Third Year Literature Review

Modeling Macroeconomic Systems

Mark Rosenstein

Introduction

The concept of simulating economic systems has a fairly long history. In the early 1960's economic simulation broke into two major branches. The simultaneous models with the name Lawrence Klein and a Nobel prize describing one branch. A second path followed by Guy Orcutt [Orcutt, 1961] and his colleagues at the Social Systems Research Institute at the University of Wisconsin is another. It is this less traveled path that will be considered. Orcutt's describes his view of the simulation process.

> One of the major objectives of the approach taken is to provide an instrument for consolidating past, present, and future research efforts of many individuals in varied areas of economics and sociology into one effective and meaningful model; an instrument for combining survey and theoretical results obtained on the microlevel into an all-embracing system useful for prediction, control experimentation, and analysis on the aggregate level. The possibilities of such a system tempt the imagination.

Orcutt's models consist of micro-subsystems generated through various technologies, such as analytic or simulation, integrated together by the simulation system to generate a predictive model. Requirements on such a system are, (1) Deep micro-level entities, (2) Extensibility. That is, if new institutions develop, or more knowledge emerges about an old one, it should be easy to augment the model. (3) It should be easy to place blame for models high level misbehavior with the assumed behavior of the micro-level entities, and (4) It should be easy to experiment with low level entities and not be concerned with changing other unrelated entities. Combining this view of macroeconomic models with more recent computer techniques make it possible to achieve the above flexibility and understandability. The computer techniques will revolve around object oriented explanatory systems.

This literature review encompasses the literature of economic systems simulation, especially macroeconomic systems. The emphasis here is more specialized then general simulations, in that the structure of the model and simulation technique must interact to provide an explanation of the interaction of the models components. This specialization eliminates estimated simultaneous equation models simulation, since in this type of model explanations are not possible. As Ackoff[1967] states:

> Furthermore, whatever else regression analysis can yield, they cannot yield understanding and explanations of phenomena. They describe, and at best, predict.

This specialization does not imply that pursuing simultaneous equation models is without interest. Instead, I claim that even in the unlikely case that the "right" model with its many equations, variables and appropriate deep parameters in the [Lucas,1976] sense were found and estimated, an explanatory model could be used to understand the ramifications of this "right" model.

The strategy for the rest of this paper is to first consider an example of a qualitative model. This model illustrates a qualitative explanatory model. The next section contains an economic model that partially avails itself to explanatory modeling techniques and includes a brief discussion of how a fully explanatory simulation would further aid in understanding the questions posed in that example. A detailed discussion of explanatory simulation techniques follows these examples. The last section indicates where future work could be done in explanatory macromodels. The paper ends with a comprehensive bibliography of this area.

Qualitative Model: An Example

As far as I can determine, the literature contains few economic examples of investigators attempting to generate explanatory models. Unfortunately, even amongst the few examples, none are clean-cut illustrations of explanatory modeling techniques. This being so, I beg the readers indulgence, while a crisp noneconomic model is discussed. Skipping this sections will substantially decrease understanding of the whole.

This material is drawn from a system designed to instruct students in the behavior of a propulsion system. The model is that of a reducing valve, whose function is to reduce pressure in a steam line, by having a larger outport than inport. The valve attempts to maintain a constant outport pressure with a system of subvalves.

From one perspective, the valve can be described by an exponential decay function. A graph of the outport pressure, subject to a sudden increase in pressure would be a curve decaying toward the equilibrium pressure. To answer many practical questions, such as what happens if a subvalve sticks this model is seriously deficient. A robust understanding of the reducing valve includes more than the exponential decay function.

The sudden increase in pressure is the analog of shocking a macroeconomic system. In a simultaneous system model, a time path of the aggregate variables is generated. In many senses, this output is superficial in the same way the exponential decay function was. To think about the system, it is necessary to have a model of the deep structure.

Consider an economic model with deep structure being shocked. Mexico defaults on its loans. There are immediate effects on the banks holding the loans and individuals holding Mexico bonds. If some of those institutions fail there are effects on their shareholders and effects on the banking industry and on the federal government. The effect is transmitted to more and more levels of the economy. This model describes how an economist would describe the deep effects of a default.

How would the default look to a simultaneous systems model. There would be graphs of GNP effects and price effects, but not the kind of story that the beginning sentences of the previous paragraph brought out. There are no internal processes, because simultaneous systems don't model them.

