014)	0.020	(0.031)
Education mother (years)	0.002	(0.004)	0.015	(0.008)*
Education father (years)	-0.001	(0.003)	-0.005	(0.007)
1st or 2nd generation immigrants	-0.003	(0.056)	0.031	(0.140)
Parents non-cohabiting	-0.041	(0.023)*	-0.100	(0.054)*
ln(pop. size school district)	0.147	(0.056)***	0.552	(0.141)***
ln(rural pop. size school district)	0.093	(0.082)	0.206	(0.160)
<i>Teacher/School characteristics</i>				
Average teacher education (years)	0.091	(0.360)	-0.211	(0.798)
Average teacher experience	0.050	(0.022)**	0.129	(0.052)**
Fraction of female teachers	-0.385	(0.638)	-0.515	(1.677)
Fraction of teachers with a temp. contract	0.817	(0.495)*	3.091	(1.329)**
Combined school	0.316	(0.184)*	0.353	(0.530)
R-squared	0.653		0.513	
N	111,463		34,780	
N Schools	781		378	
F	454.4		192.5	

Note: Standard errors are heteroscedasticity robust and corrected for school level clustering. */**/***/ statistically significant at the 10/5/1 percent level.

Table 5: Autocorrelation matrix of relative changes in school enrollment ($\Delta \log(e_{st})$)

		Correlation($\Delta \log e_{s,n}, \Delta \log e_{s,m}$)							
$n =$	t	$t-1$	$t-2$	$t-3$	$t-4$	$t-5$	$t-6$	$t-7$	
$m =$									
t	1.000								
$t-1$	-0.459	1.000							
$t-2$	-0.007	-0.464	1.000						
$t-3$	0.007	-0.011	-0.462	1.000					
$t-4$	0.040	0.013	-0.004	-0.483	1.000				
$t-5$	-0.053	0.018	0.028	0.021	-0.505	1.000			
$t-6$	0.042	-0.028	0.001	0.005	0.035	-0.504	1.000		
$t-7$	-0.031	0.039	-0.015	0.003	-0.014	0.038	-0.496	1.000	

characteristics equals 0.495. And the p-value for joint significance of student and school characteristics together equals 0.518. The only separate variable that comes in marginally significant (10%-level) is teachers' education (lower above the kinks).

5.3 Population variation

The second approach exploits demographic variation and was first proposed by Hoxby (2000). Instead of using the variation in enrollment that – in the presence of a maximum class size rule – triggers changes in the number of classes (and therefore class size), Hoxby exploits the population variation conditional on the number of classes.

Variation in enrollment may however correlate with other determinants of student achievement. This happens for example when more educated parents avoid schools where enrollment is large, or when better schools face increased demand if parents selectively choose schools based on school quality. Where part of the variation in enrollment depends on (variation in) population characteristics, there is also a random component (u) that arises from random fluctuations in timing and number of births. Since it seems natural that the share of the random component in total births does not depend on population size, it is assumed that u affects e proportionally, thus enrollment for a given school (and grade level) can be expressed by the following equation:

$$\log(e_{st}) = \log(\bar{e}_{st}) + \log(u_{st}) \quad (2)$$

where \bar{e}_{st} is the deterministic part of enrollment, and $\log(u_{st})$ the i.i.d. part which captures the random variation in enrollment caused by idiosyncratic factors such as biology.

