

**The Effects of High-Stakes Testing
On Achievement: Preliminary Findings
About Generalization Across Tests**

**Daniel M. Koretz
The RAND Corporation**

**Robert L. Linn
University of Colorado**

**Stephen B. Dunbar
University of Iowa**

**Lorrie A. Shepard
University of Colorado**

Achievement testing has been a primary instrument of educational reform in the last decade. In many states and districts across the country, testing programs have been transformed from a means of monitoring the progress of students into a mechanism for holding students, teachers, and school administrators accountable. This test-based accountability, according to its proponents, will cause real improvements in the performance of the educational system and in the achievement of students.

Most observers agree that these changes have had profound effects, but a vehement debate has arisen about their desirability. Increases in test scores have been observed in many states and districts during the first few years following the introduction of a high-stakes testing program. The degree to which those increases in scores reflect real improvements in student achievement, however--rather than gains

measure. Public reporting of these scores therefore creates an illusion of successful accountability and educational performance. (Evidence pertaining to the second major issue, the effects of high-stakes testing on instruction, are presented in another paper in this symposium [Shepard, 1991].)

CHARACTERISTICS OF THE DISTRICT

I cannot describe the district or the tests that it has used in detail because we have agreed to protect the district's confidentiality.¹ However, some basic information is needed in order to interpret the results that I will present. The district involved--"District B" in our larger design--is a large, high-poverty urban district with large numbers of both black and Hispanic students. Three-fourths of our sample schools had non-Asian minority enrollments of 70 percent or more, and about half minority enrollments above 90 percent. Three-fourths of the third-grade students in our median school received free lunch. The district as a whole was very similar to our sample in these respects.

The district's overall minority enrollment did not change much during the four years of our study, but the ethnic composition of some schools changed markedly. The minority enrollment in the typical school in our sample changed by less than 2 percentage points between 1986 and 1990. The range, however, was from a 20-percentage point decrease to a 23-percentage point increase.

The district uses unmodified commercial achievement tests for its testing program, which is perceived as high-stakes. Through the spring of 1986, they used a test that I will call Test C. Since then, they have used another, called Test B, which was normed 7 years later than Test C. The district publishes school median grade equivalents (GEs) on the tests. Our median schools had median scores on Test B in 1990 that were about average in vocabulary (GE = 3.6), below average in reading (GE = 3.1), and above average in mathematics (GE = 4.3; see Table 1). (The district tests in spring, so a GE of 3.7 corresponds to the 50th percentile.) Because the distribution of school medians on the district's test is positively skewed, the means of school medians were 1 to 5 academic months higher than the medians.

¹ Although we cannot credit the individuals who cooperated in this study by name, we gratefully acknowledge their assistance. A study of this sort is both burdensome and politically risky, and a good many districts decided that they could not participate for those reasons. Indeed, one of our primary sites pulled out of the study weeks before we would have administered tests, citing the political risks they would face if the district were identified despite our efforts to keep its identity confidential. The individuals in our participating sites deserve credit for shouldering the risk and burden that this study involved.

Table 1. School Medians on District's High-Stakes Test (Test B), GEs, 47 Schools

	Mean of School Medians	Median of School Medians
Math	4.4	4.3
Reading	3.6	3.1
Vocabulary	3.9	3.6

DESIGN AND SAMPLE ADJUSTMENTS

We sampled intact classrooms from within schools on a random basis. Sampled schools were randomly divided into subsamples, and one subsample of schools was administered Test C, the test that the district had used through 1986. Other subsamples were administered a parallel form of Test B or alternate tests constructed to match Test B or Test C in content but not in format.

Because some of the district results to which we must compare our results are reported on the level of school buildings, it was necessary to check the degree to which our classrooms were typical of the schools from which they were drawn. We did this by comparing the school-wide results on the district's test to the performance on that same test of our within-school samples. This discrepancy varied markedly and was large in some schools. C dGW our within-school sample of 0.00079 results on tg7 was .00079 Tc 0 Tw 12 0 0 1

Mathematics

The results in mathematics show that scores do not generalize well from the district's test to Test C, even though Test C was the district's own test only four years ago and is reasonably similar in format to Test B. (That is, both Test C and Test B are conventional, off-the-shelf multiple-choice tests.) The school-level results on which this conclusion rests are displayed for mathematics (total) in Figure 1.

