# Reflections on System Dynamics in 2025

The Times They Are A-Changin’

John Morecroft

Retired Senior Fellow

Management Science and Operations

London Business School

May 2025 (minor revisions in August 2025)

Paper presented at the 2025 International System Dynamics Conference

Boston, United States

August 3-7, 2025

# Reflections on System Dynamics in 2025

The Times They Are A-Changin'

## Introduction

After fifty years of system dynamics I decided that now is a good time to reflect on the field, where it has been and where it is going. I begin with comments on the character of system dynamics models gleaned from my own training and subsequent academic career. Then I examine a snippet of contemporary system dynamics in the era of AI, using a critical review of a 2024 paper about the use of generative AI for model conceptualization. Next I return to my experiences of the field in a section about 'Living with System Dynamics'. Here I insert a short paper, from 2021, to further illustrate the character of SD and style guidelines for modeling. The contrast between my characterization of the field and the selected AI snippet brings me to the final section 'Where to Next?' There's a shift underway in modeling practice, which is also the focus of ISDC2025. I consider this shift, drawing on the evocative lyrics of Bob Dylan's immortal song "The Times They Are A-Changin'".

## The Character of System Dynamics Models

Imagine you're an experienced modeler and you read a report about a system dynamics modeling project. In this thought experiment the author of the report is assumed to be anonymous. Chances are you can easily discern if the report is the work of an expert or a novice. In fact, you may even be able to name the author[1]. So, what is it about the system dynamics modeling approach that is revealed in the work of experts? Here I explore three features: style of analysis, sources of information, and portrayal of decision making. I should emphasize that this three-pronged characterization of the field applies to studies that lead to full-blown simulators. It is not intended to apply to qualitative studies in which system dynamicists focus on effective methods for group model building and participative causal mapping (Stave et al 2024), where different criteria would apply.

Style

Traditional system dynamics involves a distinctive 'style' of modeling and analysis. It lays strong emphasis on clear visualization and documentation of real-world feedback structure backed-up by rigorous yet easy-to-read equation formulations. Understanding of dynamics comes from careful narrative interpretation of simulations.

This style is clearly illustrated in a classic archive paper by Mass and Senge (1975) entitled 'Understanding Oscillations in Simple Systems'. The paper was a core reading for doctoral students

[1] I recall a situation where a degree student submitted an assignment containing some elegant formulations. It turned out they were lifted from *Urban Dynamics* without acknowledgement. The student took a roasting. And here's a little test for readers. How would you rate the expertise of the modeler who wrote the following narrative about the requirements for the growth of a firm. "Two requirements must be met by any situation if it is to qualify as one capable of exhibiting the transient phenomenon of growth: 1. A potential for growth must exist. The market must be capable of being stimulated to order more, or at a higher rate, as time goes on. Whether or not the potential is exploited, and growth does occur, is a matter determined not by the potential but by the interactions within the relevant system. The important point is that growth potential must exist in the market in order that growth might occur; 2. A desire for growth must exist within the company. No natural law forces an industrial enterprise to grow. The expansion occurs because the company responds to conditions indicative of a potential for a greater volume of business in a way which is aimed at achieving the capability for the handling of increased volume. The corner drugstore, as well as many other kinds of firms today, simply does not want to grow – thus policies are adopted which keep demand at the same level as productive capacity over the long run". Readers might also try to name the author (answer to be found in the references under 'Anonymous Working Paper').

in MIT Sloan's system dynamics program throughout the 1970s and beyond. The introduction to the paper clearly explains the authors' pedagogical intentions.

"Acquiring a firm intuitive understanding of the possible types of behavior produced by simple first, second, and third-order systems marks an important step in learning system dynamics. Such systems frequently embody generic structures that recur in a wide variety of complex systems. However, an intuitive grasp of simple oscillating systems often eludes both the beginning student and practitioner alike. Even individuals familiar with the mathematics of dynamic feedback systems often cannot provide a simple nontechnical explanation of why a continuous first-order system cannot possibly oscillate, or why a second-order system can. For example, overshoot or oscillation in a system is often explained to result from 'system delays' or 'inertia'. Vague explanations such as these impart little understanding of how decisions being made in a system generate observed problems and behavior".

The authors then go on to present arguments they used to successfully develop learners' insight into simple oscillating systems.

"The paper analyses a one-level model for the population growth of rabbits in a closed field to illustrate why a first-order negative-feedback system exhibits a smooth transition to equilibrium instead of overshoot and oscillating behavior. The paper also analyses a simple inventory-workforce model to provide an intuitive explanation of the causes of convergent, divergent, and undamped oscillations".

Readers from the system dynamics community may feel they already know all this stuff. But the Mass-Senge approach is distinctive with its emphasis on intuitive, non-technical yet rigorous explanations. Consider the depth of behavior analysis. The authors use four pages of clear narrative (accompanied by a stock-flow diagram, fully documented, readable equations and an imagined simulation of overshoot dynamics) to explain exactly why population overshoot and oscillation is impossible for a one-level rabbit model.

"On the initial upswing, rabbit birth rate exceeds rabbit death rate and the population increases. As the limited space [of the field] is filled, reduced food supply and other effects of overcrowding continuously reduce the average rabbit lifetime, causing rabbit death rate to gradually catch up to birth rate. If the overshoot in population is to occur, there must be some point at which births and deaths are balanced and population momentarily stops growing. This point is indicated by a small circle [drawn on the trajectory of the rabbit population in the imagined overshoot simulation]. However, once this point of 'temporary equilibrium' is reached the first-order system cannot move from it. Birth rate and death rate can only vary if population varies. However, population can only vary if there is an imbalance between the birth rate and the death rate. Once birth and death rate come into balance in the first-order system, the balance cannot be tipped and the system is locked into equilibrium".

This verbal argument about an equilibrium balance point for the imagined population trajectory is lucid, compelling and logically rigorous. The authors then continue their line of analysis to argue that a second level or state variable is necessary for any form of oscillation to occur.

In the next section of the paper Mass and Senge develop a similarly compelling verbal argument for the dynamics of a simple inventory-workforce model to show that a system with at least two levels can exhibit oscillatory behavior, though oscillations are not guaranteed. They consider carefully what policies within the system lead to overshoot and oscillation instead of smooth adjustment to equilibrium. The narrative analysis covers fully fourteen pages (again accompanied by a stock-flow diagram, fully documented, readable equations and an imagined simulation of cyclical dynamics in production and inventory – plus an actual simulation of undamped oscillations in inventory and production). It's an elegant argument and a pleasure to read – too long to repeat here, but worthy of close independent study. I adopt a similar style of analysis later in this paper to explain unfolding dynamics in a two-stock model of love between Romeo and Juliet (structurally equivalent to the inventory-workforce model).

Jay Forrester also uses clear narrative analysis of simulated trajectories in his important and widely cited paper 'Market Growth as Influenced by Capital Investment' (Forrester 1968). Many students and practitioners of system dynamics are familiar with this article. Me too. I studied the paper and the underlying simulation model closely during my academic career. I wrote journal articles about the model's policy structure and growth dynamics (Morecroft 1983). This work came to inform my conceptual and applied research (Morecroft 1985, Morecroft, Lane and Viita 1991). I embedded chunks of the analysis in my business dynamics teaching at London Business School for MBAs, PhDs and Executives. The resulting teaching materials underpin several chapters of my book *Strategic Modelling and Business Dynamics* (2nd edition, Morecroft 2015), especially Chapter 7 about 'Managing Business Growth'. There you will find narrative analyses of simulated trajectories, backed-up by structural diagrams, policy formulations and documented equations. Though compact (about 30 equations) the market growth model contains everything of interest in dynamical business models: multiple interacting feedback loops; non-linearities; expectation formation, delays, bias and distortion in information channels; and bounded rationality in operating policies. I used to tell MBA students that if they really mastered and understood the model's formulations and dynamics then they had grasped the nub of business dynamics.

Sources of Information

Often people expect model development to require large data-gathering exercises to compile facts and figures about a firm or industry from written and numerical data bases. However, in system dynamics it has long been recognized there is also a vast amount of relevant data about social systems that resides in the minds of experienced people (their collective mental database). These sources of information, depicted in Figure 1, reveal the inner workings of an organization, its procedures, priorities and even culture. Such operational insights are rarely found in either formal numerical or written databases, yet they are vital to a good representation of structure[2][3].

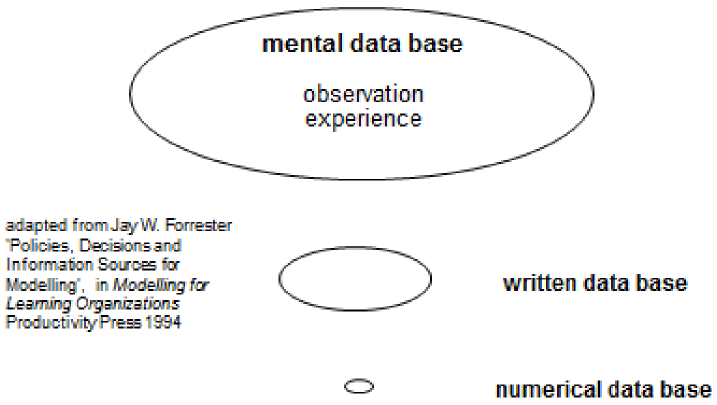


Figure 1 Sources of Information for Modelling

[2] One might argue that the internet, with its vast store of searchable information, has changed the relative size of the three 'data base ovals' and even blurred the categories. However, the three-ovals diagram should not be lightly discarded. It is a clear and powerful image about the nature of data in system dynamics modeling. It suggests that behavioral modelers, in their search for structural clues about interlocking operations, will always need tacit organizational information gleaned from the minds of experienced stakeholders.

