MAKING POLICY CHOICES:
IS CLASS SIZE REDUCTION THE BEST ALTERNATIVE?

Doug Harris

Michigan State University and
Michigan Department of Education

David Plank

Michigan State University

I. Introduction

Everyone wants to improve student performance. The question is, how? States and school districts across the country have chosen class size reductions as the answer. Anyone who has ever taught a class or read the available research knows that reducing class size does make a difference. Teachers have more time to give personal instruction to students and spend less time dealing with disciplinary issues.

Unfortunately, knowing that class size reductions make a difference is not very helpful for making school policy. Yes, lowering class sizes can help, but so can improving technology, providing better textbooks, expanding professional development opportunities for teachers, and increasing teacher salaries to attract and keep more able educators. All of these initiatives can improve student achievement. They also cost money. Information about the relative gains and costs of different policies should be accounted for in making the tough decisions about education spending.

How should these tough choices be made? In theory, all programs should be adopted if the benefits exceed the costs. In the real world, however, funding levels are fixed. This means that only some of the beneficial programs can be adopted. To get the most out of their limited resources, policymakers must instead try to find the most *cost-effective* mix of programs – seeking the "biggest bang for the buck."

The problem that policymakers face is that there is rarely good information available to compare the costs and benefits of multiple programs. These comparisons are essential for sound decision-making, yet education research usually focuses on the benefits, and only for individual programs. As a result, administrators and policymakers are forced to rely on gut instincts, intuition, the demands of special interests, and ideas

from other districts or states. It is difficult to fault them for this approach, since they have little other information on which to base their choices.

In this paper we seek to promote the use of cost-effectiveness analysis to support more informed policy decisions. We illustrate the value of this approach by providing comparisons between policies aimed at increasing the *number of teachers* (decreasing class size) with those intended to increase the *ability of teachers*. Our results call into question the conventional wisdom about class size. Analysis of the relative gains and costs of these two policies suggests that greater improvements in student performance could be achieved through sustained policies aimed at increasing teacher ability than through further class size reductions.

The remainder of this paper provides more detail about these conclusions. Technical details are provided in the appendix. Additional technical information is available from the authors.

II. **National Trends**

Policies affecting class size and teacher quality have changed dramatically in recent years. Table 1 below summarizes these changes from 1960-1990.

Table 1 – Changes in School Resources (1960-1990) [1]

Year	Class Size	Teacher Daily Wage (1990$)	Proportion of Total Spending		
			Teacher Salaries And Benefits	Other Instructional	Administration Other Costs
1960	25.8	$124	0.68	0.07	0.25
1970	22.3	$155	0.67	0.11	0.22
1980	18.7	$143	0.60	0.15	0.25
1990	17.2	$183	0.61[2]	0.12	0.27

Table 1 shows that teacher salaries (adjusted for inflation) have increased dramatically in recent years. What the table does not show is that the salaries of other college-educated workers have also increased. In fact, the relative salaries of teachers have gone down even as real salaries have increased.[3]

Table 1 also shows a significant downward trend in class sizes. Together with the increase in teacher salaries, these reforms have produced a large increase in total spending on education. Despite the increase in average teacher wages, the *proportion* of funds going to teachers decreased from 68 percent in 1960 to 61 percent in 1990. This reflects even larger percentage increases in spending for special education and support services.

III. The Effectiveness of Class Size Reductions

There are many ways to measure student performance. No single measure will capture all of the things that parents and educators consider important. For our purposes in this paper, "effectiveness" refers to changes in two common measures of student performance: test scores and future wages. Test scores are intended to reflect student achievement in knowledge, verbal and quantitative ability, thinking skills, and other education content. Future wages reflect many different aspects of performance after students have left school. These capture the likelihood of graduating from high school, motivation and ability for finishing college, people skills, and general work ability. In the following sections we focus on test scores. Wages are considered in section VIII.

