

Research Series No. 22

AN ECONOMETRIC PERSPECTIVE
ON CLASSROOM READING INSTRUCTION

Byron W. Brown and Daniel H. Saks

Published By

The Institute for Research on Teaching
252 Erickson Hall
Michigan State University
East Lansing, Michigan 48824

July 1978

Publication of this work is sponsored by the Institute for Research on Teaching, College of Education, Michigan State University. The Institute for Research on Teaching is funded primarily by the Program for Teaching and Instruction of the National Institute of Education, United States Department of Health, Education, and Welfare. The opinions expressed in this publication do not necessarily reflect the position, policy, or endorsement of the National Institute of Education. (Contract No. 400-76-0073)

Institute for Research on Teaching

The Institute for Research on Teaching was founded at Michigan State University in 1976 by the National Institute of Education. Following a nationwide competition in 1981, the NIE awarded a second contract to the IRT, extending work through 1984. Funding is also received from other agencies and foundations for individual research projects.

The IRT conducts major research projects aimed at improving classroom teaching, including studies of classroom management strategies, student socialization, the diagnosis and remediation of reading difficulties, and teacher education. IRT researchers are also examining the teaching of specific school subjects such as reading, writing, general mathematics, and science, and are seeking to understand how factors outside the classroom affect teacher decision making.

Researchers from such diverse disciplines as educational psychology, anthropology, sociology, and philosophy cooperate in conducting IRT research. They join forces with public school teachers, who work at the IRT as half-time collaborators in research, helping to design and plan studies, collect data, analyze and interpret results, and disseminate findings.

The IRT publishes research reports, occasional papers, conference proceedings, a newsletter for practitioners, and lists and catalogs of IRT publications. For more information, to receive a list or catalog, and/or to be placed on the IRT mailing list to receive the newsletter, please write to the IRT Editor, Institute for Research on Teaching, 252 Erickson Hall, Michigan State University, East Lansing, Michigan 48824-1034.

Co-Directors: Jere E. Brophy and Andrew C. Porter

Associate Directors: Judith E. Lanier and Richard S. Prawat

Editorial Staff

Editor: Janet Eaton

Assistant Editor: Patricia Nischan

Abstract

Elements of economic production theory are applied to the relation between time and student learning. The analysis takes into account that individual students matter, that students are often taught in classes, and that jointness in production is the rule rather than the exception. The values teachers place on the different possible educational outcomes are shown to be indispensable to understanding production. Some basic conclusions about the nature of accountability are derived, such as the idea that in developing measures of accountability, educators must decide just what it is they want to make the teacher accountable for: the right technology (including equipment and other inputs over which s/he may have little control), the amount of effort allocated by the teacher, or the preferences that cause him/her to produce one distribution rather than another with given technology, resources, and effort. Educators should not expect a single measure to suffice in exercising control over so many variables whose effects are hopelessly entangled.

Contents

Production Functions	2
Accountability	13
Production in the Classroom: A Time Allocation Problem	16
Future Research by Reading Specialists	27
References	29

Figures

Figure 1. Learning possibilities for two students	8
Figure 2. Attainable means and standard deviations of scores for two students.	12
Figure 3. Hypothetical relationship between time and learning for an individual student	17
Figure 4. The relationship between time allocated to two techniques of instruction and the resulting output possibilities.	22
Figure 5. Utility contours of one teacher for two different learning outcomes	24
Figure 6. Utility maximizing combination of learning outcomes when the teacher must choose along the line PP'	24

An Econometric Perspective¹
On Classroom Reading Instruction

Byron W. Brown and Daniel H. Saks²

Oddly enough, economists have had relatively little involvement with the current trend of modeling schools as if they were firms employing workers (teachers) and capital (buildings and equipment) to process raw materials (students) in producing education. Indeed, some non-economists looking at the relation between outputs of schools and inputs to schools may be surprised to know that it represents one of the principal concepts of economic analysis, the "production function." This suggests that economists might have some useful things to say about production in schools and that they might even be able to recommend some fruitful lines of inquiry for researchers on reading instruction.

