254

DIAGNOSING SURPRISE MODEL BEHAVIOR:
A TOOL FOR EVOLVING BEHAVIORAL AND POLICY INSIGHTS

by

Nathaniel J. Mass
Associate Professor of Management
Director, System Dynamics National Model Project
Alfred P. Sloan School of Management
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

August 1981

To be presented at the
1981 System Dynamics Research Conference
The Institute on Man and Science
Rensselaerville, New York
October 14-17, 1981

ABSTRACT

Both in the incipient and later phases of developing a model, unexpected behavior is frequently encountered--that is, behavior which is at odds with the initial expectations of the model builder or client. The appearance of such surprise behavior immediately raises two possibilities: either the behavior is implausible, and the model therefore must be revised; or the behavior withstands scrutiny and reveals previously unappreciated aspects of the system. In either instance, the process of diagnosing and interpreting surprise behavior gives a powerful basis for model evolution and generating policy insights. But frequently, it is quite difficult in practice to discern whether the incidence of surprise model behavior reveals errors or suggests insights.

This paper is designed to contribute to the literature on model formulation, testing, and policy analysis, by discussing the criteria for diagnosing surprise model behavior. Several case examples are presented in which appropriate resolution of surprise behavior led to significant model improvements and/or behavior insights. Moreover, operational guidelines are presented to increase the likelihood of uncovering and successfully treating surprise behavior.

1. PROCESS OF GENERATING BEHAVIORAL INSIGHTS FROM SYSTEM DYNAMICS MODELS

Mathematical models are designed for different purposes and with correspondingly different underlying approaches. At one end of the spectrum, time series models do not purport to be built up from a causal structure, but rather are designed to extract in a sophisticated way secular trends, cyclical fluctuations, or other patterns of behavior that are contained in a series of historical data, and project those patterns into the future on an assumption of continuity. In the middle range, many statistical and econometric models are designed both to replicate particular sets of time series data, but also to capture key accounting identities and behavioral relationships that characterize the system structure. System dynamics models are noteworthy, if not completely unique, in insistence on a high degree of structural realism, and most important from the standpoint of this paper, a high level of explanatory content for relating system structure to observed behavior patterns, pathologies (that is, problematic behavior), and policy alternatives. In other words, a system dynamics model is intended, beyond objectives of forecasting or prediction, to yield operational insights about the feedback relationships that can produce or contribute to problems, can counteract the efficacy of policy interventions, or alternatively, can reinforce benefits of policy actions aimed at high leverage points.

What I have said thus far about the nature of models is not new, but simply reiterates the emphasis of system dynamics models on explanatory power in practical terms and at a managerially-relevant level. On the

other hand, what is not well-documented in the literature on models, and may not even be well appreciated by many model builders, is the process through which behavioral insights are arrived at using a model. Several articles in the system dynamics literature over the years have advised the model builder to begin a new effort with a clear "reference mode" that describes the time path of the problematic behavior being addressed, and also a "dynamic hypothesis" that lays out an initial theory of the principal forces that could interact to produce the reference behavior. From these initial constructs, and from additional data, descriptive information, literature, and theory that the modeler can bring to bear, a first model is developed. The model is then improved through successive rounds of analysis and consequent refinement. This progression of problem statement, initial hypothesis, first model, and successive model versions through iterative improvement, seems logical and is in fact frequently helpful as a guide to the phases of model construction. But in my experience, and I believe as well in the experience of many other model builders, the usual description of the model building process is much too orderly and free of tumult, and thereby misses one of the most important dimensions of the model building process.

In a variety of major modeling and policy studies in which I have been involved either as a direct participant or close observer, the understanding of real system behavior held by modelers and clients alike, and sometimes even the very concept of what the modeling study is about, has changed dramatically in mid-course as a consequence of surprising behavior revealed by an early model version. As a result, the course and stated objectives of the project were altered substantially from the initial problem statements to reflect the new understanding of system functioning. Let me give three examples:

1. Early on in the development of the System Dynamics National Model, originally around 1975, we assembled versions of a "standard production sector" representing a detailed behavioral model of industrial operations, to portray a consumer goods producing sector and a capital goods producing sector, and the interactions between them as a consequence of the demand and supply for capital goods. The resulting model revealed a 50-year fluctuation of large periodicity originating in the capital producing sectors, and with recurring sharp peaks in economic activity separated by broad values of depression and subsequent recovery to a new peak. Until that time, the main objective of the National Model Project had been seen as exploring in a national context issues of inflation and the life-cycle of economic development involving resource and energy depletion. Moreover, none of the project staff had been significantly aware of any process of recurring great depressions in real life that it might be important for the National Model to address. As the reasons for large amplitude fluctuating behavior with a 50-year periodicity were understood from the perspective of interactions producing the behavior <u>within the model</u>, the behavior began to appear more as a plausible managerial and economic phenomenon, and less as an aberration. Related literature and empirical data on long-term economic behavior were marshalled as an additional medium of refutation or support for the model behavior. Over time, we have gained increasing confidence that the originally unexpected model behavior in fact represented a significant set of public and corporate policy issues that the National Model could help to expose.* Exploration of causes and implications of long-wave behavior subsequently became a major thrust, although not the only objective, of the National Model project. This brief recount provides an example of surprise model behavior that brings to light a completely new phenomenon or pattern of behavior that the model could significantly address.

