

DOCUMENT RESUME

ED 298 208

UD 026 341

AUTHOR Elmore, Richard F.
 TITLE The Problem of Quality in Chapter 1.
 PUB DATE Oct 87
 NOTE 14p.; In: Policy Options for the Future of
 Compensatory Education; see UD 026 334.
 PUB TYPE Speeches/Conference Papers (150) -- Reports -
 Evaluative/Feasibility (142)

EDRS PRICE MF01/PC01 Plus Postage.
 DESCRIPTORS *Compensatory Education; Educationally Disadvantaged;
 *Educational Quality; Elementary Secondary Education;
 Federal Aid; *Federal Programs; *Federal State
 Relationship; Program Administration; Program
 Budgeting; *Program Development; *Program
 Effectiveness; Program Evaluation; Program
 Implementation; State Programs; Urban Schools
 IDENTIFIERS *Education Consolidation Improvement Act Chapter 1

ABSTRACT

However defined, the quality of services under Chapter 1 of the Education Consolidation and Improvement Act is heavily dependent on the setting in which those services are delivered. Working definitions of the following aspects of program delivery are presented as a means of establishing evaluation criteria: (1) the resources applied to a local program (input); (2) the operating characteristics of local programs (process); and (3) the consequences of local programs for students (output). The following Federal strategies for enhancing quality are reviewed: (1) the systematic use of information to draw attention to either good or bad behavior (jawboning); (2) the use of discretionary funding to reward and claim credit for local successes (piggybacking); and (3) the use of discretionary funding to underwrite program development in the most difficult circumstances (bootstrapping). A short list of references is included. (BJV)

 * Reproductions supplied by EDRS are the best that can be made *
 * from the original document. *

ED 298208

THE PROBLEM OF QUALITY IN CHAPTER 1

by

Richard F. Elmore
Michigan State University
and
Rutgers University

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)

- This document has been reproduced as received from the person or organization originating it.
- Minor changes have been made to improve reproduction quality.

- Points of view or opinions stated in this document do not necessarily represent official OERI position or policy

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

Peggy D. Richmond
Research & Evaluation
Services, Inc.

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

UD026 341



THE PROBLEM OF QUALITY IN CHAPTER 1

The Problem

How can Federal policy enhance the quality of local programs in Chapter 1? This is a plausible question, given the amount of Federal money spent on Chapter 1; the considerable Federal, state, and local program experience accumulated over the past twenty years; and the needs that the program will be expected to meet over the next twenty years. Yet relatively little attention has been focused explicitly on the issue of quality in Title I/Chapter 1, and there is relatively little systematic understanding or analysis of the issue in the evaluation literature growing out of the program.

Everyone—members of Congress, Federal, state, and local administrators, teachers, parents, evaluators, and analysts—is concerned about quality. The rationale for a federally funded compensatory education program rests heavily on assumptions about quality: Additional money buys compensatory services for a small portion of the population for a relatively small portion of the school day, and this money is assumed to provide the difference needed to pull a significant proportion of the educationally disadvantaged into the educational mainstream. If the additional money introduced by Chapter 1 buys no more than the service a student would otherwise receive in the regular school program, then the key assumption underlying compensatory education is faulty.

For all its importance, though, we know very little about what quality is in educational programs and even less about how to produce it reliably with policy. Seemingly straightforward definitions of quality, e.g., concentrations of resources, staff characteristics, program design, and student performance, raise serious operating problems when they are translated into policy. These problems, we shall see, are intrinsic to a large-scale Federal grant program. So while quality is central to a program like Chapter 1, it is also difficult to define and even more difficult to realize in practice.

Consequently, the fiscal and administrative machinery of Chapter 1 does not deal explicitly with quality. In some cases, the program has mandated or encouraged practices that are assumed to be associated with quality, such as concentration of funds, but it is not clear that these practices actually result in higher quality local programs. In other cases, the program has encouraged local practices that are questionable under certain definitions of quality, e.g., pullout programs. In still other cases, the program takes a deliberately agnostic posture toward quality. For example, with fiscal accountability the program says, in effect, that protections against displacement of local funds by Federal funds are essential even if they make it more difficult to mount high quality local programs.

