

FAIL-SAFE vs. SAFE-FAIL CATASTROPHES

Dixon D. Jones

C.S. Holling

R.M. Peterman

August 1975

WP-75-93

Working Papers are not intended for distribution outside of IIASA, and are solely for discussion and information purposes. The views expressed are those of the authors, and do not necessarily reflect those of IIASA.

FAIL - SAFE
vs.
SAFE - FAIL
CATASTROPHES*

Dixon D. Jones
C.S. Holling
R.M. Peterman

I. INTRODUCTION.

This paper is meant to serve two purposes. First, to extend the usefulness of catastrophe theory as a tool to aid our perception of a partially known world. This theory is a newly emerged branch of topology and, as such, begins to fill a large void in our arsenal of qualitative analytical tools. It is not appropriate for all important and interesting situations, particularly those requiring precise numerical results. But it is hoped that it can provide an important missing element for our environmental management tool kit.

The second purpose is to report upon some deliberations precipitated by a recent paper of Beer and Casti (1975). We shall follow, to some degree, their development. We shall also borrow some of their examples and terminology in order to emphasize some fundamentally different strategies for managing unexpected events.

* The Fail-safe/Safe-fail dual was coined by W.C. Clark while attending the 1974 IIASA Energy Project Status Report.

Two poles on the spectrum of strategies are fail-safe and safe-fail. The goal of a fail-safe policy strives to assure that nothing will go wrong. Systems are designed to be foolproof and strong enough to withstand any eventuality. Efforts are made to radically reduce the probability of failure. Often the managers of such systems operate as if that probability were zero.

A safe-fail policy acknowledges that failure is inevitable and seeks systems that can easily survive failure when it comes. Rather than rely on reducing the occurrence of failure, this policy aims at reducing the cost of that failure.

The central aim in this paper is to emphasize that there can be quite viable alternate modes of coping with unexpected events. Our traditions generally lead us to attempt to minimize the probability of crises, failures or unexpected. There are many examples of this fail-safe approach: the engineering for safety designs of nuclear power plants, the setting of, and adherence to, fixed environmental or health standards, and, the design of dams for flood control. The words of this tradition emphasize the undesirability of step change. In the risk assessment literature, risks are labelled as acceptable or unacceptable or individuals identified as risk taking or risk averse. The words suggest that unexpected events are uniformly undesirable, and if they occur they are "lived-with" grudgingly only because that is the price required for the great benefits that accrue when our designs work well. In its most sensitive expression, well represented by the Beer and Casti paper and the nuclear safeguard approaches, the inevitability of unexpected events is explicitly recognized. The goal, then, is first to design systems with broad operational limits; second to

confine the operation of the system to a limited region well away from these limits of catastrophe. The latter requires an efficient monitoring system and feedback controls that can correct deviations. The former provides the time to detect and correct the deviations.

The undoubted success of this approach has led inexorably to the design of larger and larger systems providing enormous benefits with extremely low probabilities of failure. But in partner with this scale of design and benefits is an equally high cost if failures do occur. So much so, that the trial-and-error approach that has been at the heart of technological advances becomes increasingly dangerous. No one can now possibly propose a trial nuclear plant with the expectation that failure will provide the necessary information to fill in our gaps of knowledge. The scale of the costs of error are too great. And for the first time a moratorium has been voluntarily applied to certain genetic engineering experiments because of the scale of possible experimental failures. Technology and industrial society have expanded explicitly because failures have provided an essential probe into the unknown -- a probe that generates information that knowledge must feed upon. But if trial-and-error and the learning from mistakes is increasingly too dangerous, how can we proceed in attempts to design for the betterment of mankind? That is at the heart of the issue of "hypotheticality" raised by Haefle (1973). We are locked in a world of hypothesis because we dare not test our hypotheses.

But whatever this traditional goal, there are individuals, at least, with apparently different ones. They do not accept failure grudgingly but seem explicitly to embrace the unexpected. Individuals so constituted

are the entrepreneurs, the ones that explicitly need risks, need unexpected events for personal enrichment. Tradition would have it that such individuals weight benefits from success more heavily than cost of failure. But it could equally be argued that a certain probability and cost of unexpected events is, in itself, given high value almost irrespective of benefits. And to a degree, no one could be happy, for long, in a utopia of unlimited blessings and no disturbing unexpected.

And what is true of individuals is true of institutions. Consider a research institute. In no sense could a research institute remain productive if it explicitly avoided extremes of ideas and concepts. A consistent effort to contain activities within a narrow spectrum might be necessary during a transient phase of consolidation, but if maintained for long, normal cultural forces would gradually reduce the flexibility, the operational limits of the institute. To some degree, at least, perturbations, and partially uncontrolled and unmonitored, are healthy. Every institute needs its Beers and Marchetti.

And some societies seem to have evolved similar goals. As but one example, Rappoport (1968) presents an interesting analysis of the role of ritual in the regulation of environmental relations among a New Guinea society. In its simplest form this society obtains its food from the surrounding forest, market gardens and pigs. But there is a taboo on eating pigs except on special ceremonial occasions. These ceremonial occasions are triggered when the social temperature - conflict - reaches a critical point in the village. At this point a ceremony of propitiation to the gods occurs in which the key element is the exclusive consumption of pigs. But by and large the reason the conflict occurs is because the high pig populations

begin to interfere with the market gardens. Neighbour becomes irritated with neighbour and, magically, after the feast of propitiation the problems disappear.

This is in no sense an example of an optimal food production system that produces low degrees of fluctuation. In fact, quite the opposite. It is as if a ritual is that not only is the fluctuation assured but, more significantly, strong mechanisms are developed to turn the society away from a stability boundary as the signals are detected. Rather than minimizing the probability of difficulty this society seems to have a designed method of generating detectable but controllable "failures". They occur frequently enough to prevent stability regions from contracting by maintaining flexibility of institutional response.

These examples at least raise the possibility of an alternate goal for management and institutional design. If the traditional goal is termed fail-safe the alternate could be called safe-fail. It hypothesizes that catastrophes are not necessarily lead but can, in fact, be the source of system flexibility and the cause of its maintenance. By experiencing periodic step changes, natural or cultural selection forces can act to maintain flexibility. Eliminate those periodic "disasters" and the same forces could cause an evolution towards reduced flexibility. Just as the present danger of trial-and-error approaches lies at the heart of Haebele's hypotheticality issue, so the safe-fail strategy lies at the heart of the ecologist's resilience concept (Holling 1973).

Ecological systems have a remarkable ability to absorb unexpected events and still persist. But in partner with this ability, is a high degree of variability and periodic sharp shifts of behaviour as variables

move from one stability region to another. Such shifts are exactly congruent with the jumps of behaviour shown by folded catastrophe manifolds. The real question is whether the occasional experience of those shifts is a necessary condition in order to maintain the system's capacity to absorb the unexpected. If that is the case, then there might well be a place in environmental, institutional or societal management for disaster design -- periodic "mini-disasters" that prevent the evolution of inflexibility. That, combined with traditional fail-safe design for those parts that are more surely known, monitored and controlled could lead away from the hypotheticality trap to systems with rich options for experimentation, mistakes and hence learning.

Hypotheticality raises an issue. Resilience presents a possible conceptual framework for description and prescription. Catastrophe theory is a methodology focussed on step changes. We will, in what follows, explore the value of this methodology in illuminating the issue and in making the concept operational.

For the remainder of this paper we will focus attention mainly upon systems that are not complicated by ecological or cultural selection pressures that cause a collapse in the domain of stability. The underlying structures we examine are not static, but we do assume that they are not affected by the occurrence, or not, of occasional collapses. In a later paper we will expand the scope to include systems where the "brink of disaster" closes in if flexibility and variability are restricted.

We would hope to develop criteria for manipulating systems so as to have some degree of control over the antecedents, frequency and severity of "disaster". The form of manipulation considered by Beer and Casti is investment in selected segments of a system of organization.

The system manipulations used in this paper are not linked explicitly with investment per se. We do, however, acknowledge that in most situations allocation of capital and other resources will be required to accomplish results.

By "disaster" we mean any change in a system variable that occurs suddenly and unexpectedly and which is of sufficient magnitude to carry that variable beyond acceptable limits. "Suddenly" is relative to our perception and to other variables in the system. The element of unexpectedness relates partially to our ignorance about the system. It also implies a severe inconsistency with recent trends, in short, a discontinuity of behaviour. We restrict the term "catastrophe" to its mathematical interpretation.

The types of systems that we shall consider are assumed to be only partially known and partially influenceable. Clearly, if it is important to maintain a system variable, x , at some optimum value x_{opt} , then all that need be done is to design a system such that $\dot{x} = -(x - x_{opt})$. With this system x is completely safe from disaster since it is uncoupled from all disruptive factors and any deviations are restored by the system itself. We suggest this omnipotent example to remind the reader that in most real, complex systems such a comfortable form of behaviour is remote and such a mathematical representation so trivial as to be delusive.

