INVESTMENT AGAINST DISASTER IN LARGE ORGANIZATIONS

S. Beer

J. Casti

April 1975

Research Memoranda are informal publications relating to ongoing or projected areas of research at IIASA. The views expressed are those of the authors, and do not necessarily reflect those of IIASA.

Investment Against Disaster in Large Organizations

S. Beer* and J. Casti*

1. Introduction

Increasingly we are aware that the large systems we attempt to manage show signs of disastrous breakdown. For purposes of this discussion, we define disaster as the passage of some important systemic variable beyond a threshold of acceptability. Thus in managing the ecosystem, we may find that pollutants have killed all the fish; in managing the city, we may find that it becomes too noisy, dirty or crime-ridden to live in; a social service, such as the postal system, may break down altogether; a firm may go bankrupt.

Yet it is naive to imagine that one can draw up lists of possible disasters, and make investments that will avert them, although this is the standard managerial approach. Some cities, for example, have appointed officials to "do something about noise." But noise is a product of how the city is, just as inflation is a product of how the social economy is and the naivete lies in contemplating such abstractions as noise and inflation as dragons walking abroad, who can be cut down by sufficiently intrepid knights. We have to find an approach to the total system in which our potential disasters are embedded.

---

*International Institute for Applied Systems Analysis, Laxenburg, Austria.

The second problem is that we can never be sure that we have selected the right items to examine, or that the investments we are making will impinge on the potential disaster we most fear. Many such investments turn out to be counter-productive; this is because the behaviour of the total system in which the disaster is embedded is itself counter-intuitive.

In this paper an attempt is made to find a set of conceptual tools for handling this apparently intangible situation, and to set up hypotheses that may be falsified.

2. Dynamics of Disaster

We may approach the problem of modelling the dynamics of disaster by considering first an ecological situation in which there were two main components living in a symbiotic relationship. These components were a coral reef and a starfish population called Crown of Thorns.

It is assumed that this population varied through time; and, since it fed off the coral, it is assumed that the coral flourished in inverse proportion, thus:

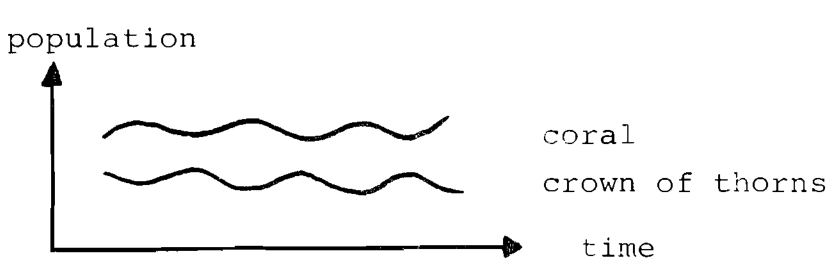

The whole point about this arrangement was its stability: nothing disturbed the homeostasis between the two protagonists. Note that we do not know why this was, because the embedding within the ecosystem was too complicated, only that it was.

At some moment, the homeostasis broke down, and disaster very rapidly supervened for the coral: many miles of reef vanished in a very short time.

It seems inevitable, then, that the cyclic dynamics of the Crown of Thorns population exploded: the amplitude of the cycle increased, until it passed a threshold that produced catastrophic collapse in the coral. Thus:

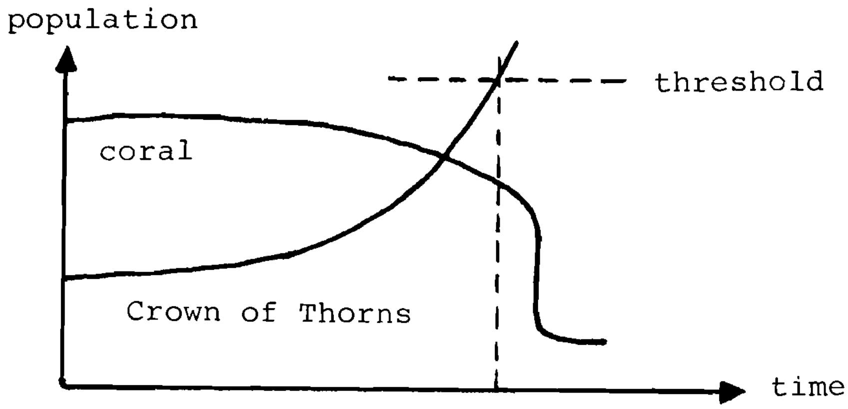

3. <u>Generalization of the Disaster Model</u>

In the management of all large organizations, homeostatic relationships may be detected between many variable factors. Indeed, the cybernetic attitude is that whatever stability they display is the outcome of the operation of a large number of interlocked homeostats, rather than the result of active managerial intervention. Active management is concerned with the objectives laid down upon the organization, and with meeting goals set for the organization. In manipulating the system for these purposes it often provokes instability in the internal homeostats rather than the reverse.