Experimental research consistently shows that class size affects how students perform on standardized tests. The Tennessee STAR study of class size reductions is the best known piece of research in the field. STAR was a large-scale experiment that showed clear achievement gains from smaller classes in lower grade levels. The experiment took place in the 1980's and included random assignment of 12,000 students to small and large classes for grades kindergarten through three (K-3). The average large class had approximately 24 students and the average small class had 15 students.[4]

A more recent experiment is Wisconsin's SAGE program, which included random assignment of 5,000 students to small and large classes with average sizes quite similar to the STAR program.[5] Many other smaller scale experiments have been conducted. Glass and Smith (1979) studied 725 separate estimates of class size effects that used varying degrees of sample control and randomization.[6]

The results from these studies are summarized in Table 2. The numbers in the Table represent the percentile gains for students who start at the 50[th] percentile and who experience an average class size reduction of five students (e.g. from 25 to 20).[7] This assumes a very flexible reform in which class sizes can be reduced for any grades and student types – as long as the average decreases by five.[8]

According to the results in table 2 from Glass and Smith (GS), the effect of such a reform in elementary grades (1-6) would help a student at the 50[th] percentile to move to the 52[nd] percentile – an increase of approximately two percentile points.

The other effects reported in table 2, such as those for "math" and "science," assume that the change in average class size occurs over all thirteen grades (K-12), and not just in elementary or secondary. We would expect these effects to be larger simply

because the smaller classes are applied over a larger number of grades, allowing the benefits to accumulate over time. Therefore, the smaller numbers for elementary and secondary grades do not reflect smaller effects.

Table 2 – Percentile Gains in Test Scores from Class Size Reductions
(Base = 50th percentile; Treatment = reduce classes by 5 students)

Study Characteristics	Study		
	STAR	SAGE	GS
Elementary Grades	2.4	1.0	1.9
Secondary Grades	---	---	4.4
Black	---	3.4	---
White	---	1.7	---
Reading	4.0	1.7	---
Language Arts	---	2.4	---
Math	6.8	4.0	---
Science	5.2	---	---

The most important idea to take away from table 2 is that reduced class sizes do produce gains in students' academic achievement. The data in the Table also suggest that the gains are larger in math than in other subjects. Black students appear to gain more than whites, which is consistent with other qualitative evidence showing that disadvantaged students gain more from class size reductions. In addition, the table suggests that gains are larger for secondary students, which departs from the conventional wisdom about policies to reduce class size.

IV. The Costs of Class Size Reductions

In this section we develop cost estimates for the simple class size reduction policy described in section III. The main cost of reducing class size is the cost of hiring more teachers. If an individual school district seeks to reduce class sizes, teachers may move

into the district from other school districts. This is unlikely to require much change in the total number of teachers in a county or state. State-level policies to reduce class size will almost certainly require additional teachers to enter the workforce. The effects of class size reduction therefore depend on how large the pool of candidates for new teaching positions is, and also on the qualifications of the teachers in the pool.

Research evidence on the teacher surplus issue suggests that most districts face a surplus of candidates for teaching positions. The size of the pool is uncertain, however, and there is great variation both within and across states.[9] In the cost model used here, we assume that the pool of candidates is 10 percent larger than the total number of teachers in our base model. We vary this assumption from 0-30 percent.

Regardless of the type of district, it is likely that new teachers hired in response to class size reduction policies will be of lower ability than those already employed. This is almost certainly true if districts are now selecting the best available teachers in their hiring processes.[10] The benefits from reducing class sizes are therefore likely to be at least partially offset by the decrease in average teacher ability. One way to prevent this would be to attract teachers with higher ability who would otherwise be working in other professions or other districts into the pool.[11] This is likely to be costly, however. Most collective bargaining agreements require that **all** salaries be increased, not just those of new teachers.[12]

To illustrate some of the trade-offs involved with these policies, consider California's recent program that puts $1 billion annually into class size reductions. In many California districts, the evidence suggests that test score gains for students in smaller classes were offset by test score losses for students who ended up with less able

teachers.[13] Wealthier districts were able to reduce class size by attracting the best teachers from low-income districts, while low-income districts were obliged to hire less qualified (and often unqualified) teachers. The net effect of the policy was to help some students at the expense of others.