[3] Figures 1 to 5 in this section of the paper, and the accompanying text, are taken from chapters 7 and 10 of *Strategic Modelling and Business Dynamics*. Those two chapters also contain full citations of the books and papers mentioned in the accompanying text.

As mentioned above, much of the data for identifying structure resides in the mental database, where it is accompanied by other kinds of system-related data as Figure 2 illustrates. The figure distinguishes three different types of information in the mental database: observations about structure and policies; expectations about system behavior; and actual observed system behavior.

The mental database is particularly rich in structural detail about operations at a functional level. The mental database also contains information about system behavior (past trends and patterns in key variables) that is useful for guiding model conceptualization and for building confidence in simulations. However, not all information in the mental database is reliable for modeling. In particular people's *expectations* about system behavior are often misleading because they cannot reliably infer the dynamics of cause-and-effect loops. Simulation is the only rigorous way to relate the structure and dynamic behavior of business and social systems.

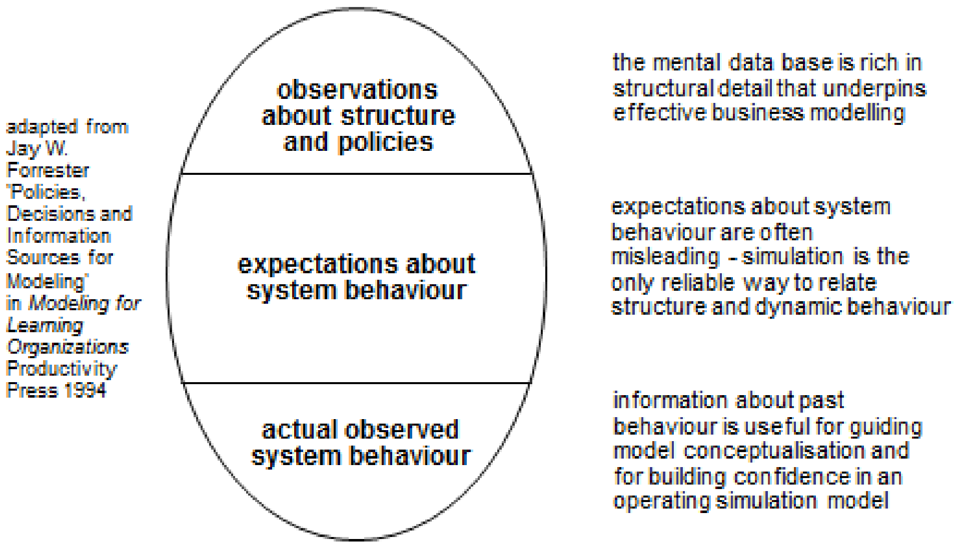


Figure 2 Role of the Mental Data Base in Modelling and Confidence Building

Portrayal of Decision Making - An Information Feedback View of Management and Policy

There is more to well-formulated system dynamics models than causal loops, feedback and stock accumulation – important though they are. There is also an underlying philosophy about management and policy (in all areas of business and society), that I want to reflect on now. A good place to start is with Forrester's original comments about the craft of management. In his seminal book *Industrial Dynamics* (Forrester, 1961, p 93) he views management as the process of converting information into action.

If management is the process of converting information into action, then management success depends primarily on what information is chosen and how the conversion is executed. Moreover the difference between a good manager and a poor manager lies right at this point between information and action.

The viewpoint is cognitive but also action oriented. As Forrester (1994, p 52) goes on to explain:

A manager sets the stage for action by choosing which information sources to take seriously and which to ignore. A manager's success depends on both selecting the most relevant information and on using that information effectively. How quickly or slowly is information converted into action? What is the relative weight given to different information sources in light of desired objectives? How are these desired objectives created from available information?

Here is a parsimonious and stylized portrayal of management, entirely consistent with the principles of information feedback systems, yet capable of capturing a very wide range of managerial behavior. Leadership, charisma and other important attributes of management (vital to business performance and explicitly addressed in the organizational and strategy literature) are represented implicitly and subtly in system dynamics through their influence on the information sources deemed important enough to justify action.

This information-processing view of management leads directly and naturally to the fundamental representation scheme in system dynamics. Again quoting from Forrester (1994, p51):

A simulation model is based on explicit statements of policies (or rules) that govern decision making. The decision making process consists of three parts: the formulation of a set of concepts indicating the conditions that are desired, the observation of what appears to be the actual conditions, and corrective action to bring apparent conditions toward desired conditions.

The verbal description above translates into Figure 3. It represents how managers make adjustments to organizational asset stocks or resources through operating policy. The policy is represented mathematically as: Corrective Action = (Desired Resource − Apparent Resource)/Time to Correct Gap. In a simple formulation, the desired resource is constant and the apparent resource is identically equal to the actual resource.

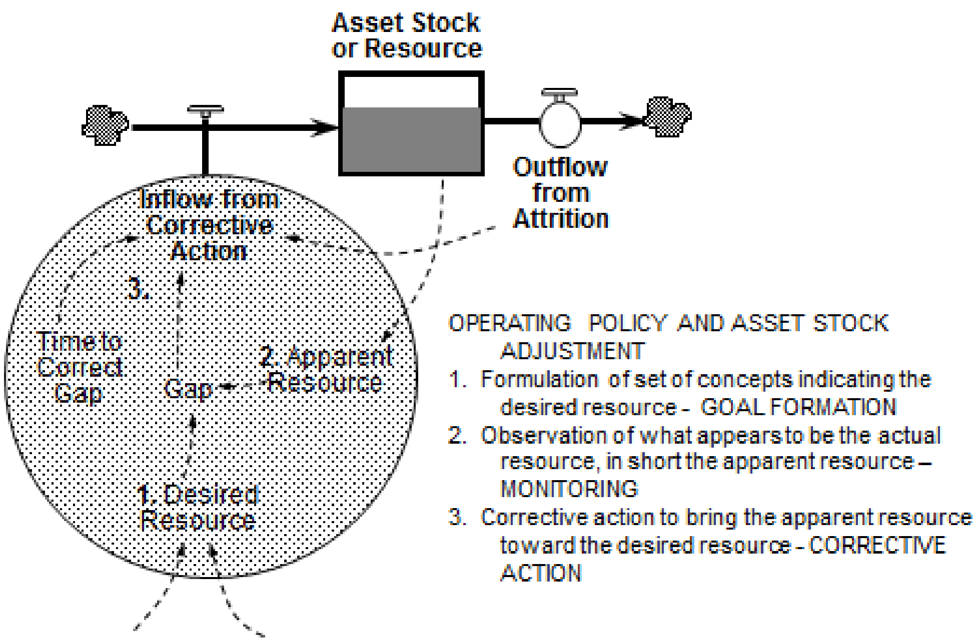


Figure 3 Converting Information into Action
An Information Feedback View of Management

More sophisticated asset stock adjustment can be portrayed too, capturing the ambiguity of resource management found in practice. For example, the apparent resource may differ from the actual resource because of reporting delays, error or bias. In extreme cases, the condition of the actual resource may not be known at all and so cannot be managed or controlled. The desired resource need not be constant but instead varies over time, through a process of goal formation, depending on conditions elsewhere in the business. Though compact, this stock adjustment formulation is very versatile. Moreover, it raises lots of interesting questions about managers' ability to build and to balance organizational resources by focusing attention on: (1) whether, for each resource, there is a clear communicable goal as the basis for corrective action; (2) whether resources are accurately known or not; and (3) whether, having sensed a gap between desired and apparent resource, management react quickly or slowly to close the gap.

Information Available to Decision Makers and Bounded Rationality

According to the feedback view of management, information drives the corrective actions of organizations as they seek to build and maintain a balanced set of resources. Which information sources, among those available, are used by managers? This question lies behind the Baker criterion in Sterman's formulation fundamentals (see Chapter 13 of *Business Dynamics* 'Modeling Decision Making', Sterman 2000). But instead of "what did the President know and when did he know it", the phrasing is modified to read "what did the policy function know and when did it know it"?

In principle, the state of every single resource in an organization is relevant to every point of decision making so that stock adjustments are fully informed. Consider, for example, a supply chain comprising customers, retailers, wholesalers, distributors and a factory. Customers buy from retailers. The retailers order from wholesalers who in turn order from distributors. The distributors then place orders with the factory. In other words, orders flow upstream; originating with customers and ending at the factory. Goods flow in the reverse direction, downstream; originating in the factory and eventually arriving at retailers for purchase by customers. What information is relevant for the factory's production planning? Obviously the factory should take account of distributors' orders because the factory supplies the distributors. The factory should also take account of its own inventory and backlog condition and more besides. Clearly there is a lot of information in the supply chain that, theoretically at least, is relevant to production planning. For example, there is customer demand and the amount of downstream inventory and backlog at every stage of the supply chain. Compiling and making sense of all this data, however, is a huge task, beyond the abilities of normally competent people, even in an era of big data and analytics.

Usually things are much simpler. Factory managers normally pay most attention to tangible information that is available locally such as distributors' orders and factory inventory. The more general point is that decision makers typically use much less information than the total available to them. Moreover, the available information is less than is commonly presumed. One way to think about this selection process is to picture operating policy surrounded by information filters, as shown in Figure 4. The figure shows five possible filters. The first and most basic filter stems from people's cognitive limitations. It doesn't take much information for us to feel overwhelmed and so we pick the signals that seem most relevant and ignore the rest. The torrent of email and social media messages is a daily reminder of our information limits. The formal name for cognitive limitations as they affect decision making is 'bounded rationality', a term first proposed by Nobel laureate Herbert Simon. The concept was developed in the influential literature of the Carnegie School as a theory of firm behavior and an alternative to conventional microeconomic theory (Simon, 1979; 1982).