Then let us make the hypothesis that a factor in a large system that undergoes catastrophic collapse (called CF, the collapsing factor) at some specific time is homeostatically related to other factors, one of which (called IF, the implicated factor) passed beyond its physiological threshold at some earlier point in time. Thus:

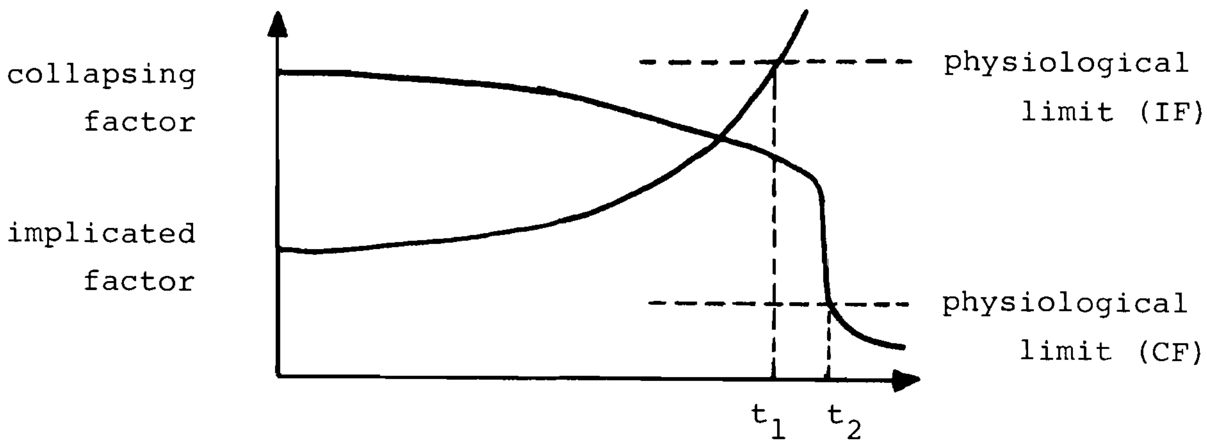

In many situations it is clear that this model holds. For example, the catastrophic collapse of a population of fish in a lake (CF) means that the fish actually die upon passing the threshold of their physiological limit. Analysis may then reveal that they died from poisioning by a particular pollutant (IF). Then it is certainly possible to specify what is the physiological limit (IF), and to observe by how much it was exceeded.

But there are complications in other cases. If the population of people living in a city is the collapsing factor, it may well be possible to specify the CF physiological limit: it is the lowest number of people that keep the city viable. The implicated factor, however, will not be a single variable, but is almost certainly a set. Then to specify the IF physiological limit we shall need a model of this set. In general, it is difficult to specify the physiological limits of homeostats, just because they are physiological. By this we mean that the limits are functions of the organization and dynamics of the homeostat. We may not succeed in understanding these well enough to specify the limits; in any case they may well be changing continuously through time thanks to the successful operation of the homeostat itself. However, in this disaster model, we shall probably be able to specify the limit for the collapsing factor at least: if we know what CF is.

The purpose of understanding disasters is not to plot their courses but to avert them. In large organizations, the catastrophic collapse of any one of many factors would count as a disaster. And although we may be aware that some are special risks (in a firm, bankruptcy, for example), there may be a disaster through another collapsing factor (in the firm, ownership, for example, where the CF physiological limit is takeover--(the loss of control) while we are busy staving off the first.

4. Catastrophe Theory

One of the principal difficulties associated with attempts at mathematically modelling social and economic phenomena has been the natural tendency of modelling too slavishly, in fact to almost religiously adhere to the modelling apparatus which has served so well in physics and engineering. Thus, very precise local interactions are postulated between identifiable components of the system and the global consequences of these postulates are then forced upon the system by the local dynamics. The critical feature of this procedure is that it is locally rigid but globally vague, requiring precise information about the local interactions.

Unfortunately, in many social science situations no such local precision is available. Various global features of the process are observed but both lack of data and inadequate understanding of the underlying mechanisms make it impossible to specify any local dynamic. Thus, such situations, in contrast to the physical sciences, are globally rigid but locally vague.

From the standpoint of modelling, the second situation has suffered not only from the psychological blocks instilled by the usual university scientific education, but also by a lack of mathematical machinery to deal with such situations. However, in recent years the picture has brightened considerably and significant conceptual and analytic advances have been made. It is of interest to note that, in contrast to classical mathematical physics which is based almost exclusively upon the gospel of analysis according to the "Old Testament," Whittaker-Watson [7], or the "New Testament," Courant-Hilbert [1], the modern global point of view is strongly biased towards algebra and geometry (topology). This is not surprising, of course, since the tools of analysis are designed exclusively for detailed local exploration of mathematical properties. But what is surprising is that it has taken so many years for the "natural" global tools--algebra and topology--to make their appearance, a fact which one is tempted to ascribe to the cult of analysis, fanaticism, and/or ignorance.

