

Designing Freedom

Stafford Beer

Contents

1	The Real Threat to “All We Hold Most Dear”	2
2	The Disregarded Tools of Modern Man	12
3	A Liberty Machine in Prototype	22
4	Science in the Service of Man	31
5	The Future That Can Be Demanded Now	41
6	The Free Man in the Cybernetic World	50
7	Biography	57

The text of six radio broadcasts given in the autumn of 1973 as the thirteenth series of Massey Lectures which were established in 1961 by the Canadian Broadcasting Corporation to enable distinguished authorities in fields of general interest and importance to present the results of original study or research.

1 The Real Threat to “All We Hold Most Dear”

The little house where I have come to live alone for a few weeks sits on the edge of a steep hill in a quiet village on the western coast of Chile. Huge majestic waves roll into the bay and crash magnificently over the rocks, sparkling white against the green sea under a winter sun. It is for me a time of peace, a time to clear the head, a time to treasure.

For after all, such times are rare events for today’s civilized man. We spend our days boxed in our houses, swarming in and out of office blocks like tribes of ants, crammed into trains, canned in aeroplanes, locked solid in traffic jams on the freeway. Our unbiblical concern for what we shall eat, what we shall drink, and what we shall put on is amplified and made obsessional by the pressure to consume—way, way beyond the natural need. All this is demanded by the way we have arranged our economy. And the institutions we have built to operate that economy, to safeguard ourselves, protect our homes, care for and educate our families, have all grown into large and powerful pieces of social machinery which suddenly seem not so much protective as actually threatening.

Mankind has always been in battle with his environment. But until quite recently in history his battles were on a reasonable scale, a human scale. He could alter his house, if he would brave the weather: he did not have to take on the whole city planning department and the owners of his mortgage and his overdraft. He could dress his children as he pleased, teaching them what he knew and how to learn: he was not flattened in this natural enterprise by educational authorities, attended by boards of experts. When he fought with danger, he matched his strength and skill with another animal of similar size to his: he was not unexpectedly knocked flying by two tons of steel travelling at sixty miles an hour. And if he faced the fact of death, that also was a personal encounter, win or lose: he did not live under the stress of a remotely threatened genocide or nuclear extermination. But this is how it is for us. We do not think much about it. When things go badly, there is all of this to blame, and not ourselves: perhaps that is some sort of consolation.

Do we indeed even want to think about such things? I believe that people increasingly do begin to question the assumptions of our society—and not because of any characteristic that I have so far mentioned. Most people alive today in urban societies settled long ago for the role of pygmy man amidst the giants of his own institutions, and for the reason that it meant apparent advance—a higher standard of living, as measured by the gross national product per head. But in the last decade or two something has come through to public consciousness. It is the doubt as to whether the whole apparatus of our civilization *actually works* any longer. Is it beginning to fail?

The evidence for this suspicion is plentiful. I instance the decay of previously rich and healthy cities from the centre outwards, creating ghettos and all the social frightfulness that goes with them, stark inequalities, private penury, social squalor, a rise in crime, a rise in violence. I instance pollution on a world-wide scale: the poisoning of the atmosphere, of seas and lakes and rivers. Then there is the widening chasm between luxury and starvation, whereby we somehow manage to concentrate more wealth with the already wealthy, and more deprivation with the already deprived. I will not go on with this baleful list, because conscientious people are already aware of these problems. The question I would like us to address in these lectures is just *why*? Because if we can fathom that, maybe we can also conduct a fruitful search for answers.

The first point to establish is the most difficult; and it is the most difficult because it sounds so easy. It is to say that all these institutions we have been contemplating—the homes, the offices, the schools, the cities, the firms, the states, the countries—are not just things, entities we recognize and label. They are instead *dynamic and surviving* systems. Well, I did say it sounds so easy. Obviously these entities are systems;

because they consist of related parts, and the relations—the connexions—between those parts. Obviously, too, they are dynamic. No-one believes that these institutions are just sitting there brooding; they are all “on the go”. Finally, if they were not surviving, they would not be there. And having taken the point that we are talking about such systems, it is too natural to pass it by—to pass over the point, pass around the point, pass through the point—without ever grappling with the real meaning of the point at all.

Although we may recognize the systemic nature of the world, and would agree when challenged that something we normally think of as an entity is actually a system, our culture does not propound this insight as particularly interesting or profitable to contemplate. Let me propose to you a little exercise, taking the bay I am looking at now as a convenient example. It is not difficult to recognize that the movement of water in this bay is the visible behaviour of a dynamic system: after all, the waves are steadily moving in and dissipating themselves along the shore. But please consider just one wave. We think of that as an entity: *a* wave, we say. What is it doing out there, why is it that shape, and what is the reason for its happy white crest? The exercise is to ask yourself in all honesty not whether you know the answers, because that would be just a technical exercise, but whether these are the sorts of question that have ever arisen for you. The point is that the questions themselves—and not just the answers—can be understood only when we stop thinking of the wave as an entity. As long as it is an entity, we tend to say well, waves are like that: the facts that our wave is out there moving across the bay, has that shape and a happy white crest, are the signs that tell me “It’s a wave”—just as the fact that a book is red and no other colour is a sign that tells me “That’s the book I want”.

The truth is, however, that the book is red because someone gave it a red cover when he might just as well have made it green; whereas the wave cannot be other than it is because a wave is a dynamic system. It consists of flows of water, which are its parts, and the relations between those flows, which are governed by the natural laws of systems of water that are investigated by the science of hydrodynamics. The appearances of the wave, its shape and the happy white crest, are actually *outputs* of this system. They are what they are because the system is organized in the way that it is, and this organization produces an inescapable kind of behaviour. The cross-section of the wave is parabolic, having two basic forms, the one dominating at the open-sea stage of the wave, and the other dominating later. As the second form is produced from the first, there is a moment when the wave holds the two forms: it has at this moment a wedge shape of 120 degrees. And at this point, as the second form takes over, the wave begins to break—hence the happy white crest.

Now in terms of the dynamic system that we call a wave, the happy white crest is not at all the pretty sign by which what we first called an entity signalizes its existence. For the wave, that crest is its personal catastrophe. What has happened is that the wave has a systemic conflict within it determined by its form of organization, and that this has produced a phase of *instability*. The happy white crest is the mark of doom upon the wave, because the instability feeds upon itself; and the catastrophic collapse of the wave is an inevitable output of the system.

I am asking “Did you know?” Not “did you know about theoretic hydrodynamics?” but “did you know that a wave is a dynamic system in catastrophe, as a result of its internal organizational instability?” Of course, the reason for this exercise is to be ready to pose the same question about the social institutions we were discussing. If we perceive those as entities, the giant monoliths surrounding pygmy man, then we shall not be surprised to find the marks of bureaucracy upon them: sluggish and inaccurate response, and those other warning signs I mentioned earlier. That is what these entities are like, we tend to say—and sigh. But in fact these institutions are dynamic systems, having a particular organization which produces particular outputs. My contention is that they are typically moving into unstable phases, for which catastrophe is the inevitable outcome. And I believe the growing sense of unease I mentioned at the start derives from

a public intuition that this is indeed the case. For people to understand this possibility, how it arises, what the dangers are, and above all what can be done about it, it is not necessary to master socio-political cybernetics. This is the science that stands to institutional behaviour as the science of hydrodynamics stands to the behaviour of waves. But it is necessary to train ourselves simply to perceive what was there all the time: not a monolithic entity, but a dynamic system; not a happy white crest, but the warning of catastrophic instability.

So far we have spoken a little of the nature of dynamic systems; but the other qualification that I used at the start was the word “surviving”. The wave is not a *surviving* dynamic system, because its destruction is built into its organization. However we certainly regard our institutions as survival-worthy. After all, they have survived until now, because they are capable of a trick we call *adaptation*, which waves are not. So why should there now be a fuss about instability and impending catastrophe?

Our institutions have already proven that they can survive, says the argument, and we can have confidence that they will continue to adapt successfully to change. Indeed, we insist that they must—for our institutions enshrine everything we hold most dear. Beginning with the family unit, based on love and mutual support; extending through the school—and perhaps that alma mater the university; bound together in the cohesion of the neighbourhood, the community, and the churches; ramifying into business and the growth of prosperity for all; exemplified, protected, and projected by the state; this—our society—is an entity that survives, albeit by adaptive change. And if this society embodied in its institutions is threatened by too rapid change, then the answer that many serious and concerned people give is to reinforce the rules of the societary game, strengthen the institutions, tighten up the criminal, social, and moral laws, and weather the storm. That is the conservative attitude. It is not mine. It is not going to work much longer.

Indeed, we ought to face the fact that this theory does not work now. People convince themselves that it does, because they see society as an entity, and its main characteristic is to be held most dear. Then they grit their teeth and declare that whatever is wrong with it must and can be put right again. Broken barriers, swept away by permissive morality, can be repaired. Departed children can be summoned home to eat the fatted calf. The majority of people, who do not attend a church, are still (surely to goodness?) fundamentally Christian. The starving two-thirds of the world will eventually be fed (well, not those two-thirds dying right now, but their descendants). And somehow a finite planet, with exhaustible resources, will be made indefinitely to support more and yet more growth. Oh no: this only even looks possible if we are dealing with a fixed entity, a society or a way of life that is held to enshrine eternal values, a golden ideal. If this has become rather tarnished, then it can be rebrunished with a little elbow grease. So some people, and especially some politicians, seem to think.

But if society is a dynamic system all these phenomena are not simply blemishes—they are its outputs. These unpleasing threats to all we hold most dear are products of a system so organized as to produce them—to produce *them*, and not their contraries. These are not accidental; and they are not mistakes. They are the continuing output of a systemic conflict which is due to specific modes of organization. And those modes of organization have currently arrived at a stage in their inexorable pattern of behaviour which, like the wedge-shaped wave of 120 degrees, is incipiently unstable—on the verge of catastrophic breakdown. Or so I think.

I hope you will accept my invitation to investigate this hypothesis with me. And let me be more cheerful right away. These dire predictions I am making have to come about, which is why I said they were inexorable, if—but only if—we continue to support modes of organization into which these outcomes are inbuilt. We do not have to do that. We really can change the whole thing. But to succeed, we must first perceive the nature of dynamic surviving systems, and the conditions they must meet to remain stable yet adaptive.

In order to get rid of the concept of an institution as a fixed entity, we have to get rid of the classical picture of its organization. You know how this looks. The institution's activity is divided into chunks, which are also perceived as entities; these chunks are divided into smaller chunks, and so on. In every chunk there is a boss-man, with lesser bosses reporting to him and running the smaller chunks. This picture looks like a family tree; and it is useful for just one purpose. If something has gone wrong, you can use the picture to trace whose fault it is. In fact, this whole picture of an institution is just like a fault-finding chart that shows how an automobile is organized. Some people, and the channels connecting them, are shown in red (as if they were the fuel system), others are shown in blue (as if they were the electrical system), and so on. But nowhere on the automobile chart can you find such a thing as *speed*, which is what the automobile is all about.

What this orthodox organization chart leaves out of account, when it comes to understanding institutions, is that we are not dealing with pistons, pumps, and distributor arms, but with people; and the connexions between the parts are not crankshafts, pipes, and electrical wires, but human relationships. What matters about the institution is not its set of dependencies but its performance—if you like, its speed. The organizational forces by which the whole institutional machinery is held together include psychological conflict, loyalty and perfidy, integrity of purpose, hard and lazy work. They also include all manner of special arrangements making cross-linkages between the cousins of the family tree, which are the formal boards and committees, the less formal liaison officers and advisory groups, the informal old-pals network—and even maybe secret connexions whose existence will be denied. How can we picture this dynamic system in our minds, and how contemplate its output and stability?

I will ask you to think first of a tennis trainer—two poles held in place by guy-ropes pegged to the ground, and connected by a piece of elastic from the middle of which hangs a tennis ball suspended by an elastic thread. Suppose that an institution has only two members. They are sitting on the top of the poles—held firmly in place by the guy-ropes, which define their formal positions. What they must do between them is the work of the institution, and for this they need a connexion, which is the elastic thread. Now the role of the suspended tennis ball in this example is not too obvious, and it is vital to understand it. The ball stands for the *output* state of the system.

However complicated a system may be, there is one output state that defines it. That state is an output of the system—not in terms of what special things or detailed consequences are flowing from the institution's activity, but in terms of its total net performance. This is rather like saying that a human being, characterized as he is by all manner of detectable outputs, is in the net state of sleep—or violent effort, or concentration, or fighting. Or again we might liken the tennis ball to the net state of a game of chess at the (let's say) 26th move. There are all sorts of tensions implicit in the relationships of the pieces, and if the game were speeded up by cinephotography we really should see a dynamic system operating under its organizational conflicts. But if instead we take a still photograph of the board at the 26th move, then the output state is a single state, and it might be called "white losing", or "checkmate".

If the men on top of the poles do their respective jobs properly, they will pull correctly on the elastic. The ball—which marks the output state of the system—will bob about for a bit, and then be still. The dynamic system is doing its work, and producing stability. If the men are inefficient, and cannot make up their minds how to pull on the elastic (especially if they keep passing the buck), then the ball will bob about for ages, and may never settle. This system is unstable. However: if we assume goodwill and reasonable efficiency on the part of the two men, so that they behave like proper elastic connexions, the ball will soon stop bobbing about. The time it takes to stop is called the *relaxation time* of the system.

Well, this picture is a bit too simple for our needs. So let us erect a lot more poles (try to imagine about forty of them) arranged in a circle, where our original elastic connexion marks a diameter. Now let us

join all the new men on the top of these poles to the system, by giving each one a piece of elastic and tying the other end to the central knot. These new members of the institution are not all equally proficient, or loyal, or hard-working; and we can represent that by giving some of them thin pieces of elastic, and some of them thick pieces. The ball was disturbed while we did that, and I think we can bet that the relaxation time will now be extremely long. In fact, and this is really rather interesting: the harder all the men try conscientiously to manipulate the system so that it settles, the more unstable it is likely to become. Just imagine the chaos. “Hey, George, stop pulling a minute;” “Harry, you pull a bit harder.” And so on; In fact if all of forty men are each trying to give instructions to each of the others, we shall have 1,560 communication channels trying to speak all the time. You are right: *it isn’t going to work*.

The reason is that this system as a whole has too many possible states. I am not talking now about the solitary output state, but about the vast number of configurations which the organization itself can assume. Every one of those men on the poles may behave in a great many alternative ways; and these are permuted together to reveal the total richness of possible organizational behaviour. If we consider the total number of behavioural configurations that are possible, we have a measure of the system’s complexity. Let us turn this perception into a formal definition. The number of possible states of a system is called its *variety*. This will be a most useful word for us, so may I repeat: variety means the number of possible states.

Suppose that each man can do only one of two things, which is an absurd simplification after all. Then between them they can produce more than a million million possible sets of conditions for the system. It is too many; and the tennis ball will never be able to settle. At least, it will in theory. But in practice the world is not going to leave the system alone for long enough. Just imagine those poor men feeling they have almost exhausted the possibilities after a week’s work, when the cat comes into the garden, and takes a playful swipe at the ball with its paw. It is back to square one.

