MODELING SCHOOL FINANCE POLICY
Using Simulation to Test A Priori Assumptions

George P. Richardson
Wheaton College

David F. Andersen
The Rockefeller College, SUNY-Albany

Robert E. Lamitie
Connecticut Department of Education

Abstract

In many of the United States, attempts to reform the funding of public education to reduce disparities among school districts have been disappointing. For this study of state aid-to-education in Connecticut, a system dynamics model of local district budgeting behavior was formulated, replicated seven times to represent all the school districts in Connecticut divided into seven clusters, and linked to a state-aid sector that dispenses aid according to the state's current guaranteed wealth formula. Simulations of a single cluster and the seven-cluster model suggest that the guaranteed wealth formula will not, by itself, narrow the gap in per-pupil expenditures in rich and poor districts, as intended.

The study described is a relatively pure case of policy modeling designed not to confirm hypotheses but rather to draw inferences from puzzling sets of assumptions about state and local behavior. Reflecting on this case, we suggest six conditions that appear to indicate that a policy problem is ripe for the sort of contingent, inference-generating analysis illuminated by the Connecticut study. Some speculative methodological connections to cognitive science are also suggested.

Introduction

In many states in the United states, reforms of funding policies for education have been disappointing to policy-makers and puzzling to policy-analysts. The experience of many states is not unlike that of Connecticut, which has implemented a funding policy designed to equalize spending in rich and poor districts, yet watched inequalities stubbornly persist and even increase.

In the analysis reported in this paper, a team of external simulation modelers worked with a team of policy analysts within the Connecticut State Department of Education to construct a dynamic simulation model to analyze the impacts of proposed changes in the funding of elementary and secondary education in

Connecticut. In the course of that analysis, a series of important assumptions were accepted with no formal, empirical verification, including: the nature of the school funding problems facing Connecticut, how local districts respond to the availability of state aid, and how local district expenditures ought to change in response to state policy initiatives. Both the client and the modeling team completed the exercise feeling that important questions had been addressed, that useful policy insights had been gained, and that in general the overall project was a success.

The simulation exercise was completed essentially in isolation from the formal analysis of time-series data drawn from the real world. Of course, the policy team from Connecticut drew upon a wealth of experience with school funding data, and current information was used to parameterize the models. However, no time-series plots were put forward as reference modes or validation data for the modeling effort. Such data is readily available (although increasingly called into question because of the disappointing results of policy changes in a number of states including Connecticut). Given that data were available and that assumptions about local fiscal response are methodologically complex and controversial, the very existence of such a policy analytic exercise raises some troubling questions. Can useful formal analysis proceed in the absence of empirical data? Can policy implications of such analyses be given any weight in shaping state policy? Finally, is there any interpretation in which such analyses can be considered good social science?

We believe the answers to these questions is a qualified yes. There is a sense in which certain formal analyses can be useful without reference to statistical analyses of patterns found in real data. Such "data free" exercises can indeed constitute good social science. They are useful, not at the stage of hypothesis construction (when observation of the real system is crucial), nor the stage of hypothesis testing or confirmation (when inferential statistical methods are applied to real data), but rather in the stage of testing the logical consistency of complex, interconnected assumptions taken as a whole. This last stage is essentially nonempirical and syllogistic in nature, involving thought experiments, and simulation experiments, to test the integrity, the consistency or wholeness, of a policy picture.

This point of view developed as we reflected upon the case described in detail below. The models built for Connecticut were a fairly pure case of testing the logical consistency of many often conflicting assumptions about how the school finance system either did function or would function under variations in state funding policies. State policy has been informed by implicit and conflicting assumptions about how local districts did or ought to respond to sustained or increasing state aid; how the system as a whole did or ought to respond to a court-mandated formula designed to equalize per-pupil expenditures in all districts across the state; and how the present pattern of unequal dollar expenditures per pupil had emerged over time. The State Department of Education, its policy staff, and policy reports presented to the courts were filled with interesting, often competing, but always intercoupled hypotheses related to these questions. Further, some of these assumptions had been the subjects of confirming empirical studies employing formal statistical techniques. What was needed from the simulation study was, in part, a logical exercise to sort out the consistencies, inconsistencies, and patterns of interaction among these assumptions. The essentially data free, syllogistic exercise was carried out through the construction of two dynamic feedback simulation models.

The Case Background

In Connecticut, as in most of the United States, the constitutional responsibility for public education has been delegated to local governments, which differ greatly in their fiscal capacity to support school services. Local schools are supported by a local property tax, supplemented by state and federal aid of various sorts. However, the ultimate responsibility for assuring the same quality of education in all the districts of a state rests with the state government.

In educational finance circles, and in the courts, the commonly accepted proxy for educational quality is the number of dollars expended in a school district per pupil per year. Prior to 1977, per-pupil expenditures in Connecticut differed greatly, with wealthier districts spending more per pupil and poorer districts spending less. The result was that the quality of education received by a child, as measured by the dollars spent on it, was highly dependent on the wealth of the community in which the child lived. The disparity in per-pupil expenditures existed even though there was a clear and strong pattern of higher school tax rates in poorer districts and lower rates in wealthier ones.

In 1977, in a decision having much in common with rulings in other states, the Connecticut Supreme Court held that the pattern of financing public schools was unconstitutional because it violated the "educational mandate" and "equal protection" clauses of the state constitution (Horton vs. Meskill: 172 Conn. 615 (1977), 187 Conn. 187 (1982), & 195 Conn. 24 (1985)). The court ordered the legislature to remedy this unconstitutional inequity. After considerable study, a blue-ribbon panel of legislators, social scientists, and policy analysts proposed to the court a solution involving a Guaranteed Tax Base Formula for state funding of elementary and secondary education.

