DOCUMENT RESUME

ED 239 378 EA 016 337

AUTHOR Bernhardt, Victoria L.; And Others
TITLE Seattle's Small-Area Approach to Forecasting Enrollments at the School Level.
PUB DATE Apr 83
NOTE 29p.; Paper presented at the Annual Meeting of the American Educational Research Association (Montreal, Quebec, Canada, April 11-15, 1983).
PUB TYPE Reports - Research/Technical (143) -- Speeches/Conference Papers (150)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS Attendance; *Computer Oriented Programs; Declining Enrollment; *Educational Facilities Planning; Elementary Secondary Education; *Enrollment Projections; Linear Programing; Mathematical Models; Regression (Statistics); *Research Methodology; Research Problems; Residential Patterns; School Closing; *School Demography; School Desegregation; School Districts; School Holding Power
IDENTIFIERS Markov Processes; *Seattle Public Schools WA; *Small Planning Units

ABSTRACT

The issues of achieving school desegregation and providing programs for students with special needs, while at the same time facing enrollment declines, have forced school administrators to improve techniques for projecting demand. This paper describes a new procedure that combines forecasting of enrollments and management of facilities and that has been successfully implemented for the Seattle, Washington, Public Schools. The Seattle system prepares its forecasts for a relatively small local unit called the "small planning unit" containing from 50-100 students each. Five-year projections are prepared for each small area, and are then aggregated to prepare forecasts for large geographical areas and for attending schools. Advantages to this approach claimed for planners are that (1) it allows for easier simulation of policy alternatives, (2) small units provide a welcome sensitivity to neighborhood activity, and (3) the use of small-area forecasts allows continuous refinement of the models in response to actual results. Some results are reported that support the planning technique. The mathematical procedures for the forecasting model used with small planning units are in the appendix. (MLF)

SEATTLE'S SMALL-AREA APPROACH TO FORECASTING ENROLLMENTS AT THE SCHOOL LEVEL

by

Victoria L. Bernhardt
Thomas W. Pullum
Stephen N. Graham

April, 1983

Center for Studies in Demography and Ecology
The University of Washington
Seattle, Washington 98195

The authors would like to acknowledge the support of the school board and staff of the Seattle Public Schools in the design and implementation of the system described in this paper. We would particularly like to thank Dave Colwell, Susan Franklin and Larry Collister on the planning staff.

Jerald Herting, a graduate student at the center, provided invaluable support in the preparation of this paper.

SEATTLE'S SMALL-AREA APPROACH TO FORECASTING
ENROLLMENTS AT THE SCHOOL LEVEL

Introduction

This paper reports on the development of a comprehensive system for facilities management based on a small-area enrollment forecasting system. Demographic techniques of age, period and cohort analysis are used to produce five year forecasts of enrollment by small geographical areas. The area forecasts are then aggregated to produce school level and larger area forecasts. This approach potentially has much value for planners in preparing simulated outcomes for facilities management purposes.

The public schools in this country have a long tradition of being "scientific" both in their administration and in their own evaluation of how they provide educational service (see Katz 1968 or Callahan 1962 for a discussion and development). Recent political events and demographic changes have forced the schools to take a more serious look at their past techniques for projecting demand, for allocating resources within the school system, and for providing equal access to all resources for all students. Most obvious are the desegregation issues that require a large amount of planning in order to provide information for decision-makers as to how a given district can achieve desegregation.

At the same time urban school districts have faced a declining enrollment. This has resulted in necessary cutbacks in capital and a reallocation of resources within the system to provide services within budget constraints. Issues concerning which schools to cut which grade levels are to remain in a given area and which will not be needed necessitate a large degree of planning and, specifically, accurate planning at smaller units with concomitant capability of forecasting over a range of several years.

A review of the literature concerning projections of school enrollment makes it clear that the emphasis has typically been on fairly large units, at the school level or larger. Forecasts produced at this level of aggregation provide little flexibility for simulation of alternative strategies. The four main techniques in current use for forecasting and planning will be discussed briefly.

Most common at the school district level or the school level are forecasts based on the grade progression ratio or GPR, also known as the cohort survival ratio or the holding power. Requiring two years of observed data, this ratio is calculated by dividing the enrollment in grade g at time t by the last year's enrollment in grade g-1 at time t-1. This ratio is then multiplied by the present enrollment in grade g-1 at time t to forecast to grade g at time t+1. The basic assumption here is that a cohort's net retention in the district or school from grade g-1 to g will be exactly the same as for the next older cohort during the most recent period of observation. In its simplest form, this procedure does not use more than two years of data, even if more is available, and does not forecast future trends in the GPR.

The planning literature also includes discussions of regression analysis and its applicability to forecasting school enrollments. The regression method is derived from the economics literature, and makes use of time series analysis. (for example see Rives, 1977). The forecasts can be stated in the form of interval estimates by shifting the estimated enrollment slope higher or lower based on standard errors and an arbitrary degree of confidence. The unit of analysis is ordinarily at a rather large level, although allocation is sometimes used to estimate smaller units. The technique as presented by Rives demands fairly long term and stable trends. Such an analysis would be quite sensitive to fluctuations in the history of enrollments.

A third method in the literature is built around linear and "goal" programming where certain constraints are placed on a system of equations and through iterative techniques an optimal solution can be attained that satisfies pre-set constraints or goals. In general this literature, as applied to school enrollment issues, has focused on the allocation of resources to best meet demand (e.g. Knutson et al., 1980). It is also used to formulate maximum integration or student allocation while minimizing transportation costs and distance. Pullum and Graham (1982) used linear programming to estimate the degree of ethnic imbalance that would remain under alternative desegregation plans. This method is useful for facilities utilization planning but nothing directly related to forecasting enrollments has appeared in the literature reviewed by the present authors.