In any event, current trends seem more promising and one of the primary contributing factors to this state of affairs is the theory of catastrophes, developed by Thom and Zeeman.

In this section, we present a very brief discussion of the basic assumptions and results of catastrophe theory in a form most useful for applications. For details and proofs, we refer to the works [3;4;6;8;9].

Let $f: R^k \times R^n \to R$ be a smooth function representing a dynamical system Σ in the sense that R^k is the space of input variables (controls, parameters) while R^n represents the space of output variables (behaviour). We assume that $k \leq 5$, while n is unrestricted. The fundamental assumption is that Σ attempts to locally minimize f. We hasten to point out that in applications of catastrophe theory, it is not necessary to know the function f. In fact, in most cases f will be a very complicated function whose structure could never be determined. All we assume is that there exists such a function which Σ seeks to locally minimize.

Given any such function f, if we fix the point $c \in R^k$, we obtain a local potential function $f_c: R^n \to R$ and we may postulate a differential equation

$$\dot{x} = - \text{grad}_x f,$$

where $x \in R^n$, $\text{grad}_x f = \text{grad } f_c = (\frac{\partial f}{\partial x_1}, \ldots, \frac{\partial f}{\partial x_n})$. Thus, the phase trajectory of Σ will flow toward a minimum of f_c, call it x_c. The stable equilibria are given by the minima of f_c and, since there are usually several minima, x_c will be a multivalued function of c, i.e. $x_c: R^k \to R^n$ is *not* one-to-one. The point of catastrophe theory is to analyze this multi-valuedness by means of theory of singularities of smooth mappings.

For completeness, and to round out the mathematical theory, we consider not only the minima, but also the maxima and other stationary values of f_c. Define the manifold $M_f \subset R^{k+n}$ as

$$M_f = \{(x,c): \text{grad}_x\, f_c = 0\} \;.$$

Let $\chi_f : M_f \to R^k$ be the map induced by the projection of $R^{k+n} \to R^k$. χ_f is called the catastrophe map of f. Further, let J be the space of C^∞ -functions on R^{k+n} with the usual Whitney C^∞-topology. Then the basic theorem of catastrophe theory (due to Thom) is

Theorem: There exists an open dense set $J_0 \subseteq J$, called generic functions, such that if $f \varepsilon J_0$

i) *M_f is a k-manifold;*

ii) *any singularity of χ_f is equivalent to one of a finite number of elementary catastrophes;*

iii) *χ_f is stable under small perturbations of f.*

Remarks: 1) Here *equivalence* is understood in the following sense: maps $\chi : M \to N$ and $\bar{\chi} : \bar{M} \to \bar{N}$ are equivalent if there exist diffeomorphisms h,g such that the diagram

$$\begin{array}{ccc} M & \xrightarrow{\;\chi\;} & N \\ {\scriptstyle h}\downarrow & & \downarrow{\scriptstyle g} \\ \bar{M} & \xrightarrow[\;\bar{\chi}\;]{} & \bar{N} \end{array}$$

is commutative. If the maps $\chi, \bar{\chi}$ have singularities at $x \varepsilon M$, $\bar{x} \varepsilon \bar{M}$, respectively, then the singularities are equivalent if the above definition holds locally with $hx = \bar{x}$.

2) Stable means that χ_f is equivalent to χ_g for all g in a neighbourhood of f in J (in the Whitney topology).

3) The number of elementary catastrophes depends only upon k and is given in the following table.

Table 1.

k	1	2	3	4	5	6
# elementary catastrophes	1	2	5	7	11	∞

A finite classification for $k > 6$ may be obtained under topological, rather than diffeomorphic, equivalence but the smooth classification is more important for applications.

Discontinuity, Divergence, and the Cusp Catastrophe

Our critical assumption is that Σ, the system under study, seeks to minimize the function f; i.e. Σ is dissapative. Thus, the system behaves in a manner quite different from the Hamiltonian systems of classical physics. In this section we shall mention two striking features displayed by catastrophe theory which are not present in Hamiltonian systems but which are observed in many physical phenomena.

The first basic feature is discontinuity. If β is the image in R^k of the set of singularities of χ_f, then β is called the bifurcation set and consists of surfaces bounding regions of qualitatively different behaviour similar to surfaces of phase transition. Slowly crossing such a boundary may result in a sudden change of behaviour of Σ, giving rise to the term "catastrophe." Since the dimension of Σ does not enter into the classification theorem, all information about when and where such catastrophic changes in output will occur is carried in the bifurcation set β which, by conclusion i) of the Theorem, is a k-manifold. Hence, even though Σ may have an output space of inconceivably high dimension, the "action" is on a manifold of low dimension which may be analyzed by geometrical and analytical tools.

The second basic feature exhibited by catastrophe theory is the phenomenon of divergence. In systems of classical physics a small change in the initial conditions results in only a small change in the future trajectory of the process, one of the classical concepts of stability. However, in catastrophe theory the notion of stability is with respect to perturbations of the system itself (the function f), rather than just the initial conditions and so the Hamiltonian result may not apply. For example, in an homogeneous embryo adjacent tissues will differentiate.

