[10]

European Journal of Operational Research 35 (1988) 301–320
North-Holland

Invited Review

System dynamics and microworlds for policymakers

John D.W. MORECROFT
Business Policy Group, London Business School, Sussex Place, Regent's Park, London NW1 4SA, United Kingdom

Abstract: In the past ten years, system dynamics has become more accessible to policymakers and to the academic community. The paper reviews four major developments in the subject that have brought about this change. There have been improvements in the symbols and software used to map and model system structure. New ideas have been adopted from behavioural decision theory which help to transfer policymakers' knowledge into computer models. There have been improvements in methods of simulation analysis that enable modelers and model users to gain better insight into dynamic behaviour. Greater emphasis has been placed on small transparent models, on games and on dialogue between 'mental models' and computer simulations. Together these developments allow modelers to create computer-based learning environments (or microworlds) for policymakers to 'play-with' their knowledge of business and social systems and to debate policy and strategy change. The paper concludes with some thoughts on future research.

Keywords: Strategic modeling, system dynamics, mapping and simulation, gaming, behavioural decision making, group strategy support

Introduction [1]

In the past ten years there have been several important developments in system dynamics which

[1] The paper explores recent developments in system dynamics for policymaking. I have tried to provide a broad survey of these developments. However, one's knowledge of even specialised topics is always incomplete, and is conditioned by his/her institutional and geographical setting. I have spent more than 10 years in the United States at MIT's Sloan School of Management and returned to the UK in mid 1986 to join the London Business School. My knowledge of system dynamics for policy-making is therefore heavily influenced by what I have seen at close hand in North America and most particularly at MIT. I am aware of several developments in system dynamics for policymaking in the UK and some in the rest of Europe, but I apologise for omissions that European readers may detect. They are not intentional.
Important developments in mathematical methods of model analysis and estimation are excluded from the paper by design, since they are contributions to basic technical methodology and not to the practice of policymaking. The conceptual developments in system dynamics inspired by work on self-organising systems (from the so-called 'Brussels school') are also excluded because the theoretical ties between the Brussels school and traditional system dynamics have only recently been examined closely. As a result, the contribution of self-organising system concepts to the practice of traditional policy modeling, though potentially significant, is not yet clear.

The author is grateful to Tony Cornford, Kees van der Heijden, Donald Kalff, Linda Morecroft and Jane Varley for comments and editorial suggestions received on an earlier draft of the paper. The author is also grateful to Arie de Geus for his support and for his commitment to the concept of microworlds for policymakers.

Received September 1987

make the subject more accessible to policymakers, more communicable to the academic community and more challenging for research. [2] There have been improvements in the symbols and software used to map and model system structure. New ideas have been adopted from behavioural decision theory which help to capture policymakers' knowledge in computer models. There have been improvements in methods of simulation analysis that enable modelers and model users to gain better insight into dynamic behaviour. Greater emphasis has been placed on small, transparent models and dialogue between 'mental models' and computer simulation models.

As a result of these developments, system dynamics can now be used, with a management team, to structure informed debate about strategic change, in a process where models and computer simulations are an integral part of management dialogue. The paper explores each of the major developments in more depth in order to show readers the range of ideas and concepts that system dynamics now encompasses. The paper concludes with some thoughts on future research to improve model supported policy dialogue and the mapping of policymakers' knowledge.

System dynamics—A microworld for policy debate

What is a 'microworld' [3] for policy debate? Figure 1 shows the elements in the microworld

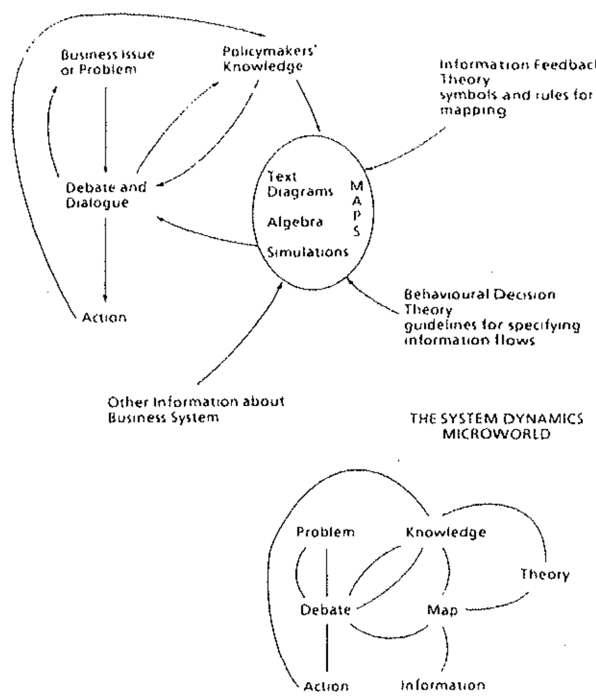

Figure 1. The microworld for policy debate provided by system dynamics

provided by system dynamics. At the top left is a problem or issue facing policymakers which initiates debate and dialogue. The debate leads to clarification of the problem or issue and eventually to recommendations for action. The microworld contains all the factors that impinge on the debate. A most important factor is the policymakers' own knowledge (or mental model) of the business or social system they manage. This knowledge provides the raw material for debate and discussion. In conventional policymaking (by argument) it is the interplay between the knowledge base and the debate that produces a consensus for action.

[2] In addition there is now an international System Dynamics Society. The Society runs an annual international conference and publishes a bi-annual journal, the *System Dynamics Review*, containing articles on the theory and applications of system dynamics, research problems, reports of meetings and book reviews. The Society also distributes a newsletter listing the universities and colleges around the world that offer courses in system dynamics. Readers should contact the System Dynamics Society, Principal Office MIT E40-294, Cambridge, MA 02139, USA.

[3] The term 'microworld' comes from Seymour Papert in a fascinating book called *Mindstorms* [45] Papert is a mathematician and computer scientist at MIT who has devoted his energy to exploring how computers can help people to learn. A fundamental premise of his work is that people learn effectively when they have transitional objects to 'play with' in order to develop their understanding of a particular subject or issue. The writer has pen, paper and word processor with which to hone his skill of composition. The very young child has building blocks to learn about sizes, sorting and simple construction. The combination of transitional objects, learner and learning process is what Papert calls a microworld or 'incubator for knowledge'. But what transitional objects can one provide for learning about 'intangible' topics like motion, geometry, mathematics and (for our purpose) policymaking? Papert suggests the computer and simulation: "The computer is the Proteus of machines. Its essence is its universality, its power to simulate. Because it can take on a thousand forms and can serve a thousand functions, it can appeal to a thousand tastes."
The combination of computer, simulation language, learner and learning process is a computer-based microworld.

When modeling and simulation enter the debate (or for that matter any other framework for policymaking) the picture becomes more complex as the interplay of knowledge, information and debate becomes richer. Policymakers' knowledge, and other information about the business or social system, (staff reports, financial documents etc.) are converted into text, diagrams, algebra and simulations. The process of mapping knowledge and information is guided by theory and concepts of system dynamics. The figure shows two main inputs from theory. The first input, from information feedback theory, provides symbols for diagramming a business or social system and rules for mapping. As explained in more detail later, these symbols include levels, action flows, flow regulators, converters and information flows to represent physical, financial and decisionmaking processes. The rules for mapping include rules for connecting the symbols, guidelines for equation formulation and guidelines for simulation testing and behaviour analysis. [4] The second input, from behavioural decision theory, improves the 'fidelity' of models. Behavioural decision theory provides the modeler with guidelines for specifying a model's information flows. It helps modelers to ask the 'right' questions of policymakers and it helps modelers to represent an organisation's decision processes accurately.

The microworld includes knowledge (K), information (I), theory (T), maps (M), debate (D) and the interplay of these factors as summarised in the inset of Figure 1. The scope of policy debate is larger (in principle) than can be achieved with conventional argument. The maps (text, diagrams, algebra, and simulations) provide policymakers with a variety of perspectives on their pooled knowledge. The maps also draw information from reports and staff specialists. So the interplay of debate and knowledge is enhanced through increased variety of representation, more information, and additional paths of interaction. Moreover, the maps themselves are created with the specialist knowledge supplied by theory.

Now let us turn to the developments in system dynamics which have made possible this microworld for policy debate.

Improvements in mapping methods and software

The origin of system dynamics can be traced to engineering control systems and the theory of information feedback systems. Indeed courses with the title system dynamics are offered in many engineering schools. But, in the setting of engineering, the subject has a rather mathematical flavour. One of the major contributions to modeling made by Forrester was to reshape sophisticated modeling and analysis methods from control engineering into a flexible form suited to modeling and debate in the business/social arena. One can think of this reshaping as a change in the 'technology' for mapping knowledge of systems into algebra and differential equations. In *Industrial Dynamics* Forrester [11] offered symbols for diagramming feedback systems and rules for connecting the symbols. Together these symbols and rules produce a diagram that shows the interconnectedness of a business or social system, that highlights feedback paths and that guides equation formulation. In addition Forrester devised a special purpose simulation language for coding symbols from diagrams into algebraic equations. Forrester's reshaping of methods from control engineering led to a visual representation of feedback systems, and through simulations, to a visual representation of feedback dynamics. These graphics provide a conduit for policymaker's knowledge and a basis for policy debate.

In the past ten years there has been renewed interest in the symbols used in system dynamics (Morecroft [38], Richardson [51], Richmond [52]). They have been simplified visually and can now be manipulated on graphical computers.

