

Systems Thinking and Business Process Redesign: An Application to the Beer Game

ANN VAN ACKERE, *Associate Professor of Decision Science, London Business School*;
ERIK REIMER LARSEN, *Co-ordinator, Learning Centre Resource Unit, London Business School*;
JOHN D.W. MORECROFT, *Associate Professor of Strategic Management, London Business School*

Managers are regrettably ignorant of the fact that their business organisations are 'designable'. But recently, concepts such as business re-engineering and systems thinking, coupled with advances in methods of quantifying business systems, have enabled managers to scrutinise their business systems afresh.

Ann van Ackere, Erik Reimer Larsen and John Morecroft use a well-known logistical system — the 'beer game' — to illustrate these re-engineering concepts and tools in a multi-stage production and distribution system involving a single brand of

beer. This business game raises the fundamental question of why it is so difficult to match shipments and factory production to consumer demand.

The authors conclude that such re-design concepts and tools can be applied successfully to full-scale business problems. Systems thinking, modelling and continuous time simulation can provide the framework for carrying the design process from mapping all the way through to redesign. The most effective CEOs of the future will be those who are competent to create corporate design in which employees are allowed to succeed.



Introduction

We are all used to the idea that automobiles, ships, aircraft, office buildings and bridges need careful design to achieve their purpose. But there is much less awareness that business organisations too are 'designable'. This blindspot about business design is all the more surprising when one considers the human and financial costs of a malfunctioning business.

Business design has been overlooked because in the past we have lacked adequate concepts and tools to represent the structure of business systems — their component parts and how they are assembled. Without representational tools we have no way to scrutinise how business operations interlock and interact. Neither can we anticipate whether a change of policy or procedure is likely to lead to improved performance, except by trial and error. There are no blueprints, no prototypes, no wind-tunnel models, to help us envisage the enterprise, pinpoint its deficiencies or experiment with alternatives.

However, in recent years design has come to the top of management's agenda with the emergence of concepts such as business re-engineering (Hammer, 1990) and systems thinking (Senge, 1990) coupled to advances in methods of mapping, modelling and simulating business systems. (Morecroft and Sterman, 1992). These approaches encourage managers to take off their functional blinkers and step back to see the coordinating mechanisms necessary to satisfy customers. In order to re-engineer or redesign an area of business, managers must work as a team, share knowledge, map out the key processes, link them across functional boundaries and then see where there is room for improvement.

In this paper we take readers inside a redesign problem to show re-engineering concepts and tools in action. We have selected for analysis a classic and well-known logistical system — a multi-stage production and distribution system. Known as the 'beer game', it is a system for producing and distributing a single brand of beer (Senge, 1990, Chapter 3). The paper first provides the background to the game and shows typical results from a session with experienced managers. The game raises a fundamental design question — why is it so difficult to match shipments and factory production to customer demand? To answer this question we scrutinise the logistical system by identifying and mapping the sequence of business processes that connect the customer to the factory. This scrutiny is typical of the discipline imposed by a redesign exercise. Once the processes are understood then it is natural to ask how they can be changed or re-configured to improve performance. The final stage of the analysis shows how modelling and simulation are used to sharpen thinking about the consequences of redesign.

The Beer Game

The production-distribution game or the beer game, as it is often called, has been used in management

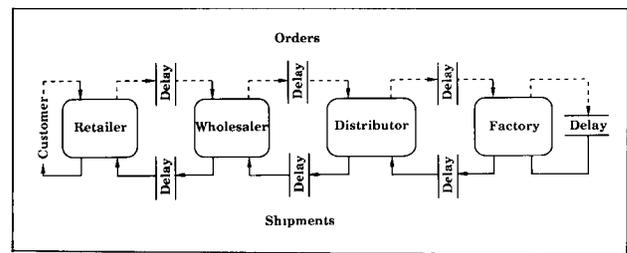


Figure 1 The Structure of the Beer Game

education and development for more than 30 years. The beer game is a role playing game where the participants have to minimise costs by managing inventory levels in a production-distribution chain. Thousands of people from undergraduates to top managers in multinational companies have played it. Most of them have discovered that what at the outset looks like a simple task is indeed almost impossible to accomplish.

The beer game consists of 4 sectors: retailer, wholesaler, distributor and factory. Figure 1 shows the basic structure of the game, and the flow of orders and cases of beer. There is an external customer whose demand is determined in advance, although the players do not know the demand pattern. Each sector has an initial small buffer inventory of 12 cases. All the retailer has to do is to fill the orders he receives from the customer, and then decide how much he wants to order from the wholesaler. The wholesaler has to fill the orders he receives from the retailer, and decide how much to order from the distributor. The distributor similarly ships beer to the wholesaler and places orders with the factory. The factory delivers to the distributor, and then decides how much to produce. The production time for beer is two weeks.

The orders from the customer are represented as a stack of cards turned upside down. Each week the retailer takes the top card which represents the customer's demand for that week. As in real life, not everything can happen at the same time, so there are mail and shipping delays built in: it takes two weeks to mail an order and two weeks to ship the requested amount of beer from one sector to the next. This delay-structure exists between all the sectors. Finally, it is *not* possible to cancel orders.

If a sector is unable to deliver the requested amount of beer to the sector downstream, the remainder of the order goes into the backlog and will be delivered when the sector receives beer from its supplier, which may take one or more weeks. It is assumed that it is more costly to end up in a stock-out situation than in a situation with a surplus of inventory.

Inventory carrying costs are £0.5 per case of beer per week. Stock-out costs, associated with the possibility of losing customers, provide an incentive to hold some inventory. In the game the stock-out costs are £1.0 per case of beer per week.

The system described here is a typical cascaded production-distribution system, where each sector has its own small buffer stock. The idea is that a structure like this should protect the factory against random fluctuations in final consumption. Only long term movements in demand should propagate along the chain and reach the factory. This type of distribution can be found in many industries ranging from automobiles to leather shoes.

