No Cohort Left Behind?

Philip Babcock
Kelly Bedard
Jennifer Schulte

University of California, Santa Barbara

September, 2010

Abstract

Much of the debate over the allocation of education resources focuses on the alleged benefits of smallness—of classroom or school—and is based on evidence from small-scale studies. This paper reframes the question in terms of cohort size. Using national data, we find that a 10-percent increase in kindergarten enrollment yields a 0.5 percent increase in cohort shrinkage across early grade transitions, which implies that larger cohorts feature higher rates of retention. Consistent with previous work on class and school size in more restricted settings, this cohort-tracking exercise provides robust evidence at the national level that smallness confers benefits.

I. Introduction

Longstanding debates over the optimal allocation of education resources have focused on the alleged links between class size or school size and student outcomes. We investigate a related topic: the potential link between entering cohort size and subsequent education outcomes—specifically, rates of grade retention. If larger class sizes or schools are less effective at producing education, or if students and teachers feel less engaged in less personalized settings, as is often argued, then one might expect a similar mechanism to operate at the cohort level. Re-framing the research question in this manner allows entrenched ideas to be tested in a new context. Moreover, it allows for an analysis of uncommon breadth and scope, as data on grade-specific enrollment counts by district for primary schools in the U.S. are available at the national level. This paper, then, offers a new and data-rich take on a long-debated theme: we investigate whether smaller is, in fact, better.

Though there is much previous work on class size and education outcomes, we have found no previous research on cohort size as a determinant of grade progression. In Project STAR, source of the most widely cited evidence on class size and student achievement, students in larger classes had lower test scores and were more likely to repeat a grade.[1] Some other studies have also found adverse student outcomes in larger classes, though the evidence varies.[2] It is often argued that teachers are more effective in smaller classes or that students are better able to attach to the education setting when they are in smaller groups. One motivation for a cohort-level analysis is that it offers a way to test whether effects implied by Project STAR and other small-scale studies are visible at the national level, i.e., whether they

[1] Project STAR findings are reported in Mosteller (1995), Finn and Achilles (1990), Hanushek (1999), Nye, Hedges, and Konstatopoulos (2000), Krueger (1999), Schanzenbach (2007), and Word et al (1990) .
[2] Hanushek (1999) and Krueger (2003) survey this extensive literature. See also Angrist and Lavy (1999), Ankerhielm (1995) and Hoxby (2000).

generalize beyond their specialized settings. Changes in entering cohort size are a fundamental source of variation in class size. If positive shocks to class size produce higher retention rates, then this should translate into declining enrollments from one grade to the next. Loosely speaking, large cohorts should produce larger classes, more retention, and thus more "cohort shrinkage" between grades. The analysis here offers a test of this implication.

Previous research also suggests that school size may influence student outcomes. A summary of the school size literature commissioned by the U.S. Department of Education concludes that smaller schools are associated with increased achievement, graduation rates, satisfaction, and behavior (Raywid, 1999).[3] Major policies promoting the construction of small schools have been motivated by this research. The Bill and Melinda Gates Foundation, in particular, embraced and funded these policies, convinced that small schools were effective because they allowed students to "know their classmates and teachers better than they could in a larger school."[4] If small schools and small classes leverage the advantages of small group settings, then the same logic should hold for students in a given district who happen to have been born into smaller cohorts. Does one observe the benefits of smallness that the logic above would seem to imply?

This paper answers the question using district level counts of students from the National Center for Education Statistics (NCES) Common Core of Data (CCD). In the main analysis, we look within a district over time to discern whether fluctuations in cohort size predict cohort shrinkage. Our main finding is that an increase in kindergarten enrollment of

---

[3] Research on the effects of small schools is also reported in Wasley and Lear (2001) and Duke and Trautvetter (2009).
[4] Linda Shaw, The Seattle Times, November 5, 2006. See also Bill and Melinda Gates Foundation, Annual Report, 2003, www.gatesfoundation.org.

10 percent is associated with shrinkage of 0.5 percentage points in the size of a cohort between first and second grade. This finding is robust across multiple specifications that control for possible confounding influences. A similar relationship between kindergarten cohort size and cohort shrinkage is also visible between second and third grade, and between third and fourth grade. County level specifications that control for changes in private school enrollment indicate that the main findings are not easily accounted for by between-district migration or the flow of students between public and private schools. We conclude that larger cohorts feature higher rates of cohort shrinkage, and that this is evidence of increased grade retention. If grade progression is an important outcome, then it is beneficial to be in a smaller cohort.

