Parameter Estimation in System Dynamics Modeling

7

Alan K. Graham

7.1 Introduction

System dynamics modeling offers an attractive tool for policy evaluation. Policy alternatives can be simulated by computer; thus the policies can be tested under completely controlled conditions. A special programming language, DYNAMO, is widely available and simplifies the mechanics of computer simulation. But simulations require a model, and model building is still an art in many respects. There is no step-by-step procedure that automatically produces a useful model. However, as more and more models are created, some steps in the modeling process have become clearer, and procedures can now be formalized to some extent. This paper lays out the main techniques of one important step in system dynamics modeling: the estimation of parameters. It is assumed that the reader is familiar with system dynamics and the DYNAMO language.

The appropriateness of the various parameter estimation techniques depends upon the entire context of the model-building process. For example, parameter estimation is closely dependent on equation formulation and model testing. Graham (1978) discusses these interdependencies in detail, so they will be touched on only briefly here. This chapter is a taxonomy of estimation techniques; the types of data, assumptions, and procedures that characterize each technique are specified. The purpose is to help the reader choose an appropriate technique and to avoid pitfalls.

Throughout the exposition of techniques, examples are drawn from a simple model of urban housing, which portrays the growth and maturity of an urban residential area. The appendix gives the complete equations for the model, which is based on a model described in Alfeld and Graham (1976, ch. 6).

7.2 Estimation Using Data below the Level of Aggregation of Model Variables

Most parameters in system dynamics studies are estimated on the basis of descriptive information obtained from participants in the system being modeled. That information is more detailed than data that corresponds directly to model variables. There are two kinds of model variables: level variables, which aggregate a collection of items like houses into a level of houses, and rate variables, which aggregate a stream of events like the construction of a house into a rate of housing construction. By contrast, data below the level of aggregation of model variables characterize the individual members of a level, or the individual events within a rate. For brevity, such data will sometimes be referred to as "unaggregate data."

For an example of setting parameters using unaggregate data, consider an equation representing housing demolition:

```
HD.K=H.K/HL                                         1, R
HL=66                                               1.1, C
     HD      - HOUSING DEMOLITION (HOUSES/YEAR)
     H       - HOUSES (HOUSES)
     HL      - HOUSING UNIT LIFETIME (YEARS)
```

The model assumes a constant average housing unit lifetime, HL, so that every year, 1/HL of the houses, H, are demolished. HL can be estimated in many ways from unaggregate data; in none is the equation used in computing the estimate. Equation (1) is used only to define the function of HL in the model. One time-consuming approach to estimating HL would be to survey a number of houses that have been demolished, take their ages at the time of demolition, and average those ages to determine an average housing unit lifetime, HL. The information used to set the parameter concerns individual houses and their demolition, which the model variables aggregate into a level of houses, H, and their outflow rate of housing demolition, HD. Thus the data come from below the level of aggregation of model variables.

As another means of obtaining information about the lifetime of houses, the modeler can examine the ages of existing houses, and observe the age at which very few houses remain standing. Or the modeler can consult someone who has observed construction and demolition of houses closely, and ask that person how long houses typically last (or survive neighborhood changes or eminent domain proceedings.) Indeed, asking experts questions of wider scope is the basis of the popular Delphi method (Turoff, 1970). The modeler can sometimes obtain descriptive

histories of particular neighborhoods that chronicle the successive waves of demolition and construction, and thus gain some idea of how long the previous houses in that neighborhood lasted. The modeler can also call upon his own experiences in observing the physical decay of one or more houses, and extrapolate to estimate how long a house will last. Finally, lacking any better information, the modeler can take two extreme values that are clearly too large and too small and pick a value somewhere between them. Modelers seldom rely on this last technique, since it is based on the implicit assumption that in the entire world there is no available information that would lead the modeler to a single estimate of the parameter in question.

These examples suggest that the sources for data below the level of aggregation of model variables are numerous and diverse. The data for the lifetime of houses came from city hall records, a history book, expert testimony, and the modeler's own day-to-day experiences. In fact, *all* factual knowledge about a system—records, books, eyewitnesses, and personal experience—falls into the category of unaggregate data; the only exception is collected statistics corresponding to model variables. Thus unaggregate data are by far the most abundant source of knowledge about real systems.