Another approach in recent years has been the application of Markov chain models to the forecasting of enrollments (or of other measures of future demand). Johnstone and Philip (1973) attempted to see how well such models perform. By setting up transition matrices of holding powers or survival ratios and then passing actual enrollments through the matrix they developed a long range forecast of district enrollments and compared these to actual enrollments. Their results and comments suggested that the method was too inflexible for widespread use. Other researchers have suggested means of altering the transition matrix used in the Markov process to enable more flexible and responsive forecasts (e.g. Britney, 1975, or Harden and Tcheng. 1971).

A major problem with Markov chains is the inability to simulate what might occur given alternative policies. No parameters are available that would allow changes (e.g. in magnitude or inclusion versus deletion) in the system to be tested and examined for their effects. In other words, its capacity for simulation is limited.

In a review of the literature by McNamara (1973) it becomes obvious that the main emphasis of planning has been on how to best allocate resources via the linear or goal programming mode, rather than to focus on forecasting demand. Even in the years since that review, little attention has been given to this problem; the citations above are exceptions.

One of the reasons for the lack of literature on demand forecasting is the fact that most forecasts are produced at such a large level of aggregation that they are of little use in answering the questions which plague most planners. Forecasts produced at the district level provide a general guideline to the level of citywide demand that is to be expected, but they cannot help to answer specific questions about which schools should be closed, how attendance boundaries should be redrawn, or how racial balance is to be maintained at the facility level.

In most school districts the lowest level used is usually the school itself, which typically contains at least two hundred children. This is done because it is felt that smaller aggregates will result in unstable or unreliable models and poor forecasts. However, even forecasts produced at the school level cannot be used for simulation very effectively. They generally assume a fixed set of attendance boundaries for the school and cannot easily be used to explore the effects of boundary changes, school closures, grade configuration changes, etc.

This paper will describe a new procedure which combines forecasting of enrollments and management of facilities into a single comprehensive system. The procedure has been successfully implemented for the Seattle (Washington) Public Schools, a medium-sized district with approximately 50,000 students. The paper will describe the application of some new small-area forecasting techniques to Seattle's particular circumstances.

Interest in population estimates for small areas is a relatively new topic in demography. A recent monograph (Lee and Goldsmith, 1982) explores a variety of methods for estimating the size and composition of small-area populations between censuses. Little has been done with the forecasting of populations at the small area level.

The Seattle system prepares its forecasts for a relatively small local unit called the "small planning unit". Forecasts for larger units are produced by aggregating the small planning unit forecasts. Thus the Seattle system produces its forecasts by predicting at the small unit level and aggregating up rather than predicting at the large unit level and allocating down.

There are some unique advantages to this approach for planners. First, it allows for easier simulation of policy alternatives. Small units can be quickly and easily aggregated in different ways to explore the effects of various management strategies. It is possible to model or to simulate different attendance boundaries for local schools, as well as different bussing and grade configuration policies at the school level.

Second, small units provide a welcome sensitivity to neighborhood activity. Relatively small local changes in enrollments can be detected and brought to the attention of planners. Even as few as ten children moving into a small area will cause a change that will register. This allows early detection of trends that would be lost at larger levels of aggregation. Models can reflect trends that are local in character, specific to a neighborhood.

Finally, the use of small-area forecasts allows continuous refinement of the models in response to actual results. If, for example, a school level forecast is seriously in error, the planners can break the forecast down into its component small units and look for errors in each unit. If a particular unit is at fault, the forecasting algorithm can be refined for that unit when the cause is located. Investigating a small unit is much easier than attempting to account for discrepancies in a much larger aggregate.

The disadvantages of the approach basically revolve around the sensitivity of small areas to noise or random movement of students. Particularly in areas of high turnover, the random component of year to year student movement can overwhelm the systematic trends which a modelling algorithm is looking for, with the result that unstable or inaccurate forecasts will be produced.

The problem discussed above is not so serious as it first appears, since there is reason to believe that errors made in one small area will be counteracted by equal and opposite errors in other areas. Many studies have shown that the vast majority of migration is short distance, much of it within a radius of a few miles. When migration occurs in the context of a large urban school district, it is most often the case that students who move out of one small area move into another, so that a positive error in one area will be counteracted by a negative error elsewhere. This helps to stabilize the aggregate totals computed by summing the small-area forecasts. There is nevertheless a tendency for errors within a small area to be relatively large in comparison to the area size. The Seattle system employs a variety of techniques to control or suppress the noise element within the small planning unit.

The following sections describe some of the concerns that led up to the implementation of the Seattle system and describe it in more detail. Some results are reported which support the notion that small area forecasting can be an effective tool for planners.

Seattle

The Seattle Public School district exemplifies many of the trends discussed earlier. Enrollment decline has been spectacular, from a high of 99,000 in 1963 to a current figure near 45,000. As a result of the decline, several high schools and middle schools have been closed along with nearly twenty elementary schools. Many of the facilities that remain open are at marginally low enrollment levels, and further closures remain likely if enrollment continues to decline.

The enrollment decline has not been uniform, however. Dramatic declines occuring in the northern and western parts of the city have been matched by steady or even increasing enrollment among the heavily minority population in the southern and central parts. The vast majority of the closures have occured in the north and west, while there is barely sufficient capacity to serve the population in the center and south. As a result of this demographic shift, the percentage minority in the school population has increased from 15% in 1963 to 49% today.

To aid in desegregation of the city schools, a mandatory bussing program was instituted in 1978, based on a feeder pattern concept.

Students are first assigned to geographic attendance area called a "neighborhood K", and all the students from the neighborhood K are then assigned to the same school at each grade level. This school assignment pattern, called a "feeder pattern," may require students from a particular neighborhood K to be bussed at particular grade levels. When bussing occurs, all students at the designated grade level are bussed together. The number of students bussed has gradually risen since 1978 until today about 7,000 students are being mandatorily bussed for desegregation purposes.

During the same period, special programs aimed at students with particular educational needs or problems have proliferated. Programs for gifted students, bilingual students and students with physical, learning and/or behavioral problems have increased in both size and scope. Magnet programs and voluntary racial transfer programs are in effect to supplement the mandatory bussing in the desegregation program. All of the programs mentioned above have the effect of assigning children outside of their regular feeder pattern school. The number of these "out of feeder" children has increased steadily until today over 18,000 children are not attending their regular feeder pattern school.