When the game is played, each sector is controlled by one or two people, who have to make the decision described above. In principle, the players of different sectors are not allowed to communicate with each other. This means that the retailer is the only one who knows customer demand. The wholesaler can try to estimate it from the orders he receives from the retailer, the distributor from the orders he receives from the wholesaler etc. Customer demand is simple: initially 4 cases per week, in week 5 a step increase from 4 to 8 cases of beer and then constant at 8 cases for the rest of the game. Managing the system might sound a relatively easy task but it is, in fact, almost impossible. Figure 2 shows some typical results from a game with

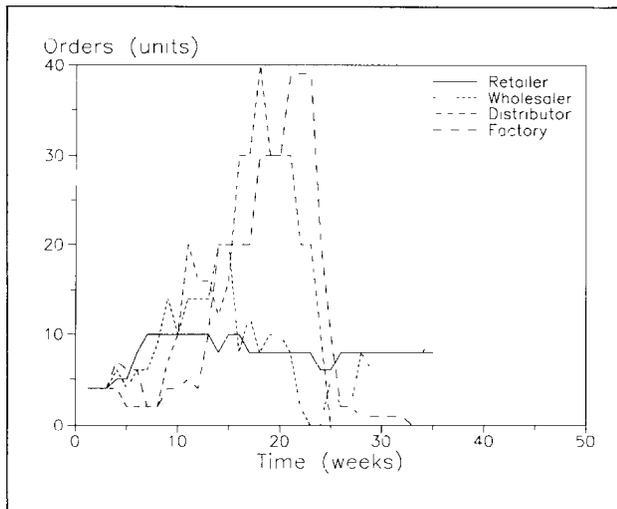


Figure 2a Orders for Beer: Experimental Data

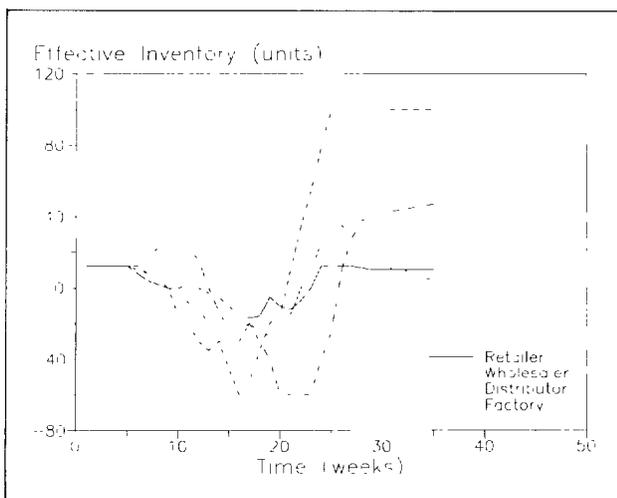


Figure 2b Effective Inventory: Experimental Data

experienced managers. Effective inventory is defined as (inventory — backlog), i.e. if the effective inventory is negative there is a backlog.

Figures 2a and b could be described in the following way: shortly after the step increase in customer orders the retailer realises that his inventory is falling and consequently starts to increase his orders to the wholesaler. As the demand from the retailer increases, the wholesaler's inventory will rapidly disappear so the retailer does not get the increased amount of beer he has ordered. From his point of view nothing is happening — he has increased his orders but receives only a fraction of what he ordered, if anything. Faced first with a rapidly falling inventory, and then an increasing backlog, he increases his orders even further as panic sets in. It is likely that he is not thinking of what he has ordered but not yet received (the supply line) or that he at least does not take this fully into account. The wholesaler goes through the same experience, but slightly worse because the retailer's panic is causing the retailer to order more than necessary. The distributor is even worse off, faced with the consequences of both the panic of the retailer and the wholesaler. As the factory finally discovers this explosively growing demand and increases production, the first 6 weeks have gone. As the factory finally increases production, all sectors discover that they have been ordering too much, and everyone's orders drop to zero. This leaves all sectors with an enormous inventory, sufficient to cover demand for weeks or months ahead.

As part of the debriefing of the game we will normally discuss how a system like this can be improved. For a long time this discussion was based on various beliefs about what changes in the structure of the game and in player behaviour might produce in terms of better or worse results. Participants were in fact trying to redesign the process by which customer demand is satisfied, but were unable to test their suggestions. Using a simulation model of the game, it is possible to test more rigorously various suggestions for how to improve the performance of the beer production-distribution system. In the next section we discuss what business process redesign is and how it has been used, before applying it to the beer game.

Business Process Redesign

The terms Business Process Redesign (BPR) or re-engineering are the buzz words of the nineties. But what do we really mean by them? To understand these concepts, it is useful to go back to Jenkin's (1971) definition of systems engineering: 'The science of designing complex systems in their totality to ensure that the component subsystems making up the system are designed, fitted together, checked and operated in the most efficient way'. As discussed in more detail by Jackson (1991), the basic idea is that the engineering approach aiming at optimising the use of resources is applicable more generally to systems made up of the interaction of many components, whether these systems

are hardware systems, departments, firms or governments.

How does systems engineering relate to business processes? Davenport and Short (1990) define a business process as 'a set of logically related tasks performed to achieve a defined business outcome' and a business system as 'a set of processes', which represent how an entity performs its business. Emphasising the role of information technology (IT) in BPR, they go on to define the interaction between IT and BPR as the new industrial engineering. When defining a business process, we prefer to focus on the underlying policy function. The policy function describes the decision process or motivation underlying a specific action. This policy function, together with its consequences represents a business process.

The beer game can be interpreted as a system, consisting of various processes. Consider for instance the process 'ordering by the distributor'. This process consists of the following steps:

- gather available information (inventory or backlog, supply line, order from wholesaler, expected future orders);
- decide how much to order;
- place card with order in appropriate box.

The key step in this process is 'decide how much to order', which hides a quite complex policy issue: how does one go from those various bits of information to a decision on the amount to be ordered? It is this decision process that we label policy function, and the amount to be ordered is its direct consequence.

What about the process by which the retailer provides information about customer demand to the wholesaler? This process is simply non-existent in the present format of the game: no explicit transfer of information takes place. The wholesaler is limited to drawing inferences from the retailer's orders.