Before any dialogue between reading researchers and economists can begin, however, they must learn to talk to each other. (Indeed, communication of this type has been promoted at the Institute for Research on Teaching.) Whether their talk makes any sense is something that educators will have to decide for themselves in a few years.

¹This paper was presented at the Annual Meeting of the International Reading Association, May 1978, in Houston.

²The authors are professor and associate professor of economics at MSU, respectively. Both are senior researchers with the Institute for Research on Teaching.

In this paper, we will try to present some of the elements of the economists' part of the dialogue. First, we will discuss the production function concept and show how it is related to statistical observations on schools and school districts. Second, we will explain the implications of work in production functions for developing measures of teacher accountability. (To give a preview, we believe that measuring teacher performance is not nearly as simple as some people believe, even when teachers are only trying to change some well-defined test score.) Third, we will explore some of the basic production issues suggested by our studies of what goes on in the classroom. (We are still engaged in such a study so, at this point, we have speculations and questions rather than answers.) We will conclude with a "shopping list" of the kinds of things we wish reading researchers could tell us and which we do not know. We even harbor the hope that some reading researchers might eventually be able to answer some of our questions.

Production Functions

Much of the post-"Coleman Report" (1966) research on schooling was aimed at finding correlates of student achievement. The hope was that if it was known how school outputs changed in response to input changes, it might be possible to increase desirable outputs by obtaining more of the relevant inputs. The search revealed that the only input which seemed to make any difference was the "quality" of the students (Averch, Carroll, Donaldson, Kiesling, & Pincus, 1972). Hence, it seemed that the best way to improve the scores of students in a given school would be to get new (and better) students (usually defined as students from higher socioeconomic status families) into the school.

The proposition that schools have no effect on student learning is rejected by teachers on grounds of self-interest, and by economists because it violates every tenet of common sense about production processes. Yet, the statistical results which point to this conclusion must be taken seriously, if only because they have been so often reproduced and cited. To understand why the conclusion is probably wrong, a deeper understanding of production functions is needed.

The economic production function³ is a device economists use to describe the technical aspects of a production process. Implicit in the production function concept are definitions of inputs and outputs. The usual textbook version of the function describes the determination of a single homogeneous output by a variety of homogeneous inputs. The function is a mathematical summary of the way inputs are converted into outputs; it reveals the maximum quantity of output which can be obtained from any combination of inputs, and in doing so, depicts the nature of the production.

Economists are most interested in knowing (1) how much output changes when a single input changes by one unit, with all other inputs held constant (this is the "marginal" output or contribution of just one factor of production), and (2) how one input may be substituted for another to produce the same level of output. A reasonable assumption is that more input will produce more output, i.e., that the marginal product is positive. However, this assumption has not been borne out by most educational production function studies (Averch et al.).

³"Function" is used here in the mathematical sense, as the description of a relationship between variables.

The single-output, multiple-input version of production has, perhaps understandably, dominated the education literature. The usual output variable is some test score which measures student achievement level. Input variables include class size and a host of teacher characteristics, ranging from years of classroom experience to verbal abilities. The variables used are often averages for groups of students. Most studies include some measure of the students' socioeconomic status and/or the SES of their classroom peers as determinants of achievement. Indeed, it is the significance of SES effects--and the apparent insignificance of the teacher/classroom effects--which has been reported most frequently in empirical studies.

Production function studies of student achievement seem to have glossed over or ignored several important aspects of the nature of production in schools. Researchers do not seem to have considered, for example, that even for an individual student, the schooling process yields many different outcomes. Pupils might acquire varying cognitive skills, such as computation or reading, or they might pick up noncognitive or affective skills which, while hard to measure, are nonetheless important. A school which organizes its resources to favor production of affective outcomes might not show a strong relationship between inputs and cognitive skills.

A second important, yet overlooked, aspect of schooling is that students are usually taught in classes. Once recognized, this factor raises important conceptual questions about how output should be defined and how inputs applied to particular students can or should be measured. How, for example, should researchers characterize the reading skills of a classroom full of students? Certainly, a single number

like the mean reading score would be inappropriate, since it would not capture the results for students at the extreme ends of the scoring distribution. With the output being the entire distribution of skills among students, the measure of the productivity of inputs must take into account not only the amount of inputs, but also how those inputs are distributed among the pupils. Simply stated, the average input might not be the amount of input applied to the student with the average output score.