The single most important characteristic of Chapter 1 is that it is a marginal program (for a more detailed discussion of the marginal role of Federal policy, see Elmore & McLaughlin, 1982). That is, Chapter 1 "works," if at all, by augmenting existing instruction in existing schools. The typical recipient of Chapter 1 services is exposed to about three-quarters of an hour of reading and math four times a week with about ten students and two adults in a separate classroom (Advanced Technology, 1983). Depending on how it is used this can be a significant increment to a student's education, but it is still marginal, in several senses. It constitutes about 12 to 15 percent of a student's time in school; it draws on available staff within the school or district, who reflect the general ability level in those settings; it augments an existing instructional program which is determined by local preferences, state requirements, and local fiscal capacity; and it works on young people who come from home environments that influence their orientation to learning in various ways. My point is simply that the "quality" of Chapter 1 services, however defined, is heavily dependent on the setting in which those services are delivered.

Quality is important to Chapter 1. But it is difficult to define and even more difficult to realize once defined. It is also, once defined and translated into administrative machinery, heavily dependent on the setting in which services are delivered.

Definitions of Quality

Acknowledging these problems is not to say that quality is unimportant, that it is unachievable on a large scale, or that concern for quality cannot play a major role in shaping the future of Chapter 1. In this section, I will develop working definitions of three types of quality that might be used to shape Federal policy toward Chapter 1—the resources applied to a local program (inputs), the operating characteristics of local programs (process), and the consequences of local programs for students (outputs). And I will speculate about the strengths and limitations of using each of these approaches as a basis for policy.

Input Standards

In its simplest terms, Chapter 1 supplies money to states and localities to purchase compensatory education. This money is spent on certain things—teachers, aides, administrators, instructional materials, tests, etc.—and these expenditures are related in some way to the services students receive. Some of the things on which money is spent—teachers, for example—have attributes—such as experience, training, and knowledge—that may also affect the services students receive. Finally, these expenditures funded by Chapter 1 supplement existing expenditures for the same things—teachers, materials, administrators, etc.—in the broader academic program. This package of things purchased by Chapter 1, their

important attributes, and the base level of resources constitutes the package of "inputs" that students receive.

One could, in theory, design a set of indicators based on inputs and use those indicators to define an acceptable range of quality. For example, these indicators might stipulate that a per student Chapter 1 expenditure of \$400-700, on a local instructional base of \$1500-\$2000 per student, would have a reasonable likelihood of purchasing the package of resources (teacher skill, experience, materials, evaluation, etc.) necessary for a program of acceptable quality for a particular number of students. One could also look in a more fine-grained way at packages of inputs and attempt to determine which packages seem to have the closest fit with student needs or outputs.

An input-driven system could have several possible consequences. First, by defining an acceptable or exemplary range of inputs, it would tend to focus the attention of Federal, state, and local administrators on the delivery of resources to students. Second, by calling attention to the mix of inputs and their characteristics in local programs, it would make explicit certain trade-offs involved in local administration. If additional years of teacher experience are costly, for example, some local administrators might trade less instructional time by more experienced staff for more instructional time with less experienced staff. While these trade-offs are inherent in the administration of any service delivery program, they are seldom made explicit, and are consequently seldom understood by policymakers or administrators. Third, by making explicit the relationship between Chapter 1 expenditures and the state and locally funded instructional base, it would call attention to distributional variations in the actual resources reaching students served by Chapter 1, rather than the distribution of Chapter 1 funds alone.

Input-driven systems raise certain problems, though. While it may be reasonable to stipulate acceptable or exemplary ranges on inputs, there is no reason to believe that inputs at the gross level measurable by per pupil funding, student-staff ratios, and Chapter 1/base program ratios will be meaningfully related to other measures of program quality. In simple terms, the same level of inputs "buys" very different program characteristics and very different student performance from one setting to another. There may be discernable correlations between inputs, programs, and performance in the aggregate, but those correlations conceal enormous variations in local practice.

Specifying acceptable or exemplary ranges of inputs also sends certain signals to state and local administrators about what Federal policymakers and administrators value. Used by itself, an input-driven system says, "if you keep certain indicators within a certain range, we don't care what else you do with the money." This philosophy has certain advantages. It does not prejudge whether certain types of instructional programs are effective. (I will use the term "effectiveness" throughout to refer only to effects on students.) But neither does it provide incentives for local administrators to look for more creative, innovative, or effective ways of using

Federal funds. In this respect, input-driven quality indicators are much like the current regulatory regime of Chapter 1, which values compliance with administrative guidelines more than the search for new or more effective programs.

