

School Size, Student Achievement, and the “Power Rating” of Poverty:

Substantive Finding or Statistical Artifact?

Theodore Coladarci
(coladarci@umit.maine.edu)

University of Maine

Paper presented at the 2006 meeting of the
American Educational Research Association, San Francisco

Abstract

The proportion of variance in student achievement that is explained by student SES—“poverty’s power rating,” as some dub it—tends to be lower among smaller schools than among larger schools. Small schools, many claim, are able to somehow disrupt the seemingly axiomatic association between SES and student achievement. Using eighth-grade data for 215 public schools in Maine, I explored the hypothesis that this in part is a statistical artifact of the greater volatility (lower reliability) of school-aggregated student achievement in smaller schools. This hypothesis was supported when the dependent variable was mathematics achievement. In contrast, this hypothesis received no support when reading achievement served as the dependent variable. Implications for subsequent research are discussed.

As every student of education research knows, the relationship between student achievement and socioeconomic status (SES) is well-established in the empirical literature: All things equal, as student SES increases, so does student achievement (e.g., Sirin, 2005; White, 1982). Further, this holds regardless of the unit of analysis employed (e.g., student, school, multilevel). The seemingly axiomatic nature of this relationship notwithstanding, a recurring finding in rural education research is that SES and school size interact in affecting student achievement (e.g., Howley, 1996; Howley & Bickel, 1999; Huang & Howley, 1993; Johnson, Howley, & Howley, 2002; McMillen, 2004; also see Friedkin & Necochea, 1988; Lee & Smith, 1997). In other words, the magnitude of the relationship between SES and achievement depends on the size of the school, or, equivalently, that the magnitude of the relationship between school size and achievement depends on the SES makeup of the school.

How is such an interaction demonstrated? With the school as the unit of statistical analysis, for example, interaction is shown by regressing achievement on SES, school size, and the mathematical product of SES and school size, and then testing the product term for statistical significance. If the slope associated with this term is statistically significant—which researchers have been reporting with remarkable consistency—there is an interaction between SES and school size. A common way to illustrate such an interaction is to show that the school-level correlation between SES and achievement is weaker among smaller schools than among larger schools. That is, SES explains less of the variance in school achievement among smaller schools than it does among larger schools. As Huang and Howley (1993) put it, smaller schools “mitigate” the effect that SES has on student achievement.

The mitigating-effect finding enjoys considerable fanfare by researchers, advocacy groups, and practitioners alike. Johnson, Howley, and Howley (2002), highly respected rural education researchers all, judged this finding to be “among the most consistent ever to be

reported in educational research” (pp. 36-37). The Rural School and Community Trust, which tirelessly advocates for rural schools and communities, crafted the phrase “poverty’s power rating” to refer to the percentage of variance in achievement that is explained by SES (i.e., the coefficient of determination). In newsletters and press releases, the Rural Trust celebrates the recurring finding that the power rating of poverty is markedly lower—sometimes negligible—among smaller schools than among larger schools. “In study after study,” the organization’s president recently announced, “small schools have been shown to cut poverty’s power over student achievement” (Tompkins, 2006). And in an op-ed published in my local newspaper, a school superintendent and his colleagues summed it up this way: “Small schools are an antidote to the impact of poverty on school achievement” (Butler et al., 2005, p. A9).

I must confess that, despite my affinity to rural education and its causes, I have always been uneasy with the mitigating-effect finding and, in particular, the markedly lower “power rating” of poverty in smaller schools. As much as I am attracted to the notion that smaller schools, by virtue of their smallness, are somehow able to disrupt the achievement disadvantage of lower-SES students, and as much as I can imagine the many ways in which smaller schools might be able to pull this off (although hard data would be helpful), my immediate suspicion was that the diluted SES-achievement correlation among smaller schools may have little to do with the educational experience characterizing such schools. Rather, I suspected a statistical artifact at play.

Loosely defined, a statistical artifact is where a research result is misleading because of an artificial or extraneous effect due to statistical considerations. For example, if X has modest variance and, further, the correlation between X and Y is $r = 0$, the absence of relationship between X and Y very well could be due to restricted range in X (a statistical artifact) rather than to an absence of relationship between the two constructs underlying X and Y . In the present

context, the putatively ameliorative role of smaller schools in the SES-achievement relationship would be a statistical artifact if, say, there were much less variability in either student SES or student achievement among smaller schools than among larger schools. Truth be told, this was my immediate suspicion, both because it is so obvious as a plausible rival hypothesis (when subgroup correlations are comparatively small) and because I saw no acknowledgment of this possibility by those who were doing (or celebrating) the research. But I was unable to find evidence of restricted variance in the statistics reported by the researchers. Nor did such evidence surface in my own analyses of Maine data that had been featured in a 2005 Rural Trust news release (Rural School and Community Trust, 2005).

My interest in the challenges that small schools face related to the “adequate yearly progress” requirement of No Child Left Behind suggested another possible statistical artifact: the greater volatility, or lower stability, of school-level student achievement among smaller schools (Coladarci, 2003). School achievement differs widely from one year to the next for smaller schools, whereas larger schools enjoy more stability in this regard (e.g., Coladarci, 2003; Hill & DePascale, 2003; Kane, Staiger, & Geppert, 2002; Linn & Haug, 2002).