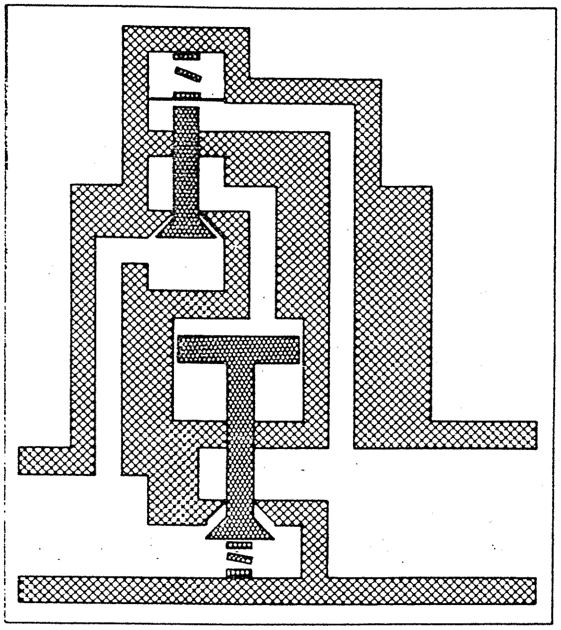

1.→>CONSIDER DEVICE SPRING-REDUCER-VALUE
WHAT SHOULD IT BE CALLED?
V1
OK

→>CONSIDER (>>PRESS OUTPUT-PORT V1)
OK

→>SUPPOSE U
WHEN THE VALUE IS INCREASING,
THE SYSTEM TRIES TO MAKE IT DECREASE.
THIS MEANS IT HAS NEGATIVE-FEEDBACK.

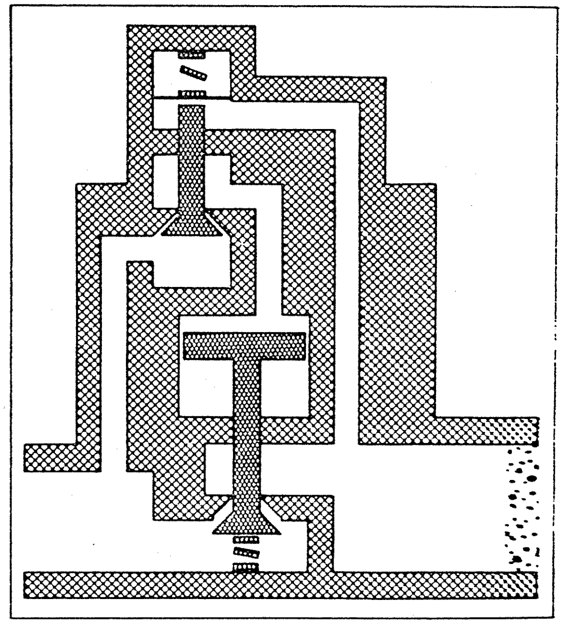

2. EXPLAIN
WHEN THE PRESSURE IN THE OUTPUT PORT
RISES,

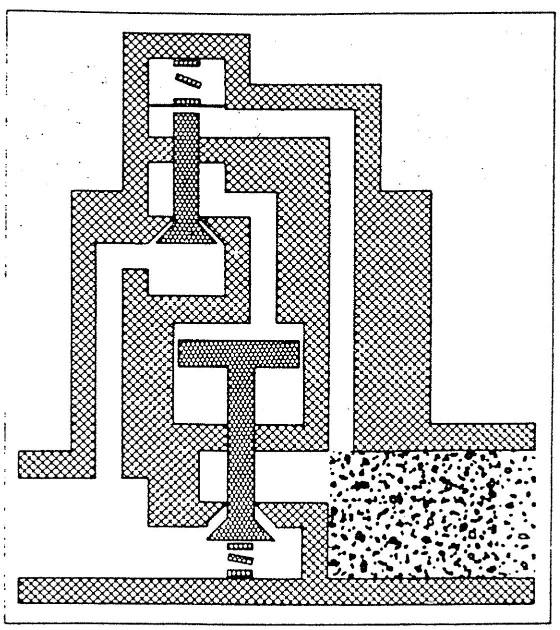

3. THE PRESSURE IN CHAMBER 0 RISES,

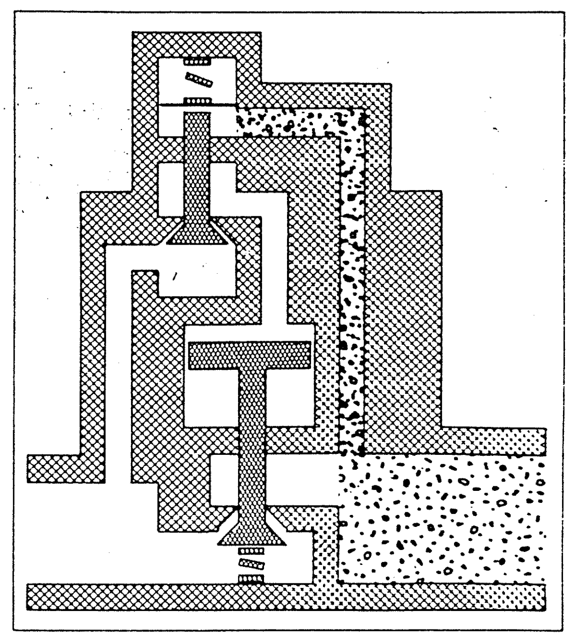

4. THE PRESSURE IN THE LOW PRESSURE PORT RISES,

Figure 1. Successive frames of the explanation generated for a spring-loaded reducing valve.

