




Learning for Action

**A Short Definitive Account of
Soft Systems Methodology
and its use for
Practitioners, Teachers and Students**



**Peter Checkland
and John Poulter**

Smyth, Frank Stowell,
Mark Winter, the late
er willing and cheerful
adwritten manuscript.

Preamble

A Ten-Minute Account of Soft Systems Methodology for Very Busy People

- We all live in the midst of a complex interacting flux of changing events and ideas which unrolls through time. We call it 'everyday life', both personal and professional. Within that flux we frequently see situations which cause us to think: 'Something needs to be done about this, it needs to be improved.' Think of these as 'problematical situations', avoiding the word 'problem' since this implies 'solution', which eliminates the problem for ever. Real life is more complex than that!
- Soft Systems Methodology (SSM) is an organized way of tackling perceived problematical (social) situations. It is action-oriented. It organizes thinking about such situations so that action to bring about improvement can be taken.
- The complexity of problematical situations in real life stems from the fact that not only are they never static, they also contain multiple interacting perceptions of 'reality'. This comes about because different people have different taken-as-given (and often unexamined) assumptions about the world. This causes them to see it in a particular way. One person's 'terrorism' is another's 'freedom fighting'; one person sees a prison in terms of punishment, another sees it as seeking rehabilitation. These people have different *worldviews*. Tackling problematical situations has

multi-modal

to accept this, and has to pitch analysis at a level that allows worldviews to be surfaced and examined. For many people worldviews are relatively fixed; but they can change over time. Sometimes a dramatic event can change them very quickly.

- All problematical situations, as well as containing different worldviews, have a second important characteristic. They always contain people who are trying to act purposefully, with intention, not simply acting by instinct or randomly thrashing about – though there is always plenty of that too in human affairs.
- The previous two points – the existence of conflicting worldviews and the ubiquity of would-be purposeful action – lead the way to tackling problematical situations. They underpin the SSM approach, a process of inquiry which, through social learning, works its way to taking ‘action to improve’. Its shape is as follows:
 1. Find out about both the problematical situation and the characteristics of the intervention to improve it: the issues, the prevailing culture and the disposition of power within the overall situation (its politics). Ways of doing these things are provided.
 2. From the finding out, decide upon some relevant purposeful activities, relevant that is to exploring the situation deeply, and remembering that the ultimate aim is to define and take ‘action to improve’. Express these relevant purposeful activities as activity models, each made to encapsulate a declared worldview, the model being a cluster of linked activities which together make up a purposeful whole. (For example, one model could express in terms of activities the notion ‘prison’ as if it were only ‘a punishment system’, another could express it as ‘a rehabilitation system’.) Such models never describe the real world, simply because they are based on one pure worldview. They are devices, or tools, to explore it in an organized

way. Techniques developed.

3. Use the models a tion. This provide both the situation surface worldview
4. In the course of t of the ‘finding ou now is to find ch models) but also situation with it process of seeki That is to say, it situation which nevertheless *live* wish that they w affairs, stimulate

- The elements (1) to (sarily been described the cycle as learning which it has been lai ent that it is essential carried out by people outside ‘expert’, thou
- Taking action to im change that situatio again. In any case th events and new idea static. In this sense It ultimately offers

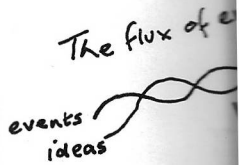
way. Techniques for building and using such models have been developed.

3. Use the models as a source of questions to ask of the real-world situation. This provides a coherent structure to a discussion or debate about both the situation and how it might be changed, a discussion which will surface worldviews and generate ideas for change and improvement.
 4. In the course of the discussion, continually bring together the results of the 'finding out' in (1) and the ideas for change in (3). The purpose now is to find changes which are both arguably *desirable* (given these models) but also culturally *feasible* for these people in this particular situation with its particular history, culture and politics. This is a process of seeking accommodations between different worldviews. That is to say, it is a process of finding versions of the to-be-changed situation which different people with conflicting worldviews could nevertheless *live with*. (Don't expect the worldviews to go away, nor wish that they would. Clashing worldviews, always present in human affairs, stimulate energy and ideas for change.)
- The elements (1) to (4) above constitute a *learning cycle*. They have necessarily been described linearly here but in use there is much iteration within the cycle as learning occurs. It is never followed in the flat-footed way in which it has been laid out here for explanatory purposes. Also it is apparent that it is essentially a group process leading to *group learning*. It is best carried out by people in the problematical situation itself, not left to an outside 'expert', though knowledgeable people can facilitate the process.
 - Taking action to improve a problematical situation will of course itself change that situation, so that the learning cycle could in principle begin again. In any case the changing flux of everyday life will itself bring new events and new ideas, so that no human situation could ever be rendered static. In this sense SSM's learning cycle can be seen as never-ending. It ultimately offers a way of continuously managing any ongoing

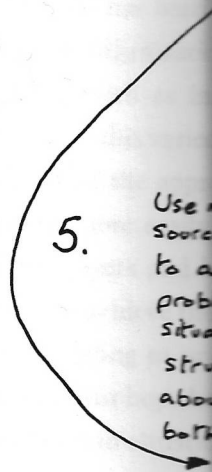
human situation. It does this by helping understanding of complex situations, encouraging multiple perspectives to be taken into account, and bringing rigour to processes of analysis, debate and taking 'action to improve'.

The seven points made above are presented pictorially in the following figure. (A diagram summarizing SSM's process of inquiry is also presented at the end of Chapter 6. That one refers to various tools and techniques described in Chapter 2.)

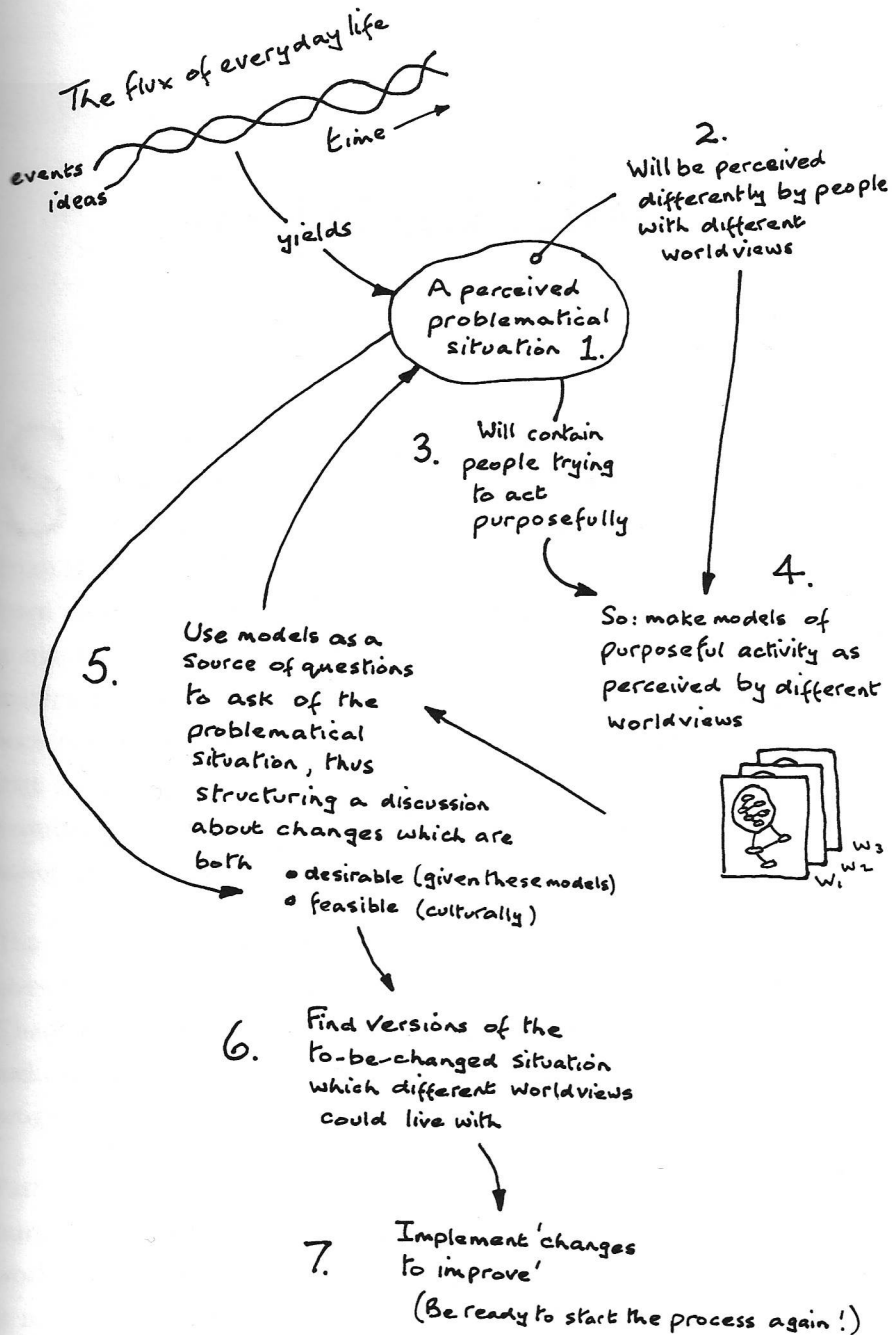
The flux of e
events
ideas

A handwritten wavy line diagram consisting of two overlapping, horizontal, wave-like shapes. The text "The flux of e" is written above the top wave, "events" is written to the left of the middle, and "ideas" is written below the bottom wave.

5. Use sources to a problem situation structure about both

A large, hand-drawn arrow that starts as a curved line on the left, loops around, and points to the right. The number "5." is written inside the loop, and the text "Use sources to a problem situation structure about both" is written along the right side of the arrow's shaft.

6



SSM's cycle of learning for action

1

A Skeleton Account of SSM

What is SSM?

The aim of the work which led to the development of Soft Systems Methodology (SSM) was to find a better way of dealing with a kind of situation we continually find ourselves facing in everyday life: a situation about which we have the feeling that 'something needs to be done about this'. We shall call such situations 'problematical', rather than describing them as 'problem situations', since they may not present a well-defined 'problem' to be 'solved' out of existence – everyday life is more complex than that! A company might feel that it needs to stimulate sales, perhaps by introducing a new product; or should they bid for the equity of a smaller rival? A university may feel that its student intake is too biased towards students from middle-class homes. What are the implications of changing that? A government may struggle to define legislation which would increase the feeling of security on the streets, given the threat of terrorism, without diminishing civil liberties. A local council may be receiving complaints that the delivery of its services is not sufficiently 'citizen-friendly'. What should it do? A head teacher may wonder how to decide whether to take on the responsibility for providing school meals (the school benefiting from any surplus generated) or to leave that function to the local education authority. An individual may develop a sense of unease about the future viability of the firm he or she works for, and wonder whether to look for a job elsewhere. All these are 'problematical situations'. They could be tackled in various ways: by appealing

to previous experience; intuitively; by randomly thrashing about (never a shortage of that in human situations); by responding emotionally; or they could be addressed by using SSM.

So what is it? It is an organized, flexible process for dealing with situations which someone sees as problematical, situations which call for action to be taken to improve them, to make them more acceptable, less full of tensions and unanswered questions. The 'process' referred to is an organized process of thinking your way to taking sensible 'action to improve' the situation; and, finally, it is a process based on a particular body of ideas, namely *systems* ideas.

That these ideas have proved themselves to be useful in dealing with the complexity of the social world is hardly surprising. Social situations are always complex due to multiple interactions between different elements in a problematical situation as a whole, and systems ideas are fundamentally concerned with the *interactions* between parts of a whole. So it is systems ideas which help to structure the thinking. (However, the way systems ideas are used within SSM is fundamentally different from the way they inform the various earlier systems approaches developed in the 1950s and 1960s, as we shall see below.)

In order to ensure that the previous two paragraphs are clear, we need to unpack them somewhat, and say a little more about the crucial elements within them, if this chapter is to fulfil its aim of presenting a broad-brush account of SSM as a whole. Four elements in the paragraphs above will be expanded: 'everyday life and problematical situations'; 'tackling such situations'; a 'flexible process', and 'the use of systems ideas'.