Consider Figure 1, for example, which shows the relationship between (a) the size of the fourth-grade cohort tested in a Maine school and (b) the one-year change in the proportion of students in that school who met or exceeded the standard on the Maine Educational Assessment reading test. Although the average change from one year to the next hovers around zero for all schools, there is considerably greater variability among smaller schools in the amount of this change. For schools having 15 or fewer fourth graders, for instance, this change ranges from $-.47$ (declining from 60% proficient to 13% proficient) to $+.83$ (increasing from 17% proficient to 100% proficient). In contrast, the corresponding figures are only $-.07$ and $+.09$, respectively,

among schools having 150 or more fourth graders.¹

At issue here is the reliability of school-aggregated student achievement. Insofar as any measure of school achievement is less reliable—i.e., more volatile—for a smaller school than for a larger school and, further, because a measure's reliability places an upper limit on its ability to correlate with any other variable (e.g., Thorndike, 1982, p. 222), a plausible conjecture is that the lower SES-achievement correlation among smaller schools is an artifact of the lower reliability of school achievement for such schools. In short, this is the conjecture I investigated in the present study.

In pursuing the statistical-artifact hypothesis, my intention was not to debunk popular opinion. Rather, I simply wished to determine whether a celebrated proposition in the rural education literature could withstand a sincere attempt to falsify it. If such an attempt were to fail, then we all are entitled to a greater confidence in this proposition—greater *warranted* confidence, I believe—than we presently can claim.

Method

Data Source and Variables

My focus is on eighth-grade achievement in Maine public schools, using reading and mathematics data from the Maine Educational Assessment (MEA) for the 2002-2003 and 2003-2004 school years. (The MEA scale ranges from 501 to 580.) For each public school having an eighth grade, I created a weighted two-year mean for both reading achievement (*reading*) and mathematics achievement (*math*). Similarly, I determined for each school the weighted two-year percentage of students receiving subsidized meals (*poverty*).

¹ The +.83 school (upper left corner) is somewhat of an outlier. The small-school range is $-.47$ to $+.46$ with this discrepant case excluded.

As for operationally defining school size, I immediately faced the distinction between a school's total enrollment across all grades and a school's mean enrollment per grade. Howley (2002, pp. 52-53) argues that the latter is the appropriate measure of school size because per-grade enrollment takes into account a school's grade configuration—that, say, a K-8 school with 270 students (30 per grade) is arguably smaller than a 6-8 school with 270 students (90 per grade). I have yet been able to appreciate the logic of this position, which inevitably must fall on how one conceptualizes “school” and its effects on students. But because most mitigating-effect studies employed the enrollment-per-grade measure of school size, I followed suit in the analyses reported below. Specifically, I determined the mean enrollment per grade for each school, averaged across 2002-2003 and 2003-2004 (*school size*). (I confess that I ran all analyses using a total-enrollment measure of school size as well, which yielded similar results to those based on enrollment per grade.)

To estimate a school's volatility in eighth-grade achievement, I determined the difference in mean achievement from 2003-2004 to 2002-2003 for reading and mathematics separately. I then recoded the absolute value of these differences to obtain a volatility rating for each school. There were separate volatility ratings for reading and math (*volatility*), and both were formed as follows:

volatility rating	change in school mean-achievement
1	0 to 2.50 points
2	2.51 to 5.00 points
3	5.01 to 7.50 points
4	7.51 to 10.00 points
5	10.01 to 12.50 points
6	12.51 to 15.00 points
7	15.01 to 17.50 points
8	17.51 to 20.00 points

Analyses

I restricted my analyses to public schools in Maine that (a) had an eighth grade in 2002-2003 and 2003-2004, (b) had data on all variables for both 2002-2003 and 2003-2004, and (c) had neither changed their grade span from one year to the next nor absorbed in 2003-2004 students from a school that had closed at the end of 2002-2003. Finally, I eliminated schools that did not have at least two eighth-grade students in each of the two school years. These restrictions resulted in a final sample of 216 schools (from a universe of 233 public schools having an eighth grade in 2003-2004).

The school served as the unit of analysis. After conducting preliminary analyses to establish the trustworthiness of the data, I began by demonstrating the aforementioned interaction between socioeconomic status and school size. I did so using ordinary least-squares regression (e.g., Aiken & West, 1991), where, in the case of two independent variables, the equation is $\hat{Y} = a + b_1X_1 + b_2X_2 + b_3X_1X_2$. Here, \hat{Y} represents the dependent variable (either reading or math); a is the intercept; X_1 and X_2 are poverty and school size, respectively; and X_1X_2 is their mathematical product. Prior to creating the product term and consistent with common practice, I centered poverty and school size at their respective means to reduce the inevitable collinearity engendered by multiplicative terms.

The statistical significance of b_3 , the slope of the product term, indicates the presence of interaction between X_1 and X_2 —that the magnitude of b_1 varies with X_2 , or, symmetrically, that the magnitude of b_2 varies with X_1 . In the present context, this means that the degree of association between poverty and achievement (b_1) depends on school size (X_2), or, equivalently, that the degree of association between school size and achievement (b_2) depends on the socioeconomic status of the school (X_1). By entering the product term on a separate step, I

obtained the increment in explained variance (ΔR^2) that is associated with the poverty-size interaction (the statistical significance of which is identical to that of b_3).

To further illustrate the degree of interaction between poverty and school size, and, in particular, to recast this interaction in terms of poverty's power rating, I fit separate achievement-on-poverty regression lines for schools falling above and below the median per-grade enrollment. That is, I did a median split on school size and then regressed reading and math (separately) on poverty for below-median schools and for above-median schools. The magnitude of interaction is shown by the degree to which the two within-group regression lines are nonparallel. From this analysis, I also obtained the within-group correlations between each achievement measure and poverty, which, when squared, represents the power rating of poverty.

