Catastrophe Theory

Things that change suddenly, by fits and starts, have long resisted mathematical analysis. A method derived from topology describes these phenomena as examples of seven "elementary catastrophes"

by E. C. Zeeman

Scientists often describe events by constructing a mathematical model. Indeed, when such a model is particularly successful, it is said not only to describe the events but also to "explain" them; if the model can be reduced to a simple equation, it may even be called a law of nature. For 300 years the preeminent method in building such models has been the differential calculus invented by Newton and Leibniz. Newton himself expressed his laws of motion and gravitation in terms of differential equations, and James Clerk Maxwell employed them in his theory of electromagnetism. Einstein's general theory of relativity also culminates in a set of differential equations, and to these examples could be added many less celebrated ones. Nevertheless, as a descriptive language differential equations have an inherent limitation: they can describe only those phenomena where change is smooth and continuous. In mathematical terms, the solutions to a differential equation must be functions that are differentiable. Relatively few phenomena are that orderly and well behaved; on the contrary, the world is full of sudden transformations and unpredictable divergences, which call for functions that are not differentiable.

A mathematical method for dealing with discontinuous and divergent phenomena has only recently been developed. The method has the potential for describing the evolution of forms in all aspects of nature, and hence it embodies a theory of great generality; it can be applied with particular effectiveness in those situations where gradually changing forces or motivations lead to abrupt changes in behavior. For this reason the method has been named catastrophe theory. Many events in physics can now be recognized as examples of mathematical catastrophes. Ultimately, however, the most important applications of the theory may be in biology and the social sciences, where discontinuous and divergent phenomena are ubiquitous and where other mathematical techniques have so far proved ineffective. Catastrophe theory could thus provide a mathematical language for the hitherto "inexact" sciences.

Catastrophe theory is the invention of René Thom of the Institut des Hautes Études Scientifique at Bures-sur-Yvette in France. He presented his ideas in a book published in 1972, *Stabilité Structurelle et Morphogénèse;* an English translation by David H. Fowler of the University of Warwick has recently been published. The theory is derived from topology, the branch of mathematics concerned with the properties of surfaces in many dimensions. Topology is involved because the underlying forces in nature can be described by smooth surfaces of equilibrium; it is when the equilibrium breaks down that catastrophes occur. The problem for catastrophe theory is therefore to describe the shapes of all possible equilibrium surfaces. Thom has solved this problem in terms of a few archetypal forms, which he calls the elementary catastrophes. For processes controlled by no more than four factors Thom has shown that there are just seven elementary catastrophes. The proof of Thom's theorem is a difficult one, but the results of the proof are relatively easy to comprehend. The elementary catastrophes themselves can be understood and applied to problems in the sciences without reference to the proof.

A Model of Aggression

The nature of the models derived from catastrophe theory can best be illustrated by example, and I shall begin by considering a model of aggression in the dog. Konrad Z. Lorenz has pointed out that aggressive behavior is influenced by two conflicting drives, rage and fear, and he has proposed that in the dog these factors can be measured with some reliability. A dog's rage is correlated with the degree to which its mouth is open or its teeth are bared; its fear is reflected by how much its ears are flattened back. By employing facial expression as an indicator of the dog's emotional state we can attempt to learn how the dog's behavior varies as a function of its mood.

If only one of the conflicting emotional factors is present, the response of the dog is relatively easy to predict. If the dog is enraged but not afraid, then some aggressive action, such as attacking, can be expected. When the dog is frightened but is not provoked to anger, aggression becomes improbable and the dog will most likely flee. Prediction is also straightforward if neither stimulus is present; then the dog is likely to express some neutral kind of behavior unrelated to either aggression or submission.

What if the dog is made to feel both rage and fear simultaneously? The two controlling factors are then in direct conflict. Simple models that cannot accommodate discontinuity might predict that the two stimuli would cancel each other, leading again to neutral behavior. That prediction merely reveals the shortcomings of such simplistic models, since neutrality is in fact the least likely behavior. When a dog is both angry and frightened, the probabilities of both extreme modes of behavior are high; the dog may attack or it may flee, but it will not remain indifferent. It is the strength of the model derived from catastrophe theory that it can account for this bimodal distribution of probabilities. Moreover, the model provides a basis for predicting, under particular circumstances, which behavior the dog will choose.

To construct the model we first plot the two control parameters, rage and fear, as axes on a horizontal plane, called the control surface. The behavior of the dog is then measured on a third axis, the behavior axis, which is perpendicular to the first two. We might assume that there is a smooth continuum of possible modes of behavior, ranging, for example, from outright retreat through cowering, avoidance, neutrality, growling and snarling to attacking. The most aggressive modes of behavior are assigned the largest values on the behavior axis, the least aggressive the smallest values. For each point on the control surface (that is, for each combination of rage and fear) there is at least one most probable behavior, which we represent as a point directly above the point on the control surface and at a height appropriate to the behavior. For many points on the control surface, where either rage or fear is predominant, there will be just one behavior point. Near the center of the graph, however, where rage and fear are roughly equal, each point on the control surface has two behavior points, one at a large value on the behavior axis representing aggressive action, the other at a small

value representing submissive action. In addition we can note a third point that will always fall between these two, representing the least likely neutral behavior.