Let us consider the same ecological example used by Beer and Casti of the coral reef and the crown of thorn starfish. The proposition is that the coral reef organisms and their predator, the crown of thorn starfish ordinarily maintain a modestly fluctuating but stable relationship, neither deviating alarmingly from their average abundance. Occasionally,

however, the starfish population increases. (We assume momentarily that the cause is related to some unknown, external influence.) Initially the coral can withstand the added pressure until the predator population surpasses a critical threshold. A rapid collapse of the coral follows soon after. The time sequence of these events is suggested in FIG. 1.

It might be that the population "explosion" of starfish stems from a similar mode, in that some lower level control variable drifted below a critical threshold level as shown in FIG. 2.

As an ecological aside it should be noted that the existing evidence is not yet adequate to say whether this is truly an "unnatural" disaster or if it is a typical and necessary event in the ecological history of coral and its associates. There is an emerging conception among some ecologists that such periodic disasters are a critical and necessary feature for maintaining the integrity and diversity of many ecosystems.

Recent travellers to Eastern Africa report vast areas of devastation to forest land by "marauding" elephants. The situation is visible enough to initiate programs on elephant control (Read: Liquidation) by some affected governments. However, this periodic tree destruction might well be a necessary force in the maintenance of the typical savannah/grassland character utilized by ungulate herds.

Fire has been cited as playing a similar role in the maintenance of grassland ecosystems (Copper, 1961; Kozlowski and Ahlgren, 1974). A permissive attitude toward fire is beginning to find its way into forest and parkland management policy.

The periodic "disasters" of spruce budworm outbreaks have also been cast in this light (Holling, 1973; Holling, et. al., 1975). Occasional

devastation of balsam fir, the preferred budworm host, robs it of its competitive advantage over other tree species and a rich forest diversity results.

We return to figs. 1 and 2 and describe a general disaster mode. Whether or not coral collapse is "good" or "bad" in the broad context, in fact whether or not figs. 1 and 2 truly represents the starfish/coral system, a simple and useful paradigm is suggested. Beer and Casti term the system variable experiencing the disaster the Collapsing Factor (CF). The collapse occurs following the passage of some Implicated Factor (IF) beyond a particular threshold value. The general time trace is shown in FIG. 3 (taken from Beer and Casti, 1975).

This figure lends itself quite easily to the introduction of the tools of catastrophe theory. A useful feature of that theory is that seemingly dissimilar and complex situations can be related to simpler, topologically equivalent forms where mathematical analysis is more convenient. Conclusions can then be related back to the original problem.

In this paper we shall investigate the so-called canonical forms of the elementary catastrophes. These are defined as the lowest degree polynomial representations that are topologically equivalent to catastrophes occurring with the same dimensionality. By focusing on a canonical form we shall have a specific object at our disposal. The purpose here is to illustrate some of the control options and trade-offs available to management.

Any real situation will of course be more complex than the simple forms used here. Also, just because our catastrophe manifolds are topologically equivalent, it does not follow that our trade-off curves will be also.

The point to be made is that any difficulties we encounter with the canonical forms will not likely be less in a real, more complex situation.

II. A MODEL FOR DISASTER

Our first example is the scenario suggested by fig. 3. The collapsing factor CF remains at an upper equilibrium until the implicated factor IF exceeds some threshold value. Thereupon CF collapses to some lower value. This leads us (following Zeeman (1972)) to the two dimensional catastrophe — the fold. This fold is shown in FIG. 4 together with the trajectories of CF and IF taken from fig. 3. Figure 5 is a more dramatic representation by the inclusion of the time axis. There is an added, and key, feature in this figure: at the end of the trajectory IF returns to a level below its threshold value but the collapse is not reversed.

The trajectories in FIG. 5 behave as they do because the system is assumed to be dissipative. That is, it moves so as to minimize some potential function f . This is a basic requisite of catastrophe theory. The canonical form of the potential for the fold catastrophe is

$$f(x,b) = \frac{x^4}{4} - \frac{x^2}{2} + b \quad (1)$$

where x corresponds to CF and b corresponds to IF. The system dynamic is

$$\begin{aligned} \dot{x} &= - \text{grad}_x f = - \frac{d}{dx} f(x,b) \\ &= - (x^3 - x + b). \end{aligned} \quad (2)$$

Stationary values of \dot{x} define the manifold M_f shown in FIG. 6. The manifold therefore represents all possible equilibria.

In standard terminology b (or IF) is the control for the behavior

variable x (or CF). In this system, if we wish to prevent a collapse we should manipulate b . The safest action would be to reduce b to keep it as far from the edge of the fold as possible. This, however, may not be a feasible solution.

Consider the situation where b cannot be manipulated by management efforts. We assume that the magnitude of b fluctuates in some manner associated with a probability distribution $p(b)$. How is this reflected in the probability of disaster?

Figure 6 has two metrics that describe the size of the manifold: the height of the fold h_f and the width W_f . The total height of "fall" is H_f . In the canonical form

$$\begin{aligned} h_f &= \sqrt{1/3} \\ H_f &= 3h_f \end{aligned} \tag{3}$$

$$W_f = 2(\sqrt{1/3})^3$$

Note that

$$W_f = 2h_f^3 \tag{4}$$

Disaster occurs whenever b exceeds W_f . Thus the probability of collapse is

$$P_C = P_C(W_f) = \int_{W_f}^{\infty} p(b) db. \tag{5}$$

It is almost by definition that $p(W_f) \ll 1$. Otherwise collapse would be a common occurrence and perceived as a nuisance rather than a disaster.

The configuration of fig. 6 invites an additional perspective. There is not only the frequency of occurrence, as measured by W_f , but also the severity, as measured by h_f . If these factors were independent, they could

be treated separately. But often this is not the case. In our present exercise with canonical catastrophe structures we can see just how interrelated these two properties are.

Associated with a collapse in the system will be a certain cost, C_c . For purposes of illustration we take this to be some increasing function of h_f :

$$C_c = \phi(h_f), \quad (6)$$

$$\phi(0) = 0 \quad (7)$$

$$\frac{d}{dh_f} \phi(h_f) > 0.$$

We define the system liability as

$$L = C_c \cdot P_c \quad (8)$$

Suppose

$$C_c = e^{\lambda h_f} - 1 \quad (9)$$

And

$$P_c = \int_{W_f}^{\infty} \frac{1}{\alpha \sqrt{2\pi}} \exp \left[-\frac{(b - b_0)^2}{2\sigma^2} \right] db \quad (10)$$

(see FIG. 7).

For a fixed system (i.e. one where $x^3 - x + b = 0$) the actual liability will be the result of the interplay between λ , the cost parameter; b_0 , the mean b coordinate; and σ , the size of deviations.

The management schemes suggested are of three types: (1) reduce σ , or otherwise distort $p(b)$ so as to prevent high values of b near W_f .

(2) Shift the mean value of b_0 . (3) Reduce the cost parameter λ .

Both (1) and (2) are aimed at P_c ; the former is the reactionary approach while the latter is cautionary. Scheme (3) is ameliorative.

We should also bear in mind that there are other price tags on collapse. One is the cost of restoration (if it isn't included in C_c). In our canonical example x must be incremented by $H_f = 3h_f$, the same distance as the fall. But if b is set less than $-W_f$, the restoration is automatic, though perhaps traumatic, because of another rapid shift in state.

A second price is not a cost, but a value -- the value of information. When a disaster occurs, we locate W_f , or at least the critical increment ($W_f - b_0$). Knowledge of where the cliff face is has value to those who would allocate resources to manipulate b . Because of perceptual time lags this information arrives too late to avert the present disaster, but it is useful for coping with future ones. This information will be of little value, however, if restoration is not possible.

III. MANIPULATIONS AND MANAGEMENT

The management strategies derived from the last section involve an acceptance of the system as it is. Changes are made through the available control variables. In this section we begin the transition to higher levels of system design and alteration. To this point we have paralleled Beer and Casti's system description as it pertains to disasters; now our paths begin to diverge.

In the fold system of the last section (eq.2) the parameter b was termed a control. But it is a control "as seen by" the system -- the collapsing factor x responds directly to the magnitude of b . However, from the manager's point of view the control variable may be something other than b . Perhaps it

is an investment level directed at the implicated factor b . Call the factor under direct managerial control β . Then (assuming some degree of effectiveness) there will be some functional "transducer"

$$b = g_b(\beta) \tag{11}$$

That translates effort β (investment, say) into its realization b .

If the function $g_b(\beta)$ changes monotonically with β , the control is well behaved (one-to-one). A typical example might appear as in FIG. 8a. A negative investment in this context is one that reduces b -- the amount spent is the absolute value of β . The use of either b or β as the implicated factor differ only by a rescaling of figs. 4 or 5. The beauty of the topological approach is that such rescalings result in equivalent manifolds and unchanged qualitative conclusions.

Attention should be given to two other forms of the function $g_b(\beta)$. In FIG. 8b. the function is no longer monotonic. An element of redundancy exists as more than one β value can produce the same b value. This redundancy produces "multiple images" of the manifold in the space of (x, β) . This complexity can be eliminated by finding the subprocesses involved in Fig. 8b that have a monotonic form. Such a step is called Component analysis by Holling (1963) and has been used effectively in studying ecological systems. If formally pursued, this technique could possibly become one of the fundamental tools of systems analysis.