This should not be interpreted to mean that the appropriate level of statistical analysis is the individual student. Although a student's ability in a particular skill might be a scalar quantity, trying to measure the inputs s/he receives would be as fruitless as trying to measure many students' scores simply by looking at the mean. In a class of students taught by one teacher, both the size and nature of the input one student receives depends on how instruction is organized. Do all students receive the same lecture all of the time? Is one tutored while the others do seat work? Or is there, perhaps, a mixture of the techniques of lecturing and tutoring? None of these organizational factors has been considered in production function studies. But, as we will show, they are factors crucial to understanding input productivity, and politically potent goals like efficiency and accountability cannot ever make sense unless the organizational factors are sensitively handled.

How can we begin to handle these factors?

Recall that in a simple production process (the economist's textbook version), a single measurable output is produced by a variety of inputs. Output changes are directly attributable to input changes. This is an important simplification, since few real-world production

processes have only one output. Most, maybe all processes have many outputs; iron mines produce ore and slag, electricity generators produce kilowatts and smoke, and so on.

The process of refining crude oil into different grades of jet fuel, gasoline, tar, and paraffin is a good example of the multiproduct production case. In this process, refinery engineers are able to change the output mix by twisting the appropriate control knobs even if the input of crude remains the same. This possibility, if ignored, can lead to serious misunderstanding about, say, the production of gasoline and in estimations of the marginal product of crude oil.

Suppose there is a coal strike and the price of fuel oil rises both absolutely and relative to other petroleum products. The refinery engineers turn the knobs so that the refinery turns out relatively more fuel oil and less gasoline. Because of the price increase, economists notice that the refinery is working at a higher capacity and processing more crude oil. Suppose, too, that economists happen to do a production function estimate of the marginal product of crude oil in the production of gasoline. What they observe is that there is a rise in the amount of crude oil being processed and a fall in the amount of gasoline being produced. Of course, the laws of physics and chemistry are still operating because fuel oil production is up, but they ignore that because they are interested only in gasoline. So what do they conclude? That crude oil has a negative marginal product in the production of gasoline?

More input gives less output! This ridiculous conclusion comes from the failure to account for the jointness of the production process. A correct analysis would consider the entire distribution of commodities

produced. The output is an array of output possibilities, not the scalar quantity of the simple production case. (Note that a disaggregation procedure will not help. Assigning various proportions of the crude oil input to gasoline output, or even pretending that this is what happens in production, cannot help identify the marginal productivities.)

When students are taught in classes, we have the same kind of joint production evident that makes the refinery problem both difficult and interesting. Indeed, it is easy to show that when the distribution of skills among students changes, the relationship between inputs and average student achievement reveals little about teacher or other school input productivity.

Figure 1 shows the reading skill possibilities (test scores in this case) for a hypothetical classroom of only two students. With the level of inputs (e.g., teacher time) devoted to reading instruction held constant, various outcomes along the line BC are possible. (The line BC is known to economists as a production possibilities curve.)

S_1 and S_2 represent the reading scores for the two students. Without any instruction the students' scores would be at Point A. If the teacher were to spend all her/his time with the first student and none with the second, the scores would fall at Point B. If all the teacher's time went to the second student, the resulting score would fall at Point C. With teaching time divided between the pupils, scores represented by any point along BC is possible. (BC is a straight line because we are making a special assumption that the production process is linear.)

Notice that scores will be equal at Point D and that any movement from D will increase the dispersion of scores. The line MM', which




Figure 1--Learning possibilities for two students.

passes through D, has slope equal to -1 and represents combinations of scores which yield the same mean scores for the students. Lines with a slope of -1 lying in a more northeasterly direction (such as NN') denote higher mean scores for the class.