Table Functions

We have shown six different ways to estimate the average housing unit lifetime using unaggregate data. All six are straightforward. This section describes a technique for estimating the 5 to 15 numbers that typically specify a table function. This may seem to be a formidable estimation problem, but it can be broken into subproblems: estimating the value and the slope of the function at one extreme, at the normal value, and at the other extreme, and connecting those known values and slopes with a smooth curve. Once these four subproblems are solved, the table function is known to within a narrow range of values.

For example, consider a group of equations that represent the effect of land availability on housing construction. The rate of housing construction, HC, is proportional to the number of houses, H, already within the urban area; thus when other things remain the same, more houses, more infrastructure, and more people create a larger market for new housing construction. The parameter that gives the constant of proportionality is the housing construction normal, HCN. The rate of housing construction, HC, is modulated by land availability; this modulation is accomplished by the housing-land multiplier, HLM:

```
HC.K=H.K*HCN*HLM.K                                        2, R
HCN=0.07                                                  2.1, C
    HC     - HOUSING CONSTRUCTION (HOUSES/YEAR)
    H      - HOUSES (HOUSES)
    HCN    - HOUSING CONSTRUCTION NORMAL (FRACTION/YEAR)
    HLM    - HOUSING LAND MULTIPLIER (DIMENSIONLESS)
```

The housing-land multiplier, HLM, responds to land availability, which is quantified by the land fraction occupied, LFO:

```
HLM.K=TABLE(HLMT,LFO.K,0,1,.1)                            3, A
HLMT=.8/.95/1.075/1.2/1.3/1.35/1.35/1.25/1/.6/0           3.1, T
    HLM    - HOUSING LAND MULTIPLIER (DIMENSIONLESS)
    TABLE  - TABLE INTERPOLATION FUNCTION
    HLMT   - HOUSING LAND MULTIPLIER TABLE
    LFO    - LAND FRACTION OCCUPIED (DIMENSIONLESS)
```

The graph for HLM is shown in figure 7.1. To estimate the table function for HLM, first consider the extreme condition of zero land occupancy, where incentives for construction should be less intense than with higher occupancy. When the land fraction occupied, LFO, approaches 0 (near the left side of the curve in figure 7.1), most of the area being modeled is vacant land. The area's viability as a future city has not yet been demonstrated. Developers cannot rely on continuing demand for the housing units they construct. Also, services taken for granted in more heavily settled areas must be installed in each successive new neighbor-

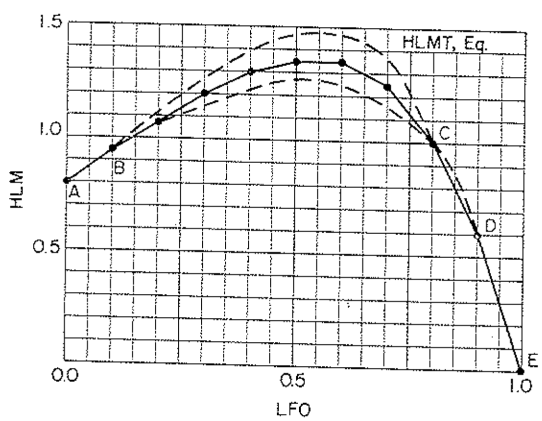

Figure 7.1
Housing-land multiplier table, equation (4)

hood: roads, sewers, electricity, gas, schools, and public transportation. Sparsely settled areas often cannot make city water or sewers (let alone public transportation) an economical proposition. So when the land fraction occupied LFO equals 0 the housing-land multiplier, HLM, should be lower than at most other values of LFO (point *A* on figure 7.1 uses the value of 0.8 for HLM).

Adding housing units to a sparsely settled area encourages more urban services. Urban services demonstrate the area's viability, and make housing construction more profitable. But adding a few houses cannot pay for the infrastructure—schools, roads, libraries, and utilities—necessary to deliver a complete ensemble of urban services. The curve for the housing-land multiplier table, HLMT, should slope upward, but not very steeply from where LFO equals 0 (see the line segment between points *A* and *B* in figure 7.1).