Perhaps the most serious problem with input-driven measures of quality is their insensitivity to local context and student background. Two of the most robust findings of research on the effects of schooling in general and Chapter 1 in particular are (1) that student performance is strongly related to the race, family income, and family resources of students; and (2) that the higher the concentration of minority, low-income students, the lower the achievement level of the school (See Kennedy, Jung, & Orland, 1986). To be sure, these findings describe overall patterns; there are many important and interesting exceptions that merit study in their own right. But the fact is that student background exerts strong influences on educational programs, both in terms of what must be taught and what it is possible to expect by way of student achievement. Input measures focus attention on the allocation of resources to schools and classrooms, rather than on the characteristics of the student population served. They carry the assumption that a certain standard of input should provide an adequate level of service to students. In fact, though, the educational problems of so-called disadvantaged students vary widely and have very different resource implications. Standardized input measures discourage attempts to treat different types of students differently.

Process Standards

Another way to think about quality is in terms of the design of local programs and the processes that surround them. Chapter 1 can be thought of as buying certain packages of instruction, for certain types of students, within certain stipulated structures and procedures. In their pure form, process standards would take as given the existing allocation of resources to districts and schools and focus on the package, or alternative packages of instruction, students, and procedures that a given level of resources would purchase.

The level of detail at which process standards can be specified is a major political and administrative issue. The Federal government has long adhered to the principle that it does not dictate curriculum content to states and localities, so requiring the selection of model curricula for local Chapter 1 programs is not a feasible option, even if it were desirable. But the Federal government does have a long tradition of sponsoring research, development, and evaluation designed to identify exemplary educational practices. These date back to the 1950s (with the Cooperative Research Act), through the Sputnik Era (with the development of model science and math curricula), into the 1960s and 1970s (with the What Works series and the National Diffusion Network), up to the most recent publication of compendia of successful educational practices. While it does not seem feasible that the Federal government could require specific

types of instructional programs, even if it were desirable, it is feasible for the Federal government to reward states and localities differentially for engaging in certain types of practices. I will return to this issue in the final section on strategy.

Process-driven standards of quality could have certain consequences for local programming. First, if they were constructed on the assumption that a "program" is a package of students, content, and structure, they would call attention to the fit between student characteristics and program content. The student population served by Chapter 1 presents relatively diverse problems—content-specific achievement (reading, math, etc.), language acquisition, behavior, motivation, etc. Different combinations and levels of problems presumably require different levels of resources, different types of content, and different structures (grouping practices, teacher-aide combinations, etc.).

Second, unlike input and output standards, process standards focus local administrators' and teachers' attention on educational practice, rather than on the allocation of dollars or on the measurement of achievement. Implementation of Chapter 1, under process standards, would increasingly consist of finding the appropriate match of students, content, and structure, rather than meeting some predetermined mix of resources or student achievement level.

Process standards have their own characteristic problems. The state of knowledge about the appropriate fit between student characteristics, content, and structure is far from amenable to straightforward, easily implemented prescriptions. There is such a base of knowledge, it does have useful implications for educational practice in local compensatory programs, and its systematic application to those programs could improve content and performance. But saying that such knowledge could be useful is something very different from saying that it should be turned into authoritative standards for local practice.

All conclusions about effective practice are statements of average relationships with large intervals of uncertainty around them. Among the most significant sources of uncertainty are the existing skill and orientation of teachers and the characteristics of the process by which new practices are introduced to schools and classroom. Certain packages of students, content, and structure will perform according to expectation in certain settings, but will not in other settings. So process standards send incomplete signals to state and local administrators. They say, "this kind of program generally works well," but they do not allow for the fact that it will fail to produce the same effect across diverse settings. One of the most robust findings of research on policy implementation is that for a given array of program models, variation across sites in degree of implementation and effectiveness for a particular model is considerably greater than variation in average performance among models. In other words, the setting in which the model is implemented exercises more influence than the model itself.