Let us review the symbols and rules of connection, and then talk about software. Figure 2 shows the main symbols. At the top of the figure, shown as a circle, is a decision function or converter. The function receives information, shown as dotted lines, processes it and generates an output in the form of action or more information (such as a command transmitted elsewhere in the system). For example, one might think of a hiring function in a sales organisation, receiving information on both the current size of the salesforce and the

[4] I am using the term 'mapping' here to mean all stages in the conversion of mental models into computer simulation models, including diagrams, text, algebra, programmed algebra, and even simulations.

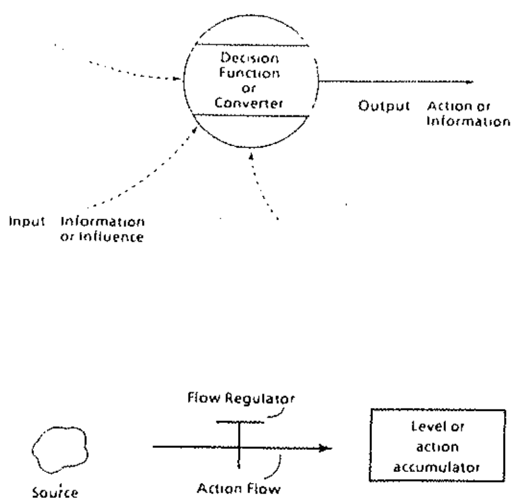

Figure 2. Symbols for mapping

force size authorised by the budget. Force planners (or a hiring committee/manager) process this information periodically in order to adjust the size of the salesforce (an action). The circular function symbol is usually identified with a specific decision making process or centre of responsibility in an organisation. [5] Within each symbol there can be quite complex information processing and therefore quite complex algebra.

The decision function or converter is a *general* function of its inputs. The detail of the function is not specified by the symbol, and this is deliberate, because a diagram made of such symbols retains a lot of flexibility in discussion—it is not overly defined and it is not too 'close' to the final algebra. The decision function or converter may contain within it 'subsymbols' such as constants, auxiliaries and information levels. The function symbol is therefore a quite-compact visual representation of the decisionmaking process under study.

A business or social system is viewed as a set of decision making 'players' whose decisions and actions are coupled (Forrester [11, Chapter 10]). Each 'player', or centre of responsibility, is represented by a decision function with information inputs and an output which is either information or action. For example, players in a sales organisation might be business planners (setting sales objectives), compensation planners (adjusting the commission paid on different product lines), salespeople (deciding how to allocate their selling time to product lines), force planners (adjusting the size of the salesforce) and customers (being influenced by salespeople to buy). But how are the decision functions and actions coupled? Here is where the remaining symbols and rules of connection are useful.

In the lower half of Figure 2 there are three symbols. The box on the right is a level that accumulates action. In the centre is a composite symbol that represents an action flow (shown as an arrow) and a flow regulator (shown as a 'T') which controls the size or volume of the flow. Finally, on the left is a source (shown as an irregular 'blob') which supplies the action flow. (If the action flow is reversed, then the source becomes a sink and the flow drains the level). In order to understand the symbols, again imagine the process of hiring in a sales organisation. The level is the number of salespeople currently employed, which accumulates the flow of people being hired. The hiring flow is controlled by the flow regulator. (which as we shall see later, is connected to the hiring function). Salespeople are hired from the labour market which is therefore the source of the flow.

System dynamics provides rules for connecting the symbols so that one can construct a network of decision functions of arbitrary complexity and thereby map and mimic an organisation. Figure 3 shows many symbols connected. The sequence of connection is: level–information (or influence)–decision function (or converter)–action–level. This is a feedback representation. A decision function (say hiring) leads to action (the arrival of salespeople). New recruits accumulate in the level of salespeople. Information from the level (the size of the salesforce) is an input to the hiring function.

The arrival of graphic computers has now made it possible to map symbols directly onto a computer screen and to benefit from interactive mapping and high quality graphics. The new modelling

[5] Sometimes a function may represent a physical process or behavioural process such as motivation. In the case of production, the function combines inputs of material, labour and capital and converts them into (say) finished goods. There is no conscious decisionmaking or information processing involved. Similarly, in the case of motivation, the function combines inputs of say stress, workload, and goal achievement and converts them into an index of individual or group motivation. In such physical or behavioural processes, the inputs to the function are best thought of as influences rather than information flows.

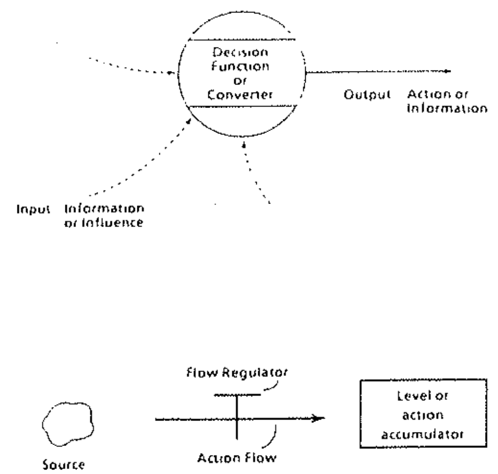

Figure 2. Symbols for mapping

force size authorised by the budget. Force planners (or a hiring committee/manager) process this information periodically in order to adjust the size of the salesforce (an action). The circular function symbol is usually identified with a specific decision making process or centre of responsibility in an organisation. [5] Within each symbol there can be quite complex information processing and therefore quite complex algebra.

The decision function or converter is a *general* function of its inputs. The detail of the function is not specified by the symbol, and this is deliberate, because a diagram made of such symbols retains a lot of flexibility in discussion—it is not overly defined and it is not too 'close' to the final algebra. The decision function or converter may contain within it 'subsymbols' such as constants, auxiliaries and information levels. The function symbol is therefore a quite-compact visual representation of the decisionmaking process under study.

A business or social system is viewed as a set of decision making 'players' whose decisions and actions are coupled (Forrester [11, Chapter 10]). Each 'player', or centre of responsibility, is represented by a decision function with information inputs and an output which is either information or action. For example, players in a sales organisation might be business planners (setting sales objectives), compensation planners (adjusting the commission paid on different product lines), salespeople (deciding how to allocate their selling time to product lines), force planners (adjusting the size of the salesforce) and customers (being influenced by salespeople to buy). But how are the decision functions and actions coupled? Here is where the remaining symbols and rules of connection are useful.

In the lower half of Figure 2 there are three symbols. The box on the right is a level that accumulates action. In the centre is a composite symbol that represents an action flow (shown as an arrow) and a flow regulator (shown as a 'T') which controls the size or volume of the flow. Finally, on the left is a source (shown as an irregular 'blob') which supplies the action flow. (If the action flow is reversed, then the source becomes a sink and the flow drains the level). In order to understand the symbols, again imagine the process of hiring in a sales organisation. The level is the number of salespeople currently employed, which accumulates the flow of people being hired. The hiring flow is controlled by the flow regulator, (which as we shall see later, is connected to the hiring function). Salespeople are hired from the labour market which is therefore the source of the flow.

System dynamics provides rules for connecting the symbols so that one can construct a network of decision functions of arbitrary complexity and thereby map and mimic an organisation. Figure 3 shows many symbols connected. The sequence of connection is: level–information (or influence)–decision function (or converter)–action–level. This is a feedback representation. A decision function (say hiring) leads to action (the arrival of salespeople). New recruits accumulate in the level of salespeople. Information from the level (the size of the salesforce) is an input to the hiring function.

The arrival of graphic computers has now made it possible to map symbols directly onto a computer screen and to benefit from interactive mapping and high quality graphics. The new modelling

[5] Sometimes a function may represent a physical process or behavioural process such as motivation. In the case of production, the function combines inputs of material, labour and capital and converts them into (say) finished goods. There is no conscious decisionmaking or information processing involved. Similarly, in the case of motivation, the function combines inputs of say stress, workload, and goal achievement and converts them into an index of individual or group motivation. In such physical or behavioural processes, the inputs to the function are best thought of as influences rather than information flows.

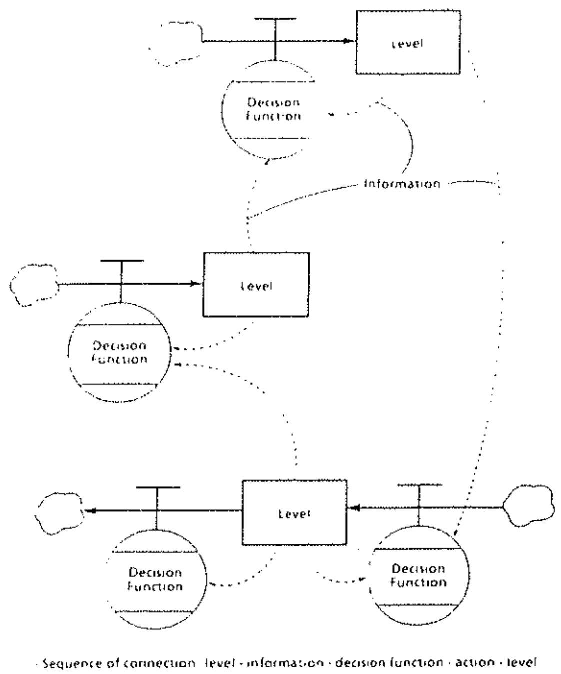

Figure 3. Rules for symbol connection

and simulation package STELLA (Richmond [53] and Richmond et al. [54]) includes very effective mapping software. STELLA provides the modeler with a menu of symbols for creating a diagram on an 'electronic' worksheet. The symbols include those shown in Figure 2 and several others (hand, ghost and dynamite) that help in editing and organising the diagram. One can select symbols from the menu, move them onto the computer screen (a small part of the available electronic worksheet) and connect them. The software 'knows' the connection rules, so modelers are constrained to produce diagrams which connect symbols in the sequence: level–information flow (or influence)–converter–action flow–level. The software provides a very effective (and entertaining) medium for mapping policymakers' knowledge of a business or social system. [6]

In addition to symbols, system dynamics provides guidelines for equation formulation. These guidelines can be thought of as rules for converting symbols, text and words into algebra. There is not the space here to examine the guidelines in depth. [7] However, let us take just a couple of examples to illustrate how methods of developing algebra are improving with the help of software.