Now consider the normal condition. Equation (2) defines the rate of housing construction, HC, so that it occurs in normal proportion to houses (specified by housing construction normal, HCN) when the housing land multiplier, HLM, equals 1.0. For consistency, HLM must equal 1.0 at whatever value of the land fraction occupied LFO is defined as normal. Normal conditions are defined to occur near the end of the area's growth when land availability begins to constrain further construction. In this model, the normal value of LFO is defined to be 80 percent land occupancy (point C in figure 7.1).

At the normal condition, land availability begins to constrain housing construction. Thus the table function must have a negative slope at the normal point, so that diminishing land availability in the model likewise begins to constrain further construction in the model. The negative slope at the normal point in turn implies that the table function must exceed a value of 1.0 just under the normal value of land fraction occupied, LFO.

Now consider the other extreme condition in which the land fraction occupied, LFO, equals 1.0. The land area within the city or district being modeled is totally occupied; even the least desirable sites have been built upon. Regardless of the incentives to construct, no housing can be constructed within the area being modeled until there is some physical space available upon which to build—that is, until LFO is less than 1.0. So the housing-land multiplier, HLM, should equal 0 when LFO equals 1.0, which establishes point *E* in figure 7.1.

If the land fraction occupied, LFO, was not 1.0 but close to 1.0 (nearly full land occupancy), urban services like sidewalks, schools, libraries, roads, and public transportation would be fully developed. To be sure, the crowding and lack of desirable construction sites implied by an LFO close to 1.0 would not cause housing construction to take place as rapidly as under normal conditions. Nonetheless, any small reduction of LFO from 1.0 opens up the possibility of appreciable housing construction. Therefore the curve of the housing-land multiplier should have a steep slope as LFO approaches 1.0 (see the line segment between points D and E on figure 7.1).

So far, values have been estimated at and near two extreme conditions and at the normal condition. Now all that remains is to draw a curve through the estimated points. Any sharply bent or kinked curve is probably not realistic. A bend or kink implies something special about the exact conditions at which the bend or kink occurs. Since the housing-land multiplier table, HLMT, represents many phenomena (prices, availability, infrastructure, and so forth), the probability is small that all phenomena would show major changes at a unique set of conditions. Accordingly, the curve for HLMT (and in general for all highly aggregated relationships) should change smoothly.

Solving the subproblems of extreme and normal conditions, and connecting known points with smooth curves, allows the modeler to estimate a nonlinear table function with confidence. The dotted lines show alternate curves for HLMT that also satisfy the constraints for extreme values, slopes, and smooth curvature; these constraints allow little latitude in specifying HLMT. The estimated table thus summarizes observations of a large number of processes below the level of aggregation of model structure. HLMT is the aggregate representation of these processes and their effect on housing construction.

Ad Hoc Computation

It sometimes happens that no unaggregated data are available that correspond directly to the parameter being estimated. In the housing example, this would happen if no information were available about the ages of houses at demolition. However, other unaggregated data may exist that describe the process of demolition, and these data may suffice to infer an estimate of the housing unit lifetime, HL. This is an ad hoc computation used to estimate a parameter from unaggregate data.

For example, suppose the modeler walks through a district and observes that of a fairly homogeneous mix of housing units of different ages and types about one in a thousand is in the process of being demolished. Suppose also that conversation with a wrecking crew reveals that it takes about three weeks to demolish a single building when neighboring buildings are left intact. From this information, the fraction of the housing stock demolished each year can be computed: it takes three weeks to demolish a building, or $(3/52) = .0577$ years. Thus, if .001 of the housing stock is observed to be in the process of demolition, and if that rate of demolition continues for a whole year, the fraction of the housing stock demolished each year is $(.001/.0577) = .01733$. This fractional demolition rate (whose unit of measure is per year) is the reciprocal of the housing unit lifetime, HL (whose unit of measure is years). Thus HL is $(1/.01733)$ or 57.7 years, which is not an unreasonable estimate.