Major details of the formula are contained below in the model-description section. Suffice to say here, the guaranteed tax base (GTB) formula provided greater aid to poorer districts and less aid to wealthy districts, in a particular way. In theory, each community would be guaranteed access to the same tax base. State aid would compensate communities that had less than this guaranteed tax base by making up the difference between what their tax rate actually yielded and what it would have yielded if they had the guaranteed level of wealth. Hence, given GTB state aid, communities that made equal tax efforts would in theory have the same total revenues available for public education.

Although the actual formula adopted involved some political compromises, it was widely expected to significantly reduce disparities in per-pupil expenditures around the state. Prior to its adoption, a series of studies demonstrated how, under various assumptions, guaranteed wealth would narrow the gap in per-pupil expenditures between rich and poor districts, without infringing on local control of public schooling. Because the new formula involved a drastic increase in state costs, a plan was worked out with the court whereby the state would move toward full funding over a period of five years. The phase-in actually took seven years, a matter that led both to further litigation and to speculation about the effects of the delay in moving to full funding. When the formula was fully funded in 1985-86, equalization aid in Connecticut had increased by $271 million per year, rising to almost two-and-a-half times the 1978-79 aid level.

However, during this same period, the spending disparities between low- and high-wealth districts did not decrease as predicted. Over the entire seven-year period, and over the entire range of measures used, disparities in per-pupil expenditures remained about the same as they had been at the start, with only minor fluctuations from year to year. Instead, disparities in school tax effort decreased. It appeared that much of the new aid was being used for tax relief. Connecticut's results were quite similar to those of other states that had reformed their school finance plans in the 1970s and '80s. Indeed several states experienced actual increases in expenditure disparities despite the adoption of finance plans designed expressly to reduce them (Goertz & Hickrod 1983).

Something was wrong. Either assumptions underlying the GTB formula itself, or assumptions of how local districts would respond to these new state dollars, or some combination of other factors had to be wrong. The system dynamics models described in the next sections were commissioned to help diagnose why the per-pupil expenditure disparities were not decreasing -- to identify what sets of assumptions, either implicit or explicit, were keeping the state from meeting its goal of minimizing inequities in per-pupil expenditures. The system dynamics models were intended to complement a host of other existing empirical studies of this same problem. The hope was that the simulation technology could provide a new point of view on a very tenacious problem.

In a very general sense, the models were designed to answer both retrospective and prospective questions. Retrospectively, why hadn't the GTB formula worked as predicted? What was wrong about the assumptions that had been made at its adoption? Prospectively, under what conditions of local response to changes in state aid might the formula work? And what adjustments in the formula itself might produce better results?

These general concerns were focused into several more specific questions. First, in a retrospective analysis, the simulated pre-1979 school finance system was moved from 50% to 100% funding of the guaranteed wealth formula. In this analysis a simulation model was used to explore how various assumptions about local response might affect inequities between low- and high-wealth districts. Second, simulations were used to help think through the possible interactions between inflation and the move to full funding. To what degree did the substantial inflation during this period impair the ability of the formula to achieve its desired results?

Finally, simulation was used to investigate the likely impact of moving toward the use of more current data in the aid computation. Under current statute, this year's aid is computed from wealth, expenditure, and tax rate data that are three years old. Static analyses performed by the Department of Education had suggested that the use of delayed data was a major factor in preventing poor districts from improving their programs, by forcing districts to pay for three years of cost increases entirely out of local funds.

The Connecticut Educational Finance Models

Two computer simulation models were developed and used for this study. The first model represents a single local school district or cluster of similar school districts, and contains a simplified state sector that dispenses aid according to the state's guaranteed wealth formula. The second model represents all the school districts in the state, and contains a full state aid sector that accurately monitors financial and demographic data across the state and dispenses aid accordingly. The more comprehensive seven-sector model is formulated by replicating the structure of the single-sector model seven times and parameterizing the seven sectors to represent clusters of Connecticut school districts.

Figure 1: Overview of the seven-cluster educational finance model

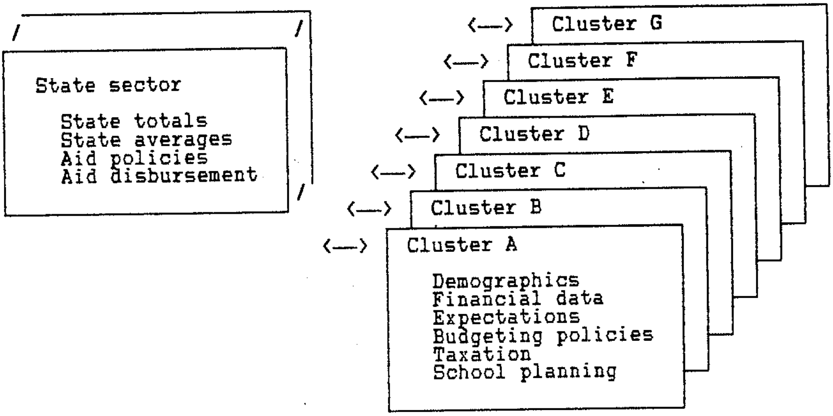

The Seven-Sector Model

Figure 1 shows an overview of the structure of the seven-sector model. Financial and demographic information from the towns is passed to the state and used to compute financial aid according to a guaranteed wealth formula. The formula computes aid per pupil in a given town as

(Guaranteed wealth - Town wealth) * (Tax effort).

The level of guaranteed wealth (per capita) is a policy parameter set annually by the state. A town's wealth per capita is determined from its equalized net grand list property value adjusted for variations in per capita income around the state. A town's tax effort is measured essentially by its school tax rate, again adjusted for per capita income. The resulting formula aid per pupil is then multiplied by the number of students in the school district, producing what is called guaranteed tax base (GTB) aid. Currently in Connecticut, local districts receive at least a minimum of $250 per pupil per year, receiving more if the computed GTB formula aid per pupil is greater than $250.