Another significant problem with process standards is that they tend to set artificial constraints on the development of new practices. For example, in environmental policy, standards for the discharge of pollutants were once determined by industry-wide agreements on "best practicable technology." After pollution abatement technology was installed the technology continued to change, but the standards reflected earlier technology. This introduced a complex set of problems about whether to update standards, whether to apply them retroactively to firms that were earlier in compliance, and whether less stringent technologies could be used in areas where pollution was not as serious a problem. In education, one could imagine the same kind dynamic. Exemplary practice at one time could become obsolete or retrograde at another time. Exemplary practices at one time could become institutionalized and rigidified, making it more difficult to introduce new practices at some future time.

Finally, process standards send an important signal to local administrators and teachers about who is responsible for generating knowledge about successful practices. In effect, they say that knowledge about what works is generated outside local settings by people with expert knowledge and that these ideas are then supplied to school practitioners who often screw up their execution because they don't have the skill or understanding to do them correctly. Expert knowledge is important to the development of educational practice, but this version of the relationship between the two is destructive to the development of professional responsibility within schools.

Output Standards

A final way to think about quality is in terms of effects on students. In this view, high quality local programs are programs that produce results, measured by such indices as students achievements, attendance, attainment, reductions in dropout rates, and the like. In this view, Chapter 1 is a way of purchasing capacity in local districts for the purpose of remedying performance problems for certain parts of the student population. In their pure form, output standards make no assumptions about the correct allocation of inputs or the best fit between students, content, and structure. They stipulate either that a given level of performance is expected from a given level of funding, or that performance above some level will be rewarded.

Output standards send a signal to local teachers and administrators that says, "do whatever is necessary to produce these effects with this amount of money." They also say implicitly that the necessary knowledge for solving detailed problems of student selection, content, and structure ultimately lies within the organizations that deliver educational services. No amount of external prescription of inputs or processes will supply the knowledge needed to mount an effective program in the absence of a strong incentive to mobilize that knowledge. Output standards supply that incentive.

The technical problems associated with output standards are similar to those associated with process standards. That is, the ability to set output standards assumes a technical knowledge of how much of a given output can be attributable to a given infusion of money. Otherwise, there would be no feasible relationship between rewards and expected effects. Empirical estimates of these relationships are statements of average relationships that conceal large variations.

Given this diversity, where should output standards be set? If they are set at the median level of performance, then half the distribution will be below standard, by definition. If they are set at a level significantly below the median, then policymakers appear to be endorsing below average performance. Another option is to set output standards by criterion levels, rather than by reference to a distribution. That is, policymakers could say that a compensatory program should, at a minimum, produce people who know how to read and write a complete sentence, do complex multiplication and division, etc. There is considerable expertise at the national and state levels in constructing such tests, so the issue is not whether it is possible to measure outputs, but rather how feasible it is to use them as indices for enhancing quality in Chapter 1.

When output standards are applied to programs like Chapter 1, which serve only a fraction of their eligibility clientele, they introduce strong selection incentives. Local programs can improve their performance by changing their student composition, rather than by improving the quality of the programs. Sometimes these selection incentives operate to focus administrators' and teachers' attention on what kind of students can best be served by a given array of services. In this case, they tend to improve the fit between program characteristics and student characteristics. On the other hand, Chapter 1 is explicitly designed to focus attention on students who are the most difficult to reach. To the extent that output standards reward selection of eligible students who present the least difficult problems, they undermine Chapter 1's central purpose.

Specifying which outputs will serve as performance measures is a complex and slippery task. Student achievement, measured by norm-referenced or criterion-referenced tests, is an obvious choice. But no responsible analyst or policymaker would advocate the exclusive use of achievement measures, since they typically tap only a narrow range of skills and they fail to capture important dimensions of what compensatory programs are trying to do. So a responsible array of output measures would include other measures—attendance, attainment, attitudes, dropout rates, and the like. But specifying multiple output measures creates the problem of how much value to attach to which measures. Will we accept a lower average achievement level in return for higher attainment and lower dropout rates? That is, is it more important for low-achieving students to stay in school or for those who stay to learn more? Relationships among output measures are very poorly understood.