This example suggests that parameters can be estimated by ad hoc computations based upon readily available unaggregate data. These computations can take many forms; see Senge (1975), for example.

Pitfalls

The greatest single pitfall in using unaggregated data lies in formulating a model structure and parameters that are aggregated to the point where the processes characterized by the parameter values cannot be reliably observed. As a result the parameters have little real-life meaning, and to estimate them, the estimator must draw conclusions based on a mental model of the behavior of the system, rather than simply reporting observations. For example, in the housing model a variety of processes determine how long it takes the system to make a transition from growth to full land occupancy—incentives to construct housing, supply and demand effects in the land market, and housing depreciation, for instance. In the model, several parameters characterize such diverse processes. An alternative formulation of the model might have contained a single parameter that specified the time constant for the transition from growth to equilibrium. Urban experts may well be willing to estimate such a quantity, but the estimate would be a conclusion or opinion drawn from their mental models of how the cities behave, rather than a report on a single cause and effect relation within the city.

Another instance of confusing conclusions with observations might occur if the simple housing model described here is expanded to include demand for housing arising from the size of a population, but expanded improperly. A properly disaggregated way of modeling demand is to represent explicitly the growth of population and its response to housing availability. But the modeler may be tempted to simplify and aggregate the model structure, perhaps by attempting to represent the influence of population with a relationship that stimulates housing construction, HC, when the stock of housing units, H, has grown slowly (representing growth of population faster than housing and thus an increase in demand). Similarly, such a relationship might retard housing construction, HC, when housing units, H, has grown rapidily (representing overexpansion of the housing stock relative to the population and thus a slackening of demand). During times of moderate growth, this relationship would have a neutral effect on HC, representing the assumption that people could be found to occupy the houses; implicitly, the population growth would keep pace with the growth in housing.

What are plausible values for the parameters in this formulation? What is "moderate growth" relative to the speed of population movements? What should the magnitude of the effect of rapid growth on HC be? These questions cannot be answered by first-hand observations of cause and effect relationships; the questions call for conclusions based on mental models of the dynamics of the city. Can the modeler predict from intuition alone and characterize with one delay and one table function the dynamic interactions among housing and population, incorporating births, deaths, incentives for migration, family formation, or the ability of construction companies to expand? If not, the model is too aggregated for parameters to be estimated reliably from the unaggregate data.

Two possible actions can be taken when the parameters of a model are too aggregated to be set reliably from available unaggregate data. One course is to use another estimation technique (usually a statistical technique) and data at the level of aggregation of model structure. It seems unwise, however, to attempt to estimate a simple relationship if the actual system is so complex that expert opinion is unreliable. A preferable course of action is to restructure (usually disaggregate) the model so that its parameters correspond directly to observable, unchanging characteristics of the system. The disaggregation usually involves not only subdivision of levels into more levels, but also explicit addition of feedback loops that control the levels. For example, consider the relationship between land availability and urban housing construction discussed at the beginning of this section. Mass (1974) and Miller (1975) disaggregate this relationship to portray the details of land pricing, speculation, rezoning, and land use.

In summary, unaggregate data are by far the most abundant source of information. There is a wide range of specific estimation techniques,

extending from direct observation to ad hoc computations based on direct observation. Of particular interest is the problem of estimating the parameters of a table function. This problem can be reduced to the subproblems of estimating extreme values and slopes, specifying the normal point, and drawing a smooth curve through the extreme and normal points. The main pitfall in estimation with unaggregate data is formulating an equation and its parameters in an aggregate, simplified manner, so that participants in the system cannot reliably observe a value of the parameter as a characteristic of the real system.

7.3 Estimation Using a Model Equation

A model equation and its parameters specify a relationship between two or more variables. Estimation using a model equation starts with statistics that aggregate individual items or events that correspond to model variables. From such data the modeler derives the parameter values that enable the model equation to match the "real" relation between the variables. Estimation using a model equation encompasses all single-equation regression techniques. Theil (1971) offers a general treatment of regression techniques. Hamilton (chapter 8) and Mass and Senge (chapter 10) discuss the application of these techniques to system dynamics models.