So far we have discussed processes and systems. Our next concept is Business Process Redesign (BPR) or re-engineering. Various authors have given a variety of definitions. Davenport and Short (1990) define BPR as 'the analysis and design of workflows and processes within and between organisations' (p. 11). In his book, Davenport (1993) uses the term process innovation. Hammer (1990) defines re-engineering as using 'the power of modern information technology to radically redesign our business processes in order to achieve dramatic improvements in their performance' as opposed to the 'use of technology to mechanise old ways of doing business' (p. 104). Morrow and Hazell (1991) refer to Hammer's definition, and rephrase it as 'the examination of the flow of activities and information that make up the key business processes in an organisation, with a view to simplification, cost reduction or improvement in quality or flexibility' (p. 38).

King (1991) states that 're-engineering aims to use the power of information technology to radically redesign business processes to improve speed, service and quality' (p. 55). She takes a more critical point of view, wondering what is new about 'the idea of analysing and then streamlining business processes before automating' (p. 56), and redefines re-engineering as 'industrial engineering concepts applied to a non-factory environment' (p. 56). Fried (1991) gives a more elaborate definition: 'BPR is a methodology for transforming the business processes of an enterprise to achieve breakthroughs in the quality, responsiveness, flexibility and costs to compete more effectively and efficiently in a chosen market. BPR uses a combination of industrial engineering, operations research, management theory, quality management and systems analysis techniques and tools' (p. 91).

Although these various definitions emphasise different aspects of BPR, two common themes emerge:

- BPR deals with major changes (several definitions use the word radical);
- BPR cuts across functional boundaries.

BPR requires starting from scratch, breaking with existing habits, and asking the question: why are we doing this? Davenport and Short (1990) contrast this approach to the continuous improvement philosophy, which implies a continuous cycle of stabilising a process, assessing its performance, and gradually improving it. BPR on the other hand, is dynamic in nature, and looks for different and better ways to perform a process. Hammer (1990) emphasises the inherent uncertainty involved in BPR: it is a jump into the unknown, rather than a sequence of cautious, carefully planned steps.

BPR can only be successful if all parties who are affected by the process being redesigned are involved. This includes the people providing inputs, those who perform the process and, maybe most importantly, those who use the output. Davenport and Short (1990), Fried (1991), Hammer (1990) and Morrow and Hazell (1991) all provide detailed discussions on what to do (and not to do) when attempting to implement BPR. All focus on the same crucial issues:

- understanding and evaluating the present process thoroughly;
- questioning every aspect of the present process: is it necessary? is it done in the most efficient way by the most suitable people?;
- using all these inputs, as well as brainstorming, to come up with a new process.

Davenport and Short (1990) and Hammer (1990) heavily emphasise the role IT can play in BPR. The former use the expression 'IT-enabled BPR' (p. 22). The recent evolution of IT has indeed created opportunities for BPR which could only be dreamt of a few years ago. Hammer (1990) mentions among others that geographical dispersion of units is no longer an issue, and the ease

with which information can be transmitted makes it possible to capture data once, at the source, and make it readily available to anyone who might need it.

It is therefore not at all surprising that a considerable amount of BPR related work is published in journals with an IT focus. The potential for successful interaction between IT and BPR is discussed at length in Davenport and Short (1990). Davenport (1993) presents IT as one of the 'enablers' of process innovation. Does this imply that sophisticated IT systems are a necessary ingredient of successful BPR projects? It would be hard to imagine an efficient process which does not rely at all on IT. But IT should remain a tool which supports and enables BPR, rather than a goal in and of itself. For instance, computer aided manufacturing has often been used solely as a tool to automate an existing manufacturing process, rather than as an opportunity to rethink the whole process. This error should not be repeated. The common error of the sixties, (attempting to make the problem fit the tool, in that case operational research tools) should also be avoided. The goal of a BPR project should not be to maximise the use of existing IT systems.

We discussed the role of IT as a support of business processes. But IT has a second, equally important role to play: Davenport and Short (1990) mention the use of IT as a design tool. As discussed further on, the use of IT as a modelling tool can greatly facilitate the process of redesigning a business process.

The process of BPR does not rely on any one particular tool. Fried's (1991) definition quoted earlier lists a whole series of possibilities. As mentioned above, Davenport and Short (1990) suggest the use of computer aided systems engineering. Morrow and Hazell (1992) advocate the use of activity mapping. This method consists of two steps:

- activity analysis (what activities, why, how frequently, what resources);
- linking the activities to the cost object.

When applying this approach to BPR, the business process is treated as the cost object. This method is an extension of activity based costing, which attempts to link every cost element to the specific activity causing that cost, to obtain a better insight into costs and profitability at both the product and the customer levels. Treating the process as a cost object and representing it as an activity map makes it more straightforward to review the process, querying the purpose of every activity. Drawing the activity map to scale also allows one to analyse the use of time as a resource. The board used to play the beer game can be thought of as a very simple activity map. As one looks at the board, some questions immediately come to mind, such as: why does it take two weeks for an order to arrive? why are there no information channels?

Another tool, which has become much more accessible due to the evolution of graphical user interfaces, is

simulation. Discrete simulation has been used successfully to design manufacturing processes for many years. At BP America (see Young, 1991) discrete event modelling has been used to gain a better understanding of various processes, and to anticipate and analyse process changes. When using this discrete event approach to tackle process-industry problems, their main problem was the reluctance of users to accept the use of a discrete event software package to model a seemingly continuous process.

In this paper, we would like to illustrate the use of an alternative approach, systems thinking, modelling and continuous time simulation, to address BPR. As mentioned earlier, BPR has its roots in systems engineering. It is therefore not surprising that the systems thinking approach is quite appropriate in a BPR context. This approach initially analyses an issue at a very aggregate level, every element of the process being represented as either a stock (accumulation) or a continuous flow. These flows are regulated by 'converters', which represent policy functions. This 'plumbing' approach enhances the understanding of the process and forces one to make explicit all policies and assumptions. This analysis can yield considerable insight into the process. Once the process as a whole is well understood, one can dig deeper into specific aspects of the process. For instance, policies can be modelled in more detail.

