

Fanfare for Effective Freedom Cybernetic Praxis in Government

Stafford Beer

1973

The Third Richard Goodman Memorial Lecture, delivered at Brighton Polytechnic, Moulsecoomb, Brighton, on Wednesday, 14th February, 1973.

This is the first memorial lecture I have given for a man I knew personally—a man whom I also loved. He was a tenacious cybernetician; the pioneer of that work here in Brighton, but one whose name at least was known throughout the cybernetic world. More than this, and more importantly than this, he had a dedication to humanity. It may not be well known, but I knew, that he was as interested in the cybernetics of society as he was in the more recondite mathematics of the science. And I also know very well that he would have been captivated by the unfinished story I am telling here formally for the first time. If I could have had his advice while the project was unfolding, it might have been a better story. But I still hope that it is worthy of his memory.

In November 1970, Dr. Salvador Allende became President of the Republic of Chile. In November 1971, after some letters had passed, a meeting held in London, and some homework done, I arrived in Santiago. There I first met the prepared group of a dozen men who formed the nucleus of a team which is now much larger, and with whom I am still working for. I have been commuting the 8000 miles between London and Santiago ever since. The charge was daunting indeed: how should cybernetics be used in the exercise of national government? You will note that the question whether cybernetics had any relevance to the problems of society and of government had already been answered affirmatively.

What was and is the situation? The answer, as I have intimately known it for these last eighteen months, is immensely complicated. Let me paint my own crude picture for you, with a rapid brush. First, more than half the total population lives an urban life in the small central region of this long, thin country—a region that perfectly balances the arid North and the wet South in a superb climate. Here the people are highly literate, and constitutionally minded; their men are frank and friendly, their women gorgeous and gay. There is as great a spirit of freedom in the air as I have sensed anywhere in the world—and decreasingly sense in so much of it today. Yet, as you must surely know, Chile is in the middle of a Marxist revolution that has so far been constitutional, so far legal, so far bloodless.

On the land, the previous government had begun a process of agrarian reform, and that policy had general agreement. Landowners would no longer control estates larger than eighty hectares—say about 200 acres. The residual land was split up, and handed to worker's cooperatives, who have the support of government agencies. In the six years of that previous government, about 20% of the programme was implemented. But the people were impatient, especially in the South, and a deeply embedded bureaucracy slowly moves. New forms of expression were given to agrarian reform; and the programme was completed, not always in good order, in the first two years of the government of Popular Unity. This rate of change has surely contributed to the current food shortage; not so much, perhaps, because the new arrangements are inefficient in themselves, but because the remaining landowners—disrupted by these events and fearful of further change—are eating their seed corn rather than investing it in production.

In industry too, the new government's policies of nationalization and worker participation have been implemented so rapidly that the control of that process was—and remains—extremely difficult. Foreign managers of expropriated firms have mostly left the country, and the problem of finding men to take temporary charge (these are the interventors) was—and remains—severe. It has been exacerbated by a brain drain of native Chileans: too many qualified professionals have left the country. That they should do so was surely implicit in their upbringing and their expectations; but their problem was much aggravated by the psychological panic induced by Opposition campaigns to spread rumours of terrors

to come. As to industrial investment, we should note that all the banks were nationalized, and those banks hold the internal assets of the landed classes.

Politically, the government's problems have been huge, all along. In the presidential election that put Dr. Allende in power, he obtained only 36% of the vote. The coalition he leads itself contains factions which struggle for influence between themselves. Throughout he has faced a hostile Congress and Senate, capable of blocking any government initiative by the Oppositions' majority of 60% to 40%. On the other hand, the government is empowered to block the majority vote of Congress—so long as its own support is at least a third. Hence the political stalemate; hence the tension of the marginal vote; hence the importance of the Congressional Election next month.

All of this is easily recognized, especially in cybernetic terms, as a grossly unstable situation. And its explosive economic tendencies were perfectly predictable when I first became involved. There had been a very large and very sudden increase in the purchasing power of the rank and file. Wages rose fast, for the land-workers in particular—who were put on the same footing as the blue collar workers. Social security benefits were much increased for everyone with young, old or incapacitated dependents. Then clearly there would be a run on stocks; clearly there would be a run on reserves. Indeed this was well understood: on my very first visit a Minister took several hours to explain the risks being run, and the political determination with which those risks were accepted as the price of rapid social progress. The question was whether the government could get a sufficient grip on the situation in time—before this inflationary time-bomb blew up in its face.

In the event it did not, and the state of the country is very precarious. It is superficial to think of this in terms of food shortages and 'housewives marches', tiresome as the food problem certainly is for the middle class. The more important fact is that Chile suffers from the effects of an economic blockade. There has been a blockade of spare parts, which has made it even harder to keep agriculture going, industry productive, transportation moving. There has been a blockade on exports, by which I refer especially to copper—which used to earn more than eighty percent of the country's foreign exchange. The attempt is being made to close world markets to Chilean copper, and the world price has fallen. Above all, there has been a blockade on foreign credit. And since Chile's natural resources will one day make it a rich country, when those resources are properly deployed, it follows that the stranglehold on credit is not a solely economic matter.

It appears to me that the government did not anticipate the full vindictiveness with which the rich world would react to its actions, which I emphasize have—so far—been perfectly legal. At any rate, a true resolution of the very potent conflicts in Chilean society is not discernible within the mounting instability, and may be long postponed. But I consider that this is largely a phenomenon of the cybernetics of international power: you could say that the Chilean people have not been given a chance. They are being systematically isolated behind those beautiful Andes mountains, and are in a state of siege. The mass media have not helped much—especially inside the country itself, where freedom of speech has been respected in very testing circumstances. Because of its ownership, this freedom is largely employed to oppose the government. Because of its prestige, the anti-government press is widely copied—embroidered even—across the world.

It says a lot for the good intentions of the government that the work I shall describe has been going on in the midst of such obvious turmoil. It wanted scientific tools to help tackle the country's problems, and it knew that their provision would take time—perhaps too long. So it may be proved. The government has so far had to work with the tools that other governments have used without success. It also wanted to work out the relationship between science and the people, and that too ought to interest us all. We have moved into an epoch in which the misuse of science has created a society that is already close to a technocracy. The very language—the dehumanized jargon—in which powerful countries talk about the wars they wage, or powerful companies talk about the people they exploit, frankly makes me vomit.

I am a scientist; but to be a technocrat would put me out of business as a man. Yet there I was, eighteen months ago, intent on creating a scientific way of governing. And here I am today, proud of the tools we have made. Why? Because I believe that cybernetics can do the job better than bureaucracy—and more humanely too. We must learn how to expunge technocracy, without rejecting science—because the proper use of science is really the world's brightest hope for stable government. Some people in Chile share that view; and they reject technocracy as strongly as do I. All of us have already been misrepresented in that respect, just as the scientific work we have done has already been misrepresented as analogous to other management control systems that have failed. Both comments miss out the cybernetics, to discuss which we are here—and a subject which, for government in general, is not at all understood.

CYBERNETICS AND FREEDOM

What is cybernetics, that a government should not understand it? It is, as Wiener¹ originally called it twenty-five years ago, 'the science of communication and control in the animal and the machine'. He was pointing, in that second phrase, to laws of complex systems that are invariant to transformations of their fabric. It does not matter whether the system be realized in the flesh or in the metal.

What is cybernetics, that government should need it? It is, as I should prefer to define it today, 'the science of effective organization'. In this definition I am pointing to laws of complex systems that are invariant not only to transformations of their fabric, but also of their content. It does not matter whether the system's content is neurophysiological, automotive, social or economic.

This is not to argue that all complex systems are really the same, nor yet that they are all in some way 'analogous'. It is to argue that there are fundamental rules which, disobeyed, lead to instability, or to explosion, or to a failure to learn, adapt and evolve, in any complex system. And those pathological states do indeed belong to all complex systems—whatever their fabric, whatever their content—not by analogy, but as a matter of fact.

With cybernetics, we seek to lift the problems of organizational structure out of the ruck of prejudice—by studying them scientifically. People wonder whether to centralize or to decentralize the economy—they are answered by dogmas. People ask whether planning is inimical to freedom—they are answered with doctrines. People demand an end to bureaucracy and muddle—they are answered with a so-called expertise which, from its record, has no effect. If dogma, doctrine and expertise fall to give effective answers, then what criterion of effectiveness shall cybernetics use? My answer to this question is: the criterion of viability. Whatever makes a system survival-worthy is necessary to it.