If $\log(u_{st})$ is not correlated with any of the determinants for student achievement (x_{ist} , w_{st} and ϵ_{ist}) in equation (1) a consistent estimate of $\log(u_{st})$ would be a valid instrument for

Table 6: First stage estimates population variation approach - Alternative instruments

	$K = 4$			$K = 0$		
	Coef	s.e.	F-stat	Coef	s.e.	F-stat
Instrument:						
$t - 5$	-0.028	(0.030)	0.8	-0.041	(0.027)	2.3
$t - 4$	0.011	(0.027)	0.2	-0.002	(0.027)	<.1
$t - 3$	0.026	(0.028)	0.9	0.022	(0.026)	0.7
$t - 2$	0.003	(0.032)	<.1	0.012	(0.031)	0.2
$t - 1$	-0.321	(0.041)	61.1	-0.252	(0.039)	41.1
t	0.714	(0.036)	385.8	0.760	(0.033)	516.2
$t + 1$	-0.265	(0.035)	58.3	-0.222	(0.033)	44.7
$t + 2$	0.047	(0.042)	1.3	0.021	(0.031)	<.1

class size since $\log(u_{st})$ is correlated with $\log(e_{st})$. Hoxby assumes that $\log(\bar{e}_{st})$ changes smoothly over time and can be approximated by a grade-school-specific intercept and a school-specific polynomial in time. In this case equation (2) can be written as

$$\log(e_{st}) = \sum_{k=0}^K \alpha_{sk} t^k + \log(u_{st}) \quad (3)$$

To investigate the dynamics of enrollment we calculated autocorrelations of $\log(e_{st})$. These were always higher than 0.95 over the period for which we have enrollment data (1992-2002). Such high persistence suggests that there is indeed at least a school specific intercept α_{s0} . As a next step we calculated the autocorrelation matrix for relative enrollment growth $\Delta \log(e_{st})$ which is reported in Table 5. It is clear from this table that after first-differencing all persistence is gone. The first off diagonal elements are approximately -0.5 and the others are close to zero. Note that if there would be a school specific linear trend first differencing would leave a school specific effect and the high persistence would remain. Table 5 is therefore consistent with the following data generating process

$$\log(e_{st}) = \alpha_{s0} + \log(u_{st})$$

which implies that $K = 0$ in equation (3).

Following Hoxby, we estimated equation (3) for each school separately to obtain the estimated residuals $\widehat{\log(u_{st})}$ both for $K = 0$ and $K = 4$, which can serve as the instrument for class size in (1). As in the maximum class size approach we base our instruments on enrollment in grade 7. Table 6 reports the coefficients from first stage regressions where we instrumented cohort's t enrollment with its own residual (t), but also with the enrollment residuals from previous ($t - 1, t - 2, \dots$) and following cohorts ($t + 1, t + 2$).

The contemporaneous residual is for both specifications of the enrollment process highly significant and has the expected sign; positive enrollment shocks increase class size. The residuals from the adjacent cohorts ($t \pm 1$) are also significant, but to a much smaller extent. The sign on these residuals is negative. This suggests that $\log(u_{st})$ is a first-order moving average. This is consistent with the biological variation motivation behind this approach: if women give birth in year t they are less likely to have another birth in year $t + 1$.¹⁵

Even if $\log(e_{st})$ is correctly modeled, the validity of the population variation approach depends on the exclusion restriction

$$E[\varepsilon_{ist} \cdot \widehat{\log(u_{st})} | \eta_s] = 0$$

which is implemented in a fixed effects 2SLS procedure which eliminates school fixed effects. As a tentative test of the validity of the exclusion restriction we regressed the estimated residuals on the individual and school observables in Table 3. There is no indication that our instruments correlate with observed determinants of students achievements when we calculate the F-test of joint significance on these regressors for the model $k = 4$ and $k = 0$. For the former the F-statistic is 1.145 with a p-value of 0.315, and for the latter the F-statistic is 1.167 with p-value = 0.296. The results will be based on the model with the quartic trend ($K = 4$) for sake of comparability with previous studies. There is substantial variation in the instrument. It is on average -0.0045 with a standard deviation of 0.19, symmetric around zero and the 1st, 5th, 10th and 25th percentiles are -0.6, -0.3, -0.2 and -0.09 respectively.

