STABILISATION IN SYSTEMS
Chreods and epigenetic landscapes

C. H. Waddington

This article is abstracted from a forthcoming book, *Tools for Thought*. The author, C. H. Waddington, died in September 1975 shortly after completing his revision of the book. It was written as a popular guide to the new ways of perceiving and thinking about the world.

THE scale of very many of the impacts of mankind on the world surrounding him is now so great that they go right below the surface of things. At the deeper level, we find that most aspects of life and its interactions with its surroundings are interconnected into complexes. No powerful action can be expected to have only one consequence, confined to the thing it was primarily directed at. It is almost bound to effect lots of other things as well. Our old-fashioned common sense has not has to face such situations before, and is not well adapted to doing so. We need nowadays to be able to think not just about simple processes but about complex systems. Many suggestions have been made, particularly in the last years when the problems have become more pressing, of different ways of trying to do this.

The book is an attempt to bring together most of these proposed "tools of thought". Many of them were originally put forward accompanied by a lavish decoration of technical jargon. Part of this may have been due to the genuine difficulty of finding ways of formulating new ideas; part perhaps for the less excusable reason that it might make the ideas look more profound and novel than they really were. However, any idea that is going to be really useful in this connection can, after adequate time to digest it, be put into reasonably simple language.

Chreods and epigenetic landscapes

A natural living system has usually acquired some degree of stability by natural selection (it would have fallen apart and died out if it wasn't stable enough); in artificial systems man commonly designs a series of checks and counter-checks

Professor C. H. Waddington FRS (1905–1975) was Buchanan Professor of Animal Genetics at the University of Edinburgh, UK. *Tools for Thought* will be published by Jonathan Cape, London in April 1977.

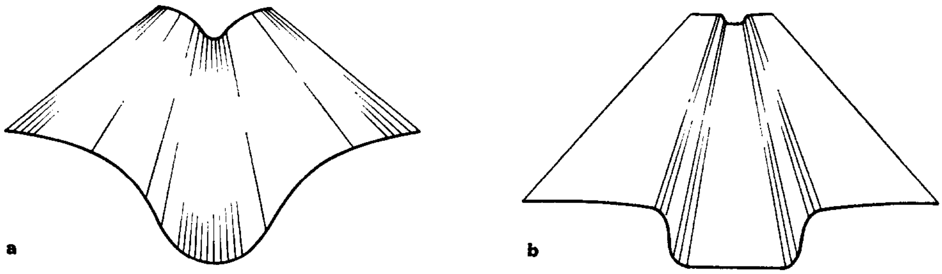

Figure 1. Chreods with different types of stability

to ensure stability. An important point to note, however, is that the stability may not be concerned to preserve the measure of some component of the system at a constant value, as in the homeostatic systems. The stabilisation of a progressive system acts to ensure that the system goes on altering in the same sort of way that it has been altering in the past. Whereas the process of keeping something at a stable, or stationary, value is called *homeostasis*, ensuring the continuation of a given type of change it is called *homeorhesis*, a word which means preserving a flow.

A phrase used to describe such systems, is to say that the pathway of change is canalised. For the pathway itself one can use the name *chreod*, a word derived from Greek, which means "necessary path". Many types of change going on in society have a more or less well-developed chreodic character; once they have got well started in a certain direction, it is very difficult to divert them.

Different canalised pathways or chreods may have rather different types of stability built into them. These can be pictured in terms of the cross-section of the valley. You may, for instance, have a valley with a very narrow chasm running along the bottom, while the farther up the hillside you go, the less steep the slope; with such a configuration of the attractor surface, it needs a very strong push of some kind to divert a stream away from the bottom of the chasm (Figure 1a). If the system is acted on by only rather minor disturbances, it is likely always to stay very close to the bottom of the valley. If one can compare several examples of such a system, there will be very little difference between them, and they will look very invariable. In contrast, we have a valley which has a very flat bottom, and the hillside gets steeper and steeper as you go away from the central stream (Figure 1b). Then, minor disturbances can easily shunt the stream from one side of the flat valley bottom to the other; it would be rather a matter of chance where in the water-meadows in the valley bottom it flows. If one looks at a number of examples of a system with this type of stability, there will be a lot of small-scale variation between them.