To explore my statistical-artifact hypothesis—that poverty's reduced power rating, when examined among smaller schools, reflects the lower reliability of school-level achievement in such schools—I repeated these analyses on successively less-volatile collections of schools. The first set of analyses included all 216 schools (i.e., volatility = 1, 2, 3, 4, 5, 6, 7, or 8); the second set included schools for which volatility = 1, 2, 3, 4, 5, 6, or 7; and so on to the final set of analyses involving the 104 least volatile schools (i.e., volatility = 1). (Again, there were separate volatility ratings for math and reading.) If, in fact, the poverty-size interaction is a statistical artifact due to the lower reliability of school-level achievement among smaller schools, then this interaction should attenuate with successively less-volatile schools—and be negligible for schools having the least volatility.

Results

I begin by portraying the achievement volatility among these schools and, in turn, the relationship between this volatility and school size. To investigate the statistical-artifact hypothesis, I then conduct the regression analyses on successively less-volatile schools.

The Volatility of School-Level Achievement

As described above, I estimated a school's volatility in eighth-grade achievement by first calculating the difference in mean achievement from 2003-2004 to 2002-2003 for reading and for mathematics. Among these 216 schools, the change in achievement from one year to the next ranged from roughly -17 to $+17$ MEA points in reading ($M = -1.56$, $SD = 4.61$) and, for math, -19 to $+16$ MEA points ($M = +1.14$, $SD = 4.79$).

The well-established relationship between school size and achievement volatility is clearly evident in the present data (Figure 2). Again, there simply is greater volatility—less stability—of school-level achievement among smaller schools than among larger schools. This also can be seen in the correlation between school size and the absolute value of a school's change in achievement from one year to the next: $r_s = -.31$ and $-.29$ for reading and math, respectively. In short, Figure 2 and these two correlations underscore the relevance of the statistical-artifact hypothesis that frames the present study.

The distribution of the 8-point volatility ratings, formed from the absolute value of a school's change in achievement from one year to the next, are shown in Figure 3 for both reading and math. Each distribution reflects extreme positive skew: While the vast majority of these 216 schools demonstrated rather stable levels of achievement (± 5 points from one year to the next), some schools evinced wide swings in this regard. Only one school fell in the highest volatility category for mathematics achievement; none did for reading achievement.

Regression Analyses: All Schools

The first set of regression analyses is based on all schools, irrespective of their volatility in achievement. Table 1 presents descriptive statistics for reading, math, poverty, and school size. Not surprisingly, schools vary considerably with respect to both poverty and size: Some schools have as few as 3 students per grade and 3% of their students receiving subsidized meals, whereas other schools have as many as 358 students per grade and 84% of their students on subsidized meals. Reading and math correlate highly ($r = .74$), as one would expect, and each correlates with poverty in the customary fashion (Sirin, 2005; White, 1982). There is some tendency for smaller schools to be located in more impoverished communities ($r = -.34$). However, school size is unrelated to achievement ($r = .07, p = .16$).

Reading. Table 2 shows the regression results for reading. Poverty significantly and independently predicts reading at Step 1, whereas the corresponding effect of school size falls short of statistical significance. An additional 2.2% of the variance in reading is explained by the introduction of the product term at Step 2, which, consistent with prior research, shows a statistically significant interaction between poverty and school size ($p = .013$).

Because the poverty-size interaction presently enjoys so much attention in the rural education literature, elaboration on the meaning of the various coefficients reported at Step 2 may be helpful to some readers. As we saw above, Step 2 estimates the effects for the full equation, $\hat{Y} = a + b_1X_1 + b_2X_2 + b_3X_1X_2$, where the last term, $b_3X_1X_2$, reflects the possible interaction of poverty and school size. As Aiken and West (1991) explain, b_1 is the reading-on-poverty slope for schools having a per-grade enrollment equal to the mean (i.e., centered $X_2 = 0$). For schools of average size, then, reading achievement decreases .127 MEA points ($b_1 = -.127$) with every one-percentage-point increase in the students receiving subsidized meals. In standardized terms, this corresponds to a decline in reading achievement of roughly half a

standard deviation ($\beta_1 = -.54$) for each standard deviation increase in poverty (again, for schools of average size). One interprets b_2 analogously: For schools at the mean for poverty, reading achievement decreases .008 MEA points ($b_2 = -.008$) for each one-student increase in school size—an achievement decline of 16% of a standard deviation ($\beta_2 = -.16$) for each standard deviation increase in school size.

The statistical significance of b_3 signals the presence of interaction between poverty and school size. Specifically, the negative coefficient for the product term X_1X_2 , coupled with the negative coefficient for poverty, means that the *simple slope* for poverty—i.e., the reading-on-poverty slope at a specified value of school size—is steeper (more negative) for larger schools than it is for smaller schools.

The concept of simple slope is central to interpreting a statistically significant interaction. The simple slope for poverty derives from the full equation, $\hat{Y} = a + b_1X_1 + b_2X_2 + b_3X_1X_2$, which, when recast as the Y -on- X_1 regression at a specified value of X_2 , looks like this:

$\hat{Y} = (a + b_2X_2) + (b_1 + b_3X_2)X_1$. The critical term here is $(b_1 + b_3X_2)$, which is the Y -on- X_1 slope for the specified value of X_2 (expressed as a deviation from the centered mean of zero). Select a deviation score to represent X_2 , plug this value into the expression $(b_1 + b_3X_2)$, and you have the simple slope for poverty at a particular school size.