2. In an industrial research project based at MIT, a model was constructed to explain the sources of long-term decline in market share experienced by a major equipment manufacturer. Early versions of the model in fact generated from internal causes the reference pattern of declining market share. The replication of actual behavior experienced in real life was, of course, significant. But the more important question was: Why does the behavior arise and what policies could be exercised to reverse the declining trend? On this question, the model suggested that declining market share occurred during periods of low overall industry demand for the company's product; and moreover, the prime cause of loss of market share during these times was a high delivery delay (meaning lack of

* Mass, Nathaniel J., and Peter M. Senge, "Behind the Clamor for Reindustrialization," <u>Technology Review</u>, August, 1981.

availability) of the company's product. The timing relationship between overall industrial downturns and loss of market share for this particular company was unanticipated; nonetheless, scrutiny of company records confirmed the correspondence. However, the incidence of high delivery delays during a period of low industry demand seemed intuitively implausible, implying that the company's product was least available when no one wanted it. However, review of the company's records again revealed that precisely this had been occurring. The resulting insights into system behavior changed dramatically the directions for further model development and ongoing policy analysis using the model. This case thus provides an example of a situation where a model did indeed replicate an anticipated reference pattern of behavior, but where the cause of that behavior was almost entirely unexpected.

3. A preliminary version of a financial model developed for a major bank suggested that the bank's policies for paying out cash and stock dividends, as they had been described by executives in the bank actually responsible for those decisions and thereby incorporated in the model, could substantially reduce, or even eliminate, growth in earnings per share. At the same time though, data showed a clear growth in actual earnings per share experienced by the bank. The process of reconciling the factual circumstances and the initial model output highlighted several key relationships that had not been mentioned by the executives as important considerations in their dividend decision process and were therefore not included in the initial model, but which later in fact appeared responsible for the continued growth in earnings per share. In this example, then, the process of understanding surprise model behavior led to important changes in model specification, as well as to the realization of potentially conflicting elements of the managerial decision process in the actual firm.

Each of the three examples cited above shares several common elements. First, behavior emerged from a preliminary model that was surprising to all participants in the modeling process, including both model builders and clients. Second, in each case, the surprising behavior could not immediately be rejected as being either factually incorrect or implausible as a prediction of future behavior. Third, the process of interpreting the surprise behavior required the development of new frameworks for viewing available data and knowledge about behavior and management policies. Fourth, the process of resolving the surprise model behavior led to appreciable shifts in the basic thrust of the modeling and policy analysis effort.

In general terms, the appearance of surprise behavior from a model immediately raises two possibilities: either the behavior is implausible, and the model therefore must be revised; or the behavior withstands scrutiny and reveals previously unappreciated aspects of the system. But even more subtly, it frequently appears quite difficult in practice to discern whether the incidence of surprise model behavior reveals errors or, alternatively, suggests insights. When confronted with surprising model behavior, and especially behavior that appears at odds with initial impressions or hypotheses about system operation, many model builders would be tempted to assume that the model is behaving unreasonably, and to "cover up" the surprise behavior through parameter changes or structural modifications. On the other side, I have seen a variety of instances where a model builder will accept surprise model behavior as providing a source of significant policy insights, where in fact the behavior points up flaws in basic model design. In many respects, it is the very behavioral richness of system dynamics that is the source of this "identification" problem.

This paper attempts to contribute to the literature on model formulation, testing, and policy analysis by discussing the criteria for diagnosing surprise model behavior. In particular, Secton 2 presents general guidelines as well as specific categories of tests for increasing the likelihood of uncovering and successfully treating surprise behavior.

2. GUIDELINES AND TESTS FOR RESOLVING SURPRISE BEHAVIOR

2.1 Follow Up All Unanticipated Behavior to Appropriate Resolution

The most basic guideline for effectively treating surprise model behavior is that whenever such behavior is encountered, it should be assessed, understood, and followed through to an appropriate resolution, whether in terms of model improvements or adoption of new perspectives on system behavior. The model builder must adopt perspective that views the encountering of surprise model behavior as a significant opportunity to be capitalized upon. In contrast, the inexperienced model builder who confronts surprise or anomalous model behavior, is prone to pursue parameter combinations that make the anomalous behavior less evident, or simply to dismiss the behavior as being outside of the intended use of the model. One of the significant aspects of system dynamics models, whether in the corporate policy or public policy realm, is that a well-structured model will frequently come to have uses beyond those originally envisioned. In other words, a effective system dynamics model is probably best viewed as a multi-purpose or general purpose model, even if it was originally designed only for narrower uses.

An important dimension of resolving surprise model behavior is to balance model-based results, empirical data, and client knowledge about system behavior. As seen in the brief examples cited in Section 1, some of the most important insights into real system behavior can arise from model results that at first appear to be at odds with knowledge of the real system, but which in fact suggest important new interpretations of perceived facts. The system dynamics model builder/analyst/consultant can have an important, though difficult to play, role as a change agent. On the one hand, the model builder must recognize and accept the possibility that much of the surprise behavior encountered, particularly in early phases of model development, may point up defects in a model more than particular insights. But on the other hand, especially as the model builder is more experienced and more knowledgeable of the real system, and as the model improves progressively over time, the likelihood increases that surprise model behavior points to new ideas that bear on policy formulation. In order to play the role of change agent effectively, the analyst must be sure that he has a broad appreciation of available data, literature, and managerial experience (including effects of previously implemented policy changes), and he must be sensitive to the actual organizational pressures and relationships. But on the other hand, the model builder must delve sufficiently deeply into the sources of model behavior to be able to explain in novel, although practical terms, the forces that may produce unexpected results in the actual system. Especially in a consulting (as opposed to research) environment, a system dynamics analyst can be rendered ineffective if he appears unaware of existing data and points of view about organizational behavior, and if he is unable to relate to that existing knowledge in a conscious and creative way. From this standpoint, a danger is that the consultant be "captured" by the client, so that significant new policy perspectives that could emerge from the model are not successfully cultivated to fruition.*

* For example, see Charles W. Gibson, "Using Models in Financial Planning," Journal of Business Strategy, Vol. , No. 4, 1981.

