

FULL FEEDBACK PARAMETER ESTIMATION

Robert L. Eberlein
Pugh-Roberts Associates

Abstract The determination of parameters within a system dynamics model is an important part of the model development and validation process. There are, however, very few useful tools available for working on this. The primary reason for this lack of tools has been the difficulty of applying the theory that has been developed for full feedback estimation. Useful tools can be based upon heuristic application of much of the theory. Rules are outlined that allow the easy determination and application of filtering techniques that deal with the problem of unobservable variables. The attributes of these techniques are discussed in different settings. Application to an example serves to illustrate a number of the issues.

INTRODUCTION

The process of model development and validation often requires adjustment of model parameters that influence behavior. This adjustment process, frequently referred to as "tuning," is slow, error prone and difficult to replicate. The application of more rigorous and automatic techniques such as regression is difficult, primarily because of the size and complexity of the problem. System dynamics models have a large number of unobservable variables and a great deal of feedback. These attributes enhance the value of models in representing processes of interest, but make the comparison of the model to what actually happened difficult. In this paper we will discuss tools and techniques to aid in the tuning process, including a usable approach to estimating parameters in the context of the complete feedback structure.

In the first section of this paper we will develop some heuristics that can be used to arrive at full feedback estimates of model parameters. These heuristics are meant to ease the task of specification required to do full feedback estimation and thereby make the techniques for such estimation more widely available. In developing these rules we will describe the theory of full feedback estimation and hence the source of our heuristics.

The second section of this paper addresses the very pertinent issues of building to a reasonably tuned model from early on in model development. Full feedback estimation requires as a

starting point a model capable of roughly replicating what happened. A model built up with arbitrary parameters is unlikely to produce sensible results, and the larger and more complex the model is the more difficult it is to arrive at sensible results. In the second section we describe a useful methodology for breaking down a model into smaller pieces which can more easily be roughly tuned. We further discuss the nature of the potential problems in using full feedback estimation techniques in what is clearly a restricted feedback setting.

FULL FEEDBACK

Full feedback estimation takes into account all of the feedback structure that defines a model in order to find the model parameters. This is the same as saying that the model is simulated in order to determine the parameters rather than being estimated statically using lagged values or other constructs. The further attribute of full feedback estimation is that the estimation of the parameters is coupled with an estimation of the model variables. Thus the simulation used is not a standard simulation, but involves adjustment of model variables on the basis of available data.

Though full feedback estimation techniques have been available for use in system dynamics for some time (Peterson 1980) their application has been limited. The primary reason for this has been the difficulty and expense of application of these techniques. The techniques require that a great deal of information about the noise structure of the model be specified, and are computationally costly. In this section we outline heuristics that can be used to accomplish full feedback estimation without requiring this detailed specification of the noise structure of the model. The rules developed are based on the assumption that it is model structure that is most important in determining the relationship of internal model variables to available data.

We will first outline the reasons for, and the nature of, full feedback estimation. We use this discussion to motivate reasonable rules for achieving full feedback estimation without having to introduce a burdensomely large quantity of information. These rules, in a sense, introduce that information - they represent simple and codified informed guesses about the stochastic nature of the system we are dealing with.

The important difficulty in the estimation of the parameters for a system dynamics model is the potentially large number of unobserved variables. Because of the lack of information on what values these variables take on, the simulation can easily diverge from what actually happened. This will be the case even when the

parameter values are exactly correct, simply by virtue of the small randomness entering the system or because the initial conditions are not correct. A striking example of this is given in Forrester (1961 appendix K) where the same model with slightly different noise inputs shows a completely different time path for its variables. If one is trying to tune such a model to data by simply simulating the model without realizing the cumulative nature of the errors the resulting parameter estimates cannot be expected to reflect correct parameter values.

While complete divergence of behavior due to the entry of noise will only occur for some models, all models will show some degree of divergence. This divergence may be restricted to differences in timing, but will likely also be important to the shapes of time paths for model variables. The greater the tendency of the model to diverge, the greater the value adjusting model variables will have.

