8

Implementation

Any attempt to implement the models and methods of this book creates in itself a process that requires analysis. We would be poor analysts indeed if the extent of our reasoned thinking about an urban problem stopped with the mathematics. Rarely are we confronted with the simple situation in which all participants in a study share the same priorities and in which the modeling and data collection tasks are inexpensive and straightforward. More likely, things are much more complicated.

Since many readers of this book will eventually attempt to develop and utilize models in an actual urban setting, we wish to close the book by presenting some views on implementation. In the first part of the chapter, to illustrate the variety of forces and factors that can come into play, we provide thumbnail sketches of several "mini case studies" or "war stories" with which we are familiar. We assume that the reader has reviewed the eight steps in an operations research study that were discussed in Chapter 1. The case studies point to the difficulties of defining performance measures, prespecifying all constraints, identifying all decision makers, and retaining sustained commitment of key personnel. Motivated by the cases, in Section 8.2 we attempt to outline in a general way the model-related issues found to be important when attempting implementation. Here, for instance, experience has shown that a model's data-base requirement is often a much more troublesome factor than is model accuracy. Finally, in Section 8.3, we discuss the all-too-critical issues related to people and institutions that affect the success or failure of a model-based analysis. For example, a potentially fatal flaw in this area is for an analyst to work solely through a single individual in any agency—one who speaks "model language" and "agency language." Such an individual may be transferred or promoted by the time the model is ready for use. With the model's major in-house advocate having "vanished" from the scene, the model, even if mathematically sound and sophisticated, may never be used.

Throughout this chapter, we think conceptually of model implementation as taking place within some "intervention program," whose purpose is to analyze and (probably) change one or more operating procedures of an urban agency. While such a program may be rather broad-based, including many types of activities unrelated to models, we naturally focus on its model-relevant attributes. The intervention program is characterized by a number of *inputs* [e.g., the model(s), the analyst(s), the user(s), a feasible set of decision options], which give rise to a *process* [e.g., the use of the model(s), the training of personnel, the collection of data], which transforms the inputs into desirable or undesirable *outcomes* (e.g., more accessible service, more efficient operation). Our focus in this chapter will be on the input and process features of a model-based intervention program.

Research on the implementation of model-based studies, and policy analytic studies in general, is a new and burgeoning field. Our purpose in this chapter is to discuss critical points directly relevant to implementing methods of this book. Hopefully, those who are interested in additional material on implementation will consult the chapter references or the general references that follow this chapter. There is no simple recipe for guaranteeing implementation success. Our hope here is to establish an awareness of the need to step beyond the mere technical features of a program and to provide a framework for formulating practicable implementation strategies.

8.1 SOME MINI "WAR STORIES"

8.1.1 Don't Take My Men Away!

As our first example, we focus on a spatial and temporal reallocation of personnel within an urban service system. Any time that allocation levels are changed, certain groups or individuals are likely to "gain" while others are likely to "lose," at least as measured on a perceptual basis. Feedback from these groups and individuals may suggest modifications of the objectives, performance measures, and/or constraints under which the analysis is carried out. This happened in the late 1960s in the New York City Police Department when queueing analysis showed that police patrol cars could be "traded" during certain hours among police precincts in southern Manhattan to reduce queueing delays. While the total number of patrol cars in the set of

precincts concerned would remain constant, Erlang's formulas (see Section 4.5) showed that, because of differences in the 24-hour demand patterns among the precincts, substantial reductions in queueing delays could be achieved. The trading was to be implemented by sending one or two of precinct A's cars to precinct B at appointed hours, when the latter precinct had unusually high demand. At other hours, when precinct A's demand was relatively high, reciprocal transfers would occur. In a pilot run, carefully planned and monitored, queueing delays were reduced by about the amounts anticipated.

However, precinct commanders, when on duty during the "depleted hours," found the plan unacceptable. Although it was shown that queueing delays had not increased measurably during the depleted hours in their precincts, the commanders felt that they would be unable to handle a serious emergency (e.g., large fire, civil disturbance) in their precincts, given its depleted resources, without requesting assistance from other precincts. Quantified, their perceived objective appeared not to be queueing delay reduction but something close to a "minimax" strategy, minimization of the deleterious consequences of the largest conceivable event requiring police services that could occur in their precincts. However, many nonquantifiable factors also pervaded the issue, including the perceived power associated with the number of officers under one's command, the dependence associated with relying on someone else's resources in times of possible large-scale emergency, the cutting of traditional officer–supervisor relationships by having the officers work temporarily under different supervisors, the resistance to change brought about by directive from headquarters, and so on. As a consequence, this effort to vary patrol levels by time of day was abandoned.

In this case, none of the model-oriented factors in the study were problematic. However, because the model users (headquarters staff members) did not encompass all relevant decision makers—some of whom had objectives strongly conflicting with the objectives of the model users—implementation of the program as originally designed was infeasible.[1]

8.1.2 What—Design a Non-perfectly Working System?

It is often difficult for an agency administrator to specify explicitly target levels for performance measures. This is particularly true when the target levels must be stated in probabilistic terms, thereby acknowledging that X percent of the time the system will not perform up to standards.

A few years ago one of us was involved in a very short term (one person-

[1] Subsequent analyses led to alternative, more politically acceptable reallocations. See, for example, [LEVI 78].

month of total effort) study for scheduling telephone operators in a new city-wide 911 system for New York City. The three-digit number 911 is now being implemented throughout the United States as a standardized emergency public safety number through which callers are linked to the local police or fire department or ambulance service. However, at the time of this study, it was a relatively novel concept. In New York's 911 installation, queueing delays had become intolerable during certain periods of the day and days of the week. Simple $M/M/N$ queueing theory provided a reasonable initial procedure for rescheduling the personnel to reduce significantly the delays experienced.[2] However, to perform the rescheduling, it was necessary to have the police commissioner (or his designated representative) state an allocation objective of the following type:

> Specify desired values for T and P, so that no more than P percent of the 911 calls will be delayed T or more seconds.

For instance, if $T = 15$ and $P = 5$, then the rescheduling based on the $M/M/N$ queue would allocate a sufficient number of telephone operators $N(t)$ during hour t so that no more than 5 percent of the calls would be delayed 15 or more seconds. While an average of 17 percent of calls were being delayed 15 or more seconds in the existing system, there were predictable periods during which 40 percent of the calls were being delayed 30 or more seconds. However, the police commissioner was quite reluctant to provide values for T and P. Finally, he settled on $T = P = 0.0$. But, given the probabilistic nature of queueing systems, such perfect operation is virtually impossible with a finite number of servers.

The commissioner's unwillingness to specify a positive value for P clearly rested in his uneasiness in acknowledging (and approving!) a system that functions imperfectly. Imagine a headline in a local newspaper, if the information were leaked that $T = 15$ seconds and $P = 5$ (assuming that 10,000 calls are received per day): "Police Commissioner Approves Excessive Delay in Answering 182,500 Police Calls This Year!"

Therefore, not only are explicit statements of objectives difficult to elicit, but they are even more difficult when they contain a certain admission of failure as a result of the probabilistic nature of the system.

This case has both a happy and a sad ending. Back at the drawing board, it was decided to *redefine* the problem as that of rescheduling the currently

[2] For the decision at hand, a more sophisticated analysis was not called for. Neither the time nor the resources were available to construct a detailed model, nor would the resulting resource allocation benefits of a detailed model have been worth the costs in time and money.

available person-hours to achieve the lowest possible mean queueing delay throughout all hours of the day. Thus, the problem became one of redeploying person-hours during relatively overstaffed periods to understaffed periods. To be sure, considerably improved values of T and P resulted from this procedure, but the police commissioner was never confronted with specifying them. The derived rescheduling was implemented in its entirety, approximately one month after the study was completed [LARS 72a].