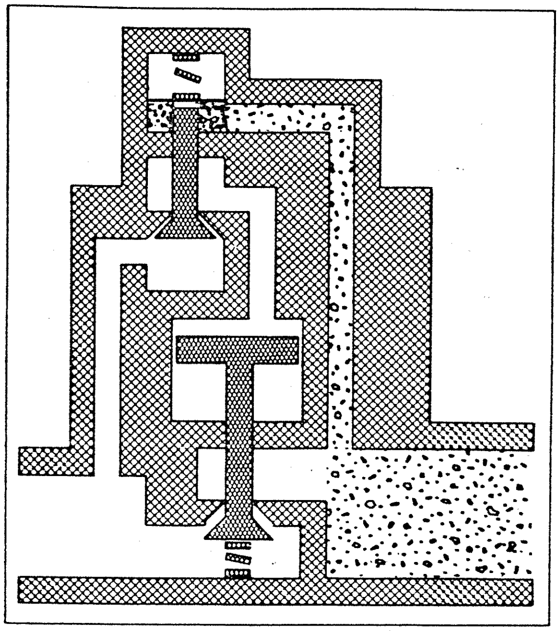

5. AND THE PRESSURE IN CHAMBER 5 RISES.

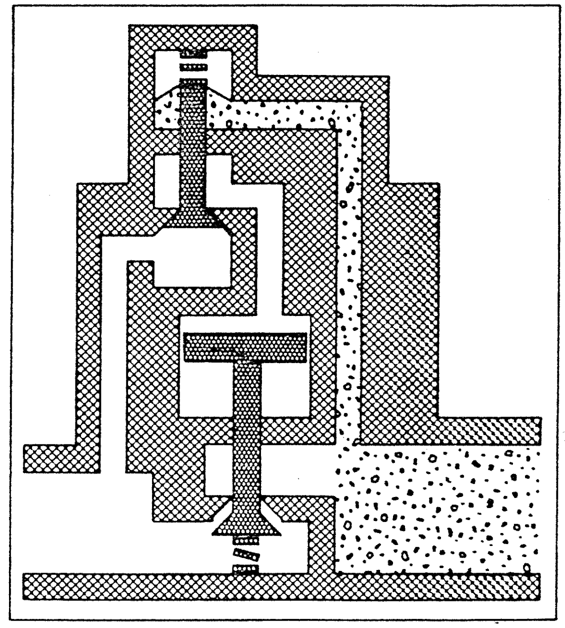

6. THE INCREASING PRESSURE PUSHES THE DIAPHRAM UP AND CLOSES THE AUX VALVE.

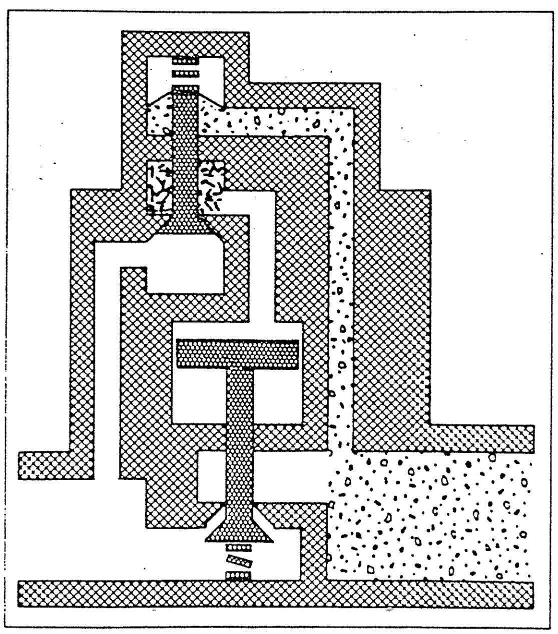

7. THE PRESSURE IN THE AUX VALVE'S OUTPUT SIDE FALLS,

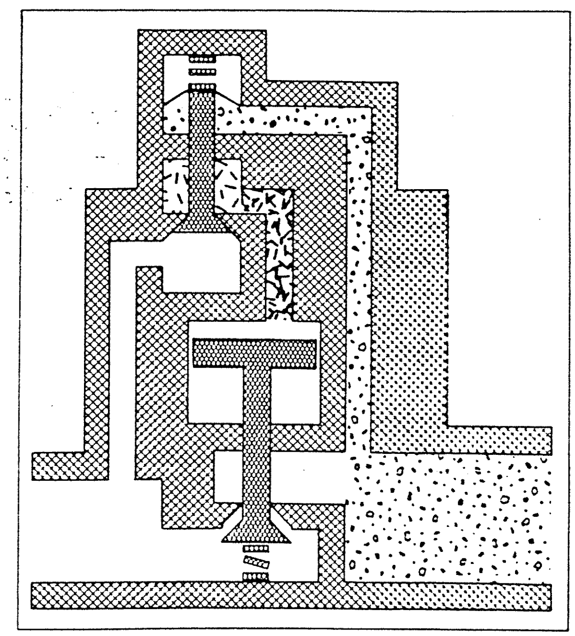

8. THE PRESSURE IN THE PISTON STEAM PORT FALLS,

Figure 1. (Continued).