All our major societary institutions are high-variety systems; all of them need to have a finite relaxation time; but all of them are subject to constant perturbation—which is the word to use for the unexpected interference of the cat’s paw. How do they cope? There is only one way to cope, and all institutions use it—although they use it in many forms. They have to reduce the variety of the system. Here are some of the ways.

They may put in four more taller poles, and connect ten of the shorter ones to each. The man on the tall pole gives instructions to his ten subordinates. That reduces the total system variety, but it also interferes with the short-pole men’s freedom to do the best they can. It is in this way that freedom starts to be subordinated to efficiency; but the only alternative—which we must face—is total anarchy.

Second, they may put in a lot of rigid connexions, called rules, between the elastic threads, so that the system looks like a spider’s web. That also reduces variety. But that confounded cat keeps coming around, and spoiling the whole effort. Or suppose that the child of the house comes into the garden and takes a tremendous crack at the ball with a tennis racket. Then the system may not have the resilience to take the strain, and may collapse altogether.

A third variety reducing method used by institutions, for example banks and insurance companies, is to shoot the cat. This works, but is no fun if you are the cat. In any case, you had better not shoot the son of the house.

We have no time to go on exploring our model (for this is the name of our elastic network) but you can do that yourself. Remember these aspects of our work together so far. A dynamic system is in constant flux; and the higher its variety, the greater the flux. Its stability depends upon its net state reaching equilibrium following a perturbation. The time this process takes is the relaxation time. The mode of organization adopted for the system is its variety controller. With these points clearly in our minds, it is possible to state

the contention of this first lecture with force and I hope with simplicity. Here goes.

Our institutions were set up a long time ago. They handled a certain amount of variety, and controlled it by sets of organizational variety reducers. They coped with a certain range of perturbations, coming along at a certain average frequency. The system had a characteristic relaxation time which was acceptable to society. As time went by, variety rose—because the relevant population grew, and more states became accessible both to that population and to the institutional system. This meant that more variety reducers were systematically built into the system, until today our institutions are nearly solid with organizational restrictions. Meanwhile, both the range and the frequency of the perturbations has increased. But we just said that the *systemic* variety has been cut. This produces a mismatch. The relaxation time of the system is not geared to the current rate of perturbation. This means that a new swipe is taken at the ball before it has had time to settle. Hence our institutions are in an unstable condition. The ball keeps bobbing, and there is no way of recognizing where an equilibrial outcome is located.

If we cannot recognize the stable state, it follows that we cannot learn to reach it—there is no reference point. If we cannot learn how to reach stability, we cannot devise adaptive strategies—because the learning machinery is missing. If we cannot adapt, we cannot evolve. Then the instability threatens to be like the wave's instability—catastrophic.

I said before that there are solutions, but I have also shown that they concern organizational modes. They concern engineering with the variety of dynamic systems. By continuing to treat our societary institutions as entities, by thinking of their organizations as static trees, by treating their failures as aberrations—in these clouded perceptions of the unfolding facts we rob ourselves of the only solutions.

In particular: by advocating a new insistence on variety reducing methods which worked in a bygone epoch, we advocate precisely the wrong thing, and seal our doom. THIS is the real threat to all we hold most dear.

NOTES IN SUPPORT OF THE FIRST LECTURE

Cybernetics

originally defined by the late Norbert Wiener as
*the science of control and communication
in the animal and the machine.*

Cybernetics is exactly thirty years old; the name was added in 1947.


This original definition points to the relationship between control and communication, and to the existence of general laws affecting equally animate and inanimate systems. The first principle of such general importance to be recognized was the significance of feedback in all systems, whatever the fabric of their components.

Thirty years on, this new definition might be preferred:

Cybernetics is the science of effective organization.

Variety


*is the measure of complexity in a system,
defined as the number of its possible states.*



How the wave becomes unstable, and finally moves into catastrophic collapse.


after René Thom

Stabilité Structurale at Morphogénèse
Benjamin, Massachusetts, 1972



The poles with their guy ropes define the **formal positions** that people hold in an institution.
The ball defines a point representing the net output state of a system.


The **relaxation time** of a system is the time it takes the **representative point** to reach stability after it has been **perturbed** (for instance, by the cat—which stands for an arbitrary interference).



The larger the organization, the longer the relaxation time is likely to be. ...

If everyone has complete **freedom** to do what he thinks is best, there will (unfortunately) be **instability**—which may feed on itself and become catastrophic.

If everyone is trying to communicate with everyone else, there will be $n(n-1)$ communication channels open. When n = only 40, $n(n-1) = 40 \times 39 = 1560$.



Three of the main ways by which institutions reduce their variety:

1. A boss constrains the freedom of his subordinates,
2. Rigid connexions called rules constrain the interactions of the elastic threads,
3. Someone shoots the cat: the institution does not accept arbitrary interference, and forces those with whom it interacts into stereotypes.

Proliferation of Variety

If there are n people in a system, and each of them has variety x (each can adopt x number of possible states), then the variety of the total system thus defined will be x^n .

So if there are only forty people ($n = 40$), each of whom has only two possible states ($x = 2$), there are still 2^{40} possible states of the system.

$$2^{40} = 1,099,511,627,776$$

In the real world, we find that increased freedom (including new opportunities) proliferates variety to a point where our respected institutions cannot any longer cope with it.

Hypotheses

The relaxation time of the institutional system is now on the average longer than the average interval between perturbations, with the result that the institutional system is permanently unstable.

Since permanent instability feeds on itself (because there is no recognizable stable condition on which to base learning and adaptation), this instability is likely to become, like the wave's instability, catastrophic.

2 The Disregarded Tools of Modern Man

If you were sitting under an apple tree, and suddenly an apple fell on your head, to what force would you ascribe the event? Is there an apple-throwing monkey in the tree? Certainly not: the force involved is called gravity. Because we understand this force to be quite general on the planet, we do not propose to experiment when asked what would happen if we dropped the glass of water we are holding, or what would happen to an aircraft whose engines suddenly stopped. We reckon we *know* what would happen.

And yet the concept of a force which affects everything on the planet, regardless of what that thing is, and affects it to the same extent (so that a pound of feathers behaves just like a pound of lead, for equal air resistance), is a very difficult concept indeed. How *can* there be a force which affects everything, that you cannot directly experience, that finds expression only in mathematical terms, and that (counter to all intuition) treats feathers and lead the same?

Come, come, you will say: this is to talk like an ignorant savage. Even a child can answer those questions; and besides you *can* experience the force of gravity directly, because it is this force that your body measures when it senses weight. To this I reply: do you think you would have answered thus if you had lived around the corner from Isaac Newton in the year 1687?

Our culture has had nearly 300 years to understand the problems of Newtonian physics. It has had more than half a century to get its grip on relativity theory and the second law of thermodynamics, knowing that it is at any rate *possible* to make general statements about the physics of the universe. Not all of us, I dare say, would care to answer basic questions about these two, although one might have supposed that the culture would have imbibed them by now. The observed fact is that the culture takes a long, long time to learn. The observed fact also is that individuals are highly resistant to changing the picture of the world that their culture projects to them.

I am trying to display the problem that we face in thinking about institutions. The culture does not accept that it is possible to make general scientific statements about them. Therefore it is extremely difficult for individuals, however well intentioned, to admit that there are laws (let's call them) that govern institutional behaviour, regardless of the institution. People know that there is a science of physics; you will not be burnt at the stake for saying that the earth moves round the sun, or even be disbarred by physicists for proposing a theory in which it is mathematically convenient to display the earth as the centre of the universe after all. That is because people in general, and physicists in particular, can handle such propositions with ease. But people do *not* know that there is a science of effective organization, and you are likely to be disbarred by those who run institutions for proposing any *theory* at all. For what these people say is that their own institution is unique; and that therefore an apple-growing company bears no resemblance to a company manufacturing water glasses or to an airline flying aeroplanes.

The consequences are bizarre. Our institutions are failing because they are disobeying laws of effective organization which their administrators do not know about, to which indeed their cultural mind is closed, because they contend that there exists and can exist no science competent to discover those laws. Therefore they remain satisfied with a bunch of organizational precepts which are equivalent to the precept in physics that base metal can be transmuted into gold by incantation—and with much the same effect. Therefore they also look at the tools which might well be used to make the institutions work properly in a completely wrong light. The main tools I have in mind are the electronic computer, telecommunications, and the techniques of cybernetics...

Now, if we seriously want to think about the transmutation of elements in physics, we will recognize that we have atom-crackers, that they will be required, and that they must be mobilized. We shall not use the atom-crackers to crack walnuts, and go on with the incantations. But in running institutions we disregard our tools because we do not recognize what they really are. So we use computers to process data, as if data had a right to be processed, and as if processed data were necessarily digestible and nutritious to the institution, and carry on with the incantations like so many latter-day alchemists.

The invitation to face up to these realities is a necessary one if there is to be any real chance of perceiving the proper role of currently available tools. For it is not something scintillatingly clever that I am proposing, not a complicated new extension of mind-blowing techniques that are already beyond most people's understanding, not a "big brother" that will alienate us still further from the monstrous electronic machinery that by now seems to govern our lives.

I am proposing simply that society should use its tools to *redesign* its institutions, and to operate those institutions quite *differently*. You can imagine all the problems. But the first and gravest problem is in the mind, screwed down by all those cultural constraints. You will not need a lot of learning to understand what I am saying; what you will need is intellectual freedom. It is a free gift for all who have the courage to accept it. Remember: our culture teaches us not intellectual courage, but intellectual conformity.

Let's get down to work, and recall where we were. A social institution is not an entity, but a dynamic system. The measure we need to discuss it is the measure of variety. Variety is the number of possible states of the system, and that number grows daily, for every institution, because of an ever-increasing range of possibilities afforded by education, by technology, by communications, by prosperity, and by the way these possibilities interact to generate yet more variety. In order to regulate a system, we have to absorb its variety. If we fail in this, the system becomes unstable. Then, at the best, we cannot control it—as happened with the bobbing ball on our elaborated tennis trainer; at the worst, there is a catastrophic collapse—as happened with the wave.

So next to something new. What is it that controls variety? The answer is dead simple: variety. Variety absorbs variety, and nothing else can.

Examine first of all the truth of that statement. Consider for example all the customers who are inside a departmental store. From the store's point of view, this represents an awful lot of variety that has to be controlled. You will notice how I use the word "control". It isn't that these ladies and gentlemen have to be told what to do, and made to do it. It is that when one of them wants to buy something, the variety of the customer/store system goes up: a possible state has been made actual. The lady has chosen a pair of shoes, but she might have chosen a fruit cake. The store has to absorb this variety: there had better be someone on hand rather quickly to take the money, and wrap up the shoes. Moreover, there will also need to be someone on hand to do the same for the fruit cake.

But not for nothing is that store called departmental. There is a shoe salesman, and a cake salesman; that is what organizational structure is *for*—to carve up the total system variety into subsystems of more reasonably sized variety. The customer who is not clear what commodity, if any, will meet her need, represents variety that cannot be trapped by this departmental arrangement; her variety will be left over, not absorbed, if the store is not careful—and we can see how this means that the situation is out of control. But if the store is careful, it will have an information bureau—which exists precisely to absorb this excess variety.

Let us return to the shoe purchaser; we observe that she is becoming angry. This is because she cannot get any attention. The shoe salesman is dealing with someone else, and four more people are waiting. The other shoe salesmen are similarly occupied. Temporarily, at any rate, the situation is out of control,

because at this moment the store has miscalculated the number of shoe salesmen needed to absorb the variety generated by the customer. Well, maybe you remember the concept we need to describe this affair, and its name. The name is relaxation time. Variety is cropping up faster in this system than the system can absorb it, and this is bad from the customer's point of view. If it happens all the time, it will be bad from the store's point of view as well: the customer will desert the store, looking for somewhere with a shorter relaxation time. So the temporary instability of service in the store will become permanent, and—at that very moment—incipiently catastrophic. The trouble with our societal institutions, of course, is that the citizen has no alternative but to use them.

Only variety can absorb variety. It sounds ridiculous, but the perfect, undefeatable way to run this store is to attach a salesman to each customer on arrival. Then we could forget about those departments, where the shoe salesmen are run off their feet, while the girls in lingerie are manicuring their fingernails, and absorb the customers' variety as we go along. For, you see, not only do we need variety to absorb variety, but we need exactly *the same amount* of variety to do it. We were speaking just now of the law of gravity in physics: it is perhaps the dominant law of the physical universe. What we have arrived at in the departmental store is the dominant law of societal systems, the Law of Requisite Variety—named Ashby's Law after its discoverer.

The example is ridiculous, because we cannot afford to supply requisite variety by this obvious expedient. We cannot give every departmental store customer a salesman, because we cannot afford it; but you may already have noticed that in very superior (and therefore very expensive) special-purpose stores, such as those selling automobiles or hand-made suits, this is exactly what happens. In fact you cannot shake the fellow off. Nor would you be able to shake off your personal policeman, if half the population were enrolled as detectives to spy on the other half. It is just because this is impracticable that we have crime. We cannot meet the demands of Ashby's Law. But we must come somewhere near it, somehow, or we are in for catastrophic collapse. How is this done?

When varieties are disbalanced, as they usually are, we structure our organizations to cope. Fundamentally there are two ways, and only two ways, of doing this. Remember: you cannot repeal the laws of nature, and Ashby's Law *will exert itself*. We met the first way in the last lecture. It is to reduce the variety generated by the system so that it matches the available supply of regulatory variety. You may recall the men we had sitting on the tall poles: this was their job. And I made a crack at the financial institutions for doing the same thing. I said they "shot the cat" that was generating variety by swinging the tennis ball. In other words, there is no way in which an insurance company can provide requisite variety for *you*—the unique human being: your variety must be attenuated, cut down so that your case fits a more general case. In theory, you can get the entirely personal quotation that corresponds to having your personal salesman or your personal policeman. Just try it (I have) and see what the insurance is going to cost. Now that is perfectly reasonable, because the claim to be able to handle every citizen's variety is as ridiculous as I admitted my example to be.

One trouble with our institutions is that they do not admit that the only full solution would be ridiculous and that therefore it is not done. They should acknowledge, at least to themselves, that they are satisfying Ashby's Law by variety attenuation. And the reason they do not admit it is that it sounds bad in the ear of our culture. Our culture insists on the uniqueness of the individual, but our society cannot live up to that. This is no criticism, it is a scientific fact. Our culture also insists on the absolute freedom of the individual, but our society cannot live up to that either. That too is a scientific fact. Then look at the mess we get into by our pretences. Instead of understanding the situation, and using science to do the best possible job of variety attenuation under the guidance of the consumer and of the electorate, we spend vast sums of the nation's wealth—our wealth—on phoney claims. On the consumer side, we put on the

advertising pressure to pretend that full account is taken of the customer's variety—which is impossible. On the electoral side, we lose the freedoms we have, when our variety is attenuated, because we are not asked how the attenuation should be done. No politician would dare to ask his electorate that question, because he is too busy standing for the inalienable rights which it is perfectly obvious we have not in any case got. Nor can we have them: let's look the facts in the face.