All of these changes in attendance patterns make it more and more difficult to plan effectively. For budget and staff planning, the administration must prepare both short range (1 year) and medium range (5 years) projections of attendance within geographic areas of the city and at the school and program level. Future utilization of facilities must be estimated so that the schools which remain open will be adequate for incoming student loads. The development of a long range facilities plan has been a high priority for Seattle but the ability to evaluate the alternatives available to the school board has not been adequate.

The key to developing a long range facilities plan is an effective enrollment projection system. This system must be able to project enrollment within the various geographical areas and to simulate the effects of various policy decisions about facilities, grade configurations, feeder patterns and optional special programs. The development of such an enrollment projection system was a high priority of the planners at Seattle Public Schools.

The Center for Studies in Demography and Ecology at the University of Washington has had a long standing involvement in population projection and became interested in the schools' problem in 1980. After some negotiation a contract was worked out to jointly develop an enrollment projection system which would meet the needs discussed above. The system was built gradually over a two year period and has only recently produced its first set of forecasts. Evaluation of the system is continuing, but the basic approach appears to be sound, and preliminary results are promising.

The system prepares five year forecasts of enrollment at the school level under a user-selected combination of attendance boundaries, feeder patterns, grade configurations and open/closed facilities. Each of the factors indicated may be varied separately for each of the five years forecast, e.g., the user may select one configuration for the first year of the simulation, another for the second, etc., and all configurations can be simulated in the same run. Simulations are quite efficient. An individual simulation can be prepared in under 5 cp seconds on the schools' computer.

STAGE 1 - Residential Forecasting

Forecasts are produced in two stages. First the system prepares an enrollment projection for the residential population. The city is divided into approximately 600 small planning units (SPU's). These units are geographical areas containing from 50 to 100 students on the average. When such a unit is divided into 13 grade levels and two ethnic groups, there are only 2 to 4 students per grade on the average. In spite of this very small cell size the system is able to make good average five year projections for each SPU by adapting some projection techniques which have hitherto been used for log linear modelling. The projection techniques are discussed in more detail in the appendix.

The SPU's are the building blocks of the simulation system. Simulations are created by aggregating the SPU forecasts in different ways to produce estimates by attendance areas. Attendance area estimates are then put through the simulated feeder patterns to arrive at enrollment projections by school. Boundaries for the SPU's are selected with likely attendance area borders in mind. SPU's must be small enough to permit useful adjustments to attendance areas yet large enough to permit adequate statistical stability in projection.

The initial set of small planning units selected for Seattle is shown in Figure 1. The light lines are the SPU boundaries. Shown in the dark lines are the large planning unit (LPU) boundaries. The large planning units are numbered and correspond to planning areas used by the city. Statistical data collected at the census tract level or at the city planning area level can enter the system via the LPU.

To produce a forecast of the residential population, the system first constructs a history of the enrollment by SPU. For each unit, the system calculates a set of tables giving the enrollment by grade and ethnic group for the most recent eight years of history. The history is for <u>residential attending</u> enrollment, i.e., it encompasses students residing in the SPU who attended the Seattle Public Schools.

ED 239 378 SEATTLE'S SMALL-AREA APPROACH TO FORECASTING ENROLLMENTS AT THE SCHOOL LEVEL.

1 OF 1 BERNHARDT, VICTORIA L.; AND OTHERS

APR 83 29P.