The remainder of the paper is organized as follows: Section II describes the data and empirical strategy; Section III presents results of the analysis; Section IV interprets the findings; Section V concludes.

II. NCES Data and Empirical Strategy

A. Data

All public school data are from the NCES CCD. The grade span and the number of full-time equivalent (FTE) kindergarten teachers for each school district are from the Local Agency (School District) Universe for the 1992/1993 through 2005/2006 school years. Grade-specific student enrollments are from the Public Elementary/Secondary School Universe for 1992/1993 through 2007/2008. Combined, these sources form a panel of U.S. public school districts that tracks fourteen kindergarten cohorts (1992/1993 through 2005/2006) through second grade. Each kindergarten cohort is identified by the year of kindergarten entry. For example, the cohort that entered kindergarten in the fall of 1992, and who were potential first and second graders in the falls of 1993 and 1994 respectively, are referred to as the 1992 cohort. Each observation matches a kindergarten enrollment (cohort size) to the corresponding change in the size of the cohort between first and second grade.

Each observation is a district that spans at least kindergarten through grade five and has a minimum of five students enrolled in kindergarten in any given year. The grade span restriction is necessary to identify a cohort through the early elementary grades. However, as a cohort progresses, students may enter or leave a district. Since the CCD does not directly measure such transfers, cohort growth may be incorrectly measured; but, under the assumption that these transfers are uncorrelated with the natural variation in cohort size, this measurement error will not bias our estimates (see Section II.B for further discussion of this issue). There are 15,310 such districts in the CCD, summary statistics for which are reported in columns 1 and 2 in Table 1.[5] All summary statistics are weighted by average district size from 1992/93-2005/06, as measured by fourth and fifth grade enrollment. As this is a 14 year panel, districts are included multiple times. The average years of inclusion in the sample is 12.5. Calculating the cohort growth rate between grades one and two and linking this growth rate to kindergarten cohort size requires that a district have complete reporting for three consecutive years (see Section II.B for more detail). This reduces the sample to the 14,373 districts described by the statistics reported in columns 3 and 4. To discern changes within districts over time, we need more than one observation per district. When we drop districts with only one observation, we have 14,052 districts. This sample, described in columns 5 and

[5] We exclude 56 observations (district-year observations, not districts) with first, second, or third grade cohort growth rates exceeding positive or negative 10 (i.e., 10 times more second graders enrolled than first graders or vice versa). All results are similar if these observations (which are clear data errors) are included or if we use other trimming rules.

6, will be used the main analysis. Notice that observable characteristics are similar across the columns of Table 1, alleviating concerns about non-random exclusion.

B. Cohort Shrinkage

The growth rate of a kindergarten cohort between first and second grade is defined as the difference between second and first grade enrollment divided by first grade enrollment:

$$(1) \qquad G_{dy} = \frac{e_{dy}^{2} - e_{dy}^{1}}{e_{dy}^{1}}$$

where G_{dy} is the enrollment growth rate between grades one and two for the cohort that entered kindergarten in year *y* in district *d*. Enrollment is denoted by *e* and superscripts denote grade. Equation (1) is written as function of enrollment because the data are available at this level. In an environment in which there is more first grade failure than second grade failure, we expect *G* to be negative – hence the term "cohort shrinkage." Interpreting changes in cohort shrinkage as stemming from changes in first grade failure involves several important subtleties that are easiest to see when Equation (1) is written as a function of kindergarten cohort size and retention (failure) levels. Assuming no inter-district transfers and no immigration or emigration (these will be discussed in detail later in this section), first grade enrollment for any given cohort *y* is:

$$(2) \qquad e_{dy}^{1} = e_{dy}^{k} - F_{dy}^{k} + F_{dy-1}^{1}$$

where e_{dy}^{k} is the size of the kindergarten cohort in year *y* in district *d*, F_{dy}^{k} is the number of cohort *y* children retained in kindergarten, and F_{dy-1}^{1} is the number of children retained in grade one from the previous (y-1) cohort.[6] Second grade enrollment can similarly be written as:

(3) $$e_{dy}^{2} = e_{dy}^{k} - F_{dy}^{k} + F_{dy-1}^{1} - F_{dy}^{1} + F_{dy-1}^{2}$$