All the school districts in Connecticut are represented in this seven-sector model, lumped into seven identifiable clusters. The parameter values that distinguish the clusters produce the following characterizations:

Cluster A:	small towns in transition from rural to suburban
Cluster B:	medium-wealth suburbs

Cluster C: very wealthy communities
Cluster D: wealthy suburbs and country estates
Cluster E: poor small towns
Cluster F: large cities
Cluster G: medium size, old industrial cities and towns

Figure 2: Overview of the single sector model, showing local planning for school staffing and taxation and interactions between the local district and state aid.

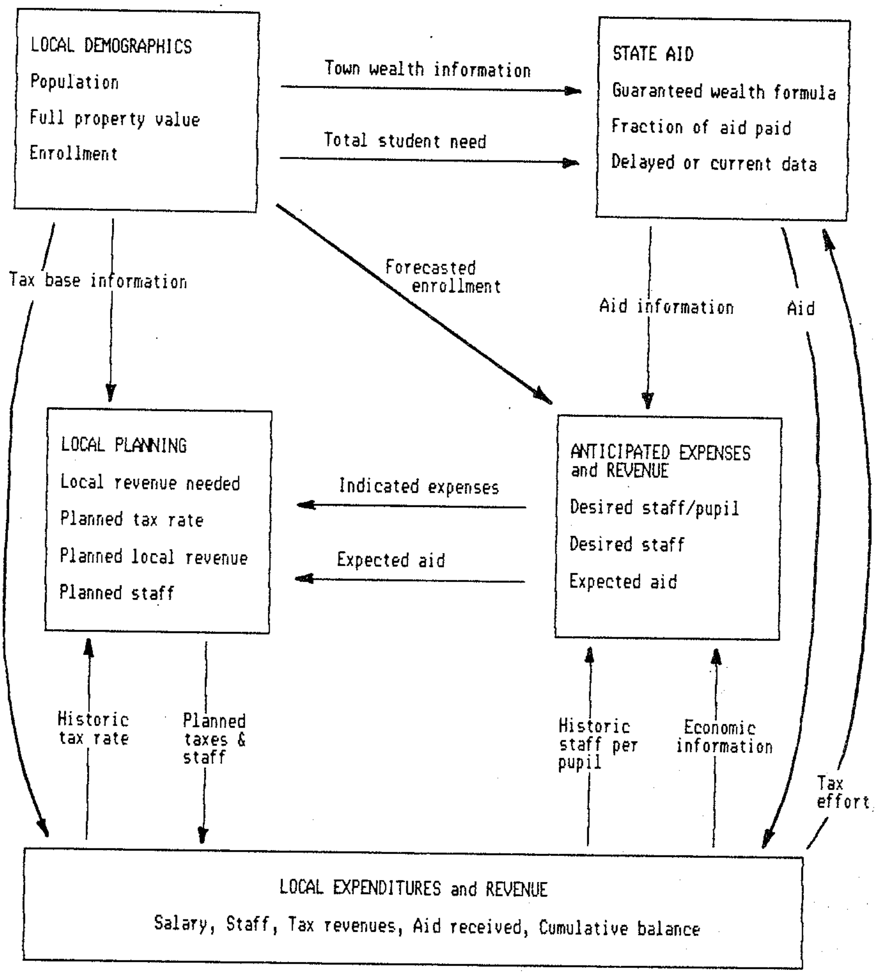

School Finance

The Single-Sector Model

Figure 2 shows an overview of the structure the single sector model. It also shows, in more detail than Figure 1, how a single cluster in the seven-sector model interacts with the state machinery for financial aid to local school districts.

The budgeting policy of the local district is captured in the variables for planned staff and planned tax rate. In essence, the local district tries to maintain its traditional staff-per-pupil ratio without veering too far from its traditional tax rate. The district projects next year's enrollment (average daily membership) and computes a level of staff necessary to maintain its traditional staff-per-pupil ratio. Then taking into account the aid it anticipates from the state GTB formula, and any other revenue, the planners in the district compute the tax rate that would be necessary to achieve the desired staffing level. If that tax rate differs too much from the district's recent historical tax rate, planners and voters are assumed to adjust the budget downward (or upward) to bring the planned tax rate more in line with past rates. School staffing is adjusted accordingly, and the planned staff-per-pupil ratio falls (or rises) relative to its desired value.

Model behavior

The models runs displayed below show the behavior of the single-sector and multi-cluster models in an inflationary scenario over twenty years. The inflationary scenario is characterized by the following exogenous growth rates:

Population	1% per year
School enrollment	1% per year
Guaranteed wealth level	6% per year
State per-capita incomes	8% per year
Local per-capita incomes	8% per year
School salaries	8% per year
Property values	6% per year
Other revenue	4% per year

With school salaries growing faster than town wealth and the state's level of guaranteed wealth, this inflationary scenario strains local district budgeting. The model runs show how the stresses play out over time.

Behavior of the Single-Sector Model

Two runs of the single-sector model are shown. For each of these runs the single-sector is parameterized to represent a relatively poor, rural school district.

Move to full funding of GTB aid over five years

Figure 3 shows the behavior of important variables in a local district as the state moves from 50% to 100% funding of GTB formula aid. The move takes place from year 5 to year 10 of the simulation. The increasing aid relieves some of the financial stress on the local district, lessening its need to raise taxes to meet rising school expenditures. The educational tax rate (Figure 3B) actually declines slightly from year 5 to 10. However, staff-per-pupil stays essentially constant throughout the run (Figure 3A).