Note that this approach is valid as long as the variation in enrollment does not trigger a change in the number of classes from 2001/02 to 2002/03 since this would violate monotonicity of the instrument. For this reason we only include schools where predicted number of classes are the same in 2001/02 and 2002/03. Finally, to take into account any remaining endogeneity – which would occur when parents transfer their children to other schools in response to the class size their children are experiencing – we follow Hoxby and also carry out the analysis at the district level in order to cancel out within school district transfers. Other sources of bias such as transfers to private schools or selective grade retention or advancement are not present in Norway because of the negligible presence of private schools and the absence of grade retention/advancement practices.

¹⁵If $\log(u_{st})$ is a first-order moving average, i.e. $\log(u_{st}) = \vartheta_{st} - \rho \vartheta_{s,t-1}$, then this does not invalidate the population variation approach as long as $E[\varepsilon_{ist} \vartheta_{s,t-1} | \eta_s] = 0$. We make this assumption in addition to the contemporaneous exclusion restriction $E[\varepsilon_{ist} \vartheta_{s,t} | \eta_s] = 0$.

5.4 Variations of population variation

A number of authors have used approaches that build on or combine the methods discussed above. We will also present results based on these approaches and therefore describe them.

Population variation and maximum class size rules In addition to exploiting population variation while conditioning on number of classes, Hoxby (2000) also uses the population variation that triggers changes in the number of classes (and therefore class size) because of maximum class size rules. Hoxby identifies all events where the number of classes changed because of maximum class size rules. For all cases where the enrollment change was not more than 20 percent, she then estimates a first differenced version of equation (1).

Small schools Urquiola (2006) focuses on small schools with only one predicted class (per grade). The idea here is that the endogeneity of class size is less of a problem in small rural schools since they are local monopolies and parents cannot enroll their children elsewhere thereby ruling out between school sorting. In addition between class sorting is not an issue since these small schools have only one class. As noted by Urquiola, class size may still correlate with unobserved factors that affect achievement since class size depends on community size, fertility, etc. All class size variation is therefore generated by differences in enrollment and this approach is like Hoxby's, but does not control for school fixed effects since it exploits cross sectional variation, moreover it does not control for trends. The implementation is straightforward and amounts to estimating equation 1 by OLS, where schools with more than one predicted class per grade are excluded. Furthermore the sample is restricted to districts with only one school. This strategy is suitable for the Norwegian context with rural settlement pattern and many small schools.

Within school between grade comparison Wößmann and West (2006) use within school and between grade variation to estimate class size effects. By comparing adjacent grades they account for between school sorting since they eliminate school fixed effects (to the extent to which they are uniform across grades). To eliminate within school sorting problems, actual class size is instrumented with average class size at the grade level. The reduced form estimates are therefore equivalent to the population variation approach of Hoxby, but without correcting for a trend, and without discarding schools around the kinks. We will implement a version of this approach where we exploit variation in enrollment between cohorts in the same school instead of variation between grades. We refer to this approach in the next section as 'population variation without trend'.¹⁶

¹⁶The reason for this deviation of Woessmann and West's approach is that we only have outcomes measured in the final grade of lower secondary school whereas they have outcome data for two adjacent

6 Results

Estimates of the class size effects based on the different approaches are presented in Table 7.¹⁷ These results are all based on specifications with the full set of control variables (e.g. the same control variables as those included in the OLS-regression, and - for the estimates based on the maximum class size rule - a third order polynomial in enrollment). Less elaborate specifications produce very similar results (and are available from the authors upon request). Moreover, for students in combined schools we also know their class size during grades 1 to 6, assuming that students do not switch between combined schools and separate schools. We have also conducted analyses using average class size during the six years in primary school as a class size measure. As mentioned above, we have chosen not to present these results since the results are very similar to the results for average class size in lower secondary schools.¹⁸