Generally speaking, enrollment in a given grade is kindergarten enrollment plus the stream of retention from the preceding cohort less retention from the cohort of interest. The difference in enrollment between grades one and two for a given cohort y therefore simplifies to:

(4) $$e_{dy}^{2} - e_{dy}^{1} = F_{dy-1}^{2} - F_{dy}^{1}\,.$$

Since our ultimate objective is to understand the effect of cohort size on the first grade failure rate we differentiate the growth rate with respect to kindergarten enrollment. This derivative helps to clarify the assumptions required to interpret changes in observed growth rates as changes in retention. If we assume that shocks to kindergarten enrollment are independent of failure rates for other cohorts[7], $\partial F_{dy-1}^{j} \big/ \partial e_{dy}^{k} = 0 \;\forall\; j = k,1,2$:

(5) $$\frac{\partial G_{dy}^{1}}{\partial e_{dy}^{k}} = -\frac{1}{e_{dy}^{1}}\frac{\partial F_{dy}^{1}}{\partial e_{dy}^{k}} - \frac{\partial e_{dy}^{1}}{\partial e_{dy}^{k}}\frac{[F_{dy-1}^{2} - F_{dy}^{1}]}{(e_{dy}^{1})^{2}}\,.$$

The change in the growth rate reflects the change in the first grade failure rate associated with an increase in kindergarten cohort size as well as the small change in the first grade reference cohort size (the denominator of G). Thus, the marginal change in failure rate captured by the first term of (5) is confounded by a second term which reflects the increase in the denominator. Assuming that the first grade failure rate exceeds the second grade failure rate, as is generally the case in retention data, the second term is positive. The existence of

[6] For simplicity we are assuming students can not fail twice. This could easily be incorporated. All equations are written in terms of kindergarten enrollment rather than first-time kindergarten size and the failure rate of the previous cohort because we can only observe kindergarten enrollment.
[7] We relax this assumption in Section III.

the second term implies, then, that cohort shrinkage slightly underestimates the response of first grade failure to a change in kindergarten cohort size.

Before proceeding to our main regressions, it is instructive to describe the cohort shrinkage measure used here to capture retention rates and compare it with those reported by Hauser, Frederick, and Andrew (2007), the only nationally representative comparable study we are aware of. The second row in column (5) of Table 1 reports that the average cohort growth rate between first and second grade for the sample is -0.014 (sd = 0.083). This says that first grade retention is 1.4 percent larger than second grade retention, given the reasoning above. In contrast, Hauser, Frederick, and Andrew (2007) estimate the difference between the first and second grade retention rates to be -4.5 percent in the CPS. There are several possible reasons for the difference between the two data sets. First, the time periods are slightly different. Second, the CCD enrollment data includes both immigrants and students who enter public school districts from private schools. Using published estimates on immigration from Homeland Security, we estimate the rate at which immigrant children flow into public schools to be about 0.4 to 0.6 percent per year.[8] Using data from the NCES Private School Universe Survey (PSS), we estimate that students from private schools enter into public schools at a rate about 0.6 percent between first and second grade. Adding these inflows to the CCD estimate, the difference between second and first grade retention is between -2.5 and -3 percent. Third, the CPS estimates may be larger in magnitude due to parental reporting errors in the CPS. The CPS estimates are too large if parents are less likely to report that their children repeat second grade compared to first grade. This could occur if there is more social stigma attached to later grade failure.

[8] This includes estimates on illegal immigrants.

C. Empirical strategy

The identification strategy used here is similar in spirit to that used by Hoxby (2000) to estimate class size effects. We begin with the following simple model:

$$G_{dy} = \beta_1 \ln e_{dy}^k + \beta_2 X_{dy} + \lambda_d + \theta_y + \varepsilon_{dy} \quad (6)$$