The second method of satisfying Ashby's Law is the one used by the departmental store. It would be madness to attenuate the customers' variety—by stocking only one kind and one size of shoe, for example, or by locking her in until she was served. The alternative is to amplify the variety of the regulatory part of the total system. You do not appoint a single salesman, but many, and you do this by departments; in each of them you try to calculate the statistical relationship between salesmen and customers, so that the relaxation time is held steady. There are scientific ways of doing this, but they are not often used. The amplification of regulatory variety comes in where one salesman handles *many* customers *effectively*. In societary systems, this is the preferable way to proceed, because it helps to preserve individual freedom. We do not do it, for several reasons—but only one is real. The others are subservient to this; Think of the outstandingly spurious reason first.

This is the most obvious reason, the cost; and there are two things I want to say about that. First: the cost of a societary system is taken as its primary measure in our society, whereas the valid measure to use is variety. The cost is not what it claims to be, namely the money that must necessarily be spent to achieve certain aims. In fact, the cost is whatever is provided in the way of funds, and this total available sum operates to constrain what we shall do—within the framework we already have. If this framework is merely a joke played on us by history, so that its structure is outmoded and its relaxation time is too long, then the system will not work: period. The thing is unstable. We pour in money to head off catastrophic collapse. But that is a silly thing to do. Because the money will—naturally enough—be spent on buttressing the framework, feeding in expensive variety to stop the relaxation time growing longer still, and making the instability actually worse (think of all those poles on the tennis trainer).

Second: the constraint of cost is entirely relative to the aims of society. The aims of society are of course a set of variety attenuators. We cannot do everything we conceive as socially desirable, and our aims—our programs—cut down the proliferating variety of societary choice. In so doing, the attenuator allocates the revenue that the taxpayer will bear. Once again, our culture prefers us not to look on the matter this way round. There is a total plausible revenue: come, then, attenuate variety—choose between developing a vertical take-off aircraft, fighting a small economic war with the United States for the ownership of Canadian industry, and doubling the price paid for bureaucracy by having everything done in two languages. ...

In the deathly silence I seem to hear following *that* solecism, I also seem to hear a faint voice (it must be a madwoman) crying: "I would sooner have a twenty-four hour child care service." But that would cost money, and there is none remaining. Don't be ridiculous madam, it's my prerogative! The point is that our variety attenuators are built into the system. Freedom of choice has gone down this particular drain. So what about the antithetical variety amplifiers? If the spurious reason for not using them is cost, the real reason is that it would mean redesigning everything—so as to get rid of the built-in attenuators, and install instead the amplifiers that could really work to achieve requisite variety, viable relaxation time, and hence some sort of social stability.

I have already suggested a list of three basic tools that are available for variety amplification: the computer, teleprocessing, and the techniques of the science of effective organization, which is what I call cybernetics. Now I am saying that we don't really use them, whereas everyone can assuredly say: "Oh yes we do." The trouble is that we are using them on *the wrong side of the variety equation*. We use them without regard

to the proliferation of variety within the system, thereby effectively increasing it, and not, as they should be used, to amplify regulative variety. As a result, we do not even like the wretched things.

If one of those unworkable institutions we were discussing buys a computer, what happens? It uses the thing, please note at enormous “cost”, to do more elaborately exactly what was done before. And, as we know, that didn’t work. One famous computer manufacturer ran an entire sales campaign for its new series of machines on the slogan that you needed simply to transfer the existing system whatever it was, lock, stock, and barrel, to the new miracle machine.

Well, that sounded good to the financial director, who had learned the hard way that immense costs are involved (yes, those same costs again) in implementing new hardware in software terms, but had not learned that the machine to do miracles is not yet invented. The effect on the institution was inevitable: to make the existing instability more unstable. It’s obvious really, once the concept of variety and the law of requisite variety are clear. The computer can generate untold variety; and all of this is pumped into a system originally designed to handle the output of a hundred quill pens. The institution’s processes overflow, just as the crest of the wave overfills, and there is a catastrophic collapse. So what do we hear? On no account do we hear: “Sorry, we did not really understand the role of the computer, so we have spent a terrible lot of money to turn mere instability into catastrophe.” What we hear is: “Sorry, but it’s not our fault—the computer made a mistake.”

Forgive my audacity, please, but I have been “in” computers right from the start. I can tell you flatly that they *do not* make mistakes. *People* make mistakes. People who program computers make mistakes; systems analysts who organize the programming make mistakes; but these men and women are professionals, and they soon clear up their mistakes. We need to look for the people hiding behind all this mess; the people who are responsible for the system itself being the way it is, the people who don’t understand what the computer is really for, and the people who have turned computers into one of the biggest businesses of our age, regardless of the societary consequences. *These* are the people who make the mistakes, and they do not even know it. As to the ordinary citizen, he is in a fix—and this is why I wax so furious. It is bad enough that folk should be misled into blaming their undoubted troubles onto machines that cannot answer back while the real culprits go scot free. Where the wickedness lies—and wickedness is not too strong a word—is that ordinary folk are led to think that the computer is an expensive and dangerous failure, a threat to their freedom and their individuality, whereas it is really their only hope.

There is no time left in this lecture to analyse the false roles of the other two variety amplifiers I mentioned—but we shall get to them later in the series. For the moment, you may find it tough enough to hear that just as the computer is used on the wrong side of the variety equation to make instability more unstable, and possibly catastrophic, so are telecommunications used to raise expectations but not to satisfy them, and so are the techniques of cybernetics used to make lousy plans more efficiently lousy.

But enough is enough. I expect that you would be more interested to hear what can be done about all this, than to hear more dreadful news. Then: may I tell you that the next three lectures will consider constructive policies for handling variety. We shall start with the state itself. This is rather bold, but you will understand by now that I believe we are all captives of gigantic systems beyond our individual control, and we need to understand them. Because, as long as we have any semblance of democracy, they are not beyond our collective control. Next we shall turn to the only thing that matters: the individual. He and she—not to mention their son and their daughter—are enmeshed in all this machinery, and they have to get out. If science cannot join with politics and management to do that, I might as well be giving the Massey Lectures in that ominous year of 1984, instead of in the year of 1973. On that reckoning we have eleven years; and frankly that is about the limit. The third of these three next sessions will return to the central topic of the institution.

And so, before I come to my last point, I would ask you—if you do me the courtesy of staying with these talks—to think over a crucial question before we meet again. It is the central importance of the law of requisite variety. Please think it

through. Think of any complex, dynamic system. How is it regulated? It will certainly be proliferating variety. Is it not true that only variety can absorb variety? How is it done? Do you not indeed find that in successful systems systemic variety is attenuated, while regulative variety is amplified? It is usually a mixture of the two. I promised you that this stuff is easy, once you break with the expositions and explanations dished out by the culture. The thinking I ask of you does not have to be done on some remote level of abstraction, or at the highest level of state affairs. Cybernetic laws are universal.

For instance, you might think over how these concepts of variety, relaxation time, stability, and potential catastrophe, work for and against you on the journey to work. You can do exactly the same exercise in the course of putting those high-variety young children to bed. Then think through how these concepts work in the big institutions of society that govern our lives. Ask yourselves how telecommunications—the telephone, television—affect those problems. Contemplate what you surely know about the role that computers play, and see if they are working on the right side of the variety equation. Finally, if it is all so easy, talk over the problem as to how we managed to get it all so wrong. Then maybe, you will see why I contend that there have to be some big changes, and that they have to come fast—before it is too late.

Then I come to the last point, which I hope will help in these deliberations. If the law of requisite variety is to be handled intelligently, and not just by leaving nature to find the variety balance (which of course can be nasty for us humans), then it follows that the regulative forces must not only dispose requisite variety—which is a number of possible states; they must also *know* the pattern by which variety in the system is deployed. On the journey to work we need to have enough options open; we also need to know the pattern of the highways—where they run, what the control points are

like, what other drivers habitually do. In the process of putting the children to bed we need several variety amplifiers at our command; but we also need to know (as we do, but let's make it explicit) the likely behaviour pattern of the children. Without these known patterns, proliferating variety looks even more threatening than it really is, which is bad enough.

What I have been calling a pattern is what a scientist calls a model. A model is not a load of mathematics, as some people think; nor is it some unrealizable ideal, as others believe. It is simply an account—expressed as you will—of the actual organization of a real system. Without a model of the system to be regulated, you cannot have a regulator. That's the point. And you can test that too.


NOTES IN SUPPORT OF THE SECOND LECTURE

1687 was the year of the completion of Newton's
Philosophiae Naturalis Principia Mathematica
containing the laws of force.


The Law of Requisite Variety (Ashby's Law)

only variety can absorb variety.


If varieties in a regulatory system are disbalanced, the system cannot attain stability. Assuming that the regulator has the smaller variety, there are only two ways of meeting the demand of Ashby's Law. One is to attenuate variety in the system, the other is to amplify variety in the regulator. These strategies can be mixed.



Examination of institutional systems often reveals that the attenuators and the amplifiers have been installed in the wrong loops—on the wrong side of the equation.




Ashby's Law Operating in a Departmental Store




Unstable: variety of store less than that of custom
= inadequate service

Unstable: variety of store more than that of custom
= unprofitability




Stable: variety matched on a one-to-one basis



Stable: variety matched on a many-to-one times occupancy basis


How the Abused Computer Replaced the Quill Pen

Year 1873




The public supplies minimal information to the institution and receives minimal information in return. The management receives minimal information too.

Year 1953




A larger public is asked for much more information, and receives much more in return. Quill pen administration continues, although the systems are mechanized. The management is threatened with an excess of information.

Year 1973



Inordinate demands for information are made on the public, which receives much less *useful* information than before. Quill pen administration continues, although the systems are computerized. The management is inundated with indigestible data.

How To Use the Computer According to Cybernetic Principles



The Public is conceived as a system, a model of which is contained in the computer. The public supplies minimal information, which the computer then synthesizes in the model. This amplifies variety as required to help the public, and attenuates variety to help the manager—thereby meeting the requirement of the law of requisite variety for each of them.

There is an evident risk in installing a model of the public in the computer, since the return loop might be misused by a despotic government or an unscrupulous management. In considering this however we need to bear in mind the cybernetic fact that no regulator can actually *work* unless it contains a model of whatever is to be regulated. Much of our institutional failure is due to the inadequacy of the contained models.

It is perhaps more alarming that private concerns are able to build systems of this type, without anyone's even knowing about their existence, than that democratically elected governments should build them in open view and with legal safeguards.

3 A Liberty Machine in Prototype

The context within which most institutions operate is still the nation state, although this situation is rapidly changing. In some ways, national sovereignty is ceded to supra-national blocs; in other ways, smaller nations find their affairs profoundly constrained by the behaviour towards them of the big powers; in yet other ways, national sovereignty is just by-passed by the world-view of their own operations taken by the giant multi-national corporations.

I have no idea what can be done to bring scientific analysis to bear on the effective organization of this global mess, unless the United Nations itself determines it. But it has seemed to me for a long time that anyone nation, thanks to the power of modern communications, could gather itself together and make that kind of effort on its own behalf. So I have spoken and written these many years. In particular, I have expressed the view that the whole business of government, that gargantuan institution, is a kind of machine meant to operate the country in the interests of individual freedom. But, for just the kinds of reason examined in the first two lectures, it does not work very well—so that freedom is in question to a greater or lesser extent in every country of the world. So, I declared, let us redesign this “liberty machine” to be, not an entity characterized by more or less constraint, but a dynamic viable system that has liberty as its output. The two conceptions, as you know from the first two lectures, are utterly different.

What then is the problem? There is no need to be overawed by the pomp and circumstance of the state, once we have found the scientific way in to the problems of effective organization, and understood the basic laws of variety. But is it lese-majesty to declare that the state runs on the same model as a departmental store, and has the same problems? Perhaps it would sound better to put it the other way round. Scientifically speaking, it just doesn’t matter—because the pattern for handling variety, which last time I called the model, is the same. If I say that it is precise to express this point by saying that the two organizations *map onto* the same model you will have no trouble. Because fortunately the phrase “map onto” in mathematics has exactly the meaning you would expect. A map is the pattern of something, represented with much attenuation of variety, but with its significant elements preserved.

Government handles its gigantic task of variety reduction by departments, just like the store. And like the store, it needs a supply of information about the on-going state of affairs. It needs a lot *more* information than the store, true; but that is irrelevant, so long as both institutions are deploying requisite variety—and this is where the mapping holds and the model is the same. The model also tells us that the relaxation time of the system must be shorter than the average interval between shocks, otherwise instability will set in. We saw how the departmental store handles that problem, and noted that if it failed there would be a catastrophic collapse, signaled by the desertion of custom. Now if government gets into that kind of difficulty, it is, more difficult to recognize. By and large, the customers will not desert. That is to say, the nation’s population will not pack up and emigrate en masse. That would be a grave decision; besides, to be a little more cynical, the people may not be able to think of anywhere else to go that is any more stable than their homeland. This problem is very general, in fact it is universal, which is why it is so important.

How do we set about the diagnosis? In government, variety is handled by attenuation in four main ways. First, models are made of the country by every government department. At least, we just have to dignify the patterns that government has in its head as models, although in all honesty the mapping onto reality is not very good. For one thing, the senior departments were created long ago, and the national variety generators onto which they are supposed to map have changed their nature, their emphasis, and their rate. So new departments have to be set up to handle the excess variety, just as happened in our store example. But in governments’ case the results cannot be so good. The store *does* change its departments with the change of life’s emphasis, and the small excess variety is absorbed by the information counter.

Government, however, does not change much; indeed it is singularly unadaptive, in my opinion because it has lost recognition of its stable state. So the excess variety for government turns out to be something really rather serious. An excellent current example is the total problem of the environment, a huge aggregation of unmapped variety. Naturally, the new department falls between the stools of all its old, component departments: it happens in every country. What we are witnessing here is the phenomenon of change that involves no actual alteration.

The second variety attenuator is the model that each department has of the component enterprises for which it is accountable. Consider the economy. Then, for example, there will have to be a model of industries, by industry, and models of those industries by product, by investment, by labour force, and so on. This is fine until we realize that these models treat the components as if they were entities, characterized by product, by investment, by labour force, and so on. But of course they are nothing of the kind. They are all, these component industries and their component firms, dynamic viable systems; and the items we were listing are continuously variable inputs and outputs. In fact, mostly the things that interest us are best described as *outputs* of these systems, since the output determines most of the input. Investment (which sounds like an input), like raw material (which sounds like an input), is attracted into the system by the size and shape of the output. At any rate, these are concepts of system we are handling now—that much is clear; and it is also clear that it is not very clever to attenuate variety by freezing the continuous variables into arbitrary time epochs (such as periods, quarters, years, and five-year plans) when the essence of the business is the way it continuously generates its output through time. We need to observe continuous time trends, that is to say: rates of change, gradients, step functions, and so on. All that is necessary because we, as government, have to look to the *interactions* of dynamic systems. The models we have were constructed primarily to inform shareholders as to the proper custody of their money. No wonder the models don't map onto each other.