2.2 Importance of A Priori Expectation of Model Behavior

If we accept the importance of surprise model behavior as a diagnostic for model improvement and policy formulation, then a basic point emerges. Appearance of "surprise" behavior implies a discrepancy between results actually produced and previous expectations of those results. Thus, it is absolutely essential that the model builder have a strong a priori expectation of model outcomes, to establish a baseline against which surprise model behavior can be recognized through the appearance of a discrepancy that evokes "cognitive dissonance."

In discussing the form that an a priori expectation of model behavior may take, I believe it is useful to distinguish three classes of models. I define a Type 1 model as a model that is addressed to a well-established set of problems or circumstances observed in the past. A Type 1 model thus corresponds most naturally to the "classical" statement of purpose for a system dynamics model, where a historical reference mode provides a basis for model development. The historical reference mode may portray declining market share in a corporation, results following implementation of a particular public policy that showed changes in the opposite direction from that intended, or similar phenomena. Such an historical reference mode provides the a priori expectation of results. If the model does not replicate the historical reference mode, then the model needs to be revised, while admitting the possiblity that model results may cast the historical circumstances in a new light.

A Type 2 model can be defined as being addressed to a defined set of policy issus rathern than to a particular historical circumstance. For example, I and several others have been working on a system dynamics model for a foreign government to help anticipate effects of alternative strategies for oil development and oil export. This particular nation has had no history of significant oil export in the past, so it is clearly not relevant to draw a historical reference mode to guide model development. Of course, other nations have gone through stages of oil development, with varying degrees of success or failure, and in some sense these alternative experience curves comprise a historical base. Experiences of nations or organizations other than the client organization being studied may have relevance to the kinds of futures that should be encouraged or avoided for a client. But I see such experiences as forming somewhat more equivocal and less direct reference for model development than established past history for a client organization. Thus, while there is probably no sharp dividing line between the reference point for starting a Type 1 or Type 2 model, the extent to which the model must replicate reference behavior is clearly different.

For developing a Type 2 model, then, experiences of related organizations or systems may comprise part of the a priori expectation of system behavior. But frequently still more important is a priori expectation of possible effects of implementing the policies that the model is being designed to analyze. In the case of the oil policy model, that a priori expectation includes guesses as to the answer of questions such as the following: What would be the effects on inflation, employment, and other economic aggregates of substantially higher or lower levels of oil

exports? What would be the short-run and long-run effects of setting the domestic price of oil at world levels, or maintaining it at substantially lower levels through explicit or implicit government subsidies? For a Type 2 model, it is important to recognize that the a priori expectations of behavior or policy impacts are not established facts that the model must replicate as a basis for validity. Rather, they comprise expectations that enable a rigorous comparison of eventual results with the results that were originally expected. If a difference arises, then the appearance of that difference calls for some resolution. The model builder and client can either adhere to the a priori expectation and elect to modify the defects in the model that cause the model to fail in producing the expected results; or alternatively, the a priori expectation may be consciously revised to conform to the new understanding of system interactions. The importance that I am attaching to the a priori expectation of behavior may seem exaggerated to some readers. But in my experience and that of others in building models of systems that lack a "hard" historical reference mode, both assessment of model validity and generation of policy insights can be impeded by the absence of criteria, however transient and subject to revision, for evaluating the plausibility and significance of model behavior at any point in time.

A Type 3 model in this classification is an extension of the Type 2 model to systems for which there is again a list of policy issues that the model should address but only a weak historical precedent for the interactions being modeled. For example, I am now involved in developing a growth strategy for a new company. The company will be a subsidiary of an existing company in a line of business that represents a significant extension of the parent company's traditional areas. Moreover, the chief executive and much of the management team for the new subsidiary has not yet been appointed. The parent company has clear ideas about the criteria for success or failure of the new subsidiary, meaning that there are identifiable patterns of behavior that would represent varying degrees of success or failure. Moreover, corporate management has a list of potential policies for the subsidiary whose effects it would like to understand better. Because the new subsidiary company has no past history, there is certainly no historical reference mode to provide a point of departure for the model building effort. Furthermore, there do not seem to be any histories of allied companies or companies in the same general area of business to the subsidiary to comprise a sharply defined historical reference mode. In my view, a Type 3 model, while much more difficult to construct and to evaluate than a Type 1 or even Type 2 model, is no less an appropriate subject for system dynamics. Interactions can be identified, albeit somewhat hypothetically, that take the form of feedback loops and tie the new company to its market, and significant policy issues can be identified for analysis using the model. In developing such a model, having a priori expectations of possible patterns of system behavior and effects of possible policy changes becomes extremely important for tying the model to reality.

My point can be summarized as follows: system dynamics models can start from different points and different degrees of historical precedent. An historical reference mode is by no means a requirement for beginning a model. Nonetheless, the model builder must have sharp a priori expectations about possible model results. These a priori expectations may

take different forms, and can be articulated in more or less creative ways. Without such expectations, there is no basis for judging when significant surprises or anomalies appear in the model building process that should motivate changes in the model or in the modeler's or client's viewpoint.

2.3 Confirm All Behavioral Hypotheses Through Appropriate Model Tests

When surprise model behavior is encountered, the model builder must identify why the model produces the unexpected results. The question of why a model produces certain patterns of behavior can always be answered with enough time and effort relative to the model framework. Once the model behavior is understood, the realism of both the behavior and the underlying mechanisms must be challenged against corresponding behavior and structure in real life.

The mechanisms that produce surprise model behavior may take several forms. For example, the model builder may identify a positive feedback loop that was not previously recognized to exist that can cause major excursions in particular variables. Alternatively, a model builder may come to identify a negative feedback loop that counteracts policy changes. As still a third possibility, while not pretending to offer an inclusive list, the model builder may identify a combined structure of positive and negative feedback loops that can diminish the effectiveness of the policy intervention, while calling for more and more of that intervention with ever-declining efficiency over time. Whatever set of feedback structures the model builder may hypothesize to yield the unexpected behavior that is observed in the model, it is important for the model builder to devise appropriate behavioral tests to confirm or reject his hypothesis.