For an example of estimation using a model equation, consider that the housing model calculates the rate of housing demolition, HD, as the number of houses, H, divided by the average housing unit lifetime, HL. Thus, if data are available for HD and H, HL can be estimated from the model equation:

$$HD.K = H.K/HL$$

$$HL = H/HD.$$

Estimation using a model equation is less frequent in system dynamics studies than estimation from unaggregate data. Nevertheless, two forms of model-equation estimation—one involving conversion factors, the other fractional rates of flow—have been useful in many studies and are therefore discussed here.

Conversion Factors

Many model parameters are conversion factors: they convert quantities from one dimension to another. For example, land per house, LPH, converts housing units to an equivalent number of acres. Equation (4) uses LPH in the definition of land fraction occupied, LFO:

```
LFO.K=(H.K*LPH)/AREA                                  4, A
LPH=0.1                                               4.1, C
AREA=9000                                             4.2, C
    LFO    - LAND FRACTION OCCUPIED (DIMENSIONLESS)
    H      - HOUSES (HOUSES)
    LPH    - LAND PER HOUSING UNIT (ACRES/HOUSE)
    AREA   - LAND AREA (ACRES)
```

These equations can be manipulated to compute the parameter as a function of the real data, in this case, the real LFO, H, and AREA:

LPH = LFO * AREA/H.

Estimating conversion factors offers a straightforward means of ensuring that the absolute magnitudes of model variables are realistic. Schroeder and Strongman (1974) show how real data were used to estimate conversion factors for a model of Lowell, Massachusetts. That model exhibits realistic magnitudes for population, housing, and employment.

Normal Fractional Rates of Flow

Equation (2) defines the rate of housing construction, HC, in terms of a level (houses, H), a normal fractional rate of flow (the housing construction normal, HCN), and a dimensionless multiplier (the housing-land multiplier, HLM). This format

rate = level * normal fraction * multipliers

is widely used (Alfeld and Graham, 1976, pp. 123–126, provides further discussion). One reason is that the multiplier can easily be estimated when it is normalized around 1.0 (see section 7.2). This format also facilitates the estimation of the normal fraction:

normal fraction = rate/(level * multipliers)

Under normal conditions (however defined), the multipliers, by definition, assume values of 1.0. So the normal fractional flow rate can be computed by dividing the observed rate by the observed level, both measured during a period of normal conditions. For example, suppose that the year 1960 is defined as the normal period for the urban area being modeled. Then, if the data are available, a value for housing construction normal, HCN, is obtained from the number of housing units constructed in the area during 1960 divided by the number of housing units in the area in 1960.

Pitfalls

Data at the level of aggregation of model variables must be collected for specific purposes. This fact may be a pitfall when data are collected for purposes different from those of the model. The completeness of a set of variables, the definition of the aggregated variables, or the time frame may render data unsuitable for estimating parameters using a model equation.

For example, the definition of the collected data may be inconsistent with model definitions. The housing-land multiplier, HLM, is defined so that housing construction is impossible when land fraction occupied, LFO, equals 1.0. Thus LFO must reach 1.0 when the area has all the housing it can hold. This means that land per house, LPH, must include not only the land directly beneath each housing unit but also the adjacent land for yards, sidewalks, driveways, roads, schools, and stores. The land per house, LPH, for a particular area could be calculated from the land area zoned for residential use (minus the area of vacant lots) divided by the number of dwelling units within the area. However, land in many cities is zoned for both residential and commercial use; some fraction of that land must be included in the residential land area as well. So the modeler must be aware of exactly what is and is not counted to form a particular piece of data, and whether that definition is consistent with the way corresponding quantities are used in the model.