The sad ending is associated with training of personnel. Approximately 8 years after the analysis reported here, insufficient training of 911 operators resulted in widely publicized inadequate response to several high-priority calls. Fully 121 operators were subsequently transferred to other duties (Figure 8.1). Thus, the successful operation of any system requires continued professional training of personnel. This is just as true for model users as for 911 operators; Section 8.3 offers some suggestions for the training of model users.

FIGURE 8.1 Recent Problems with 911. (Reprinted with permission of the *New York Post*. © 1978, New York Post Corporation.)

8.1.3 Why Spend Time and Money Making Decisions?

A decision maker often places little importance on resource allocation decisions that have profound consequences. As an example, several years ago one of us was working indirectly as a consultant on police patrol allocation in one of the largest police departments in the United States. The number of patrol officers to allocate among police commands was in excess of 10,000. The direct annual dollar costs of the patrol force exceeded $200 million at that time. After work had proceeded on this effort for one month,[3] a formal progress briefing was held before the police commissioner. Just prior to the briefing the commissioner said (in all seriousness), "What's the matter, Larson, you've been working full time on this problem for one month and you haven't yet indicated where I should put my men. Each year I assign a limited-duty sergeant to work half-time on this problem for two weeks, and at the end of the two weeks he always *gives me the numbers*" [emphasis added].

Apparently, one set of numbers was as good to this decision maker as any other set. Resource allocation had not been perceived as a problem and thus little attention had been given to it. The 911 system described above was finite; people could understand it; if no one answered the telephone for several minutes, citizens knew what to complain about. Inadequate performance of the patrol force was more difficult to identify; citizens' complaints about the police were usually targeted elsewhere, although many perceived defects were probably linked to and, in some cases, caused by inappropriate allocation of resources.

The patrol allocation method used by the sergeant entailed a simple linear hazard formula (see e.g., [LARS 72b] or [FERR 78]) containing for each patrol command about five factors, such as number of serious crimes, square mileage, street mileage, number of calls for service, and so on. Patrol officers were allocated in direct proportion to the hazard formula "scores" for each patrol command. Of course, over the years it had been found that this method produced obviously poor allocations in certain far-from-average areas (e.g., large, sparsely settled communities), perhaps resulting in unusually long travel or response times. Therefore, results of the hazard formulas were juggled by hand once a year, thus requiring one full week of a sergeant's time on that project.

As discussed in Chapter 1, performance measures are essential inputs to a resource allocation process. By setting target levels for each, a decision maker knows whether and to what extent he or she has achieved his objectives. Most entries in a hazard formula are input measures (e.g., street mileage). The formula itself mixes these measures together, much as mixing "apples,

[3]This occurred *before* the development of most of the models described in this book.

oranges, peaches, and pears," yielding nothing even resembling a performance measure. But suppose the police commissioner decides that he wants to reduce average response time to urgent calls for police service from, say, 6 minutes to 4 minutes. This, then, implies a performance measure and a revised standard of performance. A hazard formula provides no guidance in addressing such questions, whereas the methods of this book—based on performance measures—do.

As an illustrative example, suppose that the commissioner has two options: (A) add sufficiently many new patrol units, deployed under the current hazard formula's proportional allocation scheme, so that the new performance standard is met; and (B) undertake a resource allocation study to analyze alternative ways of redeploying currently existing patrol units in order to reduce response time to urgent calls. Option A is found to require N new patrol units; Option B can only reduce urgent response time to 4.8 minutes and requires $N' < N$ new units to achieve the desired 4-minute level. If each patrol unit costs $200,000 per year, option A costs $\$(200{,}000)\cdot N$; option B costs $\$(200{,}000)\cdot N'$ plus the one-time cost of the allocation study, say $100,000. As long as

$$(200{,}000)\cdot N' + 100{,}000 < (200{,}000)N$$

then the allocation study option is the more cost-effective option, even when charged entirely to the first year of operation. In fact, the first year savings[4] of ($\$200{,}000\cdot(N - N') - \$100{,}000$) could be a considerable sum of money. And each subsequent year, an additional $\$200{,}000\cdot(N - N')$ would be saved. In the absence of performance measures, the police commissioner in our study saw no such value in spending time and money "thinking about decisions."

8.1.4 Busing to Build

The causes for success or failure in implementing a study are often complicated, requiring months, perhaps years, to discover. Probably the most successful of the student projects in the MIT graduate course related to this book involved the redeployment of public school buses in a suburb of Boston. (This project occurred before the recent court-ordered busing of schoolchildren in Boston to achieve racial balance.) Each year this municipality put a contract out for bid to bus-leasing companies to provide the community's school busing services during the coming year. The cost of the contract turned out to be almost directly proportional to the number of buses required, which

[4]A more sophisticated way of computing the cost effectiveness ratio, as discussed in Chapter 1, is to use the concept of (discounted) present value. This is not necessary here to illustrate the point of the example.

was specified by the municipality in accordance with the bus routes and schedules annually developed by the Department of Schools. Over the years this community had grown considerably, with new housing occupying once-undeveloped land. As a result, the number of schools and school-bus routes had increased greatly. School-bus routes had been added on top of existing routes, with often two (and in one case three) buses stopping at the same stop at different times in the morning to transport children to the same school!

For their project, a group of three of our students worked with members of the School Committee to redraw school bus routes and reschedule their arrival and departure times (the same bus would make several trips in the morning, as elementary, junior high, and high schools each had different starting times). By employing a good deal of common sense and a number of heuristic routing techniques from network analysis, the students prepared a recommended redeployment of buses for one of the junior high schools. After giving a class presentation, which was attended by the Superintendent of Schools of that suburb, they were invited to present their results formally to the entire School Committee.

Almost immediately after their formal presentation, they were hired by the School Committee to work through the summer on implementing a new bus routing and scheduling system in September for the entire school system. This they did, but not without some problems. For instance, the newly derived bus routes suggested changing slightly the areas served by the "north" and "south" high schools. This shift caused the star football halfback for one of the schools to be reassigned to the other, the two being arch Thanksgiving Day rivals. Such a redesign was deemed politically unacceptable and resulted in gerrymandering that maintained the halfback's school allegiance (Figure 8.2). Other factors were also considered, such as safe locations for U-turns of buses and safe walkways for students, requiring extensive driving throughout the town to discover its idiosyncracies. Finally, after certain data-collection problems had been resolved in August, the plan was implemented on schedule in September. A telephone complaint service was set up to respond to calls from parents. This resulted in some further modifications of the bus deployment plan until the final plan for the year was arrived at approximately 2 weeks after the start of school.

The cost to the town for the students' summer professional services was $10,000. The savings in school-bus costs to the town for the first year alone were about $130,000, amounting to about 26 percent of its annual $500,000 bill for the school-bus services.