The resulting model can then be simulated. The continuous approach prevents one from getting bogged down in too much detail. One of the difficulties encountered is the mirror image of the one described in Young (1991): individuals who deal with a process on a day to day basis see it as being intrinsically discrete (e.g. cases of beer, arriving in truck loads) and find it hard to accept a continuous representation (a flow of cases of beer).

In the remainder of this paper we will look at the beer game as a business process and use a systems thinking approach to see how it can be redesigned to achieve lower costs (i.e. how can we limit backlogs and inventories). In its present form, the beer game has four key processes: the ordering process of the retailer, the wholesaler and the distributor, and the production process of the factory. Each of these processes is driven by an ordering policy, which can be as simple or as complex as the player wishes. There are also four delivery processes. The corresponding policies are straightforward: deliver as much of the order as the inventory position allows, and increase the backlog in case of insufficient inventory. No processes relating to information are present.

The next section discusses a continuous simulation model representing the beer game. Special attention is given to the modelling of the policies driving the ordering and production processes. This model is then used to analyse the beer game and discuss how the

system (i.e. its various processes) can be redesigned to improve efficiency.

The Model

There exist a number of ways of modelling the production-distribution game described in the section 'The Beer Game'. The model has to capture the physical flows, as well as the decision making process carried out by the participants in the game. While it is relatively easy to capture the physical flows, it is considerably harder to model the decision making processes of the participants. This is done by using a simple heuristic based on the following three criteria (Sterman 1989);

1. As customer demand increases, orders increase as well and vice versa. However, there is a lag in this response due to the time taken by decision makers to form a view of the extent and permanence of demand changes.
2. Each sector has a target inventory (given by the target coverage, i.e. the desired number of weeks of beer supply in inventory, multiplied by the expected weekly demand). It is possible to adjust the actual inventory towards the target inventory.
3. Each sector keeps track of its supply line, i.e. the quantity ordered but not yet received. If the supply line exceeds the desired supply line (given as the supply line target: desired supply line coverage multiplied by expected orders) then orders are reduced to compensate, and vice versa.

Figure 3 shows a diagram of the decision rule used for the simulation model. Similar rules have been validated and shown to be a good representation of the decision making process in the beer game (see Sterman 1989, Mosekilde *et al.* 1991).

There are three behavioural parameters which we can change in the model. The first is the desired coverage, i.e. how many weeks of expected demand we want in inventory. This parameter is fixed at 3 weeks of expected

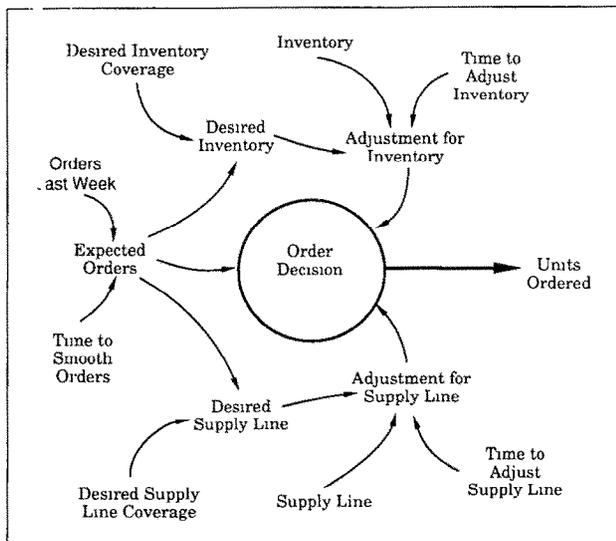


Figure 3 Order Decision

demand. The second relates to the speed with which any discrepancy between actual and desired inventory is corrected. In the present version of the model this parameter is set at 3 weeks, i.e. we want to order a third of any discrepancy between our desired and actual inventory (this could be a negative number if the inventory is too high). Finally, the third parameter determines how much emphasis is placed on the discrepancy between desired and actual supply line. This parameter is set at 0.75 i.e. the player orders three quarters of any discrepancy between desired and actual supply line. The value of the three parameters discussed above represent realistic values for participants playing the game (Sterman 1991).

The benchmark or base-case simulation of the model is shown in Figures 4a and b. As can be seen the overall pattern is similar to the experimental results shown in Figures 2a and b. Of course the exact numbers differ, as we have not attempted to estimate the parameters to fit the experimental results shown in Figures 2a and b.

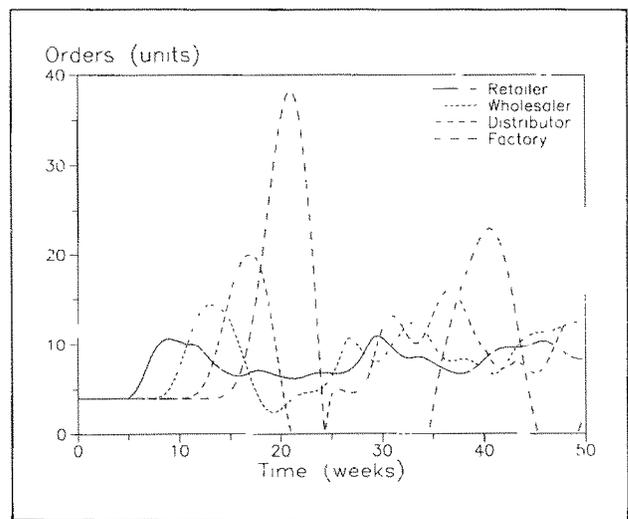


Figure 4a Orders for Beer: Base Case

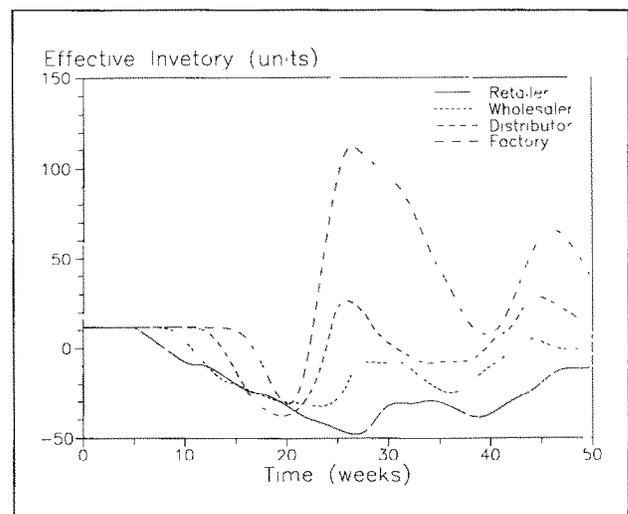


Figure 4b Effective Inventory: Base Case

Redesigning The System

When discussing how the structure of the beer game can be redesigned to reduce costs, three approaches are possible:

1. redesigning the decision process;
2. redesigning the physical process;
3. redesigning the information channels.