Cost estimates require specific assumptions about the relationships between class size reductions, the pool of candidates for teaching jobs, and average teacher ability. The "teacher supply elasticity" is the percentage change in the number of available teachers divided by the associated change in teacher salaries. The available research suggests that the elasticity ranges between 0 and 2.0, depending on the time frame.[14] A change in salaries today will obviously have little effect on the number of teachers applying tomorrow. So, the short-run elasticity is essentially zero. As information spreads through the workforce, however, more teachers may enter the pool of candidates. The long run elasticity may be as large as 2.0, meaning that a one percent increase in salaries produces a two percent increase in the number of workers willing to teach.[15]

Our model also depends on how teacher salaries affect the ability levels of teachers who may apply. Manski (1987) concluded that a ten percent increase in average teacher salary would increase the average teacher SAT score by 10 points. Ballou and Podgursky (1992, 1994) and Figlio (1997) obtain similar results.

Class size reduction policies may also require capital expenditures if more classrooms are needed to accommodate students. Two facts suggest that the cost of classroom space is small compared the costs of teachers. First, capital costs comprise only 10 percent of total education expenditures. In addition, one-half of U.S. schools have extra space available that cannot easily be utilized for other purposes.[16] For these

reasons capital costs are excluded from the model. The impact of this and other assumptions is considered in more detail below.

The cost model described above, combined with the effectiveness information in section III above, suggests that it would cost $435 annually per student to move a student from the 50[th] to the 54[th] percentile by reducing average class size by five students. These costs vary depending on assumptions about the labor supply elasticity, ability elasticity, and capital costs. Changes in assumptions do not change the main conclusions of the paper, however.

V. The Effectiveness of Policies Designed to Increase Teacher Ability

There are two main ways to increase teacher ability. The first is to hire teachers who begin their jobs with high ability. The second is through training and professional development once teachers have been hired.

According to research evidence, several characteristics of new teachers appear to improve student performance. These include verbal ability, quality of the teacher's undergraduate college[17], and certain types of undergraduate courses.[18] Some of these measures, such as verbal ability, seem to have obvious importance because these are the same abilities that education is meant to provide to students. Teachers cannot pass skills on to students that they themselves do not possess.

Other teacher characteristics, such as quality of undergraduate institution and types of coursework undertaken, can be interpreted in different ways. We view these characteristics as *signals* about expected teacher performance in the classroom – signals that administrators use in making decisions about hiring and initial compensation.

The available evidence suggests that the most important form of training for veteran teachers is *on-the-job training*, as measured by years of experience. In contrast, formal training including master's degrees appears to have little impact on student teacher performance.[19] Evidence regarding other forms of professional development is highly dependent on the specific type of training that teachers receive.

The evidence presented above focuses on administrators' decisions in hiring and compensating teachers. In fact, however, people must apply for teaching positions before they can be hired. The choices that potential applicants must make include whether to take formal teacher training (usually in college), whether to stay in a particular job or switch jobs after entering the profession; whether to stay in the teaching profession or change professions; and whether to continue working or leave the labor force. Considerable research has been done about how potential teachers make these decisions.[20] Not surprisingly, the reasons are complex and many are not under the control of school districts or state governments.

Salary is one important factor determining who teaches, as suggested earlier. In contrast to the research on class size, unfortunately, there is no experimental evidence regarding the effects of teacher salaries on student achievement. Instead, most studies on teacher salaries use a statistical tool called regression analysis that tries to imitate experimental conditions. This non-experimental technique requires that the regressions include "control variables" that affect student performance, including class size and student characteristics. Regression estimates that exclude these variables suffer from "omitted variable bias" and can yield biased estimates of the effects of the teacher salaries.