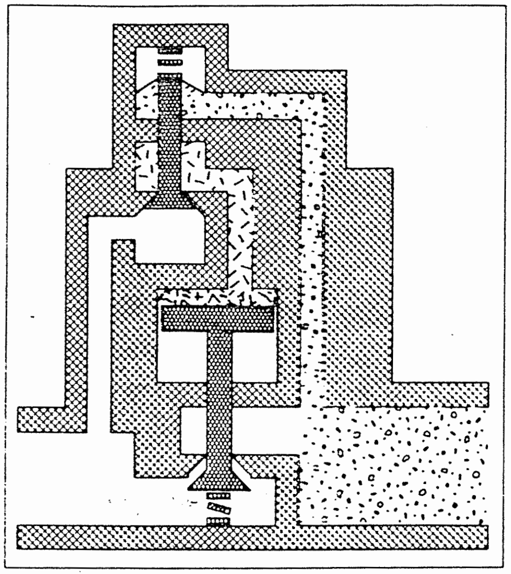

9. AND THE PRESSURE IN CHAMBER 4 FALLS.

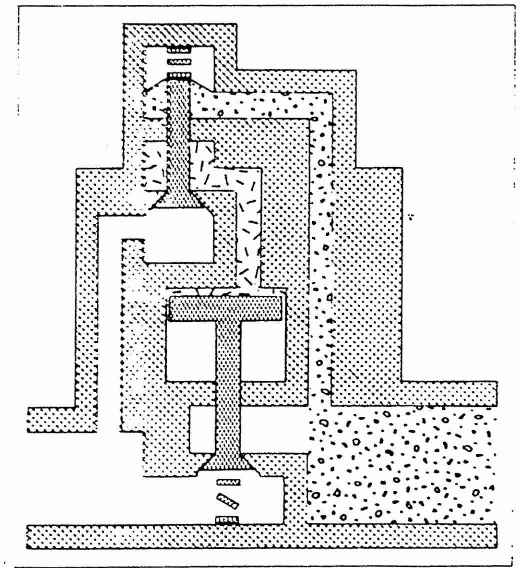

10. THE PISTON MOVES UP AND CLOSES THE MAIN VALVE.

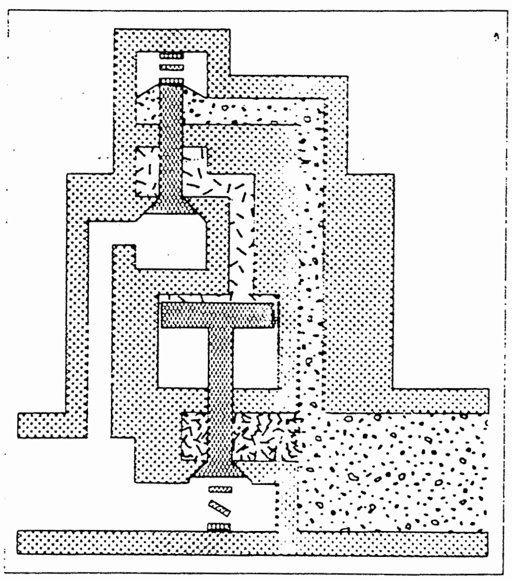

11. THE PRESSURE IN THE MAIN VALVE'S OUTPUT SIDE FALLS.

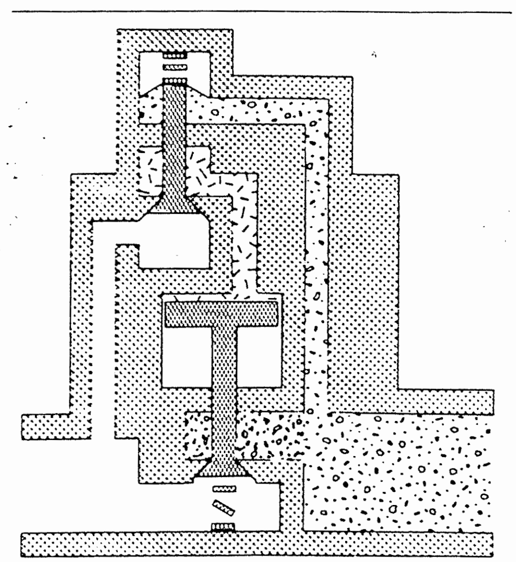

12. ALSO WHEN THE PRESSURE IN CHAMBER 0 RISES THE PRESSURE IN THE MAIN VALVE'S OUTPUT SIDE RISES.
NOTICE THAT WHEN THE PRESSURE IN THE OUTPUT PORT RISES IT MEANS THE PRESSURE IN THE MAIN VALVE'S OUTPUT SIDE RISES. HOWEVER, IT ALSO CAUSES THE SYSTEM TO ACT SO THAT THE PRESSURE IN THE MAIN VALVE'S OUTPUT SIDE FALLS. THIS MEANS IT EXHIBITS NEGATIVE FEEDBACK.

Figure 1. (Continued).

There are modeling systems that give explanations of complex systems. Forbus, 1981 describes the basis for an explanatory model of the reducing valve.

> There is a growing amount of evidence that human understanding of physical systems is based on qualitative models of those systems ... (An) explanation is comprised of a set of events, each describing a qualitative change in some part of the device. The explanation is linearized and describes how physical effect is passed from one component to another. It ignores the true temporal changes those things that are happening are happening continuously and simultaneously. The discrete and ordered nature of the events in the explanation is imposed.