We consider these approaches in isolation, to be able to assess their individual impact. Of course, one could (and should) also consider various combinations to further improve the system.

The first approach involves changing decision parameters such as the desired inventory cover and the speed at which inventory is adjusted [see Mosekilde and Larsen (1988) and Mosekilde *et al.* (1991)]. We will consider these parameters as given. The second approach raises questions such as: can we reduce the delays in the system; do we really need two intermediaries between factory and retailer? The third approach relates to the adequacy of the information channels. In the present system, only the retailer observes customer demand, and the other parties draw (delayed) inferences from their incoming orders. What would be the impact of making the retailer's information available upstream in a timely manner?

As mentioned in the previous section our focus will be on the latter two approaches. We consider four scenarios:

1. cut out the ordering delay;
2. cut out the wholesaler and distributor;
3. make customer demand data available to the factory;
4. make customer demand data available to all parties.

The first two scenarios relate to the physical process, the latter two to the information channels. We chose these specific scenarios, because they capture the changes most commonly suggested by management teams confronted with the difficulty of managing the beer game.

One question needs to be addressed at this stage: do the scenarios deserve the label 'process redesign' or do they merely qualify as steps in a continuous process improvement effort? Cutting out two intermediaries clearly is a dramatic action, and falls in the former category. Similarly, cutting out the ordering delays fundamentally affects the structure of the system, but is a less radical change. On the other hand, making demand information available to the factory may seem like a quite minor change. Still, it has important implications, as it enables the factory to base production decisions on up to date customer information rather than on a delayed and distorted picture of this information. It also represents a radical departure from

present practice. Today, information technology makes this kind of information transfer possible. There is no doubt that this would be a major factor in a comprehensive redesign effort.

We next discuss these four scenarios in more detail, giving for each a visual representation of the changes, a discussion of their impact, and the resulting inventory levels and cost figures. Caution is required when comparing the costs of the various scenarios, as there are two relevant figures: the actual cost (obtained from the simulation and which we would expect to observe in practice, resulting from the system structure and the parties' actions) and the accepted cost (intrinsic to the system structure, which would result if all parties achieved instantaneously the desired inventory level throughout the period). Typically, the actual cost will exceed the accepted cost. In isolated cases the reverse may be true, for instance when one or more parties (unintentionally!) carry a lower than desired inventory, due to ordering and shipping delays, without running into a backlog situation.

Scenario 1. No Ordering Delays

In this scenario, we consider the impact of removing the ordering delays. This allows the various parties to adapt their inventory more quickly as information about customer demand flows faster from the retailer to the factory. The retailer is still the only one to know the true customer demand, the wholesaler acts on information passed on by the retailer as before, the distributor on information from the wholesaler etc. (Figure 5).

Figure 6 shows the simulation results for scenario 1, which indicate a considerable improvement over the base case (Figure 4a). We can compare the results by measuring the amplification in the model (i.e. comparing the size of change in customer orders (from 4 to 8 cases per week, i.e. 4) to the observed change in factory orders (defined as the peak order minus the initial order level of 4). In the base case the amplification is around 900% while in scenario 1 the amplification is 'only' 500%. This is a considerable improvement. It is also worth noting that the system seems to settle down faster. In the base case 50 weeks represents a little more than 2 cycles, while in scenario 1 there are 3 complete cycles. The reduction in delays has made the system more responsive to external shocks.

How could a policy like the one described in scenario 1

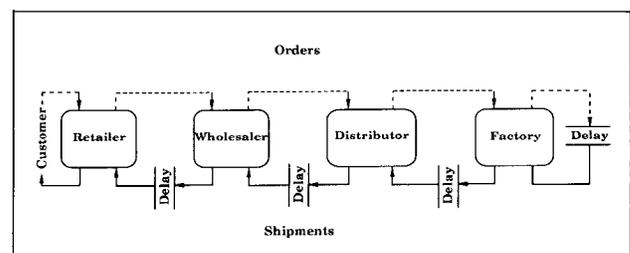


Figure 5 Structure of Model: Scenario 1

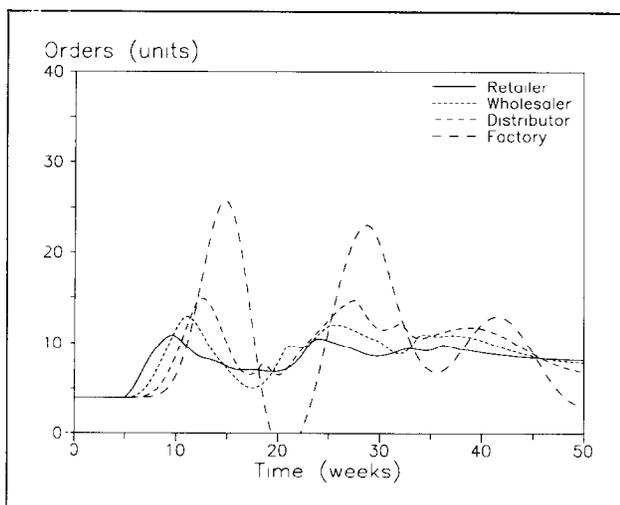


Figure 6 Orders for Beer: Scenario 1

be implemented? In the base case orders are placed simultaneously in all 4 sectors based on 2 week old information. A policy as the one discussed in scenario 1 could be implemented by sequential decision making. The retailer uses the observed customer orders to determine how much to order. He then uses a fax, phone or computer system to send his orders to the wholesaler, who uses this information to determine his order and pass it on to the distributor etc. Altogether not too difficult a system to set up, with great effect.