where districts are denoted by d and cohorts by kindergarten entry years y = 1992,...,2005. G is the rate of cohort growth between first and second grade as defined by equation (1), e^k is kindergarten enrollment, X is the available set of district characteristics, λ is a vector of district fixed effects, θ is a vector of state-by-year indicators, and ε is the usual error term. Equation (6) includes the natural log of kindergarten enrollment to take account of the fact that a one-student reduction is proportionately larger from a smaller base. The fixed effects control for unobservable differences between districts. Including a state-by-year fixed effect is necessary to remove bias created by changes in state educational policies which may be correlated with both student retention and enrollment size. X includes the fraction of students in the district that are eligible for free lunch as well as the fraction of student that are Black, Hispanic, Asian, or Native American/Alaska Native (white is the omitted race category). Summary statistics for these variables are reported in Table 1. All regressions are weighted by average district size and the standard errors are clustered at the district level. The identifying assumption is that kindergarten cohort size within a given district varies randomly with fluctuations in the birth rate. The coefficient of interest, β_1, captures within-district deviations from state trends in cohort growth associated with deviations in kindergarten cohort size.

As with examinations into the effect of class size or school size on student outcomes, one must consider the endogenous responses of parents to enrollment changes. Parents may wish to avoid large cohorts and may switch their students out of an overcrowded cohort within a school to achieve this purpose. However, the cost of switching districts is arguably much higher than the cost of switching schools within the same district. To switch school districts a family incurs the high cost of transporting a student some distance or moving the entire family into a neighborhood within the new school district. Conducting the analysis at the district level, as here, mitigates the problem of selective exit and entry into cohorts. County-level regressions, results of which we report as robustness checks, further limit the channel of migration. It would seem particularly costly to migrate to another county in response to being part of a large kindergarten cohort.

One might also be concerned about potentially endogenous decisions of parents that involve migration between public and private schools. Parents wishing to avoid large cohorts could switch their students out of the public school and into private school without having to migrate to another district or county. Using data on public and private school enrollments by grade and year, we explore whether migration of this type explains cohort shrinkage. We find little evidence that it does.

We will investigate cohort shrinkage between grades 1 and 2, grades 2 and 3, and grades 3 and 4. However, because of the potential for selective migration, we will focus on cohort shrinkage between grades 1 and 2 in the main analysis. We view these as the most reliable estimates for the purpose at hand: Less time has elapsed between kindergarten and first grade than between kindergarten and later grades, allowing less of an opportunity for selective migration between districts.[9]

[9] It might seem natural to use cohort shrinkage between kindergarten and first grade as the dependent variable. However, this would create a situation in which kindergarten cohort size appeared on both the left and right side

III. Results

A. Cohort Size and Cohort Shrinkage

The results of estimating equation (6) are presented in Table 2. Each column of the table represents a different regression. The first row of the table shows the estimated kindergarten cohort size effect on the first grade retention rate (as captured by cohort shrinkage between grades 1 and 2). Columns 1 and 2 show results from specifications that include national and state-level year indicators, respectively, and column 3 adds the full set of controls from Table 1. These estimates indicate that a cohort's growth rate is negatively affected by the cohort's kindergarten cohort size. When a kindergarten cohort in a given district is 10 percent larger, the cohort shrinks about 0.5 percentage points more between first and second grade, which would imply increased grade retention of 0.5 percentage points. Comparison of this estimate to the relationship between class size and grade retention found in Project STAR reveals that the magnitudes are virtually the same: In STAR, first grade retention rises about 0.5 percentage points in response to a 10-percent increase in class size.[10] While an increase in retention of 0.5 percentage points may appear modest, the implied impact is quite large and economically meaningful relative to the first grade retention rate (which we have estimated in Section II.B to be about 3 percent).

Results from column 4 are based on a model in which observations are not weighted by district size. Thus, the smaller districts play a greater role in the calculation of the coefficients. The coefficient on kindergarten enrollment in this column is larger in magnitude than in columns 1 through 3, suggesting that the negative association between cohort size and

---

of the specifying equation. Measurement error in kindergarten enrollment would generate a spurious negative correlation with kindergarten cohort shrinkage. Thus, we did not estimate such a model.

[10] Though the regressor of interest in STAR is class size rather than enrollment, these measures are related, as discussed in Section I.

retention is stronger for small districts. Fluctuations in enrollment tend to generate larger changes in retention for these districts, as will be shown in more detail in Table 3.