Everyday Life and Problematical Situations

As members of the human tribe we experience everyday life as being quite exceptionally complex. We feel ourselves to be carried along in an onrushing

turbulent stream, a flux of happenings, ideas, emotions, actions, all mediated through the slippery agency of language, all continually changing. Our response to our immersion in this stream is not simply to experience it. Beyond that, we have an innate desire to try to see it, if we can, as *meaningful*. We *attribute* meaning to it – the ability to do this being one of the characteristics which marks us out as human. Part of this *meaning attribution* is to see chunks of the ongoing flux as 'situations'. Nothing is intrinsically 'a situation'; it is our perceptions which create them as such, and in doing that we know that they are not static; their boundaries and their content will change over time. Some of the situations we perceive, because they affect us in some way, cause us to feel a need to tackle them, to do something about them, to improve them.

Tackling Problematical Situations

As we tackle a situation we see as problematical, we are intervening in order to take action intended to bring about improvement. In order to do that sensibly we need to have a clear idea of what it is we are intervening in. This means having a clear view of the nature of the flux which constitutes everyday life. We have already described it as complex, changing, and having multiple strands: events, ideas, emotions, actions. To this we can add an answer to the question: What then happens when we intervene in a part of the flux seen as a problematical situation?

When we interact with real-world situations we *make judgements about them*: are they 'good' or 'bad', 'acceptable' or 'unacceptable', 'permanent' or 'transient'? Now, to make any judgement we have to appeal to some criteria or standards, these being the characteristics which define 'good' or 'bad' etc. for us. For example, an 'eco-warrior' would judge any economic activity 'good' only if it met the environmentalists' criteria for 'good', namely 'environmentally friendly' and 'sustainable'. A 'capitalist' would see an economic activity as 'good' if it were 'profitable'. And where do such criteria come

from? They will be formed partially by our genetic inheritance from our parents – the kind of person we are innately – and, most significantly, from our previous experience of the world. Over time these criteria and the interpretations they lead to will tend to firm up into a relatively stable outlook *through which* we then perceive the world. We develop ‘worldviews’, built-in tendencies to see the world in a particular way. It is different worldviews which make one person ‘liberal’, another ‘reactionary’. Worldviews cause one observer’s ‘terrorism’ to be another’s ‘freedom fighting’. Such worldviews are relatively stable but can change over time. Thus a paranoid person whose worldview is ‘this hostile world owes me a living’ might become a more integrated member of society as a result of experiencing love and generosity.

This concept of worldview (the German *Weltanschauung* being the best technical word for it) is the most important concept in understanding the complexity of human situations, and indeed, the nature and form of SSM.

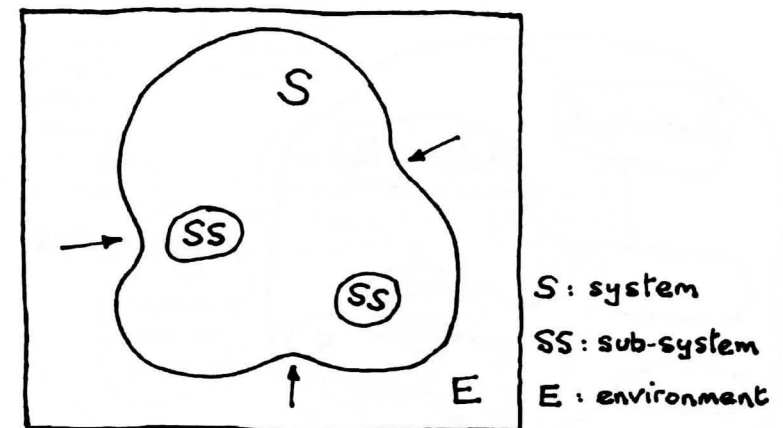
A Flexible Process

It is obvious from the argument so far that any approach able to deal with the changing complexity of real life will have to be flexible. It could never be reduced to a sequence of steps, which might be handed over to an intelligently programmed robot. It needs to be flexible enough to cope with the fact that every situation involving human beings is unique. The human world is one in which nothing ever happens twice, not in *exactly* the same way. This means that an approach to problematical human situations has to be a methodology rather than a method, or technique. A methodology, as the word indicates, is a logos of method; that is to say it is a set of ongoing principles which can be adapted for use in a way which suits the specific nature of each situation in which it is used. SSM provides a set of principles which can be both adopted and adapted for use in any real situation in which people are intent on taking action to improve it.

The Use of Systems Ideas

As stated above, systems ideas concern interaction between parts which make up a whole; also, the complexity of real situations is always to a large extent due to the many interactions between different elements in human situations. So it is not surprising that systems ideas have some relevance to dealing with real-world complexity (though they are only very rarely useful in *describing* that complexity).

The core systems idea or concept is that of an adaptive whole (a ‘system’) which can survive through time by adapting to changes in its environment. The concept is illustrated in Figure 1.1. A system S receives shocks from



Survival of S through time requires :

- o communication processes
- o control processes
- o structure in layers
- o emergent properties of S as a whole

Figure 1.1 The core systems concept: an adaptive whole

its changing environment E. If it is to survive it requires *communication processes* (to know what is going on) and *control processes* (possible adaptive responses to the shocks). Also, the system may contain sub-systems SS, or may itself be seen by a different observer as only a sub-system of some wider system. The idea of a *layered structure* is thus fundamental in systems thinking. Finally, what is said to be a system must have some properties as a single whole, so-called *emergent properties*. (Thus the parts of a bicycle, when assembled correctly, and only then, produce a whole which has the emergent property of being a vehicle, the concept 'vehicle' being meaningful only in relation to the whole.) These four italicized phrases represent the core of systems thinking. So how can it be used here?

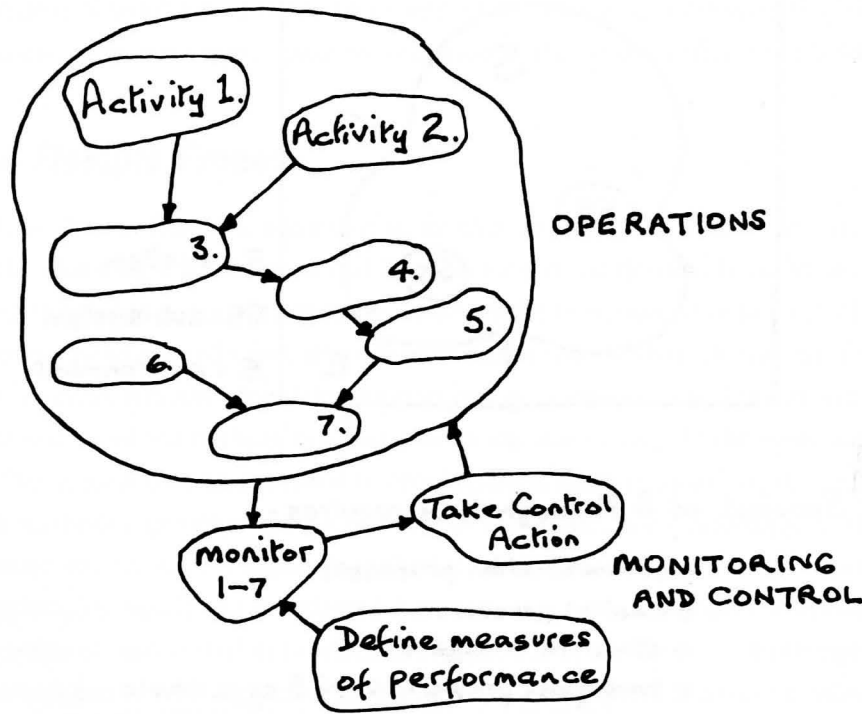


Figure 1.2 The general form of a purposeful activity model

The relevance of this kind of thinking to SSM emerged when it was realized that every single real-world problematical situation, whether in a small firm making wheelbarrows, a multi-national oil company, or in the National Health Service (which employs more than a million people) has one characteristic in common. All such situations contain people trying to act *purposefully*, not simply acting by instinct or splashing about at random. From this observation comes the key idea of *treating purposeful action as a system*. A way of representing purposeful action as a system, i.e. an adaptive whole (in line with Figure 1.1) was invented. Figure 1.2 shows its general form.

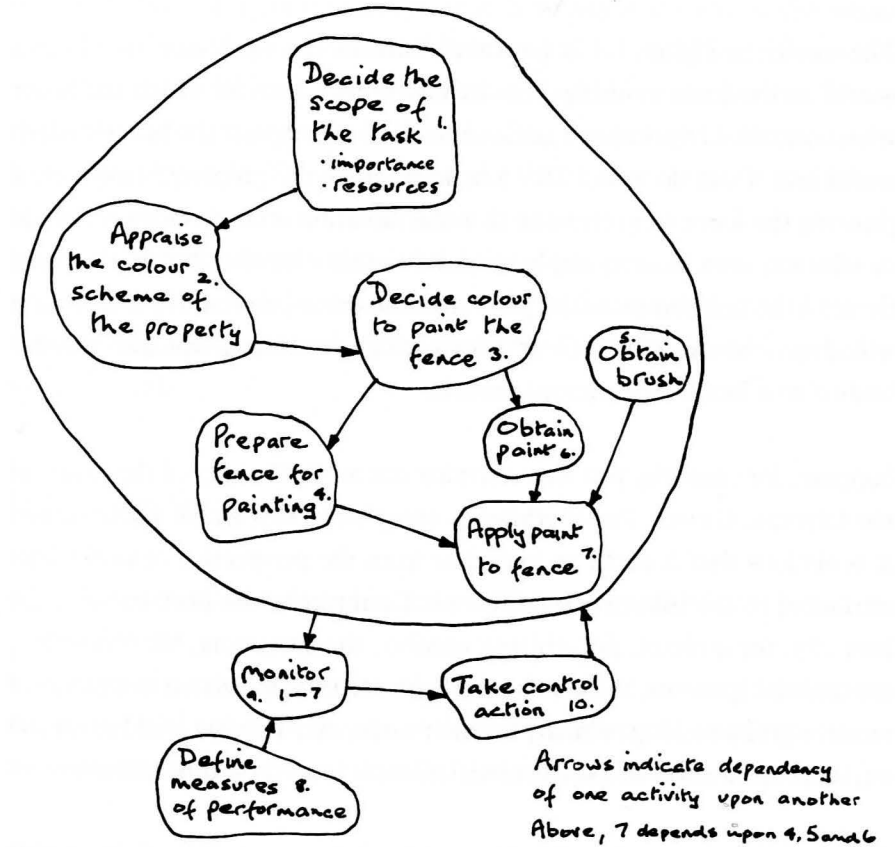


Figure 1.3 A simple example of an activity model: a system to paint the garden fence by hand painting

A logically linked set of activities constitute a whole – its emergent property being its purposefulness. The activities concerned with achieving the purpose (the operations) are monitored against defined measures of performance so that adaptive control action (to make changes) can be taken if necessary.

Figure 1.3 shows a trivial example to illustrate the concept. With regard to Figure 1.2, the ‘measure of performance’ might be the degree to which fence painting enhances the appearance of the property or, perhaps, ‘good’ or ‘bad’ might be defined according to whether or not the neighbours complain about it. This model, then, is a ‘purposeful activity model’.

The model in Figure 1.3 is essentially within the worldview of whoever would do the fence painting. It is an instrumental model which spells out what is entailed in painting a garden fence. It could express the householder’s worldview: ‘I can do useful DIY jobs to improve my property.’ However, if painting the fence were an issue in a real situation other worldviews would be relevant, even in an example as trivial as this – for example, in this case, those of the neighbours or the partner of the fence-painter. In general there will always be a number of worldviews which could be taken into account leading to a number of relevant models.

Suppose, for example, you were carrying out an SSM study of the future of the Olympic Games. For anything as complex as this global phenomenon it is obvious that it could be looked at from the perspective of worldviews attributed to the International Olympic Committee, the host country, the host city, the athletes, the athletes’ coaches, the spectators, hot dog sellers, commercial sponsors, those responsible for security, television companies, a terrorist group seeking publicity for their cause, etc. This list could go on and on; there could never be a single model relevant to all these different interests.

An important consequence flows from this: these purposeful activity models *can never be descriptions* of (part of) the real world. Each of them expresses

one way of looking at and thinking about the real situation, and there will be multiple possibilities. So how can such models be made useful? The answer is to see them as *devices* (intellectual devices) which are a source of *good questions to ask about the real situation*, enabling it to be explored richly. For example, we could focus on the differences between a model and the situation, and ask whether we would like activity in the situation to be more, or less, like that in the model. Such questioning organizes and structures a discussion/debate about the real-world situation, the purpose of that discussion being to surface different worldviews and to seek possible ways of changing the problematical situation for the better. This means finding an accommodation, that is to say a version of the situation which different people with different worldviews could nevertheless live with. Given the different worldviews which will always be present in any human situation, this means finding possible changes which meet two criteria simultaneously. They must be arguably *desirable*, given the outcomes of using the models to question the real situation, but must also be culturally *feasible* for these particular people in this particular situation with its *unique history* and the *unique narrative* which its participants will have constructed over time in order to make sense of their experience. Figure 1.4 illustrates this.