Necessary, yes, one might reply; but surely not also sufficient? The more I consider that criticism, the less I see its force. Suppose one were to say, for example (pleading necessity), that since a particular anarchic society is falling apart, a high degree of autocracy will be needed to ensure its survival. Then the critic might say: but this way lies totalitarianism and the loss of human freedom. Not so, if we adhere to our viability criterion. Because that society would be unstable also: sooner or later would come a revolution—it always does. Suppose one were to say (pleading necessity) that a particular repressive society must throw over all constraint. Then the critic might say: then you will have chaos, and no-one will be safe. But that situation would not conduce to survival either, and the pendulum would swing the other way—it always does. The point is that a truly viable system does not oscillate to those extremes, because it is under homeostatic control in every dimension that is important to its survival. Then when it comes to designing systems of government, we need to understand the cybernetic laws of homeostasis. Fortunately, and thanks mainly to Ross Ashby², we do understand.

Let me briefly explain. Homeostasis is the tendency of a complex system to run towards an equilibrial state. This happens because the many parts of the complex system absorb each other's capacity to disrupt the whole. Now the ultimately stable state to which a viable system may run (that state where its entropy is unity) is finally rigid—and we call that death. If the system is to remain viable, if it is not to die, then we need the extra concept of an equilibrium that is not fixed, but on the move. What causes the incipiently stable point to move is the total system's response to environmental change; and this kind of adjustment we call adaptation. The third notion that we need to understand homeostasis is the idea of a physiological limit. It is necessary for a viable system to keep moving its stable point, but it cannot afford to move it so far or so fast that the system itself is blown apart. It must keep its degree and its rate of change within a tolerance fixed by its own physiology. Revolutions, violent or not, do blow societies apart—because they deliberately take the inherited system outside its physiological limits. Then the system has to be redefined, and the new definition must again adhere to the cybernetic criteria of viability. Then it is useless for whoever has lost his privileges to complain about his bad luck, so long as he uses a language appropriate to the system that has been replaced. He must talk the new language or get out. This fact is the fact that is polarizing Chilean society now.

By the same token, a society that does not have a revolution, violent or not, inevitably goes on talking the inherited system's language, even though the rate of change has made it irrelevant to the problems which that society faces. Perhaps this fact is the fact that begins to polarize British society now.

At any rate, cybernetic analysis—I have tried to give you merely its flavour—enables us to study the problems of a particular society in terms of its viability. In general, I have only this to say about societal homeostasis in the nineteen-seventies:

- A homeostat works (and we know all the cybernetic rules) by moving its stable point in a very complicated response to the shocks it receives to its total system.

- Any homeostat takes a finite time to re-establish its new stable point. This is called the relaxation time of the system.
- Today it is typical of social institutions that the mean interval between shocks (thanks to the rate of change) is shorter than the relaxation time. That is because the institutions were originally designed to accept a much longer interval between shocks.
- From this it follows that societal institutions will either go into a state of oscillation, or plunge into that terminal equilibrium called death.

The cybernetician will expect the politician to adopt one of two basic postures in the face of these systemic troubles.

The first is to ignore the cybernetic facts, and to pretend that the oscillations are due to some kind of wickedness which can be stamped out. The second is to undertake some kind of revolution, violent or not, to redesign the faulty instruments of government. I do not have to relate the polarization throughout the entire world to which this cybernetic expectation is the key. But it seems very clear to me, as a matter of management science, that if in these typical circumstances you do not like violence, then you should quickly embark on a pacific revolution in government. If you do not, then violence you will certainly get.

Outstandingly it was Chile that embarked on this recommended course of pacific revolution. But, as I have already argued, the process has strained Chile's internal homeostatic faculties to the breaking point. Let me restate the reasons I gave before in cybernetic terms. Firstly it is because its minority government has been frustrated in fully restructuring the system according to the criteria of viability. Secondly it is because in the wider world system Chile's experiment was observed as an oscillation to be stamped out. How this will end I do not know. Meanwhile, however, we had set out to redefine the internal homeostasis.

I went to Chile armed with a model of any viable system, which I very well understood. It had taken twenty years to develop, in modelling, testing, and applying to all manner of organizations. The book expounding it³ was already in the press when this story started.

One of the key ideas the general theory embodies is the principle of recursion. This says that all viable systems contain viable systems, and are contained within viable systems. Then if we have a model of any viable system, it must be recursive. That is to say, at whatever level of aggregation we start, then the whole model is rewritten in each element of the original model, and so on indefinitely.

If we model the state, then one element is the economic system; if we model the economic system, then one element is an industrial sector; if we model that industrial sector, then one element is a firm. The model itself is invariant. See what happens if we go on with this recursion. An element of the firm is a plant; an element of the plant is a particular shop; an element of the shop is a section; an element of the section is a man. And the man is assuredly a viable system—as a matter of fact, the model started from the cybernetic study of man's effective neurophysiological organization in the first place.

A second key idea was that by using the viability criterion, all alone—for the reasons I gave earlier, one might succeed in identifying regions of policy in the total organizational space that represent homeostatically stable points for long term survival. I am pointing now to a possibility that it is open to mankind at last to compute a set of organizational structures that would suit the needs of actual men—as being at once themselves independent viable systems with a right of individual choice, and also members of a coherent society which in turn has a right of collective choice. Now one of the main issues identified was the issue of autonomy, or participation (these are catch words), or perhaps I mean just liberty, for whatever element within whatever viable system. Then this means that there ought to be a computable function setting the degree of centralization consistent with effectiveness and with freedom at every level of recursion. This I now believe. It is a bold claim. Let me try to give it verisimilitude.

Government and management control systems range over a fairly wide spectrum on the autocratic-permissive scale, and still remain viable. What is happening in cybernetic terms is that the homeostat connecting 'the boss' to the people's homeostat is either in high or low gear—while still operating within physiological limits. In an autocratic system, the people's homeostat is robbed of flexibility; in a permissive system, it is deprived of guidance and help. As long as oppression and freedom are seen solely as normative values, the outcome is determined by self-interest. Then we get polarization, and people will fight to the death for a prospect which is in either case ultimately not viable. But if we raise our eyes to the higher level of the total system in designing government controls, and use the viability criterion to determine the balance point, liberty must be a computable function of effectiveness for any total system whose objectives are known.

For example, when winning a war is the accepted objective—either for a nation or a guerrilla force—personal freedoms are acceptably sacrificed. But when a society fails to define its objectives, its consequent self-indulgence in freedom is met by a running tide of authoritarianism. And this is the explosive situation that so much of the world faces today, whatever its political colour, and at whatever level of recursion. Using the analysis I made a little earlier, the threat is that our world may not be viable much longer. Hence my plea for a cybernetic understanding of what is going on. I do not believe it has anything to do with genuine ethics: it is all about power.

Above all, the polarity between centralization and decentralization—one masquerading as oppression and the other as freedom—is a myth. Even if the homeostatic balance point turns out not to be always computable, it surely exists. The poles are two absurdities for any viable system, as our own bodies will tell us. And yet government and business continue the great debate, to the advantage only of those politicians and consultants who find the system in one state and promptly recommend a switch to the other.

These notions are central to the work I shall next describe. In Chile, I know that I am making the maximum effort towards the devolution of power. The government made their revolution about it; I find it good cybernetics. But the tools of science are not anywhere regarded as the people's tools; and people everywhere become alienated from that very science which is their own. Hence we are studying all these matters with the workers. Hence the systems I have to tell you about so far are designed for workers as well as ministers to use. Hence we are working on feedback systems to link the people to their government.

The enemy in all this is the image of exploitation that high science and the electronic computer by now represents. We are fighting that enemy and its ally technocracy. And so it must be only in Chile that you will find a famous folklore singer declaiming: 'Seize the benefits that science gives the people in their quest and 'Let us heap all science together, before we reach the end of our tether'.

I am proud to have worked with Angel Parra on that song, which is called Litany for a Computer and a Baby About to be Born. Contrast that title with the headline given to the first public mention of this work, which was leaked in a British newspaper last month, and has since been copied all over the world. It said: 'Chile run by Computer'. Woe to the sub-editor who wrote that.