The first row of in Table 7 repeats the OLS-results from Section 5.1. These estimates are negative but small and not significant at conventional levels. At face value these estimates would imply a 1 to 2 percent of a standard deviation improvement in achievement for a reduction in class size by 10 pupils. Out of the other fourteen effect estimates presented in the table only the two estimates based on the population variation approach at the district level have the expected negative sign. None of the effect estimates is significantly different from zero although almost all effects are estimated quite precisely, so that substantial negative effects of class size of achievement can be ruled out with high probability. Below each effect estimate we report in square brackets the largest (most negative) effect that falls in the 95% confidence interval (point estimate - 1.96*s.e.). With the exception of the population variation approach the estimates (including those in the first row) imply maximum effects of 1 percent of a standard deviation or less.

The results in Table 7 reveal no clear pattern in the estimated effect sizes across the different methods, apart from the fact that those that exploit population variation (rows 4 and 5) have substantially larger standard errors than the other methods, especially when applied to math achievement. The pooled estimates for the population variation approach shown in column (1) in Table 8 are more precise and equal -0.002 (0.006) for pooled sample, at the district level (not reported in a table) it is -0.008 (0.009). These estimates imply lower bounds of -0.013 and -0.026 respectively.

grades. Hence, their possible bias due to differences between grades is replaced by a possible bias due to differences across cohorts.

¹⁷The first stages are always quite strong, with values from F-tests for significance of the instrument at least equal to 102.8.

¹⁸There may be some concern that the results for average class size in lower secondary schools pick up effects of average class size in primary schools. This is muted, however, by the fact that the maximum class size rule in primary schools is different from that in lower secondary schools (28 versus 30).

Table 7: Estimated class size effects using different methods

	Mathematics			Language		
	Effect	s.e.	N (N _S)	Effect	s.e.	N (N _S)
(1) Conditional on observables	-0.002 [-0.008]	(0.003)	36,915 (608)	-0.001 [-0.005]	(0.002)	74,548 (752)
(2) Maximum class size	0.002 [-0.009]	(0.006)	36,915 (608)	0.008 [-0.0004]	(0.004)*	74,548 (752)
(3) Maximum class size discontinuity sample	0.006 [-0.011]	(0.009)	10,810 (190)	0.007 [-0.005]	(0.006)	23,970 (280)
(4) Population variation	0.016 [-0.031]	(0.024)	26,130 (424)	0.0004 [-0.015]	(0.008)	51,272 (549)
(5) Population variation at district level	-0.020 [-0.127]	(0.055)	240 (120)	-0.009 [-0.036]	(0.014)	420 (210)
(6) Maximum class size + population variation	0.095 [0.074]	(0.011)***	4,962 (67)	0.013 [-0.002]	(0.008)*	11,890 (84)
(7) Small monopoly schools (enrollment<25)	0.026 [0.0004]	(0.013)*	717 (42)	0.005 [-0.018]	(0.011)	1,636 (56)
(8) Population variation without trend	0.010 [-0.009]	(0.009)	36,915 (608)	0.0004 [-0.005]	(0.003)	74,548 (752)

Note: Standard errors are heteroscedasticity robust and corrected for school level clustering. ***/**/* statistically significant at the 10/5/1 percent level. In square brackets the upper bound of the 95% confidence interval. All coefficients come from different regressions. All regressions include controls for gender, age, ln(family income), education mother, education father, a dummy for immigrants, a dummy for parents living at different addresses, ln(population size of the school district), ln(rural population size in school district), average teacher education, average teacher experience, fraction of female teachers, fraction of teachers with a temporary contract, a dummy for combined schools, a dummy for year of observation. The language regressions also include dummies for the language the student took the exam in. The regressions using the maximum class size rule also include a third degree polynomial of enrollment as control. The results in rows (1) and (7) are obtained using OLS; (8) with a fixed-effects regression and the other results are obtained using IV.

