Performance Standards and Educational Cost Indexes: You Can't Have One Without the Other

William Duncombe
John Yinger
Center for Policy Research
The Maxwell School
Syracuse University

March 1998

This paper was originally prepared for a conference sponsored by the Committee on Education Finance, National Research Council, Irvine California, January 30-31, 1998. The authors are grateful to Janet Hansen, Sunny Ladd, Bob Schwab, Steve Sugarman, and other participants in that conference for comments on the original draft.

1. Introduction

Performance standards have been at the center of recent debate on educational reform. Many states have implemented new performance standards, often based on student test scores, and a district's state aid is sometimes linked to its success in meeting the standards (see Clotfelter and Ladd, 1996). National politicians have debated the merits of a nationwide testing program, which is a way to obtain comparable performance indicators across states. Several state supreme courts have ruled that their state constitution requires a system enabling all school districts to reach a minimum performance level (see Minorini and Sugarman 1997).

This focus on performance is designed to promote better performance, particularly in school districts that are currently not performing well, by holding school districts accountable. The trouble is that a district's performance is influenced not only by the actions of its administrators and teachers but also by factors outside of its control, such as the nature of its student body. A recent article in *The New York Times* expresses this concern very clearly. In a discussion of report cards and school rankings, now used in 35 states, this article points out that

> because such rankings are often based exclusively on test scores, which give only a partial snapshot of a school's performance, some educators worry that schools may be unfairly blackballed, especially those with high populations of poor children (Steinberg 1998).

Thus, a focus on performance is inevitably unfair unless it can somehow account for the impact on performance of factors that are outside the control of school officials. Without such an accounting, some schools get credit for favorable conditions that were not of their making and other schools get blamed for unfavorable conditions over which they have no control. In order to be fair, school report cards and performance-based state aid systems must distinguish between poor performance based on external factors and poor performance based on school inefficiency. Many state aid systems have taken one step in this direction by compensating districts that have poor performance because of their low wealth, a factor over which they have no control. However, school district performance is also influenced by the cost of education, which varies widely from district to district based on wage rates, student characteristics, and other factors that are outside the control of district officials. Existing state aid formulas either ignore these factors altogether or else use ad hoc corrections, such as “weighted pupil” counts, that account for them only partially at best.

In this paper, we explain in detail why a performance focus and educational cost indexes must go hand in hand, discuss the alternative methods for estimating educational cost indexes, and show how these costs indexes can be incorporated into a performance-based state aid program.

2. The Conceptual Foundations of Educational Cost Indexes

An educational cost index is designed to measure how much a school district would have to spend, relative to the average district, to obtain any given performance target. Some scholars, including Guthrie and Rothstein (1997), have used the term “cost index” to refer only to differences in input prices across districts. However, we use the term to refer to a comprehensive accounting of the reasons why some districts must spend more than others to achieve any given performance level—a definition that, as we will show, involves far more than just input prices. After an introductory example, this section explains the relationship between performance measures and cost indexes, discusses, in general terms, the factors that influence educational costs. The following section reviews alternative methods for estimating educational cost indexes.

A. An Introductory Example

Before turning to educational performance and costs, it may prove helpful to explore an example from everyday life that involves the same concepts. The example is the service of providing comfortable shelter. A natural measure of performance for this service is the indoor temperature.

This is not, of course, the only dimension of comfort. In some contexts, one might want to know about the extent to which rain leaks in; in others, humidity might be a concern. As with any performance standard, however, some simplification is necessary, and the indoor temperature appears to measure the dimension of comfort that is of most concern to most people in this country. Thus we will focus on a temperature as a measure of performance.

Now suppose some national consumer group interested in comfortable shelter sets a performance standard of 72 degrees. A natural question to ask is: How much would it cost to achieve this standard in different parts of the country? To answer this question, one must consider the technology with which comfort is provided. This technology is straightforward: a household must purchase various inputs, such as a furnace and natural gas and insulation, that can deliver or preserve heat or cold and thereby provide an indoor temperature that is different from the outdoor temperature. The impact of these inputs depends, of course, on the outdoor temperature, as well as on other environmental factors, such as the wind.

The cost of achieving the comfort standard therefore depends on two factors. First, it depends on the prices of inputs. The price of natural gas is not the same in all parts of the country, for example, and people in high-price places must pay more to obtain the same amount of natural gas. Second, it depends on the environment. During the winter, people in Minneapolis face much lower outdoor temperatures than do people in San Diego, so it costs them more to bring the indoor temperature up to the standard. To put it another way, people in Minneapolis must purchase more inputs to reach the standard. In another season (or for other cities), one obviously would have to consider the cost of bringing down an outdoor temperature above 72 degrees. A comfort cost index that reflected only gas price differences could be developed, but it obviously would provide an incomplete accounting of costs because it would ignore the outdoor temperature.

It follows that a comprehensive comfort cost index must consider both the prices of inputs and the harshness of the environment. As we will see, these are exactly the factors at work in determining the cost of education.

The factors discussed so far are all outside a household's control; the same cannot be said for everything that influences comfort. Households may, for example, buy an inefficient furnace, use a relatively expensive type of fuel, neglect to have their furnace maintained properly, or neglect to install the proper weatherstripping. These decisions all affect how much a household spends to obtain comfort but they should not influence a measure of the cost of comfort, which, as we use the term, is based solely on factors outside household's control.

This leads us to another concept, namely efficiency. In some contexts, it may be helpful to measure inefficiency and to separate its impact on comfort spending from the effect of cost factors. However, it would be inappropriate to let household choices influence a cost index, so any method to obtain a cost index must be insulated from the effects of household choices.

The key question for our purposes is: How could one determine the cost of meeting the performance standard in various locations? One way to proceed would be to select a certain type of house with certain heating and insulation characteristics as a standard and then use engineering studies to determine the amount of fuel it would take to keep this house at 72 degrees under the average weather conditions that are experienced in each location. The cost of meeting the standard is the amount of fuel required multiplied by the cost of fuel in that location. This approach has the advantages that it directly accounts for household choices (by using a standard house) and that it can be based on extensive information on weather conditions. This approach also has the major disadvantage, however, that it requires a detailed engineering study that is usually not available.

An alternative approach, which is analogous to the educational cost indexes discussed below, is to gather information for a sample of locations on (1) average household fuel bills, (2) actual indoor temperatures obtained by households, (3) a few key measures of weather conditions, such as average temperature or heating degree days, (4) the cost of the main input, namely fuel, and (5) average household choices that might influence fuel efficiency, such as whether they maintain their furnace annually and whether they use the appropriate weather stripping. An analysis of this information using regression analysis, which is a simple statistical procedure, can then reveal the extent to which input costs and weather conditions affect spending for fuel, holding constant both actual temperatures and household choices. This analysis leads directly to a cost index, which is defined as the amount a household would have to spend in each location, relative to a household in the average location, to obtain a given indoor temperature, under the assumption that the household makes efficient choices about its heating/cooling system. If, for example, Minneapolis has a cost index of 200, households there would have to spend twice as much as households in the average city to obtain the same level of comfort. This approach is more abstract than the previous one in the sense that it does not consider all the details of heating technology, but it captures the main features of the problem and has the great advantage that it can be implemented with readily available data.

In combination, a comfort standard and a cost index reveal how much households in each community would have to spend to meet the standard, based on factors outside their control. In principle, this information also could be obtained from a production study conducted in each community, which would reveal the set of inputs needed to achieve the standard, along with information on each community's input prices. The regression approach provides the same information at much lower cost by determining how spending varies with input costs and environmental conditions, controlling for actual comfort outcomes.

B. Measuring Performance

With education, as with home comfort, one cannot set a performance standard without selecting a way to measure performance. To put it another way, one cannot determine whether a school has met a performance standard unless its performance can be observed and measured. Policy makers may wish to avoid this choice, because selecting a standard is inevitably somewhat controversial. No set of performance standards can capture all aspects of learning, and schools may respond to specific standards by "teaching to the test" or otherwise shifting their resources to meet the standard at the expense of other legitimate objectives. Nevertheless, this choice cannot be avoided. Any policy to enhance school performance involves, either explicitly or implicitly, a specific performance measure. The trick is to select performance measures that are rich enough to capture success in a range of educational activities. For the most part, the selection of a performance measure is based on the judgement of politicians and educational policy officials, perhaps with some input from scholars. The most common measure is based on some kind of test score, such as an average elementary reading or math score. A drop-out rate is another widely used measure at the high school level.

To set a performance standard, policy makers must select both a measure of performance and the level of performance school districts are expected to meet. For example, all school districts might be expected to achieve a certain average test score or to ensure that a certain percentage of their students score above some standard reference point on a certain test. Standards of this type can be set for a single indicator or for a set of indicators. For example, school districts might be expected to have a certain average test score and a certain graduation rate.

We have developed an alternative approach, which selects performance indicators on statistical grounds. In particular, this approach determines which performance indicators are valued by voters, as indicated by their correlation with property values and school spending. This approach, which is explained in detail in Duncombe, Ruggiero and Yinger (1996) and Duncombe and Yinger (1997), results in an index of educational performance. This index is a weighted average of the performance indicators that are found to be statistically significant, where the weights reflect the value voters place on each indicator.[1] In the case of New York state, this approach leads to an educational performance index based on three performance indicators: the average share of students above the standard reference point on the third- and sixth-grade PEP tests for math and reading, the share of students who receive a more demanding Regents diploma (which requires passing a series of exams), and the graduation rate. These indicators reflect a wide range of school district activities, including both elementary and secondary teaching and programs designed to promote student retention, and reflect the degree of success at both the high and low ends of the student performance distribution.