An explanation from this system for the reducing valve is given in Figure 1. Forbus describes Figure 1 as follows

> Figure 1 presents an explanation generated using an underlying qualitative simulation. Each panel of the explanation is drawn from the actual computer display that a student sees. Successive panels denote successive states of the display ... The device described is a spring-loaded reducing valve, a common type of control device that serves to supply steam at a constant reduced pressure to a set of varying loads.
>
> It is instructive at this point to read the explanations. The first command, "Consider device spring-reducer-valve," indicates which device is to be studied. The controlled parameter is chosen to be the pressure in the output port by the command "Consider (>> press Output-Port V1)." The explanation is generated in response to the command "Suppose U." The system qualitatively simulates the effects of the pressure increase. By analyzing this qualitative simulation, the system discovers that the valve exhibits negative feedback. The student then asks for further clarification by typing "Explain." The system saves enough information during the qualitative simulation to reconstruct the sequence of events that led it to that conclusion. It is this saved event description that is turned into readable English and graphics and presented to the student one event at a time.

In Frame 2, the system notes that the pressure rises in the outport. Frame 3 notes the transmission of the pressure to the entire chamber. This is then transmitted to the diaphragm chamber. By this process, the system follows the changes through the system generating an explanation. It is no major stretch of the imagination to consider this modeling system with the U.S. economy. In such a context the following commands would make sense.

->> Consider Device U.S.-Banking-Industry

ok, call it B1
->>Consider (>>Cash-Flow Mexico B1),
ok
->>Suppose default

The system, if well conceived would then explain:

1. Bank of America debt/assert ratio falls

   First Bank of Chicago - Debt greater than working capital

2. Federal Intervention in First Bank of Chicago

While the Mexican example exists only as an example in this paper, it isn't unreasonable to expect a macro modeling system to produce this type of result.

Three aspects of the reducing valve simulation should be noted. The first is revealed in the organization of the explanation. Total effect is the propagation of the intial effect among simpler components of the valve. Each component receives input, then uses its state and the input to generate an output. For instance, the outport converts an increase on its out side pressure in Frame 2 to an increase in the entire chamber as shown in Frame 3. By localizing behavior to individual actors, such as a chamber, the effort required comprehend and model the system is reduced to specifying and understanding behavior of simplier components. It is only necessary to specify component's inputs, outputs, state and operating characteristics. While this last is the hardest to model, it is in understanding the operating characteristics of basic economic entities that economics earns its way.

The second is that the modeling system itself generated the explanation of the dynamic behavior. From the models of individual components, the valve's behavior could be tracked as a pressure change was propagated component by component through the valve. In this way, the user can examine how the assumptions of the component models generate the total dynamic behavior. The task of localizing assumptions with behavior is mechanized here, while in a nonexplanatory system it is a task of substantial difficulty.

Relating assumptions to behavior in an economic model is a nontrivial advantage of this form of modeling. Consider [Graifer, 1983] where the information about the banking system was summarized in literally hundreds of pieces of information. Tabulation and graphs are a partial answer, but how much more could have been learned, if the model and the modeling system provided reasonable analysis tools. This idea of filtering information in the context of management systems is discussed in Ackoff[1967].

The third aspect is the ease of modification and extensibility of models built in this environment. Since the only interaction of components is through communication with their neighbors, hooking a pipe on the outport of the valve and then asking what happens if the pressure is increased at the far end of the pipe is a simple matter. Similarly, if a model of the U.S. economy existed, but Mexico's debt wasn't considered important, it would just be a matter of changing those components of the model impacted by Mexico and rerunning the simulation to generate the banking example mentioned above. By localizing information, the model is easily changed and extended.

An Economic Example To Be

This example is drawn from [Cohen, 1981]. This paper is slightly tangential from macroeconomics, but closer to the macro case than the previous example. This work attempts to explore the space of what Orcutt calls the operating characteristic of a micro entity. There are a wide range of tools for exploring micro entities. This work is included because it is an economic attempt to use object oriented tools.

Addressed in the article is the question of how internal protocols provide an organization sufficient computational power to solve problems unsolvable by any individual component of the organization. Some of those issues are discussed in [Arrow, 1974]. The tool of choice to study this problem is simulation. As Cohen states:

> The models to be reported here are members of a very large family that can be studied using a system that has been developed for constructing computer models of organizations. Within this system one can study the effects of variation in task decompositions, patterns of incentives, level and types of environmental uncertainty, information flows and other organizational factors. The method captures much of the rigor of mathematical theory and much of the richness of verbal theory. These gains are purchased at the price of increased difficulty in the interpretation of results due to greater model complexity. For many purposes, however the trade appears to be an attractive one.

As examined with the reducing valve, there are techniques which partially overcome the last caveat on complexity. These techniques shift some of the interpretive problems, such as noting which assumptions drive the result to the modeling system.

To understand protocol questions, the internal interactions of an organization must be modeled. The experimenter wishes to set up various internal organizations, propose different problems and examine the resulting behavior. As Cohen continues to describe his experiments:

> The model individuals can be assembled into virtually any organizational structure that one wants to examine and provided with virtually any pattern of incentives and rewards for their decisions ... Other factors can also be set as an investigator's interests dictate. These include: friendship networks, agenda for meetings, rules for making collective decisions, noise, lags, or other environmental uncertainties, and organization precedents.

The problem used to demonstrate the paper's modeling system is a transportation problem with three factories and five warehouses. There are sixteen decision centers, one for each route and a general manager. The problem is to at cheapest cost maintain full warehouses, without exceeding factory output.