Another difficulty with inappropriate data occured in an attempted revision of Forrester's (1969) urban model by Babcock (1970). Babcock set the normal constants with data on levels and rates of flow, but not with data only for the normal period; data for cities near equilibrium were used also. The simple housing model presented here can show what happened as a result. The housing model reaches equilibrium after a period of growth in the housing stock: the housing stock grows until a shortage of land suppresses further housing construction. Because the normal conditions in the model are growth conditions, the housing-land multiplier, HLM, must suppress housing construction by going well below 1.0. Suppose the actual rate of housing construction, HC_a, is divided by the actual number of houses, H_a, to obtain a computed value for the housing construction normal, HCN_c. Assuming the model equations are accurate, but using actual equilibrium data to compute HCN:

$$HCN_c = \frac{HC_a}{H_a} = \frac{H_a * HCN_a * HLM_a}{H_a} = HCN_a * HLM_a.$$

Since HLM < 1 in equilibrium, then

$HCN_c < HCN_a$.

The computed value of HCN, if used instead of the actual value, reduces the model's impetus to grow, and thus reduces the extent to which HLM must drop to bring the model into equilibrium. Similarly Forrester (1969) holds that growth ceases when land shortage and unfavorable internal conditions (principally a job shortage and predominance of lower-income groups) depress construction. Using data from near-equilibrium to compute normal fractions considerably reduces the extent to which internal conditions in the model must decline to halt growth. In fact, the model will no longer reproduce and account for depressed urban conditions. Babcock's re-estimated model therefore no longer fulfills its purpose, merely because the implicit assumptions used in parameter setting are violated.

The pitfall then in estimating parameters with data at the level of aggregation of model variables is that the computations require two assumptions: accuracy of an equation and appropriateness of the data. Such assumptions always constitute "more rope to hang yourself with."

7.4 Estimation Using Multiple Equations

As just described, estimation using a model equation consists of manipulating the equation to compute a parameter value. By contrast, estimation using multiple equations consists of manipulating several equations to compute a parameter value. Both techniques use data at the level of aggregation of model variables. The two techniques are distinguished here because they are usually distinct in practice. Usually either one equation is used analytically or all equations are used in simulations. Also, the pitfalls tend to be rather different.

For example, the housing construction normal, HCN, can be estimated by finding the value of HCN that causes housing growth to fit the observed rate of growth. This estimation would use all the equations. The fitting could be performed either with repeated simulations or, if possible, by a computation. For an example of such a computation, suppose that the stock of housing grows at 5.5 percent per year under normal conditions. Also suppose that, from observation of housing demolition, the housing unit lifetime, HL, is estimated to be 66 years—that is, 1/66 of the houses are demolished each year. If the model

equations and the parameter HL are assumed to be correct, then the housing construction normal, HCN, must exceed 1/66 by 0.055 to produce the observed rate of growth during the normal period. Therefore HCN can be inferred to be $1/66 + 0.055 \cong 0.07$. Although this computation uses all of the model equations, it can be performed simply because the model is very simple; more complex models usually require more elaborate numerical computations.

For another example, suppose a real system exhibits fluctuations of some specific period. The modeler can choose the magnitudes of time constants of the system so as to produce oscillations near the real period. Forrester (1968, ch. 10) derives a simple rule of thumb: for a second-order undamped system with two time constants T_1 and T_2, their geometric average approximately equals the period divided by 2π:

$\sqrt{T_1 T_2} = P/2\pi$.

Just as the estimation using one model equation subsumes single-equation regression, so estimation using multiple equations subsumes a family of statistical techniques. The most general technique is full-information maximum likelihood (FIML) estimation. Unfortunately, for nontrivial problems FIML usually requires extravagant computation time. Therefore two families of less general techniques have evolved. One family usually requires linear formulations, and information on exogenous variables that is both complete and accurate. These are the multiple-equation regressions (Theil, 1971, ch. 9–10). They are much more efficient computationally than FIML. The other family of techniques restrict the models to be dynamic (nonsimultaneous) and only mildly nonlinear. These are the full-information maximum likelihood via optimal filtering (FIMLOF) techniques derived from control theory (Schweppe, 1973). Although they require significant computation, they offer conceptual simplicity as well as estimating nonlinear dynamic models with flawed information on only a subset of model variables (Peterson, 1975; see also chapter 11 in this volume). Software is available commercially for doing FIMLOF estimations with system dynamics models (Peterson and Schweppe, 1974).