Although these indicators are identified by an objective, statistically based procedure, they do not, of course, summarize all educational activities by a school district. Like all other approaches to measuring performance, this approach makes the problem manageable through some simplification. Moreover, this approach results in a performance yardstick, but it cannot determine the point on the yardstick that school districts should be expected to meet. As with other approaches, this step, the selection of the performance target, must be based on the judgement of public officials.

C. Separating Factors In and Outside the Control of School Officials

Either indirectly, as in the case of district report cards, or directly, as in the case of a performance-based aid system, performance standards are intended to boost a school district's incentive to improve its performance. The problem, however, is that actual performance is influenced not only by the decisions of school officials but also by factors outside their control.

Thus, some districts find it easy to meet a standard even if they are very inefficient, whereas others find it impossible to meet a standard even if they are more efficient than other districts. It is neither fair nor effective for a state to reward districts that achieve high performance (or to punish districts that perform poorly) based on factors that are outside their control. Instead, state programs need to reward districts that perform well despite external obstacles, such as concentrated poverty, and punish districts that do not perform well despite favorable circumstances.

The everyday example presented earlier may help to make this point because it explains how external factors work. Just as some communities face relatively high oil prices and harsh weather, which raise the cost of meeting any comfort standard, some school districts face relatively high input prices (such as teacher salaries) and relatively harsh educational environments, which raise the cost of meeting any educational performance standard. Thus, the key to removing external factors is to calculate a educational cost index; as we use the term, such an index measures the impact of input and environmental costs, not just input prices. To ensure fairness across districts and to encourage performance improvements, a performance standard and other performance-related state programs should reward districts that perform well (and punish districts that perform poorly) **relative to other districts with the same costs**.

D. The Role of Input Prices

In the case of education, the most important input is teachers, so in constructing a cost index, it is vital to account for teachers' salaries. Secondary inputs, such as school facilities, also play a role in delivering education, but data on the prices of these inputs are generally not available. In some cases, data on administrators' salaries are available, but these salaries are so correlated with teachers' salaries that they add little to the analysis. Like almost all the literature, therefore, we will restrict our attention to the role of teachers' salaries.

A cost index is designed to measure the impact of factors outside the control of school officials. It is not appropriate, therefore, to directly use actual teachers' salaries in constructing a cost index because those salaries reflect both the generosity of the school district, a factor over which they have control, and the underlying labor market conditions, which cannot be influenced by school officials. A cost index should reflect the fact that some school districts are located in high-wage labor markets, where they must pay high salaries to attract people away from other school districts or away from the private market; and it should reflect the fact that the external conditions in some school districts are so harsh that teachers will not come there without receiving "combat pay"; but it should not reflect the fact that some school district administrators pay higher salaries than necessary to attract their teachers, because they are poor negotiators or for any other reason.

A cost index should not be affected by school districts' choices, so the influence of school officials on teachers' salaries poses a challenge to anyone who wants to construct a cost index. Fortunately, however, well known statistical procedures can separate the impact of school officials on teachers' salaries from the impact of external factors and produce a cost index based only on factors outside the control of school officials. These procedures are discussed in a later section.

E. The Role of Environmental Factors

The home comfort example reveals that the cost of meeting a performance standard depends not only on input prices but also on the environment in which the relevant services are provided. This lesson carries over into education, as school districts with a harsher educational environment must pay more to obtain the same educational performance as other districts. This section explains the impact of environmental factors on educational costs and shows how this impact can be estimated using widely available data.

The key role of environmental factors, also called fixed inputs in the literature, was first identified in the Coleman Report (1966), which showed that a student's performance on standardized tests depended not only on his or her own characteristics and family background but also on the characteristics and backgrounds of the students in his or her class. All else equal, for example, a student's performance declines as the share of classmates from poor households increases. This finding translates into a statement about educational costs. If performance declines as student poverty increases, then a district with a high poverty rate cannot achieve the same performance as a district with a low poverty rate without running programs (which, of course, cost money) to offset the impact of poverty.

The important role of environmental factors in educational production has been verified by dozens of studies. A review of many early studies is provided by Hanushek (1986). Good recent studies include Ferguson (1991), Ferguson and Ladd (1996), and Krueger (1997). The study by Ferguson and Ladd, for example, finds that a student's fourth-grade educational performance (on reading and math tests) is affected by, among other things, the share of students receiving a free lunch (a measure of poverty), the share of adults in the district with a college degree, a measure of student turnover, and district enrollment. These studies are the analog of the engineering studies that link detailed weather conditions and indoor comfort in each type of house.

Production studies focus on the impact of environmental factors on a measure of performance, such as a test score, holding constant inputs selected by the school, such as the student-teacher ratio. These studies imply that costs are higher in school districts with a harsher educational environment, but do not estimate cost differences directly. Moreover, the results of the production studies vary significantly, depending on the methodology, the quality of the data, and other factors, so that even if the results were translated into cost differences across districts, these differences would vary widely from one study to the next.

Another set of studies shifts the focus to educational costs. These studies, which are analogous to a study of spending for home heating across a sample of communities, determine the extent to which districts with a harsh educational environment, as measured by the characteristics of their students, must pay more to achieve the same performance as other districts, where performance is measured by a set of performance indicators. These studies include Bradbury et al. (1984), which looks at all local spending, including spending on education; Ratcliffe, Riddle, and Yinger (1990); Downes and Pogue (1994); Duncombe, Ruggiero, and Yinger (1996); and Duncombe and Yinger (1997). These studies build on a well known general treatment of environmental factors by Bradford, Malt, and Oates (1969).

At one level, these cost studies are equivalent to production studies; any statement about production can be translated into a statement about costs and vice versa. In practice, however, the cost approach has several advantages over the production approach as a tool for informing state education policies. First, cost studies focus on school districts, which are the focus of state policy, instead of on individual students.[2] Second, the cost approach makes it possible to examine a range of performance indicators simultaneously instead of one performance indicator at a time. Third, the data required to implement the cost approach is widely available. This does not imply, of course, that the cost approach resolves all the controversies about the nature of educational production that have been debated in production studies. It does imply, however, that the cost approach is a practical alternative to the production approach that makes it possible for state officials to design an educational finance system that accounts, in a reasonable way, using up-to-date data, for key features of educational production in their state.

Perhaps the most crucial feature is variation in the educational environment across school districts. Existing cost studies all demonstrate that a harsher educational environment, as characterized by high rates of poverty and single-parent families, for example, results in a higher cost to obtain any given performance level. Just as the harsh weather "environment" in Minnesota ensures that people who live there must pay more during the winter time than do people in San Diego to maintain their houses at a comfortable temperature, the harsh educational "environment" in some school districts, particularly big cities, ensures that those districts must pay more than other districts to obtain the same educational performance from their students.

State educational officials are sometimes aware that environmental factors matter. For example, a report on the status of the state's schools by New York State Education Department says that "Five indicators, each associated with poor school performance, are useful for identifying students at risk of educational disadvantage: minority racial/ethnic group identity, living in a poverty household, having a poorly educated mother, and having a non-English language background" (The University of the State of New York, 1997, p. 3). However, states' performance standards and state aid programs do not take account of these environmental factors in any systematic way. As a result, these programs are, as noted earlier, both unfair and ineffective.

The role of environmental factors also is widely ignored in the debate about the relative effectiveness of public and private schools. Existing studies focus on whether differences in the performance between students in private and public schools, if any, can be explained by the possibility that students who attend private schools (or their parents) are more motivated than students who attend public schools (or their parents). See, for example, Witte (1996) and Rouse (1997). This is, of course a vital issue, but for policy purposes, an equally important, and largely ignored, issue is whether existing differences in the performance of students in private and public school are due to environmental factors or to school policies. If, for example, performance in city public schools is lower than in private schools because those public schools have more concentrated poverty among their students, then sending all city public school children to private schools would only export their poverty, undermine the educational environment in private schools, and, perhaps, have no impact on student performance. To put it another way, a finding that some private schools perform better than some public schools even after accounting for differences in environment factors (and, of course, student motivation) gives no insight whatsoever into the impact on performance of a massive move away from public schools toward private schools, which would dramatically shift the educational environment in both types of schools. Thus, more research is clearly needed on the impact of environmental factors on the cost of private education.

3. Alternative Methods for Calculating Educational Cost Indexes

Several different methods for calculating educational cost indexes have been proposed by scholars. This section explores the strengths and weaknesses of several key methods, and it compares the indexes that result when each method is applied to data for New York state.

A. Input Prices

Some scholars have proposed that educational costs be measured with an index of input prices, usually just teachers' salaries. Because teachers are by far the most important input in producing educational performance, teachers' salaries do, indeed, have a major impact on educational costs. However, a teacher salary index, by itself, has three major flaws as a measure of educational costs.

First, teachers' salaries reflect differences in teachers' experience and education, which are associated with quality differences across teachers. One cannot claim that a school district has high costs whenever it decides to hire teachers with extensive experience or with graduate degrees. Ideally, salaries that apply to teachers of comparable quality should be compared.