Subsets of decision centers conduct meetings in each period, with majority vote deciding new policy. To maintain constraints, the general manager, who also holds meeting, changes the reward structure for the decision centers, to punish nonfeasible solutions. The centers estimate rewards before acting, by calling an environment function with the proposed policy. Various organizations of meetings were compared to see if the organization reached the optimal allocation scheme.

In attempting to elucidate the internal structure of a component, this paper uses a simulation technique of interest. The presence of object oriented concepts is clear, but explanation are not present. They are absent for two fundamental reasons. First, the modeling system doesn't provide for them. Second, agents in the model are without realistic economic counterparts.

If this were a linear program, the fifteen decision centers would be variables to be solved for. They would be the optimum flow between each of the factory/warehouse pairs. From an institutional perspective, managers of the factories, warehouses, and a general manager would be the appropriate decision centers. The factory manager's task would be to prevent warehouse managers from exceeding the factories potential output. Each warehouse manager would decide order levels from each factory. Now explanations have meaning. A manager could be asked why he ordered x amount from factory Y, and an explanation would be of the form, "Well, I wanted to order Z, but the general manager said I couldn't."

Exploring the space of operating characteristics holds potential for interesting results. I believe it is possible with current technology to build a modeling system that would provide explanations of the reducing valve type for models exemplified in Cohen's work. This being so, it is time to examine this type of simulation technology.

Technical Details

Introduction

The concept of applying objects and explanatory systems was first systematically investigated in two electronic circuit instructional projects. They are EL devised by Stallman and Sussman [Stallman, 1977] and Sophie designed by Brown, Burton and de Kleer [Brown, 1982]. These projects consist of programs designed to tutor students on the behavior of specified electronic circuits. Both of these systems exhibited the explanatory power highlighted in the examples above. The following subsections consider various technologies for setting up objects and dependency networks.

Objects

In Cohen's model, each of the economically important agents was represented as an individual, an object. Similarly, in designing a macro model the economist must specify the the economically important agents and their interaction. The system, not the economist then must derive the behavior of the entire system. This is nicely described by [McArthur, 1982].

> When the programmer defines an actors's behavior, he need only be concerned with how the corresponding objects directly react to proximal inputs. When the program runs, however, complex and unforeseen distant effects of a local piece of behavior can be revealed because each local message transmitted can trigger others, and these in turn can trigger still others.

A number of computer systems support the concept of objects. One of the earliest was the simulation language Simula. Later object oriented simulation environments include Ross [McArthur, 1982], Thinglab, [Borning, 1979] and Flavors, [Weinreb, 1981]. Thinglab was written in a computer environment consisting entirely of object-oriented constructs, called Smalltalk, [Byte, 1981]. Instead of attempting to compare and contrast these systems, which has intellectual merit, I would rather look at a representative system, Flavors, describe it, and hint at how an Orcutt type of macro model could be set up in the Flavors environment.

Flavors was originally developed in LispM, a variant of the computer language Lisp, [McCarthy, 1962]. Other Lisps now have this facility. As Weinreb explains flavors:

> In this model, we think of the program as being built around a set of objects, each of which has a set of operations that can be performed on it. More rigorously, the program defines several *types* of object ... and it can create many *instances* of each type ... The program defines a set of types of object, and the operations that can be performed on any of the instances of each type.

But this is exactly Orcutt's vision of the economy.

> A distinctive characteristic of the type of model described in this chapter is that it contains components corresponding to microcomponents of the real socioeconomic system.
>
> In the real system we can identify several different types of decision units-individuals, families, firms, banks, labor unions, local governments. The model like the real system, contains a population of decision-making units composed of a *relatively small number of different types of such units and a relatively large number of units of each type*. (my italics)

So in Flavors one would say they are a small number of types of economic objects and a large number of instances of each type. In flavors, objects communicate by sending each other messages. As Weinreb describes:

> A terminology for the use of such generic operations has emerged from the Smalltalk and Actor languages: performing a generic operation is *sending a message*. The objects in the program are thought of as little people, who get sent messages and respond with answers.

Orcutt describes a first pass at a bank object. Of course banks have changed some since 1961, but since the theme of these techniques is ease of change, consider the flavor bank. Messages the bank would have to understand are:

1) Here's a check on your depositor
2) Deposit this money
3) Can I have a loan
4) Here's partial repayment of a loan

5) Your Required Reserve Ratio is

The instance variables, state of the bank would be:

1) Deposit liabilities of bank at start of period
2) Reserve assets of bank at start of period

The messages the bank would send to various other economic agents would be:

1) Here's the money for the check
2) Here is your loan
3) Cough up some money to extinguish some of your debt
4) Here is the interest rate you can have your loan at

Some of the internal characteristics are obvious. A valid check presented must be cashed. Other characteristics, such as which loans are made are less so. This is where the economic research lies, in understanding the laws, institutions, and standard operating procedures.

Given the objects, it is very natural to hook explanations onto the objects. A query to an object about an action causes an explanation as response. There is a nice symmetry here, as an object can send and receive messages to and from other objects in the model, it can similarly respond to messages sent from the experimenter.