When data have been generated by these low-variety models, covering long and static time epochs, within departments which no longer quite fit reality, they are aggregated. So this is the third variety attenuator. There are two main reasons for this aggregation. In the first place, a minister for example cannot handle all the raw information being generated, even though its variety has already been twice attenuated. His brain does not have requisite variety to match the data until they *have* been aggregated. The other reason is that commercial security demands aggregation, otherwise it may be easy to spot what a competitor is up to—if he is big enough, or localized enough, to stand out in the official statistics. Even so; both these reasons for aggregation of some kind do not justify aggregation of the kind we have—the total or the average—which, as I said, kills that subtlety of information that requisite variety demands from a dynamic system. Just imagine a doctor calling on his patient in hospital to be told by the nurse: “Over the last month his temperature has averaged 98.4 ” or: “Yesterday your thirty patients had an aggregated average temperature of 98.4.” These statements may be true, indeed it would be most surprising if they were not; and aggregation of *some kind* is needed if the doctor is not to sit all the month staring at the thermometer. But to decide which aggregations are the ones to use is a problem in cybernetics, not economics—and still less in administration.

The fourth variety attenuator of government information is by far the most dangerous. It is the delay imposed by the methods of collection and variety attenuation. How does *delay* turn out to be an attenuator of variety? Well: the situation as it really is today includes all the information that led up to it, and the most recent part of that series of data is doubtless the most important. So delayed information chops off the latest half of the variety implicit in the situation now. That means that government does not have requisite variety. It is very proper that economic statisticians should talk, as they do, about the “timeliness” of official information. But I want to dwell as before on something a little different in concept from an entity called “the statistical tables for last June”, characterized by whatever degree of out-of-dateness. I dwell

on the fact that when the government acts, it is perforce reacting to a situation where the statistical delay often happens to be half a cycle in the economic rise and fall of prosperity, so that the government may find itself doing exactly the wrong thing most of the time. (I said “happens to be” but I believe that these two facts are actually linked within the system’s dynamics, so that the machinery for taking decisions locks on to the wrong part of the economic cycle.)

Then these are some of the problems of the four variety attenuators. Of course, my friends in government will groan. I am not telling them anything new, and I know it. But I believe I am explaining this familiar problem in a new way, a way which enables us to do something about it. We have fast communications: ah, but they are not mobilized. We have computers: ah, but they are busily taking over exactly the old system—and are actually taking longer than before to do the job. So look again now at the diagnosis. When the institution of government was consolidated, we had huge problems of variety attenuation which had to be solved by very crude methods. That was because all our facilities were low-variety facilities—we had no computer with remote terminals *then*. Now we do have these tools, but they are disregarded. These essentially high-variety regulators are used on the wrong side of the equation, not to cope with the truly high variety generated by the dynamic system, but to accept much attenuated variety from poor models with a time lag, and then to generate their own variety inside government. That is to say, a minister can always call for an elaborately reorganized set of data, on which complicated mathematics have been done; but it is the computer that generates the variety, and not the real world. This is quite fundamental nonsense. We are using our powerful tools to automate and to elaborate the limited processes that we managed to achieve with the unaided brain and the quill pen-processes which our new tools were invented precisely to transcend.

With this diagnosis in our minds, the prescription for action is not difficult to understand. First of all, we need better models of the components of the economy, and they must be *dynamic* models. That simply means that we need to see how the parts are inter-related on a virtually continuous basis. Then instead of the lawyers’ and shareholders’ balance sheet model, instead of the accountants’ and managers’ profit-and-loss account model, instead of the input-output matrix beloved by economists, think of a model set down in terms of dynamic variety. This will be a simple flow chart, in which variety is symbolized by the relative thickness of the lines of flow, and the relative size of the boxes that indicate processes acting on the flow. The time lags in the system have to be shown, and this is best done by animating the flow chart so that the lines of flow move at different speeds.

All this is very easy to do inside any firm, and it is an interesting question to ask who ought to do it. Our rather technocratic culture immediately answers: a team of engineers, or accountants, or operational research men. Well, I don’t believe it. That just is not necessary; and besides no one will really be interested in the model they create. The people who know what the flows are *really* like are the people who work in the middle of them: the work-people themselves. And if their interest can be captured in putting together the total model of how the firm really works, we shall have some genuine worker participation to replace a lot of talk about worker participation.

Please note the reappearance of that basic distinction between entities and dynamic processes. By what means does the firm try to give its people a sense of participation in the business? We have seen that static entity the work’s committee, characterized by its standard set of debating points that arise at every meeting; we have seen that static entity the house journal; we have seen in each production department that static showcase in which the assemblies to which our products contribute pathetically gather dust. All of this has a certain unreality, because it freezes the dynamics of living and working in its tracks.

Then contemplate a company that is run from a control centre, in which the dynamic flow chart, continuously reflecting the world outside by teleprocessing data into it, constantly holds the pattern, and uses

the computer constantly to monitor all that variety. We are near to this concept in running a battle, or a warship, or an electricity supply system. But for some cultural reason the whole notion is alien to running the economy. Yes: despite all this talk about the firm, I have not forgotten that we are supposed to be discussing the management of the economy. The fact is that the total picture of industrial activity held by the government is made up of separate pieces, just like a jigsaw puzzle. The pieces are representations of economic sectors, or industries, and if these representations are not clear, the total picture when completed will be a total confusion. But there is a jigsaw within the jigsaw: the pieces of each industry are its component firms. So we must start in the right place. The picture of the firm must be sufficiently clear as to contribute to a clear picture of the industry. The picture of the industry must be sufficiently clear as to contribute to a clear picture of the industrial economy. And obviously, the question arises whether an autonomous firm will agree to collaborate in such a scheme. The reply is that government has many inducements to offer in obtaining the information it needs, and the greatest of these inducements is the fact that industry cannot expect sympathetic treatment from government policy if it will not contribute useful and timely information.

Then we can see what our potential model of the whole economy looks like. It consists of a dynamic system of simple models of dynamic systems, fitting into each other like Chinese boxes. Each box is called a level of recursion, because what we are doing is to reduplicate a cybernetic system of regulation recursively, that is over and over again, using the same components with appropriate variety adjustments. The law of requisite variety has to be satisfied at each level of recursion so that stability is induced, and off we go. Information continuously passes up and down this recursive system, appearing in its right form in the control room of the level concerned.

Now see what has happened to the problems of time lag and aggregation. Instead of accepting those problems, and misusing computers in the attempt to make adjustments for them by re-injecting variety on the wrong side of the equation, we have magically disposed of the problems altogether. I urge this precept on you: it is better to *dissolve* problems than to solve them. If time lags are a nuisance, don't have any; use teleprocessing to eliminate the lag. If aggregation is a nuisance, do away with it: use computers to attenuate variety more cleverly.

The vision I am trying to create for you is of an economy that works like our own bodies. There are nerves extending from the governmental brain throughout the country, accepting information continuously. So this is what is called a real-time control system. Why should governments be trying to deal today with last summer's problems—which are, in any event, settled one way or another by now? Then does this then mean that government will be flooded with masses of data that it cannot handle? Certainly not. My brain and your brain at this moment are both accepting all manner of sensory input—everything in the room is registering there, and that is good, because we may need to attend to something quite suddenly. Until that need arises, however, our brains automatically inspect all this irrelevant input, and *filter out* most of it.

This is what I mean by using computers as variety handlers on the right side of the equation. They have to accept all manner of input, and attenuate its variety automatically. What they will pass on to the control room is whatever *matters*. Now we tell our brains what matters to our bodies by detecting inputs that are deviating from what would normally be expected. Everything else maps onto the understood pattern in the model. Inputs fluctuate of course, but they fluctuate within limits that can be continuously calculated by probability theory—if you have a computer. So to recognize what *matters* the computer will need to make very very complicated calculations on every item of data coming in, and assess the chances that something novel is happening. In the huge majority of cases, nothing will be happening—in which case the input item can simply be discarded. It does not need to be stored in those gigantic data banks we keep hearing about, because it has no significance at all. We already know from our basic model what the ordinary variation

is, and this input item lies within it. So what? Unless you have shares in electronic storage equipment, or are building a career as a bureaucrat, you will see no reason for keeping it. And if you are particularly interested in freedom, you will see every reason for throwing it away.

It now becomes clear why I was making those remarks earlier about cost. As soon as you think of running the economy by computer, the culture promptly feeds you an image of acres of expensive equipment. It is not required. What is required is an ordinary computer, with teleprocessing interfaces between itself and its inputs from the country and itself and the control room, plus an extraordinarily clever program. The cost is in that software, and not in acres of hardware, its maintenance and staffing. But if the regulatory model is the same at every level of recursion in the economy, because the cybernetic models map onto each other, only one set of software is required.

So now we may visualize the control room and the decision takers within it acting together in symbiotic relationships as a brain for running the economy. The nerves that feed information continuously to that brain pass into its computing lobes, where the variety filters work. The basic model appropriate to this level of recursion—which is inhere running as an animated flow chart—is the focus of all discussion. It is constantly updated by the various kinds of alerting signal that the computers produce. These signals will cause the decision takers to call for more information, and they will use electronic storage to project that onto screens. Then, if they really wish to get down to serious decision taking, they will activate their dynamic systems models with their new data—and try out alternative policies by simulating them at very fast speed.

You probably know that it is possible by electronic simulation to make a ten-year-ahead projection instantaneously, and then to change your policy and see what difference it makes. This is to take an experimental approach to policy making, doing the experiments in the laboratory of the control room. So instead of experimenting on the poor old nation, and discovering ten years later that your policy was wrong, you can test and discard a dozen wrong policies by lunchtime without hurting anyone. After lunch maybe you will find a good policy.

Now it is true that much controversy surrounds simulation techniques. But this is mainly because of the inadequacy and the belatedness of the data fed into them. What do you expect with those low-variety models, aggregations, and time lags? But by *redesigning the whole institution* (and it is the only way) all those problems are dissolved, and the data feeding the simulations will be immediate, significant, and real.

Clearly this is no more than a brief sketch of the advocated approach, which is available in full, but of course much more lengthy, detail elsewhere. The fundamental criticisms that it has evoked are four. One says that instead of this being a Liberty Machine it means the death of individual freedom. I shall talk about this later on, because it is so important, and ask you to suspend judgment until you have heard the arguments.

The second criticism is related, but is different none the less. It says that this whole approach to running a country presupposes a regime in which the state either owns industry or intervenes massively in its affairs. This is based on a misunderstanding; it confuses the machinery of government with government policy. I am not talking about that policy at all, nor taking a particular policy for granted. I am talking only about the machinery with which any modern state is equipped in the necessary task of government, and I am arguing that it ought to work. Every government regulates, every government controls—to some degree or other. In short, it governs. These arguments have been concerned with the “how” of governing, not with the extent of intervention.

The third criticism says: it is all too simple. Real life is much more complicated than this. The only honest reply to this is a belly laugh. What is proposed here is simple enough, I hope, to be understood—once

you know some of the laws of cybernetics; and if it is not simple enough to be understood, then of what conceivable use is it? Moreover, it is a very great deal more close to real life than the system we use now, because it is dynamic instead of static, systemic instead of a collection of entities, and because it really can handle variety according to scientific principles, which the existing system cannot.

Please appreciate that once you start to use systems thinking you need fewer data than before—because the data are synthesized within the model of the system. And this relates to the previous criticism, because I by no means envisage that a private business would supply more information than it does now: probably less. Why is it that a doctor does not have to take your body to pieces, and measure everything in it, when you arrive saying: “I feel rotten”? It is because the doctor holds a good model of the human body: when he knows your temperature, pulse rate, and blood pressure, has inspected your tongue and observed the rings round your eyes, he already knows most of the usual stories. By using good recursive models of the economy we should achieve the same effect.

Ironically enough, while this third criticism says it is all too simple, the fourth criticism says it is all too complicated—you could never do it, or it would take twenty years. Maybe you have noticed that I have been describing this approach today with a force and confidence which you may have thought brash in describing a vision. Indeed, until two years ago I was much more gentle and tentative—even apologetic—in making this description. But today I am no longer guessing, because all the basic work has actually been done.

Perhaps you remember my describing how these lectures were prepared in Chile. I was not actually on holiday there, as it may have sounded, rather the reverse. Two years ago President Allende invited me to create a system of this kind for the Chilean people. Now Chile is not a rich country, and we could not afford all the apparatus we should have liked. For instance, we had to use ordinary Telex instead of teleprocessing. Even so, we made it. We had everything I have just been describing available—though not yet running the economy—in eighteen months.


In the few months that remained to us, we were teaching the workers, for whom this offering of science to the people was created, how to use the most advanced tools yet designed for national economic management. They could sit with their ministers in the economic operations room in Santiago, watching the animated screens, and discussing the alerting signals provided daily by that clever computer program. They had buttons in the arms of their chairs, so that they could command the appearance on other screens of supporting data—to the capacity of 1,200 different colour presentations, focused through sixteen back-projectors. They could also control preliminary experiments in simulation, on a huge, animated model of the dynamic system. These people, arm in arm with *their* science, were intended to become the decision machine for the economy.

You all know what happened. On 11th September 1973, Salvador Allende died in a bloody business, of which the consequences for mankind are incalculable today. I tell you solemnly that in Chile the whole of humanity has taken a beating. Of the lessons from my own work that emerge, I mention four. Firstly, it is actually possible to redesign the institutions of government according to the principles and practice of cybernetics. These are not wild dreams. Secondly, there is a long way to go in dismantling bureaucracy, and I shall discuss the problems of effecting change later in these talks. Thirdly, the possibilities propose an urgent task for our next meeting: to discuss the impact of such scientific advance as this on the status and freedom of the individual.

So I move to the fourth and final point for today. Individual freedom has been lost, momentarily at least, in Chile. I think I know how; but it was certainly not because the people became victims of technocracy. What is clear is that everything that I have described was accomplished (and ended) in two years, and *it was not fast enough*. When I drafted these lectures, and outlined the hypothesis you heard—that perhaps

our institutions could not react fast enough to avoid catastrophic collapse—I remember thinking that I should have to defend myself against a charge of sounding a premature and too scare-mongering an alarm. Do you care to make that allegation now?

NOTES IN SUPPORT OF THE THIRD LECTURE



DIFFERENT systems

map onto


the SAME model

by sacrificing whatever variety
is not needed
for the purpose
in hand.

This mapping

indicates change without alteration.

1. Models of Nation by Departments of Government:



departmental models

map

(rather loosely)

onto national variety


excess variety generates as things

is caught by some new organization without a
proper structure.


Instead, the original departmental structures should
be redesigned, and made adaptive to further change.

2. Models of Enterprises Within the Economy:

The arbitrary choice of significant characteristics and of time epochs masks important variety.




3. Aggregation of Statistical Information




4. Time Lages in Registering Information:

Perception (dotted line) of the movement of economic indicators is not only “too late”—it means that at the time of taking a decision the most relevant variety generated by the already inadequate model is suppressed. We may be led to do precisely the wrong thing.



Four Levels of Recursion : One Model: One Computer Program



A full account of this approach to management cybernetics is in the author's *Brain of the Firm* (Allen Lane, The Penguin Press, 1972). An account of the Chilean application comes at the end of *Platform for Change* (John Wiley, pending publication).

4 Science in the Service of Man

The scene is a small airport at a vacation resort where a week-end conference on automation has been held for senior businessmen. Everyone is going home. The man in front of me at the desk finds to his horror that the flight for which he has a ticket is already full. “I am so sorry,” says the desk clerk soothingly. “We are having trouble with all the Flyaway Airline’s flights—something has gone wrong with their computer.” In the quiet of the little airport lounge, everyone is listening. A large man several places back in the queue leans forward and says loudly: “Excuse me, young man; I am a director of the Flyaway Airline. We do not *have* a computer.”