Second, as noted earlier, teachers' salaries at a given quality level can be influenced by the decisions of school officials. A cost index is intended to measure factors outside the control of school officials, so it should not reflect their bargaining skill or their generosity to teachers. A cost index based solely on teachers' salaries will provide the misleading impression that generous school districts are forced to pay more than other districts to obtain the same performance, when in fact their higher spending is entirely of their own making.

Finally, a teacher salary index ignores the role of the environment altogether. A school district with a harsh educational environment must spend more than other districts to obtain the same performance, even if teachers' salaries were the same everywhere. Thus, an index based on teachers' salaries leaves out one of the key sources of variation in educational costs across districts, namely environmental factors, and is inherently unfair to districts with a relatively harsh educational environment.

B. Adjusted Input Prices

Some scholars have recognized the first two problems with a cost index based solely on teachers salaries and suggested an alternative approach based on predicted salaries. (See Chambers 1978, 1995; Wendling 1981.) This approach uses regression analysis to separate the impact on teachers' salaries of internal factors under district control from the impact of external factors, and then predicts salaries holding the internal factors constant. In a typical study, the internal factors include teacher experience, education, and certification, as well as the district's salary structure. The key external factors include the wage level in the surrounding labor market and the classroom environment that confronts teachers in each district, and then uses statistical procedures to predict wages on the basis of these external factors. This approach explicitly recognizes that conditions in some schools are so harsh that teachers must receive "combat pay" to work there. In other words, equally qualified teachers will not come to those schools unless they are paid more that they would be paid at other schools where the private wage scale is the same.

This approach solves the first two problems by constructing an index of predicted teachers' salaries. One might at first conclude that it also solves the third problem because it accounts for the impact of environmental factors on teachers' salaries. As explained earlier, however, environmental factors affect not only the price of inputs but also the quantity of inputs required.

This point can be illuminated by returning to the non-school example at the beginning of this paper. A notion similar to "combat pay" could arise in the provision of home comfort if the price of the key input, namely natural gas, depends on the weather.[3] Suppose, for example, that colder weather requires more maintenance of natural gas pipelines and hence leads to a higher price for gas in colder places. A cost index for comfort based on the price of natural gas would clearly capture this phenomenon. However, such a cost index would not capture the fact that, to obtain any given comfort standard, households in a colder climate not only must spend more per unit of gas but also must purchase more gas than households in a warmer climate. Similarly, to achieve any given performance standard, school districts with a harsh educational environment not only must pay more to attract teachers, but also must hire more teachers (or spend additional money on other educational programs) than schools with an average educational environment.

In short, a cost index based on teachers' salaries, even if it is predicted on the basis of external factors, including environmental ones, ignores an important source of variation in educational costs and is inherently unfair to districts with harsh educational environments.

C. Ad Hoc Cost Adjustments for Environmental Factors

Many state aid formulas include ad hoc adjustments for environmental factors. This type of adjustment is also proposed in the Wyoming context by Guthrie and Rothstein (1997). States may pay more for transportation in less dense districts, for example, or compensate for a concentration of students with disabilities or whose native language is not English. Some programs also provide more money to districts with more children in poverty. However, these programs inevitably are ad hoc, with no demonstrated connection between the environmental factors and educational costs.

The 1996 New York state aid programs, for example, include several provisions that could be interpreted as cost adjustments.[4] Operating aid, which provides 53 percent of the total aid paid to school districts, is based on the number of "weighted" pupils in a district. Pupils with extra weights include pupils in secondary school and pupils with "special education needs," defined as students who score below the minimum competency level on the third and sixth grade reading or math PEP tests. The first of these weighting factors is supported by some studies of school spending in other states (see, for example, Ratcliffe, Riddle, and Yinger 1990), which find a higher cost for high school than for elementary school students. However, it is not supported by our analysis of data for New York state, which finds no cost differences by grade. The second factor is undoubtedly correlated with cost variables, but we believe it is inappropriate to include a performance measure based on PEP scores in an aid formula. This approach rewards districts for poor performance and gives them an incentive to perform poorly in the future. Aid formulas should be based on factors outside a district's control, such as concentrated poverty, that make it difficult for the district to reach a high performance standard, but not on performance indicators that are influenced by the district's actions. New York also has a relatively new program, called Extraordinary Needs Aid, which gives more aid to districts with lower incomes and higher poverty concentrations. The program provides less than 5 percent of the total aid budget, however, and the formula is not based on any estimate of the relationship between educational costs and poverty.

D. Comprehensive Cost Indexes with Indirect Controls for Performance

To move beyond input prices, a cost index must consider the impact of environmental factors on a school districts' costs after accounting for teachers' salaries and for the district's performance. This step requires a statistical procedure base on data describing a district's spending and teacher's salaries, along with measures of relevant environmental factors, such as the district's poverty rate. In addition, it requires information that makes it possible to control for performance. One approach is to control for performance indirectly by controlling for factors that are known to influence performance and that are outside the control of school district officials. This method has been employed by Bradbury et al. (1984) and by Ratcliffe, Riddle, and Yinger (1990), and is illustrated by Downes and Pogue (1994) and Duncombe, Ruggiero, and Yinger (1996).

This approach draws on a large literature concerning the demand for public service outcomes, such as school district performance (see Inman, 1979; Ladd and Yinger, 1991). In particular, dozens of studies have shown that public outcomes depend, in part, on income, tax price, and measures of voter preferences, such as the percent elderly. In this context, the tax price is the cost to a voter of raising performance by one unit. It is analogous to a private price; if the tax price goes up, voters select a lower public outcome. By controlling for a school district's income, tax price, and voter characteristics, therefore, one is indirectly controlling for school district performance.

This method builds on a regression analysis in which the level of public spending is a function of income, tax price, and voter characteristics (to control for performance) and of teachers' salaries and environmental factors (to account for costs). The regression results can be used to determine how much each district would have to spend if its income, tax price, and voter characteristics (and hence its performance) were average, given its teachers' salaries and environmental factors. A cost index is simply this spending amount relative to the spending amount in the average district. In other words, the cost index indicates how much each district would have to spend, relative to the average district, to achieve an average educational performance. Because it is derived from a statistical procedure, this index is not influenced by the politics or by guesses as to the role of various environmental factors.

Three features of this approach should be emphasized. First, like the salary-based indexes discussed earlier, it can be implemented without any decision about the appropriate way to measure school district performance. However, it implicitly defines school district performance using the set of school performance measures that are highly correlated with income, tax price, and voter characteristics. Because this correlation is never observed directly, this definition is hidden from view. One might say that this is an advantage, because a potentially divisive debate about the appropriate performance measures can be avoided, but it is also a disadvantage because the performance that are selected implicitly might not, if identified, be reasonable. Moreover, even though it may be possible to avoid selecting performance standards in constructing a cost index, such a choice is, as shown below, unavoidable in the design of a performance-based aid program; the implicit standards in a cost index base on this indirect method cannot match the explicit standards in an aid program.

Second, this approach has the disadvantage is that it is relatively abstract; policy makers may not find it easy to understand why controlling for income and tax price is equivalent to controlling for educational performance. Moreover, many studies have shown that cost factors are part of the tax price, that is, they influence performance themselves. Thus, variables, such as poverty, that increase educational costs, also raise the tax price and lead to a decrease in educational performance.

The dependent variable, namely spending, picks up these two offsetting effects. Strictly speaking, therefore, one cannot determine the impact of teachers' salaries and environmental factors on costs without first accounting for their impact on tax price and hence on the demand for public services. This problem can be solved (see, for example, Ladd and Yinger 1991), but the solution is complex and difficult to explain to public officials. Fortunately, however, the tax-price effect is relatively small, so the studies that apply this approach to education ignore the tax-price effect. This implies that these studies understate the variation in costs across school districts, but probably not to a large degree.

Third, this approach does not avoid the problem that teachers' salaries are influenced by school officials. The best way to handle this problem is to use techniques for estimating what are called simultaneous equations. In this case, unobserved school district characteristics, such as bargaining skill, generosity, or managerial competence, might simultaneously influence both the dependent variable, spending, and one explanatory variable, teachers' salaries. If so, the teachers' salary variable is said to be "endogenous," and its estimated coefficient—and hence cost indexes based on it—will be biased. To eliminate this bias, a researcher must identify at least one variable that influences teachers' salaries but is not affected by unobserved school district characteristics. With such an "instrumental" variable in hand, a researcher can use a well known technique called "two stage least squares" to eliminate the effect of unobserved school district characteristics on teachers salaries—and hence to eliminate the bias. Two such instrumental variables available to us in our New York data set are the average wages of manufacturing production workers and the population of the county in which the school district is located. The first of these variables measures one dimension of the private wage scale against which a school district must compete and the second recognizes that wages tend to be bid up in larger metropolitan areas.

E. Comprehensive Cost Indexes with Direct Controls for Performance

Because many measures of educational performance, such as test scores and drop out rates, are available, an appealing alternative to indirect controls is direct controls. Under this approach, a regression analysis of educational spending uses performance measures as explanatory variables instead of income, tax price, and voter characteristics. This approach has now been used by several scholars, including Downs and Pogue (1994), Duncombe, Ruggiero, and Yinger (1996), and Duncombe and Yinger (1997). This approach makes explicit the selection of performance standards and leads directly to an educational cost index, defined as the spending a district is required to make (relative to the average district) to meet a selected value of the performance variables given its own input prices and environmental factors.