The other filters are created by the organization itself as it parcels out responsibilities across functional areas and departments. In a sense, the organization compensates for individuals' cognitive limitations by placing them in a departmental structure that means they can get on with their own job without worrying about what is happening in every other department. As Simon (1976) and Barnard (1938) have pointed out, the organization provides a psychological environment that shapes (sometimes very powerfully) how people act and it is the function of the executive to 'design' this environment so departmental efforts are coordinated. The CEO is a designer with various design tools at his or her disposal (Forrester, 1996). Most obvious are operating goals, rewards and incentives (the second filter) that direct people's attention and effort to organizational objectives. Such inducements reduce the complexity of decision making by prescribing what needs to be achieved by different departments. The potential downside is a functional mentality, though this syndrome is not necessarily a problem in a well-designed organization, it just means people are focused (and that could be an advantage).

The next filter represents the effect of information, measurement and communication systems on decision making. Undoubtedly, computers and the internet have massively boosted the amount of information flowing around and between organizations. Again supply chains are a good example. These days it is possible for factories to hold data on retail orders at point of sale and distributors' inventory for use in production planning. However, such data are not necessarily as convenient or trusted as information gleaned from informal channels such as casual conversations, ad hoc meetings and factory walkabouts. Moreover, informal 'intelligence' is often more persuasive and timely than information available on the official database.

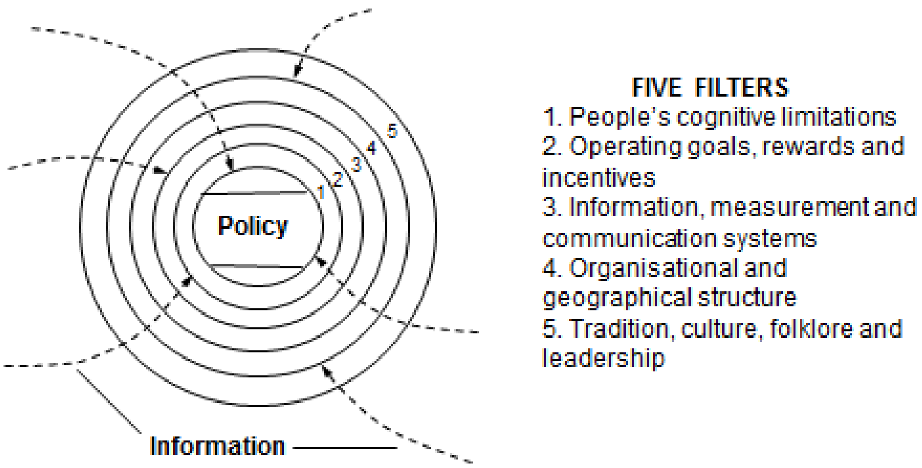


Figure 4 The Policy Function, Information Filters and Bounded Rationality – Behavioural Decisionmaking

The fourth filter represents the effect of organizational and geographical structure. Organizations frequently set up or acquire a new business unit at arms-length from the existing business. A good example is the fledgling low-cost airline Go, set up by BA in the 1990s, to compete with easyJet and Ryanair in the European short-haul business. Go was deliberately made independent of BA, operating its own planes with its own staff and information systems. The whole point was to design an enterprise that could decide and act in its own right, free from the influence of parent BA. The amount of information flowing between the two organizations was much less than if they were seamlessly merged in a single airline. Another example is the MINI car division of BMW, created as an independent business unit, able to take its own decisions (under arms-length corporate guidance) on product development and capital investment in order to develop a new and distinct brand in the highly competitive global small-car market.

The fifth filter is the most subtle yet, and also the most powerful. It captures the attenuating and amplifying effect on information of tradition, culture, folklore and leadership. From a cognitive view these intangible traits shape the psychological environment in which people take decisions and act. They define what the organization stands for and therefore what really needs attention. Consider, for example, the history of Google and its co-founders Larry Page and Sergey Brin. Commentators have said they were intellectually obsessed with an omniscient and omnipotent algorithm for mining the world's knowledge. This belief was (and perhaps still is) part of Google's culture, permeating the minds of thousands of employees and helping to coordinate their actions. Another example is MIT with its culture of technological excellence that pervades all departments including the humanities and management, shaping decisions on faculty recruitment, the curriculum and choice of students.

Nature of Decision Making and the Decision Process

An important conclusion from the discussion so far is that the feedback view of organizations incorporates behavioral decision making and assumes bounded rationality (Morecroft, 1983; Sterman, 1989). This perspective on decision making distinguishes system dynamics sharply from traditional microeconomics in which 'economic man' makes objectively rational decisions, weighing up all available sources of information to arrive at an optimal (profit maximising) configuration of resources. Figure 5 captures the essential philosophical stance to decision making in system dynamics. Any operating policy sits amid its filters, bombarded by information originating from *all* asset stocks in the system. Consistent with the Baker criterion, however, only a handful of information flows penetrate to the policy core leading to action and stock accumulation.

A corollary is that decision making is conceived as a continuous process for converting varying information flows into signals that determine action (rates of flow). In system dynamics, a decision function does *not* portray a choice among alternatives, as found in a decision tree with its various discrete nodes and branches. Neither is it the familiar logic of 'if-then-else' that applies to individual decisions like whether to take a taxi or catch the bus. The crucial point is that we are viewing decision processes from a distance where discrete choices disappear, leaving only broad organizational pressures that shape action.

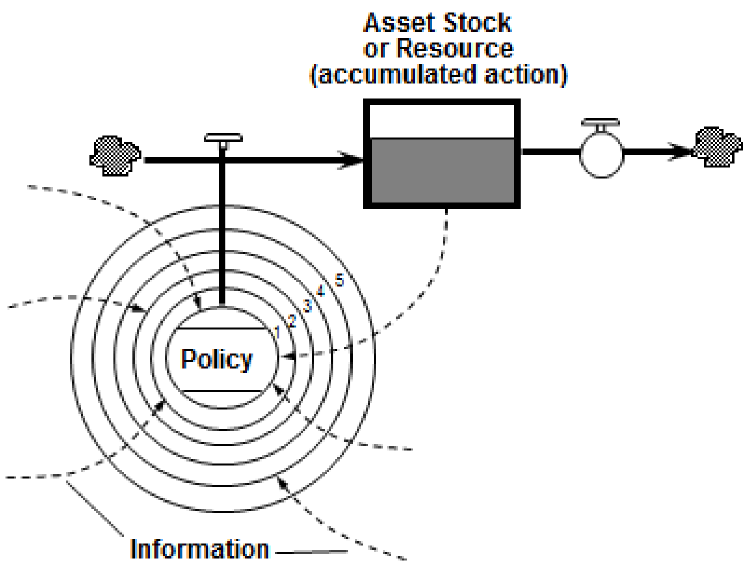


Figure 5 Behavioural Decision Making Leading to
Stock Accumulation & Feedback

Choosing the proper distance from which to view and model a system is crucial. As Forrester has noted, we are not as close as a psychologist concerned with the mechanisms of human thought, delving into the nature and sources of cognition, personality and motivation. We are not even close enough to see each separate decision but instead observe a modulating stream of decisions. We may not be close enough to care whether one person or a group creates the decision stream. On the other hand, we are not as remote as a public stockholder who is so far from the corporation as to be unaware of the internal structure, social pressures and decision points. Modelers need to adopt an intermediate position, similar to the perspective of the Board of Directors, top management team, strategy consultant, or investment banker.

## A Snippet of System Dynamics in the Era of AI

So far I've provided a character sketch of the system dynamics method as I came to know it over many years. My focus has been on the representation of behavioral decision making processes in system dynamics and their effects on the coordination of asset stocks that underpin dynamics and organization performance over time. Puzzling dynamics and underperformance in business and society often arise from poor coordination of asset stocks. Alongside portrayal of decision making, clear verbal interpretation of dynamics is also vital if insights into performance paradoxes are to be persuasively communicated.

But the times they are a-changin', in many ways (Sterman and Repenning 2018). Massive increases in easy-to-access computational power have, over the past two decades, transformed the conduct of simulation-based studies (Sterman 2018). Visually attractive gaming simulators are widely used in management education (Sterman 2014, Warren et al. SD Games Online 2025). At the same time advances in protocols for group model building have enhanced the reach of causal mapping and simulation modeling as stakeholders become more closely involved in developing models and using them (Stave et al. 2024). Black boxes are becoming whiter. Now, in the 2020s, these new pathways for modeling are being paved as AI is incorporated into dynamic modeling. Now is a timely opportunity, right here in the Boston conference, to explore the new AI-related landscape. The program chairs provide a wide gateway into the territory of machine-supported modeling. The flier and materials for ISDC2025 invite submissions that describe promising new ways for modelers to deploy AI *at all stages of model development* from conceptualization to policy evaluation.


- **Opportunities and Challenges**: AI offers opportunities to enhance model insights and identify leverage points, while addressing integration challenges.

Figure 7 Flier for the 2025 International System Dynamics Conference

My career pre-dates generative AI. Nevertheless, as a seasoned modeler, I find myself observing (at a distance) this new territory of machine-supported modeling. From this perspective I wish to comment specifically on the potential use of AI in the discovery of feedback structure and the interpretation of dynamics. By ‘discovery’ I mean the task of identifying (from verbal descriptions of a problem situation and numerical estimates), feedback loops and links that could explain dysfunctional dynamics. It’s an important step that modelers view as a creative yet disciplined process to capture structure from descriptive text and from stakeholders’ mental databases.

I should declare that I’m not an expert in the use of AI. However, I am an expert in model conceptualization and simulation-based analysis of dynamics. So I’ve chosen to critique a recent and promising system dynamics paper that demonstrates the use of AI in causal mapping. The title of the paper is ‘From Text to Map: A System Dynamics Bot for Constructing Causal Loop Diagrams’ (Anonymous Manuscript 2024). The paper was presented in a plenary session of ISDC2024 in Bergen. Here is the abstract.