Dependency Networks

Discovered in the two circuit modeling projects was that the simulation system had to keep explicit records of device interaction and have tools for reasoning with this information. These records are called dependency networks. In EL, a language called Conlan, [Forbus, 1981b] was used. The reducing valve example was modeled in this language. In Sophie, a number of techniques were used. Doyle, 1981 calls the keeping of a dependency network as reasoned deliberation.

> I have grown interested in a class of decision-making procedures I call collectively *reasoned* deliberation. Reasoned deliberation is a decision-making technique based on keeping careful records of parts of the process of decision-making, so that they can be reviewed and perhaps revised later in a flexible fashion.

The value of setting up a dependency network is also described by de Kleer, 1980.

> Dependencies are crucial in providing explanation capabilities, in transferring expertise, and in learning. In addition, these inference records can be used to maintain the currently active set of program database elements when new inferences, actions, or assumptions are made or changed. Further, dependencies can aid in controlling the program's actions by representing the inferential connections between assumptions, goals, and actions.

Dependency records themselves can be viewed as a history of the models action. This data must be organized so that the modeling system can easily find salient facts to organize an explanation around. The data structure and support facility discussed here is called a Truth Maintenance System (TMS) [Doyle, 1979].

There are two variants of TMSes. The original developed by Doyle and a second system, with different characteristics developed by McAllester[1980]. Here Doyle's system is considered. At the heart of the TMS are nodes and justifications. Nodes are actions, condition of the world or goals. Nodes can either be true or false. Two examples of condition of the world nodes are: "It is raining!", and "John said it is raining."

A justification consists of two lists of nodes called the in and the out lists. If all the nodes in the in list are true and all the nodes in the out list are false, then the node is made true. A justification for "it's raining" might be the justification with an empty out list and the one node, "John said it is raining", in the in list.

From this simple structure of nodes and justifications, it is possible to hook agents in and by making actions and goals nodes with justifications get explanations for actor's actions. The interested reader is referred to the Doyle and McAllester works, as it is impossible in this short space to give a full appreciation of the power of these ideas.

Random Observations on Explanations

Explanations are powerful tools for recognizing incomplete models. A not uncommon aspect of work done in economics is attempting to answer questions that are fundamentally impossible to answer from the given model. With an explanatory model this is more difficult, as the questions are asked directly of the model not in the sign of some parameter.

Consider de Kleer and Brown's [????] work with an electric buzzer. As a first pass, they modeled the buzzer as:

> The clapper-switch of the bell closes which causes the coil to conduct a current thereby generating an electromagnetic field which in turn pulls the clapper arm away from the switch contact thereby opening the switch which shuts off the magnetic field, thereby allowing the clapper arm to return to its closed position which then starts the whole process over again.

But if the modeling system is really to understand the circuit, this model is not sufficient. They point out the following questions that one would want the system to answer, but the current model can't.

a) What happens if we reverse the leads of the battery?
b) What happens if we switch the leads of the coil?
c) What happens if we remove the battery from the circuit?
d) What happens if we short the switch contact?
e) What happens if we make the switch arm lighter?

They then spend the rest of the paper refining their model to make it contain the essence of a buzzer. This refinement is also one of the economic tasks. Explanatory modeling systems provide tools necessary to convert our verbal intuitions into rigorous models.

Future Work

There are two major tasks that must be tackled to complete this investigation. The first is the implementation of a modeling system. The second is investigation of appropriate micro entities. At present, there are no modeling systems that encompass all the desirable properties discussed above. Parts of a number of systems exist, but it remains to integrate the parts. The theoretical modeling problems revolve around solving the problems of qualitative versus quantitative modeling. The most successful of the modeling systems, Conlan, can only do qualitative modeling.

In the reducing valve example, the modeler noted that system demonstrated negative feedback. In the study of may complex systems, including economics, it is often necessary to know more than this. The knowledge of behavior such as if the system follows a damped oscillation, exponential decay, or even if it is stable at all is vital. Unfortunately the qualitative model of the valve is not sufficient to encompass this behavior. Hollan and Hutchins, 1983 suggest a technique to overcome this problem they call quasi-qualitative modeling to generate explanations. Some of the other methodology problems in modeling continuous systems are described in Hendrix, 1973.

The other major task is to investigate micro entities in this formalism. This involves specifying inputs, outputs, and the difficult problem of understanding operating characteristics of the micro entities. For a first task, the implementation of a rational explanatory system of Cohen's model seems reasonable. I am currently working on that task. The next step is to look at real world macro systems. I feel that by facing these tasks jointly, producing the modeling system and studying a few micro entities, a reasonable product can be obtained. My hope is to use the explanatory simulation tools to converge on a realistic model. This learning cycle is described by Brown, Burton, and de Kleer as:

What?

That can't be!

Aha!

Bibliography

Ackoff, R., "Management Misinformation Systems," Management Science, Vol. 14, No. 4, December 1967

Arrow, K.J., "The Limits of Organization," W.W. Norton and Co., 1974.

Borning, A. "Thinglab--A Constraint - Oriented Simulation Laboratory," Stanford Computer Science Department Report, STAN-CS-79-746, July, 1979.