Real Time Control

All that I have so far said is a very necessary preliminary to a right understanding of the economic control system I shall describe, which in any other terms would be a nightmare. But as society becomes differently understood—cybernetically restructured, politically redefined, differently lived by our children—yesterday's nightmares may become tomorrow's dreams. That is true for the whole of technological development. Without the restructuring and the redefinition the nightmare remains, as we who live in the polluted wake of the industrial revolution ought very well to know.

The thinking begins with one very clear idea. If things are changing very fast, then government needs instantaneous information. If its information is out of date, then its decisions are worse than irrelevant. Please consider this point very closely.

In 1956, Mr. Harold MacMillan (who was at the time Chancellor of the Exchequer) complained that controlling the economy was like trying to catch a train using last year's Bradshaw (time-table). It was true: the vital statistics of the nation were twelve months out of date. Sixteen years later, Mr. Harold Wilson (at the time Immediate Past Premier and the newly elected President of the Royal Statistical Society) has recently explained that things are better, and maybe many key national statistics are now only six or eight months out of date. And of course lags of either magnitude are commonplace in government throughout the world. It will not do. This is not only because decisions taken cannot have the benefit of the latest information; there is a far more ominous reason given in cybernetics.

It is a familiar notion that economic movements operate in cycles. Then out-of-date information is not merely 'late'; it is precisely incorrects—because it represents some cyclical trend that has since been superseded, but this is not recognized. If economic cycles were regular in periodicity and amplitude there would be no problem: the delay could easily be corrected. The decision-taker would discount the time-lag, and extrapolate. Indeed he tries to do this. Please look at Figure 1. By the time we discover either of the crises depicted, those crises are actually over. But we take action without knowing that, and therefore decide on exactly the wrong action each time. Now doing this actually causes instability.

To put the point in proper scientific terms: an unstable oscillation will occur at precisely the frequency for which the time-lags cause a phase shift of 180°. The negative feedback signal reinforces—instead of corrects—the original error.

It happens that the time it takes to implement a new government economic policy is of similar order to the statistical delay in acquiring facts, and so it is very possible to have the control system completely out of phase.

Lest this explanation should sound absurdly naive, let me add two reasons why the difficulty is not as perfectly obvious as I have made it appear. In the first place, neither of the lines I have drawn in Figure 1 is clear: both are fuzzy. That is, there is a tremendous amount of 'noise' present in the system—much of it deliberately injected by economic participants who stand to gain by causing this confusion. The second point is more difficult. The controller of an economic system is not a straightforward servomechanism with a known transfer function. It is itself a complex system, with its own time-lags, which are separate from the time-lags in the economy. It too may begin to oscillate; and in my experience, it does. Then there is a distinct likelihood that there will be a resonance effect between the two loops. If so, the oscillation in the controller would actually force a new oscillation onto the already oscillating system.




Figure 1. Underlying problem in controlling an economic variable, obscured by other mechanisms (see text).

No wonder, then, that no-one can disentangle all these effects; and no wonder that we do not perceive anything as simple as Figure 1 proposes. But in the absence of a complete explanation, there is something that we can do. Instead of solving the problem, we can dissolve it. Let us get rid of all the time-lags. Indeed, we ought to break with the very idea of arbitrarily quantized managerial time. Just as lags in reporting the past produce a bogus periodicity, so quite clearly do the lags fed forward in planning the future. A year's forward projection, or five-year plan, predetermine the cycle of expenditure and investment, and betray the capability of a viable system to adapt to environmental change. We cannot afford to await 'the next quinquennial review' when someone is standing on our foot.

What is the alternative to these inherited systems of lagged, quantized reporting on what has happened and lagged, quantized response to projected change? The answer from the mid sixties onward has been and remains real-time control. We have the technology to do it. This concept was fundamental to the plan we drew up for Chile in late 1971. We would abandon the hare-and-tortoise race to make relevant statistics overtake the lag in data capture and analysis, and implant a real-time nervous system in the economy instead. We would forget about the bureaucratic planning systems that talk in terms of months and years, norms and targets, and implant a continuously adaptive decision-taking system in which human foresight would be permanently stretched as far in any context as this real-time input of information could take it. Above all, we would use our cybernetic understanding of filtration to deploy computers properly as quasi-intelligent machines, instead of using them as giant-data-banks of dead information. That use of computers, taken on its own as it usually is, in my opinion represents the biggest waste of a magnificent invention that mankind has ever perpetrated. It is like seeking out the greatest human intellects of the day, asking them to memorize the telephone book, and then telling them to man 'Directory Enquiries' at the telephone exchange.

Having advocated all these policies for many years in Britain and elsewhere before going to Santiago, I was alert to the potential objections. I knew very well what is the standard response of economists, of managers, of civil servants, of ministers, and of 'established' science to these ideas. Let me list seven of

them, and give you the answers in brief, since some (though I trust not all) of these worries may be in your minds already.

- First Objection: The boss will be overwhelmed with data.

Answer: Not so. This is what happens now, as any manager who has had a foot-high file of computer read-out slapped in front of him can attest. The idea is to create a capability in the computer to recognize what is important, and to present only that very little information—as you shall see.

- Second Objection: The management machine will over-react to such speedy signals, which may not be representative.

Answer: Not so. This also happens now, as shown embryonically in Figure 1. The objection disregards cybernetic knowledge of filtration, and damping servomechanics.

- Third Objection: Such a system would be too vulnerable to corrupt inputs.

Answer: Not so, again. Present inputs are corrupt and go undetected, because they are aggregated and because the time has passed when they could be spotted. Clever computer programs can make all sorts of checks on a real-time input to see if it is plausible.

- Fourth Objection: 'Intelligent' computer programmes to do all this are still in the science-fiction stage.

Answer: This is woolly thinking. People do not really think out what is involved, because they conceive to the computer as a fast adding machine processing a databank—instead of seeing in the computer, quite correctly, the logical engine that Leibniz first conceived. The computer can do anything that we can precisely specify; and that includes testing hypotheses by calculating probabilities—as again you shall see.

- Fifth Objection: Even so, such programmes would take hundreds of man-years to write and be debugged.

Answer: I am sorry, but they did not. That is because the people involved in both London and Santiago were first-rate programmers who understood what they were doing. Let me be brutal about this: how many managers are aware of the research done into the relative effectiveness of programmers? They should be. The best are anything from ten to twenty times as good as the worst; and when it comes to cybernetic programming, only the very best can even understand what is going on.

- Sixth Objection: A real-time system with on-line inputs? It is Big Brother; it is 1984 already.

Answer: Stop panicking, and work out the notion of autonomy. I have still more to say about this later. All technology can be, and usually is, abused. When people turn their backs on the problem, crying touch-me-not, the abuse is the worse.

- Seventh Objection: Only the United States has the money and the knowledge to do this kind of thing; let them get on with it.

Answer: 'I find that slightly boring'.

Note: This objection was voiced to me in one of the highest level scientific committees in this land. The answer came from the Chairman, and I was glad not to be in his withering line of fire at the time. But he did not prevail, and neither did I.

In Chile, it took just four months to link up the key industrial centres to computers in the capital city—using a mixture of Telex lines and microwave connections (Figure 2). Purists may well point out that this does not constitute a real-time teleprocessing network, and they will be right. However, we have used the real-time philosophy, and have simulated an on-line system. The programs are written for that; and if someone will kindly donate the teleprocessing equipment, it will soon be in action. (I have mentioned the problem of foreign exchange already.) Meanwhile, we have to use too many human interfaces. But I am not going to apologize much about that. The fact is that we can cope with daily input, and that is—relatively—very close to real-time: in normal government terms, you cannot tell the difference.




Figure 2. Nature of Telecommunications using existing public system. This cannot yet accept even low speed data. System uses Alphabet No. 2 at 50 Bauds.

This communications network was in itself a fairly simple technological manoeuvre; but even so it constitutes a big advance for government cybernetics. During the October crisis of 1972, some of the most senior people in Chilean government came fully to understand in practice what Wiener had expounded theoretically long ago: communication is indeed control.