Moving to full-funding of guaranteed tax base formula aid does not stimulate this local school district to raise its staffing standard. The goal of the GTB formula -- narrowing the differences between rich and poor districts in staffing and expenditures per pupil -- would apparently not be achieved.

Switch from 3- to 1-year lag in aid

Assuming full funding of formula aid, the simulation shown in Figure 4 explores the implications of shortening the lag in aid from three years to one year. It was conjectured that a shorter lag would help achieve the goals of GTB aid. The simulation shows that the change makes only a transient difference. The move to a one-year lag provides a one-time burst in aid (the formula aid is suddenly greater because it is based on more recent inflating quantities). In this relatively poor local district, however, the change is used mainly to obtain tax relief (see Figure 4B). Staff-per-pupil (Figure 4A) again stays essentially constant.

School Finance

Figure 3: Behavior of a local district as the state moves from 50% to 100% funding of formula aid.

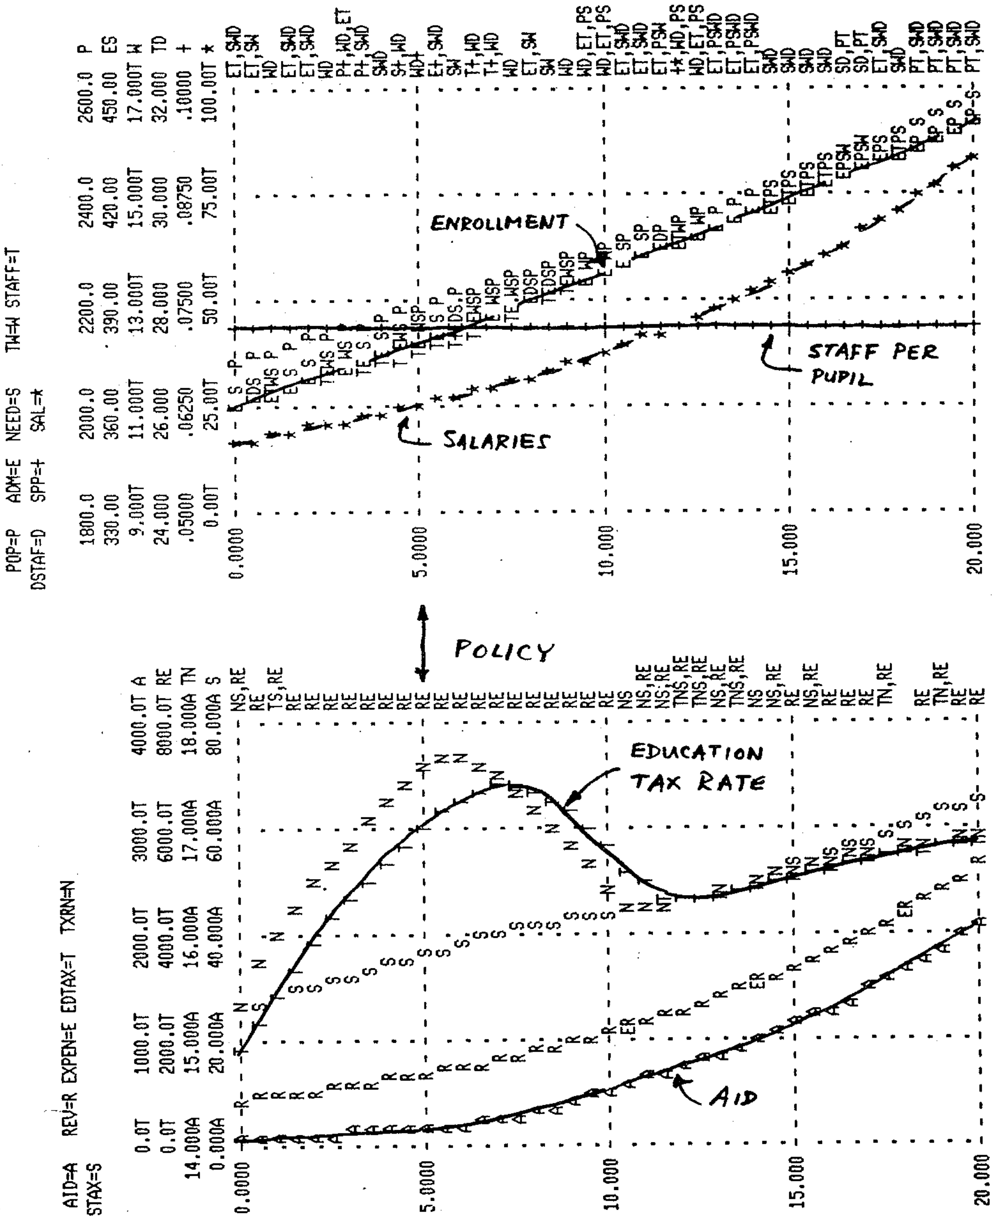

Figure 4: Moving from a three-year lag to a one-year lag in disbursing GTB formula aid.