Table 8: Class size estimates based on population variation: Alternative instrument sets and pooled estimates

		Instrument				Hansen's J	
		t	$t - 1$	$t + 1$	All		
		(1)	(2)	(3)	(4)	(5)	
- Math	b	0.016	0.042	-1.147	0.021	2.184	$p=0.336$
	s.e.	(0.024)	(0.034)	(13.453)	(0.023)		
	F	17.5	14.4	<.1	10.3		
- Language	b	0.0004	0.0002	0.008	0.001	0.255	$p=0.880$
	s.e.	(0.008)	(0.013)	(0.017)	(0.008)		
	F	85.1	29.8	13.8	29.8		
- Pooled	b	-0.002	0.003	-0.022	-0.003	3.232	$p=0.198$
	s.e.	(0.006)	(0.010)	(0.013)*	(0.006)		
	F	385.8	61.1	58.3	129.8		

Note: see Table 7.

Columns (2) to (4) of Table 8 show also class size estimates based on alternative instrument sets. Column (2) uses the predicted residuals from the previous (older) cohort (the first stage regression coefficient was reported in Table 6), column (3) uses the instrument based on the following (younger) cohort and column (4) uses all three instruments. In terms of the validity of this identification approach it is reassuring to see that we get similar point estimates when we exploit our additional instruments. This is confirmed by the Hansen J-tests reported in the last column which never reject the overidentifying restrictions.

Figure 3 illustrates the findings for the method based on the maximum class size rule. The solid line plots predicted class size as function of enrollment (in intervals of 10 pupils), while the dashed lines shows average math and language achievement as function of enrollment. If smaller class size would benefit pupil achievement we would observe a jump in achievement around the discontinuities where class size drops. Achievement does however not in any way follow (mirror) the pattern of class size (as it does in Angrist and Lavy's paper).

In addition to the sorting for which we do not find evidence above, some have argued that maximum class size rules may be inappropriate to estimate the impact of class size on achievement because of potential input substitution around the discontinuities. To investigate whether this is an issue in Norway we replaced our class size measure by teacher hours per student, measured as the total number of teacher hours in a grade (including extra education to students with specific needs) divided by the number of students in that grade. Although this does not cover resources such as materials, computers,

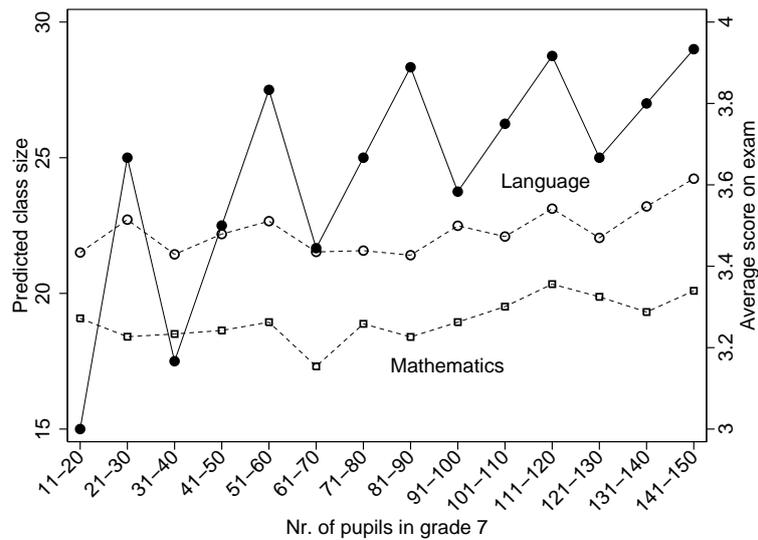


Figure 3: Predicted class size and student achievement

teacher assistants etc., it is the most important measure of resource use in Norway (e.g. Hægeland et al. (2005, 2007)). The correlation between teacher hours per pupil and our class size measure is -0.8.