We introduce and test the System Dynamics Bot, a computer program leveraging a large language model to automate the creation of causal loop diagrams from textual data. To evaluate its performance, we assembled two distinct databases. The first dataset includes 20 causal loop diagrams and associated texts sourced from the system dynamics literature. The second dataset comprises responses from 30 participants to a vignette, along with causal loop diagrams coded by three system dynamics modelers. The bot uses textual data and successfully identifies approximately sixty percent of the links between variables and feedback loops in both datasets. This paper outlines our approach, provides examples, and presents evaluation results. We discuss encountered challenges and implemented solutions in developing the System Dynamics Bot. The bot can facilitate extracting mental models from textual data and improve model building processes. Moreover, the two datasets can serve as a testbed for similar programs.

**Keywords:** Causal loop diagram, generative artificial intelligence, System Dynamics Bot, ChatGPT, mental map, systems thinking

Now here are my comments on the 2024 paper.

Causal loop diagrams (CLDs) offer a convenient and compact way for modelers to visualize feedback structure in business and social systems. Although such diagrams appear to be just word-and-arrow charts it takes skill to extract relevant concepts from text or interviews and to link the concepts into closed loops of cause and effect. So it is interesting to read about a new approach for automating the generation of CLDs from text data. The authors claim the approach can help normal human modelers to develop causal maps better and faster than existing methods. That’s quite a claim when one considers the published guidelines for modelers who wish to manually construct high quality CLDs that are compatible with diagraming rules and that fit known generic structures and system archetypes.

So what can automation add? The authors concentrate on a method to automate the extraction of causal links from text data and, importantly, proceed further to identify and map closed loops that show link polarities. The capability is embedded in a Python program called SD Bot. It’s a clever approach and one that appears to be a significant advance on existing automated methods for text processing that stop short at identifying cause-and-effect pairs. However, SD Bot’s capabilities are rather different (and less) than those of a trained human modeler who is following published guidelines for CLD construction. The trained modeler scrutinizes text data to identify credible causal loops whose interlocking structure can explain what’s going on in a dynamical problem situation. The Bot is doing something different.

So it’s not really possible to test the Bot by comparing an automated CLD of a given problem situation with a well-constructed manual CLD of the same situation. That would be a strong test. The authors instead conduct two somewhat weaker tests: 1. a comparison of SD Bot’s CLDs derived from text that accompanies already published and manually created CLDs; and 2. a comparison of causal links found by human coders and by SD Bot in text descriptions, written by student volunteers, of a documented environmental situation. These comparisons are nicely designed but have limitations that I address in more detail below.

Let me begin my criticism with comments on the extraction method used in SD Bot. The method is clearly described in section 3 entitled 'Our solution: SD BOT'. The authors begin with simple text describing causal links in the number of road-crossing chickens and their eggs based on an example from Chapter 1 of Sterman's *Business Dynamics*. It's a good way to demonstrate the basic input-output extraction capabilities of SD Bot. The resulting two-loop CLD closely matches Sterman's CLD, which is re-assuring. The next example is based on Forrester's well-known market growth model and demonstrates challenges of automated CLD extraction in more complex situations. The original model contains three interlocking loops and so does SD Bot's CLD, though the visual layout is less clear. The reasoning behind the links is well documented in a display of Bot-generated text (with a caveat - the text is lengthy and cumbersome). Nevertheless it's a good start for an entirely new method. The example demonstrates, step-by-step, how causal links are being captured from plain text. However, there's a significant limitation revealed in this input-output example (The same limitation applies to all the other CLD examples used in the first data set for SD Bot evaluation).

The SD Bot is capturing CLDs from text that describes *already conceptualized feedback loops*. But this mapping capability does not mean SD Bot can capture credible feedback structure from general descriptive text about a new and entirely different problem situation. Let me be more specific. Steps 8 to 11 of SD Bot's documented reasoning describe the capacity expansion policy portrayed in the market growth model. Capacity orders are shown to be driven by delivery delay. However, in practice, many different factors can influence a firm's capital investment of which several may be mentioned in case material or interviews with managers. It took the insight of a skilled modeler to realize that delivery delay is the dominant factor in this particular firm (and importantly the full formulation for capacity ordering includes a floating goal for delivery delay, with managerial bias).

The modeler is (and has to be) discriminating in the choice of cause-and-effect links that appear in credible CLDs and simulators. Of course, lots of potential links can be found in text or in interviews but there's a vital need for judgement on which to include or exclude. This observation brings me to a critique of Bot Testing and the datasets used. Sensibly, there are two entirely different and independent datasets from two separate experiments. The dataset for experiment 1 comprises 20 texts that describe already existing CLDs found in published articles. The results replicate what was found with the market growth model. The SD Bot is able to recover a reasonably high proportion of causal links and the loops in which they are embedded. That's good, but it's still not a strong test of Bot mapping capability. Moreover, the reader knows nothing at all about the situations described by the chosen texts and ground-truth CLDs – just that there are 20 situations in total. It would be helpful to know more.

The dataset for experiment 2 comes from cause-effect links identified by 30 volunteer students who were given a text vignette describing socio-environmental factors behind the rapid decline of the water level of Lake Urmia in north-western Iran. Although the problem situation is known, the dataset again does not provide a strong test of Bot capability. One difficulty is that the vignette is very short – less than one page of text. Another is that students have no prior training in system dynamics. As a result there are virtually no feedback loops to be found in the text produced by students. This result also holds true when the students' cause-effect links are analysed by expert human coders. The Bot does a reasonable job of identifying the same causal links as human coders but it's no surprise that it can't find loops.

Despite my criticisms, the paper contains plenty of Bot design details and experimental data to interest an SD audience. There is a novel program, SD Bot, with demonstrated capability to extract and map causal links from text – sometimes revealing closed causal loops. While that capability seems limited for now there is good reason to expect it will improve as generative AI itself improves.

The authors share their own ideas for future work in the discussion section at the end of the paper. While such concluding remarks are welcome and natural, the potential benefits of automated CLD generation seem overstated and premature without stronger evidence to support them. For example, I'm not persuaded that SD Bot can make it easier for novice modelers to develop high quality causal loop diagrams, because the judgement of a trained modeler is essential for extracting dominant causal links and loops from among many that may be found in pure text data. This need for expert judgement is even greater if the modeler's purpose is to find feedback loops that explain dynamics (as it often is). Neither do I think that a comparison of AI-generated CLDs with human generated CLDs can provide a useful check of CLD fidelity, unless it comes from extra human thought prompted by discrepancies. Moreover, the CLD comparisons in the current paper cannot demonstrate this benefit because the AI-generated CLDs are derived from text descriptions of human-generated CLDs. The authors also mention benefits of using AI-generated CLDs in group model building to support live mapping of the system. Eventually maybe – but live use is surely a long way beyond the current capabilities of SD Bot.

So there's promise in the mapping capability of SD Bot, but far less than that of a skilled system dynamics modeler. What's missing is an ability to select realistic links from descriptive data that join-up as credible closed loops. Not all links are relevant or necessary to explain a given dynamic phenomenon. Modelers are discriminating. They leave out many potential links. They exercise judgement to identify influential links that underpin dominant and enduring feedback structure.

One way to sharpen loop discovery is for the modeler to realistically portray bounded rationality in organizational decision making. As I mentioned earlier, information used by decision makers is always incomplete, so the coordinating information network is sparse. This network limitation is a direct consequence of bounded rationality and is represented in many important SD models. Incidentally, it's worth noting that feedback processes can be (and often are) developed in terms of information links (influences) rather than causal links. This distinction between influence and causality is important and is frequently overlooked. A focus on influences enables modelers to be discriminating in choosing those information links that are truly dominant in decision making[4].

## Living With System Dynamics – from Crystal Blue Skies at MIT to North London and Somerset

This year, 2025, marks my 50$^{th}$ year of system dynamics. My journey began in the autumn of 1975 when, as a young graduate student, I left England with my wife Linda for an educational adventure in the USA and MIT. I still remember vividly arriving in Cambridge, Massachusetts under the crystal blue skies of New England's fall and settling into our campus apartment at Eastgate overlooking the glistening expanse of the Charles River and downtown Boston. It was just a stroll across the road to the Sloan School of Management.

Both of us were transformed by our experiences at the Institute. I became deeply immersed in system dynamics, which at the time was a new management discipline. I completed a PhD under Jay Forrester's supervision and then joined Sloan's faculty, nurturing my interests in system dynamics and strategy. Jay was a legendary MIT figure whose career spanned the Servomechanism Lab, Lincoln Laboratory and Sloan in a dazzling series of fundamental breakthroughs in computational technology,

[4] Colleagues will no doubt point to widespread use of causal mapping for loop discovery in contemporary system dynamics, particularly in participative group modeling projects. Moreover, nowadays there are well-defined guidelines for constructing good causal loop diagrams: noun-phrases for variable names, clearly labelled links and polarities, loops that follow circular or oval paths, loop polarity indicated, loops named to match a dynamical process. The SD Bot paper implicitly assumes that modelers *always* discover feedback structure using CLDs. This assumption is wrong. In fact many published CLDs have arisen from sectorized policy maps or from detailed stock and flow diagrams, and *not* in the reverse order. Such use of CLDs in articles and books provides a convenient visual summary of feedback structure rather than a basis for conceptualization.

simulation and feedback systems (Lane 2007). In 1975 his System Dynamics Group was based in building E-40, adjacent to the main Alfred P. Sloan building. E-40 was a former factory building, spartan but spacious – a vast floor-space, brightly painted bare-brick walls, remarkable inverted-mushroom support pillars, and an office-size industrial elevator with unique "inch-up" and "inch-down" buttons to fine-tune floor alignment. E40 is much more elegant now, but it had its own special charm back then. I found myself sharing a huge chunk of this colourful working space with other newly arrived doctoral students, among whom were Barry Richmond, Khalid Saeed and Peter Senge.