If the function $g_b(\beta)$ is shaped as in Fig. 8c, there is an indeterminacy over some range of β . This figure is topologically analogous to fig. 4 and can be addressed by analogous techniques. We have one catastrophe structure embedded within another.

Beer and Casti postulate continual changes in the effectiveness of investment on the implicated factor (i.e. changes in the function $g_b(\beta)$) and changes in the sensitivity of CF to IF. "Management is ... investing resources for all purposes in such a way as to impinge on incipient disasters to a varying degree as time unfolds" (pg. 15). In their model, investments in various segments of an organization have impacts on many "organizational homeostats" and these impacts impinge through the cybernetic milieu upon the incipient disaster. In terms of the last section, the total investment activity produces changes, in the catastrophe manifold of figs. 4 and 5. Since they contend that these changes are occurring continually through time, they introduce time as the variable that alters the character of the incipient disaster.

The implication appears to be that the time course of all impacts on the "organizational homeostats" is unidirectional and irreversible. (Could it be that the authors are saying: "First the bad news. Systems are likely to evolve into a potential catastrophic configuration. But now the good news. If we wait it out, the cusp will spread and those menacing bifurcation lines will recede to the far corners of the control space.")

It is possible that system evolution at a higher level can trigger the creation of a catastrophe manifold with time as one of the control axes. In the present context there is no fold until some $t = t_0$ and then a growing fold thereafter. As suggested with the crown-of-thorns example, a catastrophe at one level (Fig. 2) can trigger a catastrophe at another level (Fig. 1). The useful manifolds of catastrophe theory can be viewed as cross-sections of manifolds of a higher dimension (Woodcock and Poston, 1974).

In this paper we shall not use time explicitly as a control variable but shall seek the causal factor that directly leads to changes in system dynamics. This is the proximate factor that impinges on the CF/IF homeostat. In the next step of added complexity we introduce the control variable \underline{a} as this impinging causal factor.

Again \underline{a} is the control "as seen by" the system. The actual control lever available to the manager may be α , which is related to \underline{a} through some function $\underline{a} = g_{\underline{a}}(\alpha)$.

In the organizational system of Beer and Casti the factor \underline{a} (or time's impact) was the net result of a complex of positively and negatively acting feed backs from competing resource allocations. For our purposes, we assume that \underline{a} wanders about, seemingly at random, under the influence of unknown interdependencies between segments of the system. We might also consider an \underline{a} factor that is at least partially controllable through the influence of some action α . We have, of course, the special case where \underline{a} increases unhaltingly into the future, or at least until some higher level change produces a new manifold form.

The next step introduces an additional factor to the CF/IF system. This factor can be completely uncontrollable, completely controllable, or as is most likely, some mixture of the two. To illustrate we use the canonical form of elementary manifold in three dimensions -- the cuspl catastrophe manifold.

IV THE CANONICAL CUSP CATASTROPHE

When there is one dynamic variable, x , and two control variables, \underline{a} and \underline{b} , the canonical form of the manifold is given by

$$x^3 + ax + b = 0. \tag{12}$$

This equation can be derived from the fold by the addition of an enhancement term $(1 + a)x$. The factor a has the required ability to alter the dimensions of the fold and thereby alter the characteristics of collapse.

An oblique perspective drawing of the canonical cusp manifold is shown in FIG. 9 for the range $-2 \leq (a,b) \leq 2$. As the origin of the coordinate system is at the center of the manifold, the control plane (a,b) has been lowered for easier visualization.

If our goal is to prevent disasters, an obvious prudent control manoeuver would be to first move $b \rightarrow -\infty$, and then, if desired, move $a \rightarrow +\infty$. (As a bonus you end up with a lot of x .)

Clearly, this is no more relevant than designing a system as $\dot{x} = -(x - x_{opt})$. The point is that one clear way to avoid disaster is to move away from the dangerous cusp region. In the present context our interest lies with cases where the manipulation of a and b are restricted due to infeasibility, inaccessibility, ignorance or extenuating circumstances.

It is illuminating to examine the case where the factor b is not available for manipulation. We assume it fluctuates with some distribution $p(b)$ with a central value $b = b_0$. We further restrict the "controllable factor" to the range $a \leq 0$. Thus a has the capability of Broadening the fold. This example allows us to further investigate the recommendations of Beer and Casti.

According to those recommendations the correct prescription is to broaden the range of the implicated factor (or b) without causing a collapse. That is, the threshold for collapse is increased and a stochastic excursion

of b will be less likely to reach the outer edge. But there is a price to pay and that price lies at the philosophical heart of the fail-safe/safe-fail dichotomy. By making collapse less likely we run the risk of making it more severe when it does occur.

As we shall be using eq.(12) as a specific vehicle for illustration, we should review its geometry. The generating potential function is

$$f(x;a,b) = \frac{x^4}{4} + a\frac{x^2}{2} + b. \quad (13)$$

The cusp manifold is defined by the set of points (x,a,b) that satisfy

$$\frac{df}{dx} = x^3 + ax + b = 0. \quad (14)$$

The fold lines occur in the manifold where tangents become vertical; that is, where

$$\frac{d^2f}{dx^2} = 3x^2 + a = 0. \quad (15)$$

Combination of (14) and (15) and elimination of x produces the image of these fold lines in the control plane (a,b) . These lines are given by

$$\left(\frac{a}{3}\right)^3 + \left(\frac{b}{2}\right)^2 = 0 \quad (16)$$

These are the cusp-shaped lines in the perspective plot, FIG. 9. They are reproduced in FIG. 10.

At any particular (negative) a value, the manifold is a fold as in Fig. 6.

Now

$$h_f = \left(-\frac{a}{3}\right)^{1/2} \quad (17)$$

And

$$\begin{aligned} W_f &= 2 \left(-\frac{a}{3} \right)^{2/3} \\ &= 2 \cdot h_f^3 \end{aligned} \tag{18}$$

The point on the lower sheet below the fold is at $x = -2 \left(-\frac{a}{3} \right)^{1/2}$

The total "fall" is always $3 \cdot h_f$. The fold height, h_f , is also shown on the same scale in FIG. 10.

In any meaningful situation there will be some trade-off between the cost of failure

$$C_c = \phi(h_f) \tag{19}$$

and the probability of failure

$$P_c = \int_{w_f}^{\infty} p(b) db = P_c(w_f) \tag{20}$$

The liability is defined as before:

$$L = C_c \cdot P_c = \phi(h_f) \cdot P_c(w_f) \tag{21}$$

(We use a zero discount rate and side step the necessary "orthodox calculations about the present worth of investments discounted up to the date of catastrophe that goes unrecognized because it does not occur.")

How does L change with changes in w_f (or h_f , or a)? Since $\phi(0) = 0$, $L(0) = 0$. If $p(w_f) > 0$ then $L'(0) > 0$. In words, when $a=0$ the liability is zero, and as the cusp is broadened (a decreased) the liability increases. Whether or not L reaches a finite maximum depends upon $\frac{dL}{dh_f}$ obtaining a zero value. The change in L is

$$\begin{aligned} \frac{dL}{dh_f} &= P_c(w_f) \cdot \phi'(h_f) + \phi(h_f) \cdot \frac{d}{dh_f} P_c(w_f) \\ &= P_c(w_f) \cdot \phi'(h_f) + \phi(h_f) \cdot 6h_f^2 \cdot p(w_f) \end{aligned}$$

(22)

A value of h_f (with $w_f = 2h_f^3$) that equates eq. (22) with zero will be the "worst" case. Things will improve for higher or lower a values. We leave it for the reader to investigate eq. (22) under various functional forms of $\phi(h_f)$ and $p(b)$.

Because of the strength of the relationship $w_f = 2h_f^3$, an extremely steep cost function $\phi(h_f)$ is required to override the diminishing probability of occurrence. In short, a broader cusp results in a lower probability of disaster but with a higher cost of that disaster.

There is an alternate perspective that supports the broad cusp recommendation. It is more closely aligned with Beer and Casti, but it depends upon different assumptions. Given that the implicated factor has been properly identified and given that it is being monitored, a wide cusp allows more time to react once aberrantly large deviations in IF are detected. If successful, one never knows how close one came to disaster, only that observed values of IF did not cross out of the cusp region.

In the canonical cusp example a broader cusp means a higher fold. To the extent that this is a model for more complex systems we might conservatively expect the same association to apply.

In the canonical form changes in the control b could affect P_c without affecting C_c because the cusp width is not affected. But changes in a affect both P_c and C_c . In any general case a and b will not be orthogonally aligned as they are in FIG. 9. We can expect changes in h_f whenever w_f changes.

To prevent disaster is not foolproof; we can only hope to delay it. One of the main points of this paper is to suggest that by postponing a disaster it may be worse when it finally comes.

The ubiquitous spruce budworm of New Brunswick has been the object of control for over 25 years. Control thus far has been fairly successful at least within the terms of reference of the managers. They have known, and have had to live with the knowledge, that if the control ceased to operate or to be effective, a "disaster" would strike that would be much worse than the one originally at hand. Recently some ominous signals point to an even higher level disaster despite continued successful control action -- 1975 or 76 could be a very bad year.