For example, consider a school having 16 students per grade—the 25th percentile in school size and roughly 57 fewer students than the mean ($\bar{X}_2 = 72.78$). The simple slope for schools of this size is $b_{.57} = -.098$, which corresponds to a standardized regression coefficient of $\beta_{.57} = -.41$.² Thus, with each standard deviation increase in poverty, reading achievement in

² In symbolizing this simple slope, I introduce the subscript $-.57$ to make explicit the particular value of X_2 at which the Y -on- X_1 slope is estimated. The specified value of X_2 is expressed as a deviation score: $X_2 - \bar{X}_2 = 16 - 72.68 = -56.58$ (rounded to $-.57$ here). Subscripts for other

these smaller schools decreases approximately 40% of a standard deviation. The simple slope is slightly steeper for schools having 42 students per grade (the median school size, or 50th percentile): $b_{.31} = -.111$ or, in standardized terms, $\beta_{.31} = -.47$. Now consider a school falling at the 75th percentile in school size, or 105 students per grade. Here, the unstandardized and standardized simple slopes are $b_{+.32} = -.144$ and $\beta_{+.32} = -.61$, respectively. For these larger schools, then, reading decreases approximately 60% of a standard deviation with each standard deviation increase in poverty. Consistent with the statistically significant interaction of poverty and school size, simple slopes estimated at various levels of school size illustrate that reading achievement is increasingly related to poverty as school size increases, and decreasingly related to poverty as school size decreases.

Figure 4 shows the within-group regression lines. As described above, I obtained these by splitting the school-size distribution at the median (42 students per grade) and, for each group of schools, fitting a reading-on-poverty regression line. These within-group regression lines further illustrate the interaction reported in Table 2: There is a flatter slope—a weaker relationship between reading achievement and poverty—for smaller schools than for larger schools. Indeed, the correlation for the former is $r = -.39$ versus $r = -.64$ for the latter, which, when squared, yield power ratings of 15% and 41%, respectively. Although there is considerable within-group variability evident in Figure 4 and, further, the nonparallel displacement of one regression line relative to the other is not great (particularly where most of the data are), there is some tendency for smaller higher-poverty schools to have reading achievement superior to that of larger higher-poverty schools.

simple slopes follow suit. Further, to minimize rounding error, I calculated simple slopes using the multi-digit values reported by the statistical software. For example, $b_{.57} = b_1 + b_3X_2 = -.127330602 + (-.000511659)(-57) = -.098$; similarly, $\beta_{.57} = b_1(S_1/S_Y) = (-.0982)(16.6250893/3.93549624) = -.41$.

Math. Table 3 shows the regression results for math, based on all schools. The pattern of results is similar to those obtained for reading. At Step 1, poverty is significantly related to math whereas school size is not ($p = .31$). And at Step 2, the interaction of poverty and school size explains an additional 5% of variance in mathematics achievement ($\Delta R^2 = .048, p = .001$): As with reading achievement, mathematics achievement is increasingly related to poverty as school size increases, and decreasingly related to poverty as school size decreases. For example, the math-on-poverty slope for median-size schools is $b_{.31} = -.086$ ($\beta_{.31} = -.33$). In contrast, the simple slope for schools at the 25th percentile in school size $b_{.57} = -.064$ ($\beta_{.57} = -.25$) and, for schools at the 75th percentile, $b_{+.32} = -.139$ ($\beta_{+.32} = -.53$).

The within-group regression lines for below- and above-median schools in per grade enrollment are presented in Figure 5, which shows the nonparallel displacement indicative of interaction. The math-on-poverty slope is flatter—signifying a weaker relationship—for smaller schools than for larger schools. The corresponding power ratings are, respectively, 4% for smaller schools ($r = -.19$) and 46% for larger schools ($r = -.68$).

The symmetry of b_3 . As noted above, the statistical significance of b_3 indicates that the magnitude of the achievement-on-poverty slope (b_1) is a function of school size (X_2) and, symmetrically, the magnitude of the achievement-on-size slope (b_2) is a function of poverty (X_1). My emphasis thus far has been decidedly on the former, given its direct relevance to the concept of poverty's power rating which frames the present study. But many writers blur the distinction between the two interpretations, referring to one and then to the other as their argument develops. Therefore, so (briefly) shall I.

Just as the simple slope for poverty (b_1) at a specified value of school size (X_2) is equal to $b_1 + b_3 X_2$, the simple slope for school size (b_2) at specified value of poverty (X_1) is equal to

$b_2 + b_3X_1$.³ At Step 2 of Tables 2 and 3, we see that school size has a negligible, if statistically significant, negative effect on both reading and math for schools of average poverty (i.e., $X_1 = 0$).⁴ But when the simple slope is calculated for a school where 23% of its students receive subsidized meals—approximately one standard deviation, or 17 percentage points, below the mean ($\bar{X}_1 = 39.52$)—school size is unrelated to achievement in either reading or math.

Specifically, $b_{.17} = .001$ and $\beta_{.17} = .01$ ($p = .91$) for reading; for math, $b_{.17} = .006$ and $\beta_{.17} = .11$ ($p = .20$). Now consider a comparatively high-poverty school in which 73% of students receive subsidized meals (roughly two standard deviations, or 33 percentage points, above the mean). Here, the effect of school size on reading is statistically significant and large: $b_{+33} = -.025$ and $\beta_{+33} = -.49$ ($p = .003$). For math, the effect is larger still: $b_{+33} = -.035$ and $\beta_{+33} = -.63$ ($p < .001$). Thus, with a standard deviation decrease in school size, reading achievement in these high-poverty schools—unlike their lower-poverty counterpart—increases by half a standard deviation, and math achievement increases almost two-thirds of a standard deviation. This finding, of course, merely restates the poverty-size interaction by focusing on the conditional effect of school size rather than the conditional effect of poverty.