As an example of what the idea of chreods, or canalisation, means in an everyday context, consider the accumulation of wealth in a community. Any individual receives some income and has some outgoings; his wealth grows or diminished according to the difference between these. In the well-known words of Dickens' (pre-decimalisation) character Mr Micawber: "Annual income £20, annual expenditure £19 19s 6d, result happiness. Annual income £20, annual expenditure £20 0s 6d, result misery".

We can draw a diagram with wealth measured upwards from the base-line, and time flowing from left to right. An individual at any given time has a certain wealth, and this is represented by a corresponding point; but his current bank balance is affected by his income and his outgoings, and as time passes it will move along a line, horizontally if he keeps them exactly in balance, upwards if he is getting richer, downwards if he is getting poorer. Such lines are known as *vectors*, and the whole area of the diagram is a *vector field* (Figure 2a).

It is unfortunately common experience that in our society, and in many like it, the rich tend to get richer, and the poor poorer. So if we start with a population of people with varied incomes, the vector field tends to look like Figure 2b.

A tendency towards canalisation, or the formation of chreods in such a system, would occur if there were forces at work tending to limit the steepest increases, or the most drastic reductions, in the wealth of individuals. In practice we do have such forces in Britain, eg steeply rising surtax of large incomes and welfare benefits added to low ones. In their simplest form, such forces might act to produce a "rich, but not too rich" chreod and a "poor, but not too poor" one. Actually, the controlling forces are more complex, and tend to result in several rather than only two wealth classes. We would like, of course, to put into operation forces acting within the vector field in such a way that no one actually

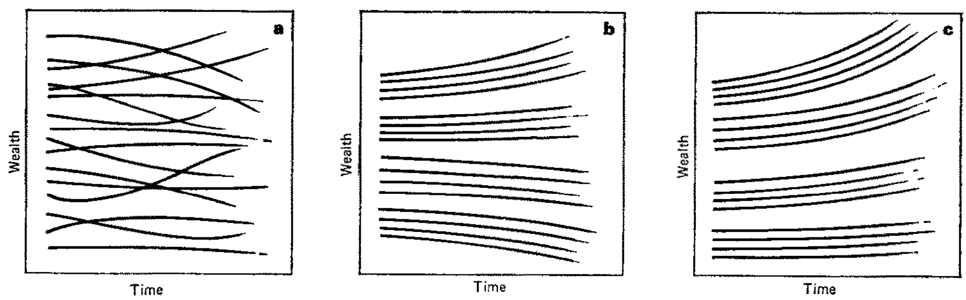

Figure 2. Vector fields showing wealth over time

ever gets any poorer (Figure 2c). And then, a further step, provisions that make it reasonably easy for individuals to move over from one chreod to another, preferably stepping up rather than down. Or, of course, we might try to set up a non-chreodic system, which did not tend to produce distinct classes at all, but at the same time prevented gross differences in wealth, and an overall upward trend.

In progressive biological systems, such as developing embryos or plants, one is usally confronted with a system which cannot be fully described in terms of a single chreod, or even collections of roughly parallel chreods, as is seen in the wealth diagrams. When an egg is developing, different parts of it will follow different courses of development, and eventually finish up forming parts of the final animal: some parts becoming muscle, some becoming nerve and so on. This can be pictured in terms of an "epigenetic landscape", in which when the process starts there is a single valley, but this later branches into two or more, and these branches split up again and agin, until they have formed a number of separate valleys corresponding to the separate parts of the adult animal.

Many progressive systems in fields other than biology behave in a similar way. For instance, when a town is beginning to grow into a city, at first it is likely that it will all be following one and the same path of change, related to a single city centre. Later on some sub-centres will develop, or suburbs will be founded, and the single pathway of change will have become diversified into a number of subsidiary paths.

Again, consider the historical development of some type of thinking, such as Christianity or Marxism. It starts with a single line of development, and later splits up into a number of more or less divergent paths of change, such as the Orthodox and Catholic churches, or Leninism, Stalinism and Maoism. Each of these not only goes on developing as time passes, but may in its turn split up into further sub-divisions. And each division has quite a lot of canalisation, or chreodic character, in the sense that someone who starts off trying to be a heretic, standing midway between Maoism and Leninism perhaps, finds that there are strong pulls trying to drag him into one or other of the orthodoxies.