Column 5 allows us to better interpret the coefficient on kindergarten enrollment in Columns 1 through 3. As described in equation (4), there are two potential sources of cohort shrinkage that are tied to retention rates. Firstly, when the first grade retention rate rises, a cohort becomes smaller between first and second grade because fewer students progress. But the cohort will become relatively larger between first and second grade if the second grade retention rate of the preceding cohort rises, i.e., if the cohort is augmented by failers from the grade ahead. It is worth distinguishing, then, whether the cohort shrinkage we observe in Table 1 is driven by increased first grade retention for the cohort in question or by a decreased flow of failers from the previous cohort. In column 5, we add kindergarten enrollment for the previous cohort to the full set of controls in column 3, in order to control for the possibility that cohort shrinkage is driven by this second form of retention.[11] The coefficient on lagged enrollment is very small, and does not significantly alter the cohort shrinkage estimate. Large cohorts, then, appear to be associated with higher rates of first grade retention, rather than with decreased retention in second grade for the cohort that precedes them.

B. Robustness and Potential Mechanisms

Table 3 tests the robustness of this main finding and investigates some potential mechanisms. The regressions in Table 2 include national or state level time trends, but one

[11] Specifically, if cohort sizes are correlated across years, the kindergarten enrollment regressor could conceivably be capturing effects of kindergarten enrollment for the previous cohort that relate to second grade retention for that cohort. Controlling directly for the size of the previous cohort is a way to address this potential problem.

could go further. One might be concerned about district-specific trends over time that could be correlated with cohort sizes and retention rates. Because there are over 14,000 districts in the main sample, it is not computationally feasible to include time trends for every district. We rely on a random subsample to circumvent this difficulty. Column 1 of Table 3, based on a model that includes district-level linear time trends along with the full set of controls from Table 1, displays results for a random subsample of 3000 districts. The effect of kindergarten enrollment on cohort shrinkage persists in this specification, and is larger in magnitude than the corresponding estimate in Table 2. The larger coefficient estimate may reflect the fact that this specification better isolates the unanticipated component of cohort size changes, and that changes of this kind are more likely to yield changes in retention.

One way that institutions may respond to fluctuations in enrollment is to add or remove kindergarten teachers. NCES data contain district-level counts of kindergarten teachers. In the model summarized in Table 3, column 2 we include as a covariate the log of the number of kindergarten teachers. If class size influences grade progression in subsequent years, then the addition of teachers ought to reduce cohort shrinkage. Consistent with this intuition, the coefficient on teachers in column 2 is positive and significant. We emphasize that we cannot see class size in first grade, as the NCES data lack teacher counts by grade for later years. As teacher counts in kindergarten are not perfectly correlated with teacher counts in later grades it is not entirely clear how to interpret this coefficient or what is says about the effect of class size. Moreover, the addition or removal of teachers is an endogenous response to enrollment and other factors. For these reasons, we do not characterize this result as a rigorous estimate of a “class size” effect. We report the regression summarized in Column 2 primarily to shed light on the effectiveness of one specific institutional response for which we have data. A comparison of Table 3, Column 2 with Table 2, column 3, reveals that including teacher counts in the regression leaves the coefficient on kindergarten enrollment almost unchanged. Large cohorts appear to produce higher retention rates, whether or not kindergarten teachers are added.

Columns 3 through 6 of Table 3 report results of models with the full set of controls, disaggregated by kindergarten enrollment size. The relationship between kindergarten cohort size and cohort shrinkage is visible in districts of all sizes, but appears strongest in smaller districts. Retention rates for districts that average more than 400 kindergarten students appear to fluctuate less with increases in cohort size. Theses districts may be better able to shift resources to address year-to-year changes in enrollment. Alternatively, it may be that many of the relevant benefits of "smallness" kick in at a lower threshold.

Table 4 extends the analysis by examining the impact of entering cohort size on cohort shrinkage in later grades. Results of kindergarten cohort size are very similar for cohort shrinkage between first and second grade, between second and third grade, and between third and fourth grade: Increases in kindergarten enrollment of about 10 percent are associated with increased cohort shrinkage of about 0.5 percentage points in all the early grades. We have emphasized shrinkage between first and second grade because of concerns about selective migration between districts, but increased cohort shrinkage for larger cohorts is visible across the early grades, and does not appear to be an idiosyncratic property of the first to second grade transition. This finding also means that the cumulative retention effect across grades is substantially larger than first grade results reported in most of this paper.