By controlling for performance directly, this approach avoids the complexity of the tax-price effect, which was discussed earlier. No fancy algebra is needed to obtain the correct cost index. However, this approach runs into another problem, namely that unobserved school district characteristics might affect both school spending (the dependent variable) and measures of performance (explanatory variables). This simultaneity problem, like the one associated with teachers' salaries, requires the use of two-stage least squares. Fortunately, instruments for this procedure are readily available, and are, in fact, defined by the indirect method just discussed. The indirect method is based on the well-established result that school performance depends on income and tax price, so income and tax price are natural instruments.

As first pointed out by Duncombe and Yinger (1997), this approach also leads to a school district performance index. This index is a weighted average of the performance measures included in the regression, where the weights are the regression coefficients. Duncombe and Yinger also point out that these weights can be interpreted as demand weights; that is, they reflect the weight households place on each of the performance measures. With this approach, the performance measures used in a final regression are those that prove to be statistically significant; under this demand interpretation, this means that this approach allows a researcher to identify the performance measures that play a significant role in households' demand for educational performance.

In addition, this performance index, with its statistically determined weights, makes it possible to set a single performance standard that covers a wide range of performance indicators. Instead of setting performance standards separately for several different indicators, such as test scores and a drop-out rate, policy makers could set a standard for this index. By setting a standard at the median of the current distribution of this performance index, for example, they would be committing themselves to bringing all districts up to the current median performance across the range of performance measures included in the index.

F. Comprehensive Cost Indexes with Direct Controls for Performance and Efficiency

One might criticize the approaches in the two previous sections because they could confuse high cost and inefficiency: large districts may not have higher costs, for example, but may instead just be inefficient. This problem is analogous to the problem of controlling for household choices, about furnace maintenance for example, in the case of home comfort. In technical terms, ignoring inefficiency could lead to "omitted variable bias" in estimating the effects of environmental factors on costs. In this section we will discuss a method for measuring district inefficiency and including it in an expenditure regression. Inefficiency is, of course, difficult to measure and the approach we will present is not the final word on the subject, but it does provide one tractable method to control for efficiency, and thereby to minimize, if not eliminate, the effect of omitted variable bias on cost indexes.

The approach we have in mind measures efficiency using what is called a "best-practice" technique (see Ruggiero 1996). With this technique, a district is said to be inefficient if it spends more on education than other districts with the same performance and the same educational costs. The degree of inefficiency is measured by the extent of this excess spending. Although the "best-practice" technique we use, called data envelopment analysis or DEA, is well known in some circles, Duncombe, Ruggiero, and Yinger (1996) were the first scholars to use it as a way to control for school district efficiency in an analysis of school district spending. It has the advantage that it requires the same data as the expenditure regression in the previous section and the disadvantage that it involves mathematical programming techniques and is therefore difficult to implement—and to explain. Moreover, it is possible that all districts in a state are inefficient relative to best practices in other states or relative to available management practices that are not used anywhere; however, no empirical method can shed light on inefficiency of this type.

This best-practice technique, like the regression approach in the previous section (and like any performance standard) requires a researcher to select specific performance measures. When the two methods are combined, that is, when a best-practice measure of efficiency is included in an expenditure regression that directly controls for performance, the same performance measures appear in the regression and in the calculations that yield the efficiency measure. Hence the choice of performance measures influences the efficiency measure, just as the inclusion of the efficiency measure in the regression might influences which performance measures turn out to be significant (and hence are included in the final regression). This linkage could be a source of contention when the focus is on a single, arbitrarily selected performance indicator; districts found to be inefficient in this case may simply be concentrating on other types of performance. This problem can be minimized by using statistical criteria, such as the ones described earlier, to select a range of performance measure for inclusion in the cost and efficiency analyses.

Strictly speaking, the DEA approach does not lead to a measure of efficiency alone but instead to a measure of the reasons why one district spends more than other districts to achieve the same performance. As a result, our best-practice variable picks up variation in educational costs as well as in efficiency. In effect, therefore, including both this variable and cost variables in an expenditure regression might, in principle, make it difficult to sort out the role of input and environmental cost factors. With our New York data, at least, this does not turn out to be the case. Even when the best-practice variable is included, the cost variables behave as expected and are statistically significant, and the best-practice variable itself has the expected impact on spending and is statistically significant, too. See Appendix Table A1.

The best-practice variable, and indeed any measure of school district efficiency, could be influenced by unobserved school district characteristics that also influence spending. As a result, this variable, like the performance measures and teachers' salaries, must be treated as "endogenous." To find the required exogenous instruments, one can turn to the large literature on the determinants of governmental efficiency. As discussed in Duncombe and Yinger (1997), for example, scholars have argued that a school district's efficiency is influenced by the state aid it receives and by education and wealth of its voters.

The process of identifying these instruments can be thought of as an analysis of the determinants of school district efficiency. Duncombe and Yinger (1997) also show that these instruments, along with controls for educational costs and for performance indicators that are not considered in deriving the best-practice efficiency index, can explain a large share of the variation in this index. For the most part, these instruments affect the best-practice measure in the direction predicted by the literature on efficiency and are statistically significant. These results provide some evidence that this approach is indeed a reasonable way to account for school district efficiency in an expenditure regression.

F. Relationship among Various Approaches

In this section we present a comparison of educational cost indexes using all of the above methods. Using data from 631 school districts New York state in 1990-91, we show how input prices vary, why it is so important to account for environmental factors, how direct and indirect controls for performance compare, and what happens when one accounts for inefficiency.

The key analytical tool in these calculations is a regression analysis of school spending. In our preferred regressions, the explanatory variables are several performance measures, teachers' salaries, environmental cost factors, and a best-practice control for efficiency. As shown in Appendix Table A1, two performance measures prove to be statistically significant: the average share of students above a standard reference point on a set of elementary math and reading tests and the share of students who receive a relatively demanding "Regents diploma," instead of a regular one. A third performance measure, the graduation rate, is close to significant and is also included. These three measures cover a wide range of school district activities both for students who are doing well and for students who are doing poorly. Input prices and the efficiency control are statistically significant with the expected sign. Moreover, four environmental variables play a statistically significant role, namely district enrollment, the percentage of children in poverty, the percentage of households headed by a single female, and the percentage of students with limited English proficiency. We find as do many previous studies, that the impact of district enrollment on costs is roughly U-shaped, with relatively high costs in both the smallest and largest districts.[5] Finally, because expenditures for students with disabilities are so high in some districts, this regression includes a disability variable, namely the percentage of students with a severe handicap, even though it is not significant at conventional levels.[6]

Table 1 describes cost indexes calculated with each of the above methods. This table shows how the cost indexes vary by region and type of district, by pupil class size, by income class, and by property value class. The regions in New York are downstate, for the New York City region, and upstate, for the rest of the state. The first column indicates the number of districts in any particular district class.

Results for the most comprehensive approach, which we believe to be the most accurate, are presented in the second column. According to this approach, upstate suburbs have the lowest costs, with an average index of 90.8. In other words, the average upstate suburb must spend $0.908 to achieve the same performance that the average district obtains for $1.00. In contrast, New York City, with a cost index of 347.6, must spend almost 3 ½ times as much as the average district to achieve the same performance.[7] Costs are also relatively high in the three large upstate cities (Buffalo, Rochester, and Syracuse), which have an average index of 175.3, and in Yonkers, which has an index of 188.1. The results in this column also reveal that costs are relatively high in the smallest and largest districts and that costs tend to increase with income class and with property value—except at the bottom of the distribution.

These results provide striking testimony to the magnitude of variation in educational costs. Even if state aid offsets all differences in district capacity to raise revenue, it is simply not fair to expect a city with a cost index of 200 or 300 to provide the same performance as a district with a cost index of 75 or 100.

An educational cost index that does not account for efficiency, which is presented in the third column of Table 1, leads to a higher apparent variation in costs because it inappropriately treats inefficiency as a cost factor, at least to the extent that inefficiency is correlated with cost variables that are included in the analysis. The key problem is a positive correlation between the wage rate and inefficiency; with no control for inefficiency, the impact of wages on costs is exaggerated and the cost index for high-wage districts is too high while the cost index for low-wage districts is too low. Compared to our preferred index, therefore, this approach drops the average cost index to 82.7 in the low-wage upstate suburbs, and boosts the cost index in downstate districts where wages are relatively high. In particular, the index goes up to 362.2 in New York City and jumps by over 20 percent in Yonkers and in the downstate small cities and suburbs. Excluding a control for efficiency also magnifies the U-shaped relationship between costs and enrollment and the positive relationships between costs and both income and property value. Because this cost index exaggerates the variation in costs across districts, it would, if used in an aid formula, overcompensate districts with high wages or enrollments at either end of the distribution.

The fourth column of Table 1 controls for efficiency, but does not treat efficiency as endogenous. With a few exceptions, the results with this approach are similar to the results in column 2. Perhaps the biggest exception is that this approach boosts the cost index for New York City by 20 percent. For most districts, therefore, this approach provides a reasonable approximation to the method with endogenous efficiency, but it may provide misleading results for a few districts.

Results based on the indirect method of controlling for service quality are presented in the fifth column of Table 1. These results exhibit considerably less variation across type of district, and in fact lead to the rather startling result that costs in New York City are only slightly higher than average. However, this approach, like the previous ones, does indicates relatively high costs in other large cities. Compared to earlier approaches, this approach also leads to less variation in costs with district size, except in the case of the smallest districts, and to less variation in costs with respect to district income and property value.