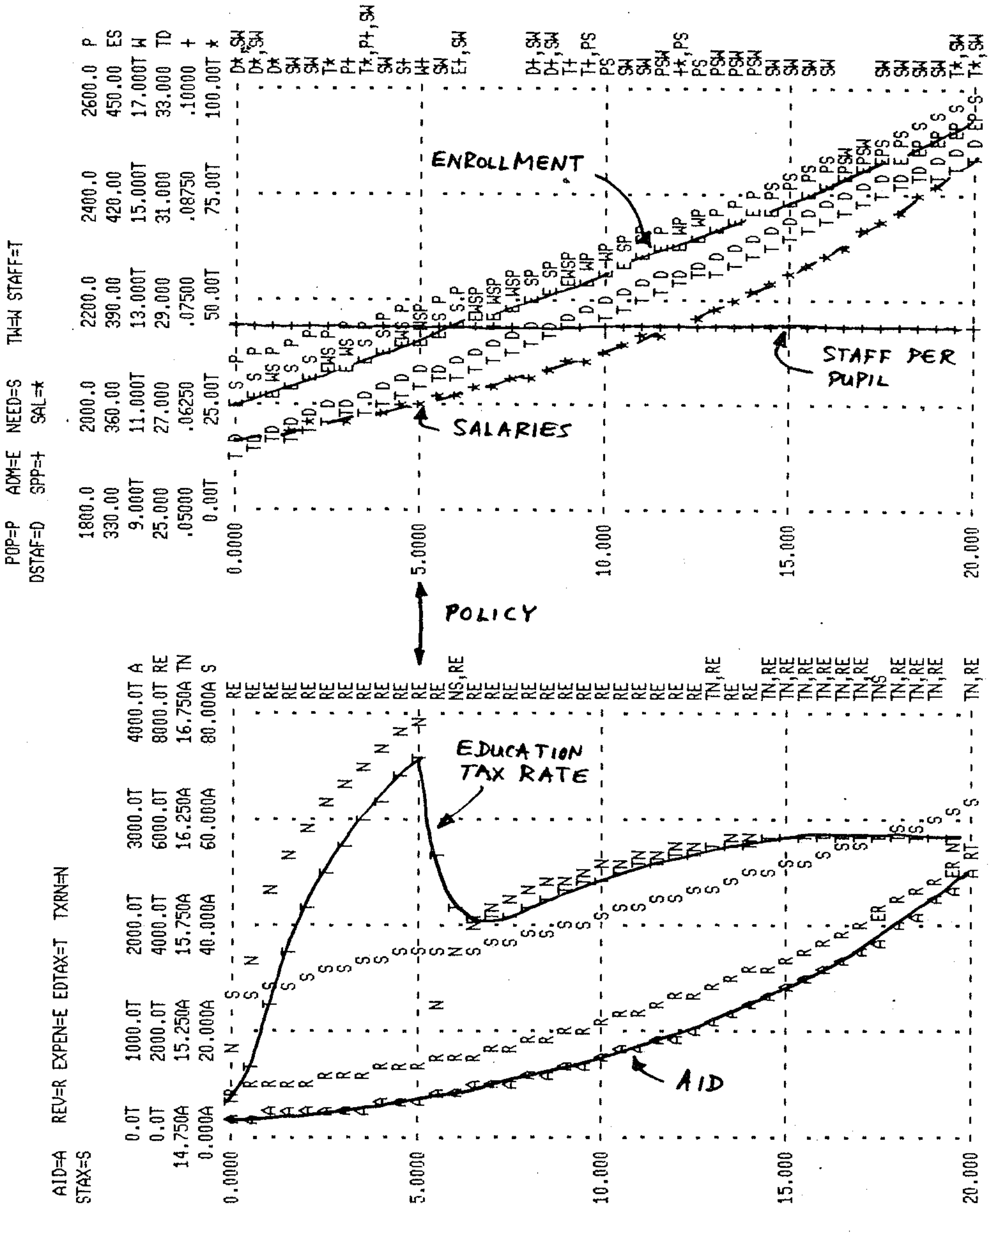

Behavior of the Multi-Sector Model

Because the full seven-sector model proved too large for the DYNAMO installation on the Wheaton computer, a five-sector model containing clusters A through E was tested for this report. The output variables show the behavior of each of the five clusters <u>in relation to</u> state totals and averages. As the tables and graphs show, these relative percentages tend to hide the dynamic patterns of growth and change shown in the output from the single-sector model and operating in both models in each district and in state aid. However, they neatly show the tendencies of the clusters to move together or apart over the course of the twenty-year simulation.

In the base run of the five-cluster model shown in Table 1 and Figure 5, relative school expenditures diverge considerably over the course of the twenty-year simulation. They range from about 70% to 155% of the state average at the beginning of the simulation and end with a range of about 60% to 180% of the state average. Clusters B (medium wealth suburbs) and D (wealthy suburbs and country estates) change places.

Compounding the diverging pattern is the fact that the state average expenditure per pupil rises dramatically in this inflationary run, moving from $3340 per pupil initially to $13700 by the end of the simulation (Table 1). The expenditure gap between Cluster E (poor small towns) and Cluster C (very wealthy communities) actually grows from about $1220 per pupil per year to over $16000! The gap in expenditures per pupil stems from differences in salaries and staff-per-pupil ratios. The expenditure gap grows faster than salary inflation because staff-per-pupil ratios diverge considerably in this simulation. In this scenario, the GTB formula is failing to meet its goal of narrowing the gap in per-pupil expenditures between rich and poor districts.

In contrast, the relative tax rates converge dramatically in the course of the simulation (Figure 5). Because the state average tax rate is rising throughout the run (Table 1), the tax rates in the various the clusters are actually rising, although at markedly different speeds. Cluster C (wealth suburbs and country estates), for example, moves from an initial tax rate of about $7.70 per $1000 to a final rate of about $12.25 per $1000. Cluster E (poor small towns), reluctant to raise taxes because of its relatively poor financial picture, moves from $11.30 per $1000 to $13.90, less of an increase but still the highest tax rate among the five clusters.

Table 1: Selected quantities in the five-cluster model in a basic inflationary scenario assuming full funding of GTB aid

State averages and totals -

	Staff per Pupil	Expenses per Pupil	Aid per Pupil	Guaranteed Wealth	Education Tax Rate	School Tax Rate (adjusted)	
TIME	SSPP	SEPP	SAPP	SGW	SEDTAX	SSTAX	
E 00	E-03	E 03	E 00	E 03	E-03	E-03	
0.000	72.823	3.340	302.88	35.00	9.058	14.341	SSPP drops.
5.000	69.350	4.757	384.96	47.22	10.089	15.973	EPP up 400%
10.000	65.625	6.735	516.14	63.70	10.967	17.364	APP up 300%
15.000	62.435	9.585	704.74	85.94	11.902	18.844	Tax effort
20.000	59.717	13.714	974.76	115.94	12.914	20.447	rises.