Looking at the first five medians starting with 1986, one sees the traditional pattern: scores dropped markedly when the new test was introduced and then rose again, particularly between the first and second years of the new test. The drop when the test was changed was about half an academic year: from a GE of 4.25 to 3.7 This drop presumably reflects two factors: lesser familiarity with the test and




Our re-administration of Test C showed that in mathematics, the schools in this sample had slipped by roughly four academic months in the four years since the district had itself used Test C. This can be seen by comparing the right-most box in Figure 1 to the left-most box. In 1990, the median school had an adjusted median of 3.83, compared to the median of 4.25 in the last year that Test C was the district's own test. For the spring of grade 3, a drop of this magnitude is not inconsequential.

Our re-administration of Test C showed that in mathematics, the schools in this sample had slipped by roughly four academic months in the four years since the district had itself used Test C. This can be seen by comparing the right-most box in Figure 1 to the left-most box. In 1990, the median school had an adjusted median of 3.83, compared to the median of 4.25 in the last year that Test C was the district's own test. For the spring of grade 3, a drop of this magnitude is not inconsequential.

While the results we have just presented are appropriate for comparing scores on Test C in 1990 and 1986 (because they are based on the identical test and norms) they understate the discrepancy between our 1990 Test C results and the district's 1990 results on Test B--that is, the right-most two boxes in Figure 1. The reason is that the edition of Test C that the district used in 1986, and that we re-administered in 1990, was normed 7 years earlier than the district's Test B. Accordingly, we adjusted our Test C medians using the publisher's conversion tables to newer norms for Test C that were set within one year of those used for Test B.

The effect of this adjustment is to further increase the discrepancy between our results and the district's own, although not dramatically. Schools' median scores in mathematics drop a month or two when placed on the newer Test C scale. With that adjustment, the median school had a median GE score of 3.7 on Test C, about 6 academic months lower than the median school on Test B (see Figure 2).

While these comparisons are bleak enough in terms of medians, they are more extreme yet if one considers the schools that score particularly well on the district's own Test B. While it would be easy to read too much into a distribution of only 36 school medians, the pattern is too striking to ignore. The distribution of medians on the district's test has more positive skew than does the distribution on Test C, and the top of the distribution on Test B is a GE of 8.4, more than three academic years above the highest median on our Test C (GE = 5.1). To some degree, this may be function of the tests; the GE scale on Test B has a substantial positive skew, and the skewness of the school

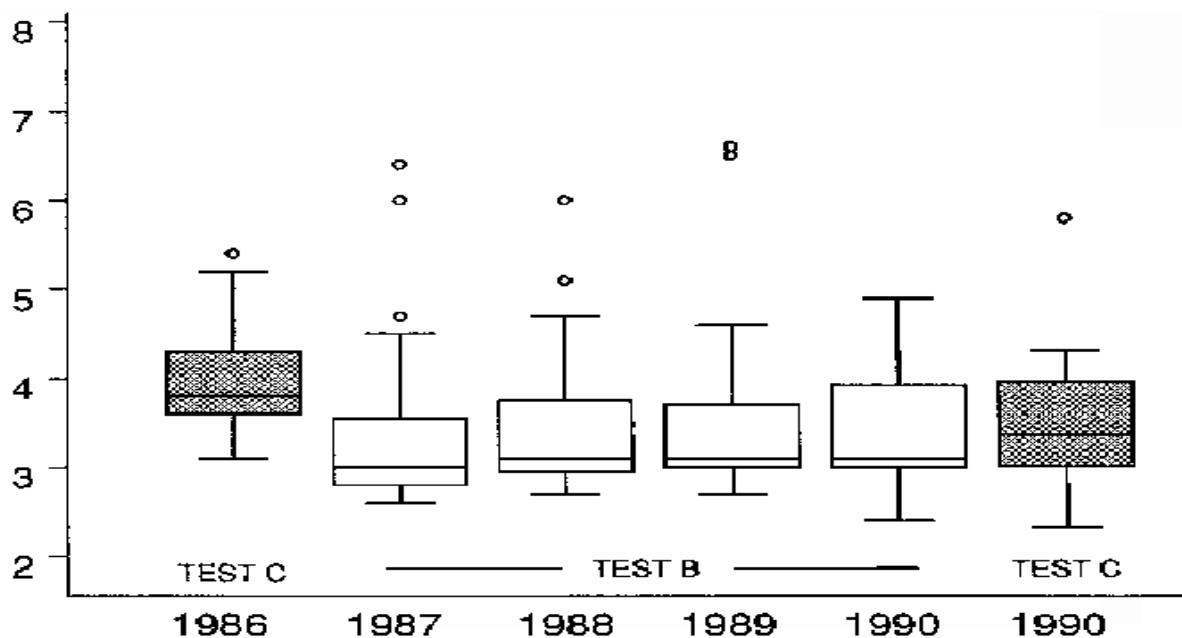


medians was apparent to some degree even in the first years the test. The movement of outliers to ever-higher values in subsequent years, however, suggests that test preparation may also be at work.