If the behavior points for the entire control surface are plotted and then connected, they form a smooth surface: the behavior surface. The surface has an overall slope from high values where rage predominates to low values in the region where fear is the prevailing state of mind, but the slope is not its most distinctive feature. Catastrophe theory reveals that in the middle of the surface there must be a smooth double fold, creating a pleat without creases, which grows narrower from the front of the surface to the back and eventually disappears in a singular point where the three sheets of the pleat come together [see illustration below]. It is the pleat that gives the model its most interesting characteristics. All the

AGGRESSION IN DOGS can be described by a model based on one of the elementary catastrophes. The model assumes that aggressive behavior is controlled by two conflicting factors, rage and fear, which are plotted as axes on a horizontal plane: the control surface. The behavior of the dog, which ranges from attacking to retreating, is represented on a vertical axis. For any combination of rage and fear, and thus for any point on the control surface, there is at least one likely form of behavior, indicated as a point above the corresponding point on the control surface and at the appropriate height on the behavior axis. The set of all such points makes up the behavior surface. In most cases there is only one probable mode of behavior, but where rage and fear are roughly equal there are two modes: a dog both angry and fearful may either attack or retreat. Hence in the middle of the graph there are two sheets representing likely behavior, and these are connected by a third sheet to make a continuous, pleated surface. The third or middle sheet, shown in gray, has a different significance from the other two sheets: it represents least likely behavior, in this case neutrality. Toward the origin the pleat in the behavior surface becomes narrower, and eventually it vanishes. The line defining the edges of the pleat is called the fold curve, and its projection onto the control surface is a cusp-shaped curve. Because the cusp marks the boundary where the behavior becomes bimodal it is called the bifurcation set and the model is called a cusp catastrophe. If an angry dog is made more fearful, its mood follows the trajectory A on the control surface. The corresponding path on the behavior surface moves to the left on the top sheet until it reaches the fold curve; the top sheet then vanishes, and the path must jump abruptly to the bottom sheet. Thus the dog abandons its attack and suddenly flees. Similarly, a frightened dog that is angered follows the trajectory B. The dog remains on the bottom sheet until that sheet disappears, then as it jumps to the top sheet it stops cowering and suddenly attacks. A dog that is angered and frightened at the same time must follow one of the two trajectories at C. Whether it moves onto the top sheet and becomes aggressive or onto the bottom sheet and becomes submissive depends critically on the values of rage and fear. A small change in the stimuli can produce a large change in behavior: the phenomenon is divergent.

66

points on the behavior surface represent the most probable behavior of the dog, with the exception of those on the middle sheet, which represent least probable behavior. Through catastrophe theory we can deduce the shape of the entire surface from the fact that the behavior is bimodal for some control points.

In order to understand how the model predicts behavior we must consider the reaction of the dog to changing stimuli. Suppose that initially the dog's emotional state is neutral and can be represented by a point at the origin on the control surface. The dog's behavior, given by the corresponding point on the behavior surface, is also neutral. If some stimulus then increases the dog's rage without affecting its fear, the behavior changes smoothly, following the upward trend of the behavior surface, to more aggressive postures; if the rage is increased enough, the dog attacks. If the dog's fear now begins to increase while its rage remains at an elevated level, the point representing its emotional state on the control surface must move across the graph toward the center. The point representing behavior must of course follow, but because the slope of the behavior surface in this region is not steep the behavior changes only slightly; the dog remains aggressive.

As fear continues to increase, however, the behavior point must eventually reach the edge of the pleat. The novel and illuminating properties of the model then become evident. At the edge of the pleat the sheet on which the behavior point has been traveling folds under and is thereby effectively annihilated; with any additional increase in fear the sheet vanishes. The behavior state must therefore fall directly to the bottom sheet of the graph, which represents quite different modes of behavior. The aggressive states of the top sheet are no longer possible; there is no alternative but a sudden, and indeed catastrophic, change to a meeker attitude. The model thus predicts that if an enraged dog is made progressively more fearful, it will eventually break off its attack and retreat. The sudden change in behavior might be called a flight catastrophe.

The graph also predicts the existence of an opposite pattern of behavior: an attack catastrophe. In an initial state dominated by fear the dog's behavior is stabilized on the bottom sheet, but with a sufficient increase in rage it passes the edge of the opposite side of the pleat and jumps up to the top sheet and a more aggressive frame of mind. In other words, a frightened dog, if it is placed in a situation in which rage steadily increases, may suddenly attack.