Most non-experimental studies find that salary matters, but the effects appear to be small.[21] Most of these studies suffer from the omitted variables problem, however, which means that their results may be misleading. To identify unbiased estimates of how teachers' salaries affect student achievement, we first selected regression estimates that include both class size and teacher salaries in the analysis. This is important because we do have experimental estimates of the class size effect. We restrict our attention to regression estimates in which the class size effect is close to the experimental estimates found in STAR, SAGE, and other studies. If one of the effect estimates is unbiased, it is more likely that the other effect estimates are unbiased as well.[22]

Twenty regression estimates from seven different studies include both a class size and a teacher salary variable. Eight of these estimates were left after restricting this sample based on the class size evidence. Taking the average of these teacher salary effects suggests that a statewide increase of 10 percent in teacher salaries will raise a student from the 50th to the 52nd percentile on a norm-referenced test.

As in our discussion of class size reductions above, this analysis does not make distinctions among the many strategies that schools could use to increase teacher salaries. These options are discussed in section VIII.

VII. The Costs of Improving Teacher Ability

The effectiveness of professional development, standards, and compensation were all considered in the previous section. The costs of professional development programs for in-service teachers vary, and must be evaluated on a case-by-case basis. The cost of obtaining a master's degree can be estimated, but it is not clear that the degree provides much benefit to students. If this is the case, then any cost is too high.

At first glance, policies that raise standards for teachers appear to have no cost at all. Districts and states can simply set rules that no one can teach who has not achieved a minimum score on some type of certification exam (e.g. the NRE). The more restrictive the rules, however, the fewer candidates are left in the pool. If standards are set so high that there are too few qualified applicants to fill the available positions at the salaries offered, then salaries may have to rise to attract a sufficient number of qualified candidates.

Our analysis considers the effects on student achievement of a statewide increase in average teacher salaries. Multiplying the change in teacher salaries by the number of teachers yields the total costs of the reform. As with the class size reduction discussed above, policies that seek to increase teacher ability by increasing salaries will also affect the number of teachers available. The final cost will depend again on teacher supply elasticity, teacher ability elasticity, and other factors. We make the same assumptions here as we did regarding class sizes. Our results suggest that raising a student from the 50[th] to the 54[th] percentile using teacher salary increases would cost about $200 per student per year.

VII. Cost Effectiveness and Simple Policy Reforms

Most research provides information about benefits *or* costs of *individual programs.* The purpose of cost effectiveness analysis is to combine information on costs *and* benefits for *various programs* to obtain the highest possible level of student achievement. Policymakers need more than just information about benefits – they need to know how they can get the "biggest bang for the buck."

The costs and benefits described above for class size reductions and teacher salary changes are presented in a way that these comparisons easy. The cost of increasing student test scores from the 50th to the 54th percentile are estimated to be $435 per student for class size, but only $200 for teacher salaries. These results vary somewhat based on assumptions about the labor supply elasticity, labor ability elasticity, and other factors. We considered ranges of values for these factors, but even values at the extreme ends of these ranges did not alter the main conclusion.

These results suggest that sustained teacher salary increases would be more cost effective than class size reduction in raising student achievement. Hiring fewer teachers with greater ability will produce greater gains at lower cost than hiring larger numbers of teachers with less ability. The short-term effect of a change in class size would almost certainly be greater, but the long-term effect would be smaller. As teachers leave the profession, higher salaries would attract better candidates for teaching positions. Over time these new teachers would have a greater impact on student performance at a lower cost.

The discussion thus far has focused on improving student test scores as the main goal of education. An additional purpose of education is to prepare students for their adult lives, including their careers. Different types of analysis are possible when the goal is identified as increasing student wages, because both costs and benefits are expressed in dollars. This contrasts with an analysis in which test scores were the outcome of interest. Test scores may be related to personal success in adult life, but they are also related to our general desire for knowledge and good citizenship. It is quite difficult to place a dollar value on these goals and values.