This approach does not provide a very compelling alternative to the approaches with direct controls for district performance. Not only does the indirect approach require more stringent assumptions than the direct approaches, but when our control for efficiency is introduced into the regression, none of the regression coefficients is statistically significant. In applying this approach, therefore, we must choose between results that are biased because they exclude a control for efficiency (the ones presented in column five of Table 1) or results based on coefficients that are not statistically significant.

A cost index based on the New York State's official weighted pupil measure is presented in the sixth column of Table 1. This index exhibits very little variation across districts; indeed, only one category of district, namely under 100 pupils, has costs more than 3 percent away from the state average. Moreover, the variation that does exist appears to miss the strong tendencies observed in the more comprehensive indexes, namely the high costs in large cities, the U-shaped relationship between costs and enrollment, and the increase in costs with district income and property value. In short, the ad hoc procedures used to determine weighted pupils in New York State bear little systematic relationship to costs, as estimated with either the direct or indirect methods.

Finally, a cost index based on teacher salaries is presented in the last column of Table 1. This index adjusts (as do the regression-based indexes) for the fact that higher salaries must be paid to attract teachers of given quality to harsher educational environments. This approach, like the direct and indirect methods, pick up the relatively high costs in New York City and in other large cities, although the magnitude of the difference between city and other districts is smaller. This approach also indicates that costs increase with district income and property value, although the differences along these two dimensions are very small. However, this approach fails to pick up the relatively high costs of small districts and finds only modestly higher costs in the large districts than in districts of average size.

Another way to compare these cost indexes is to look at the extent to which they are correlated across districts. Our correlation results are presented in Table 2. Not surprisingly, the highest correlation, 0.98, is between the direct cost index with endogenous efficiency and the direct cost index with exogenous efficiency. Leaving out the efficiency controls altogether drops the correlation to 0.91. The correlations are considerably smaller between our preferred index and both the indirect cost index (0.71) and the cost index based on weighted pupils (0.12). This, of course, reinforces the message from Table 1. The weighted-pupil approach, in particular, appears to have little to do with costs based on a comprehensive analysis. Finally, the correlation between our preferred index and the teacher salary cost index is considerably higher, 0.85. Although the teacher salary cost index exhibits relatively little variation compared to the direct index with endogenous efficiency, it appears that districts with relatively high costs according to one index also tend to have relatively high costs according to the other.

4. Cost Indexes and State Foundation Formulas

Educational cost indexes are important largely because they make it possible to design fairer and more effective educational policies. This section explores the link between educational costs and one important state policy, namely a foundation aid program. We show how to bring educational cost indexes into a foundation aid formula—and what happens when cost indexes are ignored. The issues discussed here also arise in programs designed to reward districts that meet performance standards or to punish districts that fall short. As several states have discovered, rewards or punishments that focus exclusively on performance, with no adjustment for costs, end up helping the districts that need help the least and punishing the districts that are, through no fault of their own, stuck with the harshest educational environments. See Clotfelter and Ladd (1996).

A. How to Include Cost Indexes in a Foundation Formula

About 80 percent of states use some form of a foundation grant system, which is designed to ensure that all districts meet some minimal performance standard. For the most part, however, these systems use spending as a measure of "performance," and therefore do not bring many districts up to any given performance standard defined on the basis of test scores or other reasonable performance measures. This need not be the case: cost indexes make it possible to design a foundation formula that brings all districts up to a performance standard defined in this way. We will illustrate our analysis with results from New York state.

A foundation plan is designed to bring all districts up to a minimum spending level per pupil.[8] Let V_i stand for the property tax base in district *i*, then an **expenditure-based** foundation grant per pupil is defined by

$$A_i = E^* - t^* V_i = E^* \left(1 - v_i\right) \quad , \qquad \textbf{(1)}$$

where E^* is the expenditure standard, t^* is the minimum tax rate set by the state, $V^* = E^*/t^*$ is the property value above which a district receives no aid, and v_i is a property value index. A foundation aid program is designed to provide every district with enough resources to provide the foundation level of spending per pupil at the minimum tax rate specified by policy makers. Districts that are wealthy enough to raise the required revenue by themselves simply by setting this specified tax rate receive no aid from the state.

If taken literally, (1) implies that districts with tax bases above V^* actually receive negative aid. This formula is usually modified in practice, through minimum aid amounts or hold-harmless clauses, so that all districts receive some aid, thereby reducing the equalizing power of the formula. Moreover, a foundation grant usually is accompanied by a requirement that each district levy a tax rate of at least t^*; otherwise, some districts might not provide the minimum acceptable spending level, E^*. New York and Illinois are notable exceptions; see Miner (1991) and Downes and McGuire (1994).

Because they do not systematically account for cost differences across districts, these plans do not bring all districts up to a minimum performance level. In particular, districts with relatively high costs cannot reach the standard unless they set a tax rate that is above the required minimum.

To make the switch from spending to performance, one must incorporate educational costs into the aid formula. This step can be done using one of two equivalent methods. The first method begins by selecting a performance standard, say S^*, based on a performance index derived from a cost equation, such as the one described earlier. As noted earlier, policy makers must determine which index value that is the appropriate standard for their state. They could, for example, set the standard at the 25th percentile of the current performance distribution, as measured by the index. The cost equation then indicates how much each district would have to spend, say C_i, per unit of S, given its input prices and environmental factors. Because they come from the same equation, the performance index and the cost index are consistent, that is, they involve the same performance measures. Note also that this cost measure is not the same as a cost index because it is not divided by the costs in the average district, say $\bar{C}$. The second method focuses on the amount the average district must spend to achieve the standard, E^*, which is equal to S^* multiplied by $\bar{C}$. The amount any individual district must spend to achieve the standard is then simply E^* multiplied by c_i, which is the cost index for district i. These two approaches are equivalent because, by definition, S^*C_i equals E^*c_i.

Based on these concepts, we can now see that a **performance-based** formula that brings all districts up to the selected performance standard, S^*, at an acceptable tax burden on their residents is as follows:

$$A_i = S^* C_i - t^* V_i = E^* \left(c_i - v_i \right) \quad , \qquad \textbf{(2)}$$

where the amount each district must spend to reach the standard replaces the fixed amount of spending in equation (1). See Ladd and Yinger (1994). The amount of aid this district receives equals the spending level required to reach S^* minus the amount of revenue it can raise at the specified tax rate t^*. As with equation (1), raising S^* to an extremely high level would, at great cost, result in an equal educational output in every district, and allowing negative grants would boost the equalizing impact of the grant.

In short, one cannot switch to a performance-based aid program without replacing the notion of the foundation spending level, which does not recognize cost variation across districts, with an estimate of the spending that is required to achieve the performance standard in a district with average costs. In principle, this new spending standard, E^*, could be obtained using the method described in Guthrie and Rothstein (1997), which aggregates the cost of typical input bundles, so long as the calculation is based on the inputs employed by a district with average input and environmental costs.[9] Returning to the home comfort example, our approach is analogous to a regression based approach, whereas this input-aggregation approach is analogous to a engineering study. Our approach provides an alternative to the approach in these other studies because the cost equation reveals how much a district with average costs would have to spend to meet any performance standard, so long as this standard is set at some percentile of the performance index. In other words, this approach employs a statistically based calculation of the required spending in the district with average costs instead of pricing the set of inputs that experts believe are needed for a district to meet an adequacy standard. Reasonable people may disagree about which approach is to be preferred. Whichever approach is taken, however, the use of a cost index in the aid formula is crucial.

Because some districts are inefficient, a program based on equation (2) will not bring all districts up to the foundation level (and implicit performance standard) even with a required minimum tax rate. In fact, virtually all districts fall short of the best-practice efficiency level. As a result, it seems reasonable to design a foundation formula so that every district will have enough revenue to achieve the foundation performance level at some efficiency level, say the 75^{th} percentile of the current efficiency distribution across districts, which we call the baseline efficiency level. If it falls short of this level, it will not achieve the foundation level of performance unless its tax rate is above the specified rate.

In addition, switching to a performance-based aid system undoubtedly would not immediately bring high cost districts up to the performance standard. Even if the resources needed to meet the standard were available to them, these districts would have to alter some existing practices, design new programs, and hire at least some new teachers and administrators. These steps would take time. A reasonable approach, in our view, would be for a state to move to a performance-based aid system over a period of several years, providing along the way both management assistance and new research evidence about the effects of various educational programs. In combination with these extra steps, moving to a performance-based aid system would allow a state to say that it had provided each district with the resources it needs to meet the performance standard, as determined by the best available information.

B. Simulating the Effect of a Performance-Based Aid Program

Using data from New York school districts, excluding New York City (which otherwise would dominate the aid program),[10] Duncombe and Yinger (1997) simulate the effect of different aid systems on student performance in each district. These simulations employ not only the performance and cost indexes discussed here, but also analyses of voter demand for educational performance and of the determinants of school district efficiency, which are not presented here. The key results are reproduced in Table 3.[11] This table shows the average performance, as measured by the performance index described earlier, in each class of district under foundation aid programs based on various ways of accounting for educational costs. In this table, performance is expressed relative to the actual 1991 performance level in the average district, not, as in equation (2), in terms of required spending. Thus, an index of 39.7, the current value for upstate large cities, indicates a performance that is almost 60 percent below the state average. All of the aid programs have the same overall state budget and impose a minimum-tax-rate requirement, and each type of aid system (a column) is examined for three different performance standards or foundation levels (the panels), namely the 25th, 50th, and 75th percentiles of the 1991 performance distribution in New York, as measured by the performance index. As indicated in the table, these percentiles fall at 76.2 percent, 97.4 percent, and 120.4 percent, respectively, of the current performance in the average district.