In order to deal with the issue of keeping a model on track it is necessary to be able to describe succinctly where a model is at any time. This can be done using a vector made up of all of the levels in the model which we shall refer to as the state vector. The state vector suffices to describe current conditions because the current value for every variable that is not a level can be calculated from the values for the levels. The states play the key role in determining whether a model is on track; in trying to adjust variables in a model it suffices to adjust the state vector. The model that forms the basis of our discussion is written in terms of state variables only as

$$\dot{\underline{x}} = \underline{A}\underline{x} + \underline{B}\underline{e} + \underline{\varepsilon} \quad (1)$$

with the dot (.) denoting the time derivative.¹ \underline{x} is the state vector, \underline{A} the dynamics matrix, \underline{e} the vector of exogenous variables entering the model, \underline{B} the matrix transforming those and $\underline{\varepsilon}$ the error terms influencing the process. The error term $\underline{\varepsilon}$ is assumed a stochastic process (Doob 1954) though it is easier and equivalent for our discussion to think of the error as changing at regular intervals. The above model is linear and for notational convenience all equations will be written linearly; application to nonlinear models is discussed.

Associated with the model given in equation (1) is a series of measurements of available data. When we speak of data we mean

-
1. In the standard DYNAMOTM notation the equation for the first element of \underline{x} (x_1) would be

$$L \quad x_{1.k} = x_{1.j} + dt * (a_{11} * x_{1.j} + a_{12} * x_{2.j} + \dots + a_{1N} * x_{N.j})$$

values of model variables that are measured at different times. The equations for determining these variables are written in terms of the states as

$$y = Cx + \delta \quad (2)$$

with y the vector of all observed data, C the matrix relating this to the states and δ a vector of measurement errors. Normally the matrix C has more columns than rows, that is, there are only a small number of measured variables available relative to the number of levels in a model. Thus, it is usually not possible to determine, or estimate, values for the entire state vector simply from consideration of observed variables.

The problem of determining what the actual state of a model is at any time has long been recognized in its own right, and in a seminal article Kalman (1960) derived a method to get the best estimate of a model's states given the available information. The assumptions underlying this derivation are very strong, both linearity and normality of the errors are required. However, the basic approach has been adapted to nonlinear models quite regularly (Anderson and Moore 1979, Sage and Melsa 1979, Schweppe 1974). The technique for estimating the state of a model, commonly referred to as Kalman filtering, is at the heart of the full feedback parameter estimation techniques. Kalman filtering lines the model up with what was actually happening as well as possible given the available information and model parameters.

The process of full feedback estimation makes use of Kalman filtering to get good estimates of the model variables at every time data are available. The estimates of the model's variables arrived at using the data available from one observation time are retained for further simulation. After further simulation the new values of model variables can be compared with the next available observations. Differences between the model variables and the data indicate changes for all the model variables through the Kalman gain, this is often referred to as an update. The equations for this update are quite straightforward. If we let G represent the Kalman gain then the best estimate for the state \hat{x} is given by

$$\hat{x} = \bar{x} + G (y - \bar{y}) \quad (3)$$

where \bar{x} is the model's predicted value for the state vector based on simulation from the last update, y is the actual data vector and \bar{y} is the model's predicted value for the data vector given \bar{x} . The value of \hat{x} thus obtained is used to continue the simulation to the next observation time.

After a full simulation the errors that were made can be reviewed, and the parameters changed so as to make these errors

smaller. In the model of equation 1 the parameters are the elements of the A , B , and C matrices. The errors are given by $y - \hat{y}$, the last term in equation 2. Since this error is over a number of different variables, weights have to be assigned to arrive at a summary error measure. The best parameters make the summary error measure (normally a sum of squares) as small as possible.

Essential to the Kalman filter is the determination of the Kalman gain (G) that is used to update all model states based on the errors made in predicting the available observations. Even for a linear model this gain is changing over time, though it normally goes to a steady state. The gain is based on the covariance between the state variables and the observations. The determination of the gain requires knowledge of the covariance structure of the noise entering the model, the dynamics matrix of the model, the transformation required to go from the states to the variables corresponding to the data, and the covariance structure of the measurement noise. Based on this information the calculation of the gain requires the calculation and inversion of the observation covariance matrix which postmultiplies the state-observation covariance (see Schweppe 1974 or Anderson and Moore 1979).