In using this diagram to analyze the effects of changing the amount of teacher inputs, two points become clear: (1) In general, teachers do not, and indeed should not, maximize the mean or average student score unless they feel Student Two is absolutely worthless; and (2) Observing what happens to the mean score when the amount of input changes reveals nothing about input productivity. We will try to explain these statements.

What kind of achievement levels should the teacher (or other decision maker) seek for the different pupils? Egalitarian teachers (we call them "levelers"), would choose outcomes represented by Point D, where the pupils have equal scores. Other teachers (called elitists), might prefer to maximize the difference in pupil outcomes and choose Point B. We suspect that most teachers fall somewhere between these extremes.

Suppose a teacher selected as his/her goal outcomes represented by Point E on the line NN' . An economist would conclude that the teacher arrived at this choice by applying a set of values (in economists' jargon, a "utility function") which assessed the relative tradeoffs between the different students' scores. Although the choice of point E outcomes might represent the best way for the teacher to organize classroom resources, it would not be the way for him/her to maximize the students' average score. Maximum average score would have been found at B, where all resources were devoted to Student 1.

Beginning at Point E, we will now increase the total amount of resources available. In practice, this could be done by bringing more teachers into the classroom or by allocating less time to other non-instructional activities. More resources means a new production possibilities curve, such as B'C'. (We have assumed here that resources are productive in the sense that more input gives greater output. If all resources were given to Student 1, for example, s/he could reach achievement level B' rather than B.)

The teacher must again decide how to allocate resources between the students. Suppose s/he chooses as her/his goal outcomes at Point F, more "levelling" outcomes than at Point E. At F--just as at E--resources are not being wasted, since the outcomes do not fall below the production possibilities curve.

The choice of Point F is interesting because the average score of the students is lower there even though more resources are used. (F lies below NN', so it must have lower mean scores than E.) It would be wrong, however, to conclude that because more inputs give a lower mean score they are unproductive, or even counterproductive. After all, the production possibilities clearly increased in this example when more resources were made available, and both E and F are on the production possibility curve.

What happens in cases like this is that the learning outcomes result from an interaction between the teacher's tastes or values and the technology of production. This interaction is inescapable whenever output is more than a simple scalar quantity, since the teacher must choose a mix or distribution of outputs.

The interaction can pose problems for researchers; when a given set of technologies is operated by people with different values, it is almost impossible to disentangle the effects of tastes from the effects of technologies to discover the underlying production relationships.

Figure 2 is a graphical representation of the problem; the figure indicates how the simple question of whether inputs are productive might be investigated. Line DB shows the different combinations of mean and standard deviation of scores which can be realized with the two students whose educational possibilities are shown in Figure 1. DB in Figure 2 corresponds to DB in Figure 1; the difference in the scores increases in moving from D toward B. The mean-standard deviation combinations associated with the higher level of resource use are shown as D'B'. If an outcome represented by point E is chosen when resources are relatively scarce, that outcome will yield a high mean at the cost of a high variance in student outcomes. Note that when resources are more plentiful (possibilities represented by D'B') an outcome of F lowers the differences among students, but at the cost of a lower mean score for the group.

The conditions under which a teaching resource is productive are now evident. A change in resources must change either the mean or the standard deviation of outcomes. It is possible for a resource to be productive and yet not change the mean score (for example, compare a with b in Figure 2.) But if the mean score is unchanged (or decreases) when resources increase, then the standard deviation or spread of the scores must change either up or down. This principle provided the background and basis for a statistical analysis of Michigan school district data we conducted several years ago (Brown & Saks, 1975).




Figure 2--Attainable means and standard deviations of scores for two students.

We found that resource changes often did not affect mean performance. Had that been the extent of our analysis, we would have concluded that the inputs were not productive. Yet, we also found that resource changes often affected the standard deviations of scores, leading us to modify previous conclusions on resource productivity in Michigan (see Murphy & Cohen, 1974).