The first column of Table 3 indicates district performance under the existing aid system, which includes several small lump-sum programs plus a foundation plan with the minimum expenditure (not performance) level set at approximately the 25th percentile of the current expenditure distribution and with various hold-harmless and minimum-aid provisions. The second column of Table 6 presents results for an performance-based foundation plan with a required minimum tax rate and an efficiency baseline. The performance increases above their existing levels are most dramatic for large, upstate central cities. The average performance for these three cities with the most generous plan, 104.7, is slightly over 2 ½ times as large as their current performance, 39.7. This new performance level falls short of the target *S**, 120.4, because the increased aid drives the efficiency level in these districts below the baseline level. This performance-based foundation plan also boosts performance in all other classes of district, although not by such dramatic amounts.

Implementing this aid plan requires an understanding of cost indexes, an explicit decision about the acceptable level of inefficiency, and the estimation of cost indexes controlling for efficiency. Existing state plans do not take any of these steps, so Duncombe and Yinger (1997) also simulate three alternative foundation plans based on less complete information.

The simplest foundation plan follows equation (1), with no recognition of costs or efficiency. The results for such a plan are presented in the last (fifth) column of Table 3. Because the implicit expenditure target in the current New York foundation plan is set at about the 25th percentile of the current expenditure distribution, a comparison of the first and last columns in the first panel of these tables largely reflects the impact of eliminating hold-harmless and minimum-aid provisions and pooling all lump-sum aid into a foundation formula. These steps would modestly increase aid (and performance) in upstate cities, both large and small, and decrease aid substantially (with little impact on performance) in downstate cities and suburbs. The average impact on rural districts and downstate suburbs would be minimal.

Bringing in the results in the second column, we can see that an performance-based foundation goes much farther than an expenditure-based foundation in shifting aid toward large cities and thereby boosting their performance. It does not go nearly as far, however, in shifting aid away from downstate small cities and suburbs, a result shown by their high performance. Largely because they face very high labor costs, these downstate districts tend to have high-costs, a fact that is missed by an expenditure-based plan. The current system of hold-harmless and minimum-aid provisions serves some of the same purpose as a cost correction by boosting aid to these districts, but it goes too far in this direction and does not ensure fair treatment either within these districts or between these districts and others.

Table 3 also reveals that even the most generous expenditure-based foundation plan leaves large cities far short of any performance target, even with a required minimum tax rate. In fact, the most generous such plan, in the last panel of Table 3, helps large cities but still leaves them at an performance level **well below the 25th percentile** of the current distribution!

Related simulations in Duncombe and Yinger (forthcoming b) make the key point here in a different way. Consider, the notion of a "performance gap," defined as the sum across districts of the amount by which actual district performance falls below the performance standard, weighted by the number of students in the district. Duncombe and Yinger show that with the foundation level (and implicit performance standard) set at the 25^{th} percentile of the current performance distribution and a required minimum tax rate, an expenditure-based foundation plan would close only 36 percent of the current performance gap in New York. In contrast, a comparable, and equal-cost performance-based foundation plan would close 84 percent of performance gap (and would close 100 percent of the gap if all districts met the baseline efficiency standard). The point should be clear: Expenditure-based foundation plans, which are used in most states, leave many high-cost districts short of even a minimal performance standard.

A state cannot implement a performance-based aid program without estimating a cost index. An aid program for municipal services, including education, based on an estimated cost index was implemented in Massachusetts (Bradbury et al. 1984), and school aid programs based on estimated cost indexes are presented in Ratcliffe, Riddle, and Yinger (1990), Downes and Pogue (1994).

However, the cost indexes estimated in these cases do not control for efficiency. Thus we now examine performance-based foundation programs that incorporate a cost index estimated without controlling for efficiency and that implicitly assume, following equation (2), that all districts are efficient. A cost index estimated in this way is biased, because the omission of an efficiency variable biases the coefficients of the included cost variables, but it takes a large step toward recognizing the role of input and environmental cost factors.

Results for these programs, presented in the third column of Table 3, reveal that in most cases adding a cost index closes a large share of the gap between the expenditure-based foundation in the fifth column and the complete performance-based foundation in the second column. Under the most generous plan, for example, adding a biased cost index raises performance in upstate large cities from 40.9 (column 5) to 68.7 (column 3), compared to the complete-information performance (column 2) of 68.7.

In contrast, the foundation plan based on a biased cost index leads to higher aid and higher performance for downstate small cities and suburbs than either the expenditure-based foundation or the complete-information foundation in the second column. As explained earlier, this result mainly reflects the large, negative correlation between efficiency and wage rates; because of this correlation, leaving efficiency out of the cost equation biases upward the coefficient of the wage variable and hence biases upward the cost index in places, like downstate districts, with high labor costs.[12] In effect, therefore, an aid program based on a biased cost index rewards the downstate districts for their inefficiency. This is, of course, an inappropriate outcome.

This result poses a serious challenge to policy makers and researchers. Aid formulas based on simple cost indexes of the type that have been presented in the literature appear to be a big step in the right direction, but this step has a price. To the extent that efficiency is correlated with cost factors, a standard cost index will reflect inefficiency as well as costs, and an aid formula based on it will favor inefficient districts as well as high-cost ones. In New York, this effect does not boost aid to big cities, which despite their reputation are relatively efficient, but instead boosts aid to downstate small cities and suburbs, which tend to be inefficient by our measure.

Obviously the relevant correlations could vary from state to state, so these results cannot determine whether this type of plan would reward the same types of district for inefficiency in other states. Nevertheless, the possibility that the plan rewards inefficiency clearly undercuts its appeal.

As noted earlier, one simple step a state can take to recognize the role of efficiency is to bring in the concept of baseline efficiency. All this step requires is identifying an efficiency level that is regarded as acceptable. This approach recognizes that virtually no districts will be able to achieve perfect efficiency so that spending greater than S^*C_i is needed to bring district i up to the S^* performance target. Compared to the previous approach, therefore, this approach focuses more aid on higher-need districts. The fourth column of Table 3 shows the impact of a foundation plan with a biased cost index but with a baseline efficiency level set at the 75th percentile of the current efficiency distribution. This plan takes another small step toward the complete-information plan in the second column. In the downstate districts, the entry in column 4 falls between the entry in column 3, which has no correction for efficiency, and the entry in column 2, which is based on an unbiased cost index. Moreover, performance in large cities would actually be slightly higher with this plan than with our preferred plan because the biased cost index exaggerates the cost impact of high wages in these districts.

The simulations in Table 3 all involve the same state budget, namely the actual New York State educational aid budget in 1990-91. A natural question to ask is: What would happen if the budget increased. In the case of foundation plans without a minimum tax rate, the effect of a higher state budget can be dramatic. With a foundation plan of this type, many districts set tax rates below the level needed to achieve the foundation level of spending (and, in the case of performance-based foundation plans, of performance), that is, they use some of their state educational aid to fund non-educational programs or to cut taxes. Some portion of any additional state aid will be devoted to education and will therefore boost districts educational spending and performance. In the case of performance-based foundation plans with a minimum tax rate, however, additional state aid has no such effect. Somewhat ironically, in fact, additional state aid leads to a small decrease in performance (see Duncombe and Yinger forthcoming a). The minimum tax rate requirement ensures that all districts raise enough revenue to fund the performance standard if (they meet the baseline efficiency standard). Additional state revenue is therefore not needed to meet the standard. However, a district's efficiency is influenced by the amount of aid it receives, and higher aid generally leads to lower efficiency. As a result, the increase in the state budget shifts the burden of financing education from local governments to the state, with no change in the amount of revenue available for education, and in the process makes school districts a little less efficient. This drop in efficiency results in a small drop in performance.

5. Conclusions

An extensive literature establishes that both school district and student performance depend not only on factors that school officials control, such as the student/teacher ratio, but also on factors that are outside their control, including input prices, such as regional wage rates, and environmental factors, such as concentrated poverty. It follows directly that the cost of education is not the same in every district, with higher costs in districts in higher-wage labor markets or with a harsher educational environment. A shift to educational performance standards, whether these standards are simply targets or are imbedded in a foundation aid program, can be neither fair nor effective unless it recognizes this variation in the cost of education. This shift cannot be fair because it penalizes districts that, through no fault of their own, face harsh educational environments, and it cannot be effective because it hands out rewards and punishments that are not related to the contributions of school personnel.

Scholars have identified a variety of methods for measuring the cost of education, all of which have limitations. The simplest reasonable methods, which are indexes of teachers salaries predicted on the basis of conditions in the local labor market and in a district's schools, fail to recognize that districts with the a relatively harsh educational environment must hire more teachers (or purchase more of other inputs) than other districts to achieve the same performance. The most comprehensive methods, which recognize the role of environmental factors and control for school district efficiency, involve some complex, hard-to-explain steps. Nevertheless, the literature demonstrates that cost variation across schools is very large and cannot be ignored. Policy makers and scholars need to continue the search for sensible, practical ways to measure educational costs and incorporate them into performance-based educational policies.