Consider first the simple but crucial issue of naming symbols and variables. Names are chosen to fit closely with the terminology and concepts used by policymakers. Models often go through major revisions of terminology as concepts become clearer and more precisely defined. The objective is to make diagrams and algebra readable and easy to relate to people's mental models.

Modeling software gives good support to users who wish to create readable models. For example, the modeling and simulation package Professional DYNAMO [49] has a documentation module that provides clear visual display of algebra, automatic translation of algebraic names into plain language (using a file of definitions supplied by the modeler) and comprehensive cross indexing of variables, equations and manually-prepared diagrams. STELLA allows the modeler to specify long (17 character) labels for symbols in diagrams. Once the labels are entered into a diagram they are available for use in equation writing without being retyped. Moreover, modelers must write equations that combine the inputs to symbols as specified on the diagram, otherwise the software will reject the equations. The result is a diagram with understandable labels that always match exactly with the variables used in equations.

Consider next the issue of dimensional balance in equations. The units of measure on the left-hand side of an equation should match those on the right, and any conversion coefficients (such as productivity that relates a workforce to a produc-

[6] There is one symbol provided by the STELLA software, the converter, that differs from the function symbol in Figure 2. The function symbol is a 'high level' symbol that can contain many subsymbols. In STELLA, the converter is just like a traditional 'auxiliary' symbol in system dynamics. In practice this difference is important. It causes STELLA diagrams to be more visually complex than policy structure diagrams (or policy maps). The increase in visual complexity is more noticeable the larger the model. Moreover it is difficult to identify in STELLA diagrams the major 'players' and centres of decisionmaking responsibility. However, there is no technical reason why a graphic modeling package could not employ a high-level general function symbol.

[7] Readers who would like to know more about guidelines for equation formulation are referred to Forrester [12, Chapters 3–9], Richardson and Pugh [50, Chapter 4], Coyle [6, Chapter 5] and Richmond et al. [54, pp 2.72–2.74].

tion rate) should have real meaning in the system. Dimensional analysis, if thorough, is a powerful method of rooting-out errors in formulation and for pinpointing confusion in the conversion of diagrams and verbal descriptions into algebra.

Software packages help modelers with dimensional analysis. DYSMAP (Cavana and Coyle [3], Vapenikova and Dangerfield [72]) includes a full dimensional analysis module. Professional DYNAMO and STELLA encourage dimensional analysis in that they allow one to supply dimension labels that appear in equation listings.

New concepts from behavioural decision theory

With the symbols and mapping rules of system dynamics it is possible to create quite complex networks of decision functions, actions and levels. But there are innumerable ways to link the symbols which all obey the connection rules of feedback systems. However, *only some symbol configurations* correspond to realistic organisational systems. There is a need for modelers to be quite discriminating in their choice of information links and influences if they are to produce plausible and insightful policy models.

Recently, system dynamics has adopted concepts from behavioural decision theory that are useful for specifying information links among decision functions. (Hall [25,26,27], Morecroft [39,41], Sterman [66,68,69]). Behavioural decision theory focuses on the information and heuristics in real-life decisionmaking. What information receives attention in organisational decisions? What information is ignored, and why? What factors condition the quantity and quality of this information? Behavioural decision theory concludes (with plenty of empirical evidence) that people make choices using only a few sources of information processed with simple rules of thumb. So the network of information flows in a realistic organisation is quite sparse relative to the network that would exist if each decisionmaker used information from every source in the system.

Figure 4 shows how behavioural decision theory guides the mapping of decision functions and therefore complements the rules from information feedback theory. One can see in the figure the standard feedback representation: decision function–action flow–level–information–decision function. In addition there are many other information flows and influences (originating from other levels in the system) which are shown on the outer boundaries of the decision function. Only a few of the information flows actually penetrate to the heart of the decision function where they influence the choices and actions of the 'players' (individuals, groups, subunits) represented by the function. The concentric circles surrounding the decision function represent organisational and cognitive filters which select or limit the information made available to different decisionmakers.

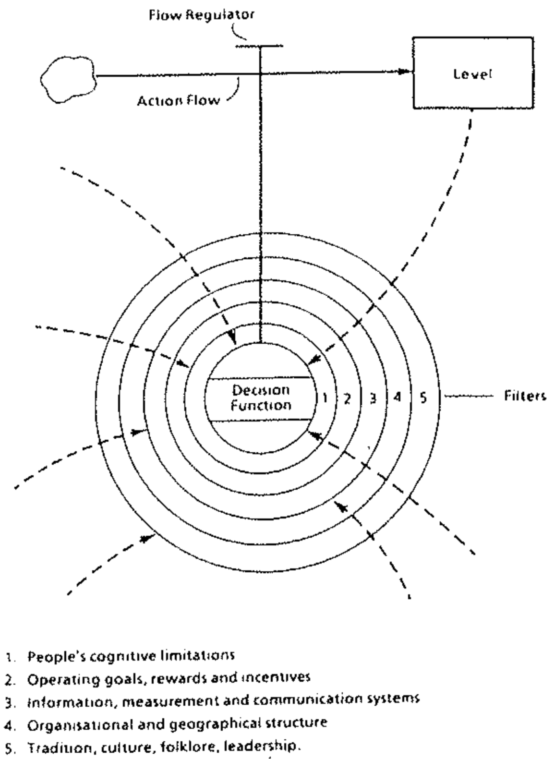

Figure 4. The behavioural decision function—Decision making and information filters

The composite symbol comprising the decision function and concentric circles (filters) is a behavioural decision function. It is a visual representation of ideas first developed by Simon [61], and the Carnegie School (summarised in Allison [1, Chapter 3] and captures aspects of modern behavioural decision theory by Hogarth [28], Tversky and Kahnemann [71].

There are five information filters surrounding a decision function. The first filter represents people's cognitive limits and Simon's notion of bounded rationality (Simon [60,62,64]). People are unable to process all the information that a business or social system may present them. They make their judgements on the basis of a few sources of information processed according to quite simple rules of thumb.

The outer filters (2, 3, 4 and 5) in Figure 4 represent the ways in which an organisation conditions the information made available to decisionmakers. This part of the figure draws particularly on Simon's *Administrative Behavior* which explains how organisations may display effective decisionmaking despite people's cognitive limits and an over-abundance of information. [8] Executives and managers (and in fact all employees) make their judgements and decisions in a 'psychological environment' provided by the organisation. The psychological environment limits the range of factors they consider and, in principle, supplies only the relevant information (a tiny subset of the total information available in the organisation) for making the correct choices at a given centre of responsibility. The filters show the components of the 'psychological environment' and they also provide a convenient basis for questioning policymakers.

Filter number 2 represents the influence of operating goals, rewards and incentives on information flow. Decisions and actions in business and social systems depend on the operating goals and rewards faced by the key players in the system. One can only understand organisational choice and action relative to these goals and rewards. So, for example, it is well-known that factory managers who are held accountable for a specific end-of-year inventory target will drastically curtail or boost production to meet the target, in defiance of 'rational' cost-minimizing scheduling criteria. For these factory managers, information about the status of inventory easily penetrates filter number 2. The filter excludes other information on future expected demand, cost structure and capacity constraints, which together with information on inventory would be required to set an 'objectively rational' production schedule. The modeler must ask questions which elicit policymakers' knowledge of goals and rewards.

Filter number 3 represents the influence of information, measurement and communication systems on information flow. To take another production example, a 'good' production schedule for a microcomputer manufacturer might require information on the status of inventory in all retail outlets. If there is no information system capable of monitoring and reporting retail inventory, then the production schedule must make do with factory information on the size of the order backlog, the amount of finished inventory and the recent shipping rate.

Filter 4 represents the influence of organisational and geographical structure on information flow. As a decisionmaker, one's position in an organisation (both geographical location and position on the organisational chart) have a profound influence on the information sources one is exposed to.

Filter 5 represents the influence of tradition, culture, folklore and leadership on information flow. Filter 5 is quite intangible yet very important in determining the factors that get the attention of decisionmakers in business and social systems. For example, suppose one is modeling the service division of a computer company and wants to understand the quality of service provided to customers. Quality of service depends on the speed with which servicemen fix customer problems. The division can respond quickly if its servicemen receive information promptly from customers. But the company also needs a 'service culture'. A customer problem which is known to a serviceman will get attention (i.e. bring about some action) if the company's 'culture' encourages good service. A culture for good service may derive from quite intangible factors such as stories and folklore which circulate the company. The stories and folklore underpin the attitudes of individuals in the service division, and condition the attention they pay to customer problems (in other words, the weight they give to information from customers requiring service).