Are several small earthquakes less devastating than one big one? The accumulated strain in the San Andras Fault System in California has been estimated to be greater than 20 feet. If this strain were to be relieved in one "event", the result would dwarf the famous 1906 earthquake. Proposals have been made to "trigger" periodically such fault systems so that dangerously high potentials do not arise. To add a bit of charm to this sensitive idea, some proposals recommend using nuclear "devices" for the trigger. Talk about hypotheticality...

V CONCLUSION.

Minimization of L is not being recommended as the best criterion. Although arguments abound that justify this measure as being optimal for society as a whole, a little reflection will show that it will lead to sub-optimal conclusions for the survivors as well as the victims.

Traditional engineering has often opted for minimizing P_c while leaving amelioration of C_c for someone else. Beer and Casti appear to be marching with this drummer. Others (cf. Haefele, 1973) see the emergence of situations where the cost of failure is above the acceptability threshold. The scale of many systems has become so large that collapse would bring

extraordinary consequences.

The preceding discussion suggests that managerial control strategies can be ranked into the following hierarchy:

1. Relocation of the control point
2. Addition of new controls
3. Distortion of the operating manifold without addition of controls.

We have not focused much upon type 3. Before it will be useful to do so, two issues must be addressed. First, we must be able to resolve the conceptual questions that arise from management at the 1 and 2 level. The issue of selecting trade-off objectives must find articulation before meaningful assessment can be made at level 3.

The second reason for the moratorium on level 3 is an uncertainty about its accessibility relative to the lower levels. In large, highly unknown systems, will management have to work its way up through levels 1 and 2 rather than jumping straight to 3? Of course, system changes can cause distortions of type 3, but if the lower levels are not understood, these distortions will be harmful or fortuitous willy nilly and beyond the repertoire of deterministic policy actions.

We close with a comment on the two auxiliary "prices" that come with collapse. First, the cost of recovery. In many situations this cost will be inseparable from the cost of collapse. In other situations this cost will involve manipulations of a , b and x in order to return x to its former level. In systems that resemble the cusp manifold this cost will increase with distance from the cusp point.

The second price is not a cost but the beneficial value of information. As one wanders around the topography of fig. 9, the only real landmark is the cliff face of the fold. If we can discover where we are in relation to that fold, wiser use can be made of resources that affect excursions in the control variables. If we can learn "experimentally" the threshold value of the implicated factor, we are in a better position to apply investments to control it. However, a onetime knowledge may not be good enough if the system is evolving and changing through extraneous and undiscovered factors. In such situations repeated monitoring of the threshold will be necessary. As one eminent scholar has recently put it: "A little disaster now and then can be good for you" (Fiering, 1975).

REFERENCES

- Beer, S. and J. Casti, (1975). Investment Against Disaster in Large Organizations. IIASA Research Memorandum RM-75-16, Laxenburg, Austria.
- Cooper, C.F. (1961) The Ecology of Fire. Scientific American 204: 150-60.
- Fiering, M.B. (1975) Harvard Gazette, May 1975.
- Haefele, W. (1973) "Hypotheticality and the New Challenges: The Pathfinder Role of Nuclear Energy", IIASA Research Report RR-73-14, Laxenburg, Austria (also Minerva 10: 303-322)
- Holling, C.S. (1963) An Experimental Component Analysis of Population Processes. Mem. Entomological Soc. Canada 32: 22-32
- Holling, C.S. (1973) Resilience and Stability of Ecological Systems. Annual Review of Ecology and Systematics 4: 1-23
- Holling, C.S., G.B. Dantzig, G. Baskerville, D.D. Jones and W.C. Clark (1975) A Case Study of Forest Ecosystem/Pest Management. Proc. Intern. Canadian Conf. on Applied Systems Analysis, May 1975, Ottawa (IIASA WP-75-60).
- Kozlowski, T.T. and C.E. Ahlgren (eds.) (1974) Fire and Ecosystems, Academic Press, New York.
- Rappaport, R.A. (1968) Pigs for the Ancestors. Ritual in the Ecology of a New Guinea People. Yale University Press, New Haven.
- Woodcock, A.E.R. And T. Poston (1974) A geometrical study of the elementary catastrophes. Springer Lecture Notes in Math. Vol. 373, New York.
- Zeeman, E.C. (1972) Differential Equations for the Heart Beat and Nerve Impulse, in Towards a Theoretical Biology 4, Edinburgh University Press.