Some people laughed, but the general atmosphere turned to acute embarrassment. I do not think that this was because the desk clerk had been caught out in his white lie. As I argued in the second lecture, the computer has come to represent a malign influence, and something of which people are frightened; therefore people are most unlikely to call on its aid, or to demand its use in the redesign of society. But the embarrassment attending this incident also suggested that the great ju-ju none the less exists, and to deny that was a somehow dangerous act. . .

The title of this lecture is *Science in the Service of Man*, and I should be terribly surprised if its announcement caused so much as a blink of scepticism. A thorough-going job has been done in putting forward the slogan that science serves mankind, and the time arrives to examine this proposition with some care. This does not mean striking a balance between the benefits of nuclear power and the risks of atomic war, between germ control for better medicine and crippling the ecology with pesticides. Those debates raise apparently insoluble dilemmas for society because they are contradictory outputs of high-variety dynamic systems whose basic institutional machinery goes unexamined.

Science has been sold as the servant of man in the wrong way, and for the wrong reasons. Let us start with the source of wealth: production. Here science is put forward as the promoter of efficiency; and I do not have to trace for you the familiar story of the industrial revolution, of the assembly lines of mass production, of the resulting fall in unit cost; nor need I point triumphantly to the widespread availability of its products. But it *is* necessary to become conscientiously aware of the alienation that all this has induced in the industrial worker, so that social scientists are now engaged in frantic attempts to restore some sense of humanity to the working situation in which so many find themselves—a situation to which numberless millions in the Third World stand to be condemned in future, as their countries struggle for their share of the earth’s fixed assets. Something has gone wrong.

Turn secondly, then, to science as the servant of the consuming man. We have been sold labour-saving devices of every kind as the fruit of science and technology. The automobile got us out into the countryside—where we met everyone else in *his* automobile. The plethora of automatic devices in the house freed the housewife from drudgery at home—where she is the victim of their very efficiency, and of the steadily increasing difficulty (and steadily increasing cost) of having increasingly tawdry things put right. And in this same home, for all its apparent desirability, she often comes to feel trapped with her young children and the television set for company, an all-too-likely prey to mental states of anxiety or depression. Again, it would seem, something has gone wrong.

It is against this background at work and at home that science and technology are driven relentlessly forward towards a society of conspicuous consumption, since this is the only development that our economic machinery can countenance. I mean that growth is the order of the day. More goods must be brought within the range of more people, which can be done by more automation, more standardization, and lower unit costs—coupled unfortunately with an increasingly noticeable fall in intrinsic worth. Therefore people must be persuaded that this is what they want, that this constitutes the good life, and this is science in the

service of man. But I believe that the society of conspicuous consumption is proving to be the most alienating force the world has ever known, and that the fantastic consumption of drugs (both legally prescribed and illegally acquired) is a useful index to the degree of alienation now in evidence. By now, something has gone very wrong indeed.

What does this brief analysis purport to show? It argues that the sense in which people accept that science serves man is a false sense, since science is in these typical ways being used to destroy man—in his humanity and in his joy of living. Moreover it is getting through to decent people that on a planet the resources of which are only now becoming recognized as finite, prosperity for all is a delusory goal. We buy increasing prosperity for we few at the expense of the many who can never attain it. As the alienation grows, there is increasing resistance to the idea of yet more science, with the result that new proposals for handling old problems by the use of computers and telecommunications are often greeted with something approaching public hysteria. I am thinking of electronic files on the citizen, or the kind of governmental control system that I described earlier. The point is that this panic is well justified, so long as society continues down the existing path, following its technological nose. Yet if societal institutions are to escape the fate of catastrophic instability, we shall very certainly need new systems of these kinds.

It follows that science has to be handled in a new way. There is only one solution that I can see. It is to remove the control of science and technology from those who alone can finance its development, and to vest its control in the people. As to scientists and technologists themselves, they are truly servants of that public; whereas the present tendency is to turn them into an elite instrument of those who have the economic power over scientific systems themselves. That way lies technocracy, and we are perilously close to it already.

How realistic can this solution possibly be? After all, people who have power simply never hand it over to others; moreover, in this case vast sums of money are involved. I reply that the solution is realistic in a democratic society to the extent that the demand to redesign societal institutions is made articulate. The process can begin by debunking the mystery surrounding scientific work. It would make a very good basic postulate for the ordinary citizen to say something like this to himself, and to discuss it with others:

“For the first time in the history of man science can do whatever can be exactly specified. Then, also for the first time, we do not have to be scientists to understand what can be done. It follows that we are no longer at the mercy of a technocracy which alone can tell us what to do, Our job is to start specifying.”

For this new channels are needed. But of course they could be set up. What is television for? Is it really a graveyard for dead movies; or animated wallpaper for stopping the processes of thought? What is the computer for? Is it really a machine for making silly mistakes at incredible expense? What will be done with cybernetics, the science of effective organization? Should we all stand by complaining, and wait for someone malevolent to take it over and enslave us? An electronic mafia lurks around that corner. These things are all instruments waiting to be used in creating a new and free society. It is time to use them.

Then as to cost: who will pay the bills? We do that already, since it is taxation and inflation that finances the schemes of governments, and a loading on the prices we pay that finances the schemes of corporations public and private. But as far as I can see, the citizens have lost control *entirely* of the choice of projects that will be undertaken on their behalf both as taxpayers and as consumers. At best, they play a defensive role in attempting to quash schemes they dislike; and that is a difficult role because it does not carry requisite variety with it. Anyone who has had dealings with a public enquiry knows only too well that the bureaucracy has the power to amplify its variety indefinitely—in terms of the time, money, and expert advice it is free to deploy against a little band of citizens who do not have access to these amplifiers.

Obviously, I am trying to dig beneath the surface layer of science and technology as we know it in society, to uncover new strata of scientific potentiality. The societary use of science we have is threatening; it becomes oppressive and alienating. The societary use of science we could have is a *liberation*. To grapple with that idea, I well understand, needs courage and resolve; the risk is that folk who see the very real dangers will turn their backs on the whole difficult business. But how safe would our great-grandparents have felt if plunged suddenly into a modern home, a modern street?

We have had three or four generations in which to adapt to a house alive with lethal electricity, a road alive with lethal trucks. We have had barely twenty years to adapt to the inventions and discoveries that these lectures discuss. Then no wonder the adaptation is not coming along too well yet. No wonder people feel at ease with an automobile that they literally dominate, and ill-at-ease with the computer which they do not. The interesting thing is that a majority (perhaps) of automobile-dominators do not understand exactly how those machines of theirs work, and yet use a similar ignorance of the computer's viscera to explain their distrust of it.

Be that as it may, the problem of rapid adaptation for the individual that has now emerged is a similar pattern—in cybernetic terms—to the problem of rapid adaptation for the institution. Let us try to analyse the modern individual's problems in the language that we have been learning, because this problem is indeed a problem of effective organization.

The first thing we have to face up to is quite a tough proposition for people reared in our culture. It is that whatever we humans can do is mediated by our brains, and those brains are finite. We have in the cranium a slightly alkaline three-pound electrochemical computer running on glucose at about 25 watts. This computer contains some ten thousand million (that's ten to the ten) logical elements called neurons, operating on a basic scanning rhythm of ten cycles per second. Then this is a high-variety dynamic system all right; but it really *is* finite. It follows from Ashby's Law that we can recognize patterns up to a certain limit, and not beyond. Thus if something is going on that involves a higher variety than the brain commands, we shall not recognize what it is. This is the old constraint of requisite variety again.

There are practical consequences to this. For instance, I am sure that the reason why we are making such a hash of the problems of global ecology is that we cannot understand them. I don't just mean that they are awfully difficult, so that understanding will take a lot of research. I mean that we can *not* understand at all, ever. Very likely this goes for many problems of government too, especially world government. It may even be true at the level of recursion where a corporation is managed. May I recall that the level of recursion is simply the focus of attention at which we contemplate any viable system, and that one level is contained within the next. So here is an unpleasing thought: maybe it is also true at our personal level of recursion. Perhaps we cannot actually understand our own lives, our own environment, any longer.

Now with or without full understanding, with or without the requisite variety to detect vital patterns, we have to cope somehow at all these levels. Of course we do it by making mental models. We simplify, so that the system we are considering will map onto our own brains. But that can be done only by attenuating variety, and we have no guarantee that we are not throwing the wrong information away. It is fairly evident that we shall become accustomed to discarding information in set ways, and to eliminating inputs that do not seem to fit very well the models we have developed. I think this must mean that what we all refer to as "reality" is a version of the universe that is very much cut off at the knees. To be rude about it, you could say that our humanity exists in sharing a delusion about the way things are.

At this point I would love to start talking about mysticism, or about psychosis, or about psychedelic drugs—especially I would like to talk about the relations between them. Because these three things have this much in common: they claim to deal with aspects of reality which our shared delusion filters out. But the point I was really after is this. The currently explosive rate of change produces perturbations at intervals

that are shorter than the relaxation time of our institutional system: that was my earlier hypothesis. I have pointed to its realization in Chile. I now extend that hypothesis to cover you and me as individuals. Can it be, perhaps, that we all suffer from a variety overload that we cannot map onto our models, and from an ungovernable oscillation in our search for mental equilibrium? In short, is our species facing the same threat of catastrophic instability as I earlier argued that our institutions are?

It could be so. One of the greatest biologists of our time, Albert Szent-Györgyi (who discovered vitamin C), believes it—and not at all from the reasoning I am using here. He calls man “the crazy ape”, and reckons that we have all gone down an evolutionary blind alley. Well, I am by nature an optimist. But I do believe in facing facts. If you suspect that my advocacy in these lectures of almost revolutionary change in our approach to running affairs is too extreme; if you consider that I have overstated the failure of our institutions; if you think that the remedies I am proposing are more dangerous to human freedom than the disease itself; then please wonder about these new and serious thoughts.

They concern the variety engineering to which our culture subjects our personal input, because for me this is what freedom is all about. I am tired of being told that the computer threatens our freedom, that cybernetics is a tool of the devil, that real-time governmental regulators are too dangerous to employ. The reason is that I reckon our existing liberty to be largely illusory: we are fooling ourselves. There is a new chance, now, to get our freedom back, even for us to bestow freedom on those who are following in our disastrous paths.

Please look at it this way. We all know that a majority of people on this planet are enslaved. I mean this in the straightforwardly physical sense. Most people alive do not have enough to eat, and must live under regimes that tell them what to do. By such standards all of us who share these broadcasts are incomparably better off, more free. The fact remains that our own relationship with our environment is governed by bank upon bank of variety attenuators, conveniently reducing a world of increasing variety to the requisite variety of our brains. *We have completely lost control of the processes by which this occurs.* Here are the two most evident examples.

The first is education. Every pupil is a high-variety organism, and the process of education essentially constrains variety. In other words, the pupil is capable of generating many responses to the question: what is six multiplied by seven; the educator will seek to attenuate this potential variety to the single answer: forty-two. But if we take a different kind of example, we may find ourselves saying something significantly different. The pupil is capable of generating many responses to the question: how should a national health service be organized? This time, however, we may hope that the educator will *not* attenuate potential variety to the single answer: like *this*. No, we say; education is a word coming from the Latin: *e-ducere*, “to lead out”. It does not mean “to push in”. And yet it remains true that in any case the process of education constrains variety. Anyone who thinks over this little paradox for a few minutes can see that, in the second example, the hope is that we can teach the pupil ways of attenuating *his own* variety. We want to offer him ways of finding answers, not of enforcing our own.

All this is perfectly obvious, you could rightly say. But do we proceed on the basis of this distinction? Of course not. Off we go once again, inserting the amplifiers and attenuators on the wrong sides of the equations. Take the question I mentioned: How should a national health service be organized? What happens? People are sent from under-developed countries to study this matter in over-developed countries—where in truth they may well discover how a health service should *not* be organized, and be robbed by the experience of the power to devise better answers. Take the question: How shall we train the work-people to engage in fully participative management, so that the outlook of the shop-floor can be represented in the board-room? Can you believe that the answer I hear advocated to this is that selected workers should be sent to business schools? The variety equation is overturned once again: here is a recipe to ensure that

on arrival in the board-room the worker's participation will be indistinguishable from that of the other directors. Take the question: What is the theory of relativity? Because of the cultural myths that all technology is dehumanizing, and that all great minds are incomprehensible, we shall prefer to sit a hundred pupils uncomfortably in front of a human teacher who hopes he understands relativity, and who roughed out some notes last night, rather than to give the individual pupil access to videotape recordings which he can replay to his hearts content, of Albert Einstein—who could be as lucid as the day.

Oh yes, I know that Einstein is dead, and did not even make those recordings. What I should like to know is how many Einsteins shall we let die, and how many theories of relativity shall we let go improperly explained, before we recognize where to use our variety amplifiers and attenuators with good sense?

The problem is nowhere clearer in the field of education than in the evolving use of the computer. Here, once again, the machine could be used as a real *liberator*. It is an instrument of colossal variety, to which each pupil could have ready access. Thanks to parallel processing, a computer can be interrogated, explored, *used*, continuously and in different ways by a few hundred pupils at once—it has requisite variety. So what happens? The variety is attenuated out of the computer, by making it operate trivial little programs that actually *condition* the pupil to give the right (in quotes) answers to a set of trivial questions.

Which brings me to the second example, namely publishing. If education begins the process of constraining our cerebral variety, publishing (whether on paper or by radio waves) continues it for ever. The editorial decision is the biggest variety attenuator that our culture knows. Then the cybernetic answer is to turn over the editorial function to the individual, which may be done by a combination of computer controlled search procedures of recorded information made accessible by telecommunications. Cable television has all the potential answers because it can command eighty channels. This offers enough capacity to circulate the requisite variety for an entirely personalized educational system, in which the subscriber would be in absolute command of his own development.

Well, we are frightened of this projection too. Someone may get inside the works, we say, and start conditioning us. Maybe we should have eighty alternative standard channels, thereby “restoring choice to the people”. Here is my third and last bit of mathematics: eighty times nothing is nothing. Meanwhile, we allow publishers to file away electronically masses of information about ourselves—who we are, what are our interests—and to tie that in with mail order schemes, credit systems, and advertising campaigns that line us all up like a row of ducks to be picked off in the interests of conspicuous consumption. I know which prospect frightens me the more.

As I try to pull together the threads of this, the most diffuse of the lectures, I would like to remind you that I have been talking about societal science and the individual—you and me—and not proposing societal solutions. I hope that some solutions will begin to appear in the final lecture, which would be the proper place for them. However, we cannot even think about finding solutions without correctly recognizing the problems, and in my opinion it is now commonplace to pose these problems incorrectly.

Certainly it is common ground that scientific techniques threaten the individual in society; and I have not used up time in talking about the obvious ones to which all governments are addressing themselves: the problems about privacy, credit-worthiness, and blackmail, problems about urban planning, and so on. Choices are available in these matters, and we may pray that the best ones will be made. The argument of this lecture is very different; it has to do with the cybernetics of our own brains, and I shall now put it together. I am sorry that it assaults what we all think of as the human prerogative, but it is best to face the truth.