The process of evaluating the hypothesis about sources of behavior will always in some way involve segmenting and neutralizing the forces in question. For example, if the model builder feels that varying prices are an indispensible part of a fluctuating mode of behavior in an economic model, he may artificially force price to be constant, through the equivalent of a full price control. Alternatively, if behavior is hypothesized to result from fluctuating adequacy of household liquidity and consequent interactions between consumer purchases, employment, and wage payments, then adequacy of liquidity may be held at a neutral value, or assumed to exert substantially greater or lesser impact on consumer purchases. Such tests are necessary to verify that direct link from liquidity to consumer purchases is in fact a sensitive point in producing the behavior in question.

Sometimes, a given hypothesis relating behavior to underlying causes can be tested from different points of view; alternatively, a given hypothesis about behavior may be difficult to test in a natural and operationally significant way. For example, in working on the System Dynamics National Model, we have variously tried to isolate particular financial or real mechanisms underlying economic behavior. As a concrete example, suppose we hypothesize that limits on the availability of short-term debt from the commercial banking system are an indispensible part of producing a particular fluctuating mode. In testing this kind of hypothesis, we have at times neutralized the availability of credit from

the banking system, so that eligible loan requests are always met with no limits from supply. But if supply of credit is unconstrained, then debt may tend to fluctuate over larger ranges, with one possible outcome being that the control of debt variations is shifted from pure availability considerations, to limitations from permissible debt to asset ratios. The analyst must then decide whether the relaxation of constraints on "availability" of bank funds means either simply the elimination of supply constraints, or whether it implies elimination of both supply constraints and limitations from credit worthiness on the eligibility of loan requests. Different ways of casting a particular behavioral hypothesis test may yield different results, so it is important that the analyst consider carefully the alternative ways of implementing a given hypothesis test, and the full dimensions of evaluating a hypothesis through a given channel of effect.

Despite the attendant difficulties in formulating hypothesis tests appropriately, I believe that it is extremely valuable to insist on behavioral confirmation of all hypotheses regarding sources of surprise model behavior. I have on any number of occasions spent substantial time diagnosing an unexpected behavior pattern, eventually arriving at an apparently satisfactory explanation, but failed to develop appropriate confirmatory tests only to discover much later that my hypothesis did not stand up as a convincing explanation of the results. In this respect, the model builder suffers the same difficulty as the policy maker: until his intuition is sufficiently honed through experience in the real system or with models of the real system, then careful analysis of model output may fail to distinguish true causes of behavior from concomitant symptoms. I find that the short-term effort devoted to model-based verification or rejection of behavioral hypotheses is well worth the time and can diminish the likelihood of long detours down particular research and policy directions that eventually are perceived as misdirected.

The remainder of this paper is devoted to developing a preliminary list of specific tests that may be helpful in increasing the likelihood of encountering and resolving surprise model behavior.

2.4 Identifying Symmetry of Policy Response

One test that is extremely valuable in model testing for revealing unanticipated behavior is to evaluate the symmetry of model response to changes in upward and downward directions. For example, if an analyst is testing a production model with upward step functions in consumer orders, he should equally test the response to downward steps, representing declines in demand. As one example of where such testing proved valuable, several years ago I was working on a financial model of how corporate policies for making investments, paying out cash and stock dividends, and pricing new issues, affected the firm's average cost of capital, and thereby the attractiveness of new investments. I started out by isolating the part of the model that represented the stock market and subjecting it to different external inputs from the standpoint of that module. These inputs represented, for example, assumptions about alternative patterns of growth in earnings per share. and investor confidence throughout the economy. As one test, I started the initial value of stock price above what I knew to be its equilibrium value, ran the model, and observed a rapid decline back to the equilibrium price, as was expected. It was only

substantially later, that I performed the opposite test, of starting the initial stock price below its equilibrium value. The model exhibited an unexpectedly slow response over five to six years through which stock price rose back toward its equilibrium value. The eventual source of this slow response turned out to be a formulation for "speculative risk," through which investors were assumed to evaluate whether or not a stock price had been driven to a value that could not be sustained on the basis of financial realities such as growth in earnings per share, but which formulation was not sufficiently robust to input conditions. In particular, speculative risk in this early model version was formulated as a simple function of the perceived growth in the stock price in relation to the growth in earnings per share. Certainly, in an equilibrium situation of stable earnings growth, stock price should grow at the same rate as earnings per share and consequently cash dividend per share. But in a disequilibrium situation that is not characterized by stable growth, faster growth in share price than in earnings does not necessarily indicate speculation. For example, as encountered in the test when the price per share was started below its equilibrium value, the total stock yield, meaning both cash dividend yield and expected capital gains, would exceed its equilibrium value. Higher total yield, in turn, should drive up the stock price. But in the faulty model formulation, increasing stock price was being taken as a signal of speculation, which was tending to increase the perceived risk associated with holding the company's shares, and thereby generating a negative feedback pressure to restrain the rise in share price. The net result was a slow rise in stock price in which the pressures of undervaluation were tending to drive price upward, and the faulty perception of speculation was exerting downward pressure to compensate in part. Once the source of this behavior was understood, a broader and more robust formulation for speculative risk was developed that encompassed not only the appropriate steady-state response, but also the response to disequilibrium pressures.