Let us now illustrate the above ideas by consideration of the cusp catastrophe.

Let $k = 2$, $n = 1$, and let the control and behaviour space have coordinates a, b, x, respectively.

Let $f: R^2 \times R^1 \to R$ be given by

$$f(a,b,x) = \frac{x^4}{4} + \frac{ax^2}{2} + bx \quad .$$

The manifold M_f is given by the set of points $(a,b,x) \subset R^3$ where

$$\text{grad}_x \; f(a,b,x) = 0 \quad ,$$

i.e.

$$\frac{\partial f}{\partial x} = x^3 + a\,x + b = 0 \quad . \tag{1}$$

The map $\chi_f : M_f \rightarrow R^2$ has singularities when two stationary values of f coalesce, i.e.

$$\frac{\partial^2 f}{\partial x^2} = 3x^2 + a = 0 \quad . \tag{2}$$

Thus, Eqs. (1) and (2) describe the singularity set S of χ. It is not hard to see that S consists of two fold-curves given parametrically by

$$(a,b,x) = (-3\lambda^2, 2\lambda^3, \lambda) \quad , \qquad \lambda \neq 0 \quad ,$$

and one cusp singularity at the origin. The bifurcation set β is given by

$$(a,b) = (-3\lambda^2, 2\lambda^3)$$

which is the cusp $4a^3 + 27b^2 = 0$. Since M_f and S are smooth at the origin, the cusp occurs in β and not in S. Figure 1 graphically depicts the situation.

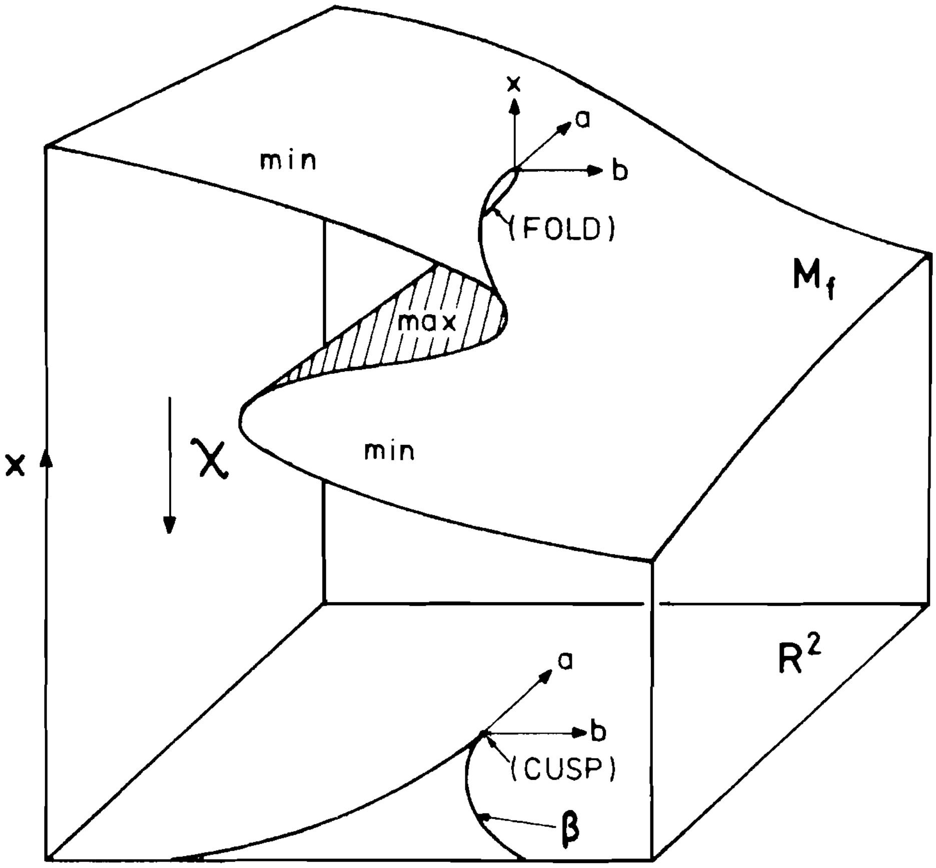

Figure 1. The cusp catastrophe.

It is clear from the figure that if the control point (a,b) is fixed <u>outside</u> the cusp, the function f has a unique minimum, while if (a,b) is inside the cusp, f has two minima separated by one maximum. Thus, over the inside of the cusp, M_f is triple-sheeted.

The phenomenon of smooth changes in (a,b) resulting in discontinuous behaviour in x is easily seen from Figure 1 by fixing the control parameter a at some negative value, then varying b. At entrance to the inside of the cusp nothing unusual is observed in x, but upon further change in b, resulting in an exit from the cusp, the system will make a catastrophic jump from the lower sheet of M_f to the upper, or vice-versa, depending upon whether b is increasing or decreasing. The cause of the jump is the bifurcation of the differential equation $\dot{x} = - \text{grad}_x f$, since the basic assumption is that Σ always moves so as to minimize f. As a result, no position on the middle sheet of maxima can be maintained and Σ must move from one sheet of minima to the other.