Figure 1. Seattle Small Planning Units showing LPU boundaries

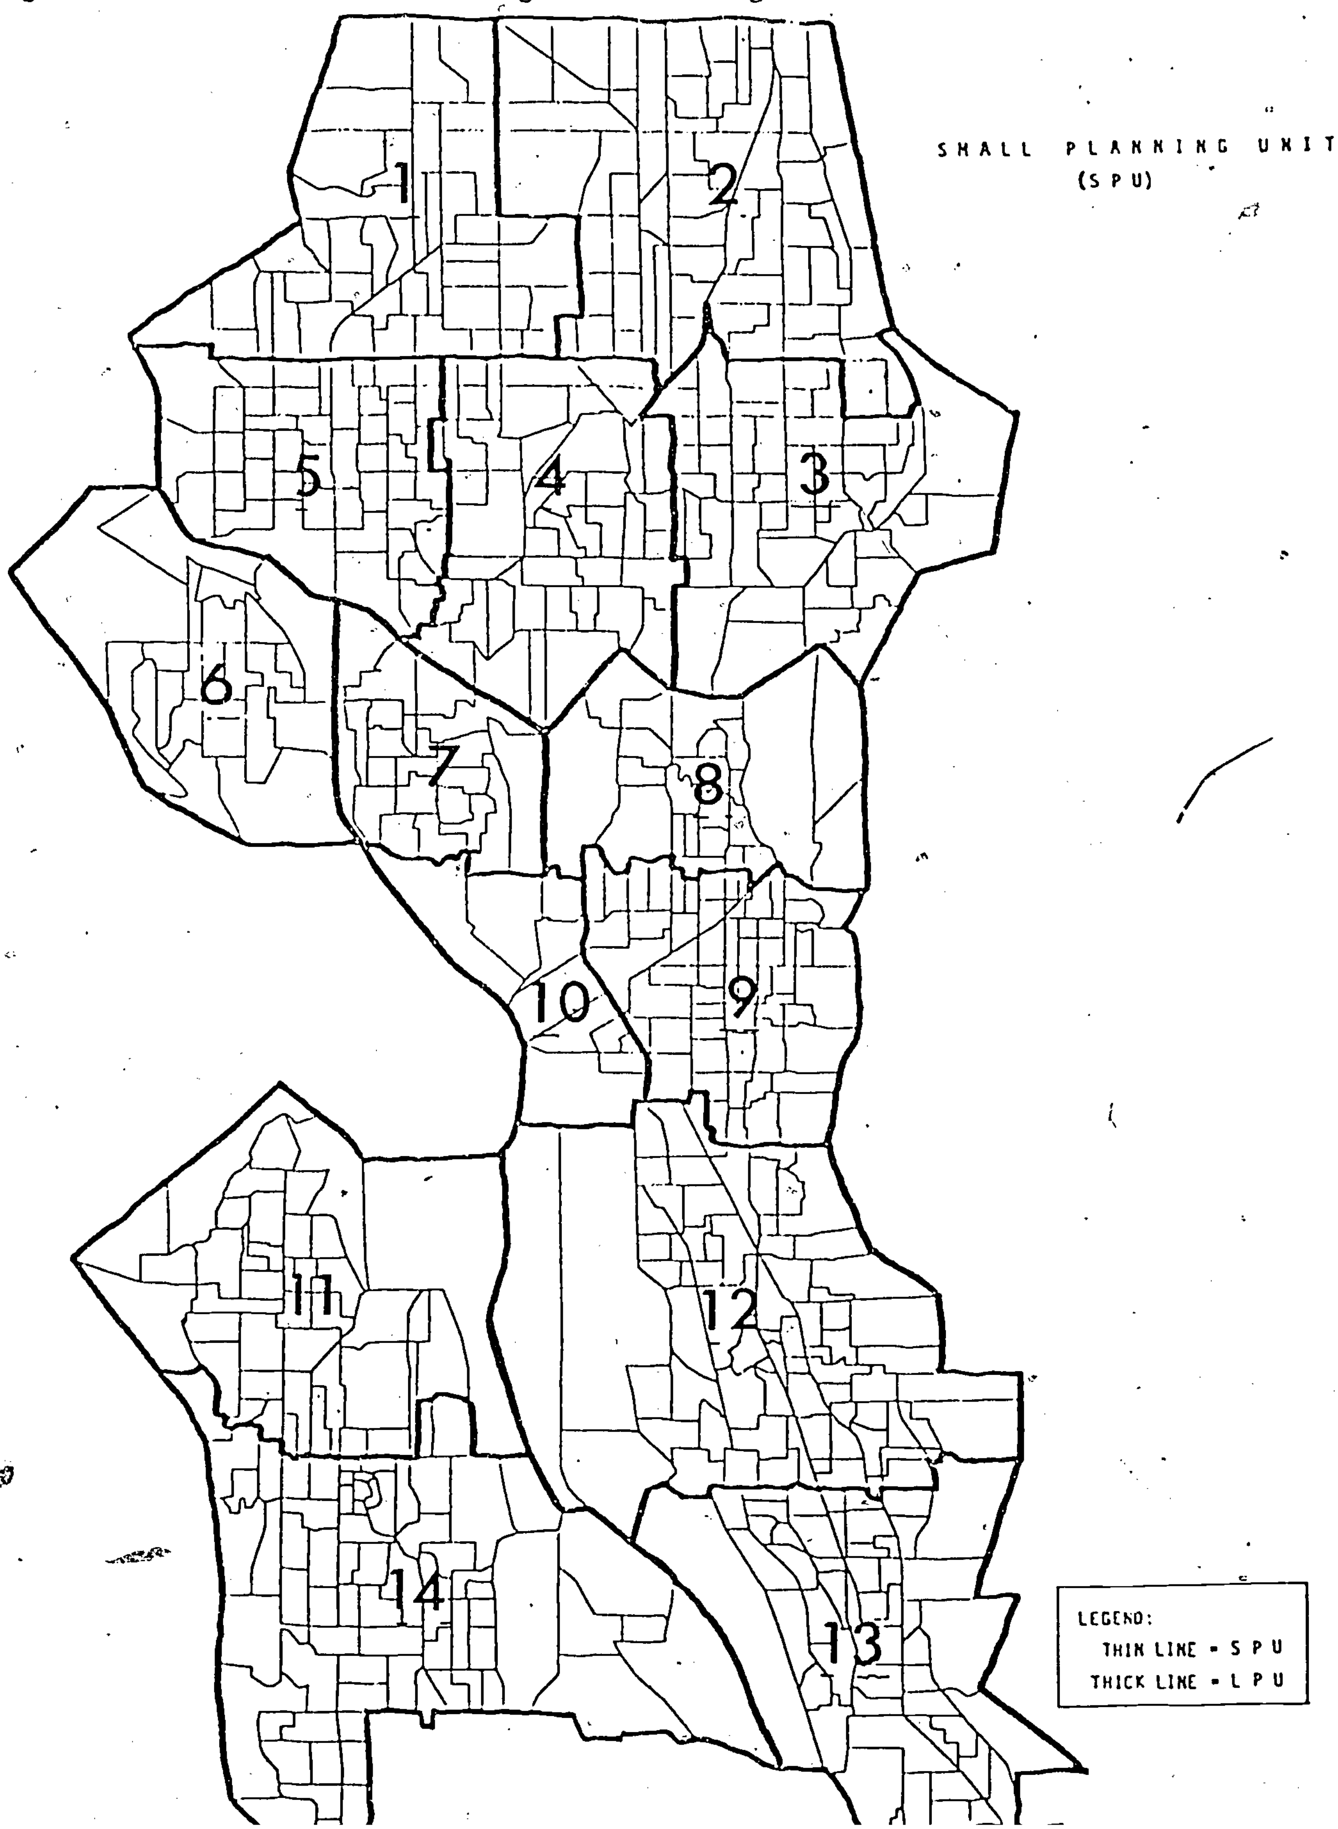

When the history has been constructed, the techniques described in the appendix are used to prepare a five year enrollment projection by grade and race for each SPU. Current enrollments are pushed forward using holding powers estimated by a linear combination of grade effects, year effects and mandatory bussing treatment effects. Kindergarten and first grade enrollments for the future five years are estimated separately, using an allocation technique based on births at the LPU level. The most important concern is not accuracy at the SPU level, but rather the prediction of groups of SPU's, so that when the SPU forecasts are aggregated and put through the feeder patterns the sums will be accurate. The first test of the system was to estimate the 1982 enrollments using history prior to 1982. The results of this forecast were then compared with actual 1982 enrollments at the LPU level. The results of this test are summarized in Table 1.