Well: to know today what was the state of the industrial economy yesterday is a considerable advance on knowing what it was six months or a year ago. But we were trying to do more than merely get up to date. Frankly, there is not much point in knowing what happened even yesterday—because even yesterday is the purest history. Nothing can be done about it any longer. But if we can get hold of a close idea of what is going to happen next week, then we have at least a chance of doing something about that. And certainly knowing what has been happening over the last few days is the best basis for estimating what is likely to happen over the next few days. . .

The question is: how? One may call for data, but he has to meet the problems I listed just now (—the 'fatal' British Objections—) if he is to make effective use of them. One may know all about yesterday; but he has to be fairly ingenious to say the right things about next week. The initial four-month plan of action, which had included setting up the communications network, tackled these problems too; and it successfully defeated them.

Systems Design And Variety Engineering

Interdisciplinary operational research teams set out to make (crude, but effective) models of all the major enterprises in the social economy. These were not to be the vast, static, historical, and essentially out of date and non-stochastic input-output matrices beloved of so many state planners. We wanted to get at the dynamic systems which made the enterprises tick; and we wanted them in a form that managers and ministers could immediately grasp. Therefore we used a visible and visualizable type of model, called a 'quantified flow-chart'. Start with production (a Marxist government has no illusions about the source of the creation of wealth). If we list the production operations of the firm, and their productive capability, we can make a map of production flow—in which the flow lines are proportional to the relative amounts of flow, using some convenient measure, and the operations themselves are boxes at the confluences—also shown in relative sizes according to their productive capability. Here is an example (Figure 3).




Figure 3. *Quantified Flowchart*: thickness of lines proportional to rate of flow; size of process boxes proportional to productive capability; productivity indicated as a level in boxes; bottleneck process highlighted at A; Customers listed C₁ to C₁₁.

Now of course if that kind of presentation can be made for the flow of production it can be made for any other kind of dynamic system in which management may be interested: cash flow, for example, or the deployment and movement of people and of goods. And although we started out on this task under the aegis of Operational Research, I am hopeful that as people become accustomed to the idea we can use a better approach. Do we really need objective, scientific enquiry to understand what the structure of the system is, and how it should best be quantified? Actually not. The people who best understand what these systems are really like are the people who operate them. You do not need a string of degrees to understand how to make a quantified flow chart of the activity that surrounds your daily life. So here I hope will be the start of 'participation' in the future, and OR expertise will be used merely in teaching and in guidance.

With this simple device we start on the road leading to the answers to those objections about overload. In cybernetics we have an actual measure of complexity, which we call variety. By devising systems in which homeostats are set up between management and whatever is managed, we embark on the process that I have labelled 'variety engineering'.

The quantified flow chart is in itself a variety-attenuating filter. In the first place, it can select its own degree of optical resolution. For example, it can show a box called simply steel production: or it can show three boxes identifying kinds of steel production—by open-hearth, electric arc, and converter, say; or it can show every individual furnace. By the same token, it can lump together all the materials that go into a steel-making furnace charge, or it can distinguish between them. This variety engineering concerns the account of the operation that has meaning for a particular management group, and the degree of optical resolution chosen depends on the level of recursion at which this operation is being considered. In the second place, iconic representation is also a variety attenuator in the suppression of words and numerical data: it is a product of gestalt psychology, in which pattern is relied upon to convey information.

The next variety attenuator involved in this representation is the concept of capability. The real-time variation in actual flows and outputs is killed in the iconic quantified flow chart, and referred instead to a relatively static idea of 'what can be done'. You might think that this would be difficult to define, but in practice it is fairly easy. Capability is a systems concept: what outputs is the total system capable of generating in each part, given the limitations imposed on any part of the system by other parts? Then 'capability' is not to be confused with 'capacity', which is not a systems concept—because it alleges that some part of the system can in theory do something that may be rendered impossible by other parts. This variety attenuator is valuable because it reflects reality for the whole system concerned, and that has meaning for the recipient of the iconic representation.

However, we could—given a breakthrough of some kind—do better than the results of which we are currently capable. After all: if capacity exceeds capability in some parts of the system, there must be other parts of the system (called bottlenecks) that are actively restricting capability. These bottlenecks

may have to do with low local capacities, or they may have to do with technological constraints. For example: a mill's engine might be perfectly adequate to drive its rolls at twice their current speed—if only we had a better lubricant. Then these considerations define potentiality, which is something better than capability. Potentiality is the performance of which the system would be capable, 'if only. . .' That does not mean that we look for pie in the sky: it means that we look for investment—in new equipment, to cure the bottlenecks, or in research to cure technological shortcomings. It is not very difficult, keeping one's feet firmly on the ground, to define a system's potentiality.

But if potentiality is better than capability, there is something worse—and that is actuality. The performance of systems cannot rise to their potentiality without investment of some kind; it cannot even rise to their capability unless activity is perfectly well organized. It never is. In consequence, what actually happens falls short of the capability expressed before. Moreover, actuality expresses that very reality of which I spoke earlier—the day-to-day viscissitudes of life. It was this continuous variation which drove our thinking down the road to real-time control. Somehow we have ended up with three versions of systemic truth: actuality, continuously fluctuating; capability, a much steadier variable; potentiality, which is absolute until the system itself is structurally changed. And it is capability which the iconic representations represent. To make them show potentiality would, for the moment, be unrealistic; to make them show actuality would, at all times, result in their dancing in perpetual fandango before our eyes. So this capability attenuator is a powerful but sensible reducer of operational variety.

So be it, in so far as iconic flow charts are concerned. But what about continuous reporting, and the problems of real-time control? Whatever information we collect, it is due to be hurled round dozens of homeostatic loops—those loops that make up the total systems design. That information has very high variety, and the analysis we have just made multiplies it by a factor of three—or so it seems, if we want a measure not only of actuality, but of capability and potentiality as well. . . But rescue is in sight. Both capability and potentiality are relatively static measures. If we take their ratio, the resulting index will also be relatively static. Moreover, such a ratio will be a massive variety attenuator—because it will be a pure number, varying between nought and one. So instead of trying to consider, all-in-one-breath, that we have a capability of 800,000 tons and a potentiality of 1,000,000 tons, we shall think of a ratio of 0.8; while the capability to use 110 men contrasted with a potentiality to use only 22 yields a ratio of 0.2; and the capability cost of an item of product at 120 escudos compared with a potential cost of 60 escudos indicates a ratio of 0.5. Well what is potential in current capability is a latent resource; and it could be freed by investment in some form. So I call the ratio between capability and potentiality the Latency Index. Looking at a new iconic diagram (Figure 4), we can see how potent a variety attenuator has been devised.




Figure 4. Iconic representation of relative Latencies (see text).

There is no need any longer to try and assimilate the numbers that characterize the units measured. That is the strength of an index—it is a pure number, varying over a fixed range. Hundreds of thousands of tons; hundreds and tens of men; units of money; there is no need to wrestle with them. Nor, if we stick to our ideas about iconic diagrams is there any real need to use digits at all. We can distinguish very clearly, using our eyes, between the levels represented in the iconic diagram. It might satisfy an accountant, but it would make no difference to a manager, to declare that a Latency Index had changed from 0.71 to 0.73. Who cares? The computers behind the manager's eyes will undertake whatever process of discrimination has meaning for his judgmental brain. Then this was the first though massive

piece of variety engineering we set out to achieve in Chile, on those initial (crude, but effective) models, contrived at an appropriate level of optical resolution, of all the firms. As I said, the Latency Index is all about investment, and we shall certainly return to it later.

Meanwhile we must consider actuality, the real-time variable in the entire system. For if a Latency difference between 0.71 and 0.73 means nothing, because both potentiality and capability are fairly static, such a difference in a fast-moving index could mean something very important. It might be part of a trend. I have already explained the arrangements by which the data representing actuality come into Santiago every day. They are used to form a second ratio, comparing actuality (the newly arrived figure) with capability (selected from the computer store). This is the Productivity Index. It is in a continual state of oscillation, which destroys the variety that is of no concern. In the next diagram (Figure 5), we can see how the three concepts of actuality, capability and potentiality are combined as two ratios to form the Latency and Productivity Indices, and how these in turn create an overall Performance Index. The reason for this iconic representation, in place of the familiar mathematical notation, lies in the fact that which part of the ratio is the numerator and which the denominator depends on what is being measured. For instance, capability is always better than actuality, but in numerical terms it may be more (e.g. output) or less (e.g. manhours per unit). Naturally enough, the smaller number in the ratio is the numerator, since the index will be less than 1.0.