Virtually no one in a decision-making position objected to this study, and most supported it vigorously—an occurrence that is almost unparalleled in the authors' experiences. Within the following year, we discovered some of the hidden ingredients. First, the school superintendent was new at the job, having been hired the preceding spring; he did not have to bear the

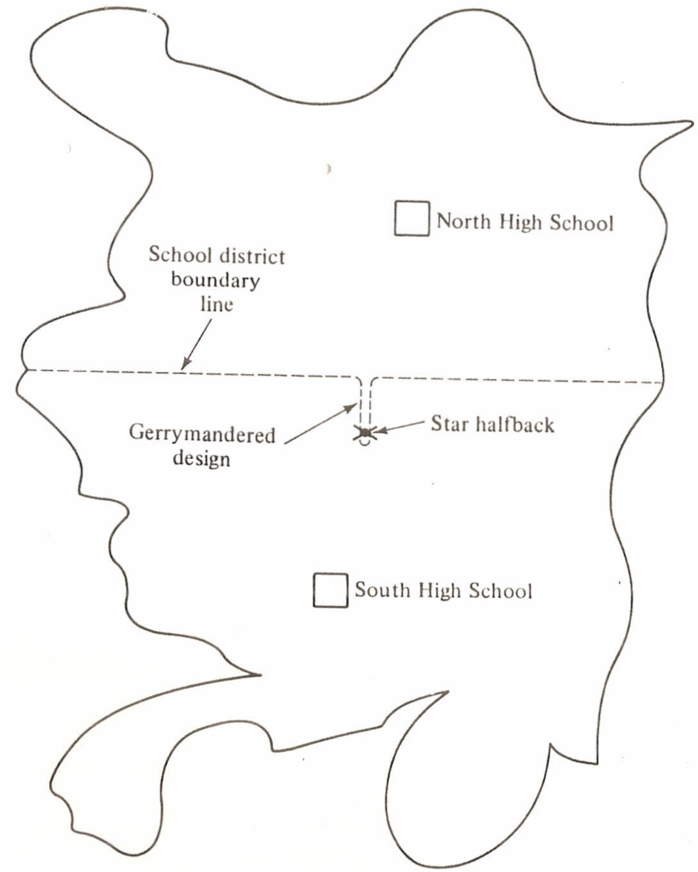

FIGURE 8.2 Saving the star halfback.

legacy for past inefficiencies. In addition, the person whose duty it had been to design school bus routes had retired at approximately the time the new superintendent took office. Thus, there was no one around to defend old policies. Most important, however, the school superintendent had a "hidden agenda": he was trying to put before the electorate a $12,000,000 school-construction bond issue, and the documented savings for school busing of $130,000 per year (and a 13:1 benefit cost ratio during that first year alone) demonstrated that he was "frugal," and, by implication, that the bond issue must be necessary.

8.1.5 Buses Come in Bunches

Other student projects did not meet with such immediate success, and some were simply ignored by decision makers. One topic that received the attention of two project groups (in two succeeding years) and resulted in two

Master's theses is the "bunching" or "clumping" problem of buses. (See Chapter 1, including the letter to the editor from an irritated bus passenger.) To restate the problem, briefly, the clumping of buses tends to occur naturally even when buses are dispatched at prescribed intervals from one end of a transportation route. Any system perturbation (e.g., more passengers than average, traffic delays) will cause one bus to slow down (operating behind schedule). This slowdown becomes more pronounced at each successive stop, as more and more passengers have accumulated since the time of the most recent bus. Meanwhile, the bus immediately behind the slowed bus "speeds up," since it finds fewer and fewer passengers waiting (because of the shortened bus interarrival time). Eventually, the buses "clump"; that is, they come together due to this accelerating process.

Our students developed both analytical and simulation models to discover ways to reduce the occurrences of clumping. Several guidelines were established that involved only local controls, such as delaying the departure of a bus at a stop if the most recent bus left the stop "too recently." However, no one could be found in the local public transportation services who viewed this as a problem deserving of their attention. So, considerably more (and quite original) analytical work was done on bus clumping than on the school busing problem, but no implementation can be reported.

8.1.6 Hypercubed Ambulances

A student project on ambulance allocation in Boston in 1972 resulted in a Master's thesis [HOEY 73] in which the student used an early version of the hypercube queueing model to help determine the required number and locations of ambulances in Boston. The city, at that time, was planning a shift from a police-operated ambulance service to one run by the Department of Health and Hospitals. While the recommendations in the study were expressed in concrete terms (e.g., "have X ambulances available during these hours and put them in these locations"[5]), the student felt frustrated by the lack of any immediate action on his plans.

As it turned out, the actual decision process was a complicated multiyear process involving diverse citizen's groups, private ambulance firms, the police department, the fire department, and the Department of Health and Hospitals. Later, a state law was passed requiring upgrading of the emergency medical skills of all ambulance attendants. However, throughout this long, drawn-out process, the student's thesis (which was published) was used as the analytical blueprint for "what should be done, if we can figure out who should do it." Thus, it contributed to the analytical side of a highly political decision-making process.

[5]Special curb-side ambulance locations were allowed, thereby eliminating any constraint to locate them in prespecified ambulance garages.

In the spring of 1978, fully 6 years after the original thesis was written, the Department of Health and Hospitals contracted with a local not-for-profit firm to implement the hypercube queueing model on its own computer system. The model is now being used to deploy Boston's 14 new ambulances and to plan for system evolution over the next 10 to 15 years [BRAN 78].

8.1.7 Vanishing Advocate (Figure 8.3)

In Boston a simulation model originally developed by one of us for the Boston Police Department was reprogrammed with a natural language interface especially tailored to the needs of the Department. Virtually all of the work had been carried out in cooperation with one civilian planner in the Department, who happened to be a former MIT student. This planner, who became a close working assistant of the Police Commissioner, resigned with the Commissioner when the mayor decided against the Commissioner's

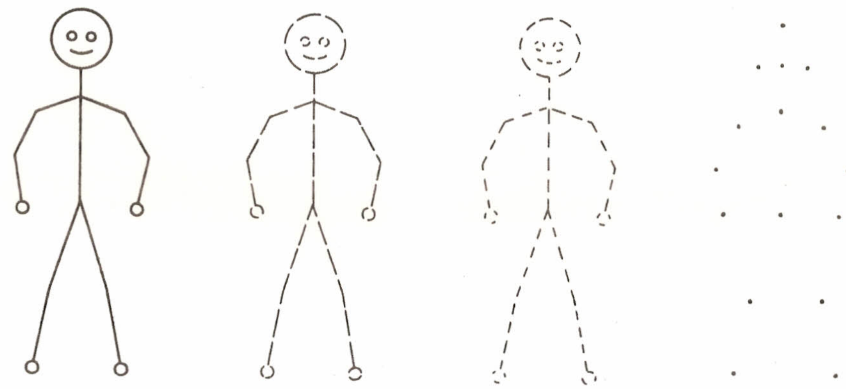

FIGURE 8.3 The vanishing advocate.

reappointment for another 5-year term. The simulation model at that time was only 4 months away from completion. Not only did no one else in the department know much about it, but the model itself took on a political identity associated with the "old regime." Despite attempts to teach others in the Department how to use the model, it was never used and went into "hibernation" after its completion. (For details, see [HEBE 78a].)

8.1.8 Model "On the Road"

In the late 1960s, Sergeant Glenn Pauly[6] played a key role in establishing a police resource allocation system based on an $M/M/N$ priority queueing model in the planning office of the St. Louis Police Department. A detailed

[6]Now a Captain, retired.

computer package was developed and, under Pauly's guidance, used by most of the district commanders in St. Louis. The package was so successful that it was picked up by IBM, modified slightly, and marketed as an IBM product.[7]

Later, IBM's package was implemented in the Los Angeles Police Department (LAPD). However, shortly thereafter the LAPD changed operating procedures to the "Basic Car Plan," which tended to queue calls in a police beat more frequently so that the beat car could respond to more of them. Thus, "beat identity" became the key objective rather than response-time reduction. The IBM model, employing as it did assumptions from the $M/M/N$ queueing model, no longer accurately portrayed operations in Los Angeles, and consequently fell into disuse. In this case, the model was not flexible enough to adapt to a new set of operating procedures nor were there personnel within the LAPD who could perform the modifications (see [HEBE 78b]).