	Average Daily Membership	Total Aid	Equalized Net Grand List total	Equalized Grand List per person	Highest Capita Income	Average School Salary	
TIME	SADM	SAID	SENGL	SENGLP	SHPCI	SSAL	
E 00	E 03	E 06	E 09	E 03	E 03	E 03	
0.000	300.88	91.13	82.74	44.35	21.402	22.93	Total aid
5.000	316.30	121.76	117.33	59.83	26.134	34.30	grows 400%
10.000	332.51	171.63	166.37	80.70	31.912	51.31	in twenty
15.000	349.56	246.35	235.90	108.85	38.968	76.76	years.
20.000	367.47	358.20	334.51	146.82	47.584	114.82	

Aid fraction - the percentage of the cluster's educational expenses met by state aid

	A cluster	B cluster	C cluster	D cluster	E cluster	
TIME	AAIDF	BAIDF	CAIDF	DAIDF	EAIDF	
0.000	11.707	5.7676	4.4344	6.5054	22.017	The richer
5.000	11.244	5.3387	3.1896	4.6860	20.969	districts,
10.000	11.569	5.5102	2.1691	3.2138	21.239	C and D,
15.000	11.888	5.7109	1.4617	2.1859	21.532	get minimum
20.000	12.163	5.8914	0.9827	1.4842	21.789	aid.

Figure 5: Behavior of the five-cluster model under a basic inflationary scenario, assuming full funding of GTB aid.

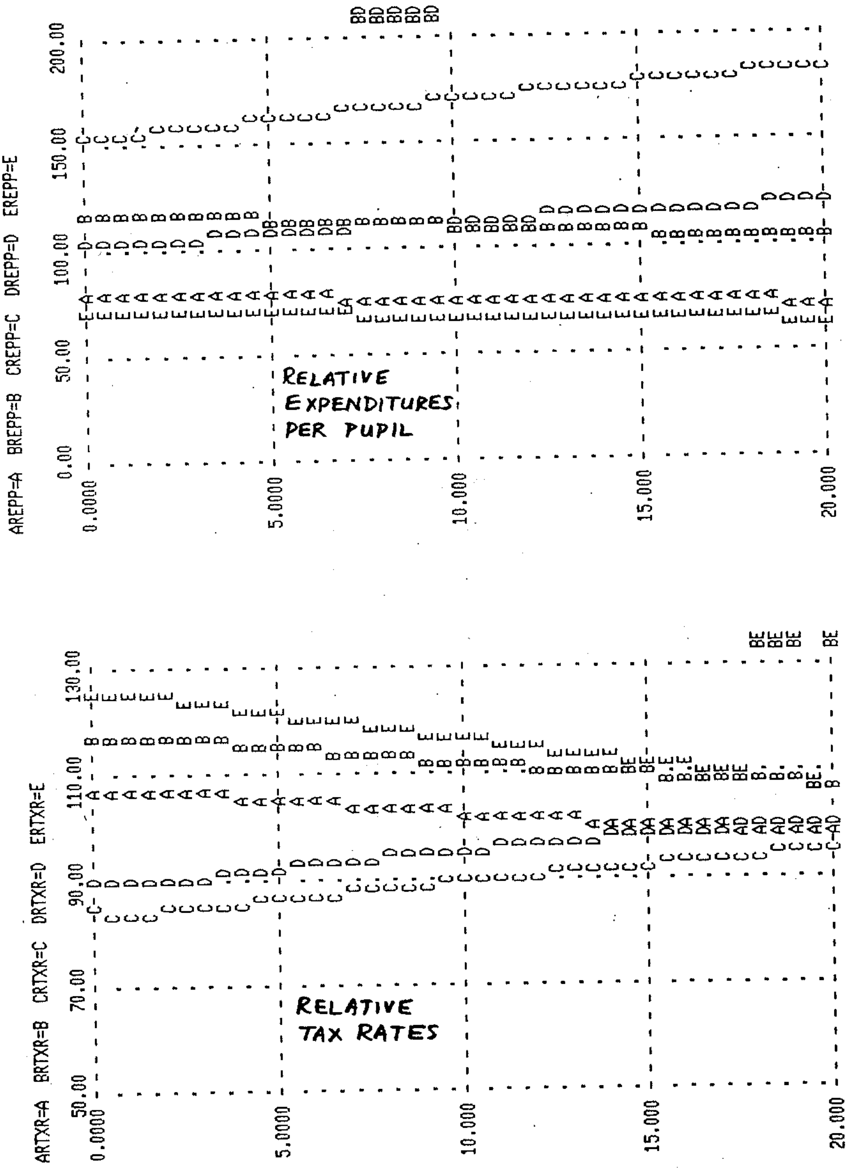

Implications for School Finance Policy

Full-funding of GTB formula aid

These simulations imply that local school districts in Connecticut have not been held back by the state's decision to move gradually to full funding of GTB aid. During the phase-in period, it was conjectured that the reason expenditures-per-pupil were not converging in rich and poor districts was that the program was not at full funding. Once full funding was achieved, according to this reasoning, the goal of equalizing per-pupil expenditures would be more rapidly approached. Unfortunately, under the school and tax planning policies assumed in these models, this reasoning is shown here to be faulty. Even with full funding of GTB aid, staff- and expenditures-per-pupil in rich and poor districts continue to diverge.