Figure 5. Three measures of capacity defining three indices of achievement and three concepts of planning.

The indices procure an enormous variety reduction; even so, we still have problems in conforming to Ashby's Law of Requisite Variety when it comes to managing the economy. The company production models for instance generate on average about ten triple-indices per plant; these always include raw material and finished stocks, the output of key production processes, and labour absenteeism. This degree of resolution is minimal, and managements have free rein to install whatever extra indicators they like. This honours the argument for autonomy, and it makes an insignificant difference to the work load of the computers, because all the numbers inside the computational system are diurnal time series of indices varying between nought and one. The programmes are therefore infinitely extensible in application. Even so, with the system in full operation, many thousands of actuality inputs will arrive daily, generating three times as many indices; and the total number could easily rise by two orders of magnitude as the autonomy criterion is understood by managements, the operational research goes deeper, and worker participation becomes real. And so we reach the more subtle notions of variety engineering.

If a particular indicator, say the rate of crushing limestone in a cement factory in Northern Chile, is generating a new Productivity Index every day, what ought to be done with it? Should be lay the new figure, each day, on the desk of the Minister of Economics? Surely not. This variety must also be filtered. There are two statistical notions involved, and the first is very simple. A population of (say) a hundred such figures generates a probability distribution. This may turn out to be oddly-shaped, rather than straightforwardly Gaussian; and especially it may be skewed to the right (since the index has a finite limit of one). It is a simple matter, however, to correct for this statistical aberration, by using a trigonometrical transformation. Then we may establish the mean and variance of this population of indices. These two statistics, all alone, characterize the stochastic behaviour of each index over time. Then if we take a running sample of the indexical figures as they are computed, it is easy to establish whether a significant change in the mean or variance of the statistical population has occurred. The statistical population characterizing each indicator is known as the taxonomic index, because it

classifies every measured activity within every operation according to its mean productivity. There is a standard computer programme that looks for changes in the taxonomic index; if such a change is found, that is notified to the management concerned, and the iconic graph is changed. Further, the history of the index over time is updated (Figure 6). These are relatively rare events, but the procedure mentioned absorbs the variety engendered perfectly well.




Figure 6. Iconic record of significant (therefore taxonomic) change.

Breaking the Time Barrier: Cyberstride

The more difficult problem, and the more sophisticated statistical notion, concerns the possible trend that each new daily figure may betray. If the economy is to be under realtime control, the government cannot wait to know that a significant change has been registered for a particular taxonomic index—although this is already much to be preferred to the orthodox system of routinely quantizing statistics, where the recognition of significant change is left first to the alertness and next to the judgment of whoever is supposed to be watching the results. No, it is much more than this: we approach the problem of breaking the time barrier. Can we tell from yesterday's figure, and the short-term run in which it participates, what will happen (unless we intervene) tomorrow and next week? It is the problem of short-term forecasting, with which a great deal of progress has been made in recent years.

Allow me once more to return to the facts of the Chilean work. Before the end of 1971 I had designed a specification for the computer programme to deal with taxonomic indices having daily actuality inputs, and it was in the hands of a team of operational research consultants in London, who had been commissioned to write the programs. We were discussing the short-term forecasting problem, when the London team discovered a brand new paper in the *Operational Research Quarterly*—hot off the press. The authors were Harrison and Stevens and they had clearly made a major advance in the field of short-term forecasting⁴. We had been talking in terms of Cusum (cumulative sum) techniques to this point, as representing the best available practice. Cusum itself was associated with the first author, who had been pressing its virtues for many years, so we were naturally impressed that this novel development came from him. The obvious power of the method (always supposing it worked), and the elegance of the mathematical demonstration behind the approach, convinced us to take the plunge. It was a noteworthy decision. The London team wrote a temporary suite of programs which included the Harrison-Stevens approach and incredibly had it working in Santiago by the March 1972 deadline of the first phase of the operation mentioned already. Meanwhile they began work on the permanent version, creating a specification that was handed over to the Chilean scientists. In the meantime, as the system was growing, experience was gained in the actual use of these very complicated program suites, and they grew in sophistication all the time. But these new developments, vitally important though they are, must await presentation by the men who made them possible in more technical papers than this.

This suite of computer programs, called CYBERSTRIDE, is the essential feature of the filtration system that achieves the variety attenuation demanded, and which breaks the time barrier of which I was speaking. It takes as input the actuality figures every day; it makes various checks on their integrity;

it computes the triple-indices; it makes statistical judgments about the taxonomic indices as I have already described. After that, using the Harrison-Stevens techniques, it really gets clever.




Figure 7. Using the computer intelligently to calculate probabilities of future alternatives.

When a new value for any index is computed, Cyberstride looks at it in the context of the recent history of that index (Figure 7). The new point might stand for any one of the four outcomes depicted. It stands for no change, or for a transient (neither of which matters to the manager); or it stands for a change of slope, or for a step function, (both of which possibilities matter very much). Using Bayesian statistical theory, the program calculates the posterior probability of each of these four events—for every index, every day. The programme is incredibly sensitive to these changes, recognizing them long before the human brain would dare to make a judgment. Cybernetically speaking, the system (as Harrison and Stevens claimed) is self-adaptive: its sensitivity increases whenever uncertainty increases—which happens whenever an apparently unusual index value is thrown up. Moreover, instead of producing merely single-figure forecasts (and who can foretell the future with that kind of precision?) it produces a joint parameter distribution that expresses the inherent uncertainty of all forecasting.

So this is what I meant in speaking of computers as quasi-intelligent machines. Cyberstride throws away the huge component of variety that has no meaning, because it represents a chance fluctuation. It is at once alert to significant changes, focussing on them an analytical eye, and capable of estimating on the strength of that analysis what will happen next. The only problem we had with Cyberstride, and it was very severe, was its calibration in terms of these posterior probabilities: how sensitive should it be made? Obviously, it could discard too much, or become overexcited about too little. The 'tuning'

subroutine that fixes these limits of excitation, so analogous to the so-called physiological limits of variation in any homeostat, was the big achievement of the Chilean scientists working on Cyberstride.

The variety engineering is complete—for the lowest level of recursion, the enterprise itself. If it would have been ludicrous to confront the Minister of Economics with the whole variety of fluctuating indices, it would still be absurd to inform him of even highly significant movements in the limestone-crushing activity of a cement plant in Northern Chile. Absurd, yes; but also ominous. I am sure you recall the argument about autonomy and overcentralization. What happens in Chile is this. The results of applying Cyberstride daily to the new inputs which quantify the iconic flowcharts are fed straight back to the managements concerned. It is their responsibility to do something about the warnings that are generated in this way by quasi-intelligent machines. No other human being than the responsible manager receives any information about this extremely elaborate piece of computation, and I attach very weighty importance to this fact.

Then, you will ask, what about the other levels of recursion? The manager of the enterprise is very well served by all of this, especially so, since he can pump any indexical series he cares to contemplate into the routine—and receive the alerting advice, whenever it is available; meanwhile he may feel perfectly confident that an absence of alerting advice means that whatever operations or activities are being monitored for him by Cyberstride are fluctuating within the physiological range of chance variation. But what about the Sector Committee, the Industrial Branch, the Minister of Economics himself? These are higher levels of recursion: how are they to be informed?

Here is the coup de grace of the cybernetician, in his role as variety engineer. All viable systems are contained within viable systems. It is the principle of recursion; the model is the same. So it is easy to see what next to do. The iconic representations, called quantified flowcharts, are to be aggregated at sector level, aggregated again at the industrial branch level, and aggregated finally at the level of total Industry. The quantifiers (those actualities, capabilities, and potentialities) are to be aggregated too—not, as is orthodox practice, in terms of averages, but in terms of new operational research models (crude, but effective) of the level of recursion concerned. In that case, raw data—heavily processed through atomic indices and through Cyberstride, which produces exceptions known only to the manager concerned—bypass that atomic level of recursion, and become raw material for a molecular level of aggregation higher up. Here they lose their identity; they merge (not by averaging but by modelling) into new molecular indices.