Scenario 2. No Intermediaries Between the Factory and the Retailer

Scenario 2 is a drastic restructuring of the supply system: the cost of buffer inventories at the intermediaries is eliminated, and both ordering and shipping delays are cut by two thirds. This results in the factory having a better picture of customer demand, and being able to react faster to changes in retail orders (and hence customer demand).

As we would expect, this redesign improves the performance of the system dramatically. The amplification is down to 350%. The production distribution system settles down after 25 weeks, whereas previously it was still fluctuating after 50 weeks. Table 1 shows that costs are considerably lower than in the base case and in scenario 1.

This reduction is due mainly to the elimination of the two middle-men in the system, and their inventory. This cost comparison does not give the whole picture, as by cutting out two sectors we have also reduced the desired inventory in the system as a whole from 12 to 6 weeks cover. To make a fair comparison we have computed a cost index, given as Total cost / Expected Cost where total cost is the actual cost and expected cost is customer demand times 3 (weeks of desired coverage) times the number of sectors. These calculations are shown in Table 1, for the various scenarios. The cost index for

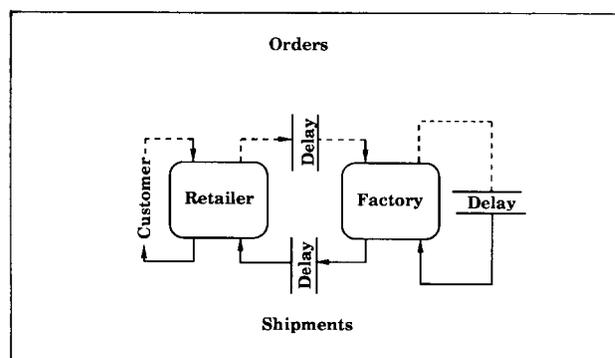


Figure 7 Structure of Model: Scenario 2

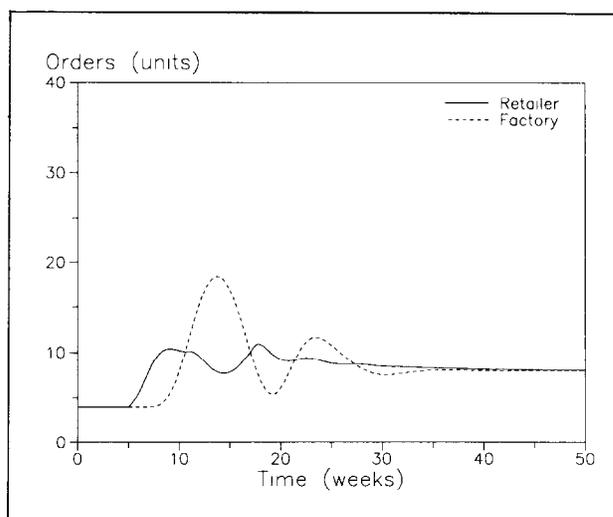


Figure 8 Orders for Beer: Scenario 2

scenario 2 is 0.82 (i.e. Total Cost is less than Expected Cost) which is considerably better than the value of 1.47 obtained in the base case.

How could Scenario 2 be implemented? We basically want to remove the buffer inventories between the retailer and the factory. This sounds like a simple thing to do, but other considerations not represented in our current model (e.g. geographic distribution) may make it impossible to get rid of all the buffer inventories. A first step could be to remove one layer of intermediaries.

Table 1 Costs and Amplification

Scenario	Total Cost (£)	Cost Index	Amplification (%)
Base case	3358	1.47	900
Scenario 1	1944	0.85	500
Scenario 2	939	0.82	350
Scenario 3	2295	1.01	425
Scenario 4	1293	0.57	200

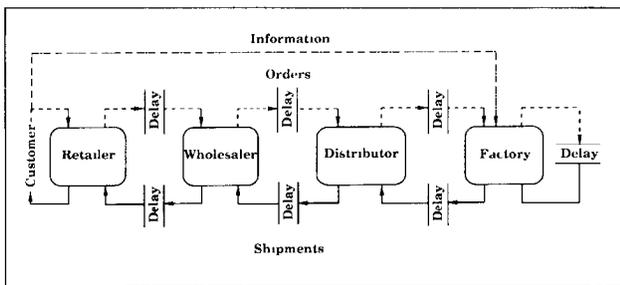


Figure 9 Structure of Model: Scenario 3

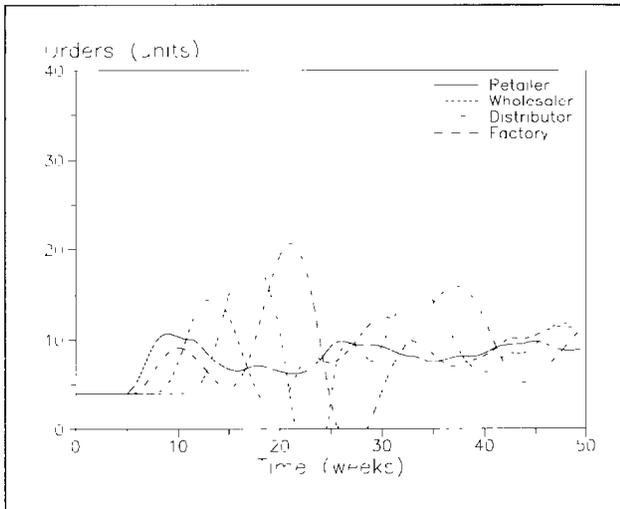


Figure 10 Orders for Beer: Scenario 3

Scenario 3. The Factory has Access to Customer Demand data

In this scenario the factory still needs to respond to the distributor's orders, but has access to up to date customer demand to plan ahead. The factory still has to ship whatever the distributor wants, but instead of using the distributor's demand to estimate expected future demand the factory can use actual customer demand.

Providing the factory with up to date information regarding customer demand has greatly improved the system performance. The amplification in the system is down to 425% compared to 900% in the base case. The cost is also significantly lower (see Table 1). Factory orders are no longer amplified by the use of the distributor's orders for decision making. This reduces huge inventories and backlogs usually found in the factory sector, resulting in lower cost. However, the time required for the system to settle down after a one-off disturbance (the increase in customer order) has not improved. The factory's more accurate production decisions result in longer delays to fulfil the orders. The factory incurs a smaller backlog, but for a longer time period.