Tables 2 through 4 show clear evidence of cohort shrinkage for larger cohorts. Whether this is due to retention may be less certain. One could still worry that selective migration out of larger cohorts may be causing the observed shrinkage. The results in Table 5 allow us to look closely at this possibility. As we have argued, it is costly for parents to switch districts in response to finding out that a child is part of a large district-level cohort. However, it is more costly to have to switch *counties*. In Table 5, Column 1 reproduces the main result (of Table 2, Column 3) while Column 2 summarizes results of the same exercise using county-level counts of students instead of district-level counts. The finding of greater cohort shrinkage for larger cohorts is robust to aggregation to the county level. Indeed, large cohorts at the county level experience greater shrinkage than do large cohorts at the district level. This would argue against selective migration as a driver of shrinkage: when this potential mechanism is choked off or limited, the result only gets stronger.

One could still worry about migration from public schools to private schools as a response to being in large cohorts. Table 5, Columns 3 through 6, address this possibility. Data on private school enrollments at the county level are available from the NCES Private School Universe Survey (PSS). This data is available in alternate years; thus we can observe changes in private school enrollment by cohort at the county level between grades 1 and 3 and include this in our county-level regression. As a first step toward this end, Column 3 estimates cohort shrinkage between grades 1 and 3 as a function of kindergarten cohort size at the district level. As expected, given our earlier finding that cohort shrinkage between grades 1 and 2 is about the same size as shrinkage from grade 2 to grade 3, the coefficient on kindergarten cohort size is about twice as large as it was in Column 1. (This is because the cohort shrinkage is calculated over two grades instead of one.) Column 4 summarizes results of an identical model aggregated to the county level. Column 5 restricts the sample to the alternating years for which we have data on private school enrollments. Estimated coefficients in columns 4 and 5 are very similar and give no indication that the sample for which we have data on private schools is idiosyncratic. Column 6 summarizes results from the regression of primary interest: a model that includes the change in private school enrollment as a control to capture migration between private and public schools. When the change in private school enrollment between grades 1 and 3 is included as a regressor, the coefficient on kindergarten cohort size changes very little. We find no evidence, then, that kindergarten cohort size is proxying for flows between public and private school in these regressions. Evidence in Table 5 suggests that selective migration to private schools does not drive the observed relationship between kindergarten cohort-size and cohort shrinkage.

IV. Discussion

Tables 2, 3, 4, and 5 provide robust evidence of higher rates of cohort shrinkage for larger cohorts, consistent with a story of increased rates of retention. One explanation for this finding is suggested by the literatures on class and school size: Smallness facilitates learning and engagement. Small cohorts may leverage the advantages of small group settings to improve student outcomes, just as small classes and small schools have been argued to do. An alternative explanation is that institutions engage in "cohort-smoothing," retaining students in larger cohorts at higher rates in order to even out cohort sizes. Though we cannot distinguish between these two potential mechanisms, both explanations imply that there are benefits associated with being in smaller cohorts, at least for those at the lower end of the ability or performance distribution. For the student marginally retained–whether because he had not learned enough, or because an institution wanted to smooth an enrollment bulge even though the student had, in fact, learned enough–the outcome would seem a deleterious one.

V. Summary and Conclusion

Evidence from district level counts of primary school students at the national level indicates that fluctuations in kindergarten cohort size within a district predict cohort shrinkage. Specifically, an increase in kindergarten enrollment of 10 percent is associated with increased shrinkage of 0.5 percent in cohort size across year-to-year transitions in the early grades. This would imply that larger cohorts feature higher rates of grade retention. Consistent with previous work on school size and class size in more restricted settings, the cohort-tracking exercise here provides robust evidence at the national level that smallness confers benefits.

References

Angrist, J. D. and V. Lavy. (1999). Using Maimonides Rule to Estimate the Effects of Class Size on Scholastic Achievement. *Quaterly Journal of Economics , 114* (2), 533-575.

Ankerhielm, K. (1995). Does Class Size Matter? *Economics of Education , 14* (3), 229-241.

Bill and Melinda Gates Foundation, Annual Report, 2003, accessed June 16, 2010: www.gatesfoundation.org.http://www.gatesfoundation.org/nr/public/media/annualreports/annualreport03/flash/Gates_AR-2003.html

Duke, D. L. and S. Trautvetter. (2009) "Reducing the Negative Effects of Large Schools" National Clearinghouse for Educational Facilities, Washington, DC.

Finn, J. and C. Achilles. (1990). Answers and Questions about Class Size: A Statewide Experiment. *American Educational Research Journal , 27* (3), 557-577.