But these new indices, although they have lost a great deal of variety in the process of molecular aggregation, have acquired variety by the sheer amalgamation of so many enterprises. How shall we deal requisitely with this new variety? Well, it is represented by triple-indices, all operating between nought and one. So although the level of recursion changes, and although the atomic index changes to a molecular index, the Cyberstride suite of programs is invariant. The whole process I have described starts again. This time and again, exceptional information is fed back to its proper level of recursion: the sector, or the branch, or the minister.

Return with me now, for the last time, to the vexed issue of autonomy. I regard the whole of this work as a fanfare for freedom—but for effective freedom. The claim was made that the degree of autonomy, and its complement the degree of centralization, are computable functions of viability. I stick to that. By separating the levels of recursion, and within those levels by preserving freedom for each separately designed interlocking homeostat, the maximum autonomy consistent with effective organization is assured. A problem remains. What happens when, for whatever reason, the appropriate homeostat at the appropriate level of recursion FAILS TO ACT?

Many a freedom must have been lost from the fear of those in power that subservient systems down the line would not do their jobs. And, if not, it makes a good excuse for the tyrant. This is a classic and intransigent problem, but we can now deal with it easily—if we keep our cybernetic heads. An autonomous unit is supposed to react to any adverse exception reports that it receives from Cyberstride. How long will that take, and how much does it matter? The answer to both questions will vary widely. In our work we have included in the operational research modelling a requirement to assess the possible rate of reaction to change, and the relative importance to the system modelled of such a change, for every indicator. When the computer sends an exception report to a manager, at whatever level of recursion, it computes for the message an acceptable delay time which is a function of both the possible reaction time and the importance, and it starts a clock. If our quasiintelligent machine fails to detect an improvement within this allotted time, it breaks with the autonomy and notifies the next level of recursion (as well as telling the responsible manager that it has done so).

These special signals are different in kind from the routine management signals. We call then 'alge-

donic'. The word means pain-and-pleasure; and it was work in neurocybernetics that taught me this answer. We rely on our bodily organs to do their jobs; but if they should fail, we get a special signal—transmitted by specially adapted neural pathways that bring the facts to our conscious attention. The mechanism is precautionary. Clearly it involves a threat to autonomy, but the body politic cannot sustain the risk of autonomic inaction any more than we can as human beings. And remember that there is nothing covert about this. The delay factors are discussed with the managers concerned, and they are informed if the algedonic signal is transmitted. Indeed, they may be very relieved—if the problem is seen as beyond their control—to know that the signal has automatically gone.




Figure 8. Cyberstride operating at four autonomous levels of recursion, with algedonic feedback (see text).

In this way, just as in the body, a sign of special distress automatically breaks through to whatever level is required to deal with it (Figure 8). For if the management group which receives the signal fails to act within its appropriate time delay, the signal will go up to the level next above. Thus the signal makes it possible for a problem concerning that limestone crusher in the cement factory to reach the President's Economic Council. Let us hope that never happens; but it would be surprising if signals of distress were never received there from the Sector level.

The Faculty of Foresight

The real-time control system I have so briefly described is founded on the following elements: a cybernetic model of any viable system; a cybernetic analysis of the real-life systems appropriate to each level of recursion, and their iconic representation; a design of a large number of interlocking homeostats; the provision of a national communications network capable of operating now on a daily basis and eventually on the basis of continuous input; variety engineering throughout the system to incorporate filtration on the human brain's scale; and the Cyberstride computer program suite capable of monitoring inputs, indexical calculations, taxonomic regulation, short-term forecasting by Bayesian probability theory, autonomic exception reporting, and algedonic feedback. It makes quite a package, and it exists. It represents a system of here-and-now management of the economy that is not based on historical records, but on an immediate awareness of the state of affairs and the projection of that awareness into the short-term future.

Let us call this whole thing the NOW system. Then clearly we also need the FUTURES system. What are we doing all this for? If government is not to be merely the management of perpetual crisis, it needs to look ahead. Party-political programmes are supposed to be all about the kind of society the people want, and the government is supposed to be dedicated to achieving that. In practice, perpetual crisis drives out mandated intentions. It even happens (dare I say?) that entire sets of electoral policies become reversed when power is granted. This can be only because government has no arrangements for realistic normative planning. It has a political theory, but it does not understand the system it is manipulating. It is just laughable to say, for example: 'the theory is all right, but the trade unions (or the City, or the banks, or the people themselves) will not operate the theory'. The unions, the City, the banks and the people are all elements of the total system that the government claims to be able to govern.

Thus I introduce what I have to say about long-range planning in terms of understanding systems and how they respond; and I do so in deliberate contrast to the many schools of thought that base their conception of inventing the future simply on forecasting it. My objection to that approach is twofold. In the first place I do not believe that we can forecast the future—and that is a fairly strong objection. The future, I reckon, is known only to God; and it seems to me that the class of men who have always come nearest to perceiving his intentions are the science fiction writers. They have usually been very close to scientific reality. The people who run society, who are famous for being 'realistic' and 'responsible', turn out to be outrageously irresponsible just because they are so unrealistic. Their unrealism consists in a refusal to notice what science is actually doing, and a refusal to think through the inevitable systemic consequences of the policies they underwrite.

These were the reasons why I was determined to provide the Chilean government with an instrument for investigating the systemic consequences of alternative courses of action. For there really are choices to be made. When you read that car prices in Chile have gone up by 900% in a single year, what is your response? Is this the inevitable result of Marxist dogma, is it just what you expect from nationalization, is it a measure of inflation, or what? To whom does it occur that it may be the result of a deliberate choice between economic desiderata? Thus are we all brainwashed by the consumer society, in which the motor car is an absolute god.

The second reason why I object to the forecasting approach to long-range planning is that it assumes that there is 'a future' out there, lying in wait for us. This is not true, surely, except in so far as larger systems beyond our own—and in which we acquiesce—take a stranglehold on us. I have already suggested that this may doom the Chilean experiment. The real freedom we have is to change our structures and our policies so that the future is different from the future we should have encountered had we not made those changes. And this is where understanding dynamic systems becomes the task. The fact is that we need not to forecast but to experiment.

Experimentation is not easily or perhaps justifiably done when we are talking of social institutions. Scientists undertake social experiments on animal populations, which they try to use as models of human populations—but the discrepancies may be very wide.

Probably the best experimental tool available is the computer simulation. According to this approach, one programmes a computer to represent the dynamic social situation, and then experiments on that. If one asks how such a model could possibly be validated—he learns that the model can be fed with historic data—on the basis of which it ought to simulate the appropriate historic outcome. That is at least a start in a demonstration of validity.

I introduce the topic of dynamic systems simulation in this way, calling it an experimental tool, because I consider there to be a great deal of misunderstanding on the subject. If we experiment on

a model, putting in possible policies and reading off possible outcomes, then of course we appear to be making predictions. Some people have been causing a great deal of public disquiet with some such predictions about the ecosystemic future of the planet. Personally, I do not mind their doing so—because I believe the public ought to be thoroughly disquieted on this score. But we must make methodological distinctions here. In so far as these models make predictions, it is vital that projections of the input variables be correctly made. There is the rub because specialists disagree quite fundamentally about the trends that have been built into some of those models. Clearly, if it is taken as input that fossil fuel will run out by a certain date, then predictions for the ecosystem incorporating this input will be falsified if that date turns out to be wrong. But suppose our objective is not to make predictions, but to make experiments to find out how the ecosystem works. That is a different matter. We should put in a whole range of possible dates for the exhaustion of fossil fuel, and find out what difference they made to total performance and by when. After that, we should have a good idea what policy to adopt towards research into novel sources of energy. And that policy would not be the fruit of predictions that might well be falsified; it would be the embodiment of our understanding as to where the system's vulnerability lies.

My belief is that government planning should be based on this same idea. If we make a dynamic model of the economy, concentrating our power of resolution on the areas in which our decisions appear most unsure or most frightening, then we shall learn how the system operates. The first task is to identify the crucial parameters, which (because complex systems are richly interactive and internally reverberating) are not always the parameters assumed to be critical. It is quite characteristic of cybernetic studies to obtain results that are counter-intuitive. Therein lies their value. The next task is to discover how those parameters may best be manipulated, which (because political dealing is a complicated business too) may be in roundabout ways rather than by direct intervention.