Brown, J., Burton, R., and de Kleer, J., "Pedagogical, Natural Language and Knowledge Engineering Techniques in Sophie I, II, and III, in "Intelligent Tutoring Systems, eds. Sleeman, Derek, Brown, Seely, Academic Press 1982.

Byte, August 1981: Entire Issue Devoted to Smalltalk.

Charniak, E. Riesbeck, C. and McDermott, D., Artificial Intelligence Programming, L.E. Erlbaum, 1979.

Cohen, M.D., "The Power of Parallel Thinking", Journal of Economic Behavior and Organization, Vol.2, No. 4., Dec. 1981.

Colby, K.M. and Smith, D.C., "Dialogues Between Humans and an Artificial Belief System" in IJCAI 69, pp 319-24.

Colby, K.M., "Simulations of Belief Systems", in R.C. Schank and K.M. Colby (ed) Computer Models of Thought and Language, pp 251-286, 1973.

Day, Richard H. "Comparative Dynamics and Development Policy Analysis: The Recursive Programming Approach in R.H. Day and I. Singh, *Economic Development as an Adaptive Process: A Green Revolution Case Study,* Cambridge University Press, 19??.

de Kleer, J., and Brown, J., "Toward a Theory of Qualitative Reasoning about Mechanisms," Xerox Parc, ????.

de Kleer, J., Local Methods for Localization of Failures in Electronic Circuits, Mit Technical Report memo, 394, 1976.

de Kleer, J. Doyle, Steele, G.L. Jr. and Sussman. G.J, Jr. "Explicit control of Reasoning" in Acm Symp. on Artificial Intelligence and Programming Languages, 1977. Also Artificial Intelligence: A Mit Perspective.

de Kleer, J. and Sussman, G.J. Jr, Propagation of Constraints Applied to Circuit Synthesis, Mit Technical Report Memo 485, 1978.

de Kleer, J. Causal and Teleological Reasoning in Circuit Recognition, Ph.D. Thesis, Mit Dept. EECS 1979.

de Kleer, J., "The Origin and Resolution of Ambiguities in Causal Arguments" in IJCAI-79 pp 197-203.

de Kleer, J. and Harris, G. Truth Maintenance Systems in Problem Solving, Xerox Parc, Technical Report Draft, 1979.

de Kleer, J. and Doyle, J., "Dependencies and Assumptions," AI Handbook, October, 1980.

Doyle, J. The Use of Dependency Relationships in the Control of Reasoning Mit, Technical Report Working Paper 133, 1976.

Doyle, J. Truth Maintenance Systems for Problem Solving, Mit, Technical Report tr-419, 1978.

Doyle, J. A Glimpse of Truth Maintenance in Artificial Intelligence: An Mit Perspective(P.H. Winston and R.H. Brown, ed).

Doyle, J. "A Truth Maintenance System, " Artificial Intelligence 12 (3) 1979.

Doyle, J. Reflexive Reasoner: A Model For Deliberation Action and Interspection, Ph.D. Thesis, Mit 1980.

Doyle, J., "Three Short Essays on Decisions, Reasons, and Logics, 1981 somewhere in, "Handbook of AI", 1981.

Fox, M.S. and Mostow, D.J., "Maximal Consistent Interpretations of Errorful Data in Hierarchically Modeled Domains," in IJCAI-77, pp 165-171, 1977.

Forbus, K. and Stevens A., "Project Steamer: III. Using Qualitative Simulation to Generate Explanations of How to Operate Complex Physical Devices", NPRDC Technical Note 81-25, August 1981.

Forbus, K., "Project Steamer: IV. A Primer on Conlan -- A Constraint - Based Language for Describing the Operation of Complex Physical Devices", NPRDC Technical Note 81-26, August 1981.

Graifer, D. Third Year Talk, February 1983.

Hendrix, G.C. "Modeling Simultaneous Actions and Continuous Processes, "Artificial Intelligence, 1, 1973, 145-180.

Hollan, J., and Hutchins, E., "Qualitative Graphical Interface to Quantitative Process Models," 1983.

Keynes, J. The Collected Works of J.M. Keynes, 1980.

Lucas, R., Jr., "Econometric Policy Evaluation: A Critique," Journal of Monetary Economics, 2:19-46, 1976.

McAllester, D.A., An Outlook On Truth Maintenance, Mit Memo 551 , 1980.

McArthur, D.A. and Sowizial, H., "An Object-Oriented Language for Constructing Simulations," IJCAI #7, pg. 809-814.

McArthur, D. and Klahr, P. "The Ross Language Manual",Rand Note n-1854-af, Sept. 1982.

McCarthy, J., et.al., "Lisp 1.5 Programmer's Manual," Cambridge, Mass, 1962.

McDermott, D., "Planning and Acting,"Cognitive Science, 2, 1978, 71-109.

Orcutt, G., Greenberger, M., Korbel, J., Rivlin, A., Microanalysis of Socioeconomic Systems: A Simulation Study, 1961.

Stallman, R.M. and Sussman, G.J., "Forward Reasoning and Dependency-Directed Backtracking in a System for Computer-Aided Circuit Analysis," Artificial Intelligence 9, 1977, 135-196.

Weinreb, D. and Moon, D., "Lisp Machine Manual", Symbolics, Inc., July, 1981, pg. 279-314.