The Viable System Model: Interpretations and Applications of Stafford Beer's VSM
Edited by R. Esperjo and R. Harnden
Published 1989 by John Wiley & Sons Ltd

1

The Viable System Model: its provenance, development, methodology and pathology*

Stafford Beer=

President of the World Organization for Systems and Cybernetics

It took the author 30 years to develop the Viable System Model, which sets out to explain *how systems are viable* – that is, capable of independent existence. He wanted to elucidate the laws of viability in order to facilitate the management task, and did so in a stream of papers and three (of his ten) books. Much misunderstanding about the VSM and its use seems to exist; especially its methodological foundations have been largely forgotten, while its major results have hardly been noted. This paper reflects on the history, nature and present status of the VSM, without seeking once again to expound the model in detail or to demonstrate its validity. It does, however, provide a synopsis, present the methodology and confront some highly contentious issues about both the managerial and scientific paradigms.

Provenance

At the end of my military service, I spent a year from the autumn of 1947 to that of 1948 as an army psychologist running an experimental unit of 180 young soldiers (a moving population, 20 of them changing every (fortnight). All these men were illiterate, and all had been graded by a psychiatrist as psychopathological personalities. They could not write a letter home, nor read a newspaper, and such sums as 4 + 3 = ? often had them fooled. But they could debate with great energy and verbal facility if not felicity; they could play darts – '21 that's 15 and a double 3 to go'; and they could state the winnings on a horse race involving place betting and accumulators with alacrity and accuracy, and apparently without working it out. they had their own conception of discipline, involving terrorism and violence in the barrack room, which met every desideratum of a military unit in its ends, though not in its means.

*Reprinted by permission from the *Journal of the Operational Research Society,* vol. 35, 7-26. Copyright © 1984: Operational Research Society.
=Visiting Professor of Cybernetics, Manchester University Business School; and Professor of Social System Sciences, University of Pennsylvania, the Wharton School.

I had a background in philosophy first and psychology second; the latter school had emphasized the role of the brain in mentation and of quantitative approaches in methodology. The analytical models that I now developed, the hypotheses set up and tested, were thus essentially neurophysiological in structure and statistical operation. The behavioural models derived mainly from experience: I had a background in the Gurkha Rifles too. What made these people, unusual as they were, tick – and be motivated and be adaptive and be happy too (for most of them were)? And how did the description of individuals carry over into the description of the whole unit, for it seemed so to do: every one of many visitors to this strange place found it quite extraordinary as an organic whole. It simply was not just a unit housing a population of unusual soldiers. The first regimental sergeant major asked for a posting.

This was the empirical start of the subsequent hypothesis that there might be *invariances* in the behaviour of individuals, whether they be 'normal' or not, and that these invariances might inform also the peer group of individuals, and even the total societary unit to which they belong. In the early 'fifties this theme constantly emerged in my operational research work in the steel industry: I used then to refer to the structure of 'organic systems'. So the viable systems model (VSM) dates back 30 years. I pursued it through neurocybernetics and social science, through the invention and study of cybernetic machines, through the mathematics of sets and stochastic processes, and at all times through the OR fieldwork in industry and government. The quest became to know *how systems are viable*; that is, how they are 'capable of independent existence' – as the dictionary has it. By the time my first book on management cybernetics was published, I had also mapped a set-theoretic model of the brain on to a company producing steel rods, and published the basis of the whole approach (Beer, 1959, 1960).

The set-theoretic model proved difficult for people to understand, and eventually a streamlined version of the model appeared called *Brain of the Firm*, using neurophysiological terminology instead of mathematics (Beer, 1972). Some commentators were offended by this and called the model analogical – despite my denials and explanations (see later) that this was so. Hence, in a still later book a new version of the VSM was developed from first principles, called *The Heart of Enterprise*, in the belief that the necessary and sufficient conditions of viability had by now been established (Beer, 1979).

The invariances that I had finally unearthed were stated; and the central principle of recursion (that every viable system contains and is contained in a viable system) stood duty as the explanation of all the observational evidence that had begun to accumulate from the military experience onward. Moreover, I developed a topological version of the original set-theoretic algebra that it seemed no-one would study properly. The drawings were now rigorous mathematics in themselves in that they offered explicit homomorphic mappings of any one VSM recursion on to the next – as may be seen in the simplified version at Figure 4. (In 1972 the drawings had given an indication of the recursion theorem and relied on the independently published mathematics.

Throughout its development, and to this day, the VSM has been in a process of continuous testing and verification. Meanwhile, however, the whole approach had its most significant and large-scale application during 1971-73 in Allende's Chile. As an outcome of this experience, five new chapters were added to *Brain*, and the overhauled and extended text was republished (Beer, 1981). Thus (the new) *Brain and Heart* stand, as complementary volumes, for the theory of the viable system and its 'laws' in management cybernetics, and a trilogy has been completed with *Diagnosing the System* (Beer, 1985). Commentators often imply that I am obsessed with this model. Well, the quest to establish *how systems are viable* and its 30-year pursuit have certainly been demanding. Even so, the three books mentioned are only three out of ten. The philosophy of science that I was simultaneously developing is expounded in *Decision and Control*, and it is from this that I draw the following methodology and apply it to the VSM (Beer, 1966).

The methodology of topological maps

When we notice similarities between two different systems, for instance between the regulatory system of an individual and a group, or between a brain and a firm, the comparison often begins in a literary manner. There is the simile: 'management communications are like the nervous system, in that…'. There is the more direct metaphor: 'the real muscle of the plant is the cogging mill'. Such comparisons may help to convey insights, although everyone knows better than to take them too seriously. but as perception of the two systems deepens, and perhaps observations are taken, we may come to hola conceptual models of both systems that become exciting and helpful. This stage is easily recognized because we find that some circumstances that we understand in one system throws light on a parallel circumstance in another. It is now worth ‘drawing analogies’; on the other hand, everyone knows that ‘analogies may be carried too far’.

The process continues, and begins to have the marks of a scientific method, when we try to develop rigorous formulations of the two conceptual models. (Figure 1 refers.) These will each be a *homomorphic mapping*, insofar as many elements in the system that is conceptually modelled will map on to one element in a rigorous model. All falling apples, and not only the particular falling apple observed by Newton, obey the law of gravitation: we *select* those mappings that exhibit mathematical *invariance*. And if we travel to Pisa, we find Galileo (who died in the year that Newton was born) supposedly dropping not-applies from the leaning tower, but determining a *constant* none the less.