8.2 MODEL-RELATED ISSUES AFFECTING IMPLEMENTATION

We now attempt to build from the case studies and related experiences to identify and discuss the general factors that can influence the chance of successful implementation of a model-based study. Our discussion starts with the central concern of this book—the mathematical model and its attributes—and broadens to include people and their institutions. In considering models for implementation, one can select preexisting models, often available at minimal cost in the public domain; or one can develop new models; or—combining the first two options—one can modify existing models to suit certain agency-specific needs. The individual making this model-selection decision may be the analyst (probably under contract to the user's agency) or the user himself, especially in cases in which an in-house study is being conducted. To address the decision in a systematic manner, the important attributes of each prospective model must be considered. No simple algorithm or formula exists for doing this. Evaluation of alternative models for use in a program is often a highly subjective process. Nevertheless, we attempt here to set out those model attributes that we and other researchers have found to be significant. Undoubtedly, in many programs other model-related factors will surface that we do not address here.

[7] Later Pauly was promoted out of the planning office, and the resource allocation package in St. Louis fell into disuse. Here an in-house advocate vanished by promotion.

8.2.1 Performance Measures

First and foremost, the outputs of a model—its computed *performance measures*—should be scrutinized. They should be understandable to both agency decision makers and the public. They should reflect to the extent possible the agency's stated objectives, relate to measurable quantitites in the actual system, be statistically stable under one set of operating policies, and if the model is to be used to compare alternatives, be dependent on the operating policy selected.

Given a meaningful and complete set of performance measures, one can begin to discuss returns on investment of a model-selected policy in an urban system, thereby bringing it slightly closer to its industrial counterparts. [DAEL 76]. For instance, in an urban services context one could begin to say: "Policy X costing $\$Y$ has brought the same improvement in operation as N additional personnel, costing $\$Z$, would bring"; thus, a fraction (Y/Z) of the "more patrol personnel" investment brings about an identical improvement in service. [Recall our patrol allocation study above ("Why Spend Time and Money Making Decisions?").] Such statements, however, may not be easy to make in systems with multiple and conflicting objectives; in such cases the model should provide the decision maker with estimated values of each of the competing performance measures, some of which are perhaps aggregate—reflecting overall efficiency—while others are distributive—reflecting equity of service among neighborhoods. Then the decision maker has to weigh factors subjectively (and objectively) to select an appropriate operating policy, perhaps using methods of multiattribute decision analysis [KEEN 76].

8.2.2 Model Accuracy

Given equally useful sets of performance measures, models are often compared on their predictive accuracy. This is fine if not taken to extremes. Model accuracy should be judged in terms of its decision-aiding utility. If a simple $M/M/N$ queueing model will suffice for determining a preferable scheduling of service personnel, an elaborate simulation is not required. [Simple models were dictated in the 911 case study, in which results were needed in one month.] All too often we see a model criticized for not depicting to the finest detail all of the operational idiosyncrasies of a system; this is a key motivation for builders of simulation models to err by inclusion much more often than by exclusion. As model creators, we may tend in our enthusiasm to transfer our priorities from the decision maker to the model itself. A model's comparative advantage in decision aiding is, in our opinion, more relevant than its predictive accuracy. In this light—to overstate the case—a model could be a factor of 2 in error on the primary performance measure, but (as long as the factor remains constant) the model would be fine for rank-ordering alternatives and assessing their relative merits. Of course, we do not advocate factor-of-2 errors. But we believe energies directed at concerns for multidecimal accuracy might be better expended in other parts of the implementation process.

8.2.3 Data for the Model

An important model-related attribute, a model's *data requirements*, can often represent major impediments to implementation. Chaiken reports in a study of urban model implementation successes and failures that the data need of a model is the most frequent model-based factor that led to failure of implementation. [CHAI 75]. One may have developed a highly efficient model, costing, say, just a few dollars per run on the computer; if, however, the cost of collecting data for the model runs in the thousands of dollars and requires months to complete, both the dollar and the time limits placed on the decision(s) at hand may be exceeded (see Figure 8.4). This factor is a highly relevant one to consider when thinking about degree of aggregation versus level of detail of a model.

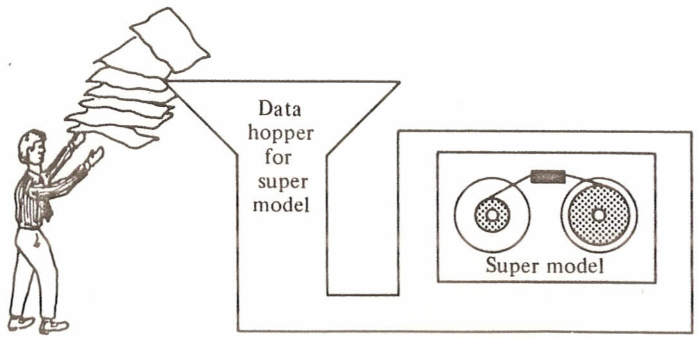

FIGURE 8.4 Even the best model can be too data hungry.

It has been our experience that for a given model, two very similar cities can incur very dissimilar data collection costs. Critical factors affecting data collection include an agency's former record keeping, both in terms of contents and of accuracy; the need for pilot data-collection efforts for new types of data; the number and types of personnel available to carry out data collection; and the amount and type of computer services required. Most often, accurate computer-encoded records also provide the *least expensive* and most timely data inputs. Boston's Department of Health and Hospitals has little trouble providing data for its hypercube model, since most required data are generated routinely by an on-line computer-aided dispatch (CAD) system.

8.2.4 Documentation

Complete and comprehensive documentation is necessary when an existing model from elsewhere is being transferred and perhaps modified. Even at the "home" for the model, the original builder and/or computer programmer will eventually leave the agency. Therefore, documentation is required for ongoing use, for various users whose identities probably change over time, and for updating to reflect changing city characteristics or agency policies.

A recent HUD-supported urban modeling effort at the Rand Corporation specified what we believe to be very reasonable documentation requirements for a transferable model [CHAI 78]:

1. An *executive summary*, containing a nontechnical introduction to the model, information to assist an administrator in deciding whether to use the model, and details about how the computer program can be obtained.
2. A *technical description*, to provide analysts with an understanding of the theoretical underpinnings of the model.
3. A *user's manual*, describing, step by step, how the model is operated once it is installed on a computer system.
4. A *description of the computer program* (a "programmer's manual"), written for data-processing personnel and providing sufficient information to permit installation of the model, construction of the required data base, and modification of the model, if desired.

Of course, simple models developed solely for local use may not require such extensive documentation.