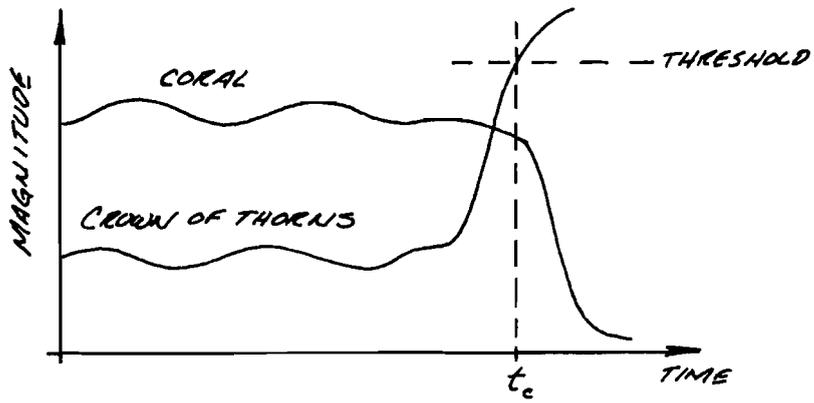


FIGURE 1

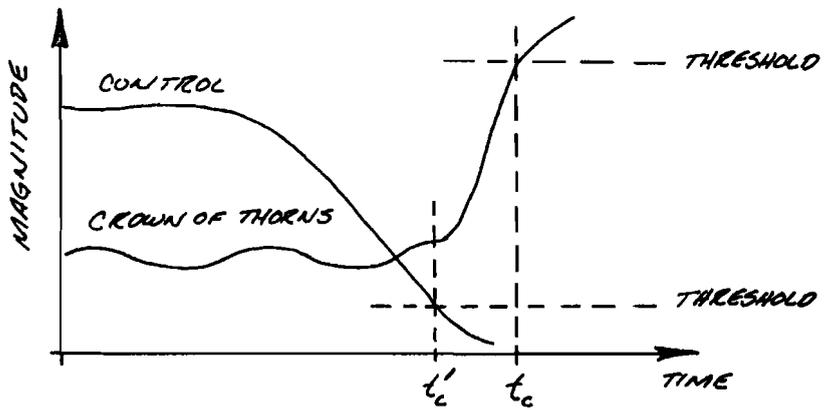


FIGURE 2

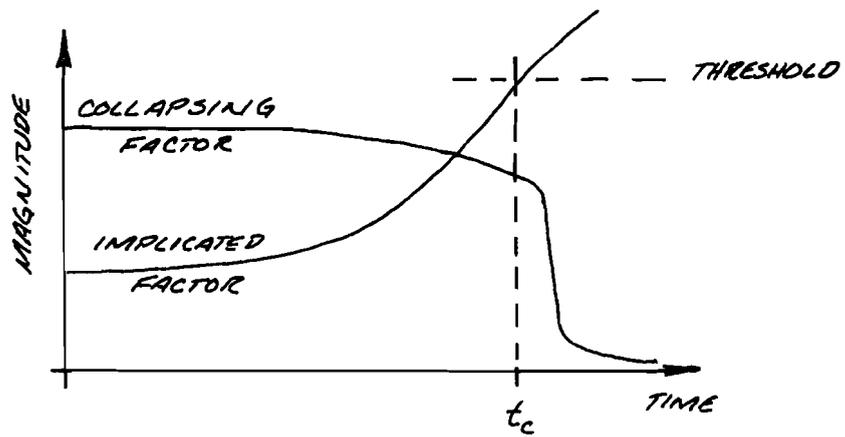


FIGURE 3