FIGURE 1 TO GO IN HERE

Now what happens if we map the two rigorous formulations of orchard systems and Pisa systems on to each other? If we find invariances between the two systems, then these are isomorphic mappings, one-to-one in the elements selected as typifying systemic behaviour in some selected but important way. The generalized system that comes out of this process, which applies to all systems of a particular class, is a scientific model – in the case just considered, a model of gravitation. The generalization of some behaviour invariably and invariantly exhibited by the system as interpreted through this systemic model we usually call a law. Nonetheless, we have made a selection; we have reduced systemic variety through our homomorphisms. But that is the very business of scientific discovery. In fact, every system can be mapped on to any other system *under some transformation*; thus Ashby was wont to say that the Rock of Gibraltar makes a good model of the brain, if your interest is exclusively in spatio-temporal extensity.

Considering these matters coolly, and handling them in a world which upholds a particular paradigm that does not compare rocks to brains, is not an easy matter. The precise difficulty that most people have arises when a breach of taxonomy is offered as between social systems, individual people and artifacts. The amalgam is seen as essentially different form the unity, and the animate as essentially different from the inanimate. But these were among the major paradigmatic distinctions that were explicitly questioned by the founders of cybernetics in the ’forties. Certainly my own methodology, especially as it relates to the class of viable systems, makes its mappings quite happily across these boundaries. Witness the very title of the most formal statement of the method: *The World, the Flesh and the Metal* (Beer, 1965). An extract from this paper, giving a group-theoretic analysis of the modelling methodology, will be found in Appendix 1.

Having said all this, there is no way of ‘proving’ a model: the by now classical criterion of ‘falsifiability’ remains instead. As experience of the VSM grew, as its format was made tidier, and as others became involved, more and more viable systems were mapped on to the model: the invariances held. The methodology at this point may be described as the *yo-yo technique*. That is to say: we have constructed a VSM by mapping (let’s say a brain on to a firm and now wish to test a second, third and so on viable system against the scientific model. We run down the chain of similes, analogies and homomorphs with one of these fresh systems until the isomorph is reached, testing the insights and invariances as appropriate on the way; then we return up the chain with another fresh system; then down again, and so on – hence the yo-yo metaphor (rather than model, note). Other scientists around the world have confirmed the VSM in various modes and situations, most but not all of them managerial. A note about these activities appears at Appendix 3.

On mapping and measuring complexity

Although we may derive a model in the manner shown, and although we may develop confidence in it through many applications over a long period, practical activity requires more than this. The management of any viable system poses the problem of managing complexity itself, since it is complexity (however generated) that threatens to overwhelm the system’s regulators. This is very obvious in biological systems, wherein there are no self-proclaimed ‘managers’; but in social systems too complexity tends to overwhelm those managers whose activities are not seriously directed towards viability but to short-term goals such as profit. A precise measure of systemic complexity had been proposed as *variety*, meaning the number of distinguishable elements in a system, or by extension the number of distinguishable systemic states (Ashby, 1965). The problem of controlling this variety is daunting indeed, if all distinguishable states are equally likely. But they are not.

We are used to suppose the variety in social systems is kept under control by a legislative mode of regulation that restrains variety proliferation. But, as Ashby learned from biological systems, something more subtle underlies any such technique. The notion of a ‘coenetic variable’ explains the delimitation of the variety of environmental circumstances and of apparently regulatory responses at the same time (Sommerhoff, 1950). Sommerhoff wrote (see Figure 2): Coenetic (pronounced ‘sennetic’, from the Greek meaning ‘common’) variables simultaneously delimit variety as shown, so that trajectories of the system *converge* on to a subsequent occurrence. Sommerhoff called this ‘directive correlation’. The schematic diagram exemplifies what I later called ‘intrinsic control’: in the very process of disturbing environmental circumstances, the coenetic variable evokes a response that converges on an adaptive outcome.

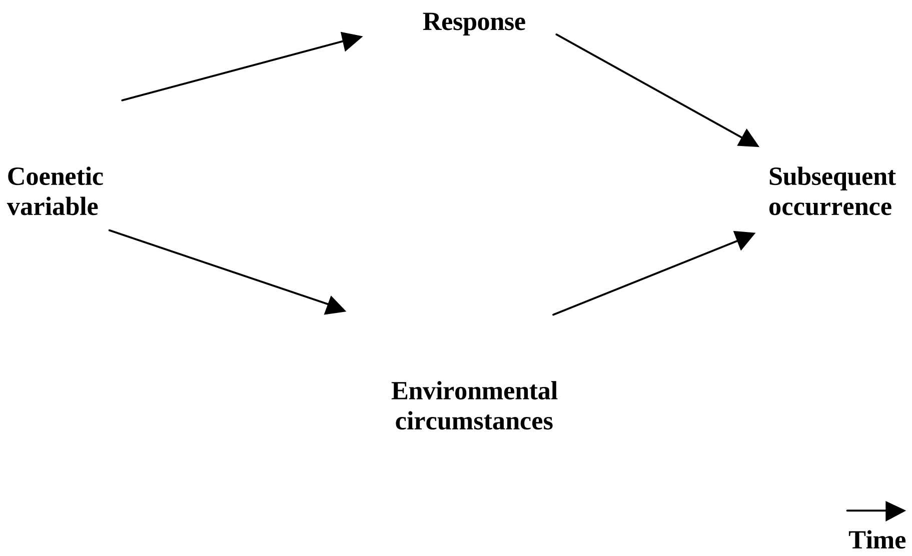

Figure 2. Sommerhoff's account of directive correlation

The Viable System Model

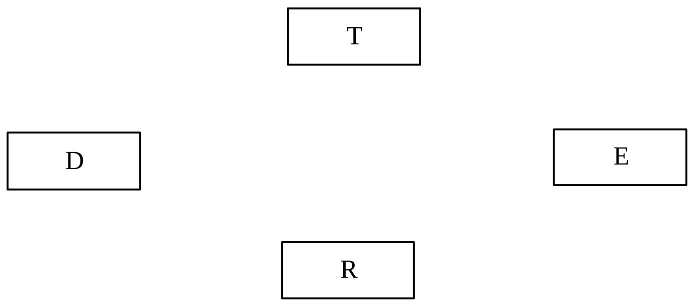

Figure 3. Ashby's account of 'requisite variety'

Ashby for his part had developed a schematic treatment based on Shannon's notation (Shannon and Weaver, 1949; see Figure 3). D stands for disturbance, and is equated by Ashby with the coenetic variable. E is still the outcome set, which is exhausted by good and not-good subsets (in relation to viability). T is a table of the transformations which D will undergo to generate E, and is equated by Ashby with the environmental circumstances of Summerhoff. But now Ashby is taking note that R may, after all, directly influence T in its task of modifying E.