Computer-related issues. Most models today are computer-implemented. This raises issues with regard to the following:

1. A model's computer requirements, including computer language (e.g., FORTRAN, COBOL, ALGOL, PL/1, SIMSCRIPT, etc.), core storage requirements, and execution-time requirements.
2. The *computer capabilities available* to the agency, including those of the city, local universities, and commercial vendors, as well as those controlled internally.
3. The *computer experience of the user.*

Many communities initiated computer usage several years ago, primarily for accounting purposes. A popular computer language for accounting is COBOL, which is implemented much more widely in municipalities than, say, PL/1 or even FORTRAN. Most models, however, are developed using a "scientific language," such as FORTRAN, PL/1, or ALGOL. Thus, the potential for a mismatch is apparent. Also, local computers often have severe core storage limitations. However, with the wide accessibility of commercial time-sharing computer services, the limitations of local computers should become less of a constraint in the future. For example, the PCAM program (Patrol Car Allocation Model), which was an outgrowth of two of the mini cases reported earlier, is supported on a commercial time-sharing system [CHAI 78].

A lack of computer experience on the part of the user may imply that more time and effort is necessary to "conquer the learning curve" of the model, as discussed in Section 8.3.

8.2.5 Model Adaptability

The model-related issues we have discussed to this point are all rather concrete and relatively easy to address. Other, more subjective model attributes are also important, however. For an agency considering implementation of an existing model, a critical question becomes: Is this model adaptable to the particular agency's needs? For any agency contemplating long-term use of a model, an analogous question can be stated: Is this model flexible in ongoing use? This second question relates to whether the model can be adjusted to depict changes in operating policies that are bound to occur from time to time.

The importance of adaptability (generalizability) in transfer has already been discussed in the context of the $M/M/N$ priority queueing model for police deployment. One city's standard operating procedure may be entirely unacceptable in another city. Even apparently subtle changes in operating policy can violate the theoretical underpinnings of an existing model (as any reader of Chapters 3–7 should be aware of by now!).

When trying to circumvent the transferability problem, one must naturally ask: "How much generality can I afford?" If a model is being funded internally, the model builder has no obligation to incorporate generalizability beyond that needed for the client city. If the model is being funded by a federal agency, more generality is certainly called for. However, even then there are limits, including those posed by analytical tractability as well as by costs.

Regarding ongoing use in a single jurisdiction, an agency's operating rules are likely to change over time. A prioritized response service may radically change its prioritization structure. A school district may split a 6-year high school into a 3-year junior high and a 3-year senior high. A bus company may introduce new types of vehicles. The question to the modeler is: To what extent will the model still be useful under various conceivable operating changes? The same comments made above about generalizability

apply here. Surely no model should be so rigid that the most minor change in an agency's operating policies makes the model obsolete. The professional model builder should incorporate sufficient flexibility into the model design to allow for reasonable changes in operating rules that do not violate the cost constraints of model development and implementation.

8.2.6 Exploring Alternatives

As a decision-making tool, a model is supposed to be a relatively accessible instrument with which to explore the consequences of alternative operating policies. Its data base and decision variables should be easy to change, its outputs easy to understand, and its use in systematic exploration of the "decision space" should be simple.

In attempting to achieve these laudable goals for a model, the first question to ask is simply this: Does the model have *understandable outputs*? As straightforward as this question may appear, output formats are quite important to the user. For instance, we have found the incorporation of a user "glossary" as an input option to be very important. Such a glossary allows the user to identify the model with his or her agency, with agency and city specific names for geographical atoms, vehicles, response areas or districts, routes, and so on. Seeing "Times Square" on the output is much more meaningful to a New York City user than "geographical atom 12." Also, simple word choices such as "average" rather than "mean" can be important. A detailed model may be improved by producing different levels of aggregation for outputs, from highly aggregate (city-wide) to highly disaggregate (neighborhood-oriented) outputs.

To assist in modifying operating policies and other input data, a *natural language interface* linking the user to the model is often appropriate. Such an interface is a computer program that "talks with" the user in his or her language, while being capable of altering all input data in the model. It is a model/user "go-between." It assists in adding units, redeploying units, rescheduling, reprioritizing, redistributing demand, and so on, none of which need cause headaches associated with data-card changes. Rather, the user can express his or her desire for an alternative policy in language with which he or she is most comfortable and the interface program does the tedious work of changing data cards. This on-line procedure can even be used in submitting batch-processing production runs.

Once data changes can be readily made, we must concern ourselves with the user's systematic exploration of alternative policies. Such user-directed exploration would not be required if he or she could use an optimization algorithm to compute "the unique answer." It is our view that few urban services can directly implement the results of an optimization model. This is due to the complexity and multifacetedness of most urban services, where objectives and constraints can rarely be uncontestedly expressed in mathematical form. Rather, it is felt that a user with an intimate knowledge of his or her own city, can be an excellent judge of the qualitative factors that are relevant. Using the computer to calculate the important performance measures of proposed alternative operating policies, model usage becomes an iterative process. First, the user proposes a particular operating policy (e.g., deployment of resources) and has the computer model calculate the resulting values of performance measures. The user then incorporates this evidence, including possible workload imbalances among workers and/or inequities in citizen accessibility to service, with the remainder of his or her knowledge of the area under consideration, and decides whether to accept the proposed operating policy or to devise an altered one. In the latter case, the entire process is repeated one, two, or several times until a satisfactory operating policy is obtained. In this way, good use is made of the user's talents and the computer's computational power.

8.3 ISSUES RELATED TO PEOPLE AND INSTITUTIONS

Although a model and its attributes are important in any intervention program, the ultimate success or failure of a program depends on the people and institutions involved. In this section we attempt to describe the important issues related to the user, the analyst, the individuals who would be affected by any implemented change, and the broader institutional setting.

8.3.1 The User

The user may in fact be several individuals, each having quite different attributes. For convenience in the following we shall refer to the singular user, recognizing that several individuals may be involved.

The user enters any implementation program with a given level of commitment, educational and experiential background, and degree of authority. His or her serious dedication to the particular problem area under study is essential for success of the program. The concern of the school superintendent in the school-busing example can be contrasted with the lack of concern of the police commissioner ("Why spend time and money remaking already adequate decisions?"). Implementation has always been much more likely in a situation in which the agency steps forward and announces a problem in need of a solution (thus reflecting commitment), rather than one in which an analyst goes to an agency with a problem solution looking for a needy user. (The latter mistake was made in the "bus clumping" mini case.)

The *training and technical background of the user(s)* is obviously important. In a technical sense, one needs to know the appropriate level of complexity for the model/user interface. How much can the user(s) be expected to know about statistics, probabilities, means, modes, medians, variances, and so on? At a more subtle level, it has been the experience of the authors that agency personnel often tend to think of their agency's operation in clinical terms (i.e., on a case-by-case basis) rather than conceptually or in terms of statistical patterns. Case recollection is often based on extremes, not averages. "I remember during the flood of 1954, when . . . ," the scenario might go. Thus, the whole concept of *models*, with their implied statistical regularities and their inability to incorporate all possible extreme events, is foreign to many urban decision makers. A model's implementor must be sensitive to this issue and attempt, through simple examples, to convey the meaning of models, together with their uses and their limitations.

The *position and power of user(s)* clearly influence the likelihood of eventual implementation. A limited-duty sergeant is not adequate as the sole contact point for a massive patrol redeployment effort. Even a high-ranking official, operating alone, may not be a sufficient resource to promote implementation.

8.3.2 The Analyst

The urban analyst assists in the technical details of model selection and development, data collection, computer implementation and usage, and—most important—is responsible for training the user(s). His or her attributes are just as important as the user's. Concerns with educational and experiential background, commitment, and so on, parallel directly those of the user.

Administratively, the analyst is typically under contract to the user's agency or is utilized as a consultant on an "as-needed" basis.