In such an epigenetic landscape, there are branching points at which a valley splits up into two or possibly more branches. What can one say about them? And what can one say about the question of whether the system as a whole goes down one branch or the other, or breaks up and part goes down each? Take the second question first. Sometimes one knows that there is a branch point ahead of the system, and that if one can give the system a push at the right time, it can be diverted into one or other of the alternatives in front of it. The point to notice here is that it is in general no use giving the push too early: if you do, the system will have got back to the middle of the valley again before it reaches the fork, and the effect of the push will have been dissipated.

The period just before the branching point, during which a push will be effective in diverting the system into one or other path, is known in biology as the period of "competence". It is no use trying to act on the system to divert it into a particular branch until it has become competent to respond, by going down the valley towards which you have pushed it. Equally, of course, it is not advisable to leave the push until too late. Once the system has started to go down one of the branch valleys, if you still want to divert it into the other you have to push it right over the hill between them. Effective revolutionaries, like Lenin, have been brilliant in choosing just the right time to give a push to a society coming up to a branch point in its stability system.

The question of how the branch points come to be there at all is a difficult one. Remember that the shape of the valley (the slope of the hillsides) represents the net result of a whole lot of controlling actions, each of which is brought about by network effects, or feedbacks. Now the strength of these feedbacks and controls would depend on the amounts of the various components present in the system. Therefore, as the system progressively changes as time passes, the strengths of the controls will also alter. As some components of the system increase in magnitude and others contract, controlling interactions which were at first of minor importance may become much more significant, and vice versa. Eventually the system has altered so much that its controls can no longer ensure the stability of the former pathway. It may then break down into a general chaotic turmoil, or it may undergo a branching into two alternative new paths, each with their own stability. It is perhaps rather surprising that so many systems we come across seem to behave in the second way, rather than the first; and I do not know that there is any good explanation yet why this should be so. The theory of such breakdowns of stability is in rather an early stage of development. It is known as catastrophe theory. It seems likely to be one of the most interesting —and quite likely the most important—types of mathematical thought in the near future.

One can again use the idea of an attractor surface to help visualise what is going on. To simplify the model let us forget about providing sloping sides to a valley to represent the canalisation of a chreod, but instead indicate this by drawing arrows on the attractor surface, indicating that things get pushed in towards the line representing the pathway of change. We can start with a more or less flat attractor surface, with a line on it, and arrows pointing towards the line. The changes that are going on in the control systems as time passes can be represented by bending the attractor surface. A branch point, or catastrophe, occurs as soon as this bending results in there being a fold which brings one part of the attractor surface vertically under some other part (Figure 3). Consider