He argues thus. If R's state is always to have the same effect on T, whatever state D may adopt, then the variety of E will be the same as the variety of D. But if R may adopt two states, then the variety at E can be halved. And so on. 'If the variety in the outcomes is to be reduced to some assigned number, or assigned fraction of D's variety, R's variety *must* be increased to at least the appropriate minimum. Only *variety in* R's *moves can force down the variety in the outcomes*.' This is the famous Law of Requisite Variety.

Now it is clear that if D is a coenetic variable, so that R and T are directively correlated, then the variety of the outcomes E will be constrained. Since in both biological and social systems there may be coenetic variables that are unrecognized as such, this would account for a more regulated system than the unrecognizing observer would have any right to expect. Even so, and as Ashby says:

> 'variety comes to the organism in two forms. There is that which threatens the survival of the gene pattern – the direct transmission by T from D to E. This part must be blocked at all costs. And there is that which, while it may threaten the gene-pattern, can be transformed (or re-coded) through the regulator R and used to block the effect of the remainder (in T).'

The model of any viable system, VSM, was devised from the beginning (the early 'fifties) in terms of sets of interlocking Ashbean homeostats. An industrial operation, for example, would be depicted as homeostatically balanced with its own management on one side, and with its market on the other. but both these loops would be subject to the Law of Requisite Variety. Since the variety generated by the market would obviously be greater than the industrial operation could contain, then 'this part must be blocked at all costs', as Ashby has said. This became in my first book (Beer, 1959):

> 'Often one hears the optimistic demand: "give me a *simple* control system; one that cannot go wrong". The trouble with such "simple" controls is that they have insufficient variety to cope with variety in the environment. Thus, so far from not going wrong they cannot go right. Only variety in the control system can deal successfully with variety in the system controlled.

This understanding came from down-to-earth experience as the production controller of a steelworks. By the same token, just as proliferating incoming variety must be blocked at all costs, so must outgoing managerial variety be enhanced – by transformation or recoding through the regulator R, as Ashby said. Looking at the variety-disbalanced homeostats of the VSM, I wrote:

> 'Each part-system provides unlimited variety… It is the function of intelligence to tap that variety, to organize it, to *select*…. What is needed, is the amplification of the primary selection.'

It has always seemed to me that Ashby's Law stands to management science as Newton's Laws stand to physics; it is central to a coherent account of complexity control. 'Only variety can destroy variety.' People have found it tautologous; but all mathematics is either tautologous or wrong. People have found it truistic; in that case, why do managers constantly act as if it were false? Monetary controls do not have requisite variety to regulate the economy. The Finance Act does not have requisite variety to regulate tax evasion. Police procedures do not have requisite variety to suppress crime. And so on. All these regulators could be redesigned according to cybernetic principles, as I have argued *passim* (Beer, 1975, especially).

For present purposes, however, I seek only to show how Ashby's Law was derived, and how it at once suggested to me that if variety were not requisite in a regulatory homeostat, then either the greater variety must be attenuated, or the lesser variety must be amplified or both. This conclusion does not appear to be novel, as has been suggested, but to be sanctioned by Ashby's own words quoted above. Certainly my own applications and extensions of homeostatic's theory in management went beyond Ashby in treating the box called T, supposedly a 'table', as a *black* box – that is to say that the box contains a table that is not available to inspection (something that I had learned in military OR, for foes do not care to make their transformation rules manifest). but Ashby was the doyen of black boxes too.

What was perhaps novel, for the record, was the recognition that in the VSM homeostats requisite variety applies in three distinct ways: to the blocks of variety homeostatically related, to the channels carrying information between them, and to the transducers relaying information across boundaries. Statements about these came to constitute my first three Principles of Organization (Beer, 1979; see Appendix 2). Ashby saw his Law as bearing particularly on the second question, that of channel capacity, probably because he had derived it from Shannon's communication model – which deals with the transmission of information. Indeed he comments that Shannon's Tenth Theorem is a special case of the law of Requisite Variety. Next, and unsurprisingly, he had no difficulty in accepting the identification of transduction as a particular aspect of transmission, and one especially important in management work. but Ashby was not satisfied that requisite variety could be contemplated in terms of relative blocks of variety generators, as my First Principle proposed. Again, it is probable that only information transmission gave operational meaning to requisite variety in his eyes; but in arguing (as he sometimes did) that therefore he had done no more than generalize the Tenth Theorem, I think that he seriously under-rated his own discovery.

Since Ashby was a psychiatrist, I put the counter-case thus. We have a set of mental illnesses, evidently of very high variety – since maybe no two people ever had exactly the same syndrome. There arises quite naturally, and this is an example of requisite variety exerting itself in informational terms, a vast number of 'names' for these illnesses; that is, if we allow that descriptive qualifiers for such generic terms as 'schizophrenia' abound. Unfortunately, however, there is no more than a handful of treatments available: psychoanalysis, convulsive therapy, tranquilization, deep narcosis, surgical intervention… it is difficult to continue. It follows that all the amplifications of channel and transduction variety in the naming is not to the purpose when it comes to managing the illness. Since the syndromes must be mapped homomorphically on to a low-variety therapeutic map, Ashby's Law asserts itself *regardless* of the operational format that is followed.

The point is important in any management process. For just as large numbers of strategies for regulating a firm or an economy can be invented to provide requisite variety, only to be proven useless because they cannot be conveyed through low-variety channels and transducers (and Ashby liked to point this out), so high-variety channels cannot enhance low-variety inputs – unless they contain the intrinsic generative power to be amplified because of the way they are organized *inside* the block. A map-reference has this quality, for instance, and so does a personal file; the policy to 'cut all stocks or costs by 10 per cent' does not.

Limitations

Analogies have limitations; but in a real sense a scientific model as defined should have few – because the transformations it covers are listed and are exactly specified. The problem with analogies is to delineate the contexts in which they are supposed to hold, and then to run the risk that elements will unexpectedly turn up in one system that have no analogues in the other. These dangers are not encountered with scientific models that are properly mapped.