On the whole the results are quite satisfactory at the LPU level and they provide support for the notion that aggregation of SPU forecasts will result in a cancellation of errors. Most of the LPU's have an error rate less than four per cent. Three LPU's have higher error rates but no error rates are extremely high.

In LPU number 1, the error rate of 6.3 per cent is largely explained by a closure that occurred the previous year. This LPU is a good illustration of the investigative power of a system like Seattle's and will be dicussed in some detail.

When the relatively large error in LPU 1 was noticed by the planning staff, the small planning units that made up the LPU were reviewed individually to see if a pattern of under-estimation could be found. It was noticed that a geographical cluster of SPU's, numbered from 0175 to

TABLE 1. Prediction Errors by Large Planning Unit for 1982. Enrollments were estimated at the Small Planning Unit level and then aggregated.

Large Planning Unit (LPU)	Estimated 1982 Enrollment	Actual 1982 Enrollment	Difference (Act - Est)	Per Cent Difference
01	2702	2874	+172	+6.3
02	3943	4019	+76	+1.9
03	2927	2971	+44	+1.5
04	2607	2620	+13	+0.5
05	3218	3175	-43	-1.3
06	1235	1165	-70	-5.6
07	1280	1288	+8	+0.6
08	1374	1402	+28	+2.0
09	4894	4830	-64	-1.3
10	628	572	-56	-8.9
11	2777	2878	+101	+3.6
12	6957	6981	+24	+0.3
13	5043	4925	-118	-2.3
14	5127	5044	-83	-1.6
TOTAL	44712	44744	+32	+0.07

0183, all had under-estimations of enrollment ranging from 20 to 50 per cent. The pattern of under-estimation was common to all grade levels and occurred among both whites and minorities.

Further investigation revealed that the SPU's indicated were part of an attendance boundary change that had occurred in the previous year. A neighborhood school, Oak Lake, had been closed and the students attending had been redistributed among nearby attendance areas. As a result of this change, an enrollment decline had been predicted, but this decline had not occurred and instead enrollment had increased slightly throughout the SPU's indicated. The Oak Lake problem accounted for about 100 of the students in the LPU discrepancy, leaving a residual much more similar to other LPU's in the system.

Similar explanations have been found for the other LPU errors above four per cent. These investigations can help to correct the forecasting algorithms or to indicate areas where manual corrections need to be made. A system such as Seattle's is uniquely able to trace errors in this way and refine and correct itself on the basis of actual attendance results.

STAGE 2 - Attendance Forecasting

When the residential forecasts are completed, the simulation phase is run to predict the attending population at each school or program. To accomplish this the forecasted enrollment in each SPU is divided into an in-feeder component and an out-of-feeder component. The in-feeder component is the estimated number of students that will follow their feeder pattern and attend the school mandated by the district. The out-of-feeder component is the residual, those students who will choose some optional program and not attend their feeder pattern school. The division into the two components is done separately by grade and ethnic group.

The in-feeder component is maintained in the system as a forecast file by SPU. Simulation is done by assigning each SPU to a neighborhood attendance area and then putting students from the attendance area through a feeder pattern to arrive at an estimated attendance of in-feeder students by school. The simulated attendance area boundaries and feeder patterns can be varied from run to run to explore the effects of various policy options. Figure 2 shows the Seattle SPU's with the 1982 neighborhood K attendance boundaries shown as dark lines.

The students from the out-of-feeder component are relegated to an out of feeder "pool". Each school or special program draws from the out of feeder pool in a proportion estimated at each grade level based on the unit's historical ability to attract students in the past. The sum of the in-feeder component and the out-of-feeder component make up the estimated enrollment at the school or program level.

The estimation of the in-feeder/out-of-feeder components has proved to be the most troublesome part of the system, and refinements are still being introduced to cope with the errors made in this stage. The problem is due to the unexpectedly high variation in out-of-feeder movement between SPU's. It appears that local variation is very high even among units with similar socioeconomic composition, racial makeup and bussing treatment.

The problem is exemplified in Table 2. The table shows the proportion in-feeder among ten small planning units from the north end of the city. Each of these units is an upper-income mostly white neighborhood. All of the units have the same assignment history, i.e., they are part of the same attendance area and have an identical bussing history. In spite of these similarities, the proportion in-feeder for whites varies from a low of .28