Accountability

Our analysis of Michigan school district data cited above indicated that researchers will probably not be able to measure the productivity or alternate resources with statistical analyses of the kind that have been so common in the past dozen or so years. Even the analysis we conducted simply allowed us to show that the marginal products of costly inputs are positive, a proposition that no sensible person ever seriously doubted anyway. Without delving much deeper into the economics of production in the actual classroom, we are not likely to disentangle the effects of technology and taste in determining outcomes for given resources. That is why such normally pristine economists as ourselves have started looking in more detail at behavior in actual classrooms. Before we continue, however, we ought to consider what implications the work we have described has for the question of "accountability" in schools.

Schools are relatively decentralized in terms of organization of production, and teachers, compared to many other workers in our society, have a fair degree of discretion in organizing and carrying out their activities. Certainly, teachers in closed classrooms work in an environment that is quite different from an automobile assembly line, for example. One of the most substantial differences lies in the employers' ability to monitor the work. Indeed, it may be control

aspects alone (and not the division of labor) which account for the apparent efficiency of assembly lines. In the industrial sector, when a process is difficult to monitor, there is a tendency to pay workers on a piece-rate or incentive system. (Of course, if workers have internalized norms of behavior which assure efficient production--i.e., professional modes of conduct--such incentives may be superfluous.) In an increasingly bureaucratic society where more and more activities are organized through markets, the pressure is likely to grow for incentive payments and sanctions for professionals like teachers. Short of installing closed-circuit television cameras and full-time monitors in every classroom, pressure emerges to base teacher pay and status on the outputs produced by the teacher.

Should teacher pay, as some have suggested, be based on an average test score for students? According to the literature, the answer is clearly "no." Such an accountability measure would be absurd, and here are the major reasons why:

1. Different teachers work with different raw materials (quality of students), and this affects output. Attempts could be made to correct for this effect, but no one seems to be making them; we observe no move for corrective measures when average test scores are announced each year for different schools and school districts. Note that we are not suggesting low scores in a district are not a problem. What we are saying is that they are no indication of how effective a school is or how hard the teachers work. Equal outcomes would probably require unequal inputs--and we are not talking about the kind of resource inequality that currently exists.

2. In measuring outputs, educators need to consider changes in the entire distribution of students' scores--even when there is interest in only one subject, such as reading skills. It is the entire distribution which the teacher is manipulating, and to pretend that the average child is getting the average input is, as we demonstrated above, not very sensible. But how can we compare distributions that overlap? That depends on the explicit weights assigned to different parts of the distributions by those who make the evaluation. While there is much that could be done to help school boards determine such weights, we do not hear anyone even mentioning the problem, much less asking for help.

3. In developing measures of accountability, educators must decide just what it is they want to make the teacher accountable for; the right technology (which includes equipment and other inputs over which the teacher may have little control); the amount of effort allocated by the teacher; or the preferences that cause him/her to produce one distribution rather than another with given technology, resources, and effort. Educators should not expect a single measure to suffice in exercising control over so many variables whose effects are hopelessly entangled.

In short, then, work on production functions has not brought us a simple measure of teacher accountability. If anything, it has shown us why such a measure is impossible, which, perhaps, is progress in itself. We now turn to our efforts to model the economics of teaching in classrooms.

Production in the Classroom:
A Time Allocation Problem

One of the notions we began with in our discussion of the economics of teaching is that the classroom can be compared to a job shop. There are two polar cases in the organization of production. When everything to be processed must go through much the same sequence of operations (as, for example, in the manufacturing of a car), production can be accomplished along an assembly line. But when each job requires different operations in possibly different sequences (as, for example, in the repair of cars in a dealer's shop), scheduling will not only be different, but extremely complex if the shop's facilities are to be used efficiently; such operations are formally called job shops. It seems to us that teaching is more like a job shop than an assembly line (although some school authorities might try to behave as if an assembly line were possible).

To schedule time on machines in a job shop, it is essential to understand the machine's production characteristics. Scheduling cars for a wheel alignment or a lube job, for example, depends on diagnosis and knowledge of the time required to set up the different tasks and perform them. In teaching, scheduling time for various tasks (decoding, phonics, spelling, etc.) is also a management problem. We imagine that lessons of different lengths produce different amounts of learning. A typical profile is shown in Figure 3.