To take an obvious example: Newton's theory of gravitation works very well inside the solar system, give or take the perihelium of Mercury. In a spatiotemporal system that is much larger, Newton must be adjusted by relativity theory. We have, in short, to nominate the context, to fix the boundaries. Now a viable system survives under considerable peturbation because it can take avoiding action, because it can acclimatize, because it accommodates, because it is adaptive, and so on. But put a human in a box, suck out all the air, and s/he dies. We know this, and do not make a lot of fuss about it, because it is an agreed aspect of the definition of viability that there should be a rather closely controlled environment. If we send an astronaut into space, therefore, we equip him with a space suit. We shall certainly not say that our whole conception of viability is faulty because s/he must wear one. On the contrary, one of the most useful products of the manned space programme was its exact specification of a life support system; this indeed fixes the physiological boundaries of viability, though (interestingly) not the psychological boundaries.

Secondly, as to elements which may be recognized in one system and not in another, let us remember that the methodology deals with formal homomorphic mappings and nominates invariances. Anything not so mapped, and anything not determined as a constant, will not be a topic of concern. If it becomes such a topic after the modelling has been done, then its mappings will have to be tested.

Two limitations of the VSM are matters of importance, but they propose no serious misgivings when examined in context and under invariance. The first is often brought up, sometimes in hysterical fashion, by those who notice that people may be the basic elements of a so-called viable system under the VXM rubric. *People* (they say) have free will. Yes, maybe; but people also have constraints laid upon their variety by upbringing, or by the roles that they agree to play in a social unit like a firm. It is true, for example, the liver cannot resign and be replaced by one less gnarled, but what about it? What matters is the functioning of an element, under whatever constraints that the job entails: not the identity of the element itself. And this is just as well for freedom-lovers - let them by all means get out, if the system is oppressive towards them, and they can. It will make no difference to the viable system, *unless* the element has special properties that cannot be replaced. Well, this is simply a matter of nominating what elements in the mapping are to count as invariances. I have known businesses fail because one man was lost, and he accounted for 85 per cent of sales. There is nothing surprising in that. So if the heart of an employee stops beating, that finishes *him* as a viable system. At the next level of recursion, whether that is considered to be his firm or his family or his church or anything else, his loss as an element of this next viable system may or may not be important to its viability. He may simply be replaced; or perhaps *that* system will die too. Obviously, all this will be of high significance to those concerned; but it has no methodological significance to the scientific model within which invariant mappings have been specified in advance.

The second limitation is of more interest, although it can be handled by similar arguments, because it seems to me to be a limitation of society itself rather than a limitation of the model. In either case, it has never been raised with me by anyone at all – at least, not in the terms that are used here. A major battle in biology concerning the possible inheritance of acquired characteristics in the individual, as conceived by Lamarck, seems to have been settled in recent years by microbiologists. There is no such inheritance, for genetic information is always carried by nucleic acid to inform the protein molecule – and never the reverse. In society, however, that is in the social group, there clearly *is* an inheritance of acquired characteristics. Therefore a major difference emerges as between the VSM of the individual and the VSM of society to constitute, at least on first sight, a limitation of the model.

However, as we saw earlier in discussing Ashby and requisite variety, there must always be a barrier (at T) to block the effects of proliferating variety (at D; otherwise results (at E) will reflect the full input variety – and are likely to be quixotic. It seems that in the case of the individual, the gene pool is protected by the encoding of the transformation table (at T). In the case of society, stability in subsequent generations must be ensured by the collaboration of the response with the transformation table (Ashby's R-and-T interaction). Experience shows that this always happens. There is always an element of tradition in the directive correlation of society – that is to say that the transformation table is acting as a block; and there is always an element of novelty coming through from recent outcomes (at E) by regulatory feedback (through R) – that is to say that the response function is acting as an amplifier. So the model can cope with these divergences. The question is whether society itself gets the (R,T) admixture right. Even if it does, it appears to be short of damping mechanisms to prevent uncontrollable oscillations – but that is another story, covered later in System Two of the VSM itself.

The viable system model (VSM)

According to the cybernetic model of any viable system, there are five necessary and sufficient subsystems interactively involved in any organism or organization that is capable of maintaining its identity independently of other such organisms within a shared environment. This 'set of rules' will therefore apply to an organism such as a human being, or to an organization consisting of human beings such as the State.

The comparison is made not by way of analogy, but, as has already been explained, because the rules were developed to account for viability in any survival-worthy system at all.

In very brief, the first subsystem of any viable system consists of those elements that *produce* it (they are the system's autopoietic generators, to use Maturana's terminology). These elements are themselves viable systems. In the limit, the citizens constitute the System One of the State. I say 'in the limit', because the citizens first produce communities and firms, cities and industries, and other viable agglomerations, which are themselves all elements to be included in the State. So a full account of the matter (see *The Heart of Enterprise*) will show how systems of increasing complexity are nested within each other like so many Russian dolls or Chinese boxes to produce the whole. Mention was made at the outset (under 'Provenance') of the discovery of the theorem of recursion and this is where it belongs. 'In a recursive organizational structure, any viable system contains, and is contained in, a viable system.' Out of a five-fold set constituting a viable system, says the model, System One is always a viable system itself. The topology is clearly visible in Figure 4, where (in the first place) one complete viable system fills the page. Inspection will show the five interacting subsystems labelled ONE, TWO, and so on, in capital letters. Among these may be discerned two Systems ONE (there could be more), each of which *contains* a complete viable system displayed at a 45 degree angle.

The whole-page viable system is shown as interacting (see above) in a precisely defined way with its environment through both its Systems ONE, and through its System FOUR, and not otherwise. Equally, the embedded viable systems are shown as interacting in exactly the same way with local environments that are peculiar to each of them – although they are (inevitably) subsets of the whole-page environment. It is vital to understand that the topology of recursion demands an exact replica in each case. In the drawing, the only discrepancy is that the connection between System 4 in the second System ONE and its sub-environment has not been completed, as its twin in the first System is correctly completed, for obvious graphical reasons.

Brief annotations are made in the diagram to indicate the roles of the five subsystems. to enlarge on these within the compass of this paper is not possible without trivializing the elaborate functions of every box and every line, and the reader wishing to investigate the theory itself must be referred to the companion volumes *Brain and Heart* previously mentioned. Some

THERE SEEMS TO BE A PAGE MISSING!!!!!