An hysteresis effect is observed when moving b in the opposite direction from that which caused the original jump, i.e. the jump phenomenon will occur only when exiting the interior of the cusp from the side opposite to that where the cusp region was entered.

To see the previously mentioned divergence effect, consider two control points (a,b) with $a > 0, \lessgtr 0$. Maintaining the b values fixed with decreasing a, the point with positive b follows a trajectory on the lower sheet of M_f, while the other point moves on the upper sheet. Thus, two points which may have been arbitrarily close to begin with, end up at radically different positions depending upon which side of the cusp point they pass.

While the cusp is only one of several elementary catastrophes, it is perhaps the most important for applications.

In Table 2, we list several other types for $k \leq 4$, but refer the reader to section 6 for geometrical details and applications.

Table 2. The elementary catastrophes for $k \leq 4$.

Name	potential function f	control space dimension	behaviour space dimension
fold	$x^3 + ux$	1	1
cusp	$x^4 + ux^2 + vx$	2	1
swallowtail	$x^5 + ux^3 + vx^2 + wx$	3	1
butterfly	$x^6 + ux^4 + vx^3 + wx^2 + tx$	4	1
hyperbolic umbilic	$x^3 + y^3 + uxy + vx + wy$	3	2
elliptic umbilic	$x^3 - xy^2 + u(x^2+y^2) + vx + wy$	3	2
parabolic umbilic	$x^2y + y^4 + ux^2 + vy^2 + wx + ty$	4	2

5. The Role of Investment

In order to make use of catastrophe-theoretic notions in our earlier context, we must first consider some basic aspects of investment.

By "investment" we mean the allocation of resources of any kind for any purpose that is relevant to the organization.

An identifiable collapsing factor clearly attracts investment that is intended to stave off collapse. But we have argued that not all collapsing factors are identifiable. Moreover, the investment must be applied in such a way that the implicated factor is held within its physiological limit, and we have expressed reservations about the identification of both of these too.

But it is interesting that all investment has an impact on some organizational homeostats. Therefore, whatever investment is going on may impinge (positively or negatively) on an incipient disaster. If management is wholly alert to the situation, it can take decisions that apply investment directly to the implicated factor, and possibly avert the catastrophe threatening the collapsing factor. But this is a trivial case, as is the case where management is wholly unaware, since the catastrophe will then happen or not depending on the influence exerted fortuitously by investment.

The middle section of the spectrum of awareness is where our interest lies, since we believe this to represent managerial reality in most cases, most of the time. In this area, management is aware of certain dangers (but not others), is more alert to such dangers at some times than at others, and is in consequence investing resources for all purposes in such a way as to impinge on incipient disasters to a varying degree as time unfolds. We may add a third dimension to our graph, thus:

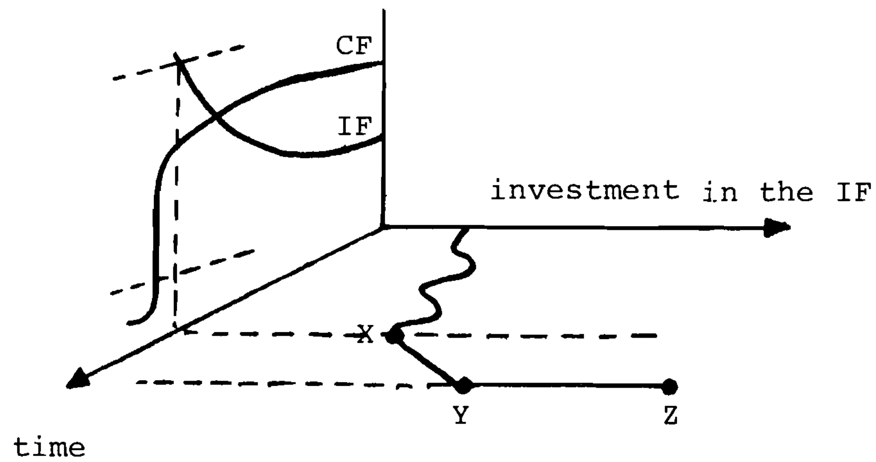

In this picture, we see the investment that is relevant to the incipient disaster varying through time up to point X, depending on the management's awareness of the risk that IF will pass its physiological limit at the time of X. The time between X and Y is a perceptual lag: it is becoming increasingly clear that CF is approaching catastrophe, and the investment to bring IF under control is increased. At point Y, when the catastrophe is evident, the investment rises very sharply in an all-out effort to retrieve the situation.

This form of the investment graph seems to be typical. And it is noteworthy that the massive investment (Y to Z) intended to retrieve the situation will only succeed to the extent that the IF set is now correctly identified. In practice, management is often hypnotized by the collapsing factor itself; it applies its investment (Y to Z) to the CF instead of the IF, because it does not understand the organizational homeostasis that underlies the disaster. This is known as "pouring good money after bad."