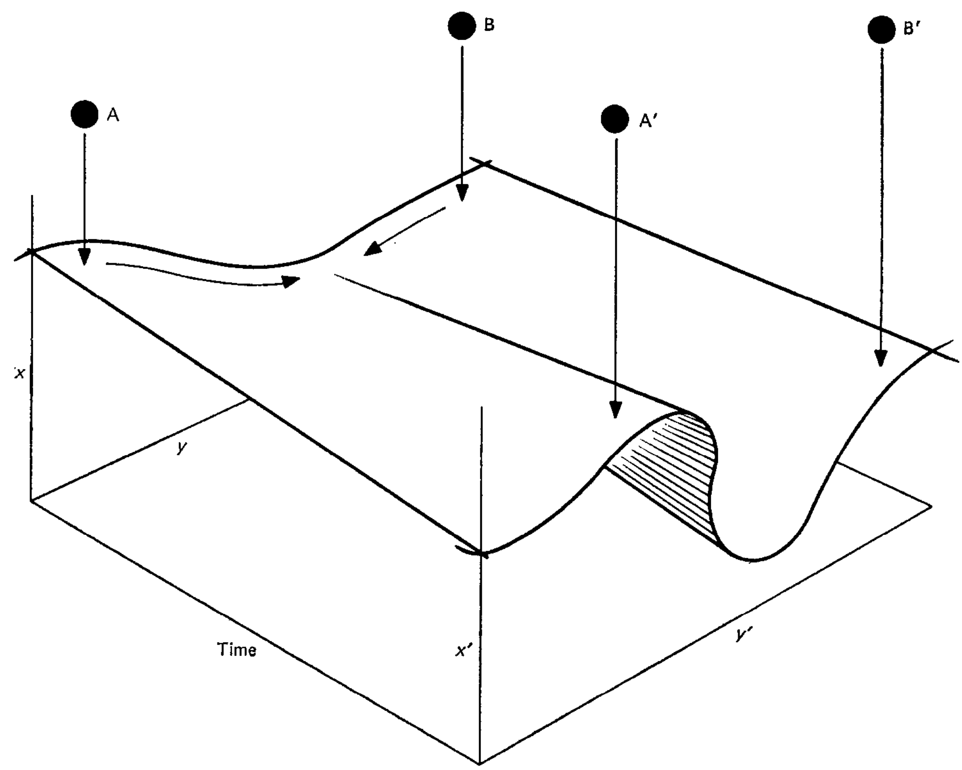

Figure 3

then what will happen if a disturbance pushes the system into an abnormal position such as those shown in A or B. The displaced point will first fall back on the attractor surface and then run along the surface to the attractor line. If, at a later time, it gets displaced to A′ or B′, it all depends whether it falls on the top surface of the fold, when the arrows will push it back into the top chreod; but if it has been displaced to B′ it will fall down on to the lower surface, and the arrows on that will bring it into the lower chreod.

Perhaps the visual model gives one some intuitive understanding why systems which reach the limit of their initial stability are often split up into two stable pathways rather than resulting in complete turmoil. The formation of such folds in a surface is the kind of thing one could reasonably expect if the surface had a certain solidarity or strength. Complete breakdown into chaos would have to be represented by total disruption of the attractor surface, which one might guess was less likely to happen.

Exploring a landscape

An important question about the epigenetic landscape and branching pathways is this: When we are confronted with an unknown system, how do we find out what the shape of the landscape is? In such parts of biology as the study of embryonic development, we can study a large number of individual eggs of the same species, do a lot of experiments on them, and explore their stability by seeing how far they can get back to normality after various sorts of disturbance. In many more important systems, such as those we meet in human and social life, there is only one example of the system, and we cannot use this way of exploring it.

One suggestion, due to the Russian mathematicians Gel'fand and Tsetlin, is to proceed as follows. We find ourselves doing something to a system which we believe has certain stability characteristics, which could be described by an epigenetic landscape; but we have no idea where we are on the landscape when we first start trying to affect the system. So what we should try to do is to try to alter it, slightly, in as many ways as possible, and observe its reactions. We will find that the system resists some types of change more than others, or restores itself more quickly after changes in some directions than in others. We can think of our actions as going out into the landscape for the same distance in every direction from where we stand, thus describing a circle. But the reactions of the system will be different in different directions, since it will be harder to push the system "uphill" than "downhill". Thus its response will be a distorted circle. From the directions of these distortions, we can deduce the main slope of the part of the epigenetic landscape where we have found ourselves. We then move down this slope—operating in the direction in which the system is most easily altered—taking quite small steps, until we begin to find a mounting resistance: then we immediately withdraw a bit, and are presumably on, or at least very near, the bottom of a valley.

We then remember where we started and from that point take quite a large jump—much bigger than the steps by which we descended the slope—not in the direction of downhill or uphill, or along the contour, where the reactions of the system were equal in response to thrusts of opposite direction, but in some direction intermediate between this contour and the downhill direction. This large jump may perhaps be expected to carry us on to the opposite hillside lower down the valley, and a local exploration around that spot may show us the slope going in the opposite direction. We can again descend this to the valley bottom, and thus trace the course of the valley we are in. That in itself will be quite useful information, but at some point, we have to make another big step, large enough to carry us out of the first valley, over a watershed into one of the other valleys of the landscape, which again we have to explore in the same sort of way.

As Gelf'fand and Tsetlin put it, we need by trial and error to fix the length of the long steps so that we "skirt high hills and step over small watersheds". One can't, of course, give any general rules for doing this. It has got to be largely a method of "suck it and see". A point of general principle is that in exploring such a landscape it would take too long to walk all over it step by step with paces of even length. It is better to alternate between

- local exploration followed by exploiting the direction of the easiest change, until the change begins to get difficult (ie opportunist reformism); and
- a jump in the dark to try to change some quite different aspect of the situation followed by another period of opportunist reform around that subject, and keeping in mind a readiness to abandon this attempt if it turns out to be a flop.