Theorem, however, each of these black boxes can next be elevated to 'whole-page' treatment – whereupon a new recursion of viable system embedments will be disclosed. The methodology resembles the movement of a magnifying class and an illuminating spotlight down the chain of embedments so the accustomed eye of which I was speaking may now review Figure 4 with its pair of recursions so far described, and discover the outsize square box at the top right-hand side which is the management element of System ONE of the *next higher* recursion; it may also discover the rudiments of the level of recursion *next below* the embedment originally discussed. Thus Figure 4 can be regarded as indicating four levels of recursion out of an arbitrary series (which descends to cells and molecules and ascends to the planet and its universe), of which the middle two recursions receive complete iconic representation.

This is not a claim that an account of a viable system's recursive embedment is ever unique, despite its progression to infinity in both directions, because each viable system figures in an infinite number of chains. Rather is it a manifestation of Hegel's Axiom of Internal Relations: the relations by which terms (or in this case, recursions) are related are an integral part of the terms (or recursions) they relate. Incidentally, if we put the Self as a viable system in the centre of the sphere generated by the infinite set of its recursive chains, then we have a model of selfhood that both expands to embrace the universe and also shrinks to a vanishing grain of sand – a model familiar in oriental philosophy.

This thought leads us conveniently to the recognition that the boundaries of any viable system are arbitrary, as is the number 'five' of its subsystems. The 'fiveness' was due to my efforts to establish the necessary and sufficient conditions of viability, and five was their number; it might have been otherise, if I had used a different rubric. What could not have been otherwise is the fact of the *logical closure* of the viable system by 'System Five', whatever its number: only this determines an *identity*. Nominating the components of System Five in any application is a profoundly difficult job because the closure identifies self-awareness in the viable system. 'What business are we in?' asks the Manager. But who are 'we'? Shareholders, employees, managers, directors, customers, taxmen, environmentalists . . . all these have different answers to offer. 'What business is the self in?' – see above.

I have repeatedly told the story (for instance, see *Brain*) of how President Salvador Allende in the Chile of 1972 told me that System Five, which I had been thinking of as himself, was in fact the people. Then perhaps the president embodies the people; or perhaps the presidency is overtaken by a gang of thugs, as was to happen in 1973. (For some recent discussions of this example, see Beer, 1983.) At any rate, it is clear that the determination of closure, and thus the recognition of identity and self-awareness, in any viable system is an outstanding example of the observer's imputations of *purpose* to that system that are probably idiosyncratic. There are ideological traps: for example, the biggest confusion in which I was ever professionally involved concerned the purpose of a health system, to which there are as many answers as interests involved. There are teleological fallacies: think once again of selfhood . . .

These difficulties are not indications that the VSM 'doesn't work': the model does not create the problems that it makes explicit. Rather does it enable managers and their consultants alike to elaborate policies and to develop organizational structures in the clear understanding of the recursions in which they are supposed to operate, and to design regulatory systems within those recursions that do not pretend (as do so many of those we employ) to disobey the fundamental canons of cybernetics.

The pathology of the viable system

Many people dislike to see the word 'pathology' written in such a context as this, because the theory of the viable system may be dealing with societary units, or even with such entirely inanimate systems as computer-based communication networks. Some of these people would be placated if the word in the title were set in inverted commas. The fact is, however, that either we have a theory of viability, meaning 'capable of independent existence', or we have not. The possibility of such a theory is anti-paradigmatic within the subculture, true; but that paradigm is overdue for change: see Capra (1982). The risk of making mistakes under any methodology of analogy is great, true; but we have been at pains to show that an heuristic such as the yo-yo technique is in search of a mathematical invariance that transcends analogy. A viable system made of metal could be melted down, true, and one made of people could be disbanded, true; but the feotus of eight months is the classic example of a viable system, and many conditions of existence are attached to its capability for independence too. In short, the opponents of 'biological analogies' are often the first to misapply them when they try to make their own case, thanks to an uncritical belief in the properties of protein-based machines which in fact work only within rather narrow physiological limits.

According to these cybernetic enquiries, practised, as has been said, in many countries over many years, viable systems of all kinds are subject to breakdown. Such breakdowns may be diagnosed, simply in the fact that some inadequacy in the system can be traced to malfunction in one of the five subsystems, where in turn one of the cybernetic features that compose the rules (cf. Appendix 2) will be found not to be functioning. To continue unabashed with medical-sounding talk that is in fact wholly appropriate to the cybernetics of viability: the etiology of the disorder may be traced, a prognosis may be prepared, and antidotes (even surgery) may be prescribed.

Subjectively speaking, confidence in the VSM as applied to societary systems derives not so much from the fact that the pathology of the viable system can be investigated with ease, as from the *speed* with which the diagnosis can be made. The knowledgeable user may expect to 'home in' on (say) half-a-dozen causes of concern within a day or so of exposure to the real-life system, and it is a frequent experience to find such danger points when they have been deliberately concealed out of embarrassment or self-serving: they tend to signal themselves. Interestingly enough, such incidents tend to enhance the confidence not only of the VSM-er, but of the client management itself.

A question often asked is this: if we are dealing with an organization that exists, that is actually there to be investigated, then surely it is by definition a viable system – and nothing remains to be said? This is where the pathological vocabulary becomes so useful. The fact that the societary system is there does not guarantee that it will always be there: its days may well be numbered, and many have been the 'buggy-whip' companies to prove it. The fact that it is there does not prove that it is effectively there, witness universities, nor efficiently there, witness hospitals. Monoliths and monopolistic systems in particular (such as these two) often operate at the margins of viability, creaking and choking like the valetudinarian organizations that they are. Moreover, many such are operating at such an enormous cost that they are becoming less and less viable in front of everyone's eyes.

One of the main reasons for this, particularly the social services, is that people looking for cheaper ways of doing things attempt to repeal the Law of Requisite Variety itself. Policing, for example, whether by the police themselves in terms of crime, or by environmental agencies in terms of pollution, or by health scanners of pre-symptoms, often fails to recognize that only variety can absorb variety. A great many examples are reproduced elsewhere (Beer, 1975).

Next, there are four diagnostic points made in a learned journal (*Brain and Strategy*, 1983). All four have been expounded in my own writings, but not I think with such pith; therefore I take leave to reproduce them here as direct quotations.

1. *Is management presiding over a 'viable system'?*

If any of Beer's five necessary functions are removed from, say, a subsidiary, then its abilities to operate successfully may well be killed. This could perhaps involve taking away a subsidiary's freedom to invest its financial surpluses or removing its sales function, for example.

2. *Does subsystem Five truly represent the entire system within the context of larger, more comprehensive and more powerful systems?*

If this function, or subsystem, is unable to find a way to represent the essential qualities of the whole system to the larger meta-system, then the system's survival is in question.