Endnotes

1. Strictly speaking, this interpretation of the weights depends on the assumption that educational performance is provided at constant cost. See Duncombe and Yinger (1997). This assumption is employed in virtually all the educational finance literature, although it has not been adequately tested. The performance indicators we selected had an adjusted R-squared of at least 0.10 with variables typically found in an education demand equation, including income and tax share. In addition, we ran a factor analysis on an array of performance indicators measures and the scree plot indicated three distinct performance dimensions, which are very similar to the three measures we actually use.

2. In principle, the methods discussed here could be applied at the school level and perhaps even at the classroom level. However, the data for such applications have not yet become available. Cost studies at the school level might prove to be helpful for understanding performance disparities within districts, which is an important issue in many large cities.

3. An analog to the first problem with using teacher's salaries could arise in this example if households could select different grades of natural gas. There is no analog here to the second problem with using teachers' salaries, however, since households do not negotiate over natural gas prices.

4. The 1996 New York State aid formulas are described in The University of the State of New York (1996).

5. The enrollment in a district is influenced, of course, by district consolidation or district splitting. The U-shaped relationship we estimate implies that it may be possible to lower educational costs by consolidating small districts and splitting up large ones. We do not consider these possibilities here, but consolidation is examined in Duncombe, Miner, and Ruggiero (1995). In the long run, enrollment also in a school district may be influenced by the decisions of school officials, as parents make residential choices and choices about private school based on school quality and property tax rates. We know of no study that considers this possibility, and it is an important topic for future research.

6. We do not attempt a full analysis of spending on students with disabilities and in fact our dependent variable, approved operating expenditure as measured by the State, does not include all such spending. For an analysis of the cost of special education, see Chaikind, Danielson, and Brauen (1993). Broader measures of students with disabilities also complicate the analysis because they may be influenced by the way a school district determines which students have disabilities. See Lankford and Wyckoff (1996). We restrict our analysis to environmental variables, such as the poverty rate and the share of students with severe disabilities, that appear to be, for all practical purposes, outside the control of school officials.

7. These results are not driven by New York City. A regression analysis that excludes New York City results in cost indexes, both for the City and for other districts, that are similar to those in Table 1. See Appendix Table A1.

8. Historically, New York has used a modified foundation formula, but the current formula mixes elements of a foundation formula and a power-equalizing formula. In effect, the current formula appears to act as a power-equalizing formula for districts with spending levels in the middle of the spending distribution. The state aid formulas are described in The State University of New York (1996).

9. The key problem with this approach is that it does not provide researchers with a method for identifying this average-cost district. This problem is not recognized by Guthrie and Rothstein (1997).

10. As noted earlier, New York City has the highest costs in the state, but it also now receives below average aid per pupil. New York City also has a very high share of the pupils in the state, so if it is included in the simulations, it receives a very large share of any performance-based aid program and little is left over for other districts. Results for other states are unlikely to be so dominated by one district, so we focus on simulations without New York City. The cost index used in these simulations is the first on in Appendix Table A1.

11. Table 3 is a revised version of Table 6 in Duncombe and Yinger (1997). The revisions were necessary to correct a minor programming error. This error did not affect any other tables, nor did it affect any of the substantive conclusions in the original Duncombe and Yinger paper.

12. All cost coefficients are biased when efficiency is left out of a cost model, but in our equation the bias in the wage variable is particularly dramatic.

References

Bradbury, Katharine L., Helen F. Ladd, Mark Perrault, Andrew Reschovsky, and John Yinger. 1984. "State Aid to Offset Fiscal Disparities across Communities," *National Tax Journal* 37 (June): 151-170.

Bradford, David, Robert Malt and Wallace Oates. 1969. "The Rising Cost of Local Public Services: Some Evidence and Reflections," *National Tax Journal* 22 (June): 185-202.

Chaikind, Stephen, Lousi C. Danielson, and Marsha L. Brauen. 1993. "What Do We Know About the Costs of Special Education? A Selected Review," *The Journal of Special Education* 26, 4: 344-370.

Chambers, Jay. 1995. "Public School Teacher Cost Differences Across the United States: Introduction to a Teacher Cost Index (TCI)." In *Developments in School Finance, 1995*. Washington, DC: National Center for Education Statistics, pp. 21-32.

Chambers, Jay. 1978. "Educational Cost Differentials and the Allocation of State Aid for Elementary and Secondary Education," *Journal of Human Resources*, 13: 459-481.

Clotfelter, Charles, and Helen F. Ladd. 1996. "Recognizing and Rewarding Success in Public Schools." In H. F. Ladd (ed.), *Holding Schools Accountable: Performance-Based Reform in Education.* Washington, DC: The Brookings Institution, pp. 23-64.

Coleman, James, 1966. *Equality of Educational Opportunity*. Washington, DC: U.S. Government Printing Office.

Downes, Thomas and Therese McGuire. 1994. "Alternative Solutions to Illinois' School Finance Dilemma: A Policy Brief," *State Tax Notes*, February 14: 415-419.

Downes, Thomas and Thomas Pogue. 1994. "Adjusting School Aid Formulas for the Higher Cost of Educating Disadvantaged Students," *National Tax Journal*, 47: 89-110.

Duncombe, William, Jerry Miner, and John Ruggiero. 1995. "Potential Cost Savings from School District Consolidation: A Case Study of New York," *Economics of Education Review* (September): 356-384.

Duncombe, William, John Ruggiero, and John Yinger. 1996. "Alternative Approaches to Measuring the Cost of Education." In H. F. Ladd (ed.), *Holding Schools Accountable: Performance-Based Reform in Education*. Washington, DC: The Brookings Institution, pp. 327-356.

Duncombe, William, and John Yinger. Forthcoming a. "An Analysis of Two Educational Policy Changes in New York: Performance Standards and Property Tax Relief." To appear in a conference volume published by the New York State Board of Regents.

Duncombe, William, and John Yinger. Forthcoming b. "School Finance Reform: Aid Formulas and Equity Objectives," *National Tax Journal*.

Duncombe, William, and John Yinger. 1997. "Why Is It So Hard to Help Central City Schools?" *Journal of Policy Analysis and Management*, 16 (Winter): 85-113.

Ferguson, Ronald. 1991. "Paying for Public Education: New Evidence on How and Why Money Matters," *Harvard Journal on Legislation*, 28: 465-498.

Ferguson, Ronald, and Helen F. Ladd. 1996. "Additional Evidence on How and Why Money Matters: A Production Function Analysis of Alabama Schools." In H.F. Ladd (ed.), *Holding Schools Accountable: Performance-Based Reform in Education*. Washington, DC: The Brookings Institution.

Guthrie, James W., and Richard Rothstein. 1998. "Enabling 'Adequacy' to Achieve Reality: Translating Adequacy into State School Finance Distribution Arrangements." Paper presented at a conference sponsored by the Committee on Education Finance, National Research Council, Irvine California, January 30-31.

Hanushek, Eric. 1986. "The Economics of Schooling: Production and Efficiency in Public Schools," *Journal of Economic Literature*, 24: 1141-1177.

Inman, Robert. 1979. "The Fiscal Performance of Local Governments: An Interpretative Review." In P. Mieszkowski and M. Straszheim (eds.), *Current Issues in Urban Economics*. Baltimore, MD: The Johns Hopkins University Press.

Kreuger, Alan B. 1997. "Empirical Estimates of Education Production Functions." Industrial Relations Section Working Paper #379, Princeton University, Princeton, N.J., May.

Ladd, Helen F. and John Yinger. 1991. *America's Ailing Cities: Fiscal Health and the Design of Urban Policy*. Baltimore, MD: The Johns Hopkins University Press.

Ladd, Helen and John Yinger. 1994. "The Case for Equalizing Aid," *National Tax Journal*, 47 (March): 211-224.

Lankford, Hamilton and James Wyckoff. 1996. "The Allocation of Resources to Special Education and Regular Instruction." In H.F, Ladd (ed.), *Holding Schools Accountable*. Washington, DC: The Brookings Institution, pp. 221-257.

Miner, Jerry. 1991. "A Decade of New York State Aid to Local Schools," Metropolitan Studies Program Occasional Paper No. 141, Center for Policy Research, The Maxwell School. Syracuse, NY: Syracuse University.

Minori, Paul, and Stephen D. Sugarman. 1998. "Educational Adequacy and the Courts: The Promise and Problems of Moving to a New Paradigm." Paper presented at a conference sponsored by the Committee on Education Finance, National Research Council, Irvine California, January 30-31.

Ratcliffe, Kerri, Bruce Riddle, and John Yinger. 1990. “The Fiscal Condition of School Districts in Nebraska: Is Small Beautiful?” *Economics of Education Review*, 9: 81-99.

Rouse, Cecilia Elena. 1997. “Private School Vouchers and Student Achievement: An Evaluation of the Milwaukee Parental Choice Program.” Unpublished Manuscript, Princeton University, May.

Ruggiero, John. 1996. “On The Measurement of Technical Efficiency in the Public Sector,” *European Journal of Operational Research*, 90: 553-565.

Steinberg, Jacques. 1998. “Underachieving Schools Are Shamed Into Improvement,” *The New York Times*, January 7, 1998.

University of the State of New York and The State Education Department. 1996. *State Aid to Schools: A Primer*. Albany.

University of the State of New York and The State Education Department. 1997. *New York: The State of Learning, Statewide Profile of the Educational System*. Albany.

Wendling, Wayne. 1981. “The Cost of Education Index: Measurement of Price Differences of Education Personnel among New York State School Districts,” *Journal of Education Finance*, 6 (Spring): 485-504.