The procedure is, the Russians suggest, a judicious mixture of mild reforms in one area, followed by letting that area lie, when it ceases to respond, and starting another programme of mild reforms in a quite different context. This, they argue is the best way to get an idea of how the behaviour of the whole system is organised, which is an essential preliminary before one can deal with the whole thing in an effective way.

The epigenetic landscape of human society

What could one sensibly mean by "stabilising" or "controlling" the human situation in the world of today or tomorrow? The type of behaviour of progressive systems which we have just been discussing—a tendency to lead rather firmly in a smallish number of alternative directions—is found not only in material systems, but in psychological and cultural ones. Christopher Zeeman has used this terminology to describe the way in which many animals seem to switch suddenly between two different types of behaviour, say aggression or fear; it is often some minor factor which decides whether a dog behaves in the "attack chreod", and bites you, or in the "frightened chreod", and runs away. Many other students of animal behaviour have described similar phenomena, when it is touch and go whether a male and female fight or mate; though they mostly do not yet use the language of chreods and catastrophes which has been described here, but speak in more particular and less general terms.

In the world of today, many people fear that the situation may be getting out of hand, and soon may run away with itself, and us, into some sort of chaos. Most conventional discussions of "stabilisation" or "control" of human situations are in terms of simple quantitative measures. For instance, it is pointed out that population numbers, consumption of power, etc, are increasing exponentially so that plotted against time they give a J-shaped curve. It is argued that we need to introduce some negative feedback which will convert the J-shaped curve into an S-shaped curve.

The first point to note is that at present we are faced with J-shaped curves if we plot against time not just population numbers, power consumption, etc, but their rate of change. We have to deal not with exponential velocities, but with accelerated exponential. We are concerned with the stabilisation/control not of things but of processes; with homeorhesis, not homeostasis.

We are also in a phase of increasing diversification or differentiation. Whereas a few decades ago mankind could be classified into blue-collared workers, white-collared workers, professionals and aristocracy or plutocracy, there is now an enormously richer diversity of lifestyles and class identifications (if we still use the concept of class). Both these points have been extensively documented by Alvin Toffler in his book, *Future Shock.*

Toffler seems to contemplate both these tendencies continuing more or less unchecked—everything going faster and faster at an accelerating rate, and differentiation between social roles and lifestyles becoming ever more diversified, to the point that any one individual will switch styles (including marriage partners, friends, etc) every few years.

Merely to try to bring these processes to a halt would, in the first place, probably be impossible; and secondly would produce a condition of stagnation and ultimately deterioration. On the other hand, I confess to some sort of (perhaps old-fashioned) intuition that if uncontrolled, they would lead to a situation of total incoherence or turbulence. But what sort of stabilisation/control is conceivable? It is not good enough to talk about improving homeorhesis or buffering of particular chreods, since this would eliminate further diversification. It could amount to the authoritarian imposition of uniformity, even if this was imposed by impersonal social forces rather than by an individual dictator like Hitler or Stalin. Diversification demands that the branching of chreods, catastrophes in the sense we discussed above, be allowed to continue. I think what we are concerned with can be loosely put by saying we want to design a system in which the catastrophes are little ones, not big ones. That is to say, when we make a switch in lifestyle (from a junior advertising executive to a pop group/lyric writer, or from an experimental biologist to a futurologist, philosopher or art critic), people do want the styles to be really *different*, genuinely alternative choices, not just the mixture as before with a trifle more or a trifle less bitters in the cocktail; but surely what we want to aim at is a "system" which allows us to do this without too much danger of our whole personality being torn into shreds in the process of transition.

One illustration would be the very final edge between the land and the sea in a great river delta like those of the Mississippi and the Nile—there are almost innumerable little separate rivulets of the fresh water running down to the sea, separated quite definitely but only by low banks of mud. This would be an "epigenetic landscape" of low profile. Or one might suggest a musical analogy. We do not want to listen to the confrontation of markedly different themes as in a Wagnerian opera, but to the running through of a gamut of slight variations on one or a few allied themes, as in much of Bach or boogie-woogie.

Or we could put it another way. Using the model of a chreod as an attractor surface in a multi-dimensional space, a catastrophe corresponds to an overlapping fold in this surface. The stabilisation/control we are looking for corresponds to one in which there are many overlapping folds; but the top and bottom surfaces of the folds are not widely separated; a crumpled surface with many little scrumplings, rather than a few large impressive folds.