Regression Analyses: Successively Less-Volatile Schools

To explore the possible operation of a statistical artifact due to the greater volatility in achievement among smaller schools, I repeated the regression analyses reported above for successively less-volatile collections of schools. Rather than exhaustively delineate these results for each value of the volatility measure, I report in Table 4 the primary statistic for each analysis:

³ As before, this derives from the full equation, $\hat{Y} = a + b_1X_1 + b_2X_2 + b_3X_1X_2$, which, when now reformulated as the Y -on- X_2 regression at a specified value of X_1 , is $\hat{Y} = (a + b_1X_1) + (b_2 + b_3X_1)X_2$.

⁴ Just as b_1 is estimated at $X_2 = 0$, b_2 is estimated at $X_1 = 0$.

the increment in R^2 at Step 2 when the product term, X_1X_2 , is introduced. I then provide additional details for the results based on the 104 least-volatile schools.

Reading. As Table 4 shows, the interaction between poverty and school size is unrelated to the volatility of school-level achievement: For each successive analysis, the increment in explained variance associated with the introduction of the product term at Step 2 is statistically significant. Although I did not conduct a trend analysis on the seven ΔR^2 values, there is no evidence that ΔR^2 , statistical significance notwithstanding, is systematically smaller when based on successively less volatile schools.

Tables 5 and 6 show, respectively, descriptive statistics and regression results for reading, based on the least-volatile schools in reading achievement ($n = 104$). Again, these are the schools for which mean achievement on the reading measure did not vary more than 2.5 points across the two years examined. The pattern of results here is similar to that reported earlier for all 216 schools, as are the within-group regression lines shown in Figure 5. Indeed, regarding the latter, poverty's power rating differential—16% for smaller schools vs. 42% for larger schools—is almost indistinguishable from the differential based on all schools (15% and 41%, respectively). With respect to reading achievement, then, my statistical-artifact hypothesis is not consistent with the data.

Math. A different picture emerges with mathematics achievement, where we see a gradual decline in ΔR^2 with successively less-volatile schools (Table 4)—to the point of statistical nonsignificance when based on the 104 least-volatile schools ($\Delta R^2 = .014$, $p = .193$). Tables 7 and 8 present the relevant statistics for the latter analysis, where, at Step 2 of Table 8, we see the statistically nonsignificant slope for the product term.

The within-group regression lines are shown in Figure 7. While the power ratings of poverty show some differential between smaller and larger schools, it derives from a poverty-

size interaction that failed to reach statistical significance and, therefore, reflects only chance variation. Between the general decline in ΔR^2 values (Table 4) and the absence of a statistically significant poverty-size interaction when based on the least volatile schools (Table 8), the hypothesis of statistical artifact in the case of mathematics achievement is consistent with the data.

Discussion

“Substantive finding or statistical artifact?” is the question posed in the subtitle of my paper, to which I can now answer with an ineluctably facetious “Yes!” When the dependent variable is reading achievement, I find no support for my hypothesis that poverty’s power rating is lower in smaller schools because of their greater volatility (lower reliability) in achievement. Thus, the celebrated interaction of socioeconomic status and school size clearly stands with respect to eighth-grade reading achievement in these Maine schools. But for mathematics achievement, the statistical-artifact hypothesis is supported. For eighth-grade mathematics achievement, poor reliability appears to be a plausible explanation of the reduced power rating of poverty among these smaller schools.

Unfortunately, the latter conclusion is complicated by plausible rival hypotheses of its own—an inevitable consequence of correlational research. Two problems immediately come to mind. First, my achievement-volatility measure does not distinguish between random variation and variation due to educational practice. Some of the high-discrepancy schools in Figure 2, as reflected in their alignment on the vertical axis, doubtless are revealing real—not random—improvement or decline in achievement. By treating all variation as *random* variation, I inevitably exclude some schools from the analysis that should have been included (were it possible to make this distinction in practice). That said, the results are not systematically biased

as a consequence, insofar as the absence of “real improvement” schools is offset by the absence of “real decline” schools, particularly regarding the poverty-size interaction.

The second problem is of greater concern. By conducting the regression analyses on successively less-volatile collections of schools, and because achievement volatility is more pronounced among smaller schools (Figure 2), I successively compromise the full representation of small schools as well. In short, I arguably exclude some of the very schools required for a fair test of my statistical-artifact hypothesis (and, in doing so, introduce a certain irony into the present study). We see the extent of this sacrifice in Figure 8, which shows the school-size distribution for all 216 schools and for the 104 least-volatile schools. Although both distributions have the expected positive skew, there are proportionately fewer small schools in the restricted sample than in the full sample. Consistent with this visual impression, the school-size mean and median are both higher in the restricted sample, and the coefficient of variation is smaller.

Yet this second problem—the successive underrepresentation of small schools—had no effect on the viability of the poverty-size interaction for reading achievement. This inconsistency presents an interesting challenge: how to explain it. If one is inclined to dismiss my findings for mathematics achievement because of this underrepresentation, then the challenge is to explain why a similar outcome was not obtained for reading achievement. After all, small-school underrepresentation operates there as well. So, what is it about reading achievement (or related instruction) that makes the poverty-size interaction immune to the successive underrepresentation of small schools in these analyses? Or, if one prefers, what is it about mathematics achievement (or related instruction) that makes the poverty-size interaction particularly vulnerable in this regard?