The brain is a finite instrument that mediates all our experience. It has high variety, but not necessarily requisite variety for handling an environment of exploding complexity. It has a relaxation time that was

fast enough to deal with a world in which perturbations came at a particular rate, but it is not necessarily fast enough to offer a guarantee of equilibril response in the current world.

This brain has certain powers, and these are essentially computational, which make it the most developed regulatory system the world knows. But my cybernetic interpretation of the evidence from biology, psychiatry, pharmacology, and criminology, is that this brain (and again that just means you and I) is by now seriously threatened by a possible catastrophic instability. Finally, this brain simply does not have the powers of untold resilience or infinite self-improvement to which three thousand years of pre-scientific culture have laid a spurious claim.

Please now hear me when I add that these considerations make no commentary whatsoever on matters that may (or may not) lie outside the physical domain. If mankind can indeed receive the divine afflatus, the point remains as I rather carefully put it just now—that the brain is a finite instrument that mediates all our experience, and is therefore limiting. As a personal aside, let me say that I am more interested in the fact that I could not recognize an angel if I met one, because my brain does not have requisite variety, than I am in the illegitimate scientific argument that angels do not exist because I have not recognized one yet.

Returning to the main argument about the limitations of the brain, I have argued that we as individuals are the unwitting victims of a cultural process which very drastically delimits variety for us. In the first place our economic environment points to an increasing use of science and technology in what is allegedly the service of man—but which I contend takes this service in a false sense. As a result, we stand, and the innocent legatees of our policies in the developing nations yet more vulnerably stand, to be exploited by whoever wields the power of science to technocratic ends. In the second place, the instruments of variety constraint turn out to be education and the communications media—both of which we culturally suppose to be variety amplifiers. This belief is as delusory as the belief that we can fully know reality.

It is entirely possible to take corrective action about all this—not the biological limitations, but the societal constraints. To do so requires that people themselves take control of the use of science, through their democratic processes. This means furnishing them and their governments with new channels of communication, and a new kind of educational system, and a new kind of publishing system.

Why are these recommendations necessary? The answer is that the necessary attenuation of variety produces in us a mere model of the world. And insofar as we wish to control the world, whether as citizens or as individuals within a personal environment, our powers of regulation are cybernetically constrained by the model we hold of what needs to be regulated. Our civilization has led us to a manifestly dysfunctional model. Then we must equip ourselves to revise it. The power to do this we certainly do possess.

NOTES IN SUPPORT OF THE FOURTH LECTURE

A Key Statement

For the first time in the history of man,
science can do
whatever can be exactly specified.

Then, also for the first time,
we do not have to be scientists
to understand what can be done.

It follows that we are no longer
at the mercy of a technocracy
which alone can tell us what to do.

Our job is to start specifying.

Some Counter-Intuitive Cybernetic Statements

Education constrains variety, because (although it may open new vistas) it leads us to reduce the alternatives that we are prepared to entertain.

Publishing (“the media”) constrains variety, because (although it too may open new vistas) it decides in practice what alternatives shall be accessible to us to entertain.

Technology is not of itself dehumanizing, though technocrats may make it so.

Great Minds are not incomprehensible, because the brains behind them have ordinary human variety, but many institutions have a vested interest in concealing this.


The Brain and Requisite Variety

WE CAN

recognize each other

do mathematics

handle perhaps three quantified
variables in our heads



CAN WE?

recognize ecological systems?

recognize an angel?

undertake world governments?

The Brain Is Finite:

The brain is a slightly alkaline three-pound electrochemical computer, running on glucose at 25 watts. It contains 10^{10} (that is, ten thousand million) nerve cells or neurons, which are its computing elements, operating on a scanning rhythm of ten cycles per second.


Why 25 watts?

About a quart of blood (about a seventh of the circulation) flows through the brain every minute, and is heated about one degree Fahrenheit. That takes about 25 watts.


Why ten cycles per second?

The resting state of the brain, called the alpha state, is characterized by the alpha rhythm, which tends to be stable at around ten cycles.


The Risk of an Electronic Mafia



buy these publications,



mail orders for these,



give this "public opinion"



watch these shows,




use this credit card?




and this?



to buy this



and is therefore exactly like this with... certain susceptibilities...



Good Morning!



Your new life style!



New politics!




New religion!



How can I ever thank you?

A Better Outlook for Cybernetic Man



- Tell me a whole lot more about news item six.
 - Show me the family tree of King Priam.
 - Get Back the work I was doing on my new house.
 - Entropy? I don't understand. Stop all this and explain.
 - Record in my tax file that today I bought a dog.
 - Tell Jack in Vancouver I'm ready for the chess match
 - Beethoven's Fifth Symphony, the composer conducting.
- oh I see.

AND DON'T TELL ANYONE ELSE UNLESS I SAY SO.

Book (not related to the above) mentioned in the lecture: Albert Szent-Györgyi, *The Crazy Ape*, Philosophical Library, New York, 1970.

5 The Future That Can Be Demanded Now

A famous summit conference, not to say confrontation, about freedom once occurred about twenty miles from my home in England. It lasted from 15th to 23rd June, in the year 1215. During those negotiations between King John and his barons, the Magna Carta was signed: a document that spoke for all time of the decentralization of power and of the rights of individuals, and is still much quoted 750 years later. I remember from my boyhood a humorous monologue explaining these events, which ended something like this:

So it's thanks to that Magna Carta
that was signed with the barons of old
that in England today we can do as we like
—so long as we do as we're told.

The blatant contradiction embodied in this joke remains the desperate problem that it always has been. How do we sustain individual liberty and societal cohesion at the same time? It is right that this problem should be incessantly discussed, and it is discussed. But the discussion always seems to lead straight into the same disastrous trap; a false dichotomy between the notions of centralization and decentralization. The vehemence with which this matter is debated is extraordinary, because the most cursory consideration of what constitutes a viable system reveals how false the dichotomy must be.

For example: if you personally were a fully centralized system, you would need to remember to tell your heart to beat. If you paid too much attention to what I am saying, and forgot, you would collapse on the floor. That would be dramatic, but it is not what I am trying to achieve with you. But if you were a fully decentralized system, you would trot off from this broadcast to investigate any sound. Neither solution would leave you a viable system for long. Let us analyse the blend of central and peripheral command and see what happens.

We discovered earlier that viable systems are bombarded continuously with high-variety stimuli, the variety of which has to be attenuated if the system is not to be overloaded. The attenuation must be done according to a pattern, if it is not merely an arbitrary discard. If that pattern is to have survival-value (which is a necessity for a viable system) then it must be a regulatory model of whatever is regulated. Then it follows that this has to be a *central function* for the system, because only the system as a whole can have a model of its own relationship with its own environment.

Nonetheless, when the central function of variety attenuation is operating, it is by definition not appropriating to itself the discarded variety. But by Ashby's Law we know that variety must be absorbed somewhere. Therefore whatever variety is not appropriated by the central function must needs be handled by a decentralized function. This variety handling is by definition autonomous. That is to say that some part of any viable system does what it likes. But of course the autonomous part of the system remains part of the system, and to do that it must take notice of the central regulatory model. To that extent, then, it does what it is told.

If we make a terrible mess of interpreting these simple cybernetic discoveries in our society, and I feel that we do, it is because there is no agreed machinery for settling clearly which parts of the system are which. To do so is indeed virtually impossible unless the models in question are made explicit for each level of recursion. To remind ourselves: a level of recursion is a level at which a viable system is in operation, as an autonomous part of a higher-level viable system, and containing within itself parts which are themselves autonomous viable systems. We spoke of this set-up before as being like a set of Chinese boxes.

We find the process of settling autonomy going on continuously within any viable family, for example. As the children grow older, they exert more and more their personal freedom of action. But this has to fit into the family's general regulatory model, at a higher level of recursion. Thus a great deal of time will be spent in discussing the notion of autonomy for the younger members, and the time required *has* to be a great deal because requisite variety must be attained. In the upshot, families often manage to preserve their model of the family, which is centralizing, and the liberty of the younger people, which is decentralizing—and forget altogether to consider the right of freedom for the parents, who then become identified with the centralizing authority. This, you will note, is not at all to the parents' advantage (since they have lost freedom in the process) although it may appear to be so to the children.

This homely example is repeated, with great force, in the yet higher levels of societal recursion. A well-intentioned corporation or national service tries to hang on to its systemic policy—because this is what makes it itself, this is what embodies its aims, this structures its regulatory model. But that policy is centralizing. Then the corporation and the national service, being well-intentioned, embark on high-variety negotiations with the parts of the system, in order to delineate autonomy—which they really wish to be maximal. But the tools they are using are not cybernetic tools, that undertake variety engineering, but administrative tools which do not. As we have several times noticed, bureaucracies install amplifiers in one loop of the homeostat when they should be installing attenuators in the other and vice versa.

The upshot now is really quite strange, but nonetheless extremely common. The parts of the system—subsidiary companies, sectors of the national service—who have been in good conscience given maximal autonomy, *believe* that they have been totally centralized. This is because their variety is attenuated by wrongly installed central amplifiers. But, on the other hand, the higher management at the centre, in conceding in good conscience maximal autonomy, *believes* that it has somehow been robbed of any role at all. This is because its variety is over-attenuated by wrongly installed peripheral attenuators. In the family, given quite a lot of hard work, everyone can just about hang on to happiness. But in big institutions, where—we have to remember—the brains of all the men and women involved are still the same size as they are in family roles, disenchantment spreads.

I often reflect that our organizations are so constructed in their typically pyramidal shapes, so that they could work only if the people in them grew bigger heads as they became more senior. In that case, of course, there would at least be a chance that they could maintain requisite variety. However, as we know, everyone's head is roughly the same size—except, perhaps, metaphorically. I also reflect upon the device whereby “decentralization” is often advocated as the solution for an institution's problems, when it is held to be over-centralized, while “centralization” is simultaneously prescribed for another institution on the grounds that it is too diffuse. And I have seen the two policies advocated alternately, and what is more alternately adopted for the *same* institution by successive groups of consultants. It is a kind of managerial madness. It leads, as it can only lead, to exacerbated oscillations in the system's search for stability.

The solutions lie, they can only lie, in good variety engineering—and here is the key point. We must not confuse the *pattern* of the regulatory model with its specific *content*. It is enough to attain requisite variety by specifying the pattern. To specify the content is too much. Yet this is what endlessly happens, and I have noted that it usually happens—in those well-intentioned institutions, in good conscience—for one fundamental reason. This is called “*fairness*”. But I believe this kind of fairness to be an excuse for avoiding responsibility.

Take the example of a big institution that has a salary policy, or an employee automobile policy, or an inventory policy. We need a way of saying what the policy—which is to say the variety-attenuating regulator model, or pattern—really is without specifying its specific content. We fail. We ought to say: this much can be spent on salaries, on automobiles, on inventories, and leave it to managers at lower levels of recursion

to apportion the money. Instead we do the variety engineering in the wrong place. “*This* is the salary scale”; “You are entitled by your job to *this* range of automobiles”; “*All* inventories must be cut by ten per cent”. It is done, as I said, in the name of fairness, but that is delusory.

It is nonsense to say that two men of the same age, with identical qualifications, with identical commitments, are necessarily worth the same wage. Obviously not. One may be useless, and the other a paragon. It is nonsense to say that my job ought somehow to determine my needs where automobiles are concerned. How and where I live, how many children I must push into the automobile—these are my own affair. It is nonsense to penalize a good manager who works on a scientifically calculated minimal inventory because his colleague managers are inefficiently tying up the firm’s capital. This is a recipe to encourage inefficient managers. Why do we blandly accept so much nonsense?

The variety attenuators to use here are not policy documents from the centre, but the managers themselves. That is what managers are for. As to the criterion of fairness, the manager—or any individual, in whatever he does—ought to be ready to take responsibility for his own decisions. Our society militates against that morality (for that is what it is) with all its force—in the name of an efficiency which is thoroughly bad cybernetics, in the name of a fairness that is manifestly *unfair*. But please remember the precept: each of us should take responsibility for his own acts. The practice is precisely the contrary .

As usual, then, we have the amplifiers and attenuators on the wrong side of the equation. But all of this is now written into our culture; all of this is now underwritten by our bureaucratic formulae. That is why I have repeatedly argued that fundamental change in our modes of organization is essential. Merely to juggle with existing forms simply increases the swing of the oscillating pendulum that never can find its stable state. And, as I have mentioned before, this means that the system is robbed of the crucial reference point without which it cannot learn, cannot adapt, cannot evolve.

How do we set about making so fundamental a change as would bring our ways of working into line with the scientific rules of this game? To answer this question it is necessary to understand the nature of resistance to change. Here is a phrase that is on everybody’s lips—“there is a resistance to change”—but it is a phrase which is not analysed according to the principles of effective organization. People seem to imagine that they are confronted with a psychological hang-up, insofar as men and women are not supposed to like change.

But pause one moment, is that true? People, considered as individuals, it seems to me, like change rather a lot. Don’t you get bored when nothing changes? I know I do. Then just why do we go around saying that there is a resistance to change? Of course, the answer is simple. It is not the living, breathing human who resists change in his very soul. The problem is that the *institutions* in which we humans have our stake resist change. Therefore we feel as individuals that we cannot afford to embrace it. And this is an extremely sound argument.

If you have spent a lifetime working your way up a ladder, you literally cannot afford to be robbed of the prizes when your turn comes to collect them. One of my earliest experiences in industry was to listen to managers explaining to senior operatives that they were to be deprived of their life-times’ ambitions, because the whole technology of the process was to be changed. That was in the steel industry, in which the skills of a first-hand melter—a job that it took a lifetime to learn—were replaced in a year by clever instrumentation. But ten years later those same managers were themselves confronted by the computer, which would have made many of *their* skills redundant in turn. However managers have power, and the computer none. It was easier to misuse the computer than to accept the institutional change—because the consequences would have been quite personal.

Now I have come to what I consider to be the explanation of the abuse of science and technology in our society. The power has remained where it resided. The tools of modern men have been disregarded at *this* level of recursion. And there is none left to say a loud NO to that—until the people themselves say No. So this is why I contend that we are considering a future that can be demanded now.

Every time we hear that a possible solution simply cannot be done, we may be sure on general scientific grounds that it can. Every time we hear that a solution is not economic, we ought to ask: “for whom?”—since it is people, just people, who will have to pay. Every time we hear that a proposal will destroy society as we know it, we should have the courage to say: “Thank God; at last.” And whenever we hear that it will destroy our freedom, we should be very cautious indeed. For such freedom as we have is our most treasured possession, and we know how to be vigilant. Yet for that very reason, this is the simplest method that the powerful have to cling to power: to convince people that any other concession of that power would be unsafe.

But I would like to stop philosophizing to you about all this, at the expense of introducing another technical term. We have not had a new one for a couple of lectures, and I hope that the very few that I have introduced make a useful vocabulary. We are by now used to the notion that institutions are not just entities, with certain characteristics. They are instead dynamic viable systems, and their characteristics are in fact outputs of their organizational behaviour. The variety that is pumped into them is absorbed by regulating variety, through an arrangement of amplifiers and attenuators. A system that, through this kind of exercise in requisite variety, achieves stability against *all* perturbations, is called a homeostat.