Hanushek, E. (1999). Some Findings from an Independent Investigation of the Tennessee STAR Experiment and from Other Investigations of Class Size Effects. *Educational Evaluation and Policy Analysis , 21* (2), 143-163.

Hauser, R. M., C.B. Frederick, and M. Andrew. (2007). Grade Retention in the Age of Accountability. In A. Gamoran (Ed.), *No Child Left Behind and Poverty* (pp. 120-153). Washington, D.C.: Brookings Institute.

Hoxby, C. (2000). The Effects of Class Size on Student Achievement: New Evidence from Population Variation. *The Quarterly Journal of Economics , 115* (4), 1239-1285.

Krueger, A. (1999). Experimental Estimates of Education Production Functions. *Quarterly Journal of Economics , 114* (2), 487-532.

Krueger, A. (2003). "Economic Considerations and Class Size," *The Economic Journal*, 113, F34-63.

Mosteller, F. (1995). The Tennesee Study of Class Size in the Early Grades. *The Future of Children , 5* (2), 113-127.

Nye, B., L.V. Hedges, and S. Konstantopoulos. (2000). The Effects of Small Classes on Academic Achievement: The Results of the Tennessee Class Size Experiment. *American Educational Research Journal, 37* (1), 123-151.

Raywid, M. A. (1999). "Current Literature on Small Schools" (ERIC Digest). Charleston, WV: ERIC Clearinghouse of Rural Education and Small Schools. (ERIC Document Reproduction Service No. ED 425 049)

Schanzenbach, D.W. (2007). "What Have Researchers Learned from Project STAR?" *Brookings Papers on Education Policy*, 205-228.

Shaw L. (2006). “Foundation’s Small-schools Experiment Has Yet to Yield Big Results,” The Seattle Times, November 5, 2006.

Wasley, P.A., and R.J. Lear. (2001). “Small Schools, Real Gains.” *Educational Leadership*, 5 (6), pp. 22–27.

Word, E., J. Johnston, H.P. Bain, B.D. Fulton, J. Boyd-Zaharias, C. Achilles, M.N. Lintz, J. Folger, and C. Breda. (1990). *The State of Tennessee's Student/Teacher Achievement Ratio (STAR) Project: Technical Report 1985-1990.* Nashville: Tennessee State Department of Education.

Table 1. Descriptive Statistics

	All		Reporting Cohort Shrinkage		Reporting Cohort Shrinkage and 2+ Observations	
	Mean (Std Dev) (1)	Sample Size (2)	Mean (Std Dev) (3)	Sample Size (4)	Mean (Std Dev) (5)	Sample Size (6)
Number of Districts	15,310		14,373		14,052	
Cohort growth	-0.0135	175,399	-0.0135	175,399	-0.0135	175,078
	(0.0828)		(0.0828)		(0.0827)	
Kindergarten enrollment	4559	210,880	4495	175,399	4497	175,078
	(12,035)		(12,053)		(12,058)	
ln Kindergarten enrollment	6.79	210,880	6.77	175,399	6.77	175,078
	(1.7650)		(1.7559)		(1.7559)	
Fraction free lunch eligible	0.2937	181,814	0.2829	149,937	0.2825	149,649
	(0.2195)		(0.2177)		(0.2173)	
Fraction black	0.1642	206,819	0.1606	172,228	0.1604	171,909
	(0.2131)		(0.2101)		(0.2099)	
Fraction Hispanic	0.1618	206,819	0.1539	172,228	0.1537	171,909
	(0.2229)		(0.2177)		(0.2176)	
Fraction Asian	0.0404	206,819	0.0390	172,228	0.0389	171,909
	(0.0740)		(0.0728)		(0.0727)	
Fraction Native American	0.0116	206,819	0.0116	172,228	0.0116	171,909
	(0.0565)		(0.0565)		(0.0565)	

Weighted by average district-level 1992/93-2005/06 fourth and fifth grade enrollment.