3. Do managers often fail to understand the need for subsystems Two and Four?

Business people have little difficulty recognizing the need for subsystems One, Three and Five. If Two is missing, activity in One can turn deadly and self –defeating as units fight for resources and against entropy; if Four is missing. Three and Five can collape into each other, leaving the critical Five subsystem a mere functionary.

4. Do the Three, Four and Five subsystems need to form a Three-Four-Five subsystem to encourage 'synergy' and interactivity?

Without a constant interaction and exchange of information between these three functions, Three is vulnerable to 'narrow tunnel' syndrome and Four is exposed to the perils of 'flights of imagination'.

Not only are these points extremely cogent and penetrating, they well illustrate how the structure and the language of the model make possible the expression of elaborate and/or subtle comments in very few words. Let me add a few remarks on each of the indicated pathologies, drawn from experience.

(i) Subsidiaries that are 'taken over' are always painstakingly assured that their individuality will be preserved, their autonomy respected, and so on. After all, the argument (very plausibly) goes, your individuality, your reputation, your goodwill, your people are all assets for which we have paid hard cash – naturally we shall nurture them. This is poppycock – although it is often believed by the takeover bidder himself. A study of the embedment of the new System One in terms of the Law of Cohesion (see Appendix 2) will reveal how the inter-connectivity between the subsystems of the two recursions inevitably takes up variety from the new subsidiary. In the VSM, 'autonomy' is a precisely defined term, and it does *not* mean zero interference. Incidentally, if the taking-over company makes the mistake of leaving intact all the new subsidiary's variety (or of handing over too much variety to an old subsidiary), this company is very likely to be the subject of a reverse takeover bid.

(ii) This is an issue of identity. The work here reported has repeatedly encountered situations in which all manner of adjustments have been necessary to make the viable system secure in a changing environment. That is, adaptation is evoked (in those situations) as a key characteristic of viability, and much change ensues. Will the system still be able to recognize itself? More particularly, will others be able to recognize it? Philosophers used to ask whether 'this apple' were still 'this apple' after a large bite had been taken out of it …. *The Heart of Enterprise* includes a highly sophisticated Test of Identity with this point in mind.

(iii) The collapse of Five into Three (in the effective absence of Four) is made particularly likely insofar as Five people have usually been promoted from Three. They are uncomfortable as demi-gods with no clear duties beyond being wise and pleasant. Thus, when something goes wrong in System Three (or even One), they are likely to dive down into the problem that they understand so well – never to emerge again. They may be seen around, but only as their previous Three incarnation – erratic and abrasive as ever. But the collapsed metasystem is a special pathology. It is a decerebrate cat, pinned out, intravenously fed. It responds *reactively*, from the autonomic command centres at three, and is incapable of planning and foresight (Four) and will and judgment (Five). but it will react to prods by a reflex kicking-back. With no apologies to those complaining about biological metaphors, who knows an organization that is a decerebrate cat?

(iv) The attention drawn to this problem is well merited. It is the intellectual springboard for recognizing the value of an operations room, or (a better term) management centre. In such an 'environment of decision', as I have called it, the Three-Four-Five metasystem has a chance to find its own cohesion, and to operate in a nutrient medium.

Obviously it would be possible to comment on every feature of the viable system from the standpoint of its pathology. But that would be boring; and perhaps the above discussion of some already profound points sufficiently gives the flavour of the pathologist's commentary.

But it may be worth ending with a suggestion which this discussion seems naturally to propose – in medical practice, there is such a thing as post mortem examination. Much knowledge of viable systems has been gained by the study of those that are viable no more. I have done some work of this kind, but only as the result of being fortuitously present at the deathbed. The suggestion would be that a small team of organizational pathologists should be formed, ready to rush to the scene of any incipient organizational demise. Of course, these people would not be loitering about, waiting for something to happen. They would be organized more like a lifeboat crew.

The first imperative would be to resuscitate the moribund victim. Failing that, however, a post mortem would be performed before rigormortis had set in, and before those nearest to the deceased had closed in like the vultures they often emulate. I have certainly noticed many times how history is rewritten in these circumstances with breathtaking speed. It happens with people too.

APPENDIX 1: The theory of the model in operation research[1]

If we call the set M of elements a totality of world events which we propose to examine then the systemic configuration of events which we know about is a sub-set A of set M. If we call the set of N of elements b the totality of systemic science, then the configuration of system which we ourselves understand is a sub-set B of set N. The process of creating a systemic model may then be described as a mapping f of A into B. By this I mean that for every element $a \in A \subset M$ there exists a corresponding element $b \in B \subset N$, and thus $b = f(a)$. The image of the sub-set A, namely, $f(A) \subset N$, is the model. If we are able to exhaust the elements of A and to nominate their images in b, we have every hope of creating an isomorphic model. This means that there exists a complete inverse image of B under mapping f in M, so that $f(A) \subset N = f^{-1}(B) \subset M$. This is the state of affairs, expressed group-theoretically, which the operational research man is trying to reach.

Now an isomorphism is important because it preserves the structure of the original group in the mapping. Typically, if it is possible to perform additions inside set M, those additions will remain valid

[1] Extract from *The World, the Flesh and the Metal* (the 1964 Stephenson Lecture). Reprinted by permission from *Nature*, vol.205, No.4968. pp. 223-231.

when the same operations are performed on the images of their elements in set N. It is this persistence of relationship when the mapping is done which makes a model operate as a model. So, if a_1 and a_2 when added together equal a_n in set M, it can be shown that $f(a_1)$ plus $f(a_2)$ must equal $f(a_n)$ in set N. Now comes the interesting comment. The conditions can be set up in which the same answers $f(a_n)$ in set N is obtained from the mapping f whether the transformation is effected before or after the mapping occurs. That is to say, we may either add the original elements in M and transform the answer under f, or we may transform the original elements first and then add them. The result will be the same. Formally: $f(a_1 + a_2) = f(a_1) + f(a_2)$. When one group is mapped into another group and this condition is generally fulfilled, the mapping is called homomorphic.

These elementary definitions are included so that the argument can be made quite clear. Because it is possible to coalesce elements of M before transforming them, without losing the capability of mapping to preserve structural relationships as discussed, it is clear that a homomorphism may have fewer elements than its inverse image. In the case of the model, then the mapping of A into B turns out to be a mapping *on* to a sub-group of B. Isomorphism turns out to be a special case of homomorphism, in that $f(A) \subset B$ turns out to mean $f(A) = B$: the one-one correspondence of elements with which we begin is maintained. But for any other sub-group of B other than B itself, homomorphism involves a many-one correspondence, and the inverse mapping $f^{-1}(B)$ will not exhaust the elements of A.