Witte, John F. 1996. “School Choice and Student Performance.” In H. F. Ladd (ed.), *Holding Schools Accountable: Performance-Based Reform in Education*. Washington, DC: The Brookings Institution, pp. 149-176.

Table A-1
Education Cost Models (with and without New York City and Yonkers),
New York School Districts, 1991[a]

Variables	Estimates without New York City and Yonkers		Estimates with New York City and Yonkers	
	Coefficient	t-statistic	Coefficient	t-statistic
Cost Equation[a]				
Intercept	-4.9550	-1.53	-4.3542	-1.44
Third- and Sixth-Grade PEP Scores (Average % above Standard Reference Point)[b]	5.1106	2.50	5.6171	2.86
Percent non-dropouts[b]	4.4757	1.62	4.2826	1.55
Percent receiving Regents diploma[b]	1.3449	3.19	0.9231	2.09
Efficiency index (percent)[b]	-1.1670	-4.87	-1.1429	-4.81
Log of teacher salaries[b]	0.6487	1.57	0.7211	1.72
Log of enrollment	-0.5680	-3.54	-1.1823	-2.67
Square of log of enrollment	0.0345	3.44	0.1171	2.32
Cubic of log of enrollment			0.0035	-1.85
Percent of children in poverty	1.0109	3.93	0.9526	3.73
Percent female-headed households	2.2260	3.85	2.0751	3.61
Percent of students with severe handicap	0.8584	1.29	0.7137	1.10
Percent of students with limited English proficiency	4.0525	2.68	4.2940	2.82
SSE	34.58		33.45	
Adj. R-square	0.32		0.34	

[a] The cost models are estimated with linear 2SLS regression. The dependent variable is the logarithm of per pupil operating expenditures. The sample size for both models is 631 districts.

[b]These variables are treated as endogenous. See the text for a discussion of the instruments.

[c]These regression results are those presented in Duncombe and Yinger (1997). Besides not including NYC and Yonkers, there are some differences in the data and measures between this regression and the regression reported in columns 3 and 4.

TABLE 1

Comparison of Education Cost Indices[a]

for New York State School Districts in 1991

Socio-Economic Characteristics	Number Of Districts	Direct Cost Indices: Endogenous Efficiency	Direct Cost Indices: No Efficiency Index	Direct Cost Indices: Exogenous Efficiency	Indirect Cost Index (no efficiency index)	Cost Index Based on Weighted Pupils	Teacher Salary Cost Index
Region Type:							
Downstate Small Cities	7	132.16	171.21	120.65	111.53	102.65	108.90
Downstate Suburbs	130	108.20	129.10	103.83	100.90	101.61	103.57
New York City	1	386.89	396.49	226.87	111.83	98.11	125.36
Yonkers	1	191.91	276.81	156.42	130.14	98.69	113.78
Upstate Large Cities	3	189.64	188.95	171.25	135.72	100.31	110.92
Upstate Rural	212	98.92	93.54	99.43	103.83	99.90	98.68
Upstate Small Cities	47	109.05	103.74	111.63	102.44	100.52	102.64
Upstate Suburbs	231	90.75	83.42	93.66	94.49	99.02	98.08
Pupil Size Class:							
Under 100 pupils	1	172.87	304.45	133.58	199.21	120.35	96.94
100-500 pupils	61	107.26	115.13	103.93	116.47	101.16	97.74
500-1,000 pupils	113	99.49	98.16	99.55	103.31	99.43	99.13
1,000-1,500 class	131	93.07	86.80	95.55	96.26	98.55	98.90
1,500-3,000 pupils	181	96.24	94.86	97.55	95.57	100.73	100.30
3,000-5,000 pupils	80	96.64	96.99	98.16	94.81	99.59	100.70
5,000-10,000 pupils	54	111.46	118.58	108.65	100.94	101.11	103.26
10,000-50,000 pupils	10	149.30	168.13	135.69	116.66	100.67	107.21
Over 50,000 pupils	1	386.89	396.49	226.87	111.83	98.11	125.36
Income Class (percentile):							
Under 10th	63	101.51	95.05	101.52	105.53	99.68	99.47
10th to 25th	95	98.14	91.82	99.88	102.03	100.49	98.93
25th to 50th	158	96.63	90.24	98.41	99.46	99.39	98.89
50th to 75th	158	99.29	98.29	100.03	98.68	99.45	99.89
75th to 90th	94	105.23	114.55	101.93	97.89	99.92	101.80
Over 90th	64	103.65	123.98	99.69	99.23	102.57	102.49
Property Values (percentile):							
Under 10th	63	99.56	90.35	101.33	102.09	100.19	99.78
10th to 25th	95	97.06	88.13	99.25	99.74	99.38	99.04
25th to 50th	158	96.50	89.95	98.74	98.20	100.28	99.19
50th to 75th	158	100.55	99.62	99.80	98.63	99.57	100.11
75th to 90th	94	103.69	115.96	101.02	100.08	100.17	101.51
Over 90th	64	106.67	129.44	101.91	106.03	100.85	101.15

[a]All indices are relative to the state average which is set at 100. Sample size for the analysis is 631 school districts. The cost index in the fourth column is a reduced form model where the demand instruments--income, taxshare and households with children--are substituted into cost model for outcome measures. The index in the fifth column is based on a ratio of weighted pupils over total enrollment. Extra weight is given to secondary, handicapped and special needs pupils. The index in the last column is based on the relationship between teacher salaries and family and student characteristics. Income is based on estimated per capita adjusted gross income in 1991and property values are per capita market value for all property in 1991.

TABLE 2

Correlations between Education Cost Indices[a] for New York State School Districts in 1991

Socio-Economic Characteristics	Direct Cost Indices: Endogenous Efficiency	Direct Cost Indices: No Efficiency Index	Direct Cost Indices: Exogenous Efficiency	Indirect Cost Index (no efficiency index)	Cost Index Based on Weighted Pupils	Teacher Salary Cost Index
Standard deviation	22.04	38.19	14.31	11.53	8.88	4.17
Maximum	386.89	396.49	226.87	199.21	264.07	125.36
75th percentile	104.39	107.38	103.96	104.56	102.90	101.71
25th percentile	88.90	78.97	92.38	92.48	97.01	97.23
Minimum	75.52	61.92	80.60	82.55	44.69	92.51
Correlations:						
Direct Cost Indices:						
Endogenous efficiency index	1.00					
No efficiency index	0.92	1.00				
Exogenous efficiency index	0.95	0.85	1.00			
Indirect cost index (no efficiency index)	0.73	0.76	0.76	1.00		
Cost index based on weighted pupils	0.12	0.15	0.14	0.15	1.00	
Teacher salary cost index	0.80	0.78	0.80	0.46	0.16	1.00

[a]All indices are relative to the state average which is set at 100. Sample size for the analysis is 631 school districts. The cost index in the fourth column is a reduced form model where the demand instruments--income, taxshare and households with children--are substituted into cost model for outcome measures. The index in the fifth column is based on a ratio of weighted pupils over total enrollment. extra weight given to secondary, handicapped and special needs pupils. The index in the last column is based on the relationship between teacher salaries and teacher, family and student characteristics.

Table 3

Comparison of Predicted Performance Under Different Foundation Formulas,

Relative to State Average Performance in 1991 for New York School Districts[a]

		Performance-Based Aid System With Direct Cost Indices			Expenditure-
Aid System	**Actual Outcomes**	**Endogenous Efficiency**	**No Efficiency Index**	**Exogenous Efficiency**	**Based Aid System**
Aid: S*=25th percentile=	**76.2**				
Average	100.0	**104.1**	102.1	100.6	103.0
Downstate					
Small cities	83.0	**85.4**	90.1	90.0	78.4
Suburbs	118.5	**123.7**	125.7	126.0	120.0
Upstate					
Large cities	39.7	**68.7**	59.7	71.0	40.9
Rural	89.4	**95.4**	89.2	91.5	93.1
Small cities	82.2	**90.1**	83.0	86.3	86.1
Suburbs	104.1	**104.8**	101.5	102.4	107.4
Aid: S*=50th percentile=	**97.4**				
Average	100.0	**122.7**	117.6	117.7	118.9
Downstate					
Small cities	83.0	**99.2**	104.6	103.1	87.1
Suburbs	118.5	**146.4**	149.4	148.7	135.8
Upstate					
Large cities	39.7	**87.0**	80.9	89.3	47.6
Rural	89.4	**115.8**	106.4	108.1	110.9
Small cities	82.2	**110.8**	101.2	103.3	100.7
Suburbs	104.1	**119.4**	114.1	113.0	122.4
Aid: S*=75th percentile=	**120.4**				
Average	100.0	**155.9**	150.8	149.0	153.3
Downstate					
Small cities	83.0	**138.2**	141.9	139.2	123.3
Suburbs	118.5	**195.6**	198.4	197.4	180.3
Upstate					
Large cities	39.7	**104.7**	100.0	106.5	64.2
Rural	89.4	**145.1**	135.7	135.0	142.3
Small cities	82.2	**136.7**	128.3	127.2	128.7
Suburbs	104.1	**148.6**	143.6	139.9	155.4

[a]All grants require approximately the same state budget to fund as the aid system in 1991, \$3.65 billion. The performance is expressed relative to the state average performance in 1991. A value of 100 equals this average. All aid plans in this table require districts to assess the minimum tax rate, t*, set by the state.