I include this detailed example here because in my experience it is very easy to be caught in the trap of subjecting a model to an insufficiently narrow set of tests, to progressively "tune" the model to fit the limited input circumstances, and thereby to miss potential defects or behavioral insights that would be quickly revealed by a different set of test conditions. (Several other of the tests described in this chapter treat analogous problems of testing that is too limited.) Sometimes, asymmetric response to different directions of model input or initial conditions may be defensible. To take a simple example, if utilization of capacity (meaning, length of work week, number of work shifts, and efficiency of utilization) is more easily reduced than it is increased, then model response to a growing demand should be slower than to a falling demand, starting from the same level of full utilization. Thus, examining symmetry of model response to upward and downward conditions may sometimes help to illuminate asymmetric time constants or other important behavioral mechanisms, besides revealing potential defects in model formulation.

2.5 Testing Large Amplitude vs. Small Amplitude Response

A form of testing that can be used to precipitate surprise behavior is to examine model response to both large and small amplitude input variations. Very frequently, different adjustment mechanisms may be

involved in regulating large and small departure from normal operating conditions. For example, consider an oligopolistic industry such as automobile manufacture. Small excess inventories of automobiles may lead producers to curtail production in order to liquidate inventories, with little force for concessions on price compared with list price, much less to change list price. On the other hand, large excess inventories such as are now being encountered in the automobile industry may force both price and output responses. We thus see an example of emphasis on different corrective mechanisms depending on the degree of disequilibrium.

Sometimes, tests of model response to large amplitude input variations can reveal important nonlinearities whose omission can lead to implausible model behavior. Alternatively, model tests in response to large amplitude input variations may make evident important behavioral mechanisms that have policy significance and that may have been obscured when model behavior was examined only over a narrow input range. For example, I have recently been working on a model of banking intermediation activity in the context of a national financial system. When the model is run with high and slowly growing rates of inflation, bank profits are seen to rise along with the higher nominal loan demands and interest rates produced by inflation. On the other hand, if government deficits accelerate rapidly, leading to rapid increase in inflation, then bank profitibility can be reduced substantially, even to zero, if a substantial fraction of loan revenues are not indexed to changes in the cost of funds. This important distinction between the effects on profitability of high rates of inflation versus large increases in the rate of inflation was always latent in the model results, but was not appreciated until a test of large amplitude response made the differing patterns of behavior evident to the eye.

2.6 Testing Policies Entering at Different Points of the System

Another important principle in model testing is to evaluate a wide range of policies whose direct effect occurs at different points in the system being modeled. For example, in testing the production sector of the System Dynamics National Model, we have tried to evaluate response of production, price, liquidity, and other variables to external (meaning external to the sector) assumptions about consumer demand, level of interest rates, availability of short-term and long-term credit, delay in filling vacancies, delivery time for capital goods, national productivity trends, increase in labor costs, and other inputs.

The model builder is often tempted to "fine tune" model response to a particular set of input conditions because the resulting outputs yield a tangible result that matches real system behavior. But I believe such testing is misdirected. In our experience in working on the National Model, we find that a more balanced testing approach of evaluating model behavior in light of various stimuli at different points of intervention is much more likely to reveal flaws or suggest insights from surprise behavior. Such emphasis is called for even if the primary objective of model analysis lies in understanding response to a particular set of input conditions, such as labor shortage or faltering productivity. Testing model response to alternative input circumstances can point up model defects or even highlight important mechanisms that bear on the primary

purpose. For example, a recent corporate model I have been developing suggests particular marketing and pricing policies to achieve a better customer mix and improve profitability. But the potential desirability of these marketing and pricing policies became most evident when realistic limitations were imposed on the expansion of primary capacity. Thus, model testing should never be limited inordinately to the immediate area of the model surrounding the point of primary issue concern.

2.7 Testing Different Patterns of Behavior

Many system dynamics models have the potential for generating more than one basic pattern of behavior. For example, the System Dynamics National Model can generate fluctuating modes ranging from the 3-7 year business cycle to the 50-year long wave, as well as a separate mode of sustained inflation from monetization of government deficits. Similarly, the stock market model mentioned earlier can exhibit both an internally generated stock market cycle as well as patterns of long-term growth or decline in share price. The importance of multiple modes of behavior can be two-fold. First, symptoms are easily confused between the separate modes, thereby complicating policy formulation; and moreover, different policies may be appropriate for treating the separate modes. If model testing is insufficiently broad, the modeler may not even be aware that a model is capable of exhibiting separate patterns of basic behavior.

Because many system dynamics models can exhibit more than one basic mode of behavior, it is important to adopt a balanced approach for testing than can expose behavioral implications of the structures that underlie different modes. In addition, the model can be subject to external inputs representing different modes of behavior. An example can help to illustrate some of the pitfalls of overnarrow model testing. In the early stages of developing the production sector of the National Model, we devoted a great deal of time to refining the model structure relative to the three cyclical modes that the model was exhibiting--the business cycle, the intermediate Kuznets or construction cycle, and the 50-year long wave. However, upon subjecting the production sector model to external growth in demand, representing a secular trend as opposed to a fluctuation around a constant average value, whole new problems were disclosed. One formulation that appeared defective and thus in need of improvement involved the ordering function for capital equipment. The original model formulation assumed that capital ordering depended first of all on the replacement of depreciating capital; second, on a correction for the existing stock of capital and for the outstanding amount of capital on order; and third, on a growth term to increase the capital stock in line with expected future growth in sales. The formulation thus recognized that in order to maintain a "neutral" growth path in the face of growing final demand and with no ongoing shortages of capital equipment, capital must be expanded at the rate of growth of demand even though at each point in time, no discrepancy exists between desired and actual capital stocks. However, when the model was run with growing demand, a shortage of capital arose, manifested in above normal delivery delays, lower than average inventory coverage, and relatively high price and profit margin. Analysis of these biases away from a neutral steady-state growth path eventually revealed that the order rate function needed to be extended to provide for expansion in the number of capital units on order in line with the rising final output demand, as

well as expanding capital stock along with growing demand. This formulation defect might have been identified through careful enough scrutiny of other model-generated modes, such as the large excursions in capital ordering that took place over the course of a long wave. Nonetheless, the point remains: behavioral implications of different structures embodied in a model may be best revealed through a broad testing approach that attempts to isolate different behavior patterns and thereby bring to the foreground structural problems that may be latent but undiscovered in other modes of model behavior. Such a testing approach can also contribute to policy analysis by revealing the relative efficacy of a given policy under different circumstances and in different modes of behavior.