As with test scores, it is important to have precise estimates of the benefits that reducing class size or increasing teacher salaries provide for future student wages. Our review of the evidence suggests that a $100 increase in teacher salaries would increase future students wages by 0.8 percent. A decrease in class size by one student in all grades would increase future student wages by 1.6 percent.

One simple and useful way to approach the analysis with future wages is simply to ask, "At what point do the extra costs of education reforms exceed the extra future wages for students due to the reforms?"[23] The results of our analysis are quite similar to those in the earlier sections. The absolute size of the gain in student achievement is larger when we reduce class size than when we increase teacher salaries, but the cost of the change is also higher. This suggests that it would be more cost-effective to raise teacher salaries than to further reduce class size. The appendix provides additional details about the model used to produce these results.

VIII.   Cost Effectiveness: The Details Matter

Real-world policies are rarely as simple as those considered in the previous sections. Instead, policies are almost always filled with fine print, restrictions, and rules that add complexity. This section focuses on how the earlier results might change based on differences in policies and different cost assumptions.

 Possible restrictions on class size policies include maximum class sizes, restrictions by grade, and unique rules for special education and other student populations. All of these restrictions would add to the cost of reform. For instance, suppose a school imposed a maximum class size of 25. This would require splitting a

class of 26 into two very small classes (e.g. 13 and 13) and hiring an additional teacher, because a class of 26 would violate the rule.[24] Imposing this rule would make class size reduction far more costly; more flexible reforms are generally cheaper.

Similarly, there are many different ways to increase teacher salaries, ranging from "signing bonuses" for new teachers to across-the-board increases for all teachers to "merit pay" based on professional development, experience, peer evaluations, or other factors. These different strategies can produce similar increases in **average** teacher salaries, while producing very different increases for particular teachers and very different consequences for the pool of candidates for teaching jobs. The details of policy design are therefore likely to make a big difference for how salary increases affect student achievement.

The available evidence suggests that increasing teacher salaries will improve student achievement. It does not provide guidance for choices among alternative strategies. As in the case of class size reductions, however, complicating salary policies by adding rules may increase the cost of reform, sometimes in unforeseen ways. For example, policies that offer "signing bonuses" or increased salaries for beginning teachers may be a good strategy for attracting talented young people into the teaching profession. Such policies may produce dissatisfaction or resentment among veteran teachers, however, leading to increased turnover or reduced retention. It would clearly be desirable to design policies that targeted salary increases in ways that would attract or retain more able or more talented teachers, but it is far from clear that administrators or policy makers are able to make the fine distinctions among prospective or current teachers that would make such policies effective.

It is also important to note that policies that significantly restrict the pool of candidates for teaching positions through testing or stricter certification requirements may require corresponding increases in salary to attract a sufficient number of teachers.

Policies to reduce class size or increase teachers' salaries could be implemented at various levels of the education system, from individual school districts to the federal government. In our analysis we assume that these are state or federal reforms. Action by an individual district may make local students better off, without much impact on the larger education system. Any single school district that reduces class sizes or increases salaries can probably attract the new teachers that it needs from nearby districts or private schools. State or federal policies that affect **all** school districts are likely to bring about major changes in the market for teachers, however.

IX. Conclusion

How can we improve student performance? Past research is useful, but falls far short of providing a definitive answer to this question. The results that we present in this paper run against the grain of recent policy initiatives. They suggest that sustained policies aimed at increasing teacher ability will be more cost effective than class size reductions in raising student achievement. This is true regardless of whether we focus on academic achievement or the future wages of students.