Figure 2. Seattle Small Planning Units showing neighborhood K boundari

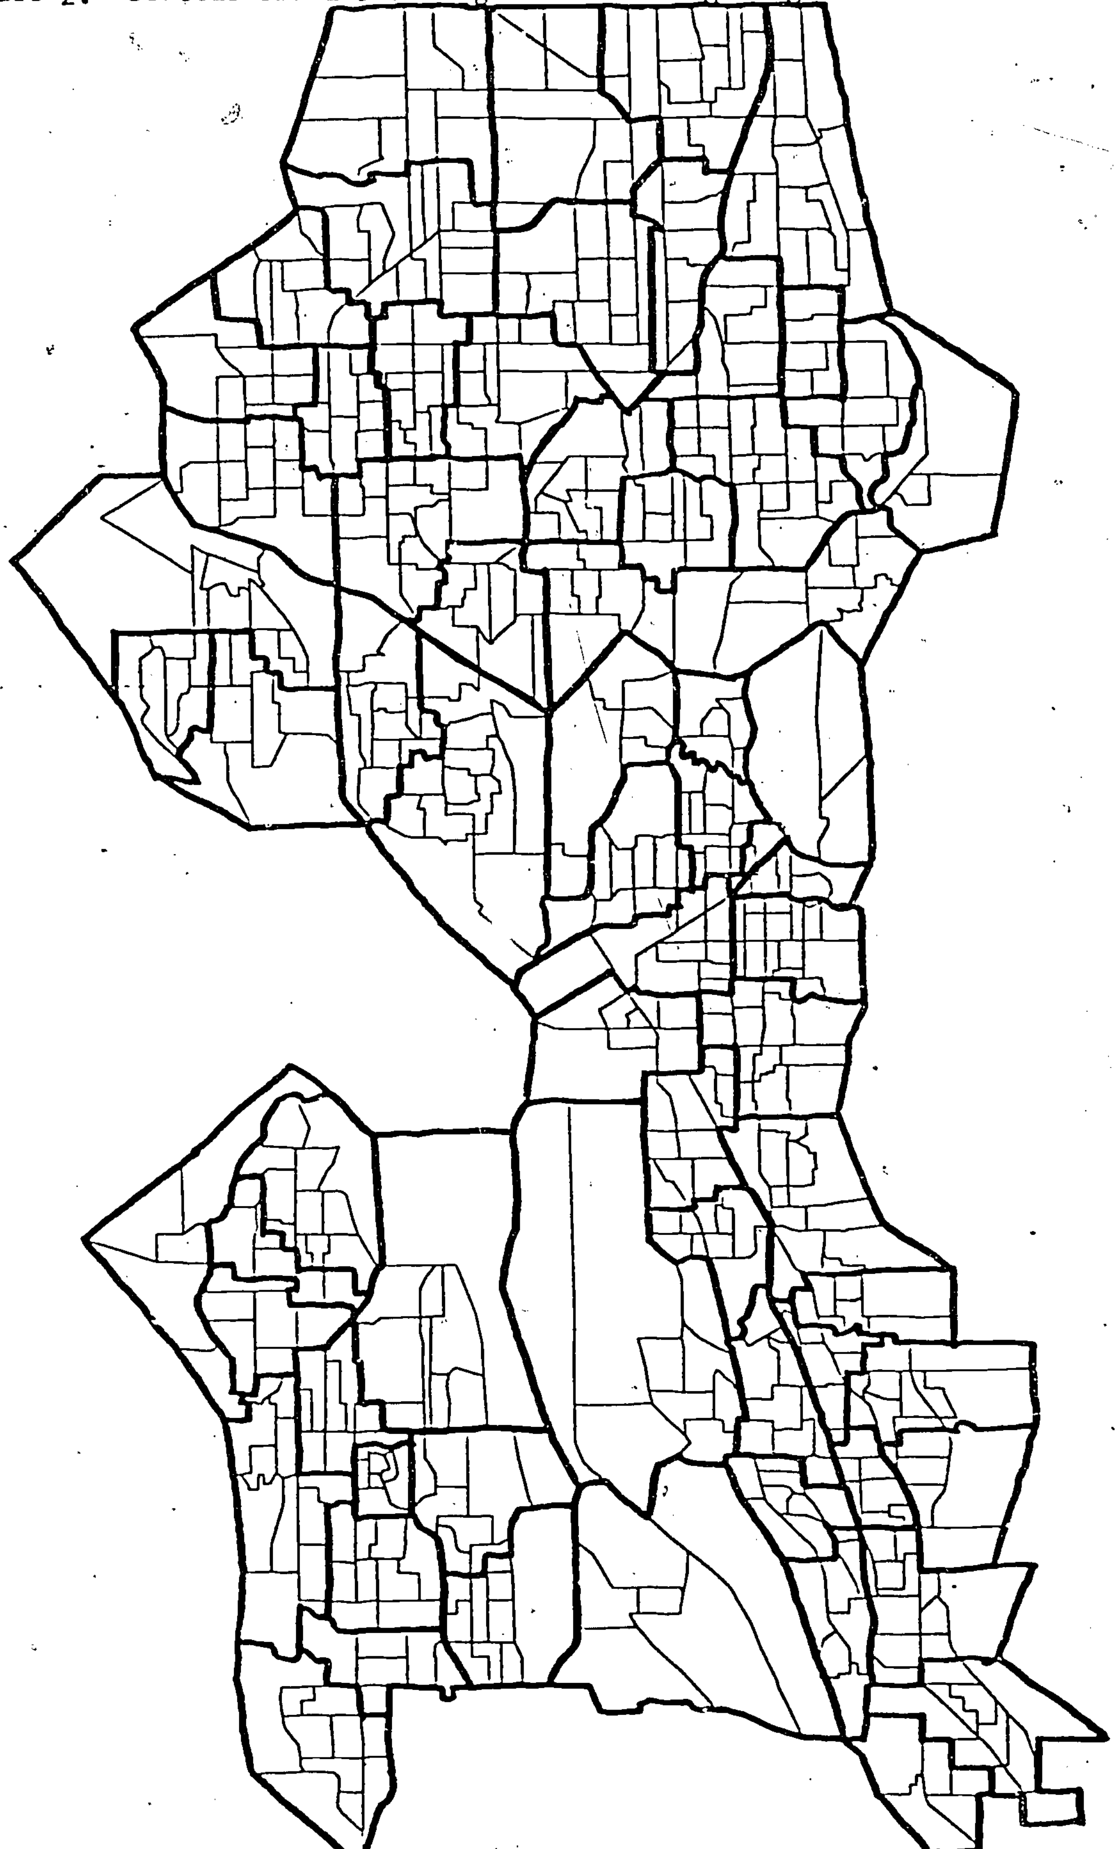

TABLE 2. Proportion in-feeder under mandatory bussing for selected northern Seattle SPU's with identical bussing histories.

SPU	Whites	Minorities
0267	.62	.55
0268	.38	----
0276	.64	.45
0277	.48	.44
0278	.73	----
0279	.84	.73
0280	.28	.28
0281	.51	----

---- indicates too few minorities to estimate to a high of .84 with an approximately random distribution between the two extremes. Among minorities, the proportions are similarly distributed.

This kind of variation makes accurate simulation of the in-feeder and out-of-feeder components quite difficult. Initial estimates have not been satisfactory, but refinements introduced later have improved the error rate until our most recent estimates are sufficiently accurate for about 70 per cent of the schools with manual corrections needed for the remainder. Development is continuing, but the expected upper limit is that about 85-90 per cent of the schools can be estimated accurately and manual corrections will be needed for the rest. Subsequent papers planned for the summer and fall will report on the final estimation techniques.