On completion of my PhD I joined the Sloan faculty where, for seven years, I taught and conducted research in business dynamics. I returned to England in 1986 and joined London Business School. At the time there was no presence of system dynamics at the School, so I set-about the delicate yet demanding task of transplanting this MIT-founded management discipline to a new and very different academic home in London. It became a life-long endeavour.

I was fortunate in the early years to collaborate on applied research with Group Planning at Royal Dutch/Shell which was then headed by Arie de Geus (author of *The Living Company,* HBS Press, 1997). Shell was already widely known for its path-breaking work on scenario planning - the craft of visualising alternative futures and communicating them throughout an organization. Arie realized that business simulators are a powerful way to generate and experience alternative futures and that modeling itself could be a kind of learning process to communicate the thinking behind scenarios. So a fruitful partnership developed which culminated in the publication of an edited book called *Modeling for Learning Organizations* (Productivity Press, Portland, Oregon, 1994) which I co-edited with MIT's John Sterman (who was by this time Director of the System Dynamics Group at Sloan). Arie wrote the foreword to the book. Over the next thirty years I continued research in business dynamics and, with LBS colleagues, developed a strong portfolio of system dynamics courses for MBAs, PhDs and Executives.

After three decades on the faculty of London Business School I moved house in 2016 from a commuter town near London and retired to a Somerset village in England's tranquil West Country. My wife and I now live in a charming seventeenth century farmhouse, close to the village centre. It's a big change from campus life at MIT and from the leafy suburb life of Boston and London's Home Counties. Nevertheless I retain a keen interest in system dynamics, even as the field evolves and grows from the roots I came to know so well. My own approach also evolved at London Business School, as a result of teaching and collaborative research projects with internationally known firms engaged in a variety of industries (manufacturing, oil, telecoms, retailing, consulting and broadcasting).

Below I share a synopsis of my approach to SD modeling, its character and style, for readers to compare with their own practice of SD. Whichever way you prefer, there is always an underpinning discipline and it's important to appreciate how that discipline is being re-shaped by ongoing developments in computing, AI and group modeling. I wrote the synopsis in 2021 when invited by the Cabinet Office of the UK Government to help them create a systems thinking toolkit for Civil Servants. I chose to emphasize model conceptualization, the modeling process with teams, and the analysis of simulations to achieve engagement with stakeholders[5]. The material further illustrates key character traits of SD modeling in action.

[5] Glossary of terms used in the synopsis: system dynamics, simulation, causal loop diagrams, sector maps, stock-and-flow diagrams, visual models, algebraic models, simulators, time charts, feedback structure, dynamic behavior, dynamic complexity, closed-loop feedback, team model building, equation formulation, dimensional consistency, interpretation of simulation runs, performance paradoxes, unintended consequences, alternative futures, other systems approaches.

## A Definition of System Dynamics

"System dynamics deals with how things change through time which includes most of what most people find important. It uses modeling and computer simulation to take the knowledge we already have about details in the world around us and to show why our social and physical systems behave the way they do. System dynamics demonstrates how most of our own decision making policies are the cause of the problems that we usually blame on others, and how to identify policies we can follow to improve our situation".

This succinct description was written by the founder of the field, MIT's Jay W. Forrester - affectionately known as JWF. It was his reply to an online request for 'elevator definitions' of system dynamics, circulated to members of the SD society back in 1997. In just a few sentences he captures the nub of the field in a way that's entirely consistent with my own 50 years' experience of teaching and using the approach.

My aim in this synopsis is to interpret and to unpack JWF''s view of system dynamics. I begin with observations on the process of using SD and the kind of insights that Civil Servants should expect. Then I develop a small SD model to demonstrate the nature and principles of the approach. Indeed the model is tiny – involving just six interlocking concepts – yet it nicely illustrates an intriguing example of dynamic complexity. Similar properties are also to be found in larger models of real-world problem situations. Finally I sample some model-based case studies.

## The Process of Using System Dynamics

In common with other systems approaches, system dynamics normally brings together a small project team of 3 to 5 people with a shared interest in an organizational 'problem situation'. Typically the situation can be expressed in terms of dynamics - graphs of performance over time. The team should include the client(s) for the project (owners of the problem situation), a facilitator, and an experienced SD modeler (preferably with an apprentice). A subset of these participants forms the technical model-building team and should include at least one of the clients/problem-owners.

A project can take as little as 12 hours (per participant, in separate meetings spread over 3 to 4 weeks) if the intention is to develop a purely visual model depicting key interdependencies in terms of causal loop diagrams. Considerably more time is needed to transform a visual model into a credible simulator. The project team as a whole may spend at least 20 hours per participant, again in spaced meetings. The model-building team (responsible for creating and interpreting the simulator) could spend an additional 80 hours or more per member spread over 3 to 6 months. The total time depends on the ultimate size of the fully-tested and calibrated simulation model. The reward for this extra effort is a highly engaging tool that enables policy makers to design and discover better future time paths along which to steer their organization.

There are a variety of activities to perform in a system dynamics project. There are team meetings /workshops for problem articulation and model conceptualization that help participants to define the boundary and architecture of the visual model. There is desk work and desk research by the technical team for the tasks of model formulation (drawing causal loop diagrams and/or stock-and-flow diagrams with matching algebraic equations).There is further desk work required for simulator calibration, testing and what-if experiments. Facilitation difficulty is moderate for purely visual models (sector maps and causal loop diagrams) and moderate-to-hard for stock-and-flow diagrams, algebra and credible simulators. Facilitation difficulty can be reduced by encouraging the client/problem-owners to attend a short executive-style course on system dynamics (typically 3-5 days) to gain familiarity with techniques for visualizing feedback loops and for interpreting simulation runs.

## What to Expect from a System Dynamics Study

The output of a system dynamics project will normally include sector maps and causal loop diagrams. They provide team members a clear and communicable visual representation of the interdependencies that shape the dynamics and performance of an enterprise. The terminology in all such diagrams should be readily familiar to participants as they themselves will have provided the key concepts. Further work to develop a simulator will likewise yield compelling visual representations with the advantage that the figures depict feedback structures which are strictly consistent with simulations of alternative futures.

The outcome of a purely visual modeling project is a broad systems view of the enterprise that can enhance participants' understanding of the problem situation and may lead to the identification of beneficial systemic interventions. The outcome of a full-blown modeling project (with simulator) is once again a broad systems view of the enterprise along with credible narratives about unfolding dynamics. These narratives or 'stories', written by the modelers, are normally very persuasive as they are based on simulated time charts and therefore reliably show the unfolding consequences of interlocking factors and assumptions in the visual model.

What new insights can system dynamics studies provide? They are good for probing situations of puzzling and problematic performance through time, arising from covert and interlocking dependencies in all manner of enterprises – both public sector and private. Models can be pitched at the level of a single organization, a set of rival organizations, an industry sector, a community or even a whole society. The resulting simulators are particularly helpful for improving coordination among disparate operations, thereby boosting performance and reducing hidden dysfunctional conflicts.

System dynamics is a distinct modeling discipline within the social sciences. Yet it overlaps with other systems approaches in seeking constructive dialogue and accommodation among interdependent stakeholders (in their quest for better outcomes from policy and strategy)[6]. It is especially appropriate for the diagnosis of performance paradoxes - when strategies fail to play-out as intended or when the consequences of policies and strategies are uncertain (demanding debate, discussion and evaluation). System dynamics can help policymakers and policy advisers to rehearse alternative time paths into the future. A note of caution however. The approach should <u>not</u> be used to model 'the system with all its parts' in the absence of dynamics from a real-world situation or case study. Without the anchor of credible dynamics there is insufficient evidence to guide the critical step of defining the model boundary – which factors to include and which, among a huge multitude, can be safely excluded.

## The Nature and Principles of System Dynamics – Unpacking JWF's View

"System dynamics deals with how things change over time in social and business enterprises". A commonly-used phrase captures this time-oriented and inertial perspective: <u>feedback structure determines dynamic behavior</u>. It is the bedrock on which the approach builds. The idea that there's an enduring structure to business and social systems, which somehow predetermines achievable futures, is not necessarily obvious. Some people argue that personalities, chance, ambition, fate and unforeseen events hold the keys to the future in an uncertain world. However, what appears to be chance may, from a different perspective, have a systemic cause. For example, when driving on a busy highway you may experience sporadic stops and starts. Does that mean you are at the mercy of random events like breakdowns or accidents? Not necessarily. Enclosed in a car at ground level you don't see the waves of traffic that arise from the collective actions of individual drivers as they try to

---

[6] Much has been written about the relationship between system dynamics and other popular systems methods. The interested reader is referred to *Systems Approaches to Making Change: A Practical Guide* (Reynolds and Holwell editors, 2020), particularly Chapter 1 'Introducing Systems Approaches' and Chapter 7 'Epilogue: Systems Approaches and Systems Practice'.

maintain a steady speed while also keeping a safe distance from the car in front. There is an invisible structure to the 'system' of driving on a crowded motorway that causes sporadic stops and starts, without the need for accidents (though of course they do happen too).

The structure of social and business systems lies in the interdependencies that bind disparate parts of any collaborative enterprise. Often these links are invisible within the flux of operations. It is the job of the project team to tease out the factors that bind the enterprise, thereby enabling - yet at the same time constraining - achievable outcomes and goals. JWF describes how. "System dynamics uses modeling and computer simulation to take the knowledge we already have about details in the world around us and to show why our social and physical systems behave the way they do".