If research suggests that salary increases and standards are so cost effective, why might actual policy decisions differ? One possible explanation is that class size reductions have a much more immediate and identifiable impact on student performance. Teachers like class size reduction because it has demonstrable effects on student

achievement and allows them to focus more closely on each of their students. Parents

like the reductions because they want their kids to have personal attention. This puts a

great deal of pressure on school leaders to reduce class sizes, even if it is not the most

cost effective use of resources. In contrast, changes in teacher quality are hard to

observe, and they may affect student achievement only over an extended period of time.

Teachers may appreciate the importance of higher salaries, but making the case to parents

and taxpayers is harder. Hiring fewer teachers and paying them more simply lacks the

intuitive political appeal of hiring more teachers and reducing class size.

There is another lesson in this paper that is even more important than the specific

results for class size reductions. Researchers and educators, devoted to finding a solution

to pressing educational needs, search very hard for programs that work. Through

innovation and creativity, many successful programs have been developed in the process.

However, these successes will remain meaningless if we cannot successfully choose

among them. The best apple in the barrel is not very useful if we cannot find it. Instead

of reaching out and picking the first one we see, it is essential that we compare at least a

handful of those that look good on the surface. These choices are not merely a matter of

economics – of mundane calculations of benefits and costs. Rather, the ability of

educators to make these tough choices will determine how much we can contribute to the

education of our students.

<u>Appendix A – Policy Reform and Students' Future Earnings</u>

A different type of analysis is possible when student's future wage is the outcome of interest. Economic theory suggests that the optimal input level occurs where the marginal benefit of the input equals the marginal cost. In the case of test scores, it is difficult to apply this rule because the "benefits" of test scores are difficult to translate into an economic value. Wages, in contrast, have a fairly straightforward interpretation, since they are expressed in monetary terms.

We start by developing functions for the present discounted value of total costs and benefits, based on the above assumptions. (All costs and benefits are discounted at a base value of 3 percent with a range of 0-6 percent.) We then consider an optimization problem in which the government maximizes a social welfare function (benefits minus costs), subject to a technology constraint on the production of education. Since the above approach does not impose any limits on spending levels, we then add a cost constraint similar to current total costs.

The optimization problems used here are partial equilibrium, ignoring some relationships across markets. However, the consistency of the results using different approaches, and the magnitude of the differences in cost effectiveness across policies, suggest that the results would not change in a more complex general equilibrium model.

Bibliography

Ballou, D. (1996). "Do Public Schools Hire the Best Applicants?" *Quarterly Journal of Economics*, Vol. 111 , no. 1, pp. 97-133.

Ballou, D. & Podgursky, M. (1994). "Recruiting Smarter Teachers." *Journal of Human Resources*, Vol. 30, no. 2, pp. 326-338.

Ballou, D. & Podgursky, M. (1992). "Buying Quality: Will Higher Pay Improve the Teaching Workforce?" *Working Paper*, University of Massachusetts at Amherst.

Brewer, D.J., Krop, C., Gill, B.P., & Reichardt, R. (1999). "Estimating the Cost of national class size reductions under different policy alternatives," *Educational Evaluation and Policy Analysis*, Vol. 21, no. 2, pp. 179-192.

Currie, J. (1991). "Employment Determination in a Unionized Public-Sector Labor Market: The Case of Ontario's School Teachers," *Journal of Labor Economics*, Vol. 9, no. 1, pp. 45-66.

Figlio, D. (1997). "Teacher Salaries and Teacher Quality," *Economics Letters*, Vol. 55, pp. 267-271.

Hanushek, E. & Rivkin, S. (1996). "Understanding the Twentieth Century Growth in U.S. School Spending," *Journal of Human Resources*, Vol. 32, no. 1, pp. 35-68.

Hanushek, E. & Pace, R.R. (1995). "Who Chooses to Teach (and Why)?" *Economics of Education Review*, Vol. 14, no. 2, pp. 101-117.

Levin, Henry (1983). *Cost Effectiveness*: *A Primer*. Sage: Beverly Hills.