One effect of this skewness is that the of school medians shows an even greater difference between Tests B and C: nearly 8 academic months, compared to the 6-month difference between medians.

Reading


In terms of two of our contrasts, the lack of generalization was as bad or even worse in reading than in mathematics. The performance of our schools in reading slipped even more dramatically when Test B was first introduced. The median of school medians dropped by a full 8 academic months (Figure 4), from about average (GE = 3.8) to well below average (GE = 3.0). The comparison of our Test C the district's final administration of that test in 1986 shows essentially the same that we found in mathematics: about 4 academic months, to a GE of about 3.4.



Our cross-sectional comparison of 1990 Test B and adjusted Test C results, however, showed relatively little difference in terms of the median schools. The of school medians was in fact albeit trivially, on our adjusted Test (by less than one academic month; see Figure 5). Here again, however, the distribution of schools was more positively skewed on the district's Test B, and accordingly the of school medians showed a moderate drop--roughly three academic months, from Tests B to Test C.


Thus we have an inconsistency: in terms of change over time, lack of generalization across tests is at least as great in reading as in mathematics, while in terms of cross-sectional comparisons, the disparity

is in reading is smaller or nonexistent, depending on the measure. This inconsistency hinges in part on the differences in trends on the district's own Test B. In mathematics, as noted earlier, the median school's score followed the trajectory we would expect: a sharp drop when the district changed tests, followed by a fairly rapid recovery as everyone gets used the new test. In reading, however, the recovery never happened. The median school score on the district's Test B rose only one month between 1987 and 1990, and the of school medians rose only 2 months.



Vocabulary

Our sample showed a modest weakness of generalization across tests in vocabulary as well. The median school was about average on Test C when the district last used it in 1986, with a median GE of 3.8. In 1990, the median school scored two months lower on the district's Test B (GE = 3.6: see Figure 6). On our re-administration of Test C, however, the median school scored nearly four months lower than the median on that test 4 years earlier (GE = 3.4). When Test C is scored against the more recent norms, the cross-sectional difference between the median schools for Tests B and C remains about two months.



RELATIONSHIPS BETWEEN CHANGE AND OTHER FACTORS

We have conducted limited analyses of the relationships between historical change on Test C and other variables. These analyses are consistent with our primary hypothesis of inflation of scores from ses of the relat

correlations between changes on Test C and change in percent minority enrollment ranged from -.24 (for vocabulary) to -.31 (for mathematics). This relationship cannot account for the overall decline on Test C, however, because on average, our schools showed almost no change in minority enrollments. Change on Test C was not appreciably related to change in percent receiving free lunch.³

test subsample, while randomly drawn, turned out not to be entirely representative of our sample as a whole.)

Table 2. Comparisons of Test B Scores and Parallel Form Scores, Median Percentile Ranks and GEs

	Reading (N = 133)		Mathematics (N = 136)		Vocabulary (N=136)	
	PR	GE	PR	GE	PR	GE
Test B	57	4.2	78	4.9	61	4.7
Parallel Form	55	4.1	78	4.8	63	4.8

The first of our substantive comparisons of student-level results contrasts score on one of our "alternate" tests to the district's Test B. we constructed two of these alternate tests. ("We" in this case is primarily the three people who follow me today: Lorrie Shepard, Roberta Flexer, and Elfrieda Hiebert.) You will hear more about these tests in the following papers. Each of the alternate tests was designed to parallel the conventional achievement test used in one of our districts in terms of content but not format. Thus these tests included item types such as open-ended questions and multiple-choice questions which permitted more than a single correct answer. One of these tests, which we call Alternate Test B, was designed to match District B's test and curriculum framework.

We assumed from the outset that to some degree, students would perform more poorly on our tests as a result of more difficult item types, regardless of any effects teaching to the test. Accordingly, our design called for equating each of our alternative tests in two other districts. The alternative test used in the district I am reporting today was equated to Test B using samples from one district where testing low stakes and a second district that has high-stakes testing but does not use Test B.