In intuitive terms, the Kalman gain is based on two things: first, how much, on average, does unexpected (noise driven) variation in the observed values tell us about variation in the states, and second, how much of the observed values are based on signal (actual system changes) and how much on noise. The stronger the link between an observed variable and a state, the more the suggested error in the value of the state for a given error in the observed value. The greater the signal relative to the noise the greater the strength of the updating.

The difficulty of application of the Kalman filter arises partly out of the need for a continual updating of the gain, but primarily from the large amount of information required to calculate the gain. The calculation of the gain requires the characterization of the noise entering the model. For every equation it is necessary to specify the variance of the noise entering the equation as well as the covariance of this noise with the noise entering all other equations. This is a formidable task for a system dynamics model and requires a great deal more information than the modeler is likely to have. While it is possible to estimate characteristics for some of the noise terms, the number that can be estimated must be less than the available number of observations (Mehra 1974). The specification of measurement noise has similar problems associated with it. Again, all the different variance and covariances must be specified, and their estimation is similarly restricted.

The determination of a Kalman gain thus requires a good deal of computation and generally a great deal of specification by the user. Because this specification is often arbitrary the value of obtaining it is also questionable. To the extent that the resulting estimation is insensitive to an incorrect specification of the noise covariance the specification is largely an inconvenience. In the approach we take we assume that the appropriate gain is fundamentally determined by model structure, and only secondarily by the covariance of the noise entering the model. We develop some simple guesses about what the states' covariances are and then use these to derive a gain.

Intuitively, the state covariance is simply a measure of to what extent, and in what direction, changes in some states are associated with changes in other states. To get at this we simply need to let a state change and observe what other states are doing. How much and in what direction change is introduced is dependent on the nature of the noise influencing a state. We specify the noise to be unrelated between the states and to have a strength (or variance) that is proportional to the variance of the original variable over the full simulation. We measure the covariance by making a simulation without this noise entering the model and then making a simulation including the noise and taking the sample variance of the difference in values between the two simulations.

Formally if we let \underline{x} represent the state variable from the model simulation with noise entering and \underline{x}' the state variable from the simulation without noise entering, the difference between these will evolve as

$$\dot{\underline{x}} - \dot{\underline{x}'} = \underline{A}\underline{x} + \underline{B}\underline{e} + \underline{\varepsilon} - \underline{A}\underline{x}' - \underline{B}\underline{e} \quad (4)$$

$$= \underline{A}(\underline{x} - \underline{x}') + \underline{\varepsilon}, \quad (4a)$$

since the exogenous variables will be the same. The stochastic structure of equation 4a is the same as that of equation 1 but has the advantage of having no influence from the exogenous inputs. Because of this the variance as measured simply by simulating the model with a noise input and without a noise input will converge to the state covariance. This is convenient, since it allows us to estimate the state covariance without having to do all the calculations normally required to do so and thereby allows us to quickly arrive at an approximation to the Kalman gain.

The measure of the state covariance arrived at in this manner is dependent on the arbitrarily chosen values of the variances. This choice of variances was intended to tie the relationships between the different states down, not to be representative of

what the different states were actually doing. The calculation of the gain involves both an inversion and a matrix multiplication. For this reason the effect of changing assumptions about the variance of each noise term is quite limited.² The dynamics matrix A is fundamental to determining how closely and in what manner the states are linked. The observation matrix (C) determines the linkage of the states back to the observations. While we do not make explicit use of the A matrix we do calculate the C matrix to determine the observation covariance.