Table 9 shows results from 2SLS estimations based on maximum class size rules as in Table 7. Note that if input substitution is complete our instrument would not affect teacher hours per pupil. As can be seen from the first column in Table 9, when predicted class size increases by one, teacher hours per pupil decrease by 0.41. Although there might be some input substitution, it is clear from these results that it is far from being complete. Columns (2) and (3) in Table 9 report impact estimates of teacher hours on mathematics and language achievement. The effects are in line with those in Table 7; they are highly insignificant and tend to have the wrong sign.

7 Heterogeneous class size effects

In this section we investigate whether the zero class size effects reported in the previous section, mask effects of class size reduction on achievement for specific sub-groups. The reason to explore this is that various studies report that effects of class size reduction are more pronounced for disadvantaged groups (Angrist and Lavy, 1999; Krueger, 1999) or depend on teacher characteristics (Wößmann and West, 2006).

Table 10 presents results for specific groups based on i) OLS (conditioning on observables), ii) the method that exploits maximum class size rules (for the full sample and around kinks) and iii) the method that exploits population variation (and includes a trend). The first sub-group is that of pupils with low educated mothers (less than 10 years of

Table 9: The impact of School Resources (Teacher Hours/Pupil) on Achievement - 2SLS estimates based on maximum class size rules

	First Stage (1)	Second Stage	
		Mathematics (2)	Language (3)
Full Sample	-0.414 (0.084)*** 111,051 (777)	-0.004 (0.013) 36,698 (605)	-0.012 (0.006)** 74,353 (749)
DS ± 5	-0.410 (0.106)*** 34,669 (377)	-0.012 (0.020) 10,701 (189)	-0.009 (0.007) 23,968 (279)

Note: See note of Table 7.

education). OLS are small and negative. Results from the maximum class size rule using the full sample show positive point estimates of the effect of class size on achievement in both language and math. The sign of the estimates obtained from the discontinuity sample and from the population variation method do, however, differ by subject matter. None of these estimates are statistically significant and there is therefore no evidence in favor of beneficial class size effects for students with low educated mothers.

As a second subgroup we consider children from immigrant families (their share in the population is only 5 percent). For this group all estimates are positive and two are significantly different from zero at the 5%-level.

The results for girls - reported in the third block - are all very small, and do not differ substantially or significantly from the results from girls and boys together. The fourth part of the table presents results for pupils in the lowest quartile of the age distribution. Here again, the emerging pattern is far from consistent. For math achievement, the population variation approach gives in a significantly positive estimate, whereas for language the point estimate is basically zero.

The final two parts of the table present effects for pupils in schools of which the teaching staff belongs to the lowest quartile in terms of education level or in terms of experience. Again the different methods fail to reveal a consistent pattern. Point estimates are all negative for mathematics achievement for pupils in schools with low educated teachers. This result concurs with the hypothesis advanced by (Wößmann and West, 2006). This finding is, however, not confirmed by the estimates on language achievement. We do not see any indication for negative class size effects in schools with less experienced teachers. Points estimates are almost all positive and rather small.

We are inclined to conclude that there is no evidence for heterogeneous class size effects. If anything, it are pupils from low educated families and pupils in schools with