Summary

Putting aside for the moment the procedures which generate forecasts, consider first the form of the forecasts which will be most useful for planning purposes. That is, what are the optimal classifications and aggregations of the student population?

Clearly, this depends upon types of decisions, policies, and information which are required. If, for example, the state government pays $2,000 from general revenues to local school districts for each K-12 student, then at the state level all that need be forecasted is the estimated total number of K-12 students. This information will suffice for preparation of the state budget. If the state subsidy varies according to a student's grade, minority status, or involvement in a special program, say, then these sub-populations must be identifiable parts of the forecast.

Operating budgets, staffing needs, a variety of contractual arrangements, and certain educational objectives rest upon estimates of future numbers of students of various types at the local district level. The district planners need to know the distribution of students across grades, ethnic groups, special programs, specific schools, and specific parts of the geographic area of the district. These numbers provide the basis for selecting certain policy options as well as for calculating the various costs.

The following kinds of facilities management requirements, all of which have arisen in recent years in the Seattle Public Schools, illustrate the demands which a forecast should satisfy.

1. School closures. Will future enrollments be low enough that certain schools should be closed? If so, which should be closed? How should the attendance area for that school be divided among other schools?

2. De-segregation planning. De-segregation efforts are often based on mandatory busing between pairs or triads of elementary schools and on heavily structured feeder patterns for middle schools/junior highs and for high schools. These patterns are intended to remain stable for several years, requiring potential ethnic breakdowns several years in advance to ensure that schools will remain racially balanced. Transportation time and cost can vary substantially depending on the selection of pairs and feeder patterns.

3. Changes in grade configurations. An option that may be considered as enrollments change is a different allocation of grades to facilities. For example, an elementary school may be converted to a middle school, or vice versa. Another system-wide type of choice is whether the sixth grade should be part of the elementary schools or of the junior highs, and whether the ninth grade should be in middle or high schools.

4. Placement of special programs. A large district can have several special programs which are partially discrete and partially combined with the general program. These include special classes or programs for students who speak English as a second language, who are emotionally handicapped, who are physically handicapped in various ways, who are gifted in various ways, etc. High schools can specialize in vocational training or college preparation. The physical placement of these programs is related to the space available in school buildings, to the existing ethnic balance in those schools, and to the costs of transporting students to these sites, to the extent that transportation is provided.

Other, related kinds of planning needs, such as planning for new construction (the converse of planning for school closures) could of course be added to the above list.

Simulation of outcomes is the major device used by many school district planners to lead to decisions in such areas. The set of policy alternatives is typically rather small, except for degrees of "fine tuning", and the measurable costs and benefits of taking each possible option can be listed and compared. For example, if the school-age population is declining in a part of the district, it may be argued that one out of four possible elementary schools should be closed and attendance areas then be modified to re-distribute the students who would have attended that school. The procedure followed by district planners is then basically to work through each alternative--including the option of not closing any school--using such factors as characteristics of the school buildings, number of staff affected, the impact on racial balance, the number of students (if any) who must be provided with transportation, etc. These calculations obviously depend heavily upon the forecasted enrollment.

In certain contexts, in which the individual school is a stable, clearly defined unit with historical continuity, planning strategies can be school-based. That is, enrollments can be forecasted by applying the Grade Progression Ratio method (or a more sophisticated variant) to the school history and producing forecasts in which the school is the basic entity, broken down by grade and perhaps ethnicity. This level is fully adequate for some kinds of simulations--e.g. comparisons of alternative strategies for pairing of schools. School-based forecasts are inadequate for simulating more general types of management options, however, and are inadequate for even the simplest types when the school reflects a history of management changes. For example, if an ESL program has been placed in a particular school building because of available space, but could easily be shifted to another building, then it is obvious that the ESL program and the rest of the school should be treated as separate populations. If a school's attendance boundaries have recently been changed, then across the period of changes "the school" is not a meaningful entity and can only be used with various adjustments.

The essence of the management procedure proposed here is the subdivision of the district's total student population into a number of small units. In the case of Seattle, the units were geographical in nature but there is no theoretical reason why units based upon other criteria would not be acceptable. The essential feature is that the units have sufficient cohesiveness and stability to allow effective forecasting, and that they be useful to planners in simulating policy alternatives.

When small units have been decided upon, forecasts can be constructed for each unit using some of the techniques outlined in the appendix. Forecasts created at the small unit level can then be aggregated to produce large unit forecasts at any desired level of size or complexity. The fact that the small unit forecasts can be aggregated in t way is what makes the technique such an effective planning tool.

APPENDIX

Forecasting Model Used with Small Planning Units

1. The procedure to forecast the number of resident children in each grade level in a Small Planning Unit (SPU) is a modification of the traditional grade progression or holding power or cohort survival procedures. Let g and t be subscripts to denote grade (0,1,2.. ...,12) and year (starting with 1 for the first year of historical data, e.g., 1975). From student history files we can reconstruct the number of students living in a specific SPU in year t and enrolled in grade g somewhere in the school district; this number will be called N_{gt}. The future estimated values will be called $\hat{N}_{gt}$.

The traditional grade progression ratio is $R_{gt} = N_{g+1\ t+1} / N_{gt}$. It links the number of students in grade g+1 in year t+1 to the number in the previous grade in the previous year by the relationship $N_{g+1,t+1} = R_{gt} N_{gt}$. Our problem is to estimate the future values of these ratios. Once we have such estimates, which will be called $\hat{R}_{gt}$, then we can apply them to the last year of observation to get $\hat{N}_{g+1,t+1} = \hat{R}_{gt} N_{gt}$ and then in succession to get estimates of subsequent numbers of resident students.