In summary, then, we have:

- a problematical real-world situation seen as calling for action to improve it;
- models of purposeful activity *relevant* to this situation (not describing it);
- a process of using the models as devices to explore the situation;
- a structured debate about desirable and feasible change.

This gives the bare bones of the process of SSM, whose shape can now be described.




Figure 1.4 SSM's basic process

What is the SSM Process?

The SSM process takes the form of a cycle. It is, properly used, a cycle of learning which goes from finding out about a problematical situation to defining/taking action to improve it. The learning which takes place is social learning for the group undertaking the study, though each individual's learning will be, to a greater or lesser extent, personal to them, given their different experiences of the world, and hence the different worldviews which they will bring to the study. Taking action as a result of the study will of course change the starting situation into a new situation, so that in principle the cycle could begin again (a relevant system then being 'a system to make




Figure 1.5 The iconic representation of SSM's learning cycle

these changes'). SSM is thus not only a methodology for a specially set-up study or project; it is, more generally, a way of managing any real-world purposeful activity in an ongoing sense.

The SSM cycle is shown in Figure 1.5, which eventually emerged as its classic representation. It contains four different kinds of activity:

1. Finding out about the initial situation which is seen as problematical.
2. Making some purposeful activity models judged to be relevant to the situation; each model, as an intellectual device, being built on the basis of a particular pure worldview.
3. Using the models to question the real situation. This brings structure to a discussion about the situation, the aim of the discussion being to find changes which are both arguably desirable and also culturally feasible in this particular situation.
4. Define/take the action to improve the situation. Since the learning cycle is in principle never-ending it is an arbitrary distinction as to whether the end of a study is taken to be defining the action or actually carrying

it out. Some studies will be ended after defining the action, some after implementing it.

This description of the cycle as activities (1) to (4) may give a false impression that we are describing a sequence of steps. Not so. Although virtually all investigations will be initiated by finding out about the problematical situation, once SSM is being used, activity will go on simultaneously in more than one of the 'steps'. For example, starting the organized discussion about the situation (3) will normally lead not only to further new finding out (1), perhaps focused on aspects previously ignored, but also to further new choices of 'relevant' systems to model. In real life, an investigation which sets out narrowly to improve, say, aspects of product distribution in a manufacturing company's distribution department, may well later sweep in issues concerning, perhaps, communications between production and marketing departments. Figure 1.6 illustrates a typical pattern of activity of the kind which emerges as an investigation digs deeper.

Figure 1.6 shows an on-going 'finding out' activity, three bursts of model building, discussion fed by both the models and the finding out, which itself leads to more finding out and more modelling. The final (fourth) burst of modelling shown here as an example follows from defining the 'action to improve' and would consist of purposeful activity models relevant to carrying out the action agreed.




Figure 1.6 A typical pattern of activity during an SSM investigation

Finally, in describing the SSM cycle, we could add (though this is really a point from the end of this book) that as users of SSM become more sophisticated they treat Figure 1.5 not at all as a prescription to be followed, but as a model to make sense of their experience as they mentally negotiate their way through the problematical situation.

What Can SSM be Used for?

The application area for SSM is very broad. This is not due to megalomania on the authors' part. Rather it stems from the wide applicability of two key ideas behind SSM. One of these is to create a process of *learning your way* through problematical situations to 'action to improve' – a very general concept indeed. The other is the idea that you can make sure this learning is organized and structured by using, as a source of questions to ask in the real situation, models (systems models) of purposeful activity. This is because every real-world situation contains people trying to act purposefully, intentionally. It is the sheer generality of purposeful action – the core of being human – that makes the area in which SSM can be used so huge.

In Part Two, the stories of SSM use come from all sizes of company from small firms to large corporations, from organizations in both private and public sectors, including the National Health Service. Chapter 4 describes uses of SSM in the world of information systems and information technology, where it is much used. This derives from the fact that for any purposeful activity model (Figure 1.3 being a noddy example) you can ask of each activity: What information would support doing this activity? And what information would be generated by doing it? Since information is what you get when you attribute meaning to data in a particular context, and meaning attribution depends upon worldview, SSM's strong emphasis on worldview explains its relevance to this field.

In summary, SSM can be used in any human situation which entails thinking about acting purposefully, and is especially useful in any situation in which it is helpful to lift the level of discussion from that of everyday opinions and dogma to that level at which you are asking: What taken-as-given worldview lies behind these assertions of opinion?

Is SSM Mature?

Obviously it is never possible to claim that the development of any approach to human inquiry is 'finished', though some features of any such process may become so taken-as-given as to appear permanent. For example, in the inquiry process of natural science, if you are testing a new drug you give some patients the drug while others receive a placebo. The difference between the group ingesting the drug and the so-called 'control' group taking the placebo tells you what effects the drug produces (given a statistically significant sample size). This pattern would seem to be a permanent feature of scientific experiment. In applied social science, where SSM sits, the situation is less definite. Nevertheless, after hundreds of studies the core processes of SSM do now appear to be well-established, though the application area continues to expand. In the early days each significant study was likely to cause some rethinking of the process itself; but such changes became increasingly rare over the 30-year development period. We now regard it as a mature process.

The most recent addition to the literature about its development describes the use of SSM both in relation to the perceived *content* of the situation in question – SSM (c) – and in relation to the *process* of carrying out the inquiry itself – SSM (p) – (described in Chapter 2). This is in a paper published in 2006. But this is a case of the literature lagging behind practice, as these twin uses of SSM have been recognized and exploited by those developing the approach since the early 1980s.

So SSM is now considered mature enough to justify writing this book.

How was SSM Created?

The classic way of doing research comes from natural science: set up a hypothesis and then test it experimentally. It is not easy to transfer this model of research to the gloriously rich social and human arena, though strenuous efforts to do that have been made over many years. SSM was developed using an alternative model of research, one more suitable for social research at the level of a situation, group or organization, namely 'action research'. In this kind of research you accept the great difficulty of 'scientific' experimental work in human situations, since each human situation is not only unique, but changes through time and exhibits multiple conflicting worldviews. Hence the pattern for the action researcher is to enter a human situation, *take part* in its activity, and use that experience as the research object. In order to do that, to do more than simply return from the research with a one-off story to tell, it is necessary to declare in advance the intellectual framework you, the researcher, will use to try to make sense of the experience gained. Given such an explicit framework, you can then describe the research experience in the well-defined language of the framework. This makes it possible for anyone outside the work to 'recover' it, to see exactly what was done and how the conclusions were reached. This 'recoverability' requirement is obviously not as strong as the 'repeatability' criterion for scientific findings within natural science. But then, human situations are very much more complex than the phenomena studied in physics and chemistry labs! It is the declared framework and the recoverability criterion which clearly separate accounts of well-organized action research from novel writing – which, alas, too much published social research resembles.

In the action research which produced SSM the initial declared framework was the Systems Engineering approach developed by the Bell Telephone Company from their own case histories. Systems Engineering (SE) is a process

of naming a 'system' (assumed to be some complex object which exists or could exist in the real world), defining its objectives, and then using an array of techniques developed in the 1950s and 1960s to 'engineer' the system to meet its objectives. This framework was rapidly found to be poverty-stricken when faced with the complexity of human situations. It was too thin, not rich enough to deal with fizzing social complexity.

The SE framework was modified (and enriched) in the light of and in direct response to real-life experiences. Eventually, we had in our hands an adequately rich framework, but it was far removed from the starting point in SE. It became known as Soft Systems Methodology. It then took some time for even its pioneers to realize just how radical the shift had been from SE to SSM. Having introduced the notion of 'worldview' – essential in dealing with human social complexity – we were thereafter thinking of systems models not as descriptions of something in the real world but simply as devices (based on worldview) to organize a debate about 'change to bring about improvement'. That was the key step in finding our way to SSM. This important shift in thinking is not abstruse, but it turns out to be very difficult for many people to grasp, simply because everyone is so used to the casual everyday-language use of the word 'system'. In ordinary talk we constantly refer to complex chunks of the everyday world as systems, even though they do not come close to meeting the requirements of that concept. We speak of 'the education system', 'health-care systems', 'the prison system', etc. using the word 'system' simply to indicate a chunk of reality which seems to be very complex but is, in some vague sense, a whole, something which might be better 'engineered'. Figure 1.7 gives a visual indication of the shift in thinking as SE was transformed into SSM.

At the starting point (SE) in Figure 1.7 (which ignores worldviews), 'systems' are names for things in the world which, given precise objectives, can be engineered to achieve them. At the end point (which accepts different world-




Figure 1.7 The shift in thinking entailed in developing SSM

views), 'systems' are devices used in a learning process to define desirable and feasible 'action to improve'.

Once the end point in Figure 1.7 was reached, and the SSM framework had been established, it was further developed, modified and honed in a few hundred new experiences. Out of this came a model which captures all of these developmental experiences. The model, known as the LUMAS model, is shown in Figure 1.8. (It is in fact a generic model for making sense of any real-world application of any *methodology*, remembering that that word covers a set of principles which need to be embodied in an application tailored to meet the unique features of a particular situation.)

LUMAS stands for Learning for a User by a Methodology-informed Approach to a Situation. In order to 'read' this model, start from the user (U) in the centre. He or she, perceiving a problem situation (S) and appreciating the methodology (M), tailors the latter to the former to produce the specific approach (A) to be used in this situation (S). This not only produces an improved situation but also yields learning (L). This will change the user,

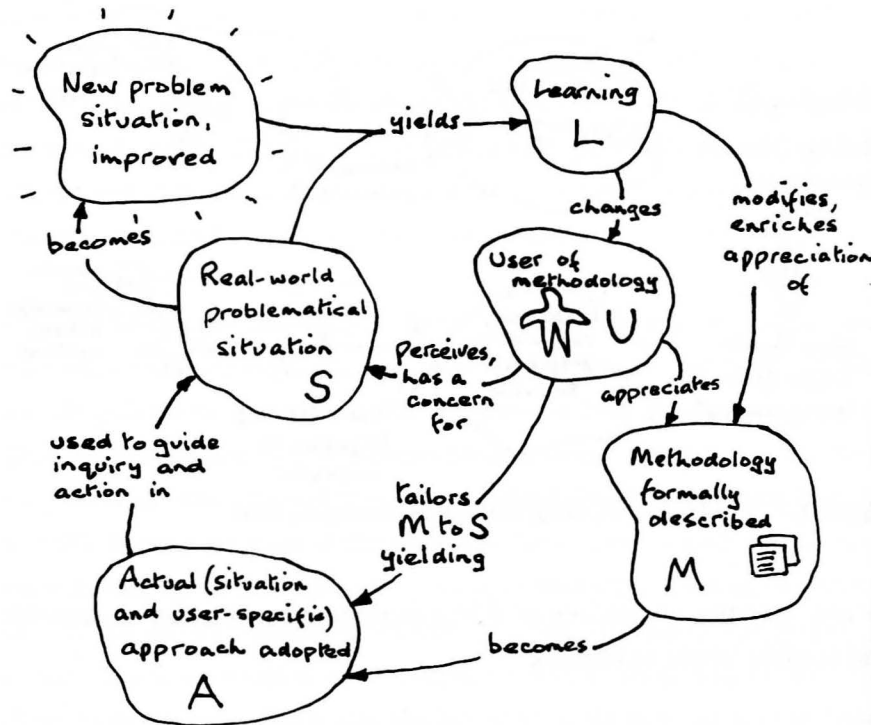



Figure 1.8 The LUMAS model – Learning for a User by a Methodologically-informed Approach to a Situation

who has gained this experience, and may also modify or enrich appreciation of the methodology. Every use of SSM can in principle be described in the language of this model. It is the gradually diminishing activity, over the years, of development occurring along the arrow which links L and M that makes it legitimate to describe SSM as mature.

How Does SSM Differ from Other Systems Approaches?

As described above, changes had to be made to Systems Engineering when it proved too blunt an instrument to deal with the complexity of human

situations. Those changes explain SSM's difference from the other systems approaches developed in the 1950s and 1960s. SE is an archetypal example of what is now known as 'hard' systems thinking. Its belief is: the world contains interacting systems. They can be 'engineered' to achieve their objectives. This is the stance not only of SE; this thinking also underpins classic Operational Research, RAND Corporation 'systems analysis', the Viable System Model, early applications of System Dynamics and the original forms of computer systems analysis. None of these approaches pays attention to the existence of conflicting worldviews, something which characterizes all social interactions. In order to incorporate the concept of worldview into the approach being developed, it was necessary to abandon the idea that the world is a set of systems. In SSM the (social) world is taken to be very



Observer 1 'I spy systems which I can engineer.'