Table 10: Class size effects for sub-groups using different methods

	Math			Language		
	Effect	s.e.	N (N_S)	Effect	s.e.	N (N_S)
<i>Low educated mothers</i>						
- Conditional on observables	-0.007	(0.006)	3,833 (537)	-0.005	(0.005)	7,936 (715)
- Maximum class size	0.008	(0.010)	3,833 (537)	0.0005	(0.008)	7,936 (715)
- Maximum class size discontinuity sample	0.023	(0.015)	1,112 (165)	-0.006	(0.011)	2,474 (267)
- Population variation	0.092	(0.052)	2,711 (356)	-0.020	(0.018)	5,516 (493)
<i>Immigrants</i>						
- Conditional on observables	0.001	(0.011)	2,000 (375)	0.017	(0.008)**	3,628 (493)
- Maximum class size	0.027	(0.020)	2,000 (375)	0.012	(0.014)	3,628 (493)
- Maximum class size discontinuity sample	0.077	(0.040)*	642 (121)	0.016	(0.017)	1,168 (198)
- Population variation	0.214	(0.202)	1,323 (190)	0.006	(0.038)	2,425 (259)
<i>Girls</i>						
- Conditional on observables	-0.003	(0.004)	17,983 (582)	0.001	(0.003)	36,864 (741)
- Maximum class size	0.005	(0.007)	17,983 (582)	0.008	(0.005)*	36,864 (741)
- Maximum class size discontinuity sample	0.015	(0.011)	5,283 (175)	0.006	(0.007)	11,784 (272)
- Population variation	0.029	(0.027)	12,719 (420)	0.012	(0.010)	25,299 (541)
<i>Youngest quartile</i>						
- Conditional on observables	-0.006	(0.005)	9,262 (576)	-0.003	(0.004)	18,602 (737)
- Maximum class size	-0.003	(0.008)	9,262 (576)	0.007	(0.006)	18,602 (737)
- Maximum class size discontinuity sample	0.004	(0.014)	2,722 (171)	0.006	(0.009)	6,082 (270)
- Population variation	0.101	(0.039)	6,519 (410)	-0.002	(0.013)	12,746 (534)
<i>Teacher education in lowest quartile</i>						
- Conditional on observables	-0.003	(0.007)	6,020 (151)	-0.005	(0.004)	13,636 (202)
- Maximum class size	-0.005	(0.010)	6,020 (151)	0.005	(0.007)	13,636 (202)
- Maximum class size discontinuity sample	-0.008	(0.083)	1,458 (37)	0.022	(0.011)*	4,382 (76)
- Population variation	-0.123	(0.088)	4,498 (108)	-0.016	(0.020)	10,475 (156)
<i>Teacher experience in lowest quartile</i>						
- Conditional on observables	0.001	(0.007)	9,416 (170)	-0.002	(0.004)	20,282 (206)
- Maximum class size	0.005	(0.012)	9,416 (170)	0.010	(0.008)	20,282 (206)
- Maximum class size discontinuity sample	0.014	(0.024)	2,577 (52)	0.009	(0.010)	6,779 (84)
- Population variation	0.158	(0.182)	7,028 (119)	0.014	(0.013)	14,077 (146)

Note: See note of Table 7

low educated teachers that benefit the most from a reduction in class size.

8 Conclusion

Based on estimation results that exploit arguably exogenous variation in class size, we find no significant effect of class size during lower secondary school on achievement in grade 9 in Norway. Depending on the identification approach used we can exclude effects as small as 1 to 1.5 percent of a standard deviation for a one student reduction in average class size during three years. Effects are rather similar for different social backgrounds groups and for schools with different teaching staffs.

Our findings contrast sharply with most of the recent studies that apply experimental and quasi-experimental methods to estimate the class size effect. Interestingly, while we applied the same identification strategy as Angrist and Lavy did in their study for Israel the findings are very different. We interpret this as evidence that there is no such thing as a universal class size effect.

Potential explanations for the negligible class size effect in Norway are substitution of parental inputs and uniform teaching styles. Substitution of parental inputs occurs if the parents of pupils who are placed in small classes reduce their own inputs in the education production function (cf. Todd and Wolpin 2003). Uniform teaching styles annihilate potentially beneficial class size effects if teachers are unable to take advantage of the extra time they could have per student. Although the fact that we find some indication that pupils of low educated teachers benefit more from class size reduction than pupils of high educated teachers, seems to contradict this explanation. Further research is required to differentiate between these various explanations. This is important because the policy implications are quite different. If the zero-effects are due to substitution of parental inputs there is not much hope that the policy effects can be improved, although the reduction of parental inputs should be included in a cost-benefit analysis. If the zero-effects are due to uniformity in teaching styles, there remains scope for improvement by teaching teachers to take advantage of smaller classes.