It is suggested, then, that the models of big systems that we entertain are homomorphisms of those systemic characteristics of the big system that we can identify. The Homomorhic group $f(A) \subset B \subset N$ is the particular model we use. It is in practice extremely difficult to include in this model all the features we recognized in A, and typically we do make the many-one reductions mentioned. Thus, for example, we undertake production costings as if the behaviour of all three shifts in a works were indistinguishable, and as if two similar products were identical, and as if materials were consistently uniform – although we actually know that none of these simplifications is true. Then the effectiveness of the model as predictive depends on the choice of an effective transformation by which to map. If we add up the outputs of three shifts and then transform the answer by some mapping into the model, it is no use supposing that any calculation, comparison or prediction undertaken in the model can be worked backwards through an inverse mapping which will distinguish between the shifts. On the other hand, it is necessary to handle only a third of the elements we know about inside the model. A definite choice has been made to jettison modelling-power in favour of economy in the recording and handling of data. This is acceptable, so long as the choice is deliberate rather than accidental, and so long as it is remembered as a limitation in the model.

Secondly, however, there is a further loss of modelling power in the facts that A is a sub-set of M and B is a sub-set of N. Now an interdisciplinary team of scientists can minimize the losses of modelling power due to $B < N$. Because such a team can examine all the major sub-sets of N before deciding to use one specific group B; it may even experiment with other groups too. But the losses due to $A < M$ are more serious, and may be disastrous to the exercise. For if what we recognize in a big system is not what is really important about its systemic character, the ability to predict A may not help much in M. In other words, A is itself a homomorphic mapping M, and one which by definition we cannot properly specify. Remember that $M - A$ was acknowledged to be systemically unrecognized from the start. We may know that our knowledge of a big system does not exhaust it, without having the faintest idea of the character of the knowledge that is missing.

It is hoped that this attempt somewhat rigorously to formulate what goes on in model-building will prove helpful in pin-pointing what we can and cannot do. The ordinary operational research exercise works, and we can see why. It is possible to advance what we understand about a stock-holding system, for example, to the point where A approaches M asymptotically. It is possible to examine most B of N, which is to say most scientific approaches to the scientific totality of understanding about such systems. If we know what the stock-holding system can do, if (as the operational research man would say) we can definite its criteria of success or objective function, then we can define a homomorphic mapping f of A ? M on to B ? N which preserves the stochastic relationships in which we are interested. More specially, we can do this in a way that the inverse image of B under mapping f yields a set of $f^{-1}(B)$ of elements in the real system M which are useful.

The difficulties about doing successful operational research in various circumstances can now be made quite specific. First, the modelling will not on the average work well if $N - B$ is large: this happens if the operational research team is not corporately versatile. Secondly, the modelling will not work at all unless f is well defined: this entails good empirical research into what the system really has to do. Thirdly, the predictions of the model will be of no actual use if a modelled outcome $f\ (b_1 \text{----} b_n)$ turns out to have a pragmatically undiscriminating inverse f^{-1} f in A. This also entails good empirical research into the forms of many-one reduction. Fourthly, the modelled predictions though useful will not exert what could be called control unless the $M \longrightarrow A$ homomorphism captures the systemic character of the big system *in extenso*. This again appears to be a matter for good empirical research, although there is more to say.

Contrary to increasingly current belief, then, operational research is empirical science above all. The mathematical models dreamed up in back rooms are useless unless they can meet the four kinds of difficulty enumerated, and this cannot be done remotely from the world.

Appendix 2: Glossary of rules for the viable system [1]

Aphorisms

The first regulatory aphorism

It is not necessary to enter the black box
to understand the nature
of the function it performs (p.40)

The second regulatory aphorism

It is not necessary to enter the black box
to calculate the variety
that it potentially may generate (p.47)

Principles

The first principle of organization

Managerial, operational and environmental varieties,
diffusing through an institutional system, tend to equate;
they should be dsigned to do so with minimum damage to
people and to cost. (p.97)

The second principle of organization

The four directional channels carrying information between
the management unit, the operation, and the environment
must each have a higher capacity to transmit a given amount
of information relevant to variety selection in a given time
than the originating subsystem has to generate it in that
time. (p.99)

[1] Extract from *The Heart of Enterprise* Beer, 1979) to which book the page numbers refer.

The third principle of organization

Wherever the information carried on a channel capable of distinguishing a given variety crosses a boundary, it undergoes transduction; the variety of the transducer must be at least equivalent to the variety of the channel (p.101)

The fourth principle of organization

The operation of the first three principles must be cyclically maintained through time without hiatus or lags. (p.258)

Theorem

Recursive system theorem

In a recursive organizational structure, any viable system contains, and is contained in, a viable system (p.118)

Axioms

The first axiom of management

The sum of horizontal variety disposed by *n* operational elements
equals
the sum of vertical variety disposed on the six vertical components of corporate cohesion (p.217)

The second axiom of management

The variety disposed by System Three resulting from the operation of the First Axiom
equals
the variety disposed by System Four (p.298)

The third axiom of management

The variety disposed by System Five
equals
the residual variety generated by the operation of the Second Axiom (p.298)

Law

The law of cohesion for multiple recursions of the viable system

The System One variety accessible to System Three of Recursion *x*
equals
the variety disposed by the sum of the metasystems of Recursion *y* for every recursive pair. (p.355)

Appendix 3: Some applications of the Viable System Model

Applications of the VSM by its author during the evolution and verification of the model have been so many and so widespread as to defy a proper listing. For the record, however, the range of amenable organizations ought to be indicated, leaving case histories to the published papers and books. Small industrial businesses in both production and retailing, such as an engineering concern and a bakery, come to mind; large industrial organizations such as the steel industry, textile manufacturers, shipbuilders, the makers of consumer durables, paper manufacturers are also represented. Then there are the businesses that deal in information; publishing in general, insurance, banking. Transportation has figured: railways, ports and harbours, shipping lines. Education, and health in several countries), the operation of cities, belong to studies of service. Finally comes government at all levels – from the city, to the province, to the state and the nation-state itself – and the international agencies: the VSM has been applied to several.