Observer 2 'I spy complexity and confusion; but I can organize exploration of it as a learning system.'

Figure 1.9 The 'hard' and 'soft' systems stances

complex, problematical, mysterious, characterized by clashes of worldview. It is continually being created and recreated by people thinking, talking and taking action. However, our coping with it, our process of inquiry into it, can itself be organized as a learning *system*. So the notion of systemicity ('systemness') appears in the process of inquiry into the world, rather than in the world itself. This shift created 'soft' as opposed to 'hard' systems thinking, the different stances adopted by the two being shown in Figure 1.9, itself another version of Figure 1.7.

This brings us to the end of a skeletal account of SSM as a whole. The next chapter expands on this, describing the techniques used in the cyclic process in detail. Meanwhile it seems worthwhile to try to summarize the broad account of SSM in a couple of sentences.

SSM is an action-oriented process of inquiry into problematical situations in the everyday world; users learn their way from finding out about the situation to defining/taking action to improve it. The learning emerges via an organized process in which the real situation is explored, using as intellectual devices – which serve to provide structure to discussion – models of purposeful activity built to encapsulate pure, stated worldviews.

2 A Fleshed-out Account of SSM

Introduction

The previous chapter has answered the basic question about SSM, namely: What is it? And it has provided some context concerning its development, its application area and its crucial difference from the earlier systems approaches from the 1950s and 1960s. In this chapter the focus is on 'how' rather than 'what': How exactly does the user move through the learning cycle of SSM, shown in Figure 1.5, in order to define useful change? Which techniques for finding out, modelling and using models to question the real situation have shown themselves robust enough to survive in many different circumstances, so that they have become part of the classic approach?

The account here will follow the four basic activities of the broad-brush account (finding out, modelling, using the models to structure debate, and defining/taking action), with the usual reminder that activity in any project using SSM will reflect the kind of pattern shown in Figure 1.6 rather than a stately linear progress.

The SSM Learning Cycle: Finding Out

Four ways of finding out about a problematical situation have survived many tests and become a normal part of using SSM. In the language of SSM they

are known as 'making Rich Pictures' and carrying out three kinds of inquiry, known as 'Analyses One, Two and Three'. These focus, respectively, on the intervention itself, a social analysis (What kind of 'culture' is this?) and a political analysis (What is the disposition of power here?). They will be described in turn.

(Readers anxious to reach the stories of SSM use might turn to the first few case histories described in Part Two, but all the accounts there use the terms and language carefully defined here, so a little patience might well be worthwhile!)

Making Rich Pictures

Entering a real situation in order first to understand it and then to begin to change it in the direction of 'improvement' calls for a particular frame of mind in the user of SSM. On the one hand the enquirer needs to be sponge-like, soaking up as much as possible of what the situation presents to someone who may be initially an outsider. On the other hand, although holding back from imposing a favoured pattern on the first impressions, the enquirer needs to have in mind a range of 'prompts' which will ensure that a wide range of aspects are looked at. Initially two dense and cogent questions were used as a prompt:

- What resources are deployed in what operational processes under what planning procedures within what structures, in what environments and wider systems, by whom?
- How is resource deployment monitored and controlled?

Certainly, if you can answer these questions you know quite a lot about the situation addressed. But these questions did not survive as a formal part of SSM. (The problem with them is that when they were formulated, in the early days of SSM development, the thinking of the pioneers had not sufficiently divorced itself from thinking of the world as a set of systems.

The questions imply intervention in some real-world system – hence the references to 'wider systems' and to monitoring and control – rather than the intervention being addressed to *a situation*.) The questions would no doubt have been changed eventually as the true nature of SSM was realized. However, what happened instead was that the questions were dropped because the phrase 'rich picture' quickly moved from being a metaphor to being a literal description of an account of the situation *as a picture*.

The rationale behind this was as follows. The complexity of human situations is always one of multiple interacting relationships. A picture is a good way to show relationships; in fact it is a much better medium for that purpose than linear prose. Hence as knowledge of a situation was assembled – by talking to people, by conducting more formal interviews, by attending meetings, by reading documents, etc. – it became normal to begin to draw simple pictures of the situation. These became richer as inquiry proceeded, and so such pictures are never finished in any ultimate sense. But they were found invaluable for expressing crucial relationships in the situation and, most importantly, for providing something which could be tabled as a basis for discussion. Users would say: 'This is how we are seeing your situation. Could we talk you through it so that you can comment on it and draw attention to anything you see as errors or omissions?'

In making a Rich Picture the aim is to capture, informally, the main entities, structures and viewpoints in the situation, the processes going on, the current recognized issues and any potential ones.

Here is a real-world problematical situation described in a paragraph of prose:

The newly appointed headteacher of an 11s-to-18s school, which has overspent its budget in the last year or two, finds herself, in her first term, facing an issue concerning the provision of school meals. Currently these are provided by the county education authority

through their catering services company, the contract being renewed annually. A member of that company who is leaving to set up her own catering company urges the headteacher to make a contract with her instead of the county, suggesting the school could save money on this. Some staff members agree with this, others want to stick with the status quo. Some parents, alerted by a national debate about school meals, want more nutritious meals as long as they don't cost more. Pupils say: 'We like burgers and chips.' The school governors are discussing this issue; the Chairman, himself MD of a catering company, is urging the headteacher to be entrepreneurial and to take on responsibility for the provision of school meals, believing this could be profitable for the school.

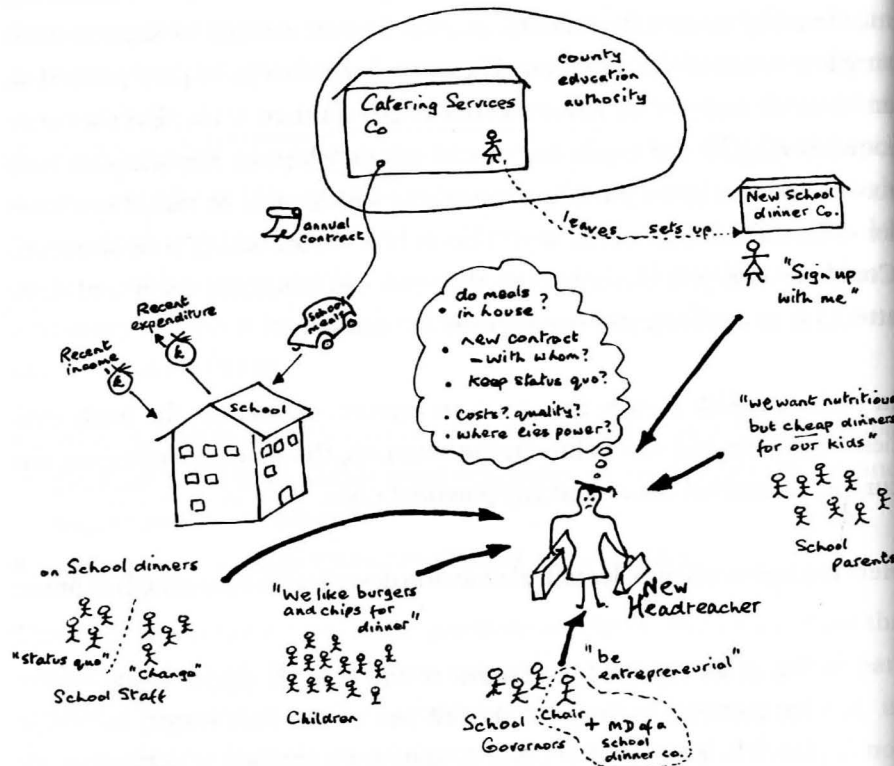


Figure 2.1 A Rich Picture of the situation described in the text

Figure 2.1 represents this situation in a Rich Picture. Our point is that this picture is a more useful piece of paper than the prose account. It could lead to a better-than-usual level of discussion because not only can it be taken in as a whole, but also it displays the multiple *relationships* which the headteacher has to manage, not just immediately, but through time. That is the power of such pictures, though we have to remember that however rich they are they could be richer, and that such pictures record a snapshot of a situation which will itself not remain static for very long. Wise practitioners continually produce such pictures as an aid to thinking. They become a normal way of capturing impressions and insights.

Carrying Out Analysis One (the Intervention Itself)

Whenever SSM is used to try and improve a problematical situation three elements – the methodology, the use of the methodology by a practitioner and the situation – are brought together in a particular relationship, namely that shown in Figure 2.2. The practitioner will adapt the principles and




Figure 2.2 The three elements in any SSM investigation

techniques of the methodology to organize the task of addressing and intervening in the situation, aiming at taking action to improve it. In developing SSM, this process was organized in a sequence of real situations, and it was quickly found useful to think about Figure 2.2 in a particular way. Three key roles were always present:

1. There was some person (or group of persons) who had *caused the intervention to happen*, someone without whom there would not be an investigation at all – this was the role ‘client’.
2. There was some person (or group of persons) who were *conducting the investigation* – this was the role ‘practitioner’.
3. Most importantly, whoever was in the practitioner role could choose, and list, a number of people who could be regarded as being *concerned about or affected by the situation and the outcome* of the effort to improve it – this was the role ‘owner of the issue(s) addressed’.

It is important to see why these are named as ‘roles’ rather than particular people. It is because one person (or group) might be in more than one role. For example, if the headteacher in the Rich Picture (Figure 2.1) were to herself carry out an SSM-based study of her complex situation, she would not only be both ‘client’ and ‘practitioner’, she would also be one of the people in the list of ‘issue owners’ who care about the outcome. Sometimes a manager who causes an intervention to take place delegates detailed involvement in it to others, and so is only in the role ‘client’. In this case the person(s) in the ‘practitioner’ role needs to take steps to ensure that the ‘client’ is kept informed about the course of the intervention so that the outcome when it emerges does not come as a big surprise. In every case the ‘practitioner’ needs to make sure that the resources available to carry out the investigation are in line with its ambition. Don’t undertake a study of ‘the future of the A-level examination in British education’ if you have only got one man and a boy to work on it between now and next Thursday.

SSM’s ‘Analysis One’, then, consists of thinking about the situation displayed in Figure 2.2 in the way shown in Figure 2.3, asking: Who are in the roles ‘client’ and ‘practitioner’? and Who could usefully be included in the list of ‘issue owner’?

Much learning came out of the simple thinking which led to this ‘Analysis One’. For example, it was always useful to think about the client’s aspirations for the intervention. They should always be taken seriously but should not be the sole focus of the work done. Thus, the person(s) in the ‘client’ role should be in the list of possible ‘issue owners’ but should very definitely not be the only one in the list. In this connection it was interesting to hear a

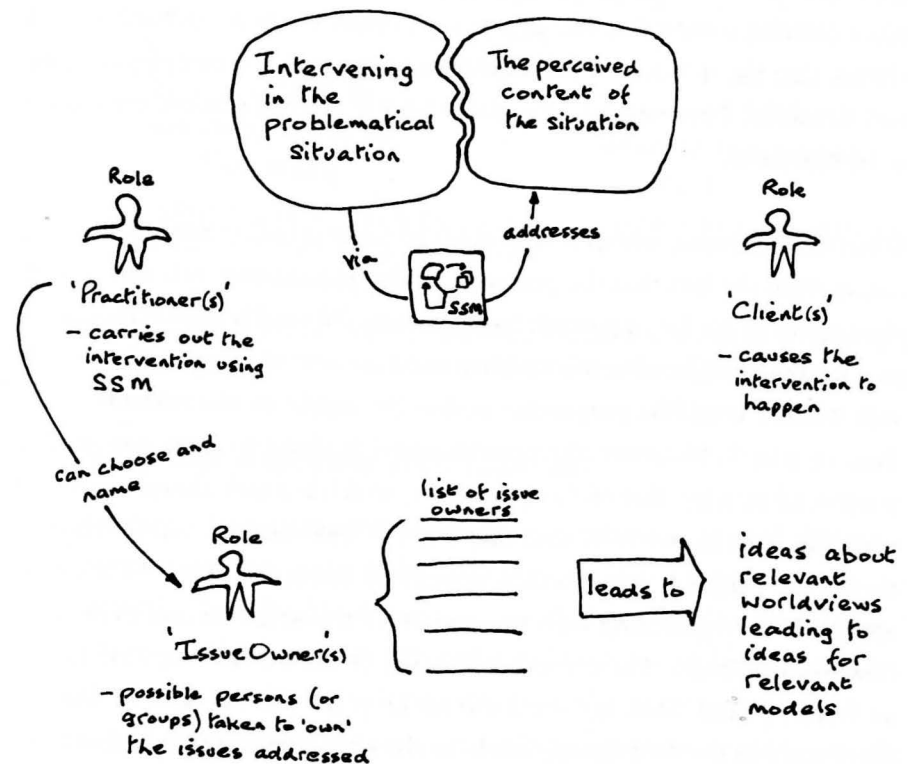


Figure 2.3 SSM's Analysis One

"Issue owner"

senior manager from the RAND Corporation declare, some years ago, 'The RAND analyst places his or her expertise at the disposal of a real-world decision-taker who has to be a legitimate holder of power.' In the language of Figure 2.3 this was to declare that for RAND the client is the issue owner, full stop. This cuts off all the richness which comes from the practitioner compiling a list of persons or groups who *could be taken to be* issue owners; for it is that list which introduces multiple worldviews. They in turn open up the chance of a richness of learning at a deep level for all involved in the intervention, leading, perhaps, to major change. The RAND manager's statement would define the practitioner as only a servant to the legitimately powerful. In the situation shown in Figure 2.1, for example, 'issue owners' might include: the headteacher; the school governors, staff and pupils; parents; the county education authority and their catering services company; other catering companies, etc. The many worldviews from such a list give a chance that the richness of the inquiry can cope with the complexity of the real situation. They suggest ideas for 'relevant' activity models, ones likely to be insightful.