2.8 Evaluate Both Real and Nominal Changes

A number of system dynamics models have now been built that treat money and financial markets, spanning from issues of corporate pricing and financial policy to evaluation of policies for controlling national inflation. These models extend the thrust of earlier system dynamics models, such as the models underlying Urban Dynamics and World Dynamics that emphasized only real or physical changes, to treat movements in real activity, as well as nominal changes in prices and financial variables.

Where both real and nominal changes are combined in the model, it is important to isolate the behavior resulting from each set of processes. An example may help to illustrate the process and the kinds of results that may emerge. In developing the National Model, we have attempted to follow an incremental approach of successively activating various model sectors and processes. At one stage, we were working with a combination of production sectors, a household sector for savings and consumption decisions, and a labor network connecting the production and household sectors. The model was being tested with active price and wage equations. However, in order to simplify the analysis conditions, we neutralized effects of money flows between production and household sectors by maintaining neutral adequacy of liquidity in each sector; by neutral adequacy of liquidity, we mean that through implicit borrowing and repayment of debt from the financial system, each sector is able to maintain sufficient liquidity to support its desired expenditure rates dictated by need for output, relative prices and wages, balance of factors of production, and other similar factors. In this configuration, we found that real variables tended to exhibit the same basic cyclical modes that had been seen in analyzing the physical system alone. However, wages and prices showed a strong tendency to drift together, either in the upward or downward direction and seemingly unpredictably as to what direction of drift would result in a given computer output. After careful study, we concluded that such behavior was in fact both necessary and appropriate given the model configuration being assumed. With an implicitly unlimited supply of credit, money supply in the model economy was likewise unlimited and indeterminate. But if money supply does not have a determinate steady state, then neither do absolute wage and price levels. In other words, suppose that from an initial equilibrium point, all prices, wages, and money stocks were doubled. Then producers would be indifferent to the higher absolute wage and prices levels, since their relative profitability would be the same, and since there would have been no shift in the relative

price/wage ratios that dictate factor intensities. Likewise, the household would be able to support the same real flows of purchases, albeit at higher prices, given twice the money level and twice the wage income and dividend stream from the production sectors. In fact, such behavior relates directly to the monetarist theories of inflation which argue that the absolute level of money supply is essential in determining the absolute level of prices, as distinguished from the behavior of relative prices (such as relative price-wage ratio).

In this instance, then, the appearance of surprise model behavior did not suggest a flaw in underlying model formulation, but rather suggested a policy insight: that control of money supply has an important impact on the absolute price level, and therefore on inflation rate. Moreover, the model results provided a vantage point for relating to an economic literature covering both theoretical and policy issues that were of obvious relevance to the project, but whose connection had never previously been so clear.

In summary, then, it is important that models that incorporate both real and nominal processes of change be analyzed so as to isolate the relative behavioral contributions of each. For example, an important issue is to understand the extent to which real and nominal changes are either separable or intrinsically connected in a given mode of behavior. Such evaluation can be performed by a variety of experiments that control the environment surrounding nominal changes: for example, prices and wage levels can be held constant; availability of credit from a financial sector supplying money capital can be held neutral so that all eligible demands are met; adequacy of money or liquidity in various sectors can be neutralized, and so on. In order for system dynamics models of economic and financial activity to exert a significant impact on the literature and on public policy, it is important that they add substantively to understanding of how real and nominal changes in the economy are related and how those processes contribute to problems such as persistent inflation, high interest rates, and periodic credit shortages.

2.9 Isolating Uniqueness of Equilibrium/Steady State

Although most system dynamics models are designed to understand disequilibrium or transient behavior of a system, equilibrium analysis of model properties can still be revealing. For example, it is a common technique to initialize a model in equilibrium, and then perturb the system through controlled exogenous inputs to understand the transient properties such as periodicity, frequency response, and damping ratio. As an additional example, in the Urban Dynamics book, Forrester uses a form of comparative equilibrium analysis of computer simulations to see how proposed policies for urban revival affect the long-run equilibrium of the city, meaning population densities, population mix, unemployment rate, and other similar indicators.

Although a model such as the Urban Dynamics model is fairly complex, containing more than 25 state variables, it has an important property that bears on both model testing and policy evaluation: the final model equilibrium for a given set of policies, constants, and exogenous inputs, is independent of the initial values of level variables such as numbers of

business firms and population levels. For example, Forrester develops a particular set of urban revival policies by applying the revival policies to a city that has already reached an equilibrium characterized by high unemployment rate. He then shows that the same policies could also be applied to a young and growing city, with identical long-run consequences:

> Policies used for reviving a decayed area should, if continuously applied, prevent decay. With rare and very special exceptions, the ultimate equilibrium in a system does not depend on the system's history. It depends only on those policies and system parameters that act during the period when equilibrium is being established. This means that the revival policies....could be applied to a city throughout its growth period and should produce the same final equilibrium conditions as they do when applied to a stagnant city... New and more satisfactory urban-development policies can be initiated at any point in the growth-maturity-stagnation cycle. Transient conditions will be affected, but the final equilibrium depends on the policies themselves and not the initial conditions at the time the policies are implemented.*

As Forrester asserts in the above quotation, the majority of state determined systems, of which system dynamics models are a subset, have the dual properties that:

a) model equilibrium is independent of initial conditions for given policies and exogenous inputs (including constants); and

b) for a given set of initial conditions, model equilibrium is independent of the time paths of the exogenous variables before those exogenous variables take on eventual constant values under which a system equilibrium is reached.