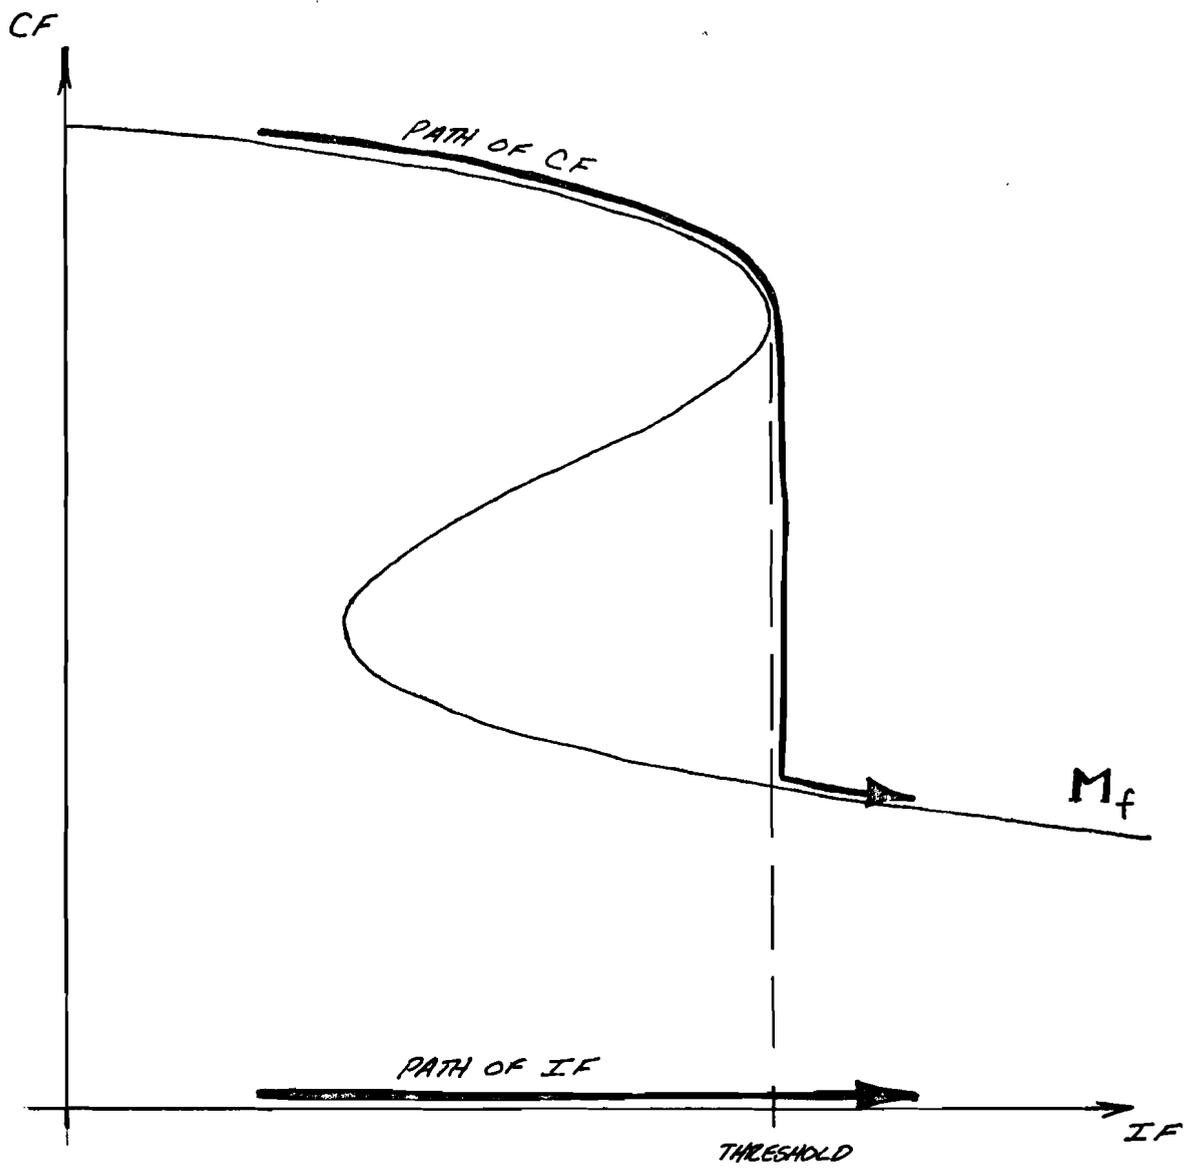


FIGURE A

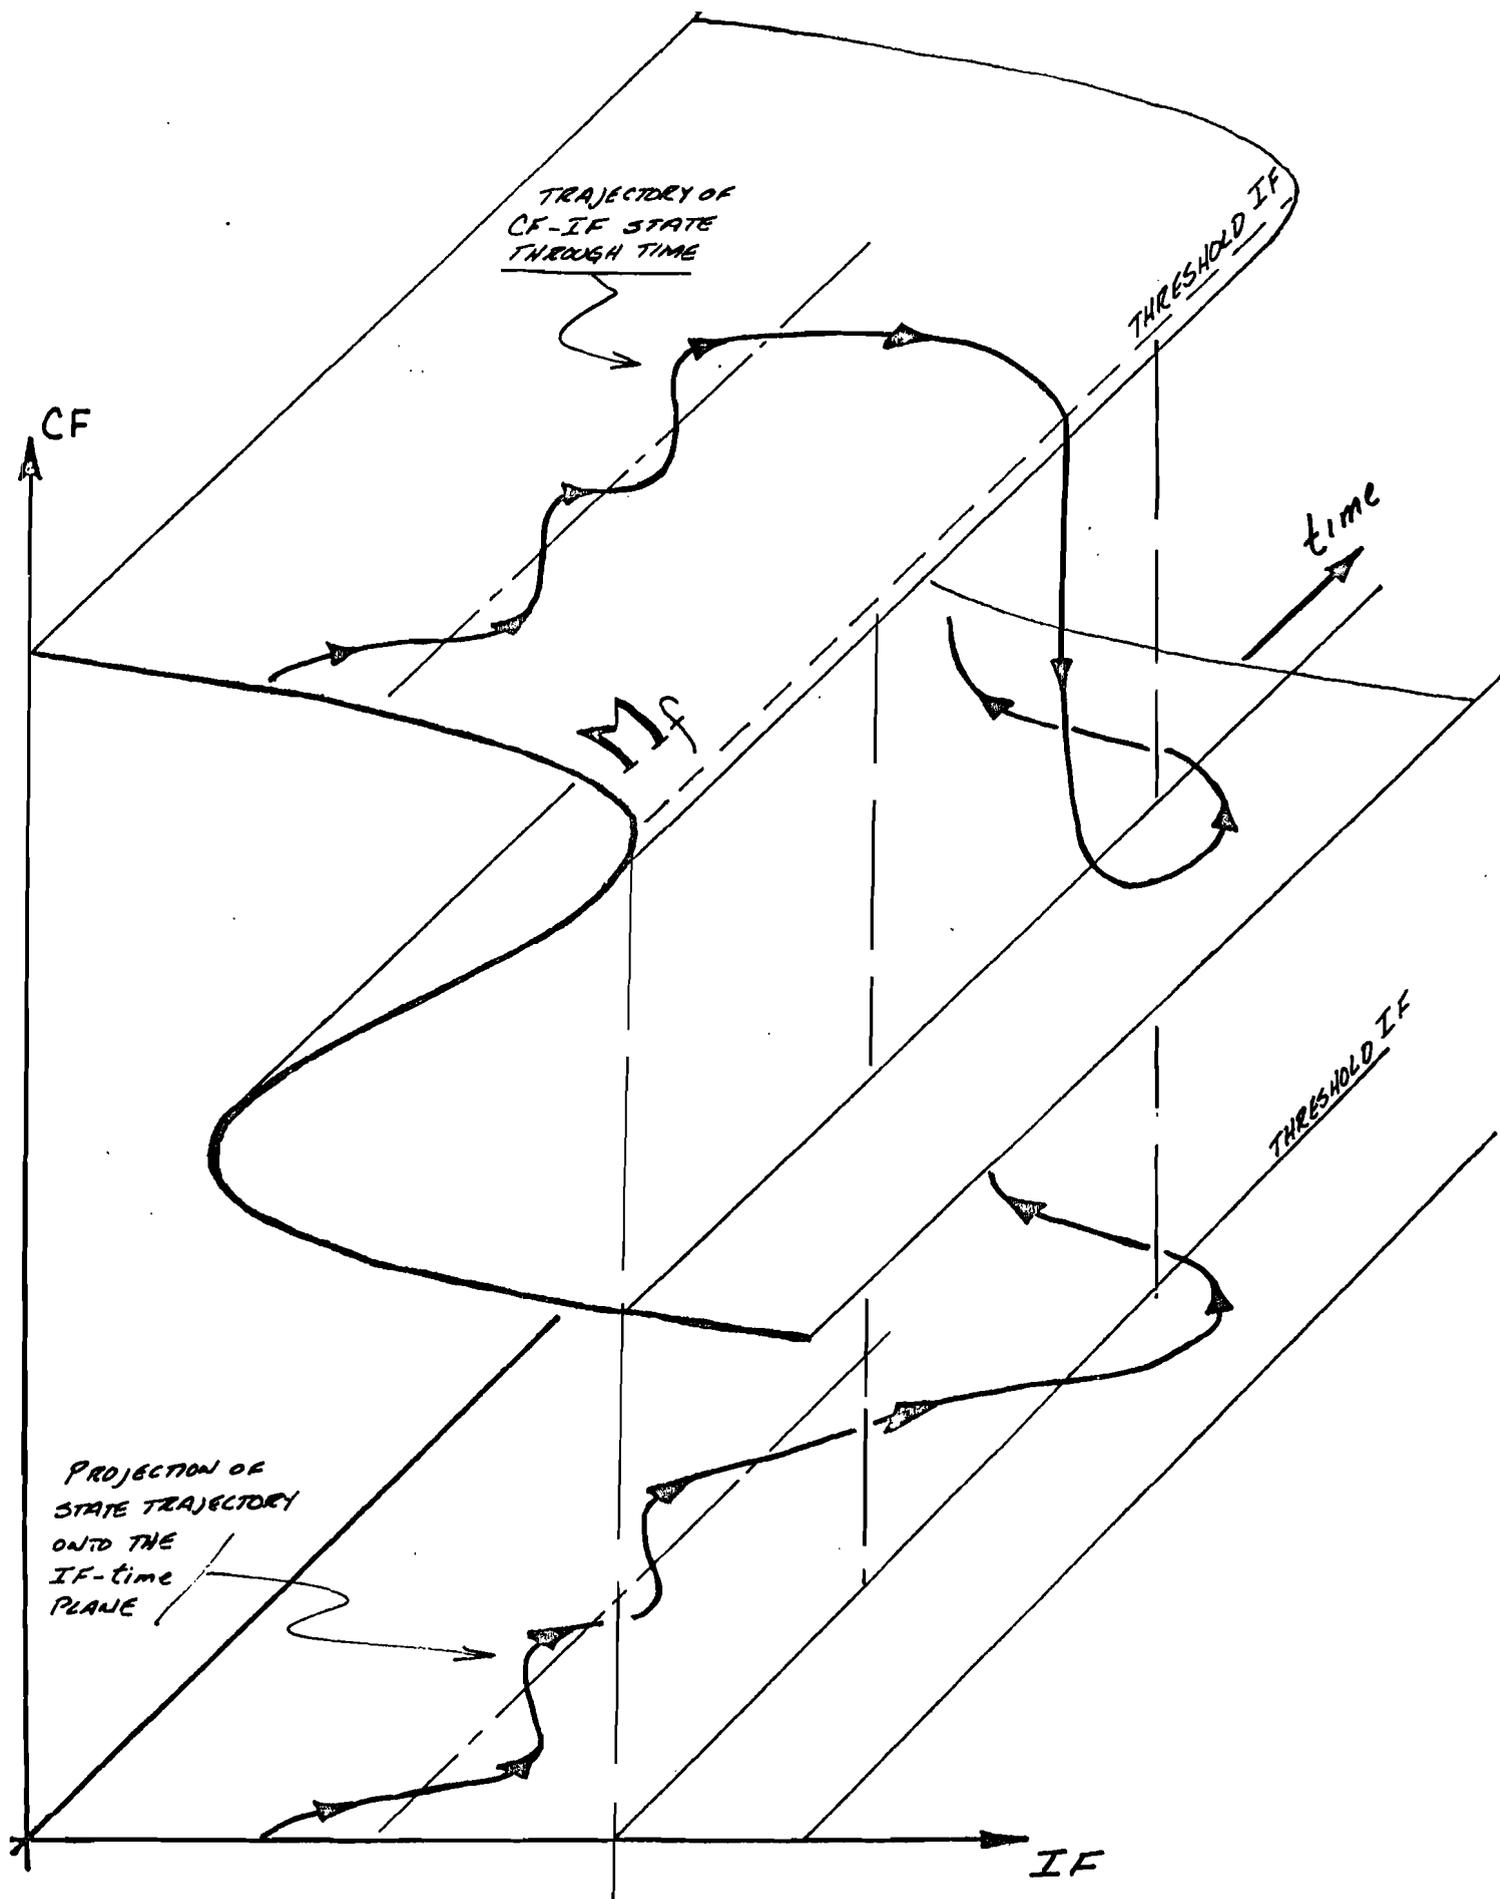


FIGURE 5

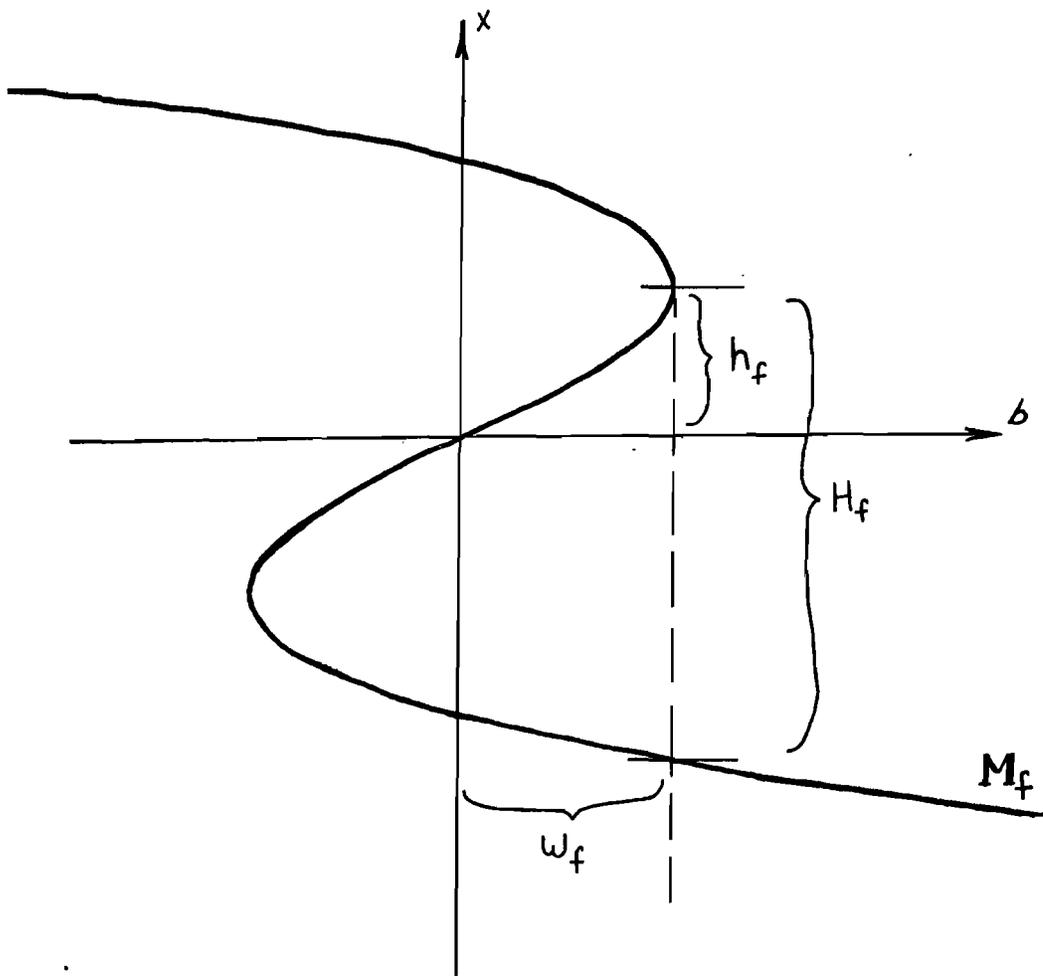


FIGURE 6