Some final learning, which is important in understanding SSM as a whole, comes from the fact that the person(s) in the 'practitioner' role can include *themselves* in the list of possible 'issue owners'. Normally SSM is thought of as a means of addressing the problematical content of the situation, which will include would-be purposeful action by people in the situation. It is that, of course. However, the practitioner(s) is about to carry out another purposeful activity, that of *doing the study*, which is a task always associated with the practitioner role. Carrying out the investigation can be thought about, and planned, using models relevant to doing this. Thus SSM can be applied both to grappling with the content of the situation and to deciding how to carry it out. These two kinds of use of the methodology are known as 'SSM (c)' and 'SSM (p)' – c for content, p for process. Use of SSM (p) often leads to the first models made in the course of an intervention being models related to doing the study. This will be illustrated in Part Two

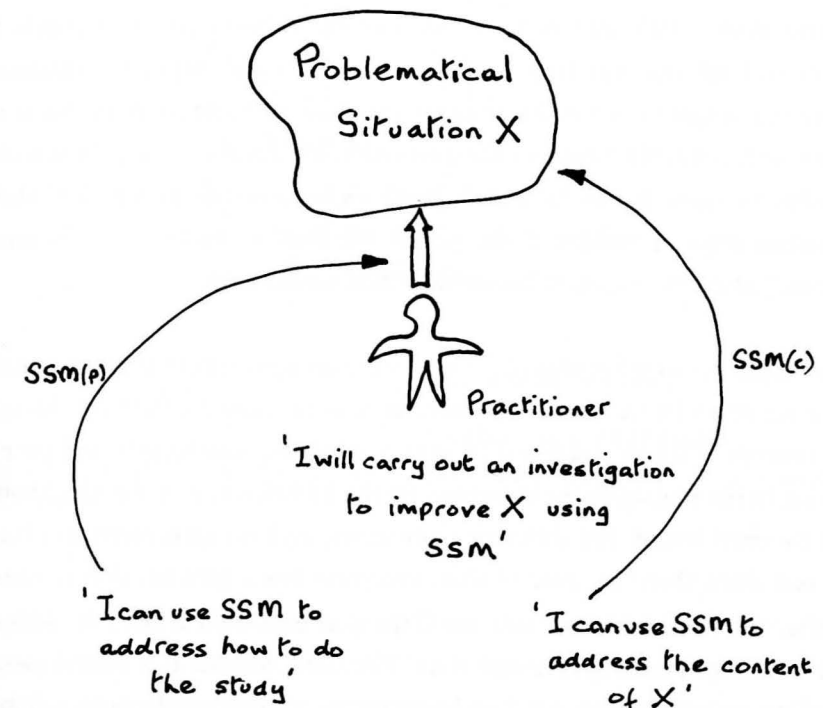


Figure 2.4 SSM(p) concerned with the process of using SSM to do the study and SSM(c) concerned with the problematical content

(Chapter 3, Case 1, Figure 3.1). Meanwhile Figure 2.4 illustrates these two ways of using SSM.

Carrying Out Analysis Two (Social)

It might seem obvious that if you are going to intervene in, and change, a human situation, you ought to have a clear idea about what it is you are intervening in. You should have some sense of what you take 'social reality' to be. However, this is not too obvious! The Management Science field, for example, tries to get by through concentrating almost entirely on the *logic* of situations, even though the motivators of much human action lie outside logic, in cultural norms or emotions. So, if we are to be effective in social

situations, we have to take 'culture' seriously and decide what we mean by it. This is especially important for SSM as an action-oriented approach. If we are to learn our way to practical action which will improve a situation under investigation, then the changes involved in 'improvement' have to be not only arguably desirable but also *culturally feasible*. They need to be possible for these particular people, with their particular history and their particular ways of looking at the world. We have to understand the local 'culture', at a level beyond that of individual worldviews.

This might be straightforward if there were an agreed definition of exactly what we mean by 'culture'. However, there is no agreed definition, though the concept is much discussed by anthropologists, sociologists and people writing in the management literature. By the 1950s, a survey (by Kluckhohn and Kroeber) found 300 different definitions, and no agreement has been reached since then! In spite of that, everyone has a general, diffuse sense of what the word means. If you say 'This is a "can-do" culture', or 'This is a buttoned-up culture', or assert that 'The Civil Service is a punishment-avoiding, rather than a reward-seeking culture' then it will be accepted that you have said something meaningful. To anyone familiar with the society in question, those statements will have conveyed some sense of the 'feel', or 'flavour', of the situation: its social texture. In order to pin down such feelings more firmly, in a way which makes practical sense, SSM makes use of a particular model. This is a model which does not claim the status of rounded theory, but it has proved itself useful in situations from small firms dominated by individuals to large corporations which develop and (partially) impose their own norms.

The model is at the same time simple (you can keep it in your head) but also subtle. It consists of only three elements – roles, norms, values – but the subtlety comes from the fact that none of these elements is static. Each, over time, continually helps to create and modify the other two elements, as shown in Figure 2.5.




Figure 2.5 SSM's model for getting a sense of the social texture of a human situation

Together the three elements help to create the social texture of a human situation, something which will both endure *and* change over time. Consider the three elements in turn.

Roles are social positions which mark differences between members of a group or organization. They may be formally recognized, as when a large organization has, say, a chief executive, directors, department heads, section heads and members of sections. But in any local culture informal roles also develop. Individuals may develop a reputation as 'a boat-rocker', or 'a licensed jester' – someone who can get away with saying things others would suppress. The informal roles which are recognized in a given culture tell you a lot about it.

Norms are the expected behaviours associated with, and helping to define, a role. Suppose you told a friend you were going to meet 'the vice-chancellor of a UK university' next day. If you returned from the meeting and said that the VC sat picking her teeth, with her feet on the table, and was very

foul-mouthed, your friend would be flabbergasted. Such behaviour is way outside the expected behaviour of someone in the role of VC in British society.

- *Values* are the standards – the criteria – by which behaviour-in-role gets judged. In all human groups there is always plenty of gossip related to this. People love to discuss behaviour in role and reach judgements which praise or disparage: ‘He’s a very efficient town clerk who services committees well’; ‘She’s an ineffective vice-chancellor who won’t take decisions.’

It is obvious from these definitions that the three elements – roles, norms, values – are closely related to each other, dynamically, and that they change over time as the world moves on. Anyone who has ever been promoted within an organization will know that occupying the new role changes them, as they adopt a new perspective appropriate to the role. Equally, how they enact the new role will have its effect, in future, on the local norm – the behaviour which people expect from whoever fills that role. The elements also change over time at a macro level. For example, when the authors were growing up in British society the worst role for a young woman to find herself in was to be an unmarried mother. At that time, society judged harshly the behaviour which led to this. Not any more; the social stigma attached to the role has disappeared in the UK over the last 50 years.

So how exactly is the model of linked roles, norms and values in Figure 2.5 used in SSM? At the start of an intervention open a file marked ‘Analysis Two’. Then, every time you interact with the situation – talking to people informally, reading a document, sitting in a meeting, conducting an interview, having a drink in the pub after work – ask yourself afterwards whether that taught you anything about the roles, norms and values which are taken seriously here and characterize this particular group. Record the finding in the ‘Analysis Two’ file. Carry on doing this throughout the engagement, and

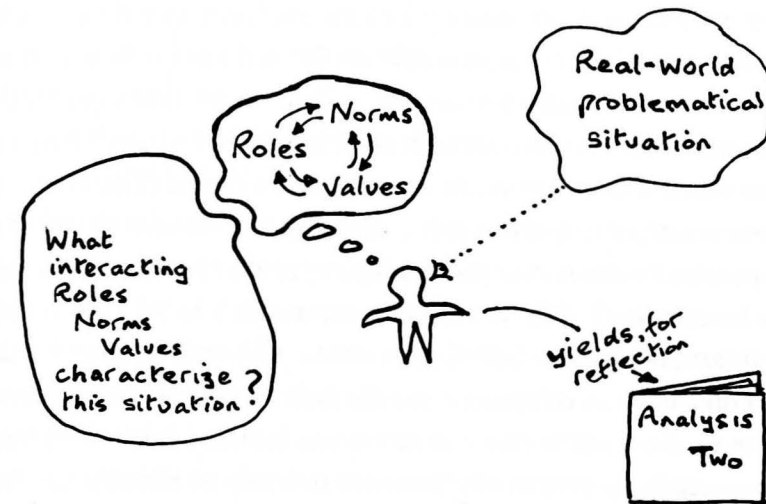


Figure 2.6 SSM's Analysis Two

put a date on every entry so that later on you can recover the progress of your learning, and reflect upon it. Figure 2.6 summarizes Analysis Two.

Carrying Out Analysis Three (Political)

The experienced reader will have noticed that so far in this discussion of ‘Finding Out’ about a problematical situation we have made no mention of the *politics* of a situation, which are always powerful in deciding what does or does not get done. That is the focus of Analysis Three: to find out the disposition of power in a situation and the processes for containing it. That is always a powerful element in determining what is ‘culturally feasible’, politics being a part of culture not addressed directly in the examination of roles, norms and values of Analysis Two.

The ‘political science’ literature contains many models – usually fairly complex ones – which set out to express the nature of politics. The model used in SSM, in Analysis Three, does not come from that literature but from some basic ideas found in the work of the founding father of the field: Aristotle.

Aristotle argues that in any society (for him, the Greek city-state) in which human beings constantly interact, different interests will be being pursued. If the society as a whole is to remain coherent over time, not breaking up into destructive factions, then those differing interests will have to be accommodated; they will never go away. Accommodating different interests is the concern of politics; this entails creating a power-based structure within which potentially destructive power-play in pursuit of interests can nevertheless be contained. This is a general requirement in all human groups which endure, not only in societies as a whole. There will be an unavoidable political dimension in companies, in international sport, in health-care provision, in the local tennis club – in fact in any human affairs which involve deliberate action by people who can hold different worldviews and hence pursue different interests.

- Analysis Three in SSM asks: How is power expressed in this situation? This is tackled through the metaphor of a 'commodity' which embodies power. What are the 'commodities' which signal that power is possessed in this situation? Then: What are the processes, by which these commodities are obtained, used, protected, defended, passed on, relinquished, etc? Figure 2.7 summarizes Analysis Three. The commodities which indicate power in human groups are, of course, many and various. There is a link here to Analysis Two, since occupying a particular role embodies power: the chief constable has more power than a detective sergeant, by virtue of his role. Other common commodities of power include, for example: personal charisma; membership of various committees in organizations; having regular access to powerful role-holders; in knowledge-based settings, having intellectual authority and reputation; having authority to prepare the minutes of meetings – a chore, perhaps, but it gives you some power! Many commodities of power derive from information. Having access to important information, or being able to prevent others from having access to certain information, is a much-used commodity of power in most organizations.