In this opening paragraph we have been talking of one man's work. Obviously, then, these were not all major undertakings, nor is 'success' claimed for massive change. On the other hand, none of these applications was an academic exercise. In every case we are talking about remunerated consultancy, and that is not a light matter. The activities did not necessarily last for very long either, since speedy diagnosis is a major contribution of the whole approach. On the other hand, some of them have lasted for years. Undoubtedly the major use of this work to date was in Chile from 1971-73: five chapters ending the second edition of *Brain* describe it in full (Beer, 1981). As this is written, however, a new undertaking on a similar scale is beginning in another country. On the question of what constitutes 'success' in consulting; reference may be made to page 211 of this book.

Of other people's work in the field of managerial cybernetics that has made application of the VSM, first mention must go to Raul Espejo. He has given his own account of the 1971-73 Chilean application that we undertook together (Espejo, 1980a). Since then, his teaching and research at Aston University in England has been centred on the VSM, and outcomes have been published in several articles and papers (especially Espejo, 1978, 1980b). His diagnoses have been profound, and he is adding to the corpus of theory.

The number of senior degrees, including doctorates, that have employed the VSM under Espejo's direction is already in double figures. Professor David Mitchell's teaching has generated a similar number of postgraduate theses using the VSM at Concordia University in Quebec, as has that of Professor Manuel Marina at the Central University of Venezuela. Several more have emerged from Brunel University, under the directiion of Professor Frank George. In the United States, Professors Richard Ericson and Stuart Umpleby (at George Washington University), Professor Barry Clemson (at the Universities of Maryland and of Maine), and Professor William Reckmeyer (at San José State University) have all made extensive use of this teaching, and others from Australia to India have reported similarly.

At Manchester University in the Business School, Geofrey Lockett (directing the doctoral programme) has sponsored whole-week 'experiences' of the VSM; and Professor Roger Collcutt has invented a unique pedagogic framework whereby MBS students undertake projects to apply the VSM to functional management, subsequently to merge the insights gained into a general management picture. Another novel development has been made by Ronald H. Anderton in the Systems Department of Lancaster University: practical applications of the VSM in the form of project work have for some years been an important part of his *under*graduate teaching.

A veritable kaleidoscope of applications of the VSM has been presented by Dr Paul Rubinyi in Canada. From penological systems to health services in the public sector, from oil companies to what cooperatives in the private sector, and from provincial planning to air transportation in federal government: every kind of organization has been mapped, in virtually continuous work over the last 13 years.

Other separate applications in Canada include the work of Walter Baker, Raoul Elias and David Griggs on the Fisheries and Marine Service, which took unique advantage of managerial involvement, and that of Raoul Elias for Gaz Metropolitan. David Beatty has used the model for educational planning in Ontario, and I believe that it has been in independent action on the West Coast as well (Baker, Elias and Griggs, 1978).

In Latin America, Professor Jorge Chapiro is a leading exponent of the VSM who consults over the whole spectrum of industrial and governmental management in several countries.

In Australia, applications in an insurance company have been made by J. Donald de Raadt; in Switzerland Dr. Peter Gomez has used the VSM in a publishsing company, making an interesting experiment in melding this methodology with the 'root definitions' of Professor Peter Checkland (Gomez, 1982). In wider fields still we find a useful VSM application in Finland by Dr. S. Korolainen to ekistics (Korolainen, 1980); and David Noor has published 'A viable system model of scientific rationality' as a working paper from the University of Western Ontario.

On the strictly biological side, but not from the original neurophysio-logical perspective, Dr. Richard Foss in England has made many mappings: for example, on the Eukaryote cell, the annual plant and the honeybee colony. He has found the VSM to hold in such diverse systems; and he is extending the work to the slime mould *Dictyoltelium*, to lichens and to vertebrates, considering both the evolution and ontogeny of each system.

It does appear that the VSM has sufficient generality to justify its origin as an attempt to discover *how systems are viable*; and that it also generates considerable power to describe and predict, diagnose and prescribe. No systematic archive of applications has been kept: perhaps it would be helpful to start one. These notes are compiled from such recollections and records as happen to be to hand.

References

Ashby, W.R. (1956). *Introduction to Cybernetics*. London: Chapman & Hall.
Baker, W., Elias, R. and Griggs, D. (1978). 'Managerial involvement in the design of adaptive systems', in *Management Handbook for Public Administration* (ed. Sutherland, J.W.). New York: Van Nostrand Reinhld.
Beer, S. (1959). Cybernetics and Management. London: English Universities Press.
Beer, S. (1960). 'Towards the cybernetic factory', in *Principles of Self Organization* (symposium).
Oxford: Pergamon Press.
Beer, S. (1965). 'The world, the flesh and the metal', *Nature*, **205**, 223-31.
Beer, S. (1966). *Decision and Control*. Chichester: John Wiley.
Beer, S. (1972). B*rain of the Firm*. Harmondsworth: Allan Lane.
Beer, S. (1975). *Platform for Change*. Chichester: John Wiley.
Beer, S. (1979). *The Hart of Enterprise*. Chichester: John Wiley.
Beer, S. (1981). *Brain of the Firm*, 2nd edn. Chichester: John Wiley.
Beer, S. (1983). 'A reply to Ulrich's "Critique of pure cybernetic reason: the Chilean experiment with cybernetics"', *Journal of Applied Systems Analysis*, **10**, 115-19.
Beer, S. (1985). *Diagnosing the system for Organizations*. John Wiley.
Brain and Strategy (1983). vol. 4, issue 9.
Capra, F. (1982). *The Turning Point*. New York: Bantam Books.
Esperjo, R. (1978). Multi-oganizational strategies; an analytical framework and case, in *Applied General Systems Research*: *Recent Developments and Trends* (ed. Klir, G.). New York: Plenum Press.
Espejo, R. (1980a). 'Cybernetic praxis in government: the management of industry in Chile 1970-1973', *Journal of Cybernetics*, **11**, 325-38.
Espejo, R. (1980b). 'Information and management: the cybernetics of a small company', in *The Information Systems Environment* (eds. Lucas, H., Land, F., Lincoln, T. and Supper, K.). Amsterdam: North Holland.
Gomez, P. (1982). 'Systems methodology in action', *Journal of Applied Systems Analysis*, **9**, 67-85.
Korolainen, S. (1980). *On the Conceptual and Logical Foundations of the General Theory of Human Organizations*. Helsinki School of Economics.
Shannon, C. and Weaver, W. (1949). *The Mathematical Theory of Communication*. University of Illinois Press.
Sommerhoff, G. (1950). *Analytical Biology*. London: Oxford University Press.