* J.W. Forrester, Urban Dynamics, (Cambridge: MIT Press, 1969), pg. 106.

Although the above conditions characterize most system dynamics models, there are models and systems that have more complex properties. For example, consider a system consisting of a flat table surface and a marble that is placed on the table. If the table is truly flat, then the marble can come to rest at an unlimited number of equilibrium points, each identical to the initial conditions. Although systems with multiple equilibria exist, and models of these systems can be usefully developed, the models will have unique properties. In my experience, it is a common hazard for model builders to produce a model that exhibits a non-unique equilibrium without appreciating the unusual nature of the results, and therefore failing to question whether the results stem from a defect in the model that fails to capture the pressures characterizing the real system equilibrium, or whether in fact the real system has unique properties that have important implications for policy design.

As one example of surprise model behavior that raised important issues about equilibrium properties, several months ago I was shown the behavior of a fairly simple teaching model that was designed to give students exercise in model formulation and analysis. The subject of the model was addict-related crime in an urban neighborhood. The model showed the surprising result with apparently significant policy implications, that an increase in police effort to control incoming drug supply, represented in the model as a step function in an exogenous level of police effort, yielded fewer addict-related crimes in the short run, but a sustained higher crime rate in the long run. As I have tried to argue throughout this paper, the appearance of such surprise behavior calls for careful scrutiny of whether the surprise behavior reveals model flaws or

alternatively policy insights. In exploring the reasons for the unexpected policy result, we examined another model simulation in which a temporary surge in police effort was represented, characterized in the model as a step-up in police effort at one point in time, followed by a step down to the original level of police effort sometime later in the same simulation. We discovered that the level of crime did not return to its original equilibrium value (an effort was made to confirm that the initial condition was indeed a sustainable equilibrium), even though all exogenous inputs including police effort returned to their initial values. The question of whether such results correspond to real life is equivocal, and probably cannot be answered on purely logical grounds. But such behavior certainly raises important issues from the point of view of model analysis and policy evaluation. For example, if model equilibrium is dependent on the time paths of exogenous variables such as police effort, then it is probably not possible in principle to say that a higher eventual level of police effort either raises or lowers the crime level: the outcome may be sensitive to the exact time path of police effort before reaching the final higher level, to the initial values of systems levels, and to the initial extent of disequilibrium. An analytical study of this simple crime model using basic algebra revealed that an unlimited number of equilibrium points could be reached as long as the ratio of drug supply to addict population attained a particular value. Thus, for example, there were no forces in the model that limited the fraction of the local population that was susceptible to drug addiction: the fraction could settle anywhere from zero to 100% with equanimity. My point is not to argue definitively whether these particular outcomes are realistic or not, but to emphasize the importance of the underlying issues that they raise for model testing and policy analysis. In other words, it is possible that surprising and seemingly provocative results about the effects of a policy on the direction of change in key system indicators may be attributable to defects in the model that distort the equilibrium outcomes.

On the other side of the issue, a number of system dynamics models have exhibited more than one equilibrium or steady-state set of conditions under circumstances that seem potentially defensible and significant for policy. For example, an unpublished Ph.D. dissertation by William Shaffer* that was done at MIT developed a model of crime rate in the state of Massachusetts and its relationship to deterrence measures in the form of police effort and eventual length of prison sentence. Under normal ranges, the model exhibited stable and well-bounded behavior. However, the model was also capable of exhibiting a very different mode in which prison capacity was significantly overloaded, and thereby rising crime rates had the potential for triggering a strong positive feedback loop that yielded exponential growth in crime: more crime led to additional arrests and additional court sentences; but in order to accommodate new prisoners in jail, average length of prison stay for existing prisoners had to be reduced; thus, turnover rate of prisoners increased and resulting lower average prison sentence reduced the deterrent effect of the prison system on crime rate, leading to further escalation in crime. While the deterrence theory underlying this particular model may be questioned on a variety of empirical grounds, the model nonetheless has important

* Shaffer, William A. "Court Management and the Massachusetts Criminal Justice System," Ph.D. dissertation, Alfred P. Sloan School of Management, M.I.T., 1976

properties. Under the overload conditions on prison capacity described earlier, the model probably exhibits exponential growth in crime away from an unstable equilibrium point. Thus, the model may have both a "normal" stable equilibrium point, as well as an unstable equilibrium point. The possibility of growth in an undesirable socio-economic variable, such as crime rate, away from an unstable equilibrium point certainly has potential policy relevance.

In summary, the issue of uniqueness of equilibrium/steady state in a system dynamics model is an important one for the analyst to evaluate. A model that exhibits multiple equilibria may be the result of insufficient structuring of the social and economic pressures that produce an equilibrium in real life, or alternatively, may reflect on the real nature of the underlying system. The whole subject of multiple equilibria in a system dynamics model (and indeed in other types of models) has barely been touched in the literature and needs further study. For the present, I would argue that the appearance of model behavior characterized by multiple equilibria is an important departure from the vast majority of models with a determinate equilibrium point, and should lead the model builder to serious and skeptical evaluation of model plausibility.