A dramatic example of an unusual commodity of power in a specific SSM project was revealed when two managers in a consultancy company were being interviewed as a pair. They began to disagree with each other and, in a deliberate bit of power-play, one of them suddenly said: 'You say that, but you're NKT; I'm KT.' This local private language within this company referred to those partners who 'knew Tom' and those, more recent joiners, who 'never knew Tom', Tom being the charismatic founder of the company, now deceased. This taught those facilitating this use of SSM that there was an unstated but very real hierarchy here. The KTs, Tom's original disciples, were much more influential than the come-lately NKTs. This indicated that the only changes likely to be culturally feasible in this situation would be those supported by the KTs, whose power stemmed from their association with the charismatic Tom. This is an interesting example of a commodity of power which would gradually fade over time. And this itself reminds us that, as with Analysis Two, Analysis Three deals with elements which are continually being redefined as life moves on.




Figure 2.7 SSM's Analysis Three

The way of doing this analysis echoes that of Analysis Two: open a file and record in it – with a date – any learning gained about power and the processes through which it is exercised. Do this, and reflect upon it, over the whole course of an investigation.

The SSM Learning Cycle: Making Purposeful Activity Models

As explained in Chapter 1, in order to ensure that learning can be captured, SSM users create an *organized* process of enquiry and learning. They do this by making models of purposeful activity and using them as a basis for asking questions of the real-world situation. This kind of model is used because every human situation reveals people trying to act purposefully. Since each model is built according to a declared single worldview (e.g. ‘the Olympic Games from the perspective of the host city’) such models could never be descriptions of the real world. They model *one way* of looking at complex reality. They exist only as devices whose job is to make sure the learning process is not random, but organized, one which can be recovered and reflected on. This section describes how to make these devices.

The task is to construct a model of a purposeful ‘activity system’ viewed through the perspective of a pure, declared worldview, one which has been fingered as relevant to this investigation. In order to do that we need a statement describing the activity system to be modelled. Such descriptions are known in SSM as Root Definitions (RDs), the metaphor ‘root’ conveying that this is only one, core way of describing the system. A too-simple example would be: ‘A system to paint the garden fence’. Here the worldview is unclear, and it is obvious that a richer description would lead to a richer outcome when the model is used as a source of questions to ask of the real situation. A number of ways of enriching an RD have shown themselves to be useful. For example, we could more richly express the RD above as:

‘A householder-owned and staffed system to paint the garden fence, by hand-painting, in keeping with the overall decoration scheme of the property in order to enhance the appearance of the property’. This makes clear that the model takes a householder’s worldview as given, and that that particular householder believes in DIY activity to improve it. In addition it not only describes *what* the system does (paint the fence); it also says *how* (by hand-painting) and *why* (to enhance the appearance of the property). (Also the worldview assumes a link between painting and improving appearance.) Clearly this would lead to a richer questioning of the real situation to which this purposeful activity was thought to be relevant as a device to structure the questioning.

The whole set of guidelines of this kind – there to help the modelling process – will now be described. They are set out in Figure 2.8; the five numbered elements in the figure will be described in turn.

1. The formula followed in enriching the fence-painting RD above is always helpful, and can apply to every RD ever written. It is known in SSM as ‘the PQR formula’: do P, by Q, in order to help achieve R, where PQR answer the questions: What? How? and Why? PQR provides a useful shape for any and every RD. Remember, though, in using PQR, that if the formula is complete, with all three elements defined, then the transforming process is captured in Q, the declared ‘how’. In the simple example above the Q is ‘hand-painting’ (not simply ‘painting’). Also, though it is not an issue in this example, the model builder has to be able to defend Q as a plausible ‘how’ for the ‘what’ defined by P. If you were to write ‘define health-care needs’ as P and then define Q *only* as ‘by asking patients for their views’ this would not be easily defensible.
2. The PQR formula allows you to write out the RD as a statement. This always describes the purposeful activity being modelled as a transformation process, one in which some entity (in the example an ‘unpainted fence’) is transformed into a different state (here, a ‘painted fence’). Any purposeful

(1) The PQR formula

Do P

By Q


In order to contribute
to achieving R

provides a
shape for

(2) Root Definition

enriches

(3) Mnemonic



Transformation
Process and
Worldview

monitored by criteria for

- Efficacy (E₁)
- Efficiency (E₂)
- Effectiveness (E₃)

leads to

(5)



may be (4)
Primary Task
Issue-based

Figure 2.8 Guidelines which help with building models of purposeful activity

activity you can think of can be expressed in this way, which is useful because it makes model building a straightforward process. For complex activities the entity being transformed will probably be best expressed in an abstract way, for example: 'the health-care needs of Coketown citizens' transformed into 'the health-care needs of Coketown citizens met'. But

the idea of purposeful activity as a transformation always holds, whether the transformation is concrete or abstract. Putting together the activities needed to describe the transforming process (i.e. 'building the model') can begin when an RD is complete, but before moving on to this, elements 3 and 4 in Figure 2.8 should be considered. They further enrich the modelling and improve it as a source of questions to ask in the real situation.

- When the idea of working with RDs as a source of models was being developed, a further enrichment of the thinking came from having, as a reference, a completely general model of any purposeful activity. (This was a way of declaring exactly what we meant by 'purposeful activity'.) The general model is shown in Figure 2.9. It contains elements which can usefully be thought about for any purposeful (transforming) activity.




Figure 2.9 A generic model of any purposeful activity, which yields the mnemonic CATWOE

The model provides the mnemonic CATWOE, defined as in Figure 2.9. The concept here is that purposeful activity, defined by a transformation process and a worldview (a T and a W):

- will require people (A) to do the activities which make up T;
- will affect people (C) outside itself who are its beneficiaries or victims (C for 'Customers');
- will take as given various constraints from the environment outside itself (E) (such as a body of law, or a finite budget);
- could be stopped or changed by some person or persons (O) who can be regarded as 'owning' it.

Many people find it useful, when model building, to start the process by defining first T and W, then the other CATWOE elements. Experience suggests, though, that it is still useful to write out the RD as a statement which gives a holistic account of the concept being modelled.

Finally, within the guideline which CATWOE provides, it is useful to think ahead to the model and ask yourself: What would be the measures of performance by which the operation of the notional system would be judged? Thinking out what those criteria would be really sharpens up the thinking about the purposeful activity being modelled. Three criteria are relevant in every case, and should always be named. We need:

- criteria to tell whether the transformation T is working, in the sense of producing its intended outcome, i.e. criteria for *efficacy*;
- criteria to tell whether the transformation is being achieved with a minimum use of resources, i.e. criteria for *efficiency*; and
- criteria to tell whether this transformation is helping achieve some higher-level or longer-term aim, i.e. criteria for *effectiveness*.

In the case of the simple fence-painting system the criteria address, respectively, the questions: Does this count as 'a painted fence' (human judgement would decide)? Is the painting being done with minimum

use of the resources of materials and time (these might be expressed as costs)? and Does the painted fence enhance the appearance of the property (again human judgement would decide)? These three criteria are always independent of each other. Thus, for example, the purposeful act of taking a drug to relieve your headache might be efficacious if the headache goes. But it could be inefficient if the drug cost too much or was very slow-acting. And it could also be ineffective, medically, if treating the symptom of the headache was unwise because the headache actually signalled a more serious complaint.

These 'three Es' will always be relevant in building any model, but in particular circumstances other criteria might also apply, such as *elegance* (Is this a beautiful transformation?) or *ethicality* (Is this a morally correct transformation?). The judgement is yours as to what criteria are needed.

4. The final consideration in Figure 2.8 when formulating RDs prior to model building concerns RDs as a whole. Are they 'Primary Task' or 'Issue-based' definitions? This useful distinction (though it does not affect model building technique) arose through experience, like most developments in SSM. In the early days, when the legacy of Systems Engineering hung heavy over the new approach, the models built were always of purposeful activity of a kind that was present in the real world in the form of departments, divisions, sections, etc.; that is to say it was institutionalized. Thus, if working in a company with functional sections – production, marketing, research and development, etc. – we would in the early days of developing SSM make models only of a production system, a marketing system, an R&D system, etc. In these cases the boundary of the models we built would coincide with internal organizational boundaries. This is not 'wrong', but it puts limitations on the thinking of the team carrying out the investigation, which may go unnoticed. Every organization has to carry out many, many purposeful activities as it goes about its business. Only a few of these can be

captured in the organization structure as departments, etc. These organizational boundaries are, in the last analysis, arbitrary, and could be changed.

Experience quickly showed that to stimulate the thinking of everyone involved in the investigation it was useful to make models of purposeful activity whose boundaries *cut across organizational boundaries*. These are 'Issue-based' models from 'Issue-based' RDs, models whose boundaries do not coincide with organizational boundaries. When such models are used to ask questions in the situation, interest and attention are always increased. This brings in broader considerations than is the case with a model which accepts organizational boundaries as a given. This is because the questions about what departments, sections, etc. should exist, and what their boundaries should be are always bound up in the power-play going on in organizations. That catches everyone's attention!

As a generalization we can suggest one choice of Issue-based RD which is always worth considering. In virtually all organized human groups there will always be contentious issues concerned with allocating resources. This is something which affects all members, leads to wide discussion, and is not usually assigned as an activity to a particular sub-group. An issue-based model based on transforming unallocated into allocated resources will be worth considering as a stimulant in most investigations. The general rule is: never work exclusively with either Primary Task (PT) or Issue-based (IB) RDs. Most investigations will best feature a mixture of both types.

5. Earlier in this chapter, in section 2 above, model building was described as 'putting together the activities needed to describe the transforming process', in other words defining and linking the activities needed to achieve the transforming process. Given the guidelines provided by PQR, an RD, CATWOE, the 3Es and PT/IB, this task should not be a difficult

one. The only skill called for is logical thinking. The most common error – even among logical thinkers – is to take your eye off the root definition and start modelling some real-world version of the purposeful activity being modelled. In work in a medium-sized manufacturing company, concerned with various issues regarding product distribution, it was easier for the SSM practitioners to build relevant models than it was for the distribution manager. He kept slipping into modelling the current ways of working in his department rather than the concepts in RDs. If you do this, of course, you find yourself not questioning current practice but comparing X with X – not very profitable!

People find their own way of making the selected relevant models, but a logical sequence to follow, or to refer to if in difficulty, is as follows:

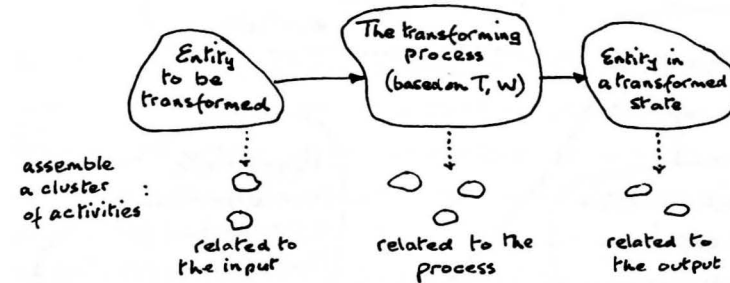
- (1) Assemble the guidelines: PQR, CATWOE, the RD, etc.
- (2) Write down three groups of activities – those which concern the thing which gets transformed (the 'unpainted fence', or the 'health needs of the citizens of Coketown', in the examples above); those activities which *do* the transforming; and any activities concerned with dealing with the transformed entity (e.g. judging if it improves the appearance of the property, in the fence-painting example); this will give you a cluster of activities.
- (3) Connect the activities by arrows which indicate the dependency of one activity upon another; for example, you can't *use* a raw material to make something before you've *obtained* it, so an arrow goes from an 'obtain' activity to the 'use' activity. In Figure 1.3 activity 7 (paint the fence) depends upon both activities 4, 5 and 6, since you can't paint the fence until you've obtained both brush and paint and prepared the fence.
- (4) Add the three monitoring and control activities, which always have the structure shown in Figures 1.2 and 1.3.
- (5) Check the model against the guidelines. Ask yourself: Does every phrase in the RD lead to something in the model? And: Can every activity in

the model be linked back to something in the RD or CATWOE, etc.? If the answer to both questions is 'Yes', then you have a defensible model. Note that the word used here is 'defensible' rather than 'correct'. This is because everyday words have different connotations for different people. Competent SSM practitioners working from the same RD might well produce somewhat different models; this is because they are interpreting the words in the RD, etc. somewhat differently. The important thing is that you can defend your model as representing what is in your RD, PQR, CATWOE, etc.

Figure 2.10 summarizes the model building process.