A homeostat can resist perturbation, not only against expected disruption, but against unexpected disruption too. For this reason it is not only stable, but *ultrastable*. Whatever happens to it—provided that its relaxation time is sufficiently short—it will not go into oscillation, and still less will it explode in catastrophic instability. The sign of this homeostasis, now so deficient in our major institutions (and perhaps, as I said earlier, even in ourselves) is that critical outputs of the system are held steady.

Why produce this extra terminology at this late stage? It is because I want to answer the question about resistance to change in a very precise way. All homeostatic systems hold a critical output at a steady level. But some of them have a very special extra feature. It is that the output they hold steady is *their own organization*. Hence every response that they make, every adaptation that they embody in themselves, and every evolutionary manoeuvre that they spawn, is directed to survival. So this special trick rather well defines the nature of life itself. It also rather well explains why we cannot change our institutions very easily. Their systemic organization is directed, not primarily to our welfare, but to their own survival.

At this point we need to draw a very careful distinction. Institutions are supposed to be homeostatic. They have been driven away from this behaviour by getting their relaxation times out of phase with explosive change. That was the argument. Yet buried inside the institution is a nucleus which retains its homeostasis by ignoring not only external change but the primary function of the institution itself. This nucleus is the special kind of homeostat that produces itself. And it is this nucleus that I call the bureaucracy. By this term I am not simply referring to paper-pushing, but to an institution within the institution that exists—narcissus-like—in self-regard.

Bureaucrats assimilate the challenge of explosive change in an essentially bogus fashion. For many years I used to make the joke that they would accept every kind of change, provided it involved no actual alteration. Indeed, I used that very phrase (but not jocularly) in the third lecture. In any case, the joke rebounded. Bureaucracies *do* accept change; they do this by acknowledging novel conditions. They are not so stupid as to pretend that such conditions have not occurred. But what changes they make are superficial, and they are made so that organization—which is what makes the system the system that it is and no other—is completely preserved. Thus there is no actual alteration, although appearances may have

been changed a great deal. When this is generally understood, it will no longer be possible to fob people off with unreal changes, masquerading as real alterations. Until then, our institutions will go on producing the social benefits of their activities simply as by-products of their major bureaucratic undertaking, which is to produce themselves.

According to the analysis of centralization and decentralization with which we began, it is clear that there should be a major devolution of power. I think it should be open to a community to organize its social services (education, health, welfare) exactly as it pleases, and to accept or reject the initiatives of local innovators. I think that goes for the local branches of national undertakings, public and private, also. I think that workers should in general be free to organize their own work, and that students (up to the age of death) should be free to organize their own studies.

The first barrier to doing any of these things is the absence of money. It is always assumed that because everything has been centralized, and because the centre makes facilities available, then a community wishing to do “its own thing” is opting out of the official plan. It may (or may not) grudgingly be allowed to do so, but it will have to finance itself. Yet the finances that have been raised at the centre are raised on behalf of the parts. It is perfectly clear that this is a monstrous infringement of liberty. The question however is not so much “how do they get away with it?”—because I continue with the hypothesis that everyone is well intentioned. The real question is “How did the system degenerate to this unviable form?”

We have done enough cybernetic thinking about dynamic viable systems by now to draw the distinction that is required to answer this question. In order to maintain viability, the total system must have a central regulatory model. This model ought to be created by democratic consultation, but we cannot dodge the truth that it will constrain variety in the parts. Put neutrally like that, all remains well: for it is essential that variety be attenuated in any case, and it is essential that services which become economic only at the total-system recursive level should be made available. But the vital distinction comes here. The precise form of variety attenuation is a matter for local decision.

The critical mistake we are making is to take the variety attenuating decisions—at the wrong level of recursion. Then this is how freedom is lost, and this is what induces the instability that threatens to become catastrophic. For the whole-system model simply does not have requisite variety to balance the local homeostats. They in their turn are robbed of the variety they need to find their own stable points.

With this insight it is possible to redesign any of our institutional systems. In my own experience of trying to do this, two major barriers to progress always appear. The first is that of bureaucracy. Bureaucracies build around any centralized system in order to administer its centrality. In decentralizing, the need for the bureaucracy disappears. But we are already in the trap. Bureaucracies exist, and are powerful: this is obvious enough. What is less obvious is the argument I used just now: they have *themselves* become viable systems that produce themselves. Now, a parasitic growth depends on its host’s continued existence it is true; but the parasite may flourish at the host’s expense.

There are two lines of approach to evaluating the facts, and you may try both on any institution you like. The first is to count heads. How many people employed are, by these definitions, of the bureaucracy? Official statistics are not collected in a form which will answer that question, either nationally or within corporations. One has to do the measuring oneself, and often it is necessary to divide an individual’s use of time between his work for the host and his work for the parasite. There is a lot of room here for self-deception. The answer is often as high as one-third, and rising. Then we need not be surprised that the bureaucracy has taken on a life of its own. The second line of approach is to examine the forms in which the bureaucracy produces itself.

For example, how much of the concern shown by the bureaucracy inside health, education, and welfare is for patients, students, and the deprived, and how much for the cybernetic machinery by which the medical, teaching, and social professions produce themselves? People are entitled to ask those questions. In asking them, they should realize that the need to maintain standards is a serious need—and also an impregnable excuse.

People are further entitled to ask whether there are not other ways of maintaining standards than by having bureaucracies. Of course there are. The trouble is that they would rely heavily on responsibility—on responsibility for one's own acts, on responsibility for one's own colleagues. As we have already seen, this morality is not favoured; it is "unfair". Besides it leaves people unprotected. Far better, then, to have a bureaucracy, which is amoral, than to depend on the morality of real human beings. How do you find yourself reacting to that? For me, it is indefensible.

Then our analysis leads to this clear conclusion. If institutions really are to be changed, then their fundamental organization really must be altered; and a major component of that alteration will be to dismantle the bureaucracy. Changes which do not dismantle the bureaucracy are unreal; they do not lead to alteration, but to the adaptation of the bureaucracy in continuing to produce itself.

You may notice that I have always said that a system of this kind produces itself, and not that it reproduces itself, which sounds much more natural. The reproduction follows, which is bad, but it is not the central point. To have the aim to produce itself is the mark of a system that cannot be dismantled without a death. When we funk the assassination, we may yet ensure much change—but the system lives on unaltered. So it has become with our societal institutions.

A little time ago I said that there were two barriers to progress, and that the first is bureaucracy. The second is the availability of money. But I have dealt with this question before, and need only summarize my answers now. Essentially, the costs associated with major projects are unreal.

- Point One: they usually represent not the actual costs, but the availability of funds.
- Point Two: the availability of funds is divided into arbitrary time epochs, which match the requirements of accountancy and not the needs of the people.
- Point Three: the people are paying for the projects anyway, one way or another, but this fact is disguised from them.
- Point Four: there is as yet no way in which the people can decide on which projects their money should be spent.
- Point Five: there is no reason why spending money according to the wishes of the people should cost more than to spend it according to the wishes of the bureaucracy, provided that the central regulatory model has been democratically composed, and is properly understood.
- Point Six: and this is new: the cost of many new societal projects could be met from savings made by dismantling the bureaucracy.

So: I am hoping that we may approach the final lecture of this series in the following state of mind. The human being is limited by his finite brain from assimilating all possible information, and from recognizing all possible patterns of the world. He is limited by his own finite resources from doing whatever he likes, and by the finite resources of the planet from demanding an endless growth in material prosperity, for all men. Indeed the pursuit of his own material prosperity, though possible, is not something that the affluent

part of the world can any longer maintain as a good, unless it is explicitly willing to declare that it will be done at the expense of the less fortunate.

Then the concept of freedom is not meaningful for any person except within measurable variety constraints: and the extent to which we have lost freedom is due to loss of control over the variety attenuators—education, publishing—and to the centralization of power at the wrong levels of recursion. This freedom could be reclaimed, using the new scientific tools at our disposal, but only if new democratic machinery is established to replace existing bureaucracies. As long as these remain cybernetically organized so as to produce themselves, our societal institutions remain set on courses that lead to catastrophic instability.

NOTES IN SUPPORT OF THE FIFTH LECTURE



For a viable system, the classic dichotomy between centralization and decentralization draws a useless distinction. Viable systems, which include institutions, require a complicated blend of central and peripheral regulatory information.

There is only one kind of information; but


- the way it is gathered and collated,
- the channels it uses,
- the filters it passes through,
- how much is remembered,
- the levels at which it provokes action,

—all these are distinct.

Studies in Neurocybernetics are especially helpful to the management scientist in designing institutional systems.

Some systems adapt to change in a **stable** way

—but only to such changes as they were actually *designed* to accept.




Most engineering systems are of the kind; and if a change affects them that the designer did not foresee, they succumb.



A **homeostat** holds its critical variables within the limits permitted by its own physiology, whether the form of the change affecting its performance was foreseen or not. Homeostasis is a natural property of living systems, which has been copied by cyberneticians to make institutional systems **ultrastable**.



Here is a drawing of an institution:



This drawing looks like an organism, living in its environment, and interacting with it. We may hope that this is a homeostatic system.

The circle labelled B *looks* like the nucleus of that cell, which ought to be directing these homeostatic operations. But it is not.

B stands for Bureaucracy, and that is self-regarding.

Bureaucracy is not concerned, as it should be, and as the cell nucleus would be, with directing the homeostatic operations, but with **producing itself**.

Because the institution is a viable system, it is at *this* level of recursion that the organism should be producing itself. How can it, if its nucleus is dysfunctional?

The new word for a homeostatic system whose major concern is to produce itself (that is, to hold constant its own organization) has been given by the man who first penetrated this phenomenon. He is Humberto Maturana, and he calls this kind of system *autopoietic*—which is the pure Greek for “making itself”. His definitive book is not yet published.

6 The Free Man in the Cybernetic World

The continuous process of liberating our minds from the programs implanted in our brains is a prerequisite of personal evolution. We can embark on that process of liberation only by constantly and consciously testing the ways in which our personal variety has been and is being constrained by the very things we tend to hold most dear.

But freedom is not pure anarchy. We are not free if we are dumped in the middle of the Sahara desert, despite the absence of walls and bars on the non-existent windows. We are free when the doors of our intellectual suite of rooms are unlocked, and we walk outside to breathe some new and fresher air. But we still need maps.

The message of hope I have to deliver is that this is a possible manoeuvre, not only for individuals—many of whom listening to me now will have discovered these things for themselves, but for society itself. And here many people are not so optimistic as I find myself to be, because the task looks too great for them. As I said at the very beginning of these talks, a lot of people feel trapped.

But I would like to make sure that you hear the following words, which form a conclusion drawn from his own cybernetic analysis of the societary condition by Sir Geoffrey Vickers. He says: “The trap is a function of the nature of the trapped.” Then I should go on to say that the failed society that I have depicted to you is not a malevolent society, not a cunning trap—for I agree with Vickers that the “trappiness” of the trap lies in our own nature. It may be nearer the mark to speak of a *Dinosaur* society.

By this epithet I mean to say that the trouble with our institutions is in their loss of the ability to respond in time, to learn in time, to adapt and to evolve. Like the dinosaurs, they cease to be viable systems. I have tried very hard to lay bare the mechanisms that appear to me to lead to this disaster, because I think they can well be understood. What we understand we can control.

“Control”: there goes that word again. I can only hope that by now you will know how I am using it. When I say that any system is “in control”, I mean that it is ultrastable: capable of adapting smoothly to unpredicted change. It has within its structure a proper deployment of requisite variety.

Just now I said: but we shall still need maps. The societary maps we need are in my view the cybernetic maps that I have tried to set forth. And so you see why I have called this series *Designing Freedom*. The contradiction built into this title is the figure of speech called oxymoron. The freedom we embrace must yet be “in control”. That means that people must endorse the regulatory model at the heart of the viable system in which they partake, at every level of recursion.

Hitherto this could be done by underwriting a constitution, or by declaring a belief in something called the democratic process. But things have changed. Constitutions, written or unwritten, turn out not to have requisite variety in a world gone crazy with its own proliferation of variety. And if the democratic process does not seize upon and use those disregarded tools of modern man, it will not itself be viable much longer.

Then all of this becomes an appeal for scientific efficiency, which belongs to the word “designing”, as providing a regulatory model to give requisite variety to human joy and fun, which belong to the word “freedom”.

The clash and conflict of these conceptions produce in all of us, and not least me, a dissonance. The idea jars, like any oxymoron. Then let us speedily reconnoitre this trap—which is a function of the nature of the trapped.

There are two things wrong with the role of science in our society. One is its use as a tool of power, wherever that is concentrated by economic forces. The other is its elite image. None of us wishes to be

manipulated by power; and if science is the tool of power, to hell with it. None of us wishes to entrust our liberty to a man in a white laboratory coat, armed with a computer and a row of ball-point pens in his pocket, if he does not share in our humanity.

The contrasting argument is just this (and I have used a lot of willpower in holding back the argument until this final lecture). Civilization is being dragged down by its own inefficiency. We cannot feed the starving; we cannot stop war; we are in a terrible muddle with education, transportation, the care of the sick and the old; institutions are failing, and often we feel unsafe in the streets of our own cities. All this is inefficient. *Then it cannot be correct to say that the only way to preserve liberty is to be so damned inefficient that freedom is not even threatened.* We have to become efficient in order to solve our problems; and we have to accept the threat to freedom that this entails- and handle it. Everything that man can do contains implicit threats. This is something written into the law of requisite variety, as far as I can see. Then we have to be knowledgeable, and we have to be untrapped.

We have to find a way by which to turn science over to the people. If we can do that, the problem of elitism disappears. For surely I do not have to convince you that the man in the white laboratory coat is human after all, and would rather use his computer to serve you than to blow the world apart? Then for God's sake (I use the phrase with care) let us create a societal system in which this kind of service is made even possible for him, before it is too late. At the moment, the scientist himself is trapped by the way in which society employs him. What proportion of our scientists are employed in death rather than life, in exploitation rather than liberation? I tell you: most of them. But that is not their free choice. It is an output of a dynamic system having a particular organization. Remember the waves.

And so my first conclusion to these lectures is: efficiency does not entail tyranny-if we can get the system right. To do so is a top priority, because some version of efficiency is required to save our dinosaur society.

The next point I would like to tackle also involves an oxymoron: you could call it "unpredictive prophecy". It would not surprise me if I sounded like a prophet, or (to use a hideous neologism) a futurologist. Let me rid myself of any such pretension; because I do not believe that we can predict the future. I believe instead that we can describe the present with perspicuity, if we use the proper instruments, and that this same present constrains future variety. This is not the same thing and I will take a few minutes to explain.

If I were to offer you a cigarette, what would happen? You might take it, or decline it. You might deliver a eulogy about the excellence of tobacco; you might read me a lecture about lung cancer. If you are a director of a tobacco company, you might insist on giving me your packet. If you are the chairman of an anti-smoking lobby, you might punch me on the nose. For all I know, you might execute a war dance—and we could offer a prize for the best explanation as to why you might. The future, in short, is unpredictable, because there is too much variety in the air. It is called freedom.

On the other hand, I may have a model of you, I may have found out how you are, and have a very good idea of what you will do. This fact does not constrain your freedom; it constrains the variety of my model of how you use it. If this distinction had been understood some time back in history, there would have been less confusion in what used to be very popular discussions of free will.