What matters about a dynamic system, if you want to understand how it behaves, is not so much noticing the sore points themselves, nor resolving the apparently insoluble politics of applying remedies to those sore points—all of which turn out to be unacceptable remedies for some segment of the population. What matters is to change the structure of the system, so that homeostatic equilibrium is restored and the sore points disappear. That involves variety engineering: it is likely to mean the redesign of institutions, the addition of informational feedbacks, and the calculated change of time-lags in various rates of flow. Economists, perhaps, would not recognize those three cybernetic prescriptions as counting towards the solution of what are regarded as economic problems. But are all our problems economic? I think there is a prior set of problems about the regulation of society (which it falls to governments to solve), which may well have economic causes and consequences, but which are themselves about effective organization.

Returning to the Chilean story, then, we wanted to create a facility for normative planning, suitable for all levels of recursion, embodying dynamic system simulation. Now the task of inventing a fresh computer compiler for this purpose was outside our time-scale. A number of compilers exists, and we chose to use the Dynamo compiler in its latest version⁵. The choice was made on the grounds of its elegance and its relatively long existence—meaning that it is very well debugged. The choice has been criticized, and will be again; because this is the very compiler used in the work that I referred to as making predictions about the planet using inputs that many ecologists regard as insecure. To me, that is like blaming the pornographic content of a book on the English language in which it is written. My defence of the compiler says nothing about my concurrence or otherwise with ecological predictions, any more than hearing my defence of English would tell you my views on pornography.

For the record, then, we again formed two teams—one in Santiago, and the other in London. These teams were organized differently from the two Cyberstride teams, and had no members in common; and this simulations pair operated in a different way. Instead of members of the London team taking program suites out to Santiago to be developed, as happened with Cyberstride, a member of the Chilean team came to London to learn a cybernetic skill. Moreover, whereas all the Cyberstride runs using actual data were undertaken on the Santiago computers, the simulation runs for a long time were undertaken on computers in London. In this way dynamic systems models for the Chilean futures system were originally developed, but by now the whole of the work is being done in Santiago.

There is much that is new about these models, but for obvious reasons I shall not discuss their content. What is worth remarking upon is the status of the information fed into them. As I said just now, models of this type are often criticized on the grounds that their inputs are suspect. Now this is not surprising; because, as I also said before, economic information at the national level is usually about a year out of date. But Cyberstride produces information that is immediate. Then there is a question about the interface between the real-time control system and the futures system. If absolutely current

information can be used continuously to update our models of the world, a new era dawns in national planning. Well... at any rate it happens just like that in the brain. We should indeed be foolish to choose between the alternatives open to us as men on the strength of knowing what our circumstances were like last year.

A Decision Machine: The Opsroom

And now we reach the final question. How do we 'get it all together', How is so much sophisticated science to be made available to those who bear the brunt? In most countries, this is a function for the civil service. Those people constitute the ultimate filter. The ministerial briefing stands, however responsibly, between the minister and all those urgent facts of the NOW system, all those experiments in foresight of the FUTURES system.

I wanted ministers to have a direct experience, an immediate experience, an experimental experience. And what goes for ministers goes also, at another level of recursion, for managers—whether the managers of the social economy, or (at yet other levels of recursion) of enterprises or of plants. Above all, if 'participation' has any meaning, no-one must be disbarred because of an inadequate grasp of jargon, of figurework, of high-level rituals. As I told you before, the workers themselves must have access to the whole of this. Let me put the point before you in two contrasting ways.

When I first expounded the cybernetic model of any viable system (which I have not expounded today) to President Allende, I did so on a piece of paper lying between us on the table. I drew for him the entire apparatus of interlocking homeostats, in terms of the neurophysiological version of the model—since he is by profession a medical man. It consists of a five-tier hierarchy of systems. I worked up through the first, second, third and fourth levels. When I got to the fifth. I drew an historic breath—all ready to say: 'And this, compañero presidente, is you'. He forestalled me. 'Ah', he said, with a broad smile, as I drew the topmost box: 'at last'—the people'.

I do not care what political outlook any of us may have: that story ought to convey a profound message. It deeply affected me, and it affects this work. The second perception of the same point that I give you comes from that Litany written by the folklore singer Angel Parra, which I quoted at the outset. This is what his song says on the subject (my translation)

Equal I say to the Minister
Selling promises forlorn
Since all of us are hostage
For that baby to be born.

Society can no more afford the alienation of the people from the processes of government than it can afford their alienation from science.

And is this really a political question any more, once we say that all of us are men? The fact is that no man, worker or minister, has more neurological jellyware than anyone else—although he may make marginally better use of his endowment. We have seen how that man, minister or worker, can be saved from drowning in a inundation of statistics and reports—through variety engineering, and the deployment of computers as quasi-intelligent machines. But how does the filtered information get into his head?

The answer to this lies in the operations room. If the connotation of that phrase reminds some people of a wartime headquarters, the allusion is quite deliberate. For in the opsroom real-time information is laid out, quite graphically, for immediate decision; and in the opsroom a synoptic view of the whole battle is made plain, so that the total system can be encompassed by human powers of foresight. We used every scrap of relevant scientific knowledge in designing the place—neurocybernetic knowledge of brain processes, knowledge from applied and group psychology, knowledge from ergonomics.

The opsroom looks like a film set for a futuristic film. But it is not science fiction; it is science fact. It exists, and it works; it exists and it works for the worker as well as the minister. [Photographs of this opsroom are reproduced on the inside of the dust jacket of this book.] There are seven chairs in there, because seven is the maximum creative group. There are various screens in there, all using iconic representations of information, because those are the sort the human brain can best handle.

The central screen—central in that all the others are referred to it—is a picture of the viable system (Photograph 3). It is eight feet high and four feet wide, and it is set up according to the recursion theorem, for whatever group happens to be using the room. The operations involved are marked in the

circles; the current molecular indexical levels for the taxonomic triple-index in each element are shown by iconic descriptions in the square boxes. The total square in each case stands for potentiality; then the green level is actuality, and the red level is capability. You can see how easy it is, if you remember the explanation about their ratios, to get an immediate grasp of the relative levels of Productivity and Latency in each case.

There are a great many interlocking homeostats operating here, which can be discussed as people come to understand the cybernetic laws that govern their behaviour. This is not readily understood from a still photograph, and in fact this screen is animated. There are no arrows to be seen, therefore: just moving lines. Scientists often suppose that to mark a line with an arrow makes it clear that the total system so encumbered with arrows is actually dynamic. Not so; people read the arrows as indicating directional, but still static relationships. Besides, the most critical loops here operate at differential speeds (they can be changed)—which tells the brain a great deal about the relative lags in the system.

In the top third of the diagram, three boxes may be noted. The lower of these controls the NOW system, and the central one controls the FUTURES system. The top box (the boss or 'the people?') monitors their interaction, to which attention is drawn in the animation by a constant movement in the big yellow circle.

The Cyberstride exception reports flow on the horizontal red lines; and when they exist an Exception Screen is lit, giving details. An algedonic signal is indicated by flashing red arrows on the vertical axis (in the photograph, such a signal may be seen emanating from the middle element); and any algedonic activity lights the Algedonic Screen.

Shown in Photograph 4 are the two screens I have just mentioned. On the left is the Exception Screen, showing two alerting signals from Cyberstride, together with a first indication of the kind of warning coming through. On the right is the Algedonic Screen, showing signals from two different levels of recursion below—each marked by a red light flashing at a different speed. As I said, attention veers to these two real-time inputs because of the clues given on the main screen first described. Obviously both these screens are currently set up by hand, whereas they could be set by direct electronic output from the computer. But I would like to repeat that this is simply an annoyance due to component shortages; it does not represent a gap in the total cybernetic system.

This, then, is the real-time input to the opsroom—its sensing devices spreading out over three thousand miles of country, and its quasi-intelligent filtration continuously reducing an immense informational variety to human proportions. Then what will our seven-man team of creative thinkers want to do next? For make no mistake; the opsroom is a decision machine, in which men and equipment act in symbiotic relationship to amplify their respective powers in one new synergy of enhanced intelligence. They have to start talking and deciding on their actions. For this purpose, they will need background information; and I need hardly explain that there are no files, libraries of reports, or minutes of the last meeting here. Paper is banned from this place. The answer is Datafeed. (Photograph 5).