Comparison of students' performance on the equated alternate test and the district's Test B showed a substantial deficiency of generalization, particularly in mathematics. In mathematics, the students who took both of these tests--who do not overlap with the sample who took ourt800011 Tw ort 300.805c - twosTea,on the

Table 3. Comparisons of Test B Scores and Equated Alternate Test B Scores, Mean Percentile Ranks and GEs

	Reading (N = 620)		Mathematics (N = 707)	
	PR	GE	PR	GE
Test B	42	3.4	35	3.1
Alternate Test	61	4.3	46	3.6

A second substantive comparison of student-level results contrasts the scores on Tests B and C for students who took both in 1990. Roughly 750 students in 34 schools are included in these results, depending on the scale.

In mathematics, the results of these student-level comparisons are quite similar to the cross-sectional school-level results already reported: they show a striking weakness of generalization. The median student in this subsample received a GE of 4.5 on the district's Test B (slightly higher than the median of our entire sample). This corresponds to a national percentile rank of 67 (Table 4). These same students' scores on Test C, once adjusted to the newer norms, were 7 academic months and 16 percentile points lower. (Note that the subsample administered Test C were similar to our total sample in reading but scored somewhat higher in mathematics.)

Table 4. Comparisons of Median Student Scores on Tests B and C, GEs and National Percentile Ranks

	Reading (N = 133)		Mathematics (N = 136)		Vocabulary (N=136)	
	PR	GE	PR	GE	PR	GE
Test B	42	3.4	67	4.5	48	3.6
Test C ^a	38	3.4	51	3.8	35	3.4

The results of this student-level comparison were quite different in reading: the median student scored as high on Test C as on the district's Test B in terms of GEs but slightly lower in the metric of national percentile ranks. (The estimated relationship between percentile ranks and GEs is not the same for Tests B and C). Recall that the school-level cross-sectional 1990 comparison between Tests B and C in reading also showed an atypically small discrepancy between the two tests.

^a Test C results are expressed with reference to new norms that are within one year of those used for Test B.

The student-level comparison of vocabulary scores was consistent with the school-level results already reported. The median student scored two months lower on Test C than on the District's Test B (GEs of 3.6 and 3.4, respectively). The difference in terms of national percentile ranks, however, was more substantial: the median student scored at the 48th percentile on Test B but only the 35th percentile on Test C.

CONCLUSIONS

In mathematics, then, all of the comparisons presented here strongly support our primary hypothesis that performance on a conventional high-stakes test does not generalize well to other tests for which students have not been specifically prepared. Three of the five primary contrasts reported here showed differences in performance of six to eight academic months between the high-stakes test and others, the fourth was just shy of that, and the fifth showed a difference of four months. In terms of estimated percentile ranks, two of the contrasts showed differences of 15 or 16 percentile points.

The evidence in reading is less consistent but nonetheless suggests significant weaknesses of generalization in some instances. The historical comparison on Test C showed a fall-off of four academic months and the change from Test C to Test B in 1987 caused a drop of eight months. Our alternative test suggests a difference about half that large. The cross-sectional comparisons of Tests B and C are the exception both the school- and student levels: they show differences ranging from near zero to three academic months, depending on the measure.

The more consistent and generally larger disparities among tests in mathematics are not surprising. Aggregate data on the "Lake Wobegon effect" show more inflation of scores in mathematics than in reading (Linn, Graue, and Sanders, 1990), and we therefore hypothesized that we would find weaker generalization in mathematics.

There is more to be done to explore this lack of generalization. Subsequent members of this panel will provide several other pieces of the puzzle: eviden

performance, say, on "mathematics as tested by Test B but not Test C." They ar

REFERENCES

Cannell, J. J. (1988). National normed elementary achievement testing in America's public schools: How all-50 states are above the national average. *Journal of Educational Research*, 80(2), 5-9.

Koretz, D. (1988). Arriving at Lake Wobegon: Are standardized tests exaggerating achievement and distorting instruction? *Journal of Educational Research*, 80(8-15), 46-54.

Linn, R. L., M. E. Graue, and N. M. Sanders, "Comparing State and District Test Results to National Norms: The Validity of the Claims that 'Everyone is Above Average'," *Journal of Educational Research*, 83(3), 1990(b), pp. 5-14.

Shepard, L. (1990). Inflated test score gains: is the problem old norms or teaching the test? *Journal of Educational Research*, 83(3), 15-22.

Shepard, L. (1991). The effects of high-stakes testing on instruction. In R. L. Linn (Chair), *High-stakes testing: A symposium presented at the annual meeting of the American Educational Research Association and the National Council on Measurement Education, Chicago, April 5, 1991.*