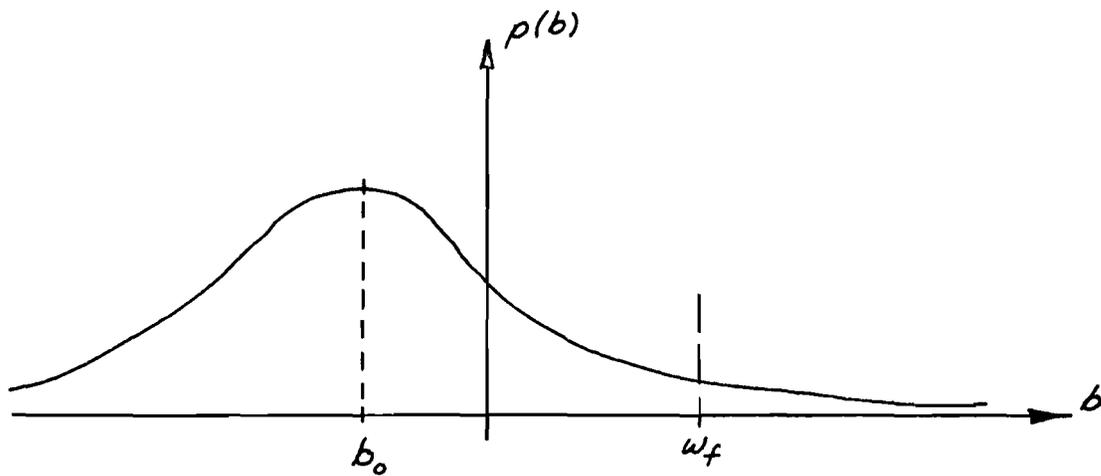


FIGURE 7

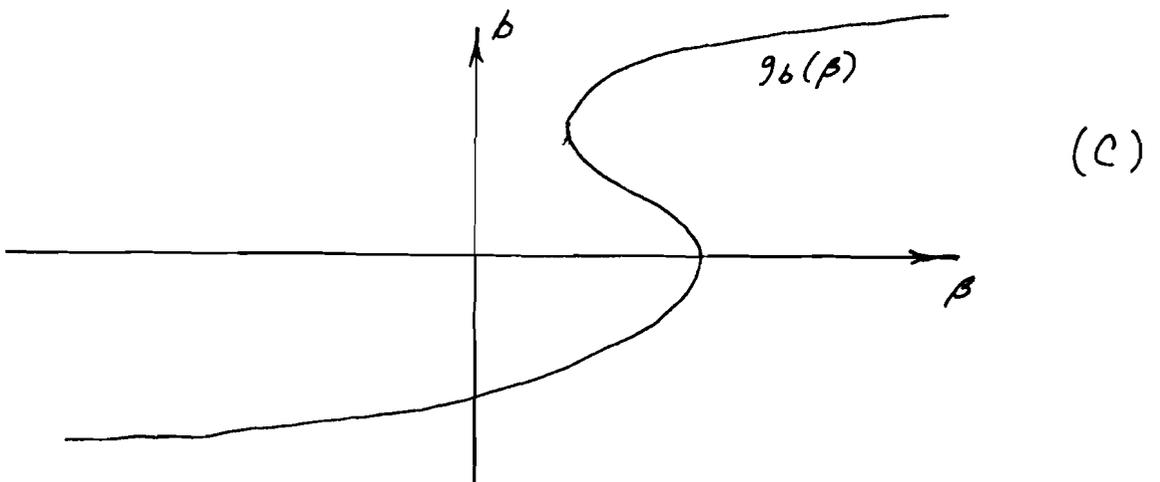
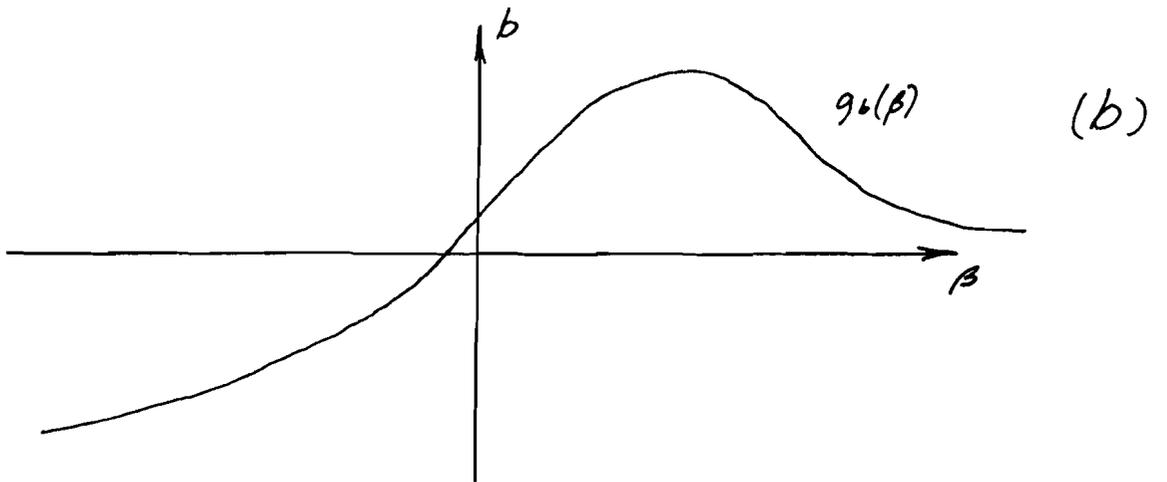
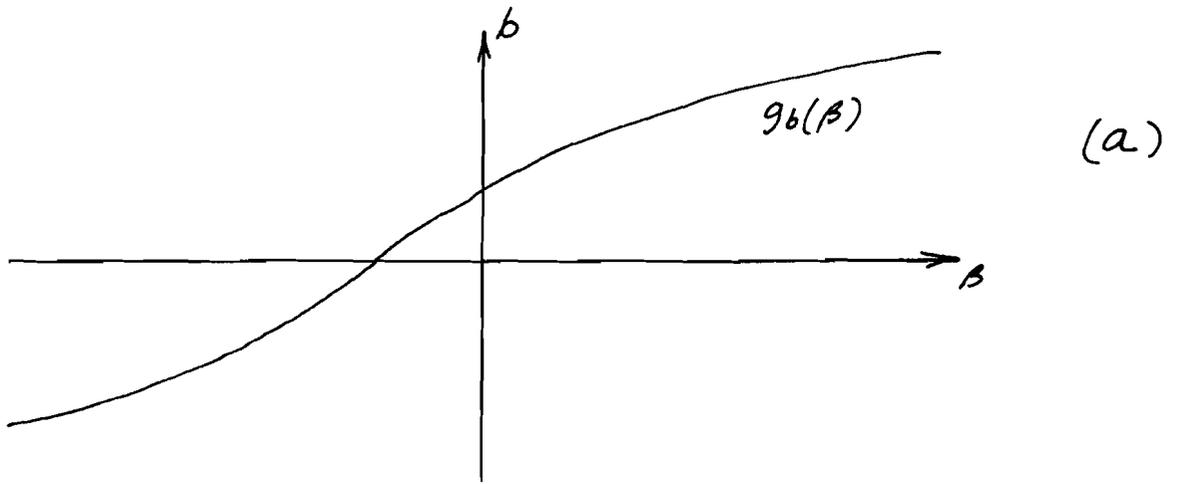