6. The Model Extended

The phase space of the model is now three-dimensional, and we turn to geometry to describe it. The following picture shows a manifold that has a continuous, infinitely differentiable surface, so folded that it has three sheets. It is possible to move over this surface in an orderly progression, because it is smooth. But it is also possible to "slip over the edge" of one sheet, and find oneself on another. This is indeed the catastrophe. Mathematically, the ability of this model to define discontinuities on a smooth surface offers both the simplest and also the most complicated apparatus needed. In other words Figure 1 is the only possible picture of a two factor catastrophe.

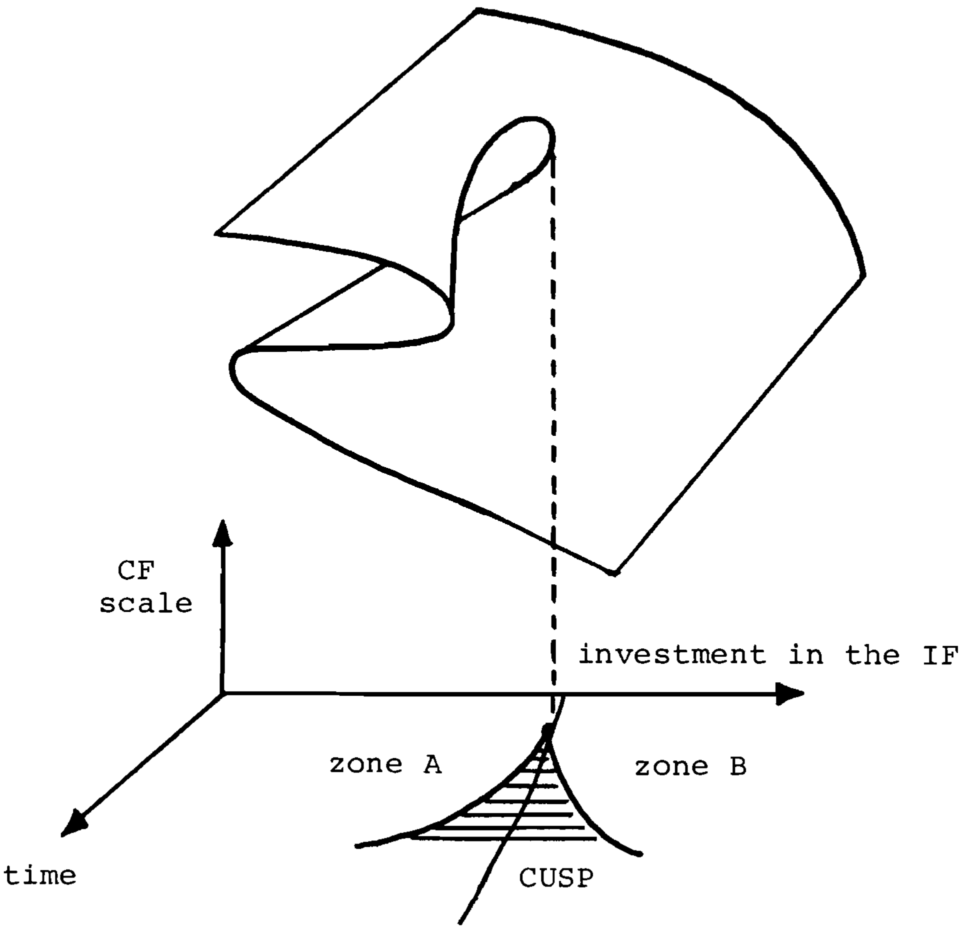

When the manifold is projected onto the two-dimensional time-investment plane, a cusp region is defined. This is the danger region where disasters can occur. The formal property is this: to cross from zone A to zone B behind the cusp is orderly; to move <u>into</u> the cusp from either zone is dangerous; to move out of the cusp into the <u>other</u> zone is catastrophic.

The precise location of the point of the cusp is central to the decision problems we are discussing. Consider, for example, two possible levels of investment, arbitrary but distinct, to be made at a steady rate through time. The levels are chosen so that they run on different sides of the cusp. Then we have Figure 2.