Romeo and Juliet

To illustrate the important idea that structure determines dynamic behavior I develop a tiny model that depicts interdependency in a love relationship between Romeo and Juliet. Admittedly this is a fanciful example in which a conjecture about varying love in a relationship substitutes for time series data found in real-world cases.  Nevertheless it demonstrates, clearly and concisely, the structural source of dynamics. The example also conveys some basic SD modeling principles and shows how to interpret simulation runs in a non-technical yet rigorous way, which fully explains the shape of simulated trajectories arising from the model's closed-loop feedback structure.

Figure 8 shows a presumed love relationship between Romeo and Juliet in three different yet compatible ways. On the left is the basic situation. There is a mutual dependency between Romeo and Juliet. The diagrams to the right depict this dependency in more specific detail.

In the center is a causal loop diagram, essentially a word-and-arrow chart whose composition follows prescribed rules for the choice of phrases and the assignment of link polarities (+ and -). Romeo's love for Juliet influences Juliet's love for Romeo, and vice-versa. When Romeo's love increases he induces greater love from Juliet (hence the + sign). Conversely when Juliet's love increases she induces disdain from Romeo and his love for her falls (hence the – sign). This pair of links forms a closed feedback loop.

On the right of Figure 8 is a stock-and-flow diagram that represents the love relationship in a form that can be converted into an algebraic simulation model that replicates *exactly* the terminology and connections in the diagram. There are two stocks of love whose gradual filling-up or depletion corresponds to variations in Romeo and Juliet's love over time. Naturally these processes of accumulation/ depletion require a specific representation of the rates of flow that impinge on love: the *change in Romeo's love for Juliet* and the *Change in Juliet's love for Romeo*. The important distinction between a stock and flow is apparent here. Love builds or withers and so a time-dependent story of love can unfold. Changes in love (whether Romeo's or Juliet's) are naturally measured in units of love per month which accumulate over time into units of love. The rate of accumulation depends also on Juliet's responsiveness to love and Romeo's propensity to spurn love (factors that also ensure dimensional consistency of the concepts in the stock and flow network). Love springs eternal from the heart, depicted as the pool or cloud from which new love flows or to which spurned love returns.

Models of social and business systems similarly contain accumulating/depleting asset stocks from which the enterprise creates its service and products. These stocks can be tangible (such as staff, equipment, inventory, cash, customers, users and premises) or intangible (like quality, reputation or even love).

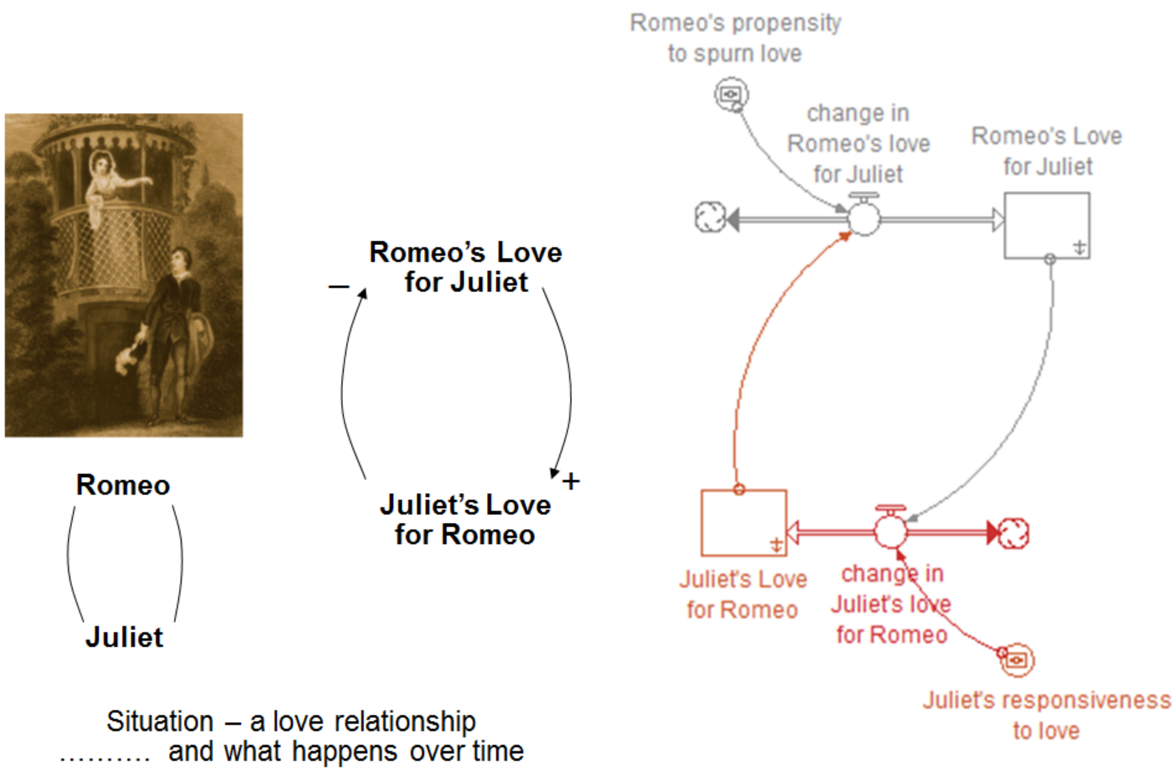


Figure 8: Interdependency of Romeo and Juliet

"System dynamics demonstrates how most of our own decision making policies are the cause of the problems that we usually blame on others, and how to identify policies we can follow to improve our situation". Here JWF re-states the core idea in system dynamics that problems we experience in organizations often stem from known interdependencies - interlocking operating policies and relationships that interact in surprising ways as time unfolds. Simulation reveals how.

Figure 9 shows both the simulated dynamics and the feedback structure in Romeo and Juliet's love relationship.  Here I've expanded the causal loop diagram to show not only love but also the rates of change of love. This extra visual detail (transposed from the stock-and flow-diagram) enables the modeler to write a clear, yet rigorous, narrative about dynamics.

It turns out that the lovers' presumed entwinement, combining love and disdain, is sufficient to produce an endless cycle of waxing and waning love. This outcome is a big surprise to many people and is a powerful example that feedback structure shapes and pre-determines dynamic behavior. The model contains a closed feedback loop that weaves its way between and among the stocks and flows of love. If this loop actually exists then there is no need for any external influence to drive the dynamics of love. The tides of love are self-generating. Similar endogenous insights are to be found in the feedback structures that underpin dynamic behavior in society and business.