FIGURE 8

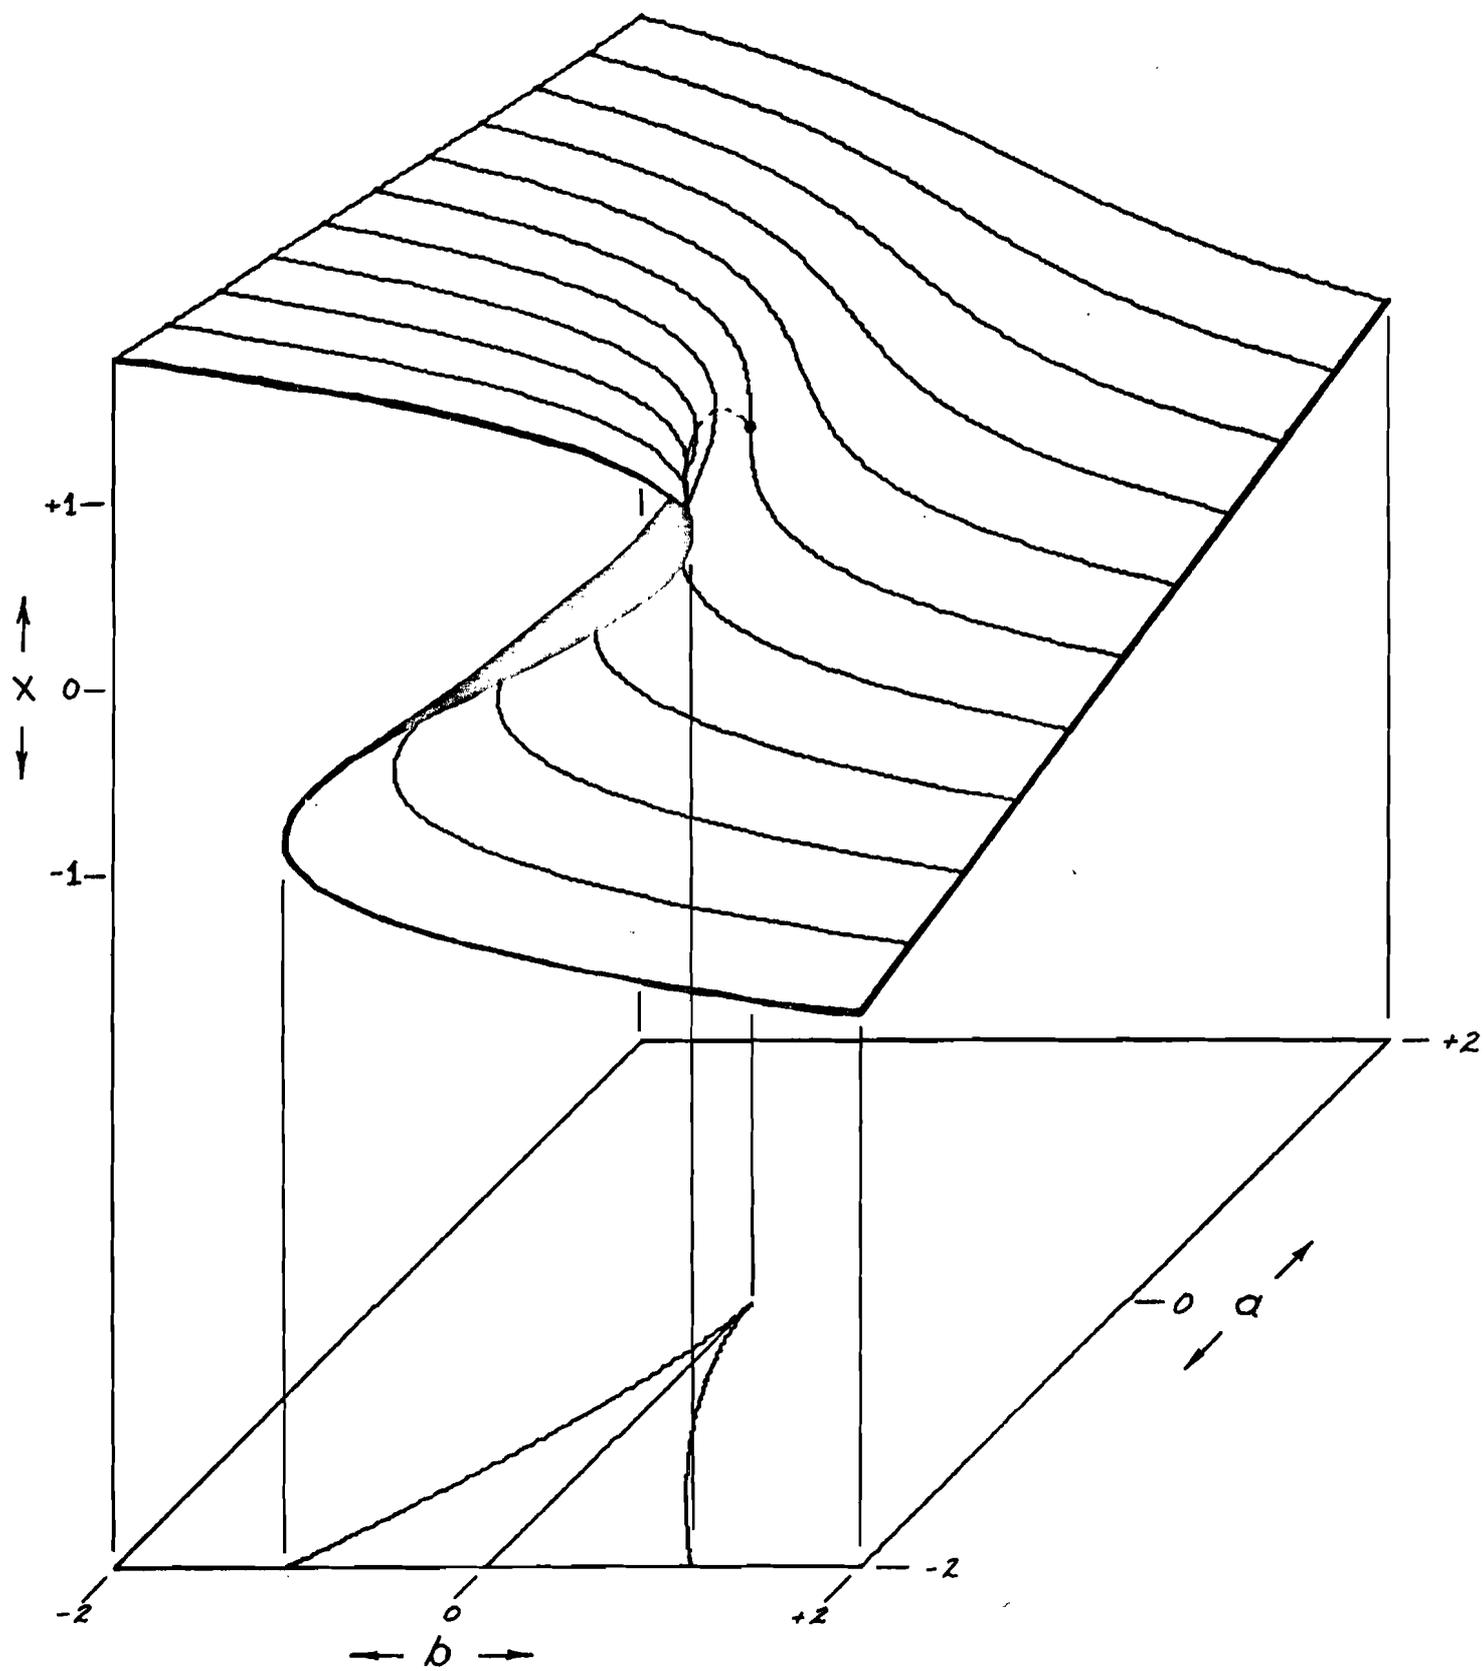


FIGURE 9

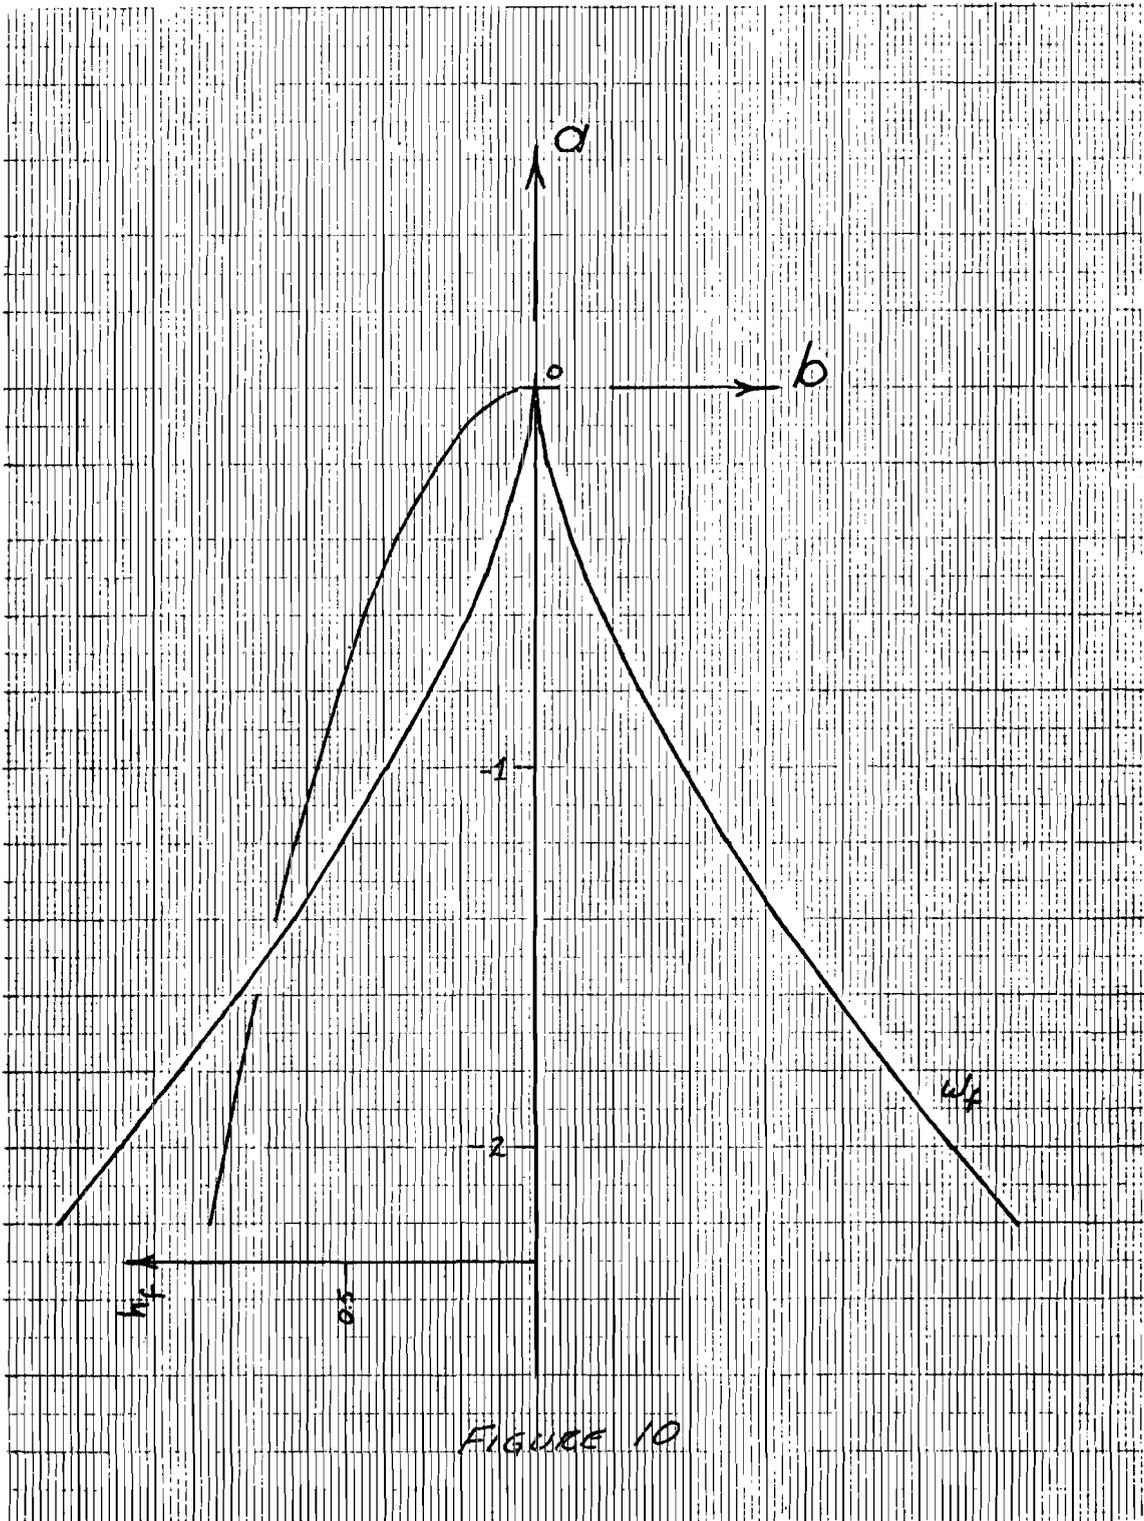


FIGURE 10