2.10 Understanding Forces Producing Equilibrium Positions

A related issue to that discussed in the previous section involves the forces that produce the one or more equilibrium positions that a system dynamics model of an actual system may exhibit. In general terms, an equilibrium point can be neutral and pressure-free, or alternatively, it can be produced by offsetting pressures. To see the difference between these two categories, consider a firm that utilizes labor and capital as factors of production to produce an output stream. Moreover, suppose for simplicity of exposition, that annual wage costs and capital charge rate for capital equipment remain constant, so that the optimal intensity of labor and capital in the production process does not change over time. A "neutral" or pressure-free equilibrium would be one in which the production sector replaces workers who quit or retire, and invests sufficiently to offset depreciation of capital equipment, but does not confront pressures to add or subtract labor or capital equipment due to shortage or excess of output or high or low marginal productivity of one factor in relationship to the other. In other words, if the sector has just sufficient output capacity to meet demand and maintain appropriate levels of output inventory and order backlog, and if each production factor is in the appropriate mix, then neither output pressures nor relative productivity of factors of production would produce incentives to change labor or capital stock over time. If, starting from a neutral equilibrium point, the sector experienced 10% more demand for its end product, it would eventually come to add 10% more labor and 10% more capital, thereby yielding 10% more output; once this point were reached, output rate would again match demand, and both labor and capital would have risen by the same percentage, so relative productivities would remain unchanged.

On the other hand, suppose that bottlenecks in the supply system for capital equipment prevent expansion of capital goods. If demand for the output of a firm that required capital goods for production were to expand, that firm could augment production only through adding labor or through

increased utilization, such as longer work week. Suppose, then, that final output demand went up by 10% and that the higher demand could be met with fixed capital stock through a 15% increase in employment and utilization. The new equilibrium that was thereby reached would no longer be pressure-free. Instead, acquisition of capital goods would be encouraged by a long output delivery delay and by a high marginal productivity of capital equipment (meaning that it would be efficient to add capital in relation to the outstanding number of employees); but capital acquisition would be restrained by insufficient supply of capital goods. On the other side, acquisition of labor would be discouraged by low marginal productivity, but encouraged by above-normal delivery delays resulting from unavailability of capital plant to supplement production capacity. Thus, the new equilibrium for each factor of production would be characterized by a balance between pressures to expand and forces to contract (or limitations on expansion) of that factor. Such occurence of offsetting equal but opposite pressures is what I mean by a non-neutral equilibrium point.

Very frequently, evaluation of the forces producing a model-generated position, be it neutral or non-neutral, can provide insight into the adequacy of the model or into the forces determining equilibrium values in real life. A good example of such analysis of equilibrium position is contained in a Ph.D. dissertation done at MIT several years ago by Barry Richmond on forces underlying the long-term expansion of government employment and expenditures in the United States.* As one thread of

---

* Richmond, Barry M. "Government Growth in a Fixed Economy," unpublished Ph.D. dissertation, Alfred P. Sloan School of Management, M.I.T., 1979.

argument in developing a theory of government growth, Richmond considers the prevalent argument that government expansion is the result of successive crises or incidents that create a temporary need for additional government intervention, but with the long-term result that government size

only grows over time, failing to fall significantly following the periods of increased intervention. This hypothesis of government expansion is sometimes called a "racheting" theory, with the term "rachet" connoting a progressive, stair-step form of increase in activity over time with growth interrupted only by periodic intervals of level activity. Richmond eventually develops simulation results to suggest that the rachet hypothesis does not provide a tenable explanation for government growth. The essential counterargument is that once the pressures producing greater apparent need for government intervention wane following a crisis, unless basic social values of the society have changed, then government activity would decline back to its original relationship to private output, although possibly with a very long downward adjustment time.

In a similar vein, in ongoing work in the National Model Project on causes of inflation and public policies to control inflation, we have argued that various "cost-push" theories of inflation do not provide a plausible theory of ongoing inflation in the absence of concomitant increase in money supply. Without increase in money and liquidity, rising price and wage levels produced from cost-push strains would eventually depress liquidity sufficiently to yield counterpressures that exactly offset the upward thrust on prices and wages due to the cost-push force.*

In summary then, evaluation of the balance of forces producing a model-generated equilibrium can be a powerful tool for evaluating model

---

* See Mass, Nathaniel J. "Cost-Push Inflation and the Politics of Monetary Expansion," Large-Scale Systems, Vol. 1, No. 2, pp. 107-115, North-Holland Publishing Company, Amsterdam, The Netherlands, 1980. (D-3098-1). Also Mass, Nathaniel J., "Monetary Sources of Inflation," System Dynamics Group Working Paper D-3254, February 1981.

adequacy and yielding policy suggestions. If, for example, a model generates a non-neutral equilibrium, the analyst should carefully consider whether the balance of opposing forces observed in the model would be likely to occur in real life, or whether in fact, mechanisms have been omitted from the model that would help to restore a neutral equilibrium from the standpoint of all impinging forces. On the other hand, if opposing forces can reasonably sustain an equilibrium, then appreciation of the nature of the balancing process may yield insights into the controlling mechanisms that may either act in concert with, or in opposition to, policy initiatives tried within the system.

3. CONCLUSIONS

As argued in the introduction to this paper, the very structural richness of a system dynamics model yields a certain degree of a priori unpredictability of model output. Certainly, as a model builder or decision maker working with a model become more familiar with the inner workings of the system being represented, the incidence of unexpected behavior may diminish somewhat. But a variety of experience suggests that some degree of unpredictability always remains. From this point of view, the experienced system dynamics model builder may indeed be more capable than the novice in anticipating behavior of a complex feedback model. But to an even greater extent, he may become more effective and creative in utilizing surprise behavior as a tool for diagnosing difficulties in basic model concept and in developing policy recommendations. There is no extensive practical literature that advises the model builder on guidelines for improving models, and even more so, on guidelines for evolving policy insights from a model. Thus, for the foreseeable future, both of these skills are likely to revolve around a high degree of art coupled with experience and good judgment. In this vein, the basic thesis of this paper is that appearance of surprise model behavior provides one of the events that can precipitate fruitful improvement and application of system dynamics models.