Is an analyst needed in a model transfer situation? The user may be tempted to "go it alone" when attempting to implement a model transferred from another jurisdiction. Can the user play the role of analyst in such cases? Eventually, this may become standard procedure. However, given the current state of urban modeling and implementation, we feel that such a user should have available one or more persons ("analysts") to provide *technical assistance*. The assistance may come from the original model builder, a local university, or even a professional consultant. Whatever the source, it is clear that any user will probably encounter a number of problems that have to be dealt with in novel ways, ranging from computer compatibility with the original program, to statistical data-collection issues, to questions pertaining to operational complexities of his or her own system. Rare is the model that survives (and thrives) in a public institution without the presence of technical assistance.

8.3.3 The User's Agency

The institutional setting in which an intervention program takes place can suggest alternative implementation strategies. An agency that is accustomed to data collection and its use in management decision making can be contrasted to one in which decisions are solely "intuitive," based on judgment, experience, and pressures from citizens' groups. Some agencies are highly professional, utilizing standards of accountability; other are fraternal, often insular, choosing to do things based on historical tradition. Particularly for the more custom-bound agencies, considerable ingenuity is required of the analyst to arrive at an appropriate implementation plan. A major component of almost any plan should be building a broad base of support for the effort, as discussed below (see Section 8.3.5).

All agencies must operate within various constraints, including those related to time, budget, and operations.

A *time constraint* is a deadline of some type. It may refer to the fiscal year, calendar year, contract obligations, elections, or reappointments. Most urban decision makers operate under tight time constraints. Many, it seems, always work in a near-crisis environment. Thus, few city-specific studies have available the time that is often desired by the model implementors. Compromises must be reached. The realities of deadlines must be faced. There are too many models now lying peacefully in the graveyard of models which consumed 80 to 90 percent of a project's efforts, leaving all too little time for the remaining processes (e.g., user training, analysis of production runs, instituting operating changes) that must be undertaken in a successful model implementation effort.

Budgeting constraints often limit the duration of an intervention program and the types of analyses that can be performed. On the positive side, limited budgets also imply ceilings on labor pools, which should enhance in the future urban managers' interest in effective utilization of existing or projected personnel.

A system's *operating constraints* limit the range of alternatives that may be examined in a model-based analysis. For instance, the State of New York had on its books for over 50 years a "three-platoon law" that limited the scheduling flexibility of police personnel. In general, the constraints with which one deals may be legal, financial, political, or technical. The highly visible constraints, such as state laws or collective bargaining agreements, are relatively easy to incorporate into an analysis. The difficult ones are invisible at first (remember the star halfback who was initially assigned to a rival school?). In an urban environment, it is our contention that virtually no change of significance can occur without some hitherto invisible constraints rising to the surface. No model can be expected to incorporate awareness of the existence of star halfbacks; rather, the model implementation *process*

should be designed to anticipate and correct for such hidden constraints when they come to light.

Finally, an institution can be considered to have *natural time constants*, reflecting typical lengths of time required for it to accept and incorporate an attempted change. A fairly simple change (e.g., a rescheduling of personnel) may be adopted fully within days or weeks. A more fundamental change, perhaps markedly affecting the institution, may require a substantial turnover in current personnel before full acceptance. This would be necessary to diminish the importance of "institutional memories," which are often supportive of traditional policies. Here, we find it convenient to think of an agency's institutional time constant as the *time to achieve a 50 percent turnover in personnel*. For a system operating in equilibrium with, say, a 25-year retirement plan, its institutional time constant would be $(25/2) = 12.5$ years (Figure 8.5). This is a sobering thought for those who expect instant cures to various urban problems.

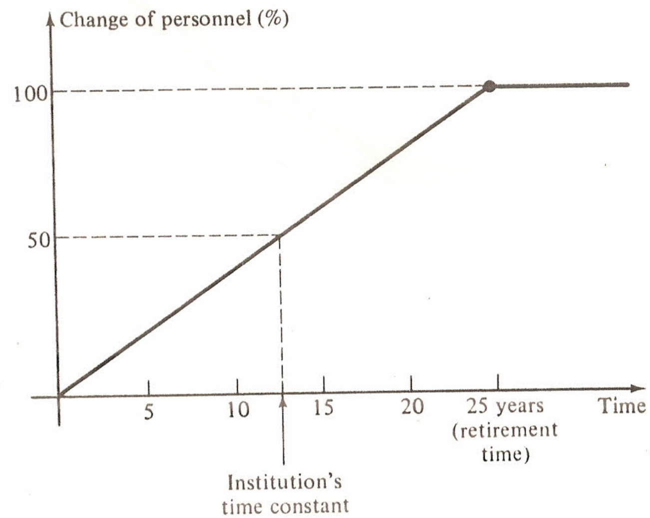

FIGURE 8.5 Institutional time constant.

8.3.4 Other Groups and Agencies

The setting of an intervention program goes beyond the user, the analyst, and the user's institution. It includes other groups and agencies likely to be affected by any proposed change. An early understanding and appreciation of this broader environment can markedly increase the chance of successful, widely supported implementation. The individuals who will be affected by a change include citizens who are consumers of the agency's services, agency workers, agency managers, and perhaps other individuals in related agencies.

The reassignment of schoolchildren to another school will affect the children's friendship networks and their parents' time investment in the previous school; even the rescheduling of school buses can upset daily life patterns. These are legitimate concerns that all too often have been ignored by narrowly focused operations research analysts trying to solve the problem by implementing the latest "transportation" or "assignment problem" computer code. The closing down of a police station, library, or "little city hall" can generate similar concerns.

Agency workers have designed their lives around current operations; any change will upset these patterns. Thus, a rescheduling of personnel or a spatial reassignment will affect a worker's temporal routine and/or commuting patterns.

Sometimes the investment in the status quo runs deep, as in a case concerning fire department personnel in New York City. Here, the New York City Rand Institute recommended an adaptive response strategy that would have resulted in relatively underutilized fire companies responding more frequently into high-alarm districts. This attempt to balance workloads, although quite rational from the point of view of the models of this book, met with stiff opposition from the firefighters' de facto union. They argued, in effect, that assignment to a light-alarm-load fire company had become a seniority right, and to change this pattern would upset the long-sought-for seniority privileges of the firefighters. Thus, the status quo had the younger, less-experienced firefighters in the high-alarm-rate districts.

Related outside agencies and groups should also be incorporated into an intervention program. For instance, a new public ambulance system should interface with private ambulance companies. A revised or initial implementation of 911 as a city's emergency number requires cooperative interaction among police, fire, and emergency medical personnel. The elimination ("consolidation") of a police precinct or a local post office will concern citizens' groups, who should be consulted prior to such a move. Development of a geographically oriented data base, perhaps reflecting partitioning into small cells or geographical atoms, should be done in concert with other city agencies that have already completed or that contemplate such a task. Lack of concern for such interaction can create opponents to a new program, much in the same manner as opponents can surface from within a user's agency.

8.3.5 Model Constituency

Because of all the diverse interests that can surface in an intervention program, it is necessary that any such program have a broad support base or constituency. In a program that is strongly model-oriented, one must build a constituency for the model itself. Such a constituency comprises the model users and other individuals and groups that support the model implementation effort.

Building a constituency for a model is often difficult. The urban analyst is likely to feel most comfortable interacting with a "bilingual person" in the agency, one who speaks "model language" as well as "agency language." Such people are rare and highly valued. They are professionally mobile—apt to be promoted or to relocate voluntarily to another agency. An urban analyst who builds his or her implementation process solely through this one person surely has manufactured a weak link in the implementation process. Any program that requires more than, say, one year to complete, runs a high risk of losing such a person, resulting in a "vanishing advocate" [CHAI 75]. (Recall the Boston police simulation mini case study.)

