

STELLA MODELLING PROCESS FOR A MANPOWER STRATEGY

Matthew Byrne
Andrew Davis

Business Consultants (ICA/41)
Shell International Petroleum Company
Shell Centre London
United Kingdom

ABSTRACT

This paper describes a process used by Shell Business Consultancy to model manpower in an IT department. The IT world has been undergoing rapid changes in the last decade due to a highly mobile job market, skill shortages and the widespread use of contractors. Given these problems, it is important that management understands the implications of personnel policies. Shell Business Consultants built a STELLA model with the managers involved to focus their attention on the key issues and to enable them to actively explore policy options. The paper focuses on the process used in this project and combines it with our other experience to draw some generally applicable lessons that should be of use to all practitioners in the field.

INTRODUCTION

This paper takes the reader through a consultancy process that was recently used by Shell Business Services. Shell Business Services is a part of the Shell International Petroleum Company, providing an internal consultancy in the field of decision support for Shell companies. A particular project will be discussed from the point at which contact was made with the client right through to the final deliverable. The discussion will focus more on the consulting process rather than a detailed explanation of the model that was built. In the first section we will discuss the client's organisation, its concerns and the process by which these were distilled into a useful STELLA model. The second section explores the issues of model complexity and puts forward a matrix to help consultants classify the type of project they are involved in and what the main thrust of their consulting process should focus on. Finally, the third section takes a look at the way in which projects are presented to the client in terms of a deliverable. It contrasts previous approaches that we have used with a new approach that was used with this client.

WHAT IS STELLA?

Stella is a software package used for modelling business thinking. It is particularly suited to modelling issues in a dynamic way and is ideal for coping with systems that contain feedback.

ENTRY TO THE CLIENT

The client was a group of managers with responsibility for managing 1400 computing staff with varied skill categories, in different geographical regions and across different sectors of the business. They have responsibilities for determining

overall numbers of staff at each level in the company, the calibre of those staff, the blend of short-term and medium-term contractors to permanent staff, recruitment and career progression, to mention but a few. As can be imagined the impact of strategic decisions in this arena can have very significant long term effects on businesses of Shell companies. The clients were interested in gaining a better understanding of this complex system and exploring the long term implications of specific strategies. Until recently, the only tool available to these managers to aid them in their planning has been a linear and fairly static view of the short to medium term. This existing spreadsheet tool was felt not to offer the kind of long-term dynamic view which they needed.

The clients had already been targetted by a piece of earlier marketing. This had involved building a small "throwaway" model to demonstrate the capabilities of system dynamics and STELLA. It had been built around a specific problem in the personnel field and was presented to the management team to interest them in the applicability of such an approach to their wider strategic concerns. Thus, the client was already familiar with our work and to a certain extent familiar with the system dynamics approach.

INITIAL MEETING

Our first meeting was with the whole of the management team consisting of about a dozen people. The original "throwaway" model was demonstrated live and it was explained how this work could be expanded to address more of their initial concerns. However, as we have often found to be the case, time, and the process of building the first model had changed the particular problems that were foremost in their minds. The main management concern was to meet future target levels of staff at the lower technical levels whilst still bringing sufficient people through the system to staff the management positions of the future. It was decided, therefore, that the best way forward would be to take one of their business sectors and use it as a guinea-pig to test the usefulness of a new model. This new model would address "softer" and more wide-ranging issues than before.

STRUCTURING MEETING

The first meeting to structure a new model took place at the client's site. It was an informal meeting in which members of the management team would wander in and out of the discussion when they had the time to attend. On the face of it, this would seem a strange way to run a meeting, but from a practical point of view it was the only way to get all of the stakeholders together in a single session, if not all of them always at the same time! At the close of the day we managed to bring all of the team together and agree a series of causal loop diagrams that should form the logic of the model. It is interesting to note that while we had complete agreement amongst the management team as to the important issues and their relationships, when we showed the model logic to other personnel professionals, they brought out other issues which they felt to be of equal importance. This emphasises the fact that any model built depends crucially on the management team that was involved in the construction. It is therefore vital to have all the major players agreeing on the logic, or the model's credibility will be immediately lost.