It consists of three data screens, as you can see, and a huge index screen. Each of the data screens is supported by five carousel projectors, each carrying eighty slides of iconic information. So we can choose three out of twelve-hundred presentations—one out of four-hundred on each screen. But it is obvious that twelve-hundred slides cannot be listed on the index screen above, however huge. How shall we get at our treasure-trove of supportive information? It is again a problem in variety engineering: select three from twelve-hundred. Of course, one could have a catalogue, and a decimal keyboard. That would have requisite variety. However, experience teaches that unskilled people will not usually agree to operate such devices. They see them as calling for a typing skill, and want to insinuate a girl between themselves and the machinery. Indeed, this fact has held up the development of on-line conversational computing very seriously indeed. We are faced with an ergonomic problem. It is vital that the occupants interact directly with the machine, and with each other.

In a creative conversation, men become very animated. They seize pieces of paper, and draw on them; they snatch the pencil, and change the drawing: 'no, no, it is like that'. The solution to the ergonomic problem takes note of all these things. We produced special chairs which swivel through 270 degrees of arc, in the arms of which are mounted panels containing large knobs of different shapes (clearly visible in Photograph 1). By thumping one of the three knobs in the top row, a screen is selected—and an index automatically appears on the control screen. The index is catalogued by the use of five symbols, which are repeated in the second row of knobs. By pressing the appropriate combination of knobs, the item selected from the first index appears on the control screen, in the form of a second index which lists the actual slides. So a second combination procures the required presentation.

The variety engineering says: there are $2^5 = 32$ ways of combining five knobs; and if that is done

twice, $2^{10} = 32 \times 32 = 1024$ alternatives are made available. That is enough selection power to handle 400 slides on each screen, with plenty to spare for control engineering purposes. (Four buttons would yield only $2^8 = 256$ alternatives.) One of the two knobs in the bottom row allows one man out of the seven to seize control of Datafeed with a thump, and the other releases the control when he says 'thump—that's what I mean'. There is no finicky skill involved in working this apparatus—and people seem to take to it very quickly indeed. As to all that thumping: I wanted to make the dramatic act of using the equipment an effective part of the creative conversation, just like seizing the pencil, or banging the table.

Thus it is that when real-time inputs indicate the need for supportive information, the decision-takers may select on the three screens (for example, as in Photograph 5) the iconic flowchart that contains the relevant input, a photograph of the plant concerned, and some indexical information. If an expansion or explanation of that information is available, for example the history of a Latency Index may well be supported by an investment plan for realizing that latency, then a direct clue is given on the screen as to how to key that new slide into place (such a clue is visible on Screen Q. In the close-up of Datafeed (Photograph 6), the picture of the plant has been replaced with a list of products.

All this supportive information is semi-permanent. It must in principle be updated, but not too often. As to its adequacy, remember that all sixteen carousel magazines can easily be changed. so we have a new set (of 1200 slide capacity) for each level of recursion. It is enough. In any case, there are two more back-projection screens in the opsroom to allow special presentations to be made. So far I have spoken about the NOW system, but certainly Datafeed supports the FUTURES system too. The relationship between the two is very clear on the huge main screen, where its corruscating homeostat is a constant reminder of the need to balance investment between what is and what will be. And that FUTURES system, with its simulation capability, has its own screen (Photograph 7). This is the flowchart of a typical Dynamo simulation—though not a very complicated one. The two points I want to make about it are unfortunately not communicated in this static picture.

The whole raison d'être for simulations is to work with them. They do not sit there 'making forecasts', as I said. The output of the model shown is a projection made by a computer of how the major variables will vary over the next ten years—if nothing changes—and that projection is illuminated on one of the spare screens. To understand how the economic system works, the people in the room need two facilities, neither of which is available on an ordinary flowchart. First, they must be able to alter the structure in front of them. That is easily done in the computer: attendant scientists can change a few equations on request, and produce a new read-out in a few minutes. But how do you alter the flowchart? The answer to that is to use flexible magnets, and we did. However, to decide how to alter the flowchart you must understand the flows—and therefore we wanted to animate this screen. The problem was how to animate a flowchart that you wish continually to reconstruct. The British suppliers of the animated equipment solved that problem; and I wish I could show you the flow-lines on this model moving, and how readily its structure can be changed.

Indeed, we could spend all day in the opsroom together without exhausting its meaning as a new tool of management, and a new route to worker participation. This is the first one ever built on these cybernetic principles, and it is only a beginning.

The room and its furnishings were designed and made in Chile. The optical system and control logic for Datafeed were designed and built in England, and both the animated screens were created by another British manufacturer. I have described such a room as this over many years, and once wrote: 'It is not the operational research, technology or experience that is lacking to produce the first (such) control centre. It is the managerial acceptance of the idea, plus the will to see it realized'⁶. I finally found both the acceptance and the will on the other side of the world.

The Inconclusion

This has been a very long lecture, but it deals with a very large subject: how the science of effective organization, which we call cybernetics, joins hands with the pursuit of effective freedom, which we call politics. What a new—and what a vital—issue those words betoken. Where have I heard them before?

'the cybernetics of men,
as you, Socrates,
often call politics. . .

You can tell from that name that I am quoting; and we seem to be up against a time-lag of two thousand years. But now we are doing something about it. Now we have some cybernetic tools.

What I have been able to tell you today, however, is plainly incomplete; please bear in mind that this whole thing began just sixteen months ago. Therefore, although the system exists, it is—in a proper academic sense—unproven. I expect that it, like any other infant, will be slapped on the wrist (if not worse) and told to toe the line—if not worse.

For during that period of sixteen months, various attempts have been made to overthrow the Chilean democracy. I have seen that, from fairly deep inside. Scientifically too, during that period, I have been told a hundred times that it would take more than twenty years to do what has now been done—during that period.

We have to take note that innovation, whether political or scientific, does not favour those who hold the real power. And if either kind of innovation stands to favour ordinary folk, and both these do, then it will be opposed.

For this reason, I am not naming here my many colleagues and collaborators. They know my feelings of esteem and affection for their ability, their dedication, and their friendship. What any of them asks of me that I can do, he should consider done.

For this reason also, I commend my compatriots here today to watch, more avidly than many doubtless have, what happens next in Chile. There will be lessons there for Britain, I believe; and for humanity.

So now good-bye.

I remember Richard Godman in this very place.

Requiescas in pace.

Notes

¹Norbert Wiener, Cybernetics, John Wiley, New York, 1948.

²W. Ross Ashby, Design for a Brain, Chapman and Hall, London., 1954.

³ Stafford Beer, Brain of the Firm, Allen Lane, The Penguin Press, London, 1972.

⁴P.J. Harrison and C.R. Stevens, 'A Bayesian Approach to Short-term Forecasting'. Operational Research Quarterly, Vol. 22, No. 4, December, 1971.

⁵Jay W. Forrester, World Dynamics, Wright-Allen Press, Cambridge, Mass., 1971.

⁶Beer, op cit.

I promised a further note
when you had read that story
but for the perspicuous
it is probably redundant.

The Goodman Lecture
was delivered in February 1973.

On the 11th September
Salvador Allende
that marvellous man
died in a bloody business.

Throughout June and July
I had been in Chile.
My last meeting with the
President
July 26th
was very strained—

it was obviously probable
by then
and evident to us both
that we would not meet again
in the presidential palace
La Moneda
in those circumstances.

La Moneda
is now a shell.

I have written and broadcast much
since then
about the assassination
of a poor country
by the rich world.

All that experience
is not part of this story

or perhaps
I do not want
to say much more
about it
yet
or perhaps
the whole story
is implied
by this very book
...

At any rate
I would invite
those who care
to read again
pages 423 to 429
with which I have not subsequently tampered

and to reflect
on what I
and you
have since learned.

Here is an undoctored quotation from page 424

I take advantage
only of this technique
of spacing:

It appears to me that the government did not
anticipate the full vindictiveness with which
the rich world would react to its actions,
which I emphasize have—so far—been
perfectly legal.

At any rate, a true resolution of the very
potent conflicts in Chilean society is not
discernible within the mounting instability,
and may be long postponed.

But I consider that this is largely a phenomenon
of the cybernetics of international power:
you could say that the Chilean people
have not been given a chance.

They are being systematically isolated behind
those beautiful Andes mountains, and are in a
state of seige.

To pick on some words:

vindictiveness

perfectly legal

very potent conflicts

the cybernetics of international power

STET
as the printers say

let it stand.