Table 2. The Impact of Cohort Size on Cohort Growth

	(1)	(2)	(3)	(4)	(5)
ln enrollment	-0.0541	-0.0578	-0.0578	-0.0942	-0.0545
	(0.0080)	(0.0089)	(0.0085)	(0.0053)	(0.0070)
lagged ln enrollment					-0.0013
					(0.0052)
Implied percentage point decrease in cohort growth associated with a 10% increase in enrollment	-0.0052	-0.0055	-0.0055	-0.0090	-0.0052
Weighted	Yes	Yes	Yes	No	Yes
National-year FEs	Yes	No	No	No	No
State-year FEs	No	Yes	Yes	Yes	Yes
Additional controls	No	No	Yes	Yes	Yes
Lagged ln enrollment	No	No	No	No	Yes
Sample size	175,078	175,078	175,078	175,078	159,454

Weighted by average district-level 1992/93-2005/06 fourth and fifth grade enrollment. All standard errors are clustered at the district level. Additional controls include fraction free lunch eligible, black, Hispanic, Asian, and Native American as well as indicator variables for missing free lunch eligible and race information.

Table 3. Robustness and Potential Mechanisms

			By Kindergarten Enrollment Size			
			<50	50-199	200-399	400+
	(1)	(2)	(3)	(4)	(5)	(6)
ln enrollment	-0.0956	-0.0577	-0.0918	-0.0892	-0.0846	-0.0302
	(0.0087)	(0.0067)	(0.0050)	(0.0173)	(0.0152)	(0.0121)
ln teachers		0.0033				
		(0.0011)				
Implied percentage point decrease in cohort growth associated with a 10% increase in enrollment	-0.0091	-0.0055	-0.0087	-0.0085	-0.0081	-0.0029
Weighted	Yes	Yes	Yes	Yes	Yes	Yes
State-year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Additional controls	Yes	Yes	Yes	Yes	Yes	Yes
Linear district trends	Yes	No	No	No	No	No
Sample size	38,480	147,002	58,379	69,170	25,445	22,084

Weighted by average district-level 1992/93-2005/06 fourth and fifth grade enrollment. All standard errors are clustered at the district level. Additional controls include fraction free lunch eligible, black, Hispanic, Asian, and Native American as well as indicator variables for missing free lunch eligible and race information. Column (1) includes 3000 randomly drawn districts.

Table 4. The Impact of Cohort Size on Cohort Growth across Primary Grades

	(1) $G_{dy}^{1 to 2}$	(2) $G_{dy}^{2 to 3}$	(3) $G_{dy}^{3 to 4}$
ln enrollment	-0.0578	-0.0473	-0.0494
	(0.0085)	(0.0062)	(0.0078)
Implied percentage point decrease in cohort growth associated with a 10% increase in enrollment	-0.0055	-0.0045	-0.0047
Weighted	Yes	Yes	Yes
State-year FEs	Yes	Yes	Yes
Additional controls	Yes	Yes	Yes
Sample size	175,078	159,441	144,361

Weighted by average district-level 1992/93-2005/06 fourth and fifth grade enrollment. All standard errors are clustered at the district level. Additional controls include fraction free lunch eligible, black, Hispanic, Asian, and Native American as well as indicator variables for missing free lunch eligible and race information.

Table 5. The Impact of Cohort Size on Cohort Growth: County Level Aggregation and Private Schools

	(1) District-Level G_{dy}^{1to2}	(2) County-Level G_{dy}^{1to2}	(3) District-Level G_{dy}^{1to3}	(4) County-Level G_{dy}^{1to3}	(5)* County-Level G_{dy}^{1to3}	(6)* County-Level G_{dy}^{1to3}
ln enrollment	-0.0578	-0.1319	-0.1226	-0.2705	-0.2524	-0.2506
	(0.0085)	(0.0301)	(0.0207)	(0.0810)	(0.0786)	(0.0783)
Private school flow from grades 1 to 3						0.00002
						(0.00002)
Implied percentage point decrease in cohort growth associated with a 10% increase in enrollment	-0.0055	-0.0126	-0.0117	-0.0258	-0.0241	-0.0239
Weighted	Yes	Yes	Yes	Yes	Yes	Yes
State-year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Additional controls	Yes	Yes	Yes	Yes	Yes	Yes
Sample size	175,078	38,837	159,441	34,657	13,033	13,033

Weighted by average district or county-level 1992/93-2005/06 fourth and fifth grade enrollment. All standard errors are clustered at the district or county level. Additional controls include fraction free lunch eligible, black, Hispanic, Asian, and Native American as well as indicator variables for missing free lunch eligible and race information. * Sample restricted to years for which private school enrollment data is available (alternating years).