An output-driven system is predicated on the assumption that the increment of funding introduced by Chapter 1 is significant enough to produce certain effects with some consistency across very diverse settings. This assumption is unlikely to be true, since variations in funding for the basic instructional program in Chapter 1 schools are likely to be considerably more than the value of the additional money introduced by the program. Other variations are also likely to overwhelm the effect of the Chapter 1 increment—teacher skill, instructional content, and the characteristics of the populations served. In theory, it is possible to produce estimates of student outputs that control statistically for these background variations. In operation, these systems introduce unresolvable methodological wrangles into debates on program effects. One can imagine the spectacle of localities suing states or states and localities suing the Federal government over the legal and statistical validity of funding decisions based on student outputs.

Several feasibility questions cut across input, process, and output standards of quality in Chapter 1. Most of these questions stem from the fact that Chapter 1 is a complex intergovernmental grant program. At least in my formulation, standards of quality apply to local program decisions, which seems logical because Chapter 1 services are mounted and delivered in local schools and school systems. But Chapter 1 as a political and administrative system is considerably more complex than this formulation suggests. States play a significant role in the routine evaluation of Chapter 1, and would have to play a significant role in any attempt to introduce quality standards to the program. States have dramatically different policies toward these quality issues and dramatically different capacities to influence local decisions about the administration of Chapter 1. Hence, federally initiated quality standards, by any definition, would be differently implemented by different states and state actions would be differently implemented in local settings.

Another closely related question is the effect of existing policy and regulations in Chapter 1 on attempts to influence quality. The existing policy and regulatory structure of Chapter 1 represents a carefully constructed resolution of a myriad of issues that have arisen over the history of the program: targeting, nonpublic recipient, displacement of local revenues by Federal revenues, etc. For the most part, these issues have been resolved by stimulating increasingly specific Federal requirements that apply to an increasingly narrower range of activities. Since the mid-1970s, the Federal government has not tried in any serious way to influence the actual instructional content of local Chapter 1 programs or the standards by which those programs are judged.

Hence, a new concern for quality, manifested in standards, would be perceived by the Chapter 1 subgovernment—interest groups, state and local administrators, and clients—as a significant shift in Federal policy. Regardless of how those standards were applied, the move would be perceived as signaling a new period of Federal activism in Chapter 1, following a long period of relative passivity. The Federal role—in Chapter 1 and in education generally—has long been

characterized by these cycles of activity and passivity. But there is some question whether the Federal government is prepared to develop the kind of administrative capacity necessary to oversee an initiative aimed at increasing quality in Chapter 1, given the foreseeable fiscal situation.

Federal Strategies for Enhancing Quality

Here is the situation confronting Federal policymakers: education policymakers at all levels have a great concern for the quality of the educational program provided to public school students. Many states and localities are already engaged in quite ambitious and detailed programs to enhance quality. Chapter 1 is unlikely to be exempt from this broader concern. By any of the definitions discussed here—inputs, processes, or outputs—serious conceptual and practical problems impede the introduction of quality standards to Chapter 1. The issue boils down to how Federal policymakers might seriously broach the issue of quality in Chapter 1, while simultaneously acknowledging the problems raised by alternative definitions of quality.

One overall conclusion seems clear. The Federal government will not muster the legislative authority or the administrative capacity to impose standards of program quality directly on local districts. In the present political and fiscal context, it is highly unlikely that the issue of quality will be addressed by writing Federal regulations that specify input, process, or output standards. For all the reasons sketched out above, writing such regulations and enforcing them would be a difficult and dubious enterprise under the best circumstances.

Ruling out direct Federal regulation of program quality does not mean that the Federal government has no resources for influencing the quality of local Chapter 1 services. A Federal concern for program quality can be manifested in ways that are consistent with a limited Federal role and with the serious practical and conceptual issues that underlie alternative definitions of quality.

By way of example, let me sketch three alternative Federal strategies for focusing increased attention on program quality in Chapter 1. I will call these strategies jawboning, piggybacking, and bootstrapping. They correspond to relatively well-established strategies already in the Federal repertoire, so they do not require extensive strategy departures from existing practice. They all take account of the limits on Federal influence imposed by the marginal nature of Chapter 1 and the wide variation in capacity at the state and local level. And they are all consistent with relatively modest changes in Federal Chapter 1 expenditures.