Finally, consider the behavior of a dog whose mood is initially neutral as its rage and fear are increased simultaneously. The behavior point is initially at the origin and under the influence of the conflicting stimuli it moves straight forward on the graph. At the singularity, however, where the behavior surface becomes pleated, the point must move onto the top sheet as the dog grows

CONTROL FACTORS in the model of aggression are rage and fear, which in dogs can be measured by facial expression. Rage is reflected by the extent to which the mouth is opened, and fear is revealed by the degree to which the ears are flattened back. From these indicators it is possible to judge the dog's emotional state, and through the model to predict its behavior.
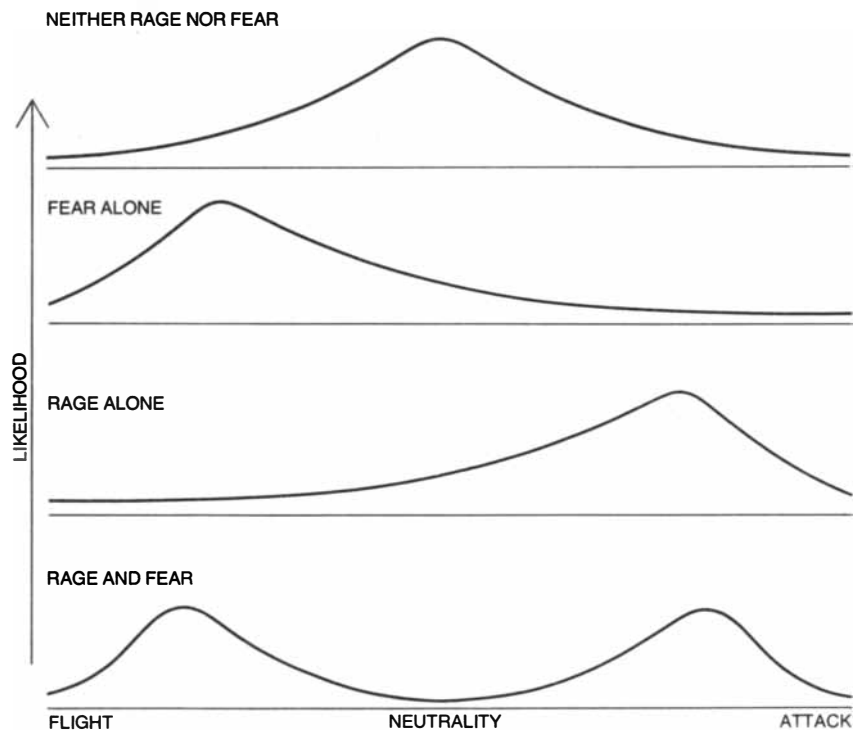
LIKELIHOOD FUNCTION determines the behavior of the dog under the conflicting influences of rage and fear. When neither stimulus is present, the most likely behavior is neutrality; rage alone elicits aggression, fear alone submission. When the dog is made both angry and fearful, the likelihood graph becomes bimodal: attack and flight are both favored, and neutrality is the least likely response. The bimodality is reflected in the behavior surface of the cusp catastrophe, which has two sheets representing most likely behavior where both stimuli are present.

67

more aggressive or onto the bottom one as it becomes less aggressive. Which sheet is selected depends critically on the dog's state of mind just before it reaches the singularity. The graph is said to be divergent: a very small change in the initial conditions results in a large change in the final state.

The Cusp Catastrophe

The line that marks the edges of the pleat in the behavior surface, where the top and bottom sheets fold over to form the middle sheet, is called the fold curve. When it is projected back onto the plane of the control surface, the result is a cusp-shaped curve. For this reason the model is called the cusp catastrophe. It is one of the simplest of the seven elementary catastrophes, and so far it has been the most productive.

The cusp on the control surface is called the bifurcation set of the cusp catastrophe, and it defines the thresholds where sudden changes can take place. As long as the state of the system remains outside the cusp, behavior varies smoothly and continuously as a function of the control parameters. Even on entering the cusp no abrupt change is

observed. When the control point passes all the way through the cusp, however, a catastrophe is inevitable.

Everywhere inside the bifurcation set there are two possible modes of behavior; outside it there is only one mode. Moreover, in the cusp there are just two modes of behavior even though the behavior surface there has three sheets. That is because the middle sheet in the pleated region is made up of points representing least probable behavior. The middle sheet is included in the graph primarily so that the behavior surface will be smooth and continuous; the behavior point never occupies the middle sheet. Indeed, there is no trajectory on the control surface that could bring the behavior point onto the middle sheet. Whenever the fold curve is crossed, the point jumps between the top and bottom sheets; the middle sheet is said to be inaccessible.

The construction of this model began with an essentially determinist hypothesis: that the behavior of the dog could be predicted from its emotional state, as reflected in its facial expression. The bimodality of the resulting graph may seem at first to undermine that hypothesis, since the exis-

tence of two possible modes of behavior for a given emotional state makes unambiguous prediction impossible. Indeed, it is true that if we know only the present emotional state (and if that state falls within the bimodal region of the graph), we cannot predict what the dog will do. The determinism of the model is restored, however, and the model is made more sophisticated, when we consider an additional factor as we make predictions. The behavior of the dog can be predicted if we know both its present emotional state and the recent history of its emotions. It should come as no surprise that the effects of frightening an enraged dog are different from those of angering a frightened dog.

Aggressiveness is not, of course, a uniquely canine trait, and the model describes a mechanism that might operate in other species as well. Consider, for example, the territorial behavior of certain tropical fishes that establish permanent nesting sites