As shown in the figure, increasing time on task results in a positive marginal product, but output levels off, displaying what economists call diminishing returns. It should be noted that the learning curve in the figure is shaped not only by time on task, but by numerous other factors, as well, including technique of instruction,

learning produced
for an
individual student




Figure 3--Hypothetical relationship between time and learning for an individual student.

teacher characteristics, and student characteristics. If we better understood the relationships between these factors and learning, we could evaluate various rules of thumb teachers use in running their classes (the instructional strategies often taught in schools of education) and perhaps develop new ones. Certainly we could imagine simulating alternative management schemes if we knew more about the process represented by Figure 3.

At present, we are concentrating on a different aspect of the time allocation problem. The essence of the problem we are working on is that teachers care about their students' performances, but the only thing they can do to affect performance measures is to allocate different amounts of time to different techniques of instruction.

Consider, for example, the manner in which teachers teach reading. Reading instruction techniques apparently can be classified into about five different types: basal, linear skills, natural language, interests, and integrated whole. (These techniques were identified by the Conceptions of Reading group at the Institute for Research on Teaching.) It has been observed that (1) different teachers in the same school spend different amounts of time on these techniques, (2) almost all use each technique to some extent, and (3) some teachers use different mixes of technique with different students (Barr & Duffy, 1978). We contend that teachers use different mixes of the techniques not only because they value the techniques differently, but also because (1) they may value the outputs differently, (2) the outputs of the various techniques might differ for different students and for different teachers, and (3) the costs of using the alternative techniques

vary among teachers depending on their training and experience.

The first step toward a solution to the time-allocation problem is to build a model of the time-allocation process. Once we have such a model we can test our predictions against data such as those being collected by the COR group and data already collected by David Berliner and his colleagues at the Far West Laboratory.

Our research relies heavily on economic analysis of hedonic price indices and also on the work of labor economists on time allocation and labor supply. For example, in purchasing automobiles, different individuals might value different features, such as speed, comfort, or sportiness. But regardless of their preferences, they must buy cars which embody all of those features to varying degrees. Similarly, teachers may value vocabulary, speed, decoding, and enjoyment of reading, but all they can do is mix time spent on basal, linear skills, and other instructional approaches.

Our model is based on two assumptions: (1) that reading achievement consists of a vector of student performance characteristics, rather than a scalar-valued variable, and (2) that there exists a set of feasible instructional techniques from which teachers may choose in producing performance characteristics.

The relevant choice variable in the model is the teacher's time allocation to different techniques of instruction. To keep our exposition simple and tractable, we offer an example involving only two techniques of reading instruction: linear skills (L) and interests (I).⁴ We also assume, again for simplicity's sake, that

⁴Linear skills is a technique based on the premise that learning to read consists of learning a set of skills (letter recognition, phonics, word order interpretation, etc.) which are acquired in a more or less fixed sequence. The interest technique is predicated upon the notion that reading skills are acquired by getting the student interested in a particular subject matter.

reading teachers value only two different characteristics of reading, decoding (D) and enjoyment or affective qualities (A). A student's proficiency in decoding is related to the amount of time spent on each of the two production techniques--skills or interest. If this production relation is linear, we might write:

$$1. \quad z_D = v_{LD}t_L + v_{ID}t_I$$

where z_D is the amount of reading decoding skills acquired, and v_{LD} and v_{ID} are the increments in decoding skills from allocation of extra time units on the linear skills technique and on the interest technique, respectively. (That is, v_{LD} and v_{ID} are the marginal products of time for the two techniques.) The number of time units spent on linear skills and on interest techniques are t_L and t_I , respectively.

A similar linear production relationship is assumed to hold for the enjoyment-affective aspect of reading outcomes:

$$2. \quad z_A = v_{LA}t_L + v_{IA}t_I.$$

Total time available for the techniques of reading instruction, \bar{T} , is the resource in limited supply in this case and we assume that

$$3. \quad \bar{T} = t_L + t_I.$$

This means that the available time must be spent on one technique or the other.⁵

⁵This is not so innocuous as it may seem at first glance. It implies, first, that time is not wasted, or at least that wasted time is not related to the particular technique chosen. In addition, it implies that a unit of time cannot be used in both techniques at once; that is, that the two technologies are independent. In economists' jargon, there is no joint production.