Jawboning is essentially the systematic use of information to draw attention to either good or bad behavior. When school systems or state agencies use student test scores to call attention to high-

or low-performing schools, or when the Secretary of the U.S. Department of Education publishes a pamphlet describing exemplary practices in schools, the expectation is that people in schools will pay serious attention and change their practices over time. The key characteristic of jawboning is that it involves no material rewards or penalties, and only occasionally praise or humiliation.

Jawboning is a weak treatment, in the sense that it does not use direct intervention to influence behavior. It is relatively powerful in shaping peoples perceptions of ideas in good currency. It legitimates certain practices by giving them the imprimatur of authority.

Chapter 1 carries the dominant image of a social pork barrel program, whose primary purpose is to distribute large amounts of money as broadly as possible among states and localities with a socially and politically defensible rationale. This image has been reinforced by the politics of Chapter 1, which consist mainly of education interest groups and congressional allies defending the program against its critics by invoking the kids, but never taking the lead in giving it a positive, ambitious new agenda. Federal leadership could play a major role in changing this image by engaging in systematic attempts to surface and publicize information about exemplary local programs and practices. Leadership might also call attention to programs and practices that undermine Chapter 1's effectiveness. The What Works series of the 1970s was a pallid version of jawboning, but it reached a fairly narrow audience and it never established real authority with the educational community.

Piggybacking is the use of discretionary funding to reward and claim credit for local successes. Like jawboning, piggybacking involves some kind of systematic surfacing of exemplary local programs, but unlike jawboning, it involves the deliberate use of financial rewards. For example, states might be asked to identify some number of exemplary local programs according to federally mandated criteria as part of their evaluation responsibility under Chapter 1. The criteria could include the nature of the student population served and the creative use of federally funded activities to complement the basic instructional program. These programs could be asked to propose a plan for how they might use some significant increment of funds over a two- or three-year period, and the progress of their efforts before and after the awards could be described in a literature designed to reach a broad audience of practitioners and policymakers.

The principle underlying piggybacking is to bankroll creative local people, call attention to their efforts, and claim credit for some portion of their success through the use of financial rewards. Publicizing their efforts lends authority to a view of Chapter 1 as aggressively searching out and rewarding creativity.

Bootstrapping is the use of discretionary funding to underwrite program development in the most difficult circumstances, with the least likelihood of success, and to claim credit for success against the odds. For example, states might be asked to nominate elementary

and middle schools feeding into high schools with high dropout rates and low academic achievement records. Those feeder schools could be the recipients of discretionary funding to develop Chapter 1 programs around the explicit objective of guaranteeing high school completion and meeting achievement standards. The discretionary funding could follow a cohort through the system, or, more likely, it could support the development of activities designed to complement the regular school program for all students. These efforts could be described in a literature designed to reach a broad audience of practitioners and policymakers.

The principle behind bootstrapping is to use discretionary funding to score successes on the most difficult problems confronting compensatory education, and to use those successes as a goad to the rest of the program. Publicizing these types of success lends authority to the view that Chapter 1 searches out the most difficult problems and finds solutions.

None of these approaches to quality is likely to be politically popular with Chapter 1's traditional political constituency. Any effort to distinguish among more or less successful programs will not be greeted with enthusiasm by a coalition whose main collective interest is in preserving its funding base. There may be a political constituency supportive of these ideas among local teachers and administrators and among friendly critics of the program who see it as having lost sight of its original compensatory purposes.

A major advantage of these approaches is that they do not require the Federal government to endorse any single operating definition of quality, but only to specify broad criteria (which may be based on inputs, processes, and outputs) and allow states and localities to grapple with the problems of defining quality operationally. This approach is, I think, consistent with a view which says quality is essential to the success of Chapter 1 even if we don't know exactly what it means and even if any single definition of it leads to consequences we may not like.

REFERENCES

- Advance Technology, Inc. (1983). Local operation of Title I, ESEA 1976-82. McLean, VA: Author.
- Elmore, R., & McLaughlin, M. (1982). Strategic choice in Federal education policy: The compliance-assistance trade-off. In A. Lieberman & M. McLaughlin (Eds.), Policymaking in education (Eighty-first yearbook of the National Society for the Study of Education. Chicago: University of Chicago Press.
- Kennedy, M. M., Jung, R. K., & Orland, M. E. (1986). Poverty, achievement and the distribution of compensatory education services (First interim report of the National Assessment of Chapter 1). Washington, DC: U.S. Department of Education.