Developing a broader constituency by holding workshops and training sessions and soliciting feedback has other important benefits. Agency personnel, through feedback, often provide inputs that can turn an otherwise naive modeling study into a realistic and sound analysis. Such feedback can relate to safety factors (e.g., permissible U-turns for school buses[8]), concerns over supervision (e.g., as in the swapping of police patrol officers in southern Manhattan), legal issues (e.g., in the selection of a jury pool), correct interpretation of recorded data (e.g., as in the accuracy of recorded travel times), and probable response of rank-and-file workers to a proposed innovation (e.g., to a redeployment of personnel). The process of eliciting feedback can transform otherwise skeptical, perhaps even hostile, personnel into fully participating supportive allies of the study. Supervisory personnel, as well as in-the-field workers, can become advocates. A broad mutual investment in a program is not likely to lead to subversion during implementation. And a model either created or at least modified in the presence of feedback from personnel is more likely to deserve broad-based support.

The needs and priorities of the user and other managers should also be considered when building a constituency. Key decision makers, including the user, have short- and long-range goals that can have a beneficial or detrimental effect on an implementation effort. Remember the commissioner of education whose short-range goal was demonstrating frugality, leading to a long-range goal (passing a multi-million-dollar bond issue)? Or a long-range goal of an urban official may be reappointment or reelection, and an urban modeling analysis may be viewed positively or negatively in this light. Some of these goals may be apparent to the urban analyst; many come to light only after the fact. A common mistake here is for the analyst to focus on long-range projects while ignoring short-term problems. Public decision

[8]The students performing the school busing study interviewed bus drivers to acquire this information.

makers are often under pressure to "put out fires," and analytic assistance is usually very much appreciated. Short-term assistance, even if it appears to delay progress in long-range projects, has two benefits: (1) it helps the decision maker address the current problem, and (2) perhaps more fundamentally, it establishes credibility for the analyst in the eyes of the decision maker, thereby enhancing his interest in the long-range projects.

A strong model constituency is no guarantee against the vanishing advocate syndrome. Sometimes a modeling effort is associated with a particular administrative regime; viewed as an instrument of reform, it develops political attributes. Should the current regime—the one supporting the modeling effort—be replaced, the effort could be discarded because of its links with that regime. However, any innovation viewed as a reform runs similar risks. The development of a wide base of support minimizes this risk.

8.3.6 Training: Conquering the Learning Curve

The analyst usually has the important responsibility of training the user. With the analyst's assistance and guidance, the user must invest time and effort to become comfortable with and knowledgeable about the model and its use. Psychologists model such a process with a learning curve, which starts low at a point of relative ignorance, often climbs steeply during the first phases of learning, and then tends to a gradually decreasing positive slope as the learner asymptotically approaches his final level of understanding[9] (Figure 8.6). The urban analyst should be aware that the model he or she developed and is attempting to implement represents a complex set of processes that even other analysts might have difficulty fully grasping. Thus, an urban decision maker, usually trained by case study (often on the job) and utilizing ad hoc techniques that have been proven historically at least not to fail, can be expected to have difficulty with models and their uses. The urban analyst should be committed to helping the user climb his or her own learning curve, beyond naive acceptance of the model, to full understanding of the uses and potential abuses of the model.

One technique that we have found useful in the learning (training) process is to develop simple examples, augmented by "rules of thumb," to aid insight and intuition. The example of Chapter 1, comparing deterministic and probabilistic reasoning for a single police patrol beat, illustrates several concepts that reappear in more complex settings. Rules of thumb include "square-root laws" for travel distances in an urban setting, which have the additional advantage of diminishing the importance of thinking purely in

[9]See, for example, L. L. Thurstone, "The Learning Curve Equation," *Psychology Monograph, 26* (114), 11–13 (1919); reprinted in *Mathematics and Psychology*, G. A. Miller, ed., Wiley, New York, 1964, pp. 128–130.

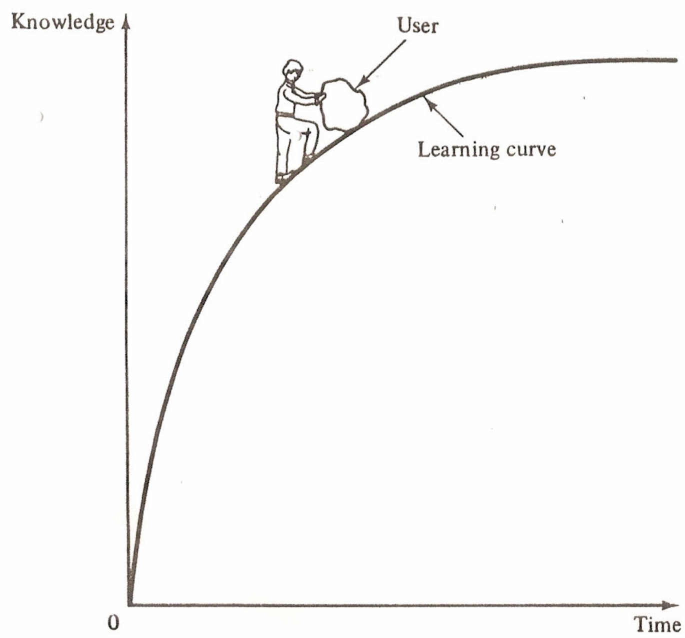

FIGURE 8.6 User conquering the learning curve.

terms of linear relationships. Simple queueing formulas also demonstrate the highly nonlinear performance of probabilistic systems. From Chapter 5 we know that the fraction of dispatch assignments that are *interresponse area assignments* is at least as great as the average system utilization factor. This can be argued intuitively in a way which (1) gives the user another rule of thumb and (2) helps him or her develop more insight into the probabilistic nature of operations. We have found such approaches invaluable in explaining the hypercube model to users [LARS 72b]. The New York City Rand Fire Project utilized this approach with simulation models [IGNA 75], and now mathematical programming enthusiasts are arguing for a similar approach [GEOF 76].

When learning to operate the model, the user must eventually confront the problem of specifying *performance standards* that he or she wishes to achieve by utilizing the results of the model. A user's ability to articulate performance standards seems to vary with the extent to which the model's performance measures capture an agency's true objectives and with the user's understanding of these measures. Some complicated weighted formula, even if derived from the latest treatise on decision theory, may be opaque to a user. In general, simple performance measures that arise naturally in ongoing operations are preferable to fabricated ones.

The analyst, when acting as trainer, may discover that the user is reluctant to articulate performance standards. Most operating urban systems reflect an evolution of decisions made in response to feedback from citizens, workers, and managers. The current picture, the status quo, is the net result of all these historical forces. Recall the police commissioner who at first refused to specify the T and P values for his 911 emergency system and who finally settled on $T = P = 0$. To articulate clearly the performance objectives of the system would almost certainly reveal the inadequacy of the status quo to some groups. To be sure, the present state of affairs has its own *implicit* performance standards. But discomfort with these standards, and hence possible disruption of the status quo, are more likely to occur only when they are made explicit. (Somewhat ironically, it is the stark political neutrality of a model, with its ever-visible performance measures, that makes it, in a sense, political.) If the user is unwilling to specify performance measures, the analyst/trainer can suggest exploring policies that maintain current manpower levels (as was done for the 911 study) or otherwise satisfy constraints with which the user is comfortable; however, his or her selection of a final policy will still require some consideration of achievable performance standards.