A very interesting problem arose from this meeting which should be of general interest to modellers. We found that there was a difference between the model which had been explained to us by the management as being the theory of how the system operated and the actual system that operated in practise. We were faced with a set of logics and guidelines which the management believed they used, albeit liberally, but in practice they palpably did not. With the first stage of modelling we

were able to show them the effects of managing the system under both sets of logics and demonstrate that if they were to use the "theoretical" logic in reality, then there would be serious problems in the future. Thus, the modelling process had already unearthed an interesting difference in perceptions of the system.

We believe that this is a major benefit of the "Modelling as Learning" approach (David Lane, 1988). The client is forced to take a clear and consistent view of their system and build a real understanding of the underlying logics. This is one of the areas where a traditional modelling approach might have failed as the client is encouraged to "leave it to the consultant". In this case, the consultant could have ended up presenting a model which gave obviously inconsistent results leading to the client rejecting the model altogether because the only client understanding of the model would have been at the level of the results it was giving rather than the logic it was using. Using the "Modelling as Learning" approach, we were able to explain the reasons for the results in terms of the causal loop diagrams that they themselves had built. This caused them to re-examine those logics and agree that they were not consistent with what they were doing in practise. They were learning!

WHAT VALUE COMPLEXITY ?

There is always a play off to be made between the complexity of the model and how understandable it is. Clients tend to want to model reality with every last detail. Unfortunately, the value of such a model (which indeed it would no longer be) would be lost in a morass of detail and complexity. The real value of STELLA models is that they can quickly communicate a set of ideas and interactions and allow the user to check the consistency or sanity of those ideas. As soon as models start to contain life size helpings of complexity, the client often loses the thread of what is actually contained in the model (not to mention the consultant!). The only person who will really be able to run and benefit from the model will be the person who built it. This surely defeats many of the reasons for using such a user friendly tool and destroys the chance for any real learning to take place.

However, if models are going to be anything but trivial they must contain some complexity. The consultant's role then becomes important in managing the interface between the client and the model. Building in enough complexity to satisfy the customers need for a representative model while balancing this with ensuring the model does not become a "black box " in the client's eyes is a critical part of the consultancy process. We have experienced both extremes of this dilemma. Having built one of the most complex STELLA models and found that no-one really understands it or uses it, has proved to be an expensive and largely pointless exercise. On the other hand, going to clients with trivial "nobby" models that everyone can quickly understand, and having them ripped apart for being of no value at all is even more soul-destroying. So clearly complexity is of value, but so is balance.

In this particular study, the client wanted issues included which required complex pieces of STELLA modelling. Though the modelling from a system dynamics point of view proved to be successful, it is our opinion that the extra complexity led to a reduction in the client's understanding of the functioning of the model. Although the model was theoretically more "correct" with the complexity included, it seems clear that the client gained no real benefit from the extra work being done. It merely led to the model being a "black box" in their eyes. Yet, whether they would have accepted a model which had not included this extra level of complexity is unclear.

Complexity is thus a double-edged sword with significant implications for the consultancy process. Adding increasing amounts of complexity into a model should be approached with care and an awareness that it might be leading the consultant into a series of different management problems. A summary of our experiences in this area is contained in Fig. 1, which we have found to be a useful model for classifying projects and making the consultant aware of the different process issues associated with different levels of complexity.