The general pitfall of multiple-equation estimations is the same as for single-equation estimations: the techniques assume the accuracy of the equation(s) and the data. The implicit assumption that most often thwarts multiple-equation estimation is that the discrepancy between real behavior and model behavior is due to the values of the parameters being estimated. In other words, the discrepancy can be misattributed. In the oscillation example cited above, if T_1 is inaccurate, and T_2 is being estimated , the estimate of T_2 will also be inaccurate as well in order that $\sqrt{T_1T_2} = P/2\pi$. A subtle type of misattribution sometimes occurs during model testing: if the model exhibits unrealistic behavior, the modeler changes a parameter to eliminate the unrealistic behavior. Thus the parameter has been estimated in the sense that its value has been chosen to allow the model equations to generate behavior that matches observed behavior. The problem is that the unrealistic behavior of the model may have been due not to the inappropriate parameter value but to an unrealistic formulation or to some *other* parameter value being awry. To attribute the unrealistic behavior to the original value of the altered parameter may indeed produce the right behavior but for the wrong reasons.

Another pitfall of both single-equation and multiple-equation parameter estimation arises from their use of data at the level of aggregation of model variables. To varying extents, both techniques force magnitudes of model variables and the relationships among them to conform to the magnitudes and relationships in the data. The forced conformity of some aspect of behavior to real data preempts the comparison of behavior to data as a validity test. For example, if the modeler estimates model parameters from data (at the level of aggregation of model variables) for housing units, housing construction, housing demolition, and land occupancy, the model is likely to replicate the overall behavior of the housing stock. But confidence in the model would be greater if the parameter estimation used unaggregate data and still resulted in a model that replicates aggregate behavior. In other words, if the model replicates real behavior when it doesn't have to, the replication is another basis for confidence in the model. Estimations from data at the level of aggregation of model variables hinders the modeler in using such a validity test.

One way to circumvent this pitfall, if there are enough aggregate data, is to follow the common econometric practice of using only part of the data from estimation and the rest for validation (Theil, 1971, pp. 603–604). This strategy is sometimes unworkable when not enough data exist, or if model equations are not general enough to replicate more than one set of data. The other way to avoid the pitfall is to use unaggregate data to estimate parameters and reserve aggregate data for validation; the latter strategy is commonly followed in system dynamics studies.

7.5 Conclusions

Three general categories of parameter estimation techniques have been presented. Estimation using data below the level of aggregation of model variables relies on observations of individual items or events that are represented in the aggregate by model variables. The principle pitfall of this technique is structuring the model on a level of aggregation too high to allow observers within the system to reliably translate their experiences into parameter values.

The two other techniques are estimation using a single model equation and estimation using several or all model equations. Both techniques assume the correctness of the given equation(s); they use the equation(s) to infer parameter values from data corresponding to model variables. These techniques share pitfalls. First, use of data at the level of aggregation of model variables diminishes the ability to validate. Second, these techniques are vulnerable to systematic errors when assumptions are violated. Econometricians encounter this pitfall in estimating simultaneous-equation models. Even though multiple-equation methods theoretically deliver greater accuracy than multiple applications of a single-equation method, the multiple-equation methods are more sensitive to minor violations of assumptions (less robust) than single-equation methods (Theil, 1971, p. 552). Similarly, parameter estimation from data at the level of aggregation of model variables is less robust than parameter estimation from data below the level of aggregation.

How should the modeler choose among the three techniques? As a point of departure, the modeler need not favor the equation-based techniques over the use of unaggregate data on the basis of accuracy: Senge (1978) shows that estimation from unaggregate data can be as accurate as other techniques. Moreover, the pitfalls or limitations of the three estimation techniques are only one consideration in choosing among them. The appropriateness of a given technique also depends on the context of the model-building effort, most notably on how the model variables have been selected, and how the model will be tested.

Variables may be selected by several criteria. Variables selected on the basis of their ability to contribute to point-predictive accuracy suggest the use of aggregate data and highly aggregated relationships. By contrast, variables in system dynamics studies are usually chosen because they can reproduce the causes of the problem being analyzed, and because they can be recognized and validated by participants in the system being modeled. These considerations favor the use of variables at a level of aggregation close to that observable by individuals, without regard to whether or not statistics are available. Goals and psychological pressures are seldom measured but can be central in the model's description of why problems arise. Thus a typical system dynamics model is formulated in a way that makes estimation from unaggregate data reliable and that circumvents the pitfall of highly aggregated, unobservable relationships.