What guidance do the filters provide the modeler? Principally they help modelers to map organisational systems accurately by forcing them to pay close attention to the information sources that are *actually used* by decisionmakers (as opposed to the information sources that are available

[8] Therefore policy modelers and policymakers would benefit from reading books like *Administrative Behaviour* (Simon [58]). *The Functions of the Executive* (Barnard [2]) *Essence of Decision* (Allison [1] and *Judgement and Choice* (Hogarth [28]), in order to cull some basic principles of human and organisational decisionmaking.

or that seem, at a distance, to be the most 'sensible') and to be aware that information deficiencies, bias and error are commonplace. [9,10] Also, the filters focus attention on the modeling of decision processes, not just causal links or influences. [11]

By being aware of the filters, modelers can ask more precise questions to draw out policymakers' knowledge, and to better specify decision functions. The result is realistic feedback structure that comes from linking well-specified decision functions.

Influence of behavioural decision theory on system dynamics

Besides helping modelers to specify decision functions and map feedback structure, the ideas adopted from behavioural decision theory have improved communication with academics, added some momentum to the 'soft-modeling' movement, and stimulated new research in gaming and the experimental study of organisational decisionmaking. These influences are examined in more detail below, particularly as they relate to policy modeling and policymaking.

Communication with the academic community

Behavioural decision theory clarifies and amplifies the conceptual content of system dynamics and provides new vocabulary for communicating with academics. System dynamics models can now be described as 'behavioural simulation models' that 'portray bounded rationality in organisations'. The models represent organisations as decisionmaking/ information processing systems involving many players, with multiple (often conflicting) goals and limited information processing capability. The feedback structure of the models emerges from the assumptions one makes about decisionmakers' access to information, the weight that decisionmakers place on different information sources and the rules of thumb they use to make judgements. Dynamic behaviour (which is often economically 'inefficient' from a system-wide perspective) is a consequence of the feedback structure and can therefore be traced to assumptions about behavioural decisionmaking. With these labels to describe models it is possible to write articles of direct interest to (and to converse with) academic economists (Sterman [66],) organisational and policy theorists (Hall [24,25,27], Morecroft [41,42], and behavioural scientists (Sterman [68,69]). Moreover, policymakers also relate to models which capture explicitly the multiple operating goals, administrative procedures and information deficiencies of real organisations.

Momentum to the soft-modeling movement

Recently interest has grown in system dynamics as a soft-modeling methodology. Work in this area has been spearheaded by Wolstenholme [73, 74 (with Coyle), 75] and is related to work in the general systems area by Checkland [4] and in cognitive mapping by Eden [7,8,9]. The thrust of the soft modeling movement in system dynamics can be seen in a portion of Figure 1 which is reproduced in Figure 5 below for convenience. The idea is to emphasize the 'front-end' of system dynamics modeling before one invokes the use of algebra or simulation. Wolstenholme and Coyle [74] distinguish 'qualitative' and 'quantitative' models and argue that 'qualitative' modeling is quite useful to policymakers in its own right, if the modeler uses the symbols and structuring rules of

[9] Forrester [11] stresses the need to model real-life decisionmaking processes and to pay attention to the information sources that are influential at different points in an organisation (see particularly Chapter 10, 'Policies and Decisions'). But *Industrial Dynamics* gives little explicit guidance to the modeler on how to select the information flows that should enter decision functions and later books by other authors give no guidance at all (because the later books show how to model causal loops, not decisionmaking processes). Behavioural decision theory (particularly from Carnegie literature) provides some conceptual apparatus for discriminating influential information sources.

[10] During the 1970's, many system dynamics models were mapped in terms of causal loop diagrams or influence diagrams without explicit representation of decisionmaking. In the case of policy models, it is much better to map decision processes, actions and levels (rather than simple influences) in order to build credible feedback structure (see Forrester [11], Chapters 6–10, and Morecroft [38]).

[11] Interestingly, the outer filters owe very little to modern behavioural decision theory, because modern theory deals principally with individual or group decisionmaking in a *given* organisational setting. The outer filters represent the conditioning of information by the organisation or, in other words, the organisational setting itself.

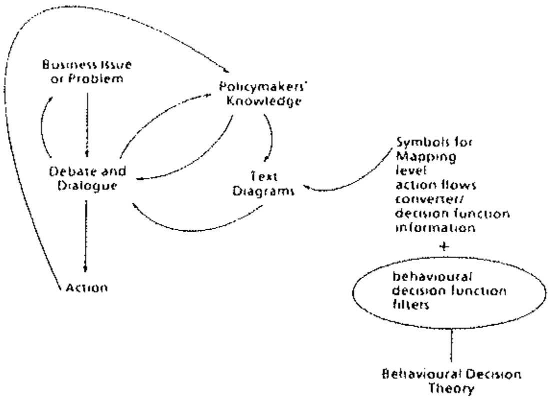

Figure 5. The soft modeling movement and the influence of behavioural decision theory

system dynamics to full advantage.

In Figure 5, a 'qualitative' model is the text and diagrams that result from mapping policymakers' knowledge into the symbols provided by system dynamics. Wolstenholme suggests a number of guidelines to enhance the mapping. Behavioural decision theory adds more symbols (the behavioural decision function and associated filters) to the qualitative modeler's graphic menu. These extra symbols enrich the diagrams, draw in more of the policymakers' knowledge, and broaden the scope of policy debate and dialogue.

Research in gaming and experimental study of organisational decisionmaking

Gaming has long been a branch of system dynamics. Recently, Meadows [35] has revitalised the topic, principally as a means of improving communication between models and policymakers. His work is reviewed later. However, building on Meadow's work, and injecting behavioural decision theory, Sterman has opened gaming as a fascinating new branch of research that promises to forge closer ties to modern behavioural decision theory and to yield new experimental methods for the study of organisational decisionmaking.

An example of Sterman's work is a behavioural simulation model of capital investment in an economy [12] (Sterman [66]). Included in the model are aggregate decision functions for production planning, inventory control, backlog control, capital ordering and capital supply, linked together by information flows, action flows and levels representing the aggregate stock of capital equipment on-hand and on-order in the economy. Sterman and Meadows [67] have converted the model into a role-playing simulation game by replacing the decision function for capital ordering with a 'decision shell' which provides to real human subjects the same information used in the capital ordering function of the full model. The design of the decision shell is guided by considering the behavioural and organisational filters in the original model's capital ordering function. Gameplayers manage a simulated economy in which they must order sufficient capital plant and equipment to satisfy aggregate demand. They have complete and perfect information on demand, backlog, delivery time and capital stock. They can view the structure of the simulated economy on a video display generated by the microcomputer-based game, and they can call up screens of information which show the past history of the system. When the game is played, the overwhelming majority of subjects generate significant and costly oscillations in capital stock and capital orders quite similar to those generated by the full model. The behaviour of subjects is far from optimal and Sterman [68] suggests that the decision-making heuristic they are following is captured by a simple decision rule which is consistent with the notion of bounded rationality.

In business policy modeling, similar gaming work has been carried out by Flint [10] who constructed a role-playing simulation game of sales planning and control for a multiproduct salesforce. Flint's game links decision functions for salesforce time allocation (salespeople's choice of the proportion of selling time they allocate to different product lines), customer ordering, sales forecasting and objective setting, to a decision shell for compensation planning. Subjects in the game play the role of compensation planners, adjusting the commission on product lines according to sales performance. Flint describes the design

[12] Sterman's model is a much simplified version of a system dynamics model of aggregate production activity in the U.S. economy, developed over the past ten years by MIT's System Dynamics Group. For more information on the model see Forrester [15,16] and Forrester et al. [18].

of the decision shell and the results of playing the game with managers and compensation planners in companies with a multi-product salesforce.

Emphasis on learning and reasoning—the microworld concept

People frequently ask what new insight can be gained into business and social problems from modeling and simulation that cannot be obtained from conventional written and verbal argument? Forrester and Simon have both made interesting statements on this question which serve to underscore the new emphasis in system dynamics on learning and reasoning. Forrester organises his paper 'Counterintuitive Behavior of Social Systems' [14] on "the basic theme that the human mind is not adapted to interpreting how social systems behave". He suggests that people misjudge the dynamic behaviour of social systems because there are:

> "...orderly processes at work in the creation of human judgement and intuition, which frequently lead people to wrong decisions when faced with complex and highly interacting systems."

He then goes on to contrast a computer simulation model with a mental model:

> "...the most important difference between the properly conceived computer model and the mental model is in the ability to determine the dynamic consequences when the assumptions within the model interact with one another. The human mind is not adapted to sensing correctly the consequences of a mental model. The mental model may be correct in structure and assumptions but, even so, the human mind—either individually or as a group consensus—is most apt to draw the wrong conclusions. There is no doubt about the digital computer routinely and accurately tracing through the sequences of actions that result from following the statements of behavior for individual points in the model system."