The adjustment for observation error is made by comparing the detrended variance of the historical and simulated variables.³ If the data values show a higher variance than the simulated values they are assumed to have some error in measurement. Because of this, they are weighted less heavily in the updating of the state vector than they otherwise would be. Specifically if we let Ω represent the state covariance and Ψ the diagonal matrix with excess of the data variance over the simulated variance then the gain we use is given by

$$G = \Omega C^T (C \Omega C^T + \Psi)^{-1} . \quad (5)$$

Thus a variance of the available data causes the gain to get smaller relative to that series (this is clear in the scalar case, and similar reasoning applies to the matrix case).

Once the best estimates of the states are available we know how much of an error the simulation model makes when the filtering tools are working to keep the model as much on track as possible. The parameter determination is then done via a goal searching technique that finds the best value of a parameter by minimizing the squared error loss.⁴

The error that is considered in performing the tuning is the squared error between the model generated value with filtering

-
2. This will not be as true when changes in the covariance between noise terms (which we assume zero) are made. When this happens the patterns of linking between states can change more dramatically.
 3. The variance for many models is a calculation about a moving mean that is the result of deterministic behavior. Taking the variance around a trend is one way to remove, approximately, as much of this deterministic component as possible.
 4. The parameter search technique employed for the results reported was from the MINPACK search algorithms, though alternatives are currently under active consideration.

and the actual data ($\bar{y} - \hat{y}$). When there are a number of different data series the series are weighted according to the variance of the data streams (highly variable streams are given low weights). It is possible and simple to change these weights, emphasizing one series and ignoring others. What weights are chosen will often just reflect the focus on tuning one aspect of a model. Adjusting some parameters to get some model variables right is a very useful way of tuning a model.

To summarize, we have described a technique that can be used to arrive at an approximate value for the gain G . This gain tells us how much to update the different states in the model in response to differences between the simulated and actual values. Simulation is performed by updating the states each time observed variables are available. This type of simulation clearly differs from the standard model simulation and there is a sense in which the model is different from the original simulation model. Using this simulation we can record what errors are made for all of the available data. Using a weighted least squares criterion to summarize this error it is possible to determine the best parameter values through an optimization search algorithm.

Estimating a Trend Function

In order to illustrate the techniques we have discussed we can apply them to a very small and simple model. This is a model of the formation of inflation expectations as described in Sterman (1986). The model is given by

$$\begin{aligned} L \text{ FINF.K} &= \text{FINF.J} + (\text{DT}/\text{TPT}) * (\text{IFINF.J} - \text{FINF.J}) & (7) \\ A \text{ IFINF.K} &= (\text{PCPI.K} - \text{RCPI.K}) / \text{THRC} \\ L \text{ RCPI.K} &= \text{RCPI.J} + (\text{DT}/\text{THRC}) * (\text{PCPI.J} - \text{RCPI.J}) \\ L \text{ PCPI.K} &= \text{PCPI.J} + (\text{DT}/\text{TPCPI}) * (\text{CPI.J} - \text{PCPI.J}) \end{aligned}$$

with FINF the forecast inflation rate, IFINF the indicated forecast inflation rate, RCPI the reference consumer price index, PCPI the perceived consumer price index and CPI the actual consumer price index. The above model is no more than a nonlinear version of the TREND MACRO (Richardson and Pugh 1981 pp 370-71) on CPI with a final smoothing of the forecast. The three time constants (TPT - time to perceive trend, THRC - time horizon for reference conditions and TPCPI - time to perceive the consumer price index) will determine how the forecast generated responds to changing conditions.

There are observations available on what people actually forecast for the inflation rate six months into the future (Carlson 1977). Thus there are measure of FINF but not of the other two states. Though the model does not possess a great deal of feedback, it clearly requires some sort of state updating to estimate the parameters. We will consider the estimation of THRC based on a

gain calculated by the rules we have outlined. It is worth noting that not all three parameters in the model can be estimated together. There is a degeneracy in this model in the sense that changing one time constant can be compensated for by changing a second. For this reason we restrict our attention to the estimation of THRC with the two other parameters specified (at TPCPI = 2 and TPT=2 months).