What tests must the model pass for it to fulfill its purpose? If the tests concentrate on the ability of the model to predict future values, then the parameter estimation techniques used must emphasize making the model variables fit aggregate data. However, the situation is different in most system dynamics studies. The ultimate aim is to predict the qualitative results of a policy change. Such predictions are difficult or impossible to evaluate quantitatively and directly. Instead, system dynamics studies tend to use a broad array of tests of structure, behavior, and policy impact (Forrester and Senge, 1978). Many of these tests do not depend on having data at the level of aggregation of model variables; hence the modeler can be flexible in choosing a parameter estimation technique. Specifically, by setting parameters from unaggregate data, the modeler can reserve any aggregate data for the purpose of validity testing—comparing model behavior to real behavior. This reservation of data is analogous to the common econometric practice of estimating parameters from data from one time interval, and testing the results with data from another time interval.

A final aspect of the relation between parameter estimation and model testing concerns sensitivity assessment: much of the simulation in a typical system dynamics study aims at identifying the equations and the parameters that are central in producing the behavior and policy results. In a model of realistic detail, only a few parameters can alter the outcome of the model if they are changed. (Graham, 1978, details the process of identifying and dealing with sensitive parameters.) Thus it is common practice to set parameters on the basis of information at hand (usually unaggregate data) and defer intensive data collection and parameter estimation until model testing reveals the parameters that require such measures.

What techniques should the modeler use to estimate the parameters of a model? The estimation techniques should facilitate, and be facilitated by, the other phases of the modeling effort: the model should be formulated in such a way as to avoid the pitfalls or limitations of the estimation techniques. The model testing should guide estimation efforts.

And information should be used in estimating in a way that facilitates validation. These considerations are reflected in the following ensemble of recommended system dynamics practices:

1. Use a model structure that is detailed and realistic enough to allow participants in the system to supply data below the level of aggregation of model variables.
2. Whenever possible, estimate parameters with data below the level of aggregation of model variables and reserve data at the level of model variables for validity testing.
3. Use techniques based on model equations only as secondary techniques since they are vulnerable to systematic error.
4. Use simulation to identify the equations and parameter values that are critical to the outcome of the modeling effort and focus subsequent efforts on those equations and parameters.

References

Alfeld, Louis Edward, and Alan K. Graham. 1976. *Introduction to Urban Dynamics*. Cambridge, Mass.: MIT Press.

Babcock, Daniel L. 1970. "Analysis and Improvement of a Dynamic Urban Model." UCLA Ph.D. dissertation. Available from University Microfilms, Ann Arbor, Mich.

Forrester, Jay W. 1969. *Urban Dynamics*. Cambridge, Mass.: MIT Press.

Forrester, Jay W., and Peter M. Senge. 1978. "Tests for Building Confidence in System Dynamics Models." System dynamics group working paper D-2926-5, MIT, Cambridge, Mass.

Goodman, Michael R. 1974. *Study Notes in System Dynamics*. Cambridge, Mass.: MIT Press.

Graham, Alan K. 1978. "Parameter Formulation, Estimation, and Assessment in System Dynamics Modeling." System dynamics group working paper D-2349-1, MIT, Cambridge, Mass.

Mass, Nathaniel J. 1974. "A Dynamic Model of Land Pricing and Urban Land Allocation." In *Readings in Urban Dynamics*, vol. 1. Edited by Nathaniel J. Mass. Cambridge, Mass.: MIT Press, 1974.

Miller, John S. 1975. "Urban Dynamics and Land Rezoning." In *Readings in Urban Dynamics*, vol. 2. Edited by Walter W. Schroeder, et al. Cambridge, Mass.: MIT Press.

Peterson, David W. 1975. "Hypothesis, Estimation, and Validation of Dynamic Social Models." Ph.D. dissertation, MIT, Cambridge, Mass.