Simon [63] devotes pages 17 to 22 of his book *The Sciences of the Artificial* to the topic 'Simulation as a Source of New Knowledge':

> "...This brings me to the crucial question about simulation. How can a simulation ever tell us anything that we do not already know? The usual implication of the question is that it can't...[However]...There are two related ways in which simulation can provide new knowledge—one of them obvious, the other perhaps a bit subtle. The obvious point is that even when we have the correct premises, it may be very difficult to discover what they imply. All correct reasoning is a grand system of tautologies, but only God can make direct use of that fact. The rest of us must painstakingly and fallibly tease out the consequences of our assumptions... The more interesting and subtle question is whether simulation can be of any help to us when we do not know very much initially about the laws that govern the behavior of the inner system?... [The question can also be answered in the affirmative]... if the aspects in which we are interested arise out of the *organization* of the parts, independent of all but a few properties of the individual components."

Increasingly system dynamicists view their models as 'sources of new knowledge' or as tools for *learning* about business and social systems. (See de Geus [21], for the views of a senior planning executive on corporate learning and the role of models). This emphasis on learning is reflected in several trends and topics: generating dialogue between mental models and simulations, using workshops and role-playing simulation games, and finally generic models.

Generating dialogue between mental models and simulations

How can one generate a dialogue between a mental model and a simulation model? In system dynamics the process for doing so was first spelled out by Mass [34] in a paper with the title 'Diagnosing Surprise Model Behavior: A Tool for Evolving Behavioral and Policy Insights'. Mass begins with the observation (like Simon) that simulations can provide new knowledge:

> "...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."

He then goes on to propose ten guidelines and tests for resolving surprise behaviour. I will not list all ten here but instead pick three which are especially relevant to starting a dialogue. First, and most important, policymakers and modelers should establish an a priori expectation of system behaviour:

> "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 recognised through the appearance of a discrepancy."

Once an expectation of model behavior is established then the second guideline is to follow up all unanticipated behaviour to an appropriate resolution:

> "The model builder must adopt a perspective that views the encountering of surprise model behaviour as a significant opportunity to be capitalized upon ... [rather than] ... to pursue parameter combinations that make the anomalous behavior less evident or to dismiss the behavior as being outside the intended use of the model."

The third guideline is to confirm all hypotheses about surprise dynamic behaviour with 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 dynamic 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."

Mass's guidelines provide a protocol for generating dialogue that many system dynamicists have since adopted. The protocol has evolved to suit the needs of policymakers (in addition to modelers) and to suit new software. Morecroft [40] describes the process of generating policy dialogue and debate by contrasting executive opinion (about the outcome of a proposed strategic move) with model-generated 'opinion'. Models which are built to facilitate this 'dialectic' are described as 'strategy support models'. Similarly, Richmond's [53] model-testing guidelines require modelers to state a hypothesis (about dynamic behaviour) before simulating, in order to 'squeeze the maximum learning' out of simulations.

For people to engage in dialogue with a model, simulation runs must readily relate to their intuition. In other words, simulation runs should be designed to correspond to scenarios whose outcome policymakers and modelers can readily imagine (express an opinion on) using their knowledge of the system.

Morecroft [41] and Sterman [66] suggest the use of partial model tests that expose the 'intended rationality' of decisionmaking in complex systems. This testing strategy is particularly effective for communicating models that exhibit counterintuitive dynamic behaviour (e.g. long term cycles in capital equipment in Sterman's model of capital investment and a productivity trap in Morecroft's model of a sales organisation). Partial model tests show that decisions and actions of players in a system are 'sensible' (intendedly rational) when the feedback setting for the players' decisions is simple. Dynamic behaviour which arises from 'sensible' decisions and actions is usually intuitively clear, and therefore conducive to dialogue.

Partial model tests are designed by cutting feedback loops in the full model (or by building a deliberately simple, incomplete model) in order to isolate a subset of the system's interacting behavioural decision functions. The simplification is carried out in such a way that one can construct plausible scenarios (scenarios that policymakers can identify with) from simulations of the partial structure. For example, Sterman conducts a test of capital ordering in an aggregate economy in which production of capital equipment is unconstrained by the economy's stock (or level) of capital equipment. Of course, in the real economy, and in the full model, the stock of capital equipment *does* constrain production. Nevertheless, removing the constraint does not alter the logic or inputs to the model's decision function for capital ordering. But it does simplify the feedback network in which the function is embedded. Two feedback loops are cut. One cut eliminates capital self-ordering, the boostrapping process in capital investment whereby, in order to boost production of capital equipment, the sector must first increase its stock of capital equipment. The other cut eliminates hoarding, the tendency for producers to order more equipment when equipment delivery time is rising. [13]

Simulations of the partial model show that capital orders and capital stock adjust in a straightforward (easy-to-understand) way to a step increase in demand for capital goods. By contrast, simulations of the full model show that capital orders and capital stock fluctuate with a long 50 year period in response to the same step increase in demand. Sterman uses several partial model tests to build a clear explanation for the complex (non-intuitive) behaviour of the full model.

Richmond et al. [54] also talk about partial model testing under the heading of 'conducting a sanity check'. Step one of the process is to construct a simple 'open-loop' model (the equivalent of a partial model). As Richmond explains, the open-loop model stimulates thinking:

> "Open loop, in this context, really means 'free lunch' i.e. you can get what you want simply by asking for it. Clearly, such an open loop model is not intended to be realistic. Rather, it is designed to indicate what's possible (under the most ideal circumstances), and also to acquaint the group of managers involved in the sanity check exercise with the basic framework and technology that will be used in the analysis. The idea is to begin simple... Then, as the exercise proceeds, the 'freebies' will be systematically removed. The operative question throughout is always the same: can we still get there from here?"

Using workshops and role-playing simulation games

It used to be common in policy modeling to develop models containing several hundred or even several thousand equations. These large models (which are still built today and have an important role to play in policymaking) are large because their creators want them to accurately replicate historical time series and to provide good short-term predictions, in addition to generating dialogue. It usually takes lots of equations to write a model that will accurately simulate history.

Now, much smaller models of thirty, forty or fifty equations are commonly presented to policymakers. The purpose of these models is to 'prime' policymakers' for debate. Much less emphasis is given to replicating time series.

In order to stimulate debate, a model should be transparent so that policymakers can see their knowledge reflected in the model's assumptions. The model should also be presented in a way that dramatises its assumptions and relates them to policymakers' experience.

The idea of 'dramatising' a model has led to the development of 'policy workshops' and has brought renewed interest in role-playing simulation games. [14] In both cases the modeler (perhaps best thought of as a facilitator/modeler) creates a 'learning environment' for policymakers that makes them feel part of the model situation. In principle, policymakers who are placed in such an environment come to relate their own experience more closely to the model than they would in a conventional model presentation. They also internalise more readily the 'lessons' about dynamic behaviour that the model contains.

There are several people paying close attention to workshops and games who are discovering how to design effective learning environments. Kreutzer

[13] Strictly speaking, the hoarding loop is not cut but instead neutralized, because equipment delivery time stays constant in the partial model.

[14] Another way to provide a context of drama and realism for a policy model is to integrate it with a conventional case study. The case, which let us suppose is a business policy case, provides general information on a selected company: the situation or problems it faces; the industry it competes in, its products and markets, its history, its organisation and administration and the personalities of its leaders (see for example Christensen, Andrews and Bower [6]). The model provides discussants with additional information on the business which may include maps, descriptions of policy functions, algebra and simulations. The case and model are complementary. The case reveals the strategy, structure and operational decisionmaking of the business, while the model shows graphically (with maps and simulations) how the company's resistance to strategic change arises from its structure and the inertia of its operational decisionmaking. The case/model combination can help probe the crucial (and often weak) link between strategy and operating policies. Work in this interesting area is only just beginning.

[30] has developed a workshop for educators and students to explore the dynamics of an arms-race. The workshop builds on a small (20 equation) dynamic model of an arms-race created by Forrester [19]. Forrester's model represents in outline the decisionmaking processes used by two countries, X and Y, for estimating their opponent's stock of arms, for judging the adequacy of their own stock of arms and for procuring arms from industrial military suppliers. The model also includes levels that represent the stock of existing arms and new arms under development. The decisionmaking network of the model captures in very interesting ways the lags, distortions and biases that occur in the transmission and processing of sensitive military and political information. The model's treatment of information flow is a fascinating example of bounded rationality in the military and political domain. The dynamic properties of the model (exponential growth in the stock of arms of both countries X and Y) arise from the imperfections assumed in the system's decisionmaking processes.

Kreutzer's workshop immerses participants in the realities of military and political decisionmaking. They are provided with articles on the arms-race from magazines like *Newsweek* and the *Economist*. They are presented with charts showing the history of the Soviet–U.S. arms race. They are given cartoon illustrations from magazines like *Punch* or the *New Yorker* which portray (in amusing but memorable and usually realistic ways) the imperfections of military intelligence (for example, an illustration showing large crates being shipped to Cuba on anonymous freighters, and two military officers debating the likely contents of the still-closed crates). All this material activates participants' mental models of the arms-race and highlights the role of information processing and information feedback in arms control. With this preparation participants are able to relate their knowledge and experience to the model and to appreciate the assumptions that underlie the model's feedback structure and dynamic behaviour.

In his workshop, Kreutzer also uses a series of partial model tests to show how the decisionmaking processes that generate an arms-race are quite 'sensible and benign' when the imperfections and biases in information processing are eliminated. (e.g. simulations which assume that decisionmakers in countries X and Y have perfect knowledge of their opponent's stock of arms (both installed and in development) exhibit much slower exponential growth, or in some cases, no growth at all.

Role-playing games fulfill a very similar function to workshops by providing a context of realism and drama to relate policymakers' knowledge to simulation models. In the case of games the drama is provided by making participants play the role of a selected decisionmaker/s in the model system.