We can substitute Equation 3 into Equations 1 and 2 to get the corresponding production relations:

$$1A. \quad z_D = v_{ID}\bar{T} + t_L (v_{LD} - v_{ID}), \text{ and}$$

$$2A. \quad z_A = v_{LA}\bar{T} + t_I (v_{IA} - v_{LA}).$$

Equations 1A and 2A make it clear that the effects of changing time allocation depend upon the relative productivities of the techniques in enhancing the two output characteristics. For example, increasing the time spent on the linear skills technique (t_L) will increase decoding skills only if the marginal product of the skills technique in teaching decoding exceeds that of the interest technique in teaching decoding.

Equations 1, 2, and 3 (or 1A, 2A, and 3) can be manipulated, by eliminating the variables t_L and t_I , to indicate the array of outcome possibilities (z_D and z_A) for any given level of total time available, \bar{T} .

Figure 4 shows the solution. In the graph, all axes measure positive quantities starting from the origin. The southwest quadrant shows the graph of Equation 3 for a given \bar{T} and indicates the possible time allocations for linear skills and interest technologies. The northwest quadrant shows Equation 2A (assuming that $v_{LD} > v_{ID}$). For any choice of the independent variable t_L , we can find the resulting level of the decoding score. Equation 2A falls in the southeast quadrant where the enjoyment-affective outcomes, z_A , are a function of time spent on the interest technique T_I .

The linear relation PP' in the northeast quadrant represents the outcomes that can be obtained from different time allocations. Points on this line are derived as follows. A time allocation between




Figure 4--The relationship between time allocated to two techniques of instruction and the resulting output possibilities.

techniques (e.g., point s) is chosen. According to the production relations, that allocation will yield unique amounts of decoding and affective outcomes (e.g., points m and n). In this case, allocation s yields production mix r on the locus of production outcomes, PP' .

Which time allocation and set of outcomes are best? We assume that teachers value both kinds of reading outcomes and that their preferences for different combinations can be summarized by a utility function of the kind shown in Figure 5. Each curve in the figure represents output combinations about which the teacher is indifferent and to which she attaches the same level of utility. The contour U_2 represents better outcomes; that is, it represents higher utility values than those along U_1 . The task here, then, is to find the output combination (and accompanying time allocation) which will get the teacher to the highest contour.

The solution to the problem is shown in Figure 6, in which the best set of outcomes lies at point a. An interesting feature of this solution is that the teacher optimally employs both instructional techniques; neither reading technique dominates the other. In an optimally run classroom we would expect to find a variety of instructional methods being used. (This theory, however, is completely dependent upon our assumption about tastes and on our assumption that each technique has an absolute advantage in producing one of the valued characteristics.)

Two important insights can be derived from this discussion of the time allocation problem. One is that variations in time allocations depend on variations in the outputs of alternative techniques and on how teachers value those different outcomes. The other is that mixtures of techniques will be found only in situations where each




Figure 5--Utility contours of one teacher for two different learning outcomes.




Figure 6--Utility maximizing combination of learning outcomes when the teacher must choose along the line PP' .

technique has some absolute advantage in producing a particular characteristic.

As these insights suggest, observation of time allocations can provide important information about the valuation of different characteristics and about the productivity of different instructional techniques. We must observe time allocations (and output changes) in different settings and be clever enough to extract information from these observations through various statistical procedures. If, for example, we can assume that technology is fixed for some group of observations, but that tastes vary, then we might be able to identify the slope of the production frontier and the relative productivities implied. If we can use techniques developed by psychologists to identify some parameters of the utility function, we might do an even better job of identifying the production technology.

By estimating this underlying structure, we might be able to find more efficient mixes of techniques. Furthermore, if new techniques are developed, and if their characteristics and the values different teachers attach to those characteristics are known, it might be possible to predict and even influence how the techniques are adapted. We do not want to suggest that obtaining information about the underlying structure will be easy; we do believe, however, that it is desirable and that it is possible, using the kind of data which could be made (or already are) available.