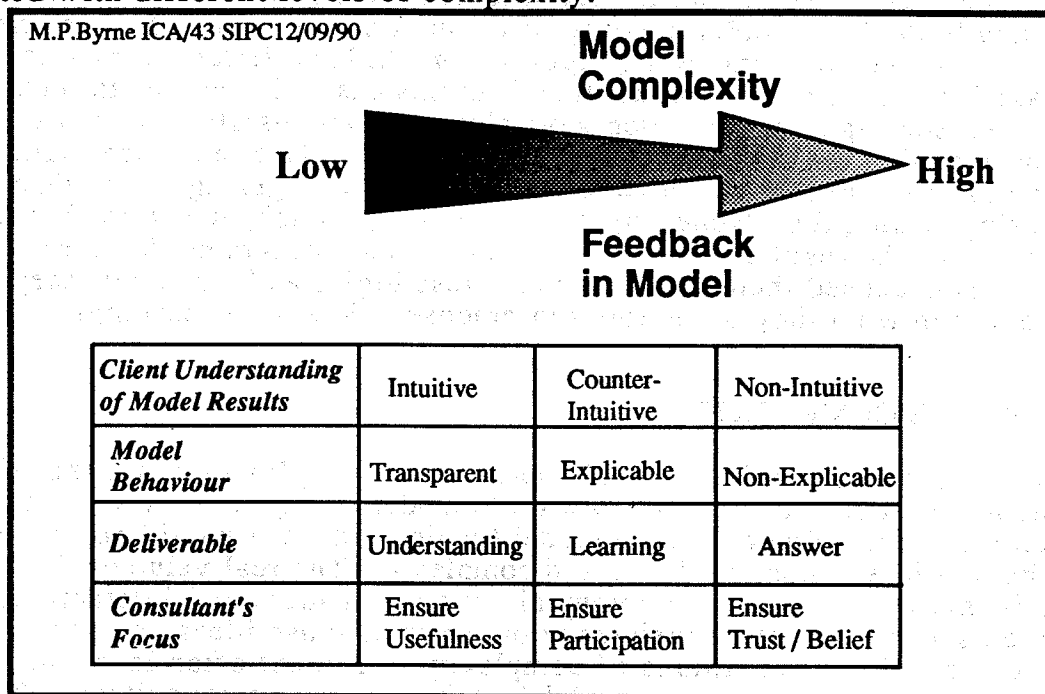


Fig.1. The Value and Problems of Model Complexity

As explanation of this matrix, we will discuss a couple of cases.

The small “throwaway” model that was built for the management team in order to sell the system dynamics concept would have fallen somewhere in the intuitive column. The model had almost no feedback and was very simple with only a small number of variables. When we ran examples through the model, the clients were able to predict the result with ease due to the model’s linearity and simplicity, thus their understanding of the results was intuitive and to them the model was transparent. For these sorts of models the deliverable is a reinforcement of the client’s understanding of the system and the modelling technique. As the consultant, you should be concerned with ensuring that the model is performing some useful function. The danger with over-simplified models is that they are useless.

The first significant model that was built for this study would fall into the middle column. The model had a significant amount of feedback and was reasonably complex. When we ran examples through this model, the clients were surprised by the results that emerged. They had predicted that a variable would move in one direction and the model showed it moving in the opposite direction, the results were counter-intuitive. This is often due to the amount of feedback in the model. However, the results were explicable because we could refer to the causal loop logics that were developed with the client and trace the reason for results being counter-intuitive. This process of going back to the logics, explaining and then refining generates client learning about the system. Yet, we were only in the

position to generate this learning because the client had built the logics and been actively participating in the whole process.

By the time we came to build the final the model, the clients had decided that they needed many more variables in the model which further increased it's complexity and took us into the final column. Now we had a very large model that was giving us results that were non-explicable due to the large amount of complex logic. The results were not only counter-intuitive they had become completely non-intuitive and all we could do was communicate an answer to the clients that they either believed or didn't. By this stage we had to ensure that they trusted us enough to trust the model. Thus, the management of the client had shifted from getting them to believe the model to much more focus on getting them to trust us and our ability.

Our feeling is that the importance of STELLA as a learning tool demands that the client has a grasp of the model. Once that has been lost or significantly reduced then the value of the process diminishes significantly. Thus, it is important that a balance be struck between complexity and value. Too little complexity and the model will be of no value as the client won't buy it, too much and the client won't understand it and the resultant value from the running of the model will be lost. It is all very well asking a large number of "what if ?" questions and seeing graphs generated in front of your eyes, but this is not the real learning experience. That comes when the user sees the results of a run and then asks "why am I getting this behaviour?". This circular process of running the model, comparing the results with the real world situation, going back to the model logic, analysing and UNDERSTANDING why this particular behaviour pattern has been obtained and then proceeding, is where the real learning is generated. This "Learning Wheel" is shown below. If any of the six spokes in the wheel are not present then the wheel is likely to buckle or collapse and destroy the chance of generating any real learning.