8.3.7 Politics and Improper Reasons for Analysis

Occasionally, an urban analyst may be brought into a situation for the wrong reason. As one example, any decision, program, or operational change supported by operational analysis can, for city and agency officials, help to create an *improved public image* as a result of the presence of technical support. Although this visibility aspect of a program can be positive, creating the desire for bona fide commitment from officials, it can also be detrimental if it creates an atmosphere in which predetermined policies are to be legitimized by analysis. Thus, when officials sense the *necessity for technically "proving" an already selected policy*, they may "call in the analysts."

Another example is that "putting a problem out to analysis" is a sophisticated way of tabling discussion and otherwise delaying action on an issue. These and other concerns are discussed at length by Brewer [BREW 73].

8.4 CONCLUDING REMARKS

We do not intend to discourage prospective urban analysts by focusing on difficult issues in this chapter. At the same time, we do not want to help train eager analysts who are unaware of the difficulties of the implementation process or the limits of quantitative methods. Analytical competence must be balanced with mature judgment. The uncritical application of any of the methods of this book will almost surely lead to frustration and failure. It has been our observation that modelers who have pursued implementation with the same energy as they applied to modeling have been successful in the urban scene.

As mentioned in Chapter 1, increasingly severe financial constraints are confronting cities throughout the industrialized world. Yet urban residents are demanding more services, in type, quantity, and quality. Hence, administrators of urban services are being forced to become managers, weighing the relative costs and benefits of alternative allocations of resources. They are having to spend time and money making, hopefully, improved decisions. The significant returns on investment of competent studies are now becoming apparent to citizens as well as to agency administrators. Thus, the demand is increasing for quantitative tools to assist in decision making. The methods of this book provide a beginning. The coming years should be exciting ones, both for methodology development and for the intelligent implementation of urban operations research.

References

[BRAN 78] BRANDEAU, M., AND R. C. LARSON, *Implementing the Hypercube Queueing Model to Plan Ambulance Districts in Boston*, Report PPR 7804, Public Systems Evaluation, Inc., Cambridge, Mass., September 1978.

[BREW 73] BREWER, G. D., *Politicians, Bureaucrats and the Consultant: A Critique of Urban Problem Solving*, Basic Books, New York, 1973.

[CHAI 75] CHAIKEN, J., T. CRABILL, L. HOLLIDAY, D. JAQUETTE, M. LAWLESS, AND E. QUADE, *Criminal Justice Models: An Overview*, Rand Corporation, R-1859-DOJ, Santa Monica, Calif., 1975.

[CHAI 78] CHAIKEN, J., "Transfer of Emergency Sevice Deployment Models to Operating Agencies," *Management Science*, *24*, 719–731 (March 1978).

[DAEL 76] DAELLENBACK, H. G., "Note on a Survey on Evaluation of Operations Research Projects," *Interfaces*, *6* (2) (February 1976).

[FERR 78] Ferreira, J. M, Jr., "Comparing Patrol Unit Allocation Methods," in *Police Deployment: New Tools for Planners*, R. C. Larson, ed., Lexington Books, Lexington, Mass. 1978.

[GEOF 76] Geoffrion, A. M., "The Purpose of Mathematical Programming Is Insight, Not Numbers," *Interfaces*, 7 (1), (81–92) (November 1976).

[HEBE 78a] Hebert, S., "The Introduction of Sophisticated Allocation Techniques in the Boston Police Department," in *Police Computer Technology*, K. W. Colton, ed., Lexington Books, Lexington, Mass., 1978, pp. 97–110.

[HEBE 78b] Hebert, S., and K. W. Colton, "The Use of Resource-Allocation Models in the Los Angeles Police Department," in *Police Computer Technology*, K. W. Colton, ed., Lexington Books, Lexington, Mass., 1978, pp. 111–138.

[HOEY 73] Hoey, J. M., *Planning for an Effective Hospital Administered Emergency Ambulance System in the City of Boston*, IRP Report TR-01-73, MIT Operations Research Center, Cambridge, Mass., September 1973.

[IGNA 75] Ignall, E. J., P. Kolesar, and W. E. Walker, *Using Simulation to Develop and Validate Analytical Emergency Service Deployment Models*, P-5463, Rand Corporation, Santa Monica, Calif., August 1975.

[KEEN 76] Keeney, R. L., and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, Wiley, New York, 1976.

[LARS 72a] Larson, R. C., "Improving the Effectiveness of New York City's 911," in *Analysis of Public Systems*, A. Drake, R. Keeney, and P. Morse, eds., MIT Press, Cambridge, Mass., 1972.

[LARS 72b] Larson, R. C., *Urban Police Patrol Analysis*, MIT Press, Cambridge, Mass., 1972.

[LEVI 78] Levi, M., "And the Beat Goes On: Patrolmen's Unionism in New York City," in *Police Accountability*, R. C. Larson, ed., Lexington Books, Lexington, Mass., 1978.

[WALK 78] Walker, W., "Changing Fire Company Location: Fire Implementation Case Studies," Document 023-000-00456-9, U.S. Government Printing Office, Washington, D.C., 1978.

General References

The following books and papers contain case studies, theories of implementation, and other information useful to those seeking to learn more about the implementation process.

Baer, W. S., L. L. Johnson, and E. W. Merrow, "Analysis of Federally Funded Demonstration Projects: Executive Summary," *Final Reports, Supporting Case Studies*, Reports R-1925, 1926, 1927-DOC, Rand Corporation, Santa Monica, Calif., 1976.

Beltrami, E., *Models for Public Systems Analysis*, Academic Press, New York, 1977.

Blumstein, A., "Operations Research in the Public Sector: Frustrations, Failings and Future," invited paper presented at the Eighth International Conference of the International Federation of Operation Research Societies (IFORS), Toronto, June 19, 1978.

Brewer, G. D., *Politicians, Bureaucrats and the Consultant: A Critique of Urban Problem Solving*, Basic Books, New York, 1973.

Brewer, G. D. and M. Shubik, *The War Game: A Critique of Military Problem Solving*, Harvard University Press, Cambridge, Mass., 1979.

Colton, K. W., ed., *Police Computer Technology*, Lexington Books, Lexington, Mass., 1978.

Drake, A. W., "Quantitative Models in Public Administration: Some Educational Needs," in *Analysis of Public Systems*, A. W. Drake, R. L. Keeney, and P. M. Morse, eds., MIT Press, Cambridge, Mass., 1972.

Fromm, G. W. Hamilton, and D. E. Hamilton, "Federally Supported Mathematical Models: Survey and Analysis," Document 038-000-0021-0, U.S. Government Printing Office, Washington, D.C., June 1974.

Gass, S. I., "Evaluation of Complex Models," *Computers and Operations Research*, 4, 27–35 (1977).

Gass, S. I., and J. Dawson, "An Evaluation of Policy-Related Research: Reviews and Critical Discussions of Policy-Related Research in the Field of Police Protection," prepared for the National Science Foundation by Mathematica, Inc., 1974.

Gass, S. I., and R. L. Sisson, eds., *A Guide to Models in Governmental Planning and Operations*, Sauger Books, Washington, D.C., 1975.

Gass, S. I., et al., *A Study for Assessing Ways to Improve the Utility of Large-Scale Models*, Control Analysis Corp., Palo Alto, Calif., December 1978.