The future ratios $\hat{R}_{gt}$ are obtained in two stages: (a) first, fitting the historical ratios with an additive model of "grade effects" and "year effects", and then (b) second, projecting the "year effects" and inserting the projected values in the model in (a). We shall describe these two steps briefly. If g = 0,1,. ....,12 and t = 1,.... 8 then the additive model

$$R_{gt} = M + G_g + T_t$$

is fitted with M, the main effect (or "average holding power") calculated as

$$M = \left(\sum_{g=0}^{11}\sum_{t=1}^{7} N_{g+1,t+1}\right) / \left(\sum_{g=0}^{11}\sum_{t=1}^{7} N_{gt}\right)$$

The grade effect for year y, called G_g, is a deviation from M calculated as

$$G_g = \left(\sum_{t=1}^{7} N_{g+1,t+1}\right) / \left(\sum_{t=1}^{7} N_{gt} - M\right)$$

and the year effect for year t is a similar deviation,

$$T_t = \left(\sum_{g=0}^{11} N_{g+1,t+1}\right) / \left(\sum_{g=0}^{11} N_{gt}\right) - M$$

By pooling the numerators and denominators of the grade- and year-specific holding powers, we are able to circumvent some of the instability in the specific ratios. The additive model is based on the assumption that there is no interaction between grade effects and year effects.

The main effects and grade effects are assumed not to change during the limited period (5 years) of our projection into the future. A variety of assumptions about the future year effects have been considered, but our preferred assumption is that they will remain at the average level of the past three years. A number of mechanisms are included in the procedure to alert us to recent holding powers which are so far from unity that it is implausible to assume that they will continue.

2. The above description bypasses the forecasting of the kindergarten (grade 0) enrollment. This is obtained for the immediate future from birth counts available for each census tract in Seattle. A ratio of kindergarten students to the original birth count is calculated from historical data and

pushed ahead at the LPU level, and the students are then allocated to each SPU in proportion to its historical share of the total kindergarteners in the LPU.

3. If all students attended the school or program to which they were assigned on the basis of residence, then the forecasting procedure for the resident population of SPU's would translate directly into school enrollments. However, as remarked in the body of this paper, a very substantial proportion (over one third) of Seattle Public School students are out of their feeder pattern. At present, the treatment of out-of-feeder students is only partially automated and relies heavily on the specific knowledge of the planning staff regarding the flexibility and attractiveness of the various alternative programs. The following paragraphs describe the current method being used to forecast out-of-feeder populations, but development is continuing in this area and the final system may differ in specific details.

The observed and forecasted proportions of students who are out of feeder in grade g and year t in a specific SPU are labelled P_{gt} and $\hat{P}_{gt}$ respectively. Unlike R_{gt} and $\hat{R}_{gt}$, which are generally close to 1.0, it is possible for P_{gt} and $\hat{P}_{gt}$ to be anywhere within a range from 0.0 to 1.0, but never outside of that range. Therefore, following the logic of log linear analysis, and to avoid the possibility of projecting a proportion outside of the legal range, we use here an additive model analogous to that used for R_{gt}, but of the log-odds form

$$\log [P_{gt} / (1 - P_{gt})] = M + G_g + T_t + B X_{gt}$$

B is a "treatment" effect, which arises from a distinction between grade levels which require bussing and those which do not. If bussing was required in order to stay in feeder, then $X_{gt} = 1$; otherwise, if the feeder pattern was into the local school, then $X_{gt} = 0$. The proportion out-of-feeder is consistently higher in grade levels which require bussing

Pooled estimates of M, G_g, T_t, and B may be calculated from the historical data. Future values of the log-odds of P_{gt} are calculated by assuming a linear trend through T_t but constant values of the other parameters. Alternative feeder patterns will change the value of X_{gt} for future values of t. $\hat{P}_{gt}$ itself is obtained from the inverse of the log-odds, and then the forecasted number out of feeder in grade g and future year t in an SPU will be $\hat{P}_{gt} \hat{N}_{gt}$; the forecasted number in feeder will be $(1 - \hat{P}_{gt}) \hat{N}_{gt}$.

4. Finally, we must consider the allocation of the out of feeder population. This is perhaps the most difficult of all aspects of the procedure. All of the students in grade g and year t who are projected to be out of feeder are aggregated across all SPU's in the city. This giant pool of students is then allocated to the programs which are open to them as alternatives to their feeder pattern. Some of these options, such as the gifted student programs, have many applicants but strict size limits, so that their forecasts are wholly a result of policy decisions. Other programs, which are essentially voluntary racial transfer programs, are much more flexible, but still subject to capacity constraints at the receiving schools. Estimates of attendance provided by the computer need to be adjusted manually by the planning staff to be consistent with some of the concerns expressed above.

Bibliography

Britney, Robert R., "Forecasting Enrollments: Comparison of a Markov Chain and Circuitless Flow Network Model", Socio-Economic Planning Sciences, Vol. 9, 1975.

Callahan, Raymond E., Education and the Cult of Efficiency, University of Chicago 1962.

Harden, Warren R. and Tcheng, Mike T., "Projection of Enrollment Disribution with Enrollment Ceilings by Markov Processes", Socio-Economic Planning Sciences, Vol. 5, 1971.

Johnstone, James N. and Philip, Hugh "The Application of a Markov Chain in Educational Planning," Socio-Economic Planning Sciences, Vol 7, 1973.

Katz, Michael B., The Irony of Early School Reform, Beacon Press, 1968.

Knutson, Dennis L. with Marquis, Linda M., Rischiute, David N. and Saunders, Gary T. "A Goal Programming Model for Achieving Racial Balance in Public Schools," Socio-Economic Planning Sciences, Vol 14, 1980

Lee, Everett S. and Goldsmith, Harold F., Population Estimates - Methods for Small Area Analysis, Sage Beverly Hills, 1982.

McNamara, James F., "Mathematical Programming Applications in Educational Planning," Socio-Economic Planning Sciences, Vol 7 1973.

Pullum, Thomas W. and Graham, Stephen N., "Measuring the Implementation of Racial Balance under Formal Constraints," Educational Evaluation and Policy Analysis, Vol 4(1), 1982.

Rives, Norfleet W. "Forecasting Public School Enrollment," Socio-Economic Planning Sciences, Vol 11, 1977.