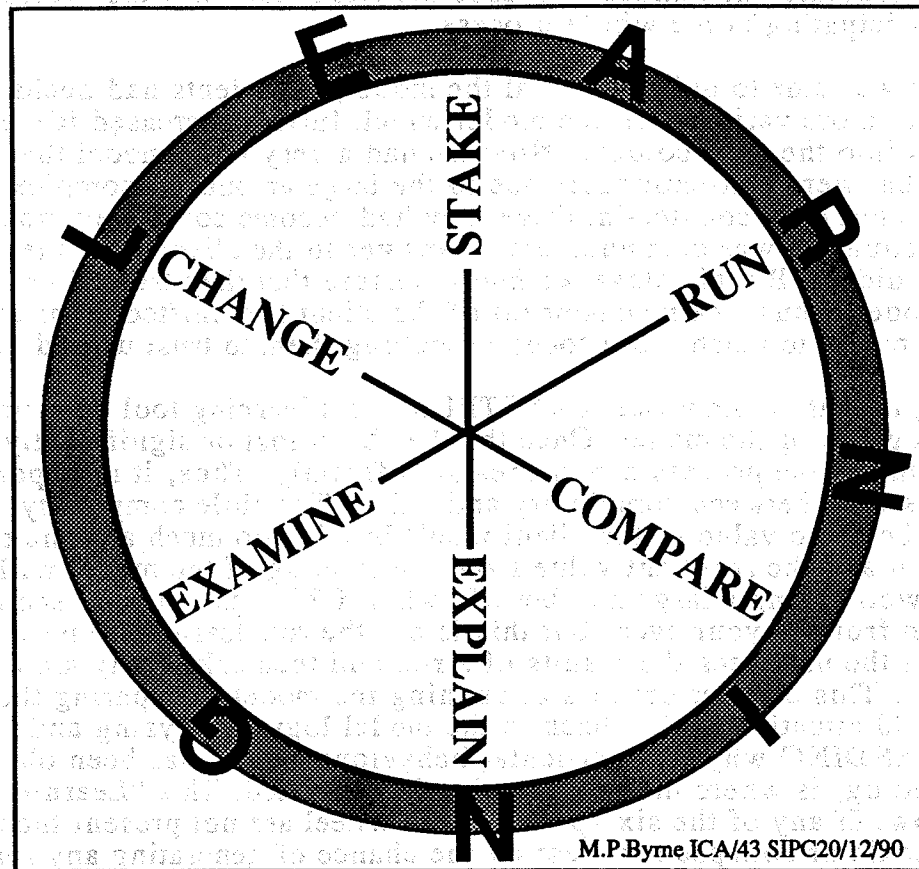


Fig. 2 The Learning Wheel

Expanding briefly on these six stages:

1. The first step should be to get the client to “put a stake in the ground”. By this we mean that they should commit themselves before any run of the model is made to what they think the results will be. This has two functions. It stops the inevitable comment after they have seen the results of “Well, that’s obvious..”. Indeed with 20/20 hindsight it often is. It also gets them to really think before making a run rather than just treating it like a game of space invaders.
2. After discussing what the results might look like, make the run.
3. Then compare the actual results with those that were expected. Is there any difference?
4. If there is a difference, then go back to the causal loop logic and explain why these results were obtained.
5. If the logic is found to be wanting then go back and examine the assumptions that went into it.
6. Change the model to fit the new logic described in phase 5.

This process continues until the team are satisfied with the model.

If the model has been built without any real participation of the client, then the consultant will find it very difficult to execute any of these phases except the RUN element. It then becomes almost impossible to generate any significant learning.

DELIVERABLES

The real aim of any management or business study should be to initiate tangible improvements in the business. This implies that the decision maker or power broker in the system must understand, believe and trust the results of the study. There is no better way of ensuring this than making the problem owner the problem solver. The consultants role is then to manage the thinking and problem solving processes rather than the actual problem solving itself.

In previous system dynamics projects it has always been deemed necessary to produce a report of the project. This has been the result of existing culture; the feeling that a report demonstrates by its sheer size the importance of the work that has been carried out. Yet, the very nature of a report suggests that the client needs to learn about what has been done. If the project has been carried out successfully, then this reason is redundant. The client should understand what has been done, because he/she did most of it. The participative nature of the process should be aimed at involving the client in the work. That is where the learning and understanding is really generated, not in a report to be filed in a cupboard, never to surface again.

Indeed the whole reason for carrying out a system dynamics project according to the Modelling as Learning principle, is to ensure that the results from the study initiate action and real business benefit. The problem is, what is the most effective way of ramming home the issues and learning that are generated by a study?