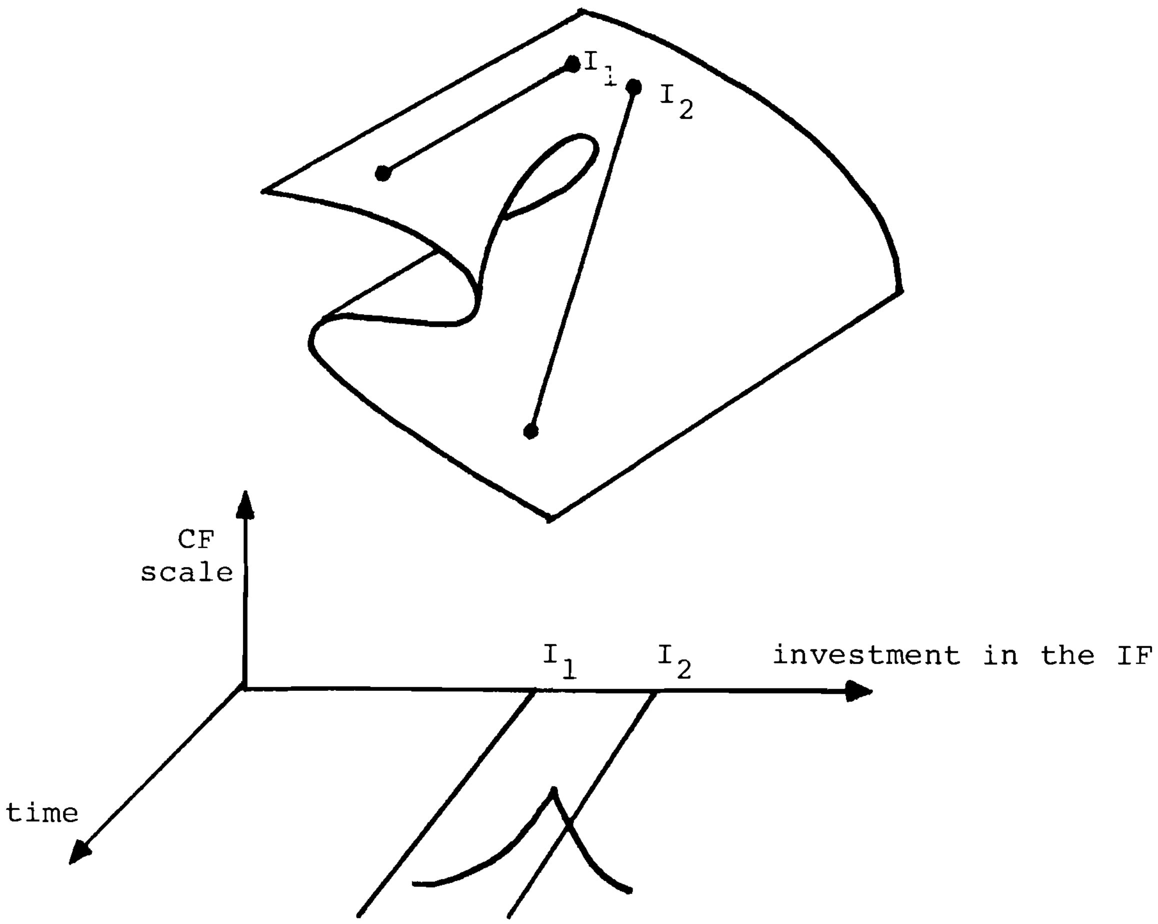

Figure 2.

Neither investment provokes a catastrophe (although, as we saw before, either may have averted one). But, as can be seen from the projection onto the manifold, the trajectory on the CF scale generates two totally different futures: the end points are on different sheets.

Next, we should reconsider the vacillating mode of investment discussed in Section 4, since this is common in practice (remember that investment refers to resources at large, not simply to agreed budgets).

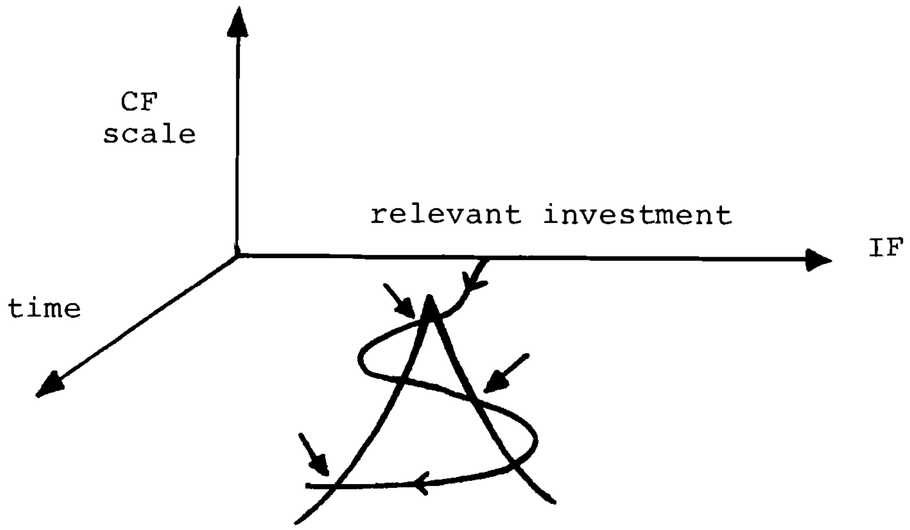

There are three catastrophes, marked with arrows, entailed by this diagram, and this may be a correct picture where the collapsing factor is morale. Reviewing several experiences of this situation, one of us (S.B.) notes that in each case the implicated factor was the availability of information; the management thought that the implicated factor was salary level; and after the catastrophes, at the time for making a massive investment, they made the further mistake of applying it directly to the collapsing factor--in terms of exhortation. Moreover, that statement also fits the facts of Revans' impressive studies of morale in hospitals [2], made nine years ago, in which he found that the implicated factor for catastrophes in nursing morale was the information needed to look after patients, and not the payment system whereby nurses earn less than shop assistants (as the management believed); again the X to Y investment was primarily exhortative.