Given this result, would it not make sense to let all sectors have the information the factory has i.e. let all sectors know what the real customer demand is? This

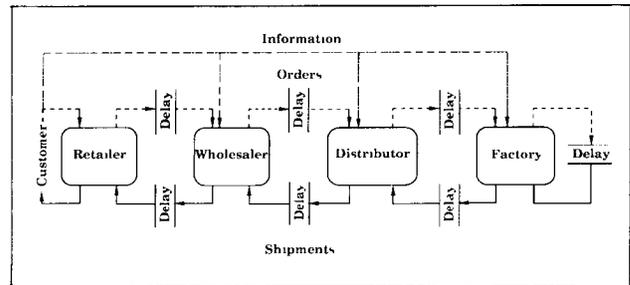


Figure 11 Structure of Model: Scenario 4

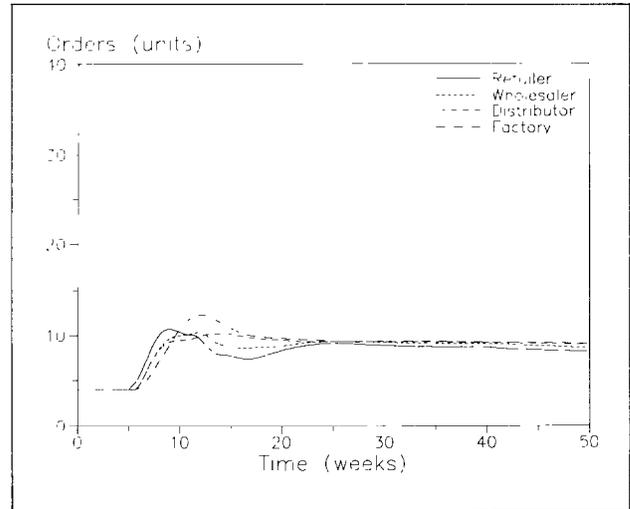


Figure 12 Orders for Beer: Scenario 4

is exactly what scenario 4 does. Comparing scenario 3 and 4 will enable us to evaluate the incremental advantage of a 'full information' system over the case where only the retailer and the factory have up to date information.

Scenario 4. All Parties Have Access to Customer Demand Data

This additional information also enables the wholesaler and distributor to plan ahead, although they still need to satisfy the orders of the retailer and wholesaler respectively. The factory, distributor and wholesaler will in this case use the customers' orders to calculate their expected orders, instead of the orders coming from their downstream customer. They still have to ship whatever the downstream sectors want, but make their order decisions based on the customer demand and the adjustment from inventory and supply line.

Figure 12 shows the simulation results. The amplification has almost disappeared, being only 200% compared to 900% in the base case. Furthermore, the system stabilises after one cycle, i.e. there is no over nor undershooting. Orders raise to about 12 and then slowly adjust to the 'correct' value of 8. This was to be expected, as providing all sectors with the 'true' customer demand has eliminated the cause of the cycle. The remaining amplification results from the need to

adjust the inventories to the new higher level, i.e. it is necessary to increase production to raise the inventory from 12 to 24 cases. The total cost is almost halved compared to scenario 3, showing that the 'full' information system has a considerable incremental value. The cost index of only 0.57 indicates that the cost is lower than expected. This is by far the most successful way of redesigning the system if we do not want (or are unable) to change the number of sectors. The total cost is approximately 40% of the cost in the base case. Only the two-sector model (scenario 2) achieves a lower cost.

How could we implement a system as in scenario 4? Possibilities include installing a sophisticated management information system, or an automated inventory control system. This seems to be what a number of large retailers are just starting to do. A computer system takes care of ordering when a pre-determined lower limit for the inventory is reached. Such a system contains the information requested for scenario 4. The complexity in this scenario lies in the need to deal with a large number of retailers, wholesalers and distributors.

Conclusion

The beer game is a microcosm of how real organisations function. It provides considerable insight into the web of business processes that lie behind routine organisational tasks like matching production and shipments to customer orders. Most people are shocked when they discover the gyrations that the typical factory has to endure in order to accommodate a simple step change in demand. If demand increases by 50 per cent then why doesn't production just do the same? The answer lies in the processes that link the customer and the factory. Usually we don't think about all the intermediate steps, judgements and actions.

The model of the beer game shows graphically the major processes and how they are linked. Just filling customer orders is really quite complex. The retailer has to forecast demand, control inventory, track the supply line, manage costs and place orders. So too do the wholesaler and distributor. Moreover there are shipping and order delays at each stage that complicate the management task. Altogether there are more than a dozen loosely coordinated judgements and actions. Once the structure is modelled and simulated, the factory's behaviour becomes understandable. Further modelling and simulation then open the way for informed redesign.

Can the same concepts and tools be applied to full-scale business problems? Evidence from numerous projects suggests they can, though proponents may argue about the value of different mapping tools (Richmond *et al.* 1992, chapters 1 and 2) and the importance attached to simulation modelling.

A recent article in *International Business Week* (1992)

reports a re-engineering intervention by Hammer at ITT Sheraton. Twenty two top operating executives attended Hammer's three-day seminars. The result according to the management team: 'We threw away the book and invented a new hotel. The typical 300-room Sheraton Hotel had required up to 40 managers and 200 employees. By eliminating narrowly defined jobs and rethinking antiquated procedures, ITT found it could run a re-engineered version of 250 suites with only 14 managers and 140 employees — with higher customer satisfaction. We redesigned the processes of the company and eliminated everything we didn't need to do'.

Re-engineering projects involve management teams closely in the design process. As in the ITT example, the managers themselves identify key processes, map them, discuss them and then think about new ways to organise. Such team mapping breaks down functional barriers and enhances buy-in to change. However, mapping alone cannot provide a means to challenge preconceptions about the consequences of redesign. Here's where simulation comes in.