The workshop environment represents an interactive and dynamic learning experience rather than the one-way, static report method. As such, the workshop is much closer to the philosophy of Modelling as Learning. The approach is also much better suited to the capabilities of STELLA. One of the real plus points of STELLA is that it is so easy to make changes to a model, run it and then watch that particular scenario unroll before your eyes. The process of then comparing the results of a run with perceptions of the real world, using the underlying logic in the model to explain why it is happening generates discussion, new ideas and learning. Not only that, it's exciting !! The package is an interactive learning tool. As such, we feel that the only way to make the project and model really come alive is to allow it to do just that. Make it come alive in a live environment, run a workshop!

Running a workshop is however no bed of roses for the consultant. It is far easier to sit down in a nice cosy office environment and write a safe report of the study. No one to ask silly awkward questions, no-one trying to do things to your model that you never intended in fact no-one really challenging you and your work to prove their worth. Also from the client's point of view it involves a commitment of often scarce time and energy to something which won't give them the answer. Trying to sell the concept of learning rather than solving, understanding rather than mastery and generating more questions and ideas rather than giving a solution is difficult.

So clearly there is a problem here. How can we get the decision makers to commit the necessary time and energy to a study? If they don't make that commitment and choose instead to delegate the responsibility, then how do we get them to buy into what has been done? In other words, how do we make sure that our work is the catalyst for real change in the business?

The approach we took with the clients in Aberdeen was to offer them belt and braces. The safe and instantly recognisable report, a presentation and the more challenging danger of a workshop.

With a model that is addressing strategic issues, it is very important to have the decision makers understanding and using the model. However, in this case the management had delegated much of the modelling process to their staff due to the scarceness of their own time.

Therefore, for the work to make any impact it was important that the managers first understood the underlying logic in the model. The presentation was used for this purpose. We decided that this presentation would have much more impact if their staff were to explain the model rather than ourselves. As one of the ideas behind Modelling as Learning is that the process generates learning and understanding, we thought it only appropriate that the client's team should be presenting. At first this may seem like us taking the easy ride. Being paid all that money and not even making our own presentation. However, it was far more of a test of our work and the success we had achieved than giving the presentation ourselves. If the clients stood up and made a mess of things then it would be a bad reflection on our work. We were putting our reputation in their hands! Yet it worked well. Not perfectly but the effect was as important as the content. It meant that the management team felt their department was still the problem owner. They hadn't handed over the problem to anybody else.

Having now read the report and listened to an explanation of the model logic, the management team were prepared for a workshop. The model was used live with an overhead projection screen so that the outputs from the model could be seen clearly by all. The first session of the workshop was spent developing two scenarios that the management felt were possible futures for their system. The scenarios were described in terms of the model variables; how many people they would need in this future, what balance they wanted between permanent and contract staff, how many people they would recruit each year and at what level. Once there was general agreement on the details of each scenario, we discussed what sort of management problems each scenario might develop and whether the system would be able to cope. Having detailed what we thought would happen under each scenario, we then ran the cases through the model and watched the results unfold before our eyes. As the graphs were produced an instant discussion would develop to explain why we were getting that type of behaviour and why it was different from our expectations. The STELLA model was acting as both a catalyst and a focus for those discussions. Several runs were made over a period of two hours as the management explored different options and discussed possible policies. Each time we followed the "learning wheel" round, and each time some new aspect of the system was discovered and debated. The workshop had succeeded. Managers left the session saying they had moved further forward in this one session than they had in the whole of the last year. Where there had been arguments over what the implications of a strategy might be, they were now able to run that strategy through a model of consistent logic that they had agreed upon. The process turned argument into discussion and understanding.

Putting together a package of a workshop, presentation and report certainly seems to have worked. The clients can get much more of a "feel" for their problem by interacting with the model directly rather than reading about the benefits that the consultants gained from it. And by running the model themselves they are a lot more likely to understand and believe the results (not to mention criticise them!).

CONCLUSION

The modelling process that was used with these clients proved to be successful. Managers are focussing their attention on the key issues and are actively exploring policy options with the aim of developing a robust strategy for the future. We believe that much of the success of this project was due to the consultancy process rather than the specific modelling work itself which is why much of this paper has concerned itself with a discussion of these issues. We would be very interested to hear from other STELLA users as to how they have approached the problem of adding real business value.

REFERENCES

Lane, D.C., 1989, Modelling as Learning, Proceedings of the European Simulation Congress