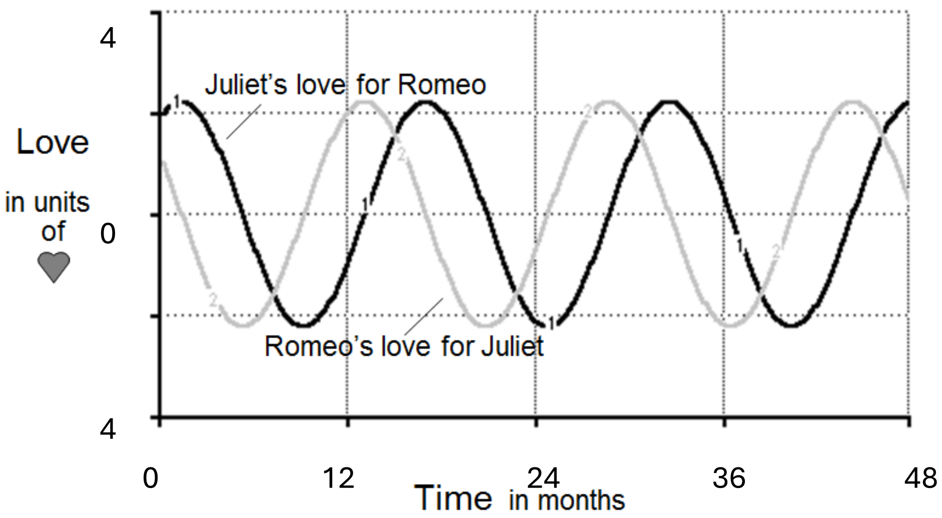


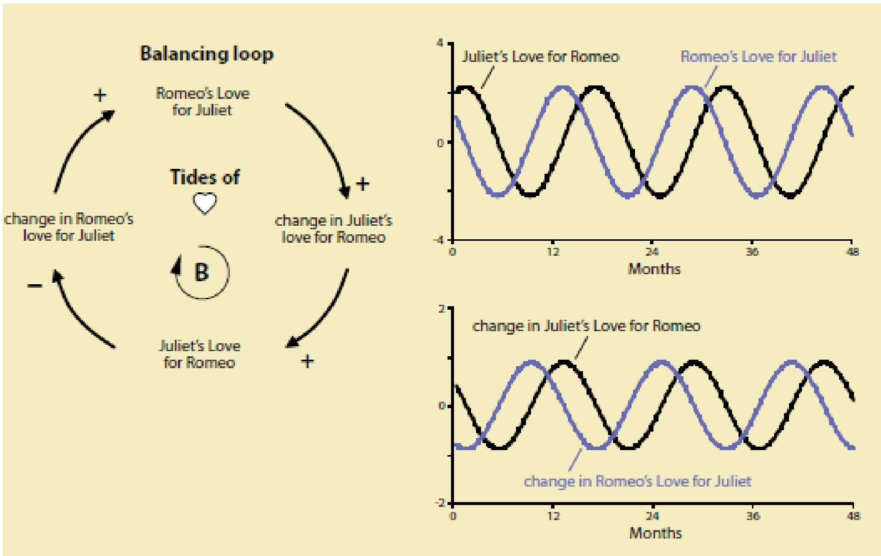


Figure 9: Structure and Simulated Dynamics in Romeo and Juliet's Relationship

The simulation runs for 48 months and is shown in two separate but related time charts. The top chart shows stocks of love and the bottom chart shows changes in love. At the start of the simulation, *Juliet's Love for Romeo* is set at two units of love while *Romeo's Love for Juliet* is set at one unit. Romeo's affection causes Juliet's love to grow while Romeo's love falls since he spurns affection. After one month Juliet's love reaches a peak of slightly more than two units. The reason for this peak can be seen in the bottom chart where the *change in Juliet's Love for Romeo* has fallen to zero, meaning that there is no longer cause for her love to change. She is in a period of contented and stable love. But this state does not and cannot last. By month 1, Romeo's love has fallen to zero and the rate of change of his love is at a minimum of minus one units of love per month. He does not share Juliet's stable contentment and his love for Juliet continues to fall, reaching a low of slightly more than minus two units of love by month 5. Of course Juliet notices this marked decline in affection over a period of three months. Her love for Romeo gradually recedes and reaches zero by month 5. By now she is quite upset with Romeo and the rate of change of her love reaches a low of minus one units of love per month. Inevitably, this further deterioration in their relationship causes her love to fall still further, becoming negative in value and reaching a minimum of slightly more than minus two by month 9. Meanwhile, Romeo (the contrarian) feels freed from unwelcome affection and his disdain for Juliet begins to lessen, so that, by month 9, his love returns to zero and is growing at its fastest rate just as Juliet's love reaches its nadir.

Further shifts in affection are already in train. *Romeo's Love for Juliet* continues to rise for four months, reaching a peak of slightly more than two units by month 13. This surge of affection induces new romantic feeling in Juliet. Her love climbs from its negative depths to a calm neutrality of zero by month 13. But this seeming neutrality does not last as her love continues to grow, spurred on by Romeo's evident and sustained affection. Just before month 16, Romeo and Juliet find themselves back in exactly the same romantic state they were at the start of the simulation. *Juliet's Love for Romeo* is at two units of love and rising, while *Romeo's Love for Juliet* is at one unit and falling. The stage is set for another identical cycle of love. The tiny model of "Romeo and Juliet," when developed and completed, presents a good example of closed-loop feedback structure that pre-determines dynamic behavior. There is dynamic complexity even in small closed-loop models.

Note the style of simulation analysis: plain English, non-technical narrative that is nevertheless dynamically rigorous. Similar clarity is possible, and should be expected, from SD models of social and business systems.

The algebraic equations for the Romeo and Juliet simulator are shown in Figure 10. Note the matching terminology and concepts in the stock-and flow diagrams and the equation listing. The diagram provides a client-friendly visual bridge to the algebra and a helpful framework for the modeling team to develop equation formulations consistent with the interdependencies identified by the project team.

```
Romeo's_love_for_Juliet(t) = Romeo's_love_for_Juliet(t - dt) +
(change_in_Romeo's_love_for_Juliet) * dt
INIT Romeo's_love_for_Juliet = 1  {unit of love, on a scale from 4 to minus 4}

change_in_Romeo's_love_for_Juliet = Juliet's_love_for_Romeo *
Romeo's_propensity_to_spurn_love  {units of love per month}

Juliet's_love_for_Romeo(t) = Juliet's_love_for_Romeo(t - dt) +
(change_in_Juliet's_love_for_Romeo) * dt
INIT Juliet's_love_for_Romeo = 2  {units of love, on a scale from 4 to minus 4}

change_in_Juliet's_love_for_Romeo = Romeo's_love_for_Juliet *
Juliet's_responsiveness_to_love  {units of love per month}

Juliet's_responsiveness_to_love = 0.4  {fraction per month}
Romeo's_propensity_to_spurn_love = - 0.4  {fraction per month}
```

Figure 10: Equations for the Romeo and Juliet Simulator

A Sample of Model-Based Policy Studies

Imagine a spectrum of model fidelity (in terms of detail and realism) spanning a range from small and metaphorical (5 to 10 concepts) to large and detailed (hundreds or even thousands of concepts). Between these two extremes lie illustrative models containing 10 to 100 concepts. Over the years system dynamics studies have led to models and simulators that cover the entire range. Here I will sample some of these model-based studies. Illustrative models have proven effective and sufficient to address a wide variety of issues and problem situations. Sometimes metaphorical models enable client engagement. One could say that 'small is beautiful' in the world of policy and strategy modeling, though there are plenty of studies that have resulted in models containing one thousand or more concepts (often carefully calibrated too).

Many system dynamics models have been built to address public policy issues. The topics read like a list of political speeches: health care reform, prison overcrowding, drug-related crime, transportation, urban renewal, environmental policy and fisheries regulation. An early and highly influential model, dating from the early 1970s was JWF's *World Dynamics* (Forrester 1971) and the closely-related Club of Rome study *Limits to Growth* (Meadows et al 1972)[7]. JWF's World 2 model contains just 5 stock accumulations to portray the interlocking factors behind global industrial growth: population, industrial capital, natural resources, capital in agriculture and pollution (a covert yet potent accumulation). It's a bold sketch on a compact canvas, bordering on metaphorical[8].

More recently there have been in-depth studies of climate change with accompanying models and simulators (Sterman et al 2012)[9]. The Covid-19 pandemic has spawned several insightful SD models about disease transmission and policies for mitigation (see for example Struben 2020). These models are of intermediate size, less than 100 concepts, to portray aggregate stock accumulations of susceptible, infected and recovered people whose interactions govern the epidemiological spread of the infection. They are illustrative models of sufficient fidelity to advise politicians and policy advisers as well as inform the general public[10]. Likewise the transmission of norovirus in the UK has been modeled in a prize-winning SD study conducted by Lane et al (2019). The study also nicely illustrates the art and craft of team model building involving a combination of well-grounded causal loop diagrams, stock-and-flow diagrams and a carefully calibrated simulator. In chapter 9 of my 2015 SD textbook *Strategic Modelling and Business Dynamics* there is a section devoted to a model-based study of medical workforce dynamics and patient care during the transition to new work patterns for junior doctors (prescribed by the European Working Time Directive EUWTD). The project was conducted by Dr Mark Ratnarajah, a pediatric specialist registrar based in London, who, at the time was enrolled on the Executive MBA program at London Business School. Here is an example where the client was also the modeler – a situation that usually leads to a well-grounded model and thought-provoking simulations (prize-winning in this particular case). Chapter 9 also contains references to healthcare modeling by Hirsch et al 2015, Wolstenholme 2006 and others.

Of course, given the industrial origins of system dynamics, there are many studies and models of real-world business situations in a broad range of firms and industries from oil-and-gas to high-tech start-ups, R&D ventures, motorcycles, electric vehicles, fast-moving consumer goods, pharmaceuticals, manufacturing supply chains and more. Examples can be found in the books and edited collections listed at the end of the paper.

To appreciate the positioning and use of system dynamics within the context of other systems and OR methods I strongly recommend *Systems Approaches to Making Change* (Reynolds and Holwell 2020) and *System Dynamics – Soft and Hard Operational Research (*Kunc 2018) . System dynamics is moving with the times, adopting and adapting experimental and computational methods from related branches of the social sciences. These developments are reviewed in Sterman 2018 'System Dynamics at Sixty: The Path Forward' Finally, readers who wish to know more about the upbringing, life,

---

[7] NOTE: For the sake of brevity I will not here provide an exhaustive list of original references. Instead see the list of further readings at the end of the paper. The list includes contemporary textbooks that contain more details of these important studies along with a full set of references and original source materials.

[8] Meadows et al's World 3 model contains 5 sectors that portray, in more detail, the same interlocking factors that JWF identified in World 2. Both models were capable of generating plausible scenarios of societal growth followed by surprise collapse. They could also delineate policy changes to enable a smooth transition from growth to a sustainable industrial society. The models remain relevant and influential today - though, at the time, they were deemed controversial and overly simplistic by some academics.

[9] The research team has developed detailed and carefully calibrated models of numerous vital, interlocking industrial sectors that are responsible for generating climate-modifying carbon emissions. These models have been further crafted to yield interactive web-based simulators (described in Sterman 2014) that are being used in COP climate negotiations and in public education about climate policy. The models contain thousands of concepts.

[10] In addition to aggregate SD models there are also numerous agent-based simulators that capture realistic and detailed spatial interactions among populations of susceptible and infected people. These models may contain thousands of variables capable of drilling-down to the spread of infection at the level of regions, local communities and public events.

academic career, achievements and aspirations of JWF should read Fisher 2005 'The Prophet of Unintended Consequences' and/or Lane 2007 'The Power of the Bond Between Cause and Effect'.

## Where to Next?

Many conference delegates may be unfamiliar with the 1960's folk music scene. That's all changed this year, 2025, with James Mangold's acclaimed movie 'A Complete Unknown'. Timothee Chalamet as Bob Dylan vividly rekindles that special 60's era in a compelling time capsule of Dylan's music, and turbulent life in the company of other legendary folk singers (Woodie Guthrie, Pete Seeger, Joan Baez and Sylvie Russo) who nurtured his prolific singer-songwriter talent. Whether or not you like Dylan's performances there is surely something special and enduring about his lyrics.

My conference paper was inspired by the movie and by the poetry of 'The Times They Are A-Changin'.There's relevance in both these sources to ISDC2025: revisiting creative power (of Dylan's early acoustic folk); momentous and controversial transition (to electric instruments), and the lyrics of the song itself that emphasize the inevitability of change and the need for individuals – writers, politicians, and the public – to adapt or risk obsolescence.