Manski, C.F. (1987). "Academic Ability, Earnings, and the Decision to Become a Teacher: Evidence from the National Longitudinal Study of the High School Class of 1972." In *Public Sector Payrolls*, Ed. David A. Wise, University of Chicago Press: Chicago.

Nye, B., Hedges, L., and Konstantopoulos, S. (1999) "The Long-Term Effects of Small Classes: A Five-Year Follow-Up of the Tennessee Class Size Experiment," *Educational Evaluation and Policy Analysis*, Vol. 21, No. 2, pp. 127-142.

Norris, G., and R. Richburg (1997) "Hiring the Best," *American School Board Journal*, Vol. 184 no.11, p.46, 48, 55.

Endnotes

[1] The class size numbers in table 1 are from the *Digest of Education Statistics* (1999). The other numbers and interpretation of table 1 are based on Hanushek and Rivkin (1997).

[2] This number is estimated because information about teacher benefits was not available for 1990.

[3] See Hanushek and Rivkin (1996).

[4] Nye, Hedges, and Konstantopoulos (1999) describe the study in greater detail. Other papers focus on the weaknesses of the project and previous interpretations of results, including attrition and some violations of random assignment.

[5] The SAGE program, which started in 1996, also included staff development, after-school programs, and a new curriculum for kids in the treatment group. While this might appear to complicate the analysis, Molnar et al (1999) found that these other programs had no significant effect after controlling for class size.

[6] Hedges and Stock (1983) reanalyzed the Glass and Smith sample with some modifications, but found that these changes did not affect their results.

[7] All calculations based on test scores assume that student performance is normally distributed.

[8] More precisely, the poicy experiment we propose assumes that school will reduce class sizes in the same they as those schools used in the studies. This approach also ignores possible non-linear effects. For instance, the effect of going from a class of 30 to 25 may be different from deceaing class sizes from 25 to 20.

[9] See, for example, Ballou (1996).

[10] This might not be true if the new teachers are expected to provide instruction on computers or other skills with which newer, younger teachers may be more familiar. Other reasons may also prevent this from being the case and these are discussed in later sections. However, these do not appear to be strong enough to contradict the general conlusion that teachers on waiting lists are less able on average

[11] This assumes that schools are hiring the most able teachers, which may not be a true assumption, as mentioned earlier.

[12] One way around this rule is the signing bonus, which provide one-time payments only to new teachers.

[13] Class Size Reduction Research Consortium (2000). http:\\www.classize.org.

[14] It is extremely difficult to estimate elasticites for particular time frame – e.g. short run versus long run. However, the distinction is important.

[15] See Manski (1987), Currie (1991), Ballou and Podursky (1994).

[16] See the *Condtion of America's Public School Facilities: 1999* from the U.S. Department of Education. Also, note that if extra space could be rented out, then the cost would include this foregone rental value.

[17] See, for example, Ehrenberg and Brewer (1994).

[18] See, for example, Goldhaber and Brewer (1997).

[19] See, for example, Hanushek (1986, 1994) and Goldhaber and Brewer (1997).

[20] These exams may also exclude teachers who lack verbal and quantitative skills, but who have other valuable abilities.

[21] See, for example, Hanushek and Pace (1997), Ballou (1992), Ballou and Podgursky (1994).

[22] To check the validity of these estimates, we also estimated the effect that teacher salaries have on average teacher test scores (e.g. SAT scores from college or certification exams). We then reviewed the evidence regarding the effect of teacher ability on student outcomes, where teacher ability is measured by teacher test scores. The results were quite similar to thosse obtinaed using the procedure described in the text.

[23] Economists use the terminology, "marginal costs and marginal benefits." The point at which these are equal is called the "optimal" level. See the appendix for more details.

[24] The costs of class size reductions are interrelated with two other polices of many schools and states. First, schools of choce policies make it increasingly difficult for schools to shift student populations across schools to avoid these class splits. In middle and high schools, an additional problem is that students request specific classes, and teachers specialize in ways that prevent splitting classes at all.