7. The Recognition of Incipient Disaster

Aside from setting objectives, striving for goals, and sheer administration, management operates by changing the system of interlocking homeostats that define the organization.

This is reflected in our model by the stretching or compressing of the manifold in any direction. Such distortions will of course affect the "cusp region."

In principle it is possible to make organizational changes to the structure of the interlocking homeostats, so that the cusp region is widely broadened, or so that it is narrowed. If the narrowing brings the arms of the cusp region together, then we simply have one line to cross which is catastrophe.

We discussed earlier the difficulties of identifying the two key factors CF and IF, and of specifying the physiological limits of their homeostasis. This is, in our model, the problem that management faces if it is to avert disasters. How do successful managers do it?

Suppose that they so organize that the cusp region contracts to a single line. Then their lives are plain sailing in either zone A or zone B, and they are not bothered by any signals of incipient disaster. At first sight, this is an attractive managerial scenario. However, if the trajectory they are on crosses the cusp line, which counts as both entering and leaving the cusp region simultaneoulsy, catastrophe hits them without warning.

Suppose that they so organize that the cusp region is broadened as far as possible. Then the likelihood that the trajectory will enter the region is much higher. They are, in this scenario, apparently dicing with death. For the cusp region is unstable, and the trajectory may easily be thrown out into another zone--entailing a catastrophe.

There are many, many instances of managerial failure where disaster strikes without either the CF or the IF having been recognized in advance. (The paper by Walters "Foreclosing of Options" [5] gives six perfect examples.) This suggests that the cusp region is too narrow. Danger and instability have been minimized, but the "plain sailing" feeling is the managerial trap. Walters writes: "We were over a year along into a happy exercise...."

Then the hypothesis is that successful managers BROADEN THE CUSP. This increases danger and instability; but according to the hypothesis it is exactly by operating in this zone that successful managers provide themselves with the sensitivity to the risk of disaster that enables them to adjust their investment to avert disasters in unidentified collapsing factors.

In other words, we are postulating a self-adapting feedback system for continuous policy adaptation in successful management, that depends on a broad cusp.

8. The Monetary Investment

In any work done to quantify this model and to make it useful, the measurement of investment in the full sense will be very difficult. However, the monetary component of investment is extremely easy to measure for any programmed budget.

Consider a program directed to some objective. Does this program identify a collapsing factor? If so, does it identify an implicated factor? How much money is allocated to the manipulation of the CF/IF homeostat? Are organizational changes in prospect in consequence of this program that will induce distortion of the manifold, and if so what effect will this have on the cusp region? These are the kinds of questions we should like to ask; and if there is a suitable program beginning somewhere we could perhaps model it continuously on the lines of this embryonic approach.

Since investment against disaster has no payoff unless a disaster is averted that would otherwise have occurred, the degree of belief in potential disaster must govern what investment is made against it. The catastrophe lies in the future; and maybe the investment is perceived not so much as averting disaster as holding up the collapsing factor, thereby pushing the incipient collapse indefinitely into the future. On this hypothesis, it becomes difficult to make any orthodox calculations about the present worth of investments discounted up to a date of catastrophe that goes unrecognized because it does not occur.

References

[1] Courant, R. and D. Hilbert. Methods of Mathematical Physics. New York, Interscience Publishing Company, 1953.

[2] Revans, Reginald W. "Research into Hospital Management and Organization." Milbank Memorial Fund Quarterly, 44, 3, 2 (1966), 207-248.

[3] Thom, R. Stabilité Structurelle et Morphogenèse. Reading, Massachusetts, Addison-Wesley Co., 1972.

[4] Thom, R. "Topological Models in Biology." Topology, 8 (1969), 313-315.

[5] Walters, C. "Foreclosure of Options in Sequential Resource Development Decisions." Internal paper. February 1975.

[6] Wasserman, G. Stability of Unfoldings. Springer Lecture Notes in Mathematics, Vol. 393. New York, Springer-Verlag, 1974.

[7] Whittaker, E.T. and G.N. Watson. A Course of Modern Analysis. London, Cambridge University Press, 1902.

[8] Zeeman, E.C. "Applications of Catastrophe Theory." Coventry, England, University of Warwick, Mathematical Institute, 1973.

[9] Zeeman, E.C. "Differential Equations for the Heartbeat and Nerve Impulse." In C. Waddington. Ed. Towards a Theoretical Biology - 4, Edinburgh, Edinburgh University Press, 1972.