Come gather 'round people wherever you roam
And admit that the waters around you have grown
And accept it that soon you'll be drenched to the bone
If your time to you is worth savin'
And you better start swimming or you'll sink like a stone
For the times they are a-changin'

….. jumping ahead

Come mothers and fathers throughout the land
And don't criticize what you can't understand
Your sons and your daughters are beyond your command
Your old road is rapidly agin'
Please get out of the new one if you can't lend a hand
For the times they are a-changin

The line it is drawn, the curse it is cast
The slow one now will later be fast
As the present now will later be past
The order is rapidly fadin'
And the first one now will later be last
For the times they are a-changin'

Of course there's value in the past traditions of system dynamics. It's not uncommon in contemporary SD discourse to hear skilled modelers recommend that learners study JWF's original ground-breaking book *Industrial Dynamics* (Forrester 1961, Lane 2022). More recent is *Business Dynamics* (Sterman 2000) which combines the territory covered in *Industrial Dynamics* with important new ground from the 1980's and 1990's. It is the most widely used and influential text book in the field. Both JWF and Sterman were writing in the pre-AI era. The waters around their books may have grown but they're not yet drenched to the bone.

The same applies to my own book *Strategic Modelling and Business Dynamics* (Morecroft 2015), originally published in 2007, with a second edition SMBD2e in 2015. I'd like to say a few words about SMBD2e and its website materials for learners and instructors. The book illustrates all the principles for modeling style mentioned in this paper and connects them with underpinning SD literature. As I state in the preface to the second edition "I want the book to be an enduring bridge from traditional to contemporary system dynamics" – a way to cross the growing computational waters, without losing sight of past territory. Each of the book's chapters lays strong emphasis on clear visualization and documentation of real-world feedback structure, backed-up by rigorous yet easy-to read equation formulations. Understanding of dynamics comes from careful narrative interpretation of dynamics. The same style applies to models of fishery dynamics found in Chapters 1 and 9, to manufacturing/supply-chain dynamics in Chapter 5, and to business growth dynamics in Chapters 6 and 7. Chapter 8's one-hundred equation model of oil price dynamics and the global oil producers shows that the style guidelines also work well with larger models. Additional electronic topics on the learners' website cover the dynamics of firms' diversification strategy - represented in behavioral SD models and interpreted using gaming simulators.

To reinforce the book's guidelines I provide, chapter-by-chapter on the Learners' website, selected articles from the working paper archives of the MIT System Dynamics Group (the so-called D-memo series) covering the 1960s, 70s and 80s. On the Instructors' website I provide annotated and graded solutions to course assignments showing exemplary work by students who followed the style guidelines[11].

From this foundation of rigor-with-accessibility it is then possible for learners to bridge securely to complementary methods (established and new) whatever form they might take. For example Chapters 6 and 7 connect well with contemporary behavioral and resource-based views of the firm. These ideas, from modern economics and strategy, fit neatly with asset stock accumulation and the information feedback view of the firm found in traditional system dynamics. In addition, and scattered throughout the book there are references to analytical methods for dynamic models, precursors of AI's uses in dynamic modeling. And finally, for model conceptualization, there is mention of tried-and-tested protocols for group model building.

The old road is indeed agin'. System dynamics from the pre-millennium (the era I know best), is already past, the order is gradually fadin'. But bear in mind too, in the excitement of the present, that the first ones now will (themselves) later be last, for the times they are a-changin'.

Quoting more of Dylan – "I won't speak too soon for the wheel's still in spin, and I won't criticize what I don't understand". Well, I admit there's been gentle criticism of SD Bot. But that stems from a desire to see continued effort devoted to machine support of model conceptualization and analysis. Augmented creativity. Help in setting a model's boundary and identifying key decision making processes. Help to accurately portray the sparse information feedback network that characterizes decisionmakers' bounded rationality, thereby realistically capturing societal structure and architecture. And assistance with the interpretation of simulations. Suggestions for clear and rigorous phrasing which can help modelers explain simulated trajectories. These added capabilities for conceptualization, formulation and analysis go beyond causal loop mapping. They would sharpen the discovery of links-that-matter and help modelers to identify enduring, dominant feedback loops that shape unfolding futures in business and society.

The times they are a-changin'. Over to you!!

[11] In 2018 I added to the Instructors' site a complete set of 30 video lectures based on SMBD as delivered at London Business School and for an online course at WPI. There are sufficient lectures and supporting materials for a full-semester.

## Further Reading

Books and Edited Collections (Note: full references to model-based policy studies mentioned on pages 19 and 20 can be found within these books and edited collections.)

Ford, A. (2010). *Modeling the Environment*, 2nd edition, Washington: Island Press.

Forrester, J.W. (1961). *Industrial Dynamics*, available from the System Dynamics Society www.systemdynamics.org; originally published by MIT Press 1961.

Hovmand, P.S. (2014). *Community Based System Dynamics*, New York: Springer Science + Business Media.

Kunc, M. editor (2018). *System Dynamics – Soft and Hard Operational Research*, in the OR Essentials Series, London: Palgrave-Macmillan.

Morecroft, J.D.W. (2015). *Strategic Modelling and Business Dynamics*, 2nd edition, Chichester UK: Wiley.

Morecroft, J.D.W. and Sterman J.D. editors (1994). *Modeling for Learning Organizations*, Portland, Oregon: Productivity Press.

Rahmandad, H. Oliva, R. and Osgood, N.D. editors (2015). *Analytical Methods for Dynamic Modelers*, Cambridge MA: MIT Press.

Reynolds, M. and Holwell, S. editors (2020). *Systems Approaches to Making Change: A Practical Guide*, 2nd edition,London: Springer-Verlag in collaboration with the Open University.

Richmond, J. Stuntz, L. Richmond, K. and Egner, J. (2010). *Tracing Connections: Voices of Systems Thinkers*, Lebanon NH: isee systems inc. in collaboration with The Creative Learning Exchange.

Stave, K., Zimmermann N., and Kim H. editors (2024). Qualitative Aspects of System Dynamics Modeling, Special Issue of the *System Dynamics Review*, 40, e1793. (Note: The literature on qualitative system dynamics is substantial and growing. Qualitative SD studies follow different style guidelines than those covered here in my Reflections paper. Readers can find all they need to know about qualitative SD, both contemporary and historical, in Stave, Zimmermann and Kim's edited collection. As a starter read the editors' 18 page Introduction).

Sterman, J.D. (2000). *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Boston, MA: Irwin McGraw Hill.

Sterman, J.D. and Repenning, N.P. editors (2018). Celebrating the 60th Anniversary of the System Dynamics Field, Special Issue of the *System Dynamics Review*, 34 (1-2).

Contemporary and Historical Journal Articles and Working Papers

Anonymous Manuscript (2024). From Text to Map: A System Dynamics Bot for Constructing Causal Loop Diagrams, Manuscript R1059 submitted to (and accepted for) the 2024 International System Dynamics Conference in Bergen. (Note: A copy of R1059 is available as supplementary material to accompany to this Reflections paper. I have deliberately chosen the anonymised manuscript as it was this version of the document I originally read and on which my comments are based).

Anonymous Working Paper. (1962). Introduction to the Industrial Dynamics Study of Corporate Growth, MIT System Dynamics Group Working Paper, D-433, available online from the System Dynamics Society www.systemdynamics.org in the MIT System Dynamics Group Literature Collection. (Note: The paper was in fact written by David W. Packer who collaborated closely with Jay Forrester on the important topic of corporate growth in early pioneering work of the MIT SD Group).

Fisher, L. (2005). The prophet of unintended consequences. *strategy+business*, 40, 1–12.

Forrester, J.W. (1968). Market growth as influenced by capital investment, *Industrial Management Review*, 7(1), 5-17. Journal retitled the *Sloan Management* Review in 1986. Re-published in *Collected Papers of Jay W. Forrester*, Chapter 7, 111-132, Cambridge MA: Wright-Allen Press, 1975. Also available online from the System Dynamics Society www.systemdynamics.org in Eberlein R., Richardson G.P. et al. editors *Social Dynamics*, Chapter 5 – Corporate Applications, article 2.

Forrester, J.W. (1994). Policies, decisions and information sources for modeling. In Morecroft, J.D.W. and Sterman J.D. editors (1994). *Modeling for Learning Organizations*, Portland, Oregon: Productivity Press.

Lane, D.C. (2007). The Power of the Bond Between Cause and Effect: Jay Forrester and the Field of System Dynamics, *System Dynamics Review*, 23 (2-3), 95-118.

Lane, D.C. Husemann, E, Holland, D. and Khaled, A. (2019). Understanding Foodborne Transmission Mechanisms for Norovirus: A Study for the UK's Food Standards Agency, *European Journal of Operational Research*, 275 (2), 721-736. Winner of the System Dynamics Application Award 2020.

Lane, D.C. (2022). Fons et origo: reflections on the 60th anniversary of *Industrial Dynamics*, *System Dynamics Review*, 38 (3), 292-324.

Mass, Nathaniel J. and Senge, Peter M. (1975). Understanding Oscillations in Simple Systems. D-2045-2, 1975. Available online from the System Dynamics Society www.systemdynamics.org in the MIT System Dynamics Group Literature Collection.

Sterman, J. D. (2014). Interactive Web-Based Simulations for Strategy and Sustainability; The MIT Sloan *Learning Edge* Management Flight Simulators, Part 1. *System Dynamics Review,* 30 (1–2), 89–121.

Sterman, J.D. (2018). System Dynamics at Sixty: The Path Forward, *System Dynamics Review*, 34 (1-2), 5-47.

Struben, J. (2020). The Coronavirus Disease (COVID-19) Pandemic: Simulation-based Assessment of Outbreak Responses and Postpeak Strategies, *System Dynamics Review*, 36 (3), 247-293.

Warren, K. et al. (2025). SD Games Online, website for serious games about business strategy at sdgamesonline.com, accessed March 2025.