References

- Angrist, J. D. and Lavy, V. (1999). Using Maimonides' rule to estimate the effect of class size on scholastic achievement. *Quarterly Journal of Economics*, 114(2):533–575.
- Bingley, P., Jense, V. M., and Walker, I. (2005). The effects of class size on educational attainment: Danish quasi-experimental evidence and evidence that controls for family, school and neighbourhood effects. Working Paper.

- Bonesrønning, H. (2003). Class size effects on student achievement in Norway: Patterns and explanations. *Southern Economic Journal*, 69(4):952–965.
- Boozer, M. and Rouse, C. (2001). Intraschool variation in class size: Patterns and implications. *Journal of Urban Economics*, 50(1):163–189.
- Browning, M. and Heinesen, E. (2007). Class size, teacher hours and educational attainment. *Scandinavian Journal of Economics*, 109(2):415–438.
- Gary-Bobo, R. J. and Mahjoub, M. B. (2006). Estimation of class-size effects, using ‘maimonides’ rule’: The case of French junior high schools. Discussion Paper No. 5754, CEPR.
- Hægeland, T., Raaum, O., and Salvanes, K. G. (2005). Pupil achievement, school resources and family background. Discussion Paper 1459, IZA, Bonn.
- Hægeland, T., Raaum, O., and Salvanes, K. G. (2007). Pennies from heaven: Using exogenous tax variation to identify effects of school resources on pupil achievement. Unpublished working paper.
- Hahn, J., Todd, P., and Van der Klaauw, W. (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1):201–209.
- Häkkinen, I., Kirjavainen, T., and Uusitalo, R. (2003). School Resources and Student Achievement Revisited: New Evidence from Panel Data. *Economics of Education Review*, 22(3):329–335.
- Hanushek, E. A. (1986). The economics of schooling: Production and efficiency in public schools. *Journal of Economic Literature*, 24:1141–1177.
- Hanushek, E. A. (2003). The failure of input-based schooling policies. *Economic Journal*, 113:F64–F98.
- Hattie, J. (2005). The paradox of reducing class size and improving learning outcomes. *International Journal of Educational Research*, 43:387–425.
- Hoxby, C. M. (2000). The effects of class size on student achievement: New evidence from population variation. *Quarterly Journal of Economics*, 115(4):1239–1285.
- Krueger, A. B. (1999). Experimental estimates of education production functions. *Quarterly Journal of Economics*, 114(2):497–532.
- Krueger, A. B. (2003). Economic Considerations and Class Size. *Economic Journal*, 113:F34–F63.

- Krueger, A. B. and Whitmore, D. M. (2001). The effect of attending a small class in the early grades on college-test taking and middle school test results: Evidence from project star. *Economic Journal*, 111:1–28.
- Lindahl, M. (2005). Home versus School Learning: A New Approach to Estimating the Effect of Class Size on Achievement. *Scandinavian Journal of Economics*, 107(2):375–394.
- Piketty, T. (2004). L'impact de la taille des classes et de la ségrégation sociale sur la réussite scolaire dans les Écoles françaises : une estimation à partir du panel primaire 1997. Working Paper.
- Strøm, B. (2004). Student achievement and birthday effects. Unpublished manuscript, Norwegian University of Science and Technology.
- Todd, P. and Wolpin, K. I. (2003). Towards a unified approach for modeling the production function for cognitive achievement. *Economic Journal*, Features:F3–F33.
- Urquiola, M. (2006). Identifying class size effects in developing countries: Evidence from rural bolivia. *Review of Economics and Statistics*, 88(1):171–176.
- Wößmann, L. (2005). Educational Production in Europe. *Economic Policy*, 20(43):445–504.
- Wößmann, L. and West, M. (2006). Class-size effects in school systems around the world: Evidence from between-grade variation in timss. *European Economic Review*, 50(3):695–736.