Then let us extend the analysis to planning. If I set out to catch the 8:32 train tomorrow, then maybe you will find me on it. It would be absurd to say that if I were a free man I might just as well be at home in bed, or flying the Atlantic. Planning is a variety attenuator. What is planned tends to come about—but often rather shakily, so that perhaps we make a loss when the whole idea was to make a profit. If so, variety has sneaked back in again, when we thought we had rid ourselves of it. In either case, the act of planning does not rob us of free will.

Then why has planning such a bad reputation? For it has. People talk about "the planners" in a very

pejorative tone of voice. The reason surely is that our plans are not adaptive, and the institutions charged with making them grind on with their implementation long after it has become obvious to everyone who will be affected that the plans are inappropriate. We are back to the unviability of the institution again. Institutions are stuck with their ponderous machinery, while the newspapers reflect the public rage. ...

Planning should be continuous and adaptive. Societary plans should continuously abort, and be recast, before they give birth to a monster. If this is true, there is no need to base them on the predictions that no-one can correctly make in any case, but only on the analysis of an unfolding situation in which every decision constrains future variety. In that statement the unpopular notion of planning is turned on its head, and deserves to become popular again. Because it means that the future is something we use our freedom to determine, rather than something that is lurking out there, and will happen to us, unless we are mighty smart. We can make, rather than prophesy, the future.

As to variety sneaking back in again, we can keep an eye on that. Again, this is hardly forecasting. It is an analysis of current patterns of variety, so as to assess the probabilities that a system will next move to one state (indicated by a representative point) rather than another. This process has no bearing on the problem of freedom either; it is simply quantified business acumen. Science can do something about that, through operational research. But I must add that I always laugh when I hear a businessman or a politician talk about “a calculated risk”, because this invariably means that he is taking a risk that he cannot calculate.

Then let me sum up my next key points, not as predictions at all—and therefore not as doom-laden—but as analyses which indicate firstly where things are wrong and secondly how they could be put right.

Civilization operates through a set of institutions with a particular organization. This organization appears to be an anachronism. It worked well enough in a more leisurely age; but now its relaxation times no longer match the rate of perturbation. Therefore these systems are actually designed to have unstable outputs. There is evidence that the outputs really are unstable, a fact which tends to confirm the hypothesis; and there is no cybernetic regulation in the design to stop the instability feeding on itself to the point of catastrophe.

Then we can see what to do. We cannot grab hold of explosive variables and drag them down to earth again. If we get tough, and also expensive, and reinforce the whole machinery (which is what we tend to do) we stand to lose our freedom. Moreover, and absurdly enough, this approach simply makes the machinery heading for catastrophe more efficiently catastrophic. What we do is to redesign the system itself, so that its outputs are no longer unstable.

To do this we need much faster communications inside the machinery, and these are readily available. It means using telecommunications properly, in high-variety, real-time, broad-band circuits available to all. To be available to all, they may very well need to be free of charge—like the air and the view, on which our humanity and survival also depend. I see this large expenditure as quite proportional to the threat we have to meet, and far less absurd than equivalent expenditures on which we needlessly but cheerfully embark, and which it would be embarrassing to list.

Next we need to use the computer properly inside this network; not as a device to make silly mistakes, not as a calculator to do cheap sums expensively, and not as an invigilator of the people’s free expression of themselves. Those proscriptions would knock out ninety-five per cent of current applications, and free computer power so that people could engage in their personal evolution—by guiding their own learning, and editing their own input.

Very likely computer power should be free of charge as well. Let us note that it becomes increasingly expensive to monitor charges for high-variety services. Each consumer absorbs these to a different degree and in a different pattern—and it all has to be logged. If a toll-road is opened, so that the cost of building

the road may be met by the toll, we shall need an organization of requisite variety to monitor the use of the road; and we may find that all of that costs more than the toll is raising. If that can happen in a relatively low-variety system, the situation is far more ridiculous in the high-variety systems that I adumbrate.

So we should beware of precedents in these matters. There is hallowed machinery built into all our institutions which knows how things are paid for from the public purse. Well, maybe all that too is out of date. If I lived in an isolated prairie community, and discovered that no-one in the capital was taking the question of my isolation seriously, so that my telecommunication circuits and computer power would have to be paid for as a function of my distance from the city, I would form a local committee and propose to charge the city dweller on his holiday for looking at our local view.

Given all this technology, we need new institutions for handling it; which brings me to summarize what I have said about freedom. So many people seem to think that to advance along these necessary paths might cost our liberty.

To this I have replied that our personal freedom is not the absolute we take for granted. We are profoundly constrained by the limitations of our brains and by the inexorable attenuation of our input variety. That is how we are, and we ought not to start our thinking from a worldly minded pretence. Secondly, I do not forecast or predict that such freedom as is our natural right will be imperilled: I say with passion that it is imperilled *now*, but we are too complaisant to face up to this. We live in too cosy a world. This is not the real world, uncomfortable and discomfiting, where so many people are enslaved and dying; it is a variety-attenuated model of the real world, in which these stark horrors acquire that air of unreality which our television screens know well how to bestow.

Therefore the argument is that something must be done to redesign our institutions, boldly using science in that very cause. Society, in the form of its own institutions, public and private, is making a bold use of science now—not to redesign, but to reinforce itself in what may turn out to be its most oppressive aspects.

Conspicuous consumption is an oppressive cause if it means robbing the Third World. Science is behind this, primarily because of the way telecommunications are used. Not only does television serve the cause of spurious growth; it has become little short of optical imperialism. Please contemplate those plays in which “bandits” are trying to overthrow the rightful king, only to be put down by clean-cut heroes sent in to help by the First World, from the point of view of men and women fighting dangerously for their liberation from a tyrannical dictatorship. I would like to remake one of those plays, using the same cast, from the standpoint of the so-called bandits.

Next I draw attention to computer-driven systems that compile dossiers on the individual, to rob him of his credit and his good name. That is oppression. If multi-national companies are allowed to use science on a global scale to exploit the planet’s dwindling finite resources for the benefit of the few in whom the power to do this is concentrated, then that will be oppressive. And if the might of military science is used, or even threatened to be used, against the democratic choice of any nation, then that is oppression indeed. This last example, unlike all the others, is not new; but if we are going to pour so much science into that oppressive purpose, at least let us use science in the service of freedom too.

All the oppressive uses of science that I have mentioned are in full deployment now. So science is not a neutral thing, as many scientists themselves try to believe. As for the public, I sometimes think they just hope that all this power implanted in our institutions will not hurt them, if they are quiet as mice. But the mouse trap is loaded with cheese, called growing prosperity, conspicuous consumption; and the destructive force stored in the wound-up spring is the economic power that underwrites technocracy. Then we can lose our freedom ... *snap!*

The intuition that this could be the case is there all right. It is built into that alienation of which I have

spoken. But alienation leads to impotent rage, perhaps to violence; it is an excess of human variety that is blocked off, and is explosive; alienation of itself does not lead to new constructions. Nor are we led to new constructions simply by dismantling the bureaucracies, although I have advocated this. Besides, how does one do it?

Requisite variety for running the world does not exist in any man's headful of ten thousand million badly programmed neurons. Requisite variety for running things properly exists with the people who generate the world's variety in the first place, and that means everyone. Whoever opts out of his or her regulatory role is robbing the total system of its power to be stable. Therefore it is not for me to specify the content of the total regulatory model, but only to point to the need for it. But if this stricture applies to me, it also applies to you. The requisite variety for being messianic belongs only to the genuine Messiah.

I suggest that the first thing to note is that most of us have done what I just said we should not do: we have robbed society of regulatory variety by our passivity. The occasional democratic exercise of a vote is not a big enough variety amplifier; and besides many of the most thoughtful people I know have given up voting anyway. This is on the grounds that to choose between alternatives to which one is indifferent does not increase regulatory variety at all. Then people will need to abandon their cynicism, and become active.

Their accepted course is to get into societary institutions and to try to change them. Again, many thoughtful people have given that up- because they perceive the effort as a losing battle. And if the analysis I offered of bureaucracy is correct, they are probably right; especially if the relaxation time hypothesis is correct as well. The only conclusion that I am able to draw is that we must start again. If that is not to result in anarchy, then the institutions themselves (including of course government) must take a hand.

That would sound like the kiss of death, to any good revolutionary. But I persist in that other hypothesis: that institutions, including government, operate with good intentions, in good conscience. If you and I have understood the problems, why not they?

Then suppose that groups of people draw together to consider the problems of society, and what kind of society they want. I cannot tell you the content of their deliberations. But the regulatory model will have to do with the control of variety attenuators and the provision of variety amplifiers, at various levels of recursion; and it will have to do with the way in which science should be harnessed to these ends. I do not think that the problems of acquiring scientific tools are nearly as difficult as they sound, despite the expense. The greater problem is the alienation from science that has already set in, and needs to be reversed. For I should be quite content if these groups of which I speak considered my views about the need for science to be quite wrong, so that they decided on a craft culture if they reckoned it would work, but only if they had free minds about it. Knowledge is a human possession, and that includes science-which is only ordered knowledge.

Science makes bold use of experiment: I mean the crucial experiment-something that may fail, and thereby falsify a theory. In attempting social advance, we work in an evolutionary fashion, testing the route with a toe all the way. Now of course I believe that this is much too slow. We do not have that much time. I advocate the bold experiment, but on condition that it is recognized to be just that. For here is a key thought: we can very well afford to pay ourselves for being wrong. To be wrong slashes variety; one thing the scientist knows full well is that, in experiments, it is just about as useful to be wrong as to be right. Both outcomes attenuate variety, until the search homes onto the answer that we seek.

So I would say that it would pay to set up experimental institutions, deliberately antithetic to the existing ones-and with their full support. The objection is immediate and clear. Just who, and just whose children, would be the guinea pigs? I tell you that the answer to this is a great many volunteers, for they could have

the safeguards. They would design those, like the experiments, themselves. The reason why I feel so sure about this, is that so many people are doing it already- without any permission, without any safeguards, and also without any call on funds to which I reckon they are entitled. For this is liberation.

The rest of the design is simple. If science can do whatever can be exactly specified, and if people really do start specifying, then they will need recourse to science. It needs only a tiny team, and no bureaucracy of any kind, to make the links. The levels of recursion must be got right. But it could easily be done. One team for every province, one in Ottawa—and that not to tell anybody anything at all, but to co-ordinate the efforts, to communicate (by videotape, of course) the results.

But I have said enough. It is not for me to project my own imaginings upon the world, although it seems legitimate to try to release untapped and perhaps frustrated variety. These things cannot be forced, but perhaps they can be freed.

Why freed, you ask? Why do they not happen of their own accord, if they are good? The answer lies, I think, in mass effect. Because to use science is expensive, the little group—however fervent—finds science difficult to command. When the movement is general, however, the cost is shared and becomes manageable. This is the reason for my little teams. And who should pay for those? Come, well-intentioned institutions, in good conscience: If one of your staff has a natural place in such a team (and his or her election by the social group would be a genuine honour), why not let go? *Second* the person: you have much to gain. And it is you, after all, you institutional man, who has tied up this person (perhaps he is yourself)—with the high salary and the fringe benefits you pay- and robbed him of mobility. Could you not make this gesture to freedom, and indeed survival?

But when I speak of mass effect, I could point to no more potent an example than that of a country, acting through its democratically elected government, that turns its whole self into an experimental society- and of course, I am citing Chile once again. In the third lecture I discussed a system designed for economic regulation; but this was an almost incidental feature of the Chilean Experiment. That began with agrarian and industrial reform, with making food and clothes available to the poor, and continued in a surge of enthusiasm for what even the main opposition party would refer to calmly as the Chilean Process. It was the middle class who had to pay for this: they knew it, and pulled a wry face. But they were mostly well-intentioned people in good conscience, and mostly they behaved decorously. I knew many who voted for Allende. They made up jokes about the shortages and queues, and carried on.

In the two years of my own work in Chile, I witnessed several attempts to pull the government down: one very serious attempt was made in October 1972. To this, which produced high stress and great difficulties, the Chilean people responded the following March by turning out to the polls and increasing Allende's vote by an amazing seven per cent. But he was still a minority government, a fact which tied his hands; and now he looked as though he might succeed. It was time to halt the great experiment.

As I see it, the rich world would not allow a poor country to use its freedom to design its freedom. The rich world cut off vital supplies- except for the armaments that eventually reduced La Moneda to a smoking shell. The rich world cut off vital credit, so that there was no hard currency—except for the illegal flows of it that financed the contrived paralysis of the distribution system to justify the coup.

Then let us not say, as we hear said, that Allende reduced his country to chaos, and destroyed the economy. A system of world forces acting upon Chile reduced his economy to chaos, and destroyed *him*. Allende understood that his country was losing its freedom in the oppressive grip of that external system, and went and said as much to the United Nations. The free world, as it likes to call itself, heard what he said, and waited until his own prophetic words were fulfilled: “They will only drag me out of La Moneda in wooden pajamas.” At that point it offered muted protests, and set about recognizing the military junta.

Thus is freedom lost; not by accident, but as the output of a system designed to curb liberty. My message is that we must redesign that system, to produce freedom as an output. If we are inefficient about that, on the grounds that scientific efficiency threatens liberty, then the institutional machinery that acts in our name will fail to prevent the spread of tyranny, war, torture, and oppression. We speak of the growth of prosperity; but the growth of those four things throughout the world today is yet more real.

Let us use love and compassion. Let us use joy. Let us use knowledge. These qualities are in us, obscured though we may let them be by the lethal strategies of our dinosaur society. And let us use that acquired and ordered knowledge: science. This too is in our heritage. If it has been seized by power, then seize it back. Expect it of statesmen and politicians who represent us that they should, on our behalf; or demand a new breed of statesmen and politicians. Expect it of educators that they should change the institutions of education not to train crazy apes; or start new schools and universities instead.

Above all, let us all expect it of each other that we find ways to use the power of science in better cause. It is no more sensible to say that we cannot, because ordinary folk do not understand science, as it would be to say we cannot sail a boat, because we cannot understand the wind and the sea and the tide-race.

Men have always navigated those unfathomable waters. We can do it now.

This is the book reference you need to follow up the quotation in this lecture: Geoffrey Vickers, *Freedom in a Rocking Boat*, Allen Lane, The Penguin Press, London, 1970

7 Biography



Stafford Beer

Stafford Beer is an international consultant in the sciences of management and effective organization. He is visiting professor of cybernetics at Manchester University in the Business School, adjunct professor of statistics and operations research at Pennsylvania University in the Wharton School, lately also visiting professor of general systems at the Open University. He is past-president of the Operational Research Society in Great Britain, and holds the Lanchester Prize of the Operations Research Society of America. He is past-president of the Society for General Systems Research in Washington D.C., and holds the McCulloch Memorial Award of the American Society for Cybernetics. He has also received the silver medal of the Royal Swedish Academy for Engineering Sciences, and a resolution of thanks from the United States' House of Representatives.

After the Second World War he headed Britain's largest operational research group in United Steel for thirteen years and then was Managing Director of SIGMA and built up its international consultancy. Then followed four years as Development Director of the world's largest publishing company, IPC. Since then he has been an independent consultant.