The production distribution 'hand simulation' game is a good example of a game that promotes learning and policy debate. It is a board game played by teams of four players. Each player takes a role as either retailer, wholesaler, distributor or manufacturer in a vertically integrated manufacturing and supply system (a beer production and distribution system is usually selected). A player is responsible for managing inventories and backlogs at one point in the system (e.g. wholesaler) and for placing orders with the adjacent player downstream (e.g. distributor) in the supply chain. The objective of the players is to minimise the team's inventory and backlog costs in the face of exogenous customer orders. The volume of orders is not known in advance by any player, and is revealed week-by-week to only the retailer. The game shows the difficulty of coordinating decisionmaking and action in a system with decentralised decisionmaking and imperfect information processing. Almost all teams that play the game incur inventory and backlog costs which are much higher than the 'theoretical' cost minimum. Recently, Senge [15] and Sterman [65] have refined the game instructions, improved the board layout and documented an effective protocol for debriefing so that the game now serves as a general purpose introduction to feedback systems and modeling.

The production-distribution game uses coins, paper, and a plastic printed board. But many new electronic or semielectronic role-playing games are being developed. Meadows is spearheading the development of games for policymakers and has

[15] Senge uses the game regularly with chief executives in a short course on Systems Thinking. The game materials (board, cards, instructions, debriefing protocol) are available from Innovation Associates, PO Box 2008, Framingham MA 01701, USA.

devoted much of his research effort to the topic. His games are aimed not only at policymakers in business, but also public policymakers as for example with the game STRATAGEM 1 (Meadows [35]). He has also collaborated with Sterman to produce STRATAGEM 2, the role-playing simulation game of economic investment cycles mentioned earlier [67].

More policy games have been developed such as Hall's magazine publishing game [23], Flint's multi-product salesforce game [10] and Habibe's arms-race game [22]. The management consulting company Pugh-Roberts Associates has developed a role-playing game for managers of large-scale projects (based on a comprehensive simulation model of project management [48]). Further games are under development at several modeling centres both as learning environments and as laboratories for research in behavioural decisionmaking.

Generic policy models

Generic policy models are (usually small) models which display important dynamic processes that occur frequently in business and social systems. In addition generic policy models often encapsulate some 'managerial wisdom' in the form of 'principles' for effective policymaking in situations like the one represented by the model.

Examples of important and general dynamic processes are captured in such words and phrases as growth, decline, saturation, goal-seeking, fluctuation, goal erosion, worse-before-better, better-before-worse, vicious cycle, snowball effect, bandwagon, productivity trap, poverty trap, policy resistance, compensating feedback and others. For each dynamic process there is a corresponding feedback structure. For example, Forrester's 'market growth model' [14] contains three important feedback loops that interrelate production capacity, salesforce size and customer orders in a growth company (or business unit). One loop—which combines policies for budgeting, salesforce hiring and customer ordering—generates exponential growth in sales. A second loop—which combines policies for capacity expansion and for control of delivery time—adjusts production capacity to changes in customer orders. The third loop–which links customer's ordering and company production—ensures a long run balance between customer orders and production capacity. The complete model can exhibit dynamic behaviour ranging from exponential growth in sales and capacity, to stagnation and decline. This range of dynamic behaviour comes from combining dynamic processes (and associated feedback structures) of growth, saturation, fluctuation and goal erosion.

Generic models offer modelers and policymakers a way of collecting and storing knowledge about feedback structure and dynamics of social and business systems. [16] Each generic model is a self-contained 'behavioural theory' of the dynamic processes it illustrates. Because generic models store dynamic theories (and 'insight' into dynamic behaviour) they have attracted research effort, especially in the last five years. But, progress so far has been disappointingly slow, partly because it is difficult to prove that a given model is truly general, and partly because the word 'generic' means different things to different people. [17]

Despite these problems, research on generic models remains an important topic, (see for example Paich [44]) particularly given the growing emphasis on models and simulations for learning and reasoning. A comprehensive 'library' of generic models (Forrester [17]) would help modelers organise descriptive information about a system. One could identify important feedback loops in policy diagrams and then, from knowledge of simulations of similar structures in the 'library', analyse growth, fluctuations, decline and saturation in system performance.

[16] Lyneis' book *Corporate Planning and Policy Design* [33] contains several well-documented generic models of manufacturing systems.

[17] Some modelers think of generic models in terms of feedback structures, dynamic behaviour and insights which are of use *across the board* in social, physical, ecological and biological systems. These people tend to focus on patterns and types of interacting feedback loops, and on abstract dynamic behaviour. Others think of generic models in terms of policy structure, dynamic behaviour and insights which are of use in *many policymaking situations*. These people tend to focus on policy interactions and feedback loops that arise from connecting policy functions, and on managerial principles that stem from dynamic behaviour.

Future research—Improving model supported 'dialogue' and the mapping of policymakers' knowledge

An important objective of future policy-related research in system dynamics is to improve the quality of dialogue and debate among policymakers and between policymakers and models. Better dialogue comes from capturing accurately in maps and models policymakers' knowledge of business, and from strengthening the influence of model-generated opinion in policy debate. Many research paths are open to improve model-supported dialogue. They include field experiments, behavioural decisionmaking, game design and mapping technology.

Field experiments

Field experiments are already underway to explore the process for generating effective model supported policy dialogue. The experiments are taking place in both large and small business organisations in the United States and Europe. Researchers and consultants are experimenting with the content and sequence of model development to better understand which modeling activities should be conducted during meetings and which beforehand; to better understand what balance to strike between qualitative mapping and simulation; and to better understand how to use partial model tests and simple scenarios to challenge policymakers' intuition.

Researchers and consultants are also experimenting with the composition of the project team (the mix of policymakers, modelers and facilitators),[18] the format of meetings (how frequent, how long, and what mix of discussants), and the 'technology' for presenting and recording policy debate (flip-charts, blackboards, paper, overheads, video projectors and computers (with word-processing, diagramming and modeling software)).

Several recent papers describe the style and direction of the field work. Richmond [55] and Senge [59] describe a 'Strategic Forum' which they view as a 'process' to enable a cross-functional management team to improve the match between operating policies and stated strategic objectives.[19] A forum involves several work steps for a management team: articulating current vision and strategy, developing simple 'reality check' models, developing more complex models by closing feedback loops, conducting 'what-if' policy testing and defining action steps. Morecroft [40,42,43] describes 'strategy support models' which are intended to 'provide executives with insight into whether the policies and programs (of a business strategy) are properly coordinated and whether they are in fact capable of achieving the market and financial objectives called for by the strategy'. He describes two phases of modeling, a first qualitative mapping phase to identify 'players', policies, and feedback structure, and a second simulation modeling phase to develop equations and concepts and to debate the outcome of simple simulated scenarios.

It is interesting to note that research and consulting on the process of model-building with management teams is already well-established outside the system dynamics field. Well-known work has been carried out by Phillips [46,47] and Eden [8,9] and the topic is receiving increasing attention in the area of decision support systems (Land et al. [31]), Keen and Scott-Morton [29], Lorange et al. [32]). Some cross-fertilisation of research and methods would likely be fruitful.

Behavioural decisionmaking and gaming

The value of behavioural decision theory to system dynamics is clear enough: its ideas can help modelers to ask better questions of policymakers, to specify decision processes more accurately, and to capture more or policymakers' knowledge in maps and algebra. It is likely that modelers can pull still more ideas from the combined literature of the Carnegie school and mod-

[18] Roberts [56] wrote an influential paper in 1972 that highlights several key issues in the modeling process: project selection, project team composition, pace of model development, model detail and communication/implementation of model-based recommendations. New field research in system dynamics is examining these issues in more depth.

[19] Senge is using the strategic forum in a research program at MIT's Sloan School of Management called 'Systems Thinking and the New Management Style'. The program, which involves leaders of some of America's most innovative corporations, is exploring how systems thinking can be developed within the participating organisations (Senge [57,58]).

ern behavioural decision theorists. An important extension to this bridge-building is to embody the new ideas explicitly into symbols for mapping (say by including information filters in maps) and into protocols for questioning policymakers.

Another significant area for research is game design. Behavioral decision theory gives some guidance to game design by focussing the game-builders' attention on the design of the 'decision-shell' in which human subjects will role-play. Immediately one thinks of 'designing a decision shell' then game-building takes on many interesting research dimensions (that go well beyond the purely technical issue of outfitting a simulation model with the capability for occasional human intervention). There is the question of how one 'replicates' the organisational, cultural and administrative filters (of information) that condition choice and action. What information (from the vast matrix of simulated data available) should be presented to gameplayers? How should screens of information be organised? What balance of graphic, verbal and visual displays is appropriate? How much leakage of information between players should be allowed in multi-player games? What is an appropriate protocol for gaming-decisions? How should one gauge the adequacy and fidelity of the decision-shell? The research questions are numerous. At a more technical level one might consider the merits of different programming environments and computers for developing behavioural decision shells.