Finally, on model building, there is one more guideline worth taking seriously. *Aim* to capture the activity in the operational part of the model in 'the magical number 7 ± 2 ' activities (but do break the 'rule' if necessary). This famous phrase comes from a celebrated paper in cognitive psychology. George Miller, based on laboratory work, suggests that the human brain may have the capacity to cope with around seven concepts simultaneously. Whether or not this is true it is certainly the case that a set of 7 ± 2 activities can be thought about holistically. If the number seems low, this is not a problem. Any activity in a model can itself, at a more detailed level, become the source of an RD and a model. Thus, in Figure 1.3, activity 6 (obtain paint) could itself be expanded into a model which set out the connected, more-detailed activities which together combine to constitute 'obtain paint' – activities concerned with checking out suppliers, their prices, selecting one, etc. If this model were built, its activities would be numbered 6.1, 6.2, 6.3, etc. since they all derive from activity 6 in the parent model. In this way coherence is maintained no matter how many levels it may be necessary to go to in a particular investigation. In the authors' experience of more than a hundred studies it has never been necessary to expand beyond two levels below that of the parent model, and even then expanding only a few activities at the lower levels.

1. Assemble guidelines : T and W
PQR ; PT/IB
CATWOE , E₁ E₂ E₃
2. Starting from T and W name the purposeful action as a transformation:



3. Structure the activities according to dependency of one on another



4. Add the monitoring and control activities



5. Check the mutual dependency of guidelines and model

Figure 2.10 A logical process for building SSM's activity models

The first model presented here, to illustrate the idea of purposeful activity models, was that in Figure 1.3. This was presented without a Root Definition, but now that this has been defined (above) we can present part of the model in a more developed form. This is done in Figure 2.11 which makes one particular change. It would have been possible to include in the 'operations' part of the model an activity such as 'ascertain the judgement about the enhanced appearance of the property'. Another way of bringing in the R of PQR (the higher-level, or longer-term aim of the transforming process,




Figure 2.11 A variant of part of the model in Figure 1.3

judged by the criteria for effectiveness) is shown in Figure 2.11. The monitoring and control activity has been split into two, with the monitoring for effectiveness having the added activity: 'Appreciate householder's aspirations for the fence painting.' This leaves open who would make the judgement about the hoped-for enhancement of the appearance of the property – the householder? his or her partner? the neighbours? a prospective purchaser? This is probably, in this instance, the most elegant way of bringing all the elements in the guidelines into the model.

Appendix C carries an example of model building which starts from a Root Definition and 'talks through' the whole process for that definition. In general, the best way to learn about activity modelling is to have a look at examples and then have a go. And do remember that even rough-and-ready models can be helpful in real situations.

The SSM Learning Cycle: Using Models to Structure Discussion about the Situation and its Improvement

When we enter a problematical situation and start drawing rich pictures and carrying out preliminary versions of Analyses One, Two and Three, we begin to build up what can become a rich appreciation of the situation. This appreciation – helped especially by the list of possible 'issue owners' from Analysis One – enables us to begin to name some models which might be helpful in deepening our understanding of the situation and beginning to learn our way to taking 'action to improve'. Having built a hopefully relevant model or two, we are then ready to begin the structured discussion about the situation, and how it could be changed, which will eventually lead to action being taken. The models are the devices which enable that discussion to be a structured rather than a random one.

In everyday situations, typical discussions among professionals are characterized by a remarkable lack of clarity. In a typical 'management' discussion in an organization, unless there is a chairperson of near-genius, different voices will be addressing different issues; different levels, from the short-term tactical to the long-term strategic, will be being addressed; different speakers will assume different timescales. The resulting confusion will then provide splendid cover for personal and private agendas to be advanced. Use of the models to help structure discussion enables us to do rather better than this.

Structure to the discussion is provided by using the models as a source of questions to ask about the situation. This phase of SSM has usually been referred to as a 'comparison' between situation and models, but this wording is truly dangerous if it is taken to imply that the discussion focuses on deficiencies in the situation when set against the 'perfect' models. *The*

models do not purport to be accounts of what we would wish the real world to be like. They could not, since they are artificial devices based on a pure worldview, whereas human groups are always characterized by multiple conflicting worldviews (even within one individual!) which themselves change over time – sometimes slowly, sometimes remarkably quickly. (It is those conflicting worldviews which are the fundamental cause of the confusion in most ‘management’ discussion.)

No, the purposeful activity models simply enable our organized discussion to take place. From the model we can define a set of questions to ask. For example: ‘Here is an activity in this model; does it exist in the real situation? Who does it? How? When? Who else could do it? How else could it be done?’ . . . etc. Or: ‘This activity in the model is dependent upon these other two activities; is it like this in the real situation?’ There is no shortage of possible questions, and practitioners quickly develop the knack of passing in a light-footed way over many possibilities and resting on those questions which are likely to generate attention, excitement or emotion. The questions can be about activities or the dependence of one activity upon another or upon the measures of performance by which purposeful activity is judged.

A general finding is that groups find it very difficult to answer questions derived from the measures of performance in a model. ‘What criteria would indicate the degree to which this activity (either individual, or the set of operational activities as a whole) is efficacious, efficient and effective?’ This is usually a difficult question to answer in most real-world situations, due to their complexity, but it usefully draws attention to the need for organized processes of monitoring, something which is often given scant attention in organizations of all kinds. At a broader level, the fact that a given model is based upon a declared (pure) worldview will draw attention to other, usually implicit, worldviews which may underlie what is actually going on in the situation. This may serve to define other relevant models worth building and also helps to raise the level of discussion to that at which

previously taken-as-given assumptions are now questioned. This will usually wake up anyone who is sleep-walking through the discussion, not least because differences of worldview always provoke *feelings*, not simply mental activity. (Also, incidentally, experience in developing SSM suggests that the stimulation of emotion is probably, for most people, a powerful trigger for significant learning to occur.)

In practice, several ways of conducting the questioning of the situation have emerged. An informal approach is to have a discussion about improving the situation in the presence of the models. If some relevant models are on flip charts on the wall, they can be referred to and brought into the discussion at appropriate moments. This has been found useful in situations in which detailed discussion of the SSM approach is inappropriate or is not feasible for cultural reasons. It was effective in a situation in a giant publishing/printing company which was characterized by an operation – publishing, printing and selling consumer magazines – which combined two very separate cultures who found it difficult to appreciate each other’s worlds. The editor/publisher culture contained people very different from those in the printing culture, though they worked in the same company. Models which related to the whole operation of commissioning material, editing and assembling magazine issues, printing them and marketing them, proved useful here as a background, rather than as a source of specific detailed questioning. They were on flip charts on the wall, and could be referred to during discussion.

A more formal approach, probably the most commonly used, is to create a chart matrix as in Figure 2.12. The model provides the left-hand column, consisting of activities and connections from the model, while the other axis contains questions to ask about those elements (which may vary depending on the investigation underway). The task is then to fill in the matrix by answering the questions. Case 1 in Chapter 4 illustrates a question-matrix from an SSM investigation.




Figure 2.12 A formal process for using models to question the real-world situation

An important warning here is that this process should not be allowed to become mechanical drudgery. This is where a light-footed approach is needed, glancing quickly at many activities and questions, making judgements, and avoiding getting bogged down. Experience quickly develops this craft skill. In fact, experience suggests that this business of seeking to avoid plodding through every cell in the matrix itself helps develop insights into 'the real issues in this situation' – though such judgements have to be tested.

A third way of using models to question reality is to use a model as a basis for writing an account of how some purposeful action would be done

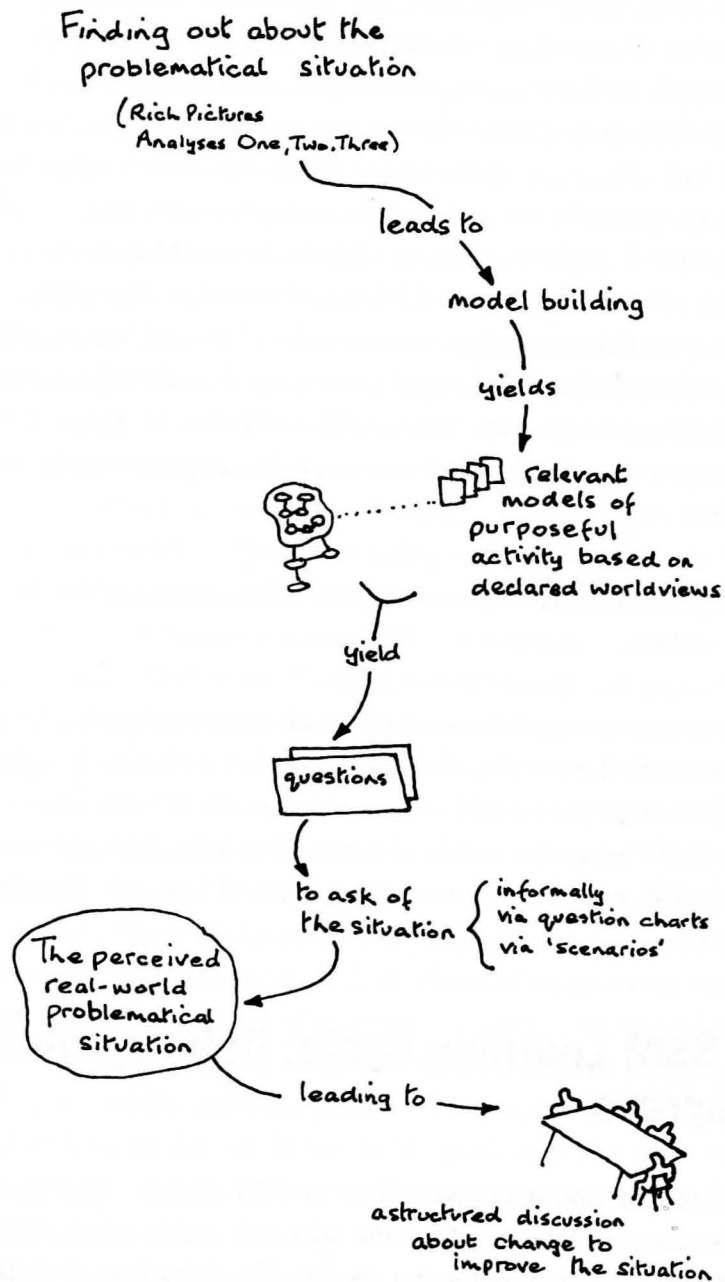


Figure 2.13 The role of models in SSM summarized

according to the model, and comparing this story, or scenario, with a real-world account of something similar happening in the real world. For example, work with SSM was carried out in a chemical company which treated every plant start-up as if it were the first they had ever carried out. It was very useful in that situation to make a basic generic model of 'a system to start up a new chemical plant' and then write a story from this pure (instrumental) model which could be compared with the real-world stories of previous plant start-ups, usually stories of delays and cock-ups. The company was right in saying that every plant start-up revealed unique features. But this work also showed that it was useful to have a generic model to hand when planning for a new start-up. This model could then be enriched by new experiences, so that the chance of future surprises in plant start-up could be diminished.

Figure 2.13 summarizes different ways of using models in the context of SSM as a whole.

Whichever way the models are used to structure discussion, the aim is the same: to find a version of the real situation and ways to improve it which different people with different worldviews can nevertheless *live with*. Outside of the arbitrary exercise of power, this is the necessary condition which must be met in any human group if agreed 'action to improve' is to be defined.

The SSM Learning Cycle: Defining 'Action to Improve'

When describing the discussion/debate in SSM, much – perhaps most – of the secondary literature about the approach makes a remarkable and fundamental error. It assumes that the purpose of the discussion/debate is to find *consensus*. It is a 'remarkable' mistake in that anyone who had read

the primary literature with care would not make it, and it is 'fundamental' because, in order to cope with the complexity of human affairs, SSM uses a much more subtle idea than 'consensus'. It works with the idea of finding an *accommodation* among a group of people with a common concern. This does not abandon the possibility of consensus; rather it subsumes it in the more general idea of accommodation. A true consensus is the rare, special case among groups of people, and usually occurs only with respect to issues which are trivial or not contentious; issues which people do not feel particularly strongly about. In the general case, however, because individuals enter the world with different genetic dispositions and then have different experiences in the world, there will always be differences of opinion resulting from different worldviews. So, if a group of people are to achieve agreed corporate action in response to a problematical situation, they will have to find an accommodation. That is to say they will have to find a version of the situation which they can all live with. These accommodations will of course involve either compromise or some yielding of position. A compromise may give no member of the group all they personally would look for in action to improve the situation. But finding an accommodation is usually a necessary condition for moving to deciding 'what we will now do' in the situation.