On the other hand, for those whose confidence in the statistical-artifact results for mathematics achievement is unshaken by the underrepresentation problem—after all, the bottom distribution in Figure 8 still shows positive skew and healthy variance—the corresponding challenge is to explain why the statistical-artifact hypothesis did not prevail for reading achievement. After all, reading achievement is not appreciably less volatile than mathematics achievement. So, what is it about reading achievement (or related instruction) that explains this apparent invincibility—a greater robustness—of the poverty-size interaction?

Unfortunately, I cannot answer these questions. At least not yet. But insofar as I cannot explain, even with the benefit of hindsight, a statistical-artifact finding that would surface only for mathematics achievement, I am inclined to attach greater import to the successive underrepresentation of small schools in these analyses than I had at the outset. Although I cannot explain why this underrepresentation has no concomitant effect on the poverty-size interaction with respect to reading achievement, this anomaly presently perplexes me less than does a mathematics-specific statistical artifact. Furthermore, it is only in the final, most restrictive analysis—where a sizeable number of small schools are lost—that the poverty-size interaction for mathematics achievement fails to reach statistical significance (Table 4).

In view of these considerations, then, I conclude that my results are insufficient to support the statistical-artifact hypothesis with respect to mathematics achievement. Although this conclusion is not as unequivocal as that for reading achievement, I nevertheless believe it is the reasonable conclusion given the considerations above. In short, the celebrated interaction of poverty and school size has survived a sincere attempt to empirically cast doubt on it.

Consequently, we can have greater confidence in this interaction than was warranted before.

Implications for Research

First, further tests of the statistical-artifact hypothesis would be informative, if only to show that my somewhat equivocal results for mathematics achievement are a mere anomaly. Replications should involve Maine data drawn from other years, but replications also should involve data beyond Maine. In this spirit, I am hopeful that other researchers who have explored the mitigating-effect phenomenon will, where possible, conduct (re)analyses of their own with the inclusion of an achievement-volatility measure.

Second, if we are inclined to take the interaction of socioeconomic status and school size as an established phenomenon, we nonetheless are left wanting for a credible explanation of it. Such an explanation seemingly would draw on the mechanisms through which smaller schools facilitate student achievement and related outcomes, but, unfortunately, we are wanting there as well. As Fowler and Walberg (1991) said in reference to the then-extant research,

“[a]lthough these studies show a positive relationship between small school size and student outcomes, they do not suggest *why* this may occur. In other studies, which only peripherally included school size, researchers have *suggested* reasons for the beneficial effect that small school size has upon student outcomes” (p. 191; emphasis added).

A decade later, Howley (2002) offered a similar conclusion in his synthesis of the school size research:

Many, many other unanswered questions exist. For instance, *why* is smaller school size (variously defined) associated with higher . . . levels of achievement for individuals, schools, and districts? Hypotheses abound, with most having to do with the care, attention, and respect enabled by smallness in the conduct of personal relations. (p. 62; emphasis in original)

As an influence on student achievement, school size clearly is a proxy rather than a causal force in and of itself. To offer credible explanations for the poverty-size interaction, then, we first need stronger evidence regarding the mechanisms—the mediating variables—through which school size putatively influences student achievement (McMillen, 2004, p. 20). Howley (2002, p. 62) refers to “care, attention, and respect.” Lee and her colleagues refer to “the academic and social organization and functioning of schools” (Lee & Smith, 1997, p. 219). Doubtless there are other context- and process-related forces at play as well. Whatever the focus, a warranted claim about its relationship to both school size and student achievement must be based on careful empirical investigation, not on casual observation, anecdotal reports, reasonable (but untested) hypotheses, popular opinion, or the will to believe. We need additional descriptive research like that conducted by Howley and Howley (2006) and Lee, Smerdon, Alfeld-Liro, and Brown (2000), which should be followed up by analyses that exercise the statistical control necessary to test hypotheses that fundamentally get at cause-and-effect relationships.

Equipped with empirically established mediating variables regarding the relationship between school size and student achievement, we can then craft defensible conjectures regarding the poverty-size interaction. In this regard, of course, one’s central obligation will be to argue *why* a mediating variable would be expected to differentially affect student achievement as a function of student SES. For example, if the accumulation of evidence from sound empirical research were to show that smaller schools are characterized by more personalized social relations and, in turn, that these more personalized social relations improve student achievement, our obligation is to cogently argue why lower-SES students would benefit from such social relations more than higher-SES students would. These conjectures should then be subjected to empirical tests of their own. For example, one could introduce a set of social-relations variables

into the full regression equation (in the tradition above) to see whether the poverty-size interaction disappears—as it would if the poverty-size interaction is in fact due to social relations.

In any case, well-crafted arguments followed by equally well-crafted investigations—both premised on warranted claims regarding the mechanisms through which school size influences student achievement—should be the direction of future research on the poverty-size interaction.