Finally, there is a challenging, and potentially large, research topic in the use of gaming to *link experimentally* the behavioural decisionmaking of individuals and groups to the dynamics of large organisations. In this kind of research a simulation game becomes a laboratory for 'testing' cognitive limitations of individuals and groups in environments that 'simulate' large organisations. Subjects make choices in an experimentally controlled setting (the decision shell) that provides operating information. The operating information is generated by a simulation model that 'surrounds' the decision shell. Subjects are free to make any choice they consider appropriate, given the available operating information, their knowledge of operating goals and incentives, their 'mental model' of how the rest of the organisation operates, and also given their own cognitive limitations. The actions and reactions of the rest of the organisation (comprising several behavioural decision functions, actions and levels) are represented by algebraic functions and simulated during the game. Since the business situation is entirely experimental, one can replace the decision shell and human decisionmaker/s with an algebraic decision rule and discover (through analysis or simulation) an 'optimal' decision rule. Knowing an optimal decision rule and the results of many game trials with many different players, one can discover if and when people use systematically poor decisionmaking heuristics. One can also 'model' the players' heuristics and compare them with the optimal decision rule in order to probe the link between cognitive limits and observed dynamic behaviour.

Research along these lines is being carried out by Sterman [69]. It is a fascinating area that promises to yield better understanding of the reasons for (economically) inefficient dynamic behaviour in business and social systems; experimental methods for validating model assumptions; and new insights into the design of roleplaying simulation games.

Better mapping technology

There is a large potential for research which leads to better mapping technology and therefore to a richer flow of policymakers' knowledge into maps and models. The most direct research path leads straight to improvements in software. A more ambitious research path leads into aspects of modern computer science and artificial intelligence.

Software for mapping, modeling and simulation has improved over the past five years, as outlined earlier. However, there is room for still more improvement. Mapping (of the kind allowed by STELLA) should permit word-and-picture maps to be built at the level of policies (Morecroft [38]) rather than at the present level of algebraic converters. Word-and-picture maps would allow better communication with policymakers (because the maps are readable, visually compact and easily changeable) and would guide equation formulation without constraining conversation (because they stand in a natural hierarchy above equation formulation). The needed software should combine the flexibility of drawing and writing packages (say like MacDraw and MacWrite) with the modeling capability of STELLA.

New software should also help modelers write good clear algebra that a policymaker can (almost literally) read! A simple step is to allow much longer labels so that equations look like sentences. Also needed, but more difficult to provide, is guidance for equation formulation—a computer environment for developing equations that weeds-out poor formulations. Here is an ambitious but clear research challenge: to capture in a software package (at least some of) the expert modelers' rules of formulation (for example, dimensional consistency checks and extreme-condition tests).

Finally, new software should give modelers more simulation power and flexibility. Given a credible model, one should be able to probe 'policy parameter' space as quickly as one can envisage and articulate meaningful policy scenarios. The required flexibility here is not only for rapid re-simulation, but more important, for rapid reformating and reorganisation of simulated graphs and charts. Some original thinking on the 'visual display of quantitative information' (Tufte [70]) is called for.

The most ambitious research path leads into modern computer science and artificial intelligence. Here I will speculate from the perspective of a knowledgeable modeler and policy analyst but a relative novice in modern computer science.

The challenge is to better understand how to elicit and reconstruct policymakers' broad business knowledge into meaningful word-and-picture maps, algebraic 'sentences', models and simulations. It seems to me that an important prerequisite is to discover more precisely what we mean by the phrase 'policymakers' knowledge'. Branches of Artificial Intelligence (AI) may provide some answers (for an eloquent and authoritative introduction to AI see Minsky's *Society of Mind* [36]). However, there is a need for focus. The likely criterion for achieving focus is to select the work that is most informative on how symbols (words, charts, pictures, etc.) can be used to provide a 'framework' on which to hang policymakers' knowledge. [20]

[20] A recent article by Geoffrion [20] on structured modeling provides some stimulating ideas on frameworks for modeling and representation schemes (the use of graphs, charts, text-based schema and 'elemental detail tables'). The article also provides references to AI literature on knowledge representation.

I do not know which branch of artificial intelligence/computer science would be most useful though I can think of areas that are related. For example, the branch of expert systems has made headway in coding specialists' (narrow domain) knowledge into collections of facts and inference engines for relating the facts. One might say that expert system designers have created frameworks that help (narrow domain) experts articulate their specialist knowledge and transfer it into computer models. In addition one might note that the structure of the frameworks (the facts and the rules of inference) both guides *and limits* the form in which knowledge is articulated and collected. Expert systems then provide examples of frameworks for mapping knowledge. I expect that deliberately designed frameworks to capture policymakers' knowledge will differ radically in content and structure from expert systems (since policymakers' judgements draw on broad (rather than specialist) experience and knowledge). Nevertheless some clues to the design of policy frameworks might emerge from expert systems. Almost certainly there are other branches of research in artificial intelligence which could offer more (and perhaps more relevant) clues.

Let us suppose now that there exist powerful qualitative frameworks to capture policymakers' knowledge in words-and-pictures. Now consider the more specific research issue of converting words-and-pictures into simulations that can be used to challenge policymakers' intuition. System dynamics has a particular 'conversion technology'. It is quite effective, but it has remained fundamentally unchanged for the past twenty five years. Is there room for improvement, and where should one look for new ideas?

In my view, a weakness in the existing conversion technology is in the link between word-and-picture maps and algebra. It is difficult to write good algebra that means mathematically what you intend with words, and some ideas just don't seem to fit very well into the mould of algebra and differential equations. Are there alternatives to algebraic equations and would they be better? I don't pretend to know the answer to this question, but I think it is worth exploring as a research topic. Moreover, new programming/modeling approaches from modern computer science may prove helpful.

I have some familiarity with the programming language LOGO (Papert [45]). LOGO is used to create learning environments (microworlds) to help schoolchildren (and university students) to understand 'difficult' or abstract concepts such as force, momentum and energy in physics. A LOGO microworld is built using LOGO primitive commands each labeled in plain English. LOGO commands can be grouped together and given a name to form a new command. New commands built in this way can be abstractions for complex concepts. For example, if one is learning about relative motion one can use a relative-motion microworld (Linda E. Morecroft [37]). Such a microworld is constructed from a set of commands that simulate motion. The commands are used to set objects in motion on a computer screen and then combined to study relative motion. Thus, objects set in circular motion can be used to investigate complex patterns of relative motion that occur when one object moving in a circle is viewed from another. The point here is to show that the structure of the programming language allows the construction of micro-worlds in which complex situations can be studied by playing with and combining simple building blocks. I do not know what the structure of an analogous language for policy systems might be, but I am sure it differs from conventional simulation modeling languages.

I have outlined some promising paths for future research in system dynamics. A lot has been accomplished over the last ten years, but the remaining opportunities and challenges are enormous. Future research should provide the technology, theory and group processes for policy microworlds which will (in Richmond's words) 'help organisations design their own future'.

References

[1] Allison, Graham, T., *Essence of Decision*, Little Brown, Boston, MA, 1971.

[2] Barnard, Chester I., *The Functions of the Executive*, 29th printing, Harvard University Press, Cambridge MA, 1982.

[3] Cavana, Robert Y., and Coyle, R. Geoffrey, *Dysmap User Manual*, University of Bradford Publications, Bradford, Yorkshire, UK, 1982.

[4] Checkland, Peter, *Systems Thinking, Systems Practice*, Wiley, Chichester, 1981.

[5] Christensen, C. Roland, Andrews Kenneth R. and Bower, Joseph L., *Business Policy: Text and Cases*, Irwin, Homewood, IL, 1978.

[6] Coyle, R. Geoffrey, *Management System Dynamics*, Wiley, Chichester, UK, 1977.

[7] Eden, Colin, Jones, Sue, and Sims, David, *Messing About in Problems*, Pergamon Press, Oxford, 1983.

[8] Eden, Colin, "Perish the Thought", *Journal of the Operational Research Society* 36(9) (1985) 808–819.

[9] Eden, Colin, "Managing Strategic Ideas: The Role of the Computer", *ICL Technical Journal* 5/2 (1986) 173–183.

[10] Flint, Brilsford, B., "A role-playing simulation for sales planning and control", Unpublished Master's Thesis, Sloan School of Management, MIT, Cambridge, MA, May 1986.

[11] Forrester, Jay W., *Industrial Dynamics*, The MIT Press, Cambridge MA, 1961.

[12] Forrester, Jay W., *Principles of Systems*, The MIT Press, Cambridge MA, 1969.

[13] Forrester, Jay W., "Market growth as influenced by capital investment", *Sloan Management Review* 9/2 (1968) 83–105. Also in *Collected Papers of Jay W. Forrester*, pp 111–132, MIT Press, Cambridge, MA, 1975.

[14] Forrester, Jay W., "Counterintuitive behavior of social systems", pp. 211–237 in *Collected Papers of Jay W. Forrester*, MIT Press, Cambridge, MA, 1975.

[15] Forrester, Jay W., "Business structure, economic cycles and national policy", *Futures* 8 (1976) 195–214.

[16] Forrester, Jay W., "An alternative approach to economic policy: Macrobehavior from microstructure" in: N. Kamrany and R. Day (eds.), *Economic Issues of The Eighties*, Johns Hopkins University Press, Baltimore, MD, 1979.

[17] Forrester, Jay W., "System dynamics—Future opportunities", in *System Dynamics, TIMS Studies in the Management Sciences*, Vol. 14, North-Holland, New York, 1980, 7–21.

[18] Forrester, Jay W., Graham, A., Senge, P., and Sterman, J., "An integrated approach to the economic long wave", Working paper D-3447-1, System Dynamics Group, MIT, Cambridge, MA, 1983.