The gain when calculated by the above described rules is given by

$$\underline{G} = \begin{array}{ll} 1.0 & \text{on FINF} \\ -0.8 & \text{on PCPI} \\ -72. & \text{on RCPI} \end{array} \quad (8)$$

where the price indices are normalized about 100, and the inflation forecast is done as a fraction. What the gain says is that if the model simulates too high a value for the inflation forecast it is because the model's value for the inflation forecast is too high and because the reference price index is lower than the model indicates (the perceived price index will not be appreciably changed). The magnitude of the numbers suggests that if the model is high by 1% (.01) it will cause RCPI to decrease by about .75, the effect on PCPI is negligible. This adjustment is approximately the same as a .5% change in the reference consumer price index.

When the model is estimated using this gain to update the state values, the best parameter value that can be determined for THRC is about 16 months. Without any filtering the best value that can be determined for THRC is approximately 5 years, which is not very reasonable. The reason for this long estimate without filtering is that the model consistently over estimates what the forecast will be. Since the majority of the data are over a period of rising inflation the long time constant keeps the model generated forecast low. As argued in Sterman (1986) it is likely that this over-estimation indicates a shortcoming in the model. The long time constant tries to correct for this. With filtering, on the other hand, the reference conditions are updated continually, and most frequently lowered. This accomplishes the same thing as the long time constant but for different, and more appropriate, reasons.

Time plots of the residuals and the original series are shown in Figure 1 for the 16 month time constant with and without filtering active. Without filtering, the residuals are simply the actual forecast less the simulated forecast. The residuals when filtering is active are smaller and less systematic in nature. In both cases, however, the model consistently over predicts what the inflation forecast should in fact be. The filtering equations tries to correct this and to some extent do. The fact that the bias is still there indicates that the problem

is not transient, not the effect of incorrect initial conditions or some unexpected change that was not modeled. Rather the error seems to be systematic and thus an indication of specification problems with the model being used.




Figure 1 Errors with (FILTE) and without (NOFILTE) filtering.

SECTORIZATION

Full feedback estimation is best viewed as a tool that can be combined with a number of others in a complete tuning process. The process of tuning requires parameter adjustment, model reformulation and in some cases extensive reconceptualization. For many stages of the tuning process it is neither feasible nor desirable to proceed with the complete feedback structure. An important technique for tuning a model involves tuning a variety of components of the model. Homer (1983) discusses some of the issues in tuning by the more traditional approaches in System Dynamics. Eberlein and Wang (1985) suggest utilization of readily available statistical techniques on model equations and subsectors wherever possible.

In this section we will outline some procedures for breaking down a model into different subsectors. The important feature of work on a subsector of a model is that feedback links are cut. This is valuable because serious problems in one area will not

compound themselves to yield worthless results in other areas. In addition, where tuning requires repeated simulation, working with a small subsector can be much faster. Given such a sectorization the full feedback estimation techniques can be applied to the partial model. When this approach to tuning is taken there are some potential problems that need consideration.

The procedures we consider for breaking down a model are based on a straightforward analysis of the feedback structure of the model which, because it is comprehensive yields a clean, orderly and easily reproducible sectorization of a model. We will generate a sectorization that is complete in the sense that the sector developed will give the same simulation results as the original model. This, in most circumstances, will require that time series for variables used by, but not part of, a sector be generated. These time series will form part of the exogenous input into the sector.

The technique that is employed for sectorization is the exact analogue of going to the stock and flow diagram and cutting links between variables. What is to be retained within a sector is determined by first looking at all the things that are desired for consideration, and then at all the links that have been cut. All variables continuing to lie within a feedback path among the desired variables are retained. Any variables not in the feedback path but directly affecting a retained variable will be made exogenous. An exogenous variable will be kept as an input into the sector, but it will not be generated within the sector. Variables that are not part of the feedback structure and are not needed as exogenous inputs are simply discarded.

The dynamic feedback structure of the model is the one that is used to determine what variable should be kept. Variables that are part of the initialization structure for any included variables are also kept, but any variable retained only for the purpose of initialization can be made nondynamic (if this is not already the case). This approach allows a complete feedback model except that some exogenous variables are added. If all exogenous variables are given the same values as they had in the original model, the sector thus arrived at will yield the same simulation results as the original model for all retained variables. As we shall discuss, for the purposes of tuning, retaining the original model values is not desirable and the use of historical data is preferable. In either case though, the subsector derived can be simulated and should be easier than the full model to tune.