References

- Aiken, L. S., & West, S. G. (1991). *Multiple regression: Testing and interpreting interactions*. Newbury Park, MA: Sage.
- Butler, L., Carr, R., Cook, K., Gould, R., Keenan, F., MacArthur, S., Power, L., & Ritchie, K. (2005, June 8). Consolidation can't save money. [Letter to the editor.] *Bangor Daily News*, p. A9.
- Coladarci, T. (2003). *Gallup goes to school: The importance of confidence intervals for evaluating "Adequate Yearly Progress" in small schools*. A policy brief of the Rural School and Community Trust. Washington, D.C.: The Rural School and Community Trust. Available online at <http://files.ruraledu.org/docs/nclb/coladarci.pdf>
- Fowler, W. J., & Walberg, H. J. (1991). School size, characteristics, and outcomes. *Educational Evaluation and Policy Analysis*, 13, 189-202.
- Friedkin, N. E., & Necochea, J. (1988). School system size and performance: A contingency perspective. *Educational Evaluation and Policy Analysis*, 10, 237-249.
- Hill, R. K., & DePascale, C. A. (2003). Reliability of No Child Left Behind Accountability Designs. *Educational Measurement: Issues and Practices*, 22(3), 12-20.
- Howley, C. B. (1996). Compounding disadvantage: The effects of school and district size on student achievement in West Virginia. *Journal of Research in Rural Education*, 12(1), 25-32.
- Howley, C. B. (2002). Small schools. In A. Molnar (Ed.), *School reform proposals: The research evidence* (pp. 49-77). Greenwich, CT: Information Age Publishing.
- Howley, C. B., & Bickel, R. (1999). *The Matthew Project: National Report*. (ERIC Document Reproduction Service No. ED433174)

- Howley, A., & Howley, C. B. (2006). Small schools and the pressure to consolidate. *Education Policy Analysis Archives*, 14(10). Retrieved March 29, 2006 from <http://epaa.asu.edu/epaa/v14n10/>.
- Huang, G., Howley, C. B. (1993). Mitigating disadvantage: Effects of small-scale schooling on student achievement in Alaska. *Journal of Research in Rural Education*, 9(3), 137-149.
- Johnson, J. D., Howley, C. B., & Howley, A. A. (2002). *Size, excellence, and equity: A report on Arkansas schools and districts*. (ERIC Document Reproduction Service No. ED459987)
- Kane, T. J., Staiger, D. O., & Geppert, J. (2002). Randomly accountable. *Education Next*, 2(1), 57-61. Retrieved January 15, 2003, from <http://www.educationnext.org/20021/56.html>
- Lee, V. E., Smerdon, B. A., Alfeld-Liro, C., & Brown, S. L. (2000). Inside large and small high schools: Curriculum and social relations. *Educational Evaluation and Policy Analysis*, 22, 147-171.
- Lee, V. E., & Smith, J. B. (1997). High school size: Which works best and for whom? *Educational Evaluation and Policy Analysis*, 19, 205-227.
- Linn, R. L., & Haug, C. (2002). Stability of school-building accountability scores and gains. *Educational Evaluation and Policy Analysis*, 24(1), 29-36.
- McMillen, B. J. (2004, October 22). School size, achievement, and achievement gaps. *Educational Policy Analysis Archives*, 12(58). Retrieved October 22, 2004, from <http://epaa.asu.edu/epaa/v12n58/>
- Rural School and Community Trust. (2005). *Maine's smaller schools cut poverty's power over student achievement*. [January 11, 2005 news release.]
- Sirin, S. R. (2005). Socioeconomic status and academic achievement: A meta-analytic review of research. *Review of Educational Research*, 75, 417-453.

Thorndike, R. L. (1982). *Applied psychometrics*. Boston: Houghton Mifflin.

Tompkins, R. (2006). Small schools, small districts: Good for rural kids, economies, and democracy. *Rural Americans* (Issue 14). Retrieved February 22, 2006, from <http://www.demos.org/aroundthekitchentable/article.cfm?edition=014&article=education>

White, K. (1982). The relation between socioeconomic status and academic achievement. *Psychological Bulletin*, *91*, 461-481.




Figure 1. The relationship between (a) the number of fourth-grade students tested in a school and (b) the one-year change in the proportion of students who met or exceeded the standard. (Source: Coladarci, 2003, Figure 4)





Figure 2. School size and the volatility of achievement in reading (top) and mathematics (bottom).

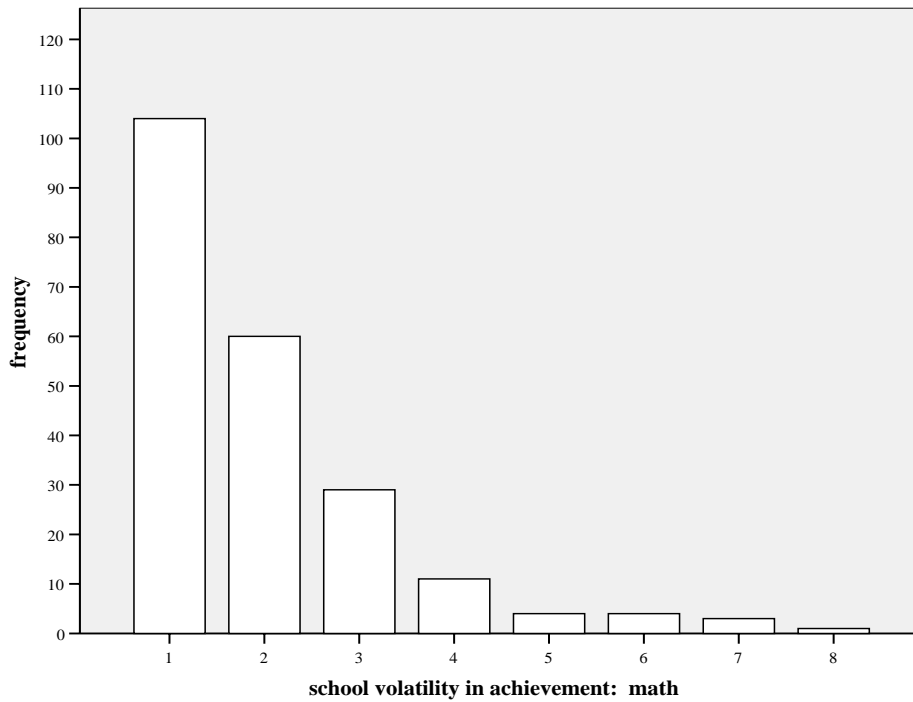



Figure 3. The distribution of volatility in school achievement: reading (top) and math (bottom.)