The idea of finding accommodations is probably most familiar to us in our personal lives. Any family, as long as it is not of the classic Victorian kind, run by a (male) tyrant who decides everything, will have to continually find versions of the family situation which the different members can accept and live with. This is a necessary characteristic if families are to stick together over a long period. But the idea is also relevant to our professional lives, and to public life. A dramatic illustration of the latter is provided by some British political history. In the UK in the 1970s there were a number of major strikes in the coal industry, the disputes usually involving pay. One of those strikes lasted for a year. Now, the interesting thing about these disputes was that they were conducted within an accommodation between the two sides, the Coal Board and the National Union of Miners (NUM). Although

the miners were on strike, members of the NUM nevertheless went down every mine in the country, every day, in order to keep the pumps running, since if you don't continuously pump water out of a coal mine you lose the mine. Although both sides regretted, but were prepared to have the dispute, there was an accommodation between them at a higher level: neither was prepared to live with the idea of the conflict destroying the whole industry. (It took political action to do that some years later!)

This view taken within SSM – that consensus is rare in human affairs, due to clashing worldviews – is not to be regretted. Clashing worldviews, always present, are a source of strong feelings, energy, motivation and creativity. If you find that the models you've built are not leading to *energetic* discussion, abandon them and formulate some more radical Root Definitions.

As discussion based on using models to question the problematical situation proceeds, worldviews will be surfaced, entrenched positions may shift, and possible accommodations may emerge. Any such accommodation will entail making changes to the situation, if it is to become less problematical, and discussion can begin to focus on finding some changes which are both arguably desirable and culturally feasible. In practical terms it is a good idea not to try and discuss the abstract idea 'accommodation' directly. It is best approached obliquely through considering what changes might be made in the situation and what consequences would follow. The relations between accommodations, consensus and changes is summarized in Figure 2.14, and the practical way forward in seeking accommodation is by exploring possible changes and noting reactions to them.

In doing this it is best to think richly about change in human situations, separating the concept into three parts for analytical purposes, even though any significant change in real situations will usually entail all three elements. These are: making changes to *structures*; changing *processes* or procedures; and changing *attitudes*.

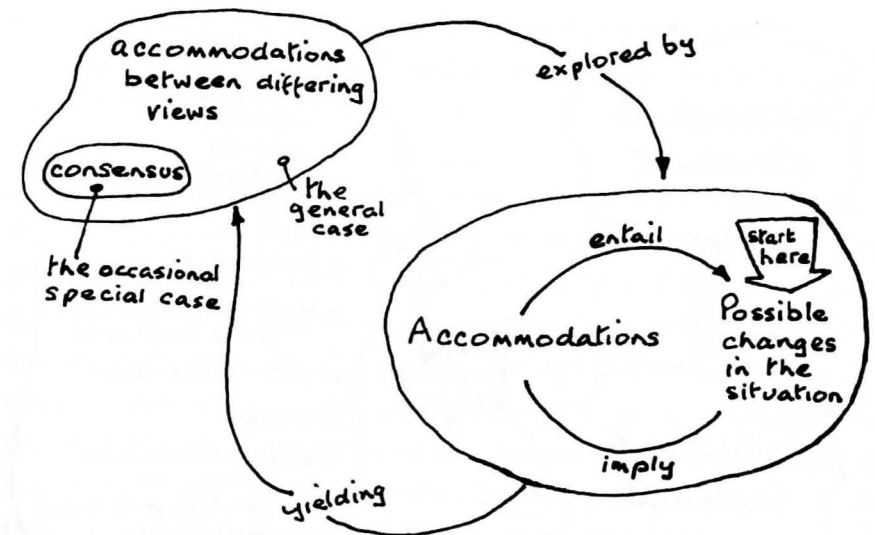


Figure 2.14 Seeking accommodations or (rarely) consensus by exploring implications of possible changes

Obviously the easiest element to change is structure, which can often be done by decree through the exercise of legitimate power. Researchers have noted, for example, that large organizations tend to reorganize themselves structurally about every 18 months to two years. In the UK, governments have imposed structural change upon the National Health Service more than 20 times since it was established in 1948. That is the easy part, for governments. But of course new structures usually require both new processes and new attitudes on the part of those carrying out the processes or being affected by them. Organizations (and governments) find it much harder to think out the necessary new processes; and no one can be sure, in a unique social situation, about what to do to change attitudes in a particular direction. (In our current culture, obsessed with economics, the usual mechanism for trying to change attitudes is to provide material incentives, but this reflects acceptance of a bleak model of human beings as creatures responding only to sticks and carrots. Human beings are more complex than that.)

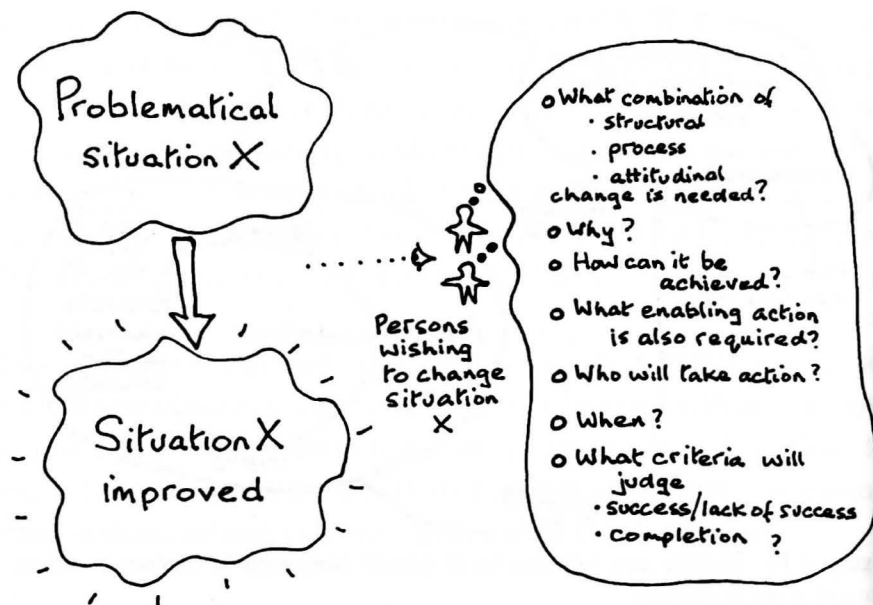


Figure 2.15 SSM's stance on introducing change in human situations

Figure 2.15 illustrates the stance on 'change' taken within SSM. It represents a reminder of things to think about when considering changes which are both desirable and feasible. It is self-explanatory, but two points are worth making. There is a question concerning the 'enabling action' which may be necessary if a potential change is to be accepted. This recognizes the social *context* in which any change will sit. Because of this context, introducing the change may require other action, enabling action, which is not directly part of the change itself. For example, when working within the UK National Health Service for the first time, in the early 1970s, the authors quickly found that in an acute hospital no proposed change would get accepted unless it had the support of senior hospital consultants. Shifts in the disposition of power have now modified that, but at that time in the history of the NHS, enabling action to secure the support of senior doctors was essential if any change of any kind was to occur in a hospital! The second point concerns trying to define the criteria by which a change can be judged as 'completed' and 'successful/unsuccessful'. This point has already been made

above in connection with asking about 'monitor-and-control' activities in a real situation: well worth doing, but don't expect people in the situation to have any ready answers.

As we come to the end of this chapter's exposition of a 'fleshed-out' account of SSM, the discussion has become less detailed, in the sense that there are more detailed guidelines for finding out about a real-world situation, and building models used to question it, than there are for taking action to improve the situation. This is inevitable, and is due simply to the fact that no human situation is ever exactly the same as any other. Once we start exploring the real complexity of a human situation, not simply its logic, then formulae, algorithms and ready-made solutions are not available. Even guidelines become fewer. That being so, it seems helpful to give here, ahead of Part Two of this book, a real example of these ideas about change in action.

In the work mentioned above in the publishing-printing industry, the company carried out both of these major activities in selling a large range of consumer magazines. Publishing and printing were organizationally separate, and were in the hands of two very different cultures: on the one hand 'media-folk', on the other 'technologists'. There were many issues in the company concerning investment, pricing, and the placing and scheduling of work. For example, the printers thought of themselves as 'jobbing printers', making no distinction between printing one of the company's titles or that of a competitor. Publishers had ill-defined freedom to print within the company or externally. There were many rows about 'where to print', for example. This was an occasion in which the least-formal way of using models to question the situation was used: discussion in the presence of the models, which were on flip charts on the walls. In the discussion stimulated by the models the end point finally reached, subsequently approved by the board, was that there should be structural change. A new unit within the company was set up. This unit was centrally placed, and was staffed (part-time – it was not

permanently in session) by people from both publishing and printing. This structural change was just about culturally feasible (where a fully integrated magazine-producing operation was out of the question) and the processes within the new unit were defined. As far as changes of attitude were concerned, the chief executive, who understood the difficulties of forcing change of that kind, wrote in the in-house company 'newspaper': 'Primarily the new unit is concerned with trying to develop a more effective relationship between our publishers and printers.' He was hoping that each of the two cultures would, through working together on some issues, begin to see the world through the eyes of the other.

The Whole SSM Learning Cycle Revisited: Seven Principles, Five Actions

Before moving on to accounts of SSM in action, in Part Two, we can now summarize the whole learning cycle of the SSM approach. In a concise account of SSM, which is as spare as we can make it, seven principles lead to five actions. These are based only on findings which, through many experiences over a long period, always turned out to be helpful. They are the end product of the several hundred cycles through the LUMAS model in Chapter 1 (Figure 1.8).

The seven principles which underlie SSM are set out first.

1. The idea 'real-world problem' is subsumed in the broader concept of 'real-world *problematical situation*'; that is to say, a real situation which someone thinks needs attention and action.
2. All thinking and talking about problematical situations will be conditioned by the *worldviews (Weltanschauungen)* of the people doing the thinking and talking. These worldviews are the internalized taken-as-given assumptions which cause us to see and interpret the world in a particular way (one observer's 'terrorism' being another's 'freedom fighting').

3. Every real-world problematical situation will contain people trying to act *purposefully*, with intent. This means that *models of purposeful activity*, in the form of systems models built to express a particular worldview, can be used as *devices* to explore the qualities and characteristics of any problematical human situation.
4. *Discussion and debate* about such a situation can be *structured* by using the models in (3) as a source of questions to ask about the situation.
5. Acting to improve a real-world situation entails finding, in the course of the discussion/debate in (4), *accommodations* among different worldviews. An accommodation entails finding a version of the situation addressed which different people, with different worldviews, can nevertheless live with.
6. The *inquiry* created by principles (1) to (5) is in principle a *never-ending process of learning*. It is never-ending since taking action to improve the situation will change its characteristics. It becomes a new (less problematical) situation, and the process in (3), (4) and (5) could begin again. Learning is never finished!
7. Explicit organization of the process which embodies principles (1) to (6) enables and embodies *conscious critical reflection* about both the situation itself and also about the thinking about it. This reflection, which leads to learning, can (and should) take place prior to, during and after intervening in the situation in order to improve it. The process thus itself virtually ensures *reflective practice* by those who make use of it. Once the practitioner has internalized the SSM process, so that he or she no longer has to stop and ask questions about it ('Remind me again, what did PQR stand for?') then reflective practice becomes built-in too. The SSM user becomes a reflective practitioner.

These seven principles clearly underlie the four actions which define the classic shape of SSM in Figure 1.5: finding out about a problematical situation;

making models relevant to exploring it, based on different worldviews; questioning the situation using the models, in order to find desirable and feasible change; and defining/taking action to change the situation for the better. The seventh principle itself defines a fifth action which ensures cycling round the primary four, namely critical reflection on the whole process. This fifth action is at a different level from the other four. It is *about* the other four, i.e. at a meta-level. It is the activity which ensures that the lessons learned are captured, in the way that the LUMAS model of Figure 1.8 indicates. Figure 2.16 expresses these five activities at their two levels.




Figure 2.16 The five activities which flow from SSM's seven principles

Finally, in completing this more detailed account of SSM, it is worth re-emphasizing some of its core ideas. It does not seek 'solutions' which 'solve' real-world problems. Those ideas are a mirage when faced with real-life complexity, with its multiple perceptions and agendas. Instead SSM focuses on the process of engaging with that complexity. It offers an organized process of thinking which enables a group of people to learn their way to taking 'action to improve'; and it does that by means of a well-defined, explicit process which makes it possible to recover the course of the thinking which leads to action. This makes sure that every use of the approach produces learning which will accumulate over time, leaving the user better equipped to cope with future complexities.

Part Two now provides examples of these ideas in action in situations of many different kinds.