Before we can move further toward such an estimation, however, we must consider some extensions of our model.

The first and most important deals with joint production. We have outlined a simple model for one student, but we need to know what happens when students affect and are affected by other students and when time on task is spread across groups of students in non-uniform ways. We know from our previous work that the solution to joint production problems depends largely on aggregating up to the level where externalities are internalized.

It is also important to learn how set-up time and the diminishing marginal product of time spent on any activity affect the time-allocation problem. Perhaps this could be accomplished by having the v 's depend non-linearly on student and teacher characteristics, time of day, and time already spent.

It is essential, too, that we generalize each part of the model to make it realistic. There are more activities that must be considered (e.g., hugging the kids, teacher resting at desk); more characteristics that might be valued (e.g., motivation, discipline, speed, and perhaps the techniques themselves which have psychic and actual costs associated with teacher training and experience); and there is likely to be variation in the utility function itself, depending on the teacher, school district, principal, organization of the school, and other factors.

Finally, judgments must be made about the appropriate time period for the analysis and to determine whether or not adjusting time allocations will cause teachers to be far from their optimal allocations when we observe them.

We do not expect to come to any conclusions about the minute-to-minute operation of the classroom. We do hope, however, that we will be able to say something about weekly, monthly, or even yearly time-allocation strategies. Reading instruction experts will still have to devise the specific techniques; economists might help in evaluating those techniques, in suggesting how they may be more profitably mixed, and in suggesting how adoption of such mixtures might be achieved in relatively decentralized organizations such as schools.

Future Research by Reading Specialists

We are uncertain about the long-term value of studying reading instruction from an econometric perspective. Obviously, we think there is some promise in this line of research or we would not be spending our time this way. Based on our past experience, we expect that some reading researchers will agree with us and that they will consider the implications of the econometric research for their own research agenda. This work must also involve practitioners, who are best qualified to help answer the questions which econometricians need answered before they can make any progress with their studies.

What are some of the questions which econometricians need to be answered? First, we need to know what the time-learning profiles (Figure 3) look like and what they depend upon. Do they vary with the content of the material taught, the sequence, the materials already learned, the teacher's characteristics, the school characteristics, the specific techniques of instruction, the organization of the classroom, the lapse in time since the last lesson, the other subjects being taught? We are not likely to obtain clear answers to these questions, but we need some helpful taxonomy before we can even

begin any research. We need to know, for example, how many different components there are to reading instruction. Can we really classify the types of techniques used for instruction in the way we have? How can we get into the classroom to observe these things without hopelessly altering normal activities? Are there cheap ways of identifying the components of reading instruction (by analyzing teacher logs, for instance)? Are there adequate output measures for the different components of the taxonomy of reading? In one form or another, reading researchers have been working on these problems for a long time, but we need to attain some consensus or at least a probable range of variation.

Econometricians also need to understand how teachers form preferences so that we can make some progress in disentangling the taste and technology question. We need to know how difficult it is for a teacher to adopt alternative technologies (i.e., what the cost of changing methods is) and what determines those costs. We need to know what rules of thumb teachers use when making allocation decisions and how flexible those rules are.

Most reading researchers are probably not used to asking questions in quite the form we are asking them and in making the kinds of distinctions we like to make. But if we can obtain answers to these questions, and if we can make these distinctions, we might be able to give some useful structure to the massive quantities of research that have been generated through the years. Maybe in the long run the value of our research will be in helping to form some consensus, for a few years at least, on the important questions which need to be asked. But, of course, nothing can happen unless we listen to each other--and foreign languages are always tough at first.

References

- Averch, H.A., Carroll, S.J., Donaldson, T.S., Kiesling, H.J., & Pincus, J. How effective is schooling? A critical review and synthesis of research findings (Santa Monica, CA: Rand, 1972).
- Brown, B.W., & Saks, D.H. The production and distribution of cognitive skills within schools, Journal of Political Economy, 1975, 83(3), 571-593.
- Murphy, J.T. & Cohen, D.K., Accountability in education--the Michigan experience. The Public Interest, No. 36 Summer 1974: No. 36, 53-81.