Applying Full Feedback Estimation to a Sector

Though there is a sense in which a model sector intrinsically lacks full feedback, the tools of full feedback estimation can be

applied. The application of such tools in this setting needs some discussion since the idea behind full feedback estimation is that nothing is left out. The direct application of the tools we have discussed is subject to some caveats and these are worth considering even when the full model is being used. In this section we give some guidelines as to the nature and seriousness of the errors that may arise when using full feedback estimation in a partial feedback framework.

The model upon which the techniques of full feedback estimation are based is given in equations (1). The noise term (ϵ) is assumed to display no serial correlation. A sectorization of the model effectively breaks the state variable vector \underline{x} into two components. We can rewrite equation (1) with this break in mind as

$$\begin{vmatrix} \dot{\underline{x}}_1 \\ \dot{\underline{x}}_2 \end{vmatrix} = \begin{vmatrix} \underline{A}_{11} & \underline{A}_{12} \\ \underline{A}_{21} & \underline{A}_{22} \end{vmatrix} \begin{vmatrix} \underline{x}_1 \\ \underline{x}_2 \end{vmatrix} + \begin{vmatrix} \underline{B}_1 \\ \underline{B}_2 \end{vmatrix} e + \begin{vmatrix} \underline{\epsilon}_1 \\ \underline{\epsilon}_2 \end{vmatrix} \quad (9)$$

The sectorization of this set of equations concentrating on the first sector yields

$$\dot{\underline{x}}_1 = \underline{A}_{11}\underline{x}_1 + \underline{A}_{12}\underline{x}_2 + \underline{B}_1e + \underline{\epsilon}_1 \quad (10)$$

with the \sim on \underline{x}_2 denoting the fact that \underline{x}_2 is now an exogenous variable. Equation 10 is different from equation 1 only in that some of the exogenous variables were part of the original model. To the extent that these variables are known exactly, this does not cause any difficulty. However, if it is necessary to introduce exogenous variables which can be known no more than approximately, then equation 10 is not an exact analogue of equation 1.

The problem in using incorrect exogenous variables is very intuitive. If the inputs into a model are not correct, the model with the right parameters will not appear to be well tuned. It is possible, and likely, that a model with incorrect parameters will appear to be better tuned. This is true for the model of equation 1 as well as that of equation 10. However, if the errors in the exogenous variable measurements are not related to the endogenous variables the only parameters affected will be those for the exogenous variables. That is, the manner in which the exogenous variables influence the model will be incorrect, but the variables affecting the feedback structure of the model will not. Model tuning will tend to adjust the parameters on the

exogenous variables to compensate for their incorrect measurement (Thiel 1971 section 12.2).

When the errors in exogenous variables are related to the endogenous variables the results are less reliable. It is possible in this case to attribute the effects of an endogenous variables to an exogenous variable. This problem can be accentuated by the use of a full feedback estimation technique, because errors in the exogenous variables may cause unneeded and incorrect adjustment of the endogenous variables. This will affect the endogenous feedback structure of the sector that results from parameter estimation. The severity of the error is difficult to predict, and the consequences of integrating sectors with such errors for full model simulation are also unpredictable.

The full feedback estimation of a model sector will do a reasonable job of identifying the endogenous feedback structure of that sector if the needed exogenous variables do not have errors correlated with the endogenous variables. This will likely be the case if the exogenous variables needed have historical values available. If there are no historical data available it is necessary to generate values by some means. One obvious method is through simulation of the full model. If this is done, however, the stronger the dependence of the exogenous variables being created on the endogenous variables, the less reliable the results will be. For this reason it may be preferable to make up plausible values (possibly based on a model simulation) for important unmeasured quantities in the early tuning stages. At later stages in the tuning these values can be made endogenous.