Table 1. Descriptive statistics: All schools ($n = 216$).

	<i>M</i>	<i>SD</i>	range	intercorrelations		
				reading	math	poverty
reading	535.96	3.94	522.72, 547.69			
math	528.16	4.36	514.51, 542.17	.74*		
poverty	39.52	16.63	2.68, 83.86	-.48*	-.37*	
school size	72.78	77.31	2.94, 358.00	.07	.07	-.34*

Note. For the purpose of this table, poverty and school size are in their original uncentered form (which affects only the mean and range).

* $p < .01$.

Table 2. Regressing reading on poverty, school size, and their product:
All schools ($n = 216$).

	b	s.e.	β	t	p	ΔR^2
Step 1: (constant)	535.962					
poverty	-.122	.015	-.51	-8.07	< .001	
school size	-.006	.003	-.11	-1.71	.089	
Step 2: (constant)	535.738					
poverty	-.127	.015	-.54	-8.45	< .001	
school size	-.008	.003	-.16	-2.40	.017	
poverty x size	-.001	.0002	--	-2.52	.013	.022

Note. Poverty and school size were centered for this analysis.




Figure 4. The interaction of poverty and school size ($p = .013$), reading: All schools ($n = 216$).

Table 3. Regressing math on poverty, school size, and their product:
All schools ($n = 216$).

	b	s.e.	β	t	p	ΔR^2
Step 1: (constant)	528.161					
poverty	-.103	.018	-.39	-5.78	< .001	
school size	-.004	.004	-.07	-1.02	.310	
Step 2: (constant)	527.796					
poverty	-.112	.017	-.43	-6.40	< .001	
school size	-.008	.004	-.14	-2.05	.042	
poverty x size	-.001	.0002	--	-3.53	.001	.048

Note. Poverty and school size were centered for this analysis.




Figure 5. The interaction of poverty and school size ($p = .001$), math: All schools ($n = 216$).

Table 4. Volatility in school achievement and the magnitude of ΔR^2 .

Reading				Math			
Volatility	<i>n</i>	ΔR^2	<i>p</i>	Volatility	<i>n</i>	ΔR^2	<i>p</i>
≤ 8	–	–	–	≤ 8	216	.048	.001
≤ 7	216	.022	.013	≤ 7	215	.047	.001
≤ 6	214	.022	.014	≤ 6	212	.042	.001
≤ 5	208	.029	.005	≤ 5	208	.039	.002
≤ 4	204	.029	.006	≤ 4	204	.038	.002
≤ 3	188	.030	.008	≤ 3	193	.026	.011
≤ 2	166	.042	.002	≤ 2	164	.027	.018
1	104	.031	.027	1	104	.014	.193

Note. ΔR^2 is associated with the introduction of the product term (poverty x size) at Step 2 of each regression analysis.

Table 5. Descriptive statistics: Least volatile schools, reading achievement ($n = 104$).

	<i>M</i>	<i>SD</i>	range	intercorrelations	
				reading	poverty
reading	535.95	3.76	527.99, 545.95		
poverty	38.78	15.98	2.68, 78.52	-.59*	
school size	89.19	79.67	2.94, 358.00	.09	-.35*

Note. For the purpose of this table, poverty and school size are in their original uncentered form (which affects only the mean and range).

* $p < .01$.

Table 6. Regressing reading on poverty, school size, and their product: Schools having minimal volatility in achievement ($n = 104$).

	b	s.e.	β	t	p	ΔR^2
Step 1: (constant)	535.945					
poverty	-.149	.020	-.63	-7.45	< .001	
school size	-.006	.004	-.14	-1.61	.110	
Step 2: (constant)	535.721					
poverty	-.141	.020	-.60	-7.067	< .001	
school size	-.007	.004	-.16	-1.875	.064	
poverty x size	-.001	.0002	--	-2.237	.027	.031

Note. Poverty and school size were centered for this analysis.




Figure 6. Interaction of poverty and school size ($p = .001$), reading: Schools having minimal volatility in achievement ($n = 104$).

Table 7. Descriptive statistics: Least volatile schools, math achievement ($n = 104$).

	<i>M</i>	<i>SD</i>	range	intercorrelations	
				math	poverty
math	527.60	4.26	514.51, 542.17		
poverty	38.25	14.71	7.99, 73.89	-.41*	
school size	82.28	81.72	3.39, 327.50	.06	-.30*

Note. For the purpose of this table, poverty and school size are in their original uncentered form (which affects only the mean and range).

* $p < .01$.

Table 8. Regressing math on poverty, school size, and their product:
Schools having minimal volatility in achievement ($n = 104$).

	<i>b</i>	s.e.	β	<i>t</i>	<i>p</i>	ΔR^2
Step 1: (constant)	527.475					
poverty	-.125	.028	-.43	-4.54	< .001	
school size	-.004	.005	-.07	-.76	.448	
Step 2: (constant)	527.305					
poverty	-.131	.028	-.45	-4.709	< .001	
school size	-.007	.006	-.13	-1.265	.209	
poverty x size	-.001	.0004	--	-1.309	.193	.014

Note. Poverty and school size were centered for this analysis.




Figure 7. No interaction of poverty and school size ($p = .193$), math: Schools having minimal volatility in achievement ($n = 104$).





Figure 8. Variability in school size: All schools ($n = 216$) versus least-volatile schools ($n = 104$).