Rather, it is tax rates in rich and poor districts that move closer together as a result of GTB aid. The reason is relatively straightforward: Striving to achieve staffing goals in an inflationary scenario, rich and poor districts differ in their willingness to raise taxes. Rich districts, which start with lower taxes rates, have more leeway to raise tax rates; poor districts, which already have relatively high tax rates, are less willing to raise taxes. Thus over a time of budgetary strain, in which expenditures rise faster than property values, tax rates move closer together. State aid is not useless here, however. It is being used by the less wealthy districts to help hold down tax increases.

One-year versus three-year lag

Simulations of both the single-sector and multi-sector models suggest that the three-year lag in aid does not significantly stand in the way of progress in expenditures per pupil. Understandably, switching from a three-year lag to a one-year lag in an inflationary scenario has a short-run beneficial effect -- the aid makes a two-year jump in one year. However, if the local district is already making a high tax effort, it will likely use the short-term increase in aid to reduce taxes. Even if it does use the increase to increase staff or salaries, the increase will occur in just a year or two, and then the district reverts to its previous pattern of performance.

More General Policy Implications for State Aid to Education

The models, and the thinking they reflect and stimulate, strongly suggest several principles or working hypotheses of state aid:

> Untargeted financial aid that can be anticipated and planned for is likely to be simply incorporated into the

local budgeting process as general revenue.

Significant <u>changes</u> in aid have effects in raising (or lowering) expenditures per pupil. Sustained aid is not as effective as a sudden unexpected increase in aid. The change must exceed what local planners and voters perceive as a reasonable drop in expenses and tax rates, so that the money is actually used to increase expenditures.

The observed effects of a change in the aid lag from three to one year are apparently a corollary of sorts to the first working hypotheses. The change appears to be significant only in creating a short-term increase in aid, which can raise per-pupil expenditures. The switch is most helpful in districts biased toward spending extra revenues rather than lowering existing taxes.

If local districts set their goals based upon their recent past history and operate with local budgeting discretion, then raising standards in the long run requires a sustained period of increasing aid, preferably targeted to staff-per-pupil or salary standards the aid is intended to improve.

Aid can be addictive. The local district will come to depend upon its sustained presence -- a case of shifting the burden (school finance) to the intervener (the state). Yet if goals are set on recent past history, the removal of aid may not result in the collapse of a school system, but rather a scrambling for local ways to hold onto the standards the district has become used to. Quality schooling is addictive, too.

If it could be done, state aid should be given or increased in <u>bursts</u> lasting as long as necessary to raise traditional standards. Then aid could be lowered in that district. Performance will fall somewhat, but not to the level from the which district started before the burst in aid.

Methodological Implications of the Case

While the simulation exercise arrived at several interesting substantive conclusions concerning school finance reform in Connecticut, it also raises a series a methodological questions. Strictly speaking, the modeling effort could be said to violate many of the norms of positive social science. The models and their conclusions were not formally grounded in empirical observation. They were contingent in nature, providing knowledge of the form "If you assume A, B, and C concerning local district budgeting behavior, then the long range implications for GTB aid and equity will be thus and so." Moreover, the validity of A, B, or C, the central premises of the logical syllogism, were never empirically tested or even seriously questioned during the modeling exercise itself.

On the other hand, policy makers who worked with the models report that analyses derived from them served to sharpen their insights and to help frame concrete policy options. As a result of the modeling effort and their continuing study, they have increasingly turned toward policy proposals involving specific standards of adequacy and the targeting of aid directly toward the achievement of state policy goals. They have proposed new aid formulas for teacher salaries, for example, which will make aid unavailable unless used for its state-mandated purposes. These policy makers see the implementation of the new aid formulas as a real-world test of policy implications of this modeling study. Nonetheless, they perceive that the highly contingent nature of the simulation results would pose some difficulties if they were to try to use the models to communicate model-based insights to people not directly involved in the modeling effort.

Quite obviously, this study, and other modeling efforts similar to it, are subject to methodological critiques asserting that such models reduce either to tautological thinking or unsubstantiated speculation. According to such critiques, the models either reflect merely what is already known or suspected to be true, or they push forward a series of contingent assumptions unchecked by empirical verification. Either way, such exercises could be accused of being incapable of generating any new knowledge. Yet we believe that there is an important place in the social and policy sciences for the sort of analyses performed in this study. Hence, we are led to ask,

> What are the circumstances under which policy analytic exercises designed to <u>draw inferences from assumptions</u> are as useful as exercises designed to <u>test or confirm assumptions</u>?

The Connecticut case suggests that a policy problem may be ripe for a contingent, inference-generating analysis (as contrasted with a hypothesis-confirming analysis) when the following conditions are met:

<u>An unresolved puzzle:</u> The most important aspect of the Connecticut case is the fact that inequities in per-pupil expenditures have continued even after substantial and sustained policy interventions designed to reduce them. That is, inequity is spending is a persistent policy problem that is well studied, much thought about, and apparently resistant to corrections that are substantial, intuitively appealing, and grounded in empirical study.

Existing pool of hypotheses: Another important aspect of the Connecticut case that made the simulation study possible and helpful was that many formal propositions about what policies should work had been advanced. Years of research and attention had advanced a rich set of hypotheses concerning the nature of the problem and the types of likely solutions.

Existing intuitive experts as clients: Key policy makers within the state had been working with the school finance problem for years. As the primary authors of many detailed empirical studies and specific legislative proposals, these intuitive experts carried a wealth of both general knowledge and specific anecdotal knowledge related to the system under study. Their knowledge, embodying the pool of existing hypotheses, provided the "data" from which the model was built.

Existing empirical studies: Not only was knowledge embodied in expert intuition, but in the Connecticut case numerous empirical studies existed, which had been designed to confirm several competing hypotheses of the intuitive experts. Without such an extensive background of empirical studies, it seems likely that the modeling effort would have been much more difficult and much less believable.