Even when exogenous values are not suspected of having errors correlated with the endogenous variables coefficients determining the impact of these exogenous variables may be incorrect. The errors in the exogenous variables introduce a bias in parameters determining their impact. The simplest example of this would be an exogenous variable always measured 20% higher than its actual value. Parameters on this variable would be scaled down accordingly. When the variable becomes endogenous and takes on lower values, the parameters thus arrived at will no longer be valid.

All of the discussion of errors in exogenous variables can be applied to equation 1. To the extent that the kinds of problems likely to exist in a model sector are to be found in the full model there will be difficulties. As long as the exogenous variables are not measured with an error that is correlated with the endogenous model variables this problem is restricted to the direct effects of the exogenous variables. The parameters on these direct effects are essentially adjusted to compensate for

the errors. As long as new values for the exogenous variables continue to have these properties, their effect on the endogenous variables will be appropriate.

CONCLUSIONS

We have outlined a series of rules that can be used to specify the stochastic nature of a dynamic system. These rules allow the determination of a gain that can be used to get improved estimates of all model variables over the course of a simulation. Using the improved estimates of the models variable over the course of the simulation the error made by the model is calculated. The model parameters are then chosen so as to minimize this error. This technique can aid in the determination of model parameters and has the advantage of being easily reproducible.

The development of a model normally proceeds in a number of steps and it is important to know the role of full feedback estimation in these steps. We have set out some guidelines for isolation of sectors and structures within a model for the purposes of parameter determination. In the context of these sectors there are some cautions necessary as to the meaning of parameters resulting from the tools discussed.

The theory underlying full feedback estimation is based on strong assumptions of linearity and normality of errors. We have extended the rules to deal with nonlinear problems essentially by analogy. The results will not have the optimal properties that will be the case for linear models. However, the results are likely to be a substantial improvement over less sophisticated techniques that make no attempt to correct for the fact that a model may be off track at any time.

REFERENCES

- Anderson, B.O. and J.B. Moore (1979). *Optimal Filtering Techniques*. Prentice Hall, Englewood Cliffs, New Jersey
- Carlson, J.A. (1977). A Study of Price forecasts, *Annals of Social and Economic Measurement*, Vol 6, 27-56
- Doob, J.L. (1954). *Stochastic Processes*. John Wiley and Sons, New York.
- Eberlein, R.L. and Q. Wang (1985). Statistical Estimation and System Dynamics Models, in *The Proceedings of the 1985 International System Dynamics Conference*, Keystone, Colorado. pp. 194-205
- Forrester, J.W. (1961) *Industrial Dynamics*. MIT Press, Cambridge, Massachusetts
- Homer, J.B. (1983) Partial Model Testing as a Validation tool for System Dynamics, *The proceedings of the 1983 International System Dynamics Conference*, Boston, Massachusetts pp 919-932
- Kalman, R. (1960) A New Approach to Linear Prediction and Filtering Problems. *Journal of Basic Engineering*, Series D, 83 pp. 95-103
- Mehra, R.K. (1974). Identification in Control and Econometrics: Similarities and Differences. *Annals of Social and Economic Measurement*, Vol 3
- Peterson, D.W. (1980). Statistical Tools for System Dynamics. in J. Randers ed. *Elements of the System Dynamics Method*. MIT Press, Cambridge, Massachusetts. pp 224-245
- Richardson, G.P. and A.L. Pugh (1981). *Introduction to System Dynamics Modeling With Dynamo*, MIT Press, Cambridge, Massachusetts.
- Sage, A.P. and J.L. Melsa (1979) *Estimation Theory With Application to Communications and Control*, McGraw Hill, New York
- Schweppe, F.C. (1973) *Uncertain Dynamic Systems*, Prentice Hall Inc., Englewood Cliffs, New Jersey
- Sterman, J.D. (1986) Expectations Formation in Behavioral Simulation Models. Working paper D-3821, System Dynamics Group, Sloan School of Management, MIT, Cambridge, Massachusetts
- Thiel, H. (1971). *Principles of Econometrics*, John Wiley and Sons, New York.