[19] Forrester, Jay W., "Dynamic modeling of the arms race", System Dynamics Group Working Paper D-3684-3, Sloan School of Management, MIT, Cambridge, MA, 1985.

[20] Geoffrion, Arthur M., "An introduction to structured modeling", *Management Science* 33/5 (1987) 547–588.

[21] Geus, Arie P. de, "Planning as learning: The adaptive corporation", presentation paper from the *Shell Planning Conference*, Banff, Alberta, Canada, May 1986.

[22] Habibe, Tommy O., "The arms race game", System Dynamics Group Working Paper D-3836, Sloan School of Management, MIT, Cambridge, MA, June 1986.

[23] Hall, Roger I., "Managing a magazine publishing company: A decision making game", in: T. Carney (ed.), *Constructing Instructional Simulation Games*, University of Manitoba, Winnipeg, Manitoba, 1974, 22–29.

[24] Hall, Roger I., "A system pathology of an organization: The rise and fall of the old Saturday Evening Post", *Admin. Sci. Quart.* 21 (1976) 185–211.

[25] Hall, Roger I., "Decision making in a complex organization", in: G.W. England, A. Neghandi and B. Wilpert (eds.), *The Functioning of Complex Organizations*, Oelgeschlager, Gunn and Hain, Cambridge, MA, Chapter 5, 1981, 111–144.

[26] Hall, Roger I., and Menzies William, "A corporate system model of a sports club: Using simulation as an aid to policy making in a crisis", *Management Science* 29 (1983) 52–64.

[27] Hall, Roger I., "The natural logic of management policy making: Its implications for the survival of an organization. *Management Science* 30/8, (1984).

[28] Hogarth, Robin M., *Judgement and Choice*, Wiley, New York, 1980.

[29] Keen, Peter G.W., and Scott-Morton, Michael S., *Decision Support Systems*, Addison-Wesley, Reading, MA, 1978.

[30] Kreutzer, David P., "A microcomputer workshop exploring the dynamics of arms races", System Dynamics Group Working Paper D-3689-1, Summer 1985.

[31] Land, Frank, Gall, Michael, Hawgood, John, Miller, Gordon, and Mundle, Fred, (editors and organisers) "Knowledge Based Management Support Systems", *Proceedings of International Business Schools Computer User's Group and Information Systems Association*, Joint European Meeting, London Business School, April 1987.

[32] Lorange, Peter, Scott-Morton, Michael S., and Ghoshal, Sumantra, *Strategic Control Systems*, West Publishing Company, St. Paul, MN, 1986.

[33] Lyneis, James M., *Corporate Planning and Policy Design: A System Dynamics Approach*, MIT Press, Cambridge, MA, 1980.

[34] Mass, Nathaniel, J. "Diagnosing surprise model behavior: A tool for evolving behavioral and policy insight", *Proceedings of the 1981 System Dynamics Research Conference*, Rensselaerville, NY, 14–17 October 1981, 254–272. Also available as *MIT System Dynamics Group Working Paper D-3323*, MIT Sloan School of Management, Cambridge, MA 02139, 1981.

[35] Meadows, Dennis L., "STRATAGEM 1: A resource planning game", *Environmental Education Report and Newsletter* 14/2 (1985) 9–13.

[36] Minsky, Marvin, *The Society of Mind*, Simon and Schuster, New York, 1986.

[37] Morecroft, Linda E., "A relative-motion microworld", SM Thesis, Laboratory for Computer Science, Publication no. MIT/LCS/TR-347, MIT, Cambridge, MA, September 1985.

[38] Morecroft, John D.W., "A critical review of diagramming tools for conceptualizing feedback system models", *Dynamica* 8, Part 1 (1982) 20–29.

[39] Morecroft, John D.W., "System dynamics: Portraying bounded rationality", *Omega* 11/2 (1983) 131–142.

[40] Morecroft, John D.W., "Strategy support models", *Strategic Management Journal* 5/3 (1984) 215–229.

[41] Morecroft, John D.W., "Rationality in the analysis of behavioral simulation models", *Management Science* 31/7 (1985) 900–916.

[42] Morecroft, John D.W., "The feedback view of business policy and strategy", *System Dynamics Review* 1/1 (1985) 4–19.

[43] Morecroft, John D.W. and Paich, Mark, "System dynamics for reasoning about business policy and strategy", Centre for Business Strategy Working Paper 30, London Business School, Regent's Park, London, UK, March 1987.

[44] Paich, Mark, "Generic structures", Research problems section *System Dynamics Review* 1/1 (1985) 126–132..

[45] Papert, Seymour, *Mindstorms*, Basic Books, New York, 1980.

[46] Phillips, Lawrence, D., 'Computing to consensus', *Datamation* 68 (1986) 2–6.

[47] Phillips, Lawrence D., "Decision support for managers", forthcoming in: Harry J. Otway and Malcolm Peltu (eds.), *The Managerial Challenge of New Office Technology*, Butterworths, London, 1987.

[48] Pugh-Roberts Associates, *Project Management Modeling System PMMS*, Pugh-Roberts Associates, Five Lee Street, Cambridge, MA, 1986.

[49] Pugh-Roberts Associates, *Professional DYNAMO Introductory Guide and Tutorial*, and *Professional DYNAMO Reference Manual*, Pugh-Roberts Associates, Five Lee Street, Cambridge, MA, 1986.

[50] Richardson, George P. and Alexander L. Pugh, *Introduction to System Dynamics Modeling with DYNAMO*, MIT Press, Cambridge, MA, 1981.

[51] Richardson, George, P., "Problems with causal loop diagrams", Archives section, *System Dynamics Review*, 2/2 (1986) 158–170.

[52] Richmond, Barry M., "STELLA: Software for bringing system dynamics to the other 98%", *Proceedings of the 1985 International Conference of the System Dynamics Society*, Keystone, CO, July 1985.

[53] Richmond, Barry M., *A Users Guide to STELLA* (2nd printing), High Performance Systems Inc., 13 Dartmouth College Highway, Lyme, NH, November 1985.

[54] Richmond, Barry M., Peter Vescuso and Steven Peterson, *STELLA for Business*, High Performance Systems, 13 Dartmouth College Highway, Lyme, NH, 1987.

[55] Richmond, Barry M., "The strategic forum: From vision to operating policies and back again", *High Performance Systems Publications*, 13 Dartmouth College Highway, Lyme, NH, 1987.

[56] Roberts, Edward B., "Strategies for effective implementation of complex corporate models", in: Edward B. Roberts (ed.), *Managerial Applications of System Dynamics*, Chapter 4, MIT Press, Cambridge, MA, 1978, 77–85.

[57] Senge, Peter M., "Systems thinking in business: An interview with Peter Senge", *ReVISION* 7/2 (1984).

[58] Senge, Peter M., "Systems principles for leadership", in J. Adams (ed.), *Transforming Leadership*, Miles River Press, Alexandria, VA, 1986.

[59] Senge, Peter M., "Catalyzing systems thinking within organizations", System Dynamics Group Working Paper D-3877-2, Sloan School of Management, MIT, Cambridge, MA, March 1987.

[60] Simon, Herbert A., "Rationality and Decisionmaking", in: *Models of Man*, Wiley, New York, 1957.

[61] Simon, Herbert A., *Administrative Behavior*, third edition, The Free Press, New York, 1976.

[62] Simon, Herbert A., "Rational decision making in business organizations", *The American Economic Review* 69/4 (1978).

[63] Simon, Herbert A., *The Sciences of the Artificial*, MIT Press, Cambridge, MA, 1982.

[64] Simon, Herbert A., *Models of Bounded Rationality Vol. 2: Behavioral Economics and Business Organization*, MIT Press, Cambridge, MA, 1982.

[65] Sterman, John D., "Instructions for running the beer

distribution game", System Dynamics Group Working Paper D-3679, October 1984.

[66] Sterman, John D., "A Behavioral Model of the Economic Long Wave", *Journal of Economic Behavior and Organization* 6/1 (1985) 17–53.

[67] Sterman, John D. and Meadows, Dennis L., "STRATAGEM-2: A microcomputer simulation game of the kondratiev cycle", *Simulation and Games* 16/2 (1985) 174–202.

[68] Sterman, John D., Testing behavioral simulation models by direct experiment", System Dynamics Group Working Paper D-3783-1, Sloan School of Management, MIT, March 1986, forthcoming in *Management Science*.

[69] Sterman, John D., "Misperceptions of feedback in dynamic decisionmaking", System Dynamics Group Working Paper, D-3876-1, Sloan School of Management, MIT, May 1987.

[70] Tufte, E.R., *The Visual Display of Quantitative Information* Graphics Press, Cheshire, CT, 1983.

[71] Tversky, A., and Kahneman, D., "Judgement under uncertainty: Heuristics and biases", *Science* 185 (1974) 1124–1131.

[72] Vapenikova, O. and Dangerfield, B., *DYSMAP2 User Manual*, University of Salford, 1987.

[73] Wolstenholme, E.F., "System dynamics in perspective", *The Journal of the Operational Research Society* 33/6 (1982) 547–566.

[74] Wolstenholme, E.F. and Coyle, R.G., "The development of system dynamics as a methodology for system description and qualitative analysis", *The Journal of the Operational Research Society* 34/7 (1983) 569–581.

[75] Wolstenholme, E.F., "System dynamics: A system methodology or a system modeling technique?", *Dynamica* 9, Part II (1983) 84–90.