Interactions between multiple hypotheses: In the Connecticut case, multiple sets of hypotheses existed concerning many different parts of the system under study. Assumptions had been made about how the GTB formula ought to equalize aid; how local districts would react to changes in state aid; how the existence of delays or other defects in the "pure" formula would impact local response; and so on. Quite obviously, these hypotheses or assumptions are not all independent. They interlock to form a network of propositions that, taken as a whole, provides a complex, interconnected explanation of how the school finance system allegedly operates. The purpose of the modeling exercise was to work out the logical implications of this complex set of propositions.

Hypotheses involving both causal structure and behavior over time: In the Connecticut case, some of the existing propositions about the system involved statements about time-series behavior, such as "implementing the guaranteed wealth formula will, over time, create a convergence in per-pupil expenditures across the state." Other statements had more of the character of chains of causal influence, such as "the greater the tax effort, the greater will be the state aid (after a delay)," and "the greater the state aid to education, the greater will be local expenditures for education." In a situation involving many interconnected propositions about a persistent problem, it seems quite likely that *a priori* statements about how a system will behave over time could be in conflict with statements about causality underlying the system. A complete statement of the propositions about the system's network of causal influences could imply behavior over time that contradicts behavioral propositions arrived at in other ways. By linking underlying causal structure and time-series behavior, a simulation model can be used to identify and sort out these inconsistencies.

These six characteristics appear to be prerequisites for success in the sort of policy-analytic modeling study represented by the Connecticut school finance case. It is interesting to note that these characteristics also describe the prevailing situation in cognitive science, the emerging field of study combining psychology, linguistics, neurophysiology, logic, and computer simulation (Miller, Galanter, & Pribram 1960; Feigenbaum & Feldman 1963; Newell & Simon 1972). Thus we find precedents and some philosophical justification for our approach in the literature associated with artificial intelligence.

Computer simulation comes to the aid of cognitive theorists because it provides a way of testing whether hypotheses about unobservable cognitive machinery produce, or are consistent with, observations of cognitive behavior. Moreover, in extremely complex situations, such as cognition (or school aid policy analysis), simulation is thought by its practitioners to be more reliable than thought (Miller, Galanter & Pribram 1960, 182-183; Simon 1981,26); Forrester 1961,14). There is, of course, some disagreement about the claim (e.g., Weizenbaum 1976).

The approach in the Connecticut case was "artificial," in the same sense of that word used by Simon (1981). It was artificial not just because it involved a computer model, but more so because it was hypothetical. We investigated the dynamic implications of a set of reasonable hypotheses about local school planning and taxation, without direct detailed knowledge of time-series representations of that budgeting behavior.

Such a simulation approach can be empirical and scientific, in two senses. It is obviously empirical in the classical sense when simulation results are compared with observations of the real system. But running a computer simulation by itself is also an empirical enterprise. We make hypotheses about model behavior expected, we perform simulations, we compare the observed results with our expectations, and we use such empirical experiments with a simulated world to gain insight.

From different directions, both Forrester (1961,17-18) and Simon (1981,17-26) argue that such an artificial simulation approach gives us a way of scientifically studying puzzling human phenomena that can not be experimented with directly. These observations, prompted by our reflections on the Connecticut case, suggest that system dynamicists may find some fruitful associations and philosophical foundations among the literature in cognitive science and artificial intelligence.

School Finance

Summary

Constructing formal mathematical models of policy systems involves a blend of three distinct types of activities -- creating propositions about how the policy system actually functions, confirming or casting doubt on these propositions by formal statistical tests using empirical observations, and drawing inferences from interconnected sets of propositions. The Connecticut case discussed in this paper is a nearly pure case of the third type of analysis.

The Connecticut study yields several interesting substantive conclusions concerning the financing of public school education using a guaranteed wealth formula designed to equalize per-pupil spending in rich and poor districts. However, upon reflection the case also yields some methodological insights about the nature and value of what might be called syllogistic policy modeling.

This policy-analytic exercise relied on a formal inference-generating technology (a system dynamics model), which may be contrasted to a hypothesis-confirming technology such as formal significance or goodness-of-fit tests. The two technologies differ substantially in the methods used to calibrate the formal models they employ, the standards and approaches used to judge the quality of the models, and the types of conclusions they can substantiate.

Bibliography

Cyert, R.M. and J.G. March. 1963. A Behavioral Theory of the Firm. Englewood Cliffs, NJ: Prentice-Hall.

Feigenbaum, E. and J. Feldman. 1963. Computers and Thought. New York: McGraw-Hill.

Forrester, J.W. 1961. Industrial Dynamics. Cambridge, MA: The MIT Press.

Forrester, J.W. and P.M. Senge. 1980. Tests for Building Confidence in System Dynamics Models. System Dynamics. TIMS Studies in the Management Sciences, 14:209-228.

Goertz, M. and G.A. Hickrod. 1983. Evaluating the School Finance Reforms of the 1970s and Early 1980s. Journal of Educational Finance, 8,4 (spring).

Miller, G.A., E. Galanter and K.H. Pribram. 1960. Plans and the Structure of Behavior. New York: Henry Holt.

Newell, A. and H.A. Simon. 1972. Human Problem Solving. Englewood Cliffs, NJ: Prentice-Hall.

School Finance

Richardson, G.P. and D.F. Andersen. 1986. Connecticut Educational Finance Modeling Project, Final Report. Wheaton College, Norton, MA 02766.

Simon, H.A. 1981. The Sciences of the Artificial. Cambridge, MA: The MIT Press.

Weizenbaum, J. 1976. Computer Power and Human Reason. San Francisco: W.H. Freeman.