As the beer game example shows, systems thinking, modelling and continuous time simulation provide a consistent framework for carrying the design process from mapping all the way through to redesign. The underlying field known as system dynamics has a long history of application to practical problems of business design in companies ranging from biotechnology (Morecroft, *et al.*, 1991), to insurance (Senge and Sterman 1992). The typical approach to a project is to form a cross functional management team to supply information about the company's operating policies from which to construct a simulation model. As with the beer game model, simulations provide a basis for understanding the company's present performance and for redesigning policy.

The links between system dynamics and business process redesign are only just beginning to be grasped and exploited. But the challenge for the future is already clear — better designed organisations where senior management play a direct role in design activity. The founder of system dynamics MIT's Jay Forrester provides a vision for the future (in Keough and Doman 1992):

'The role of senior management, especially of chief executive officers, should be that of corporate designers and not corporate operators. CEOs may need assistance with the details of simulations. But they will eventually have to be competent in the creation and use of models if they are to lead such planning activity effectively.'

... There are not yet many CEOs who are serious about corporate design. In the future, however, the most effective CEOs will not be those making day-to-day decisions, but those who are designing their corporations — that is, structuring which information is available and to whom, establishing

policies that govern decisions, and deciding what decisions should be made where.

Future CEOs will focus on creating corporate designs in which ordinarily competent people can succeed. Too often, we see people in corporate positions repeatedly blamed for failure and replaced when the fault lies in the situation into which they have been put. Much of the time, it is the design of the organisation that is defective'.

References

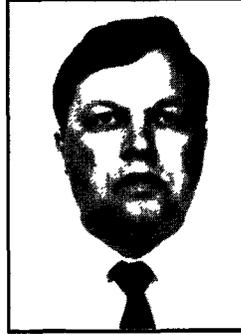
- Byrne, J.A., 'Management's New Gurus', cover story, *International Business Week*, 42-50, August 31, 1992.
- Davenport, T.H. and J.E. Short, The New Industrial Engineering: Information Technology and Business Process Redesign, *Sloan Management Review*, Summer, 1990, 11-27.
- Davenport, T.H., *Process Innovation*, Harvard Business School Press, Cambridge, MA, 1993.
- Fried, L., A Blueprint for Change, *Computerworld*, December 2, 1991, 91-95.
- Hammer, M., Reengineering Work: Don't Automate, Obliterate, *Harvard Business Review*, (July-August) 1990, 104-112.
- Jackson, M.C. *Systems Methodology for the Management Sciences*, Plenum Press, 1991.
- Jenkins, G.M., The Systems Approach, in J. Beishon and G. Peters (eds.), *Systems Behaviour*, OUP, London, 1971, 78-104.
- Keough M., and A. Doman, 'The CEO as Designer', an interview with Jay W. Forrester, *McKinsey Quarterly*, 2, 1992, 3-30.
- King, J., Rip it Up, *Computerworld*, (July 15), 1991, 55-57.
- Morecroft, J.D.W., D.C. Lane and P.S. Viita, 'Modelling Growth Strategy in a Biotechnology Startup Firm', *System Dynamics Review*, 6,2, 1991, 93-116.
- Morecroft, J.D.W. and J.D. Sterman (eds), 'Modelling for Learning', special issue *European Journal of Operational Research*, 59,1, 1992.
- Morrow, M. and M. Hazel, Activity mapping for Business Process Reengineering, *Management Accounting* February 1991, 36-38.
- Mosekilde, E. and E.R. Larsen, Deterministic Chaos in the Beer Production-Distribution Model, *Systems Dynamics Review*, 4,1-2, 1988, 131-147.
- Mosekilde, E., E.R. Larsen and J.D. Sterman, Coping with Complexity: Deterministic Chaos in Human Decision Making Behaviour, in J. Casti and A. Karlquist (eds), *Beyond Belief: Randomness, Prediction and Explanation in Modern Science*, CRC-Press, Boston, 1991, 199-229.
- Richmond, B.M., S. Peterson and C. Charyk, *Think Continuous Process Improvement Module*, High Performance Systems Publications, Hanover, NH 03755, 1992.
- Senge, P.M. (1990), *The Fifth Discipline*, Doubleday, New York.
- Senge, P.M. and J.D. Sterman, 'Systems Thinking and Organizational Learning: Acting Locally and Thinking Globally in the Organization of the Future', 'Modelling for Learning', special issue *European Journal of Operational Research*, 59,1, 1992, 137-150.
- Sterman, J.D., Modelling Managerial Behaviour: Misperception of Feedback in Dynamic Decision Making Experiment, *Management Science* 3 vol. 35, 1989, 321-339.
- Young, M.B., Reengineering the Process, *Manufacturing Systems*, (May), 1991 48-51.



ANN VAN ACKERE, London Business School, Sussex Place, Regent's Park, London NW1 4SA

Ann van Ackere is Associate Professor of Decision Science and currently Chairman of the PhD Programme. She holds a PhD from Stanford and a degree in Applied Economics from UFSIA, Belgium. Her research focuses

on the use of simulation and queuing models to evaluate manufacturing and services operations. She is investigating the use of simulation methods for business re-engineering and business process redesign. She has published in Operations Research and Management Science.



ERIK REIMER LARSEN, London Business School, Sussex Place, Regent's Park, London NW1 4SA

Erik Reimer Larsen is Coordinator of the Learning Centre Resource Unit. He obtained his PhD from the institute of Economics, Copenhagen Business School. He has published in the areas of

mathematical chaos, nonlinear dynamics and system dynamics. His doctoral studies included appointments as visiting scholar at London Business School and the Sloan School of Management, MIT.



JOHN MORECROFT, London Business School, Sussex Place, Regent's Park, London NW1 4SA

John Morecroft is Director of the Learning Centre, Associate Professor of Strategic Management and a leading expert in strategic modelling and systems dynamics. He holds a PhD from MIT and was on the

faculty at the Sloan School of Management. He has published widely on the topics of computer-based strategy support, behavioural modelling, and simulation analysis. His latest publication is an edited volume Modelling for learning (co-edited with MIT's John Sterman) and in 1990 he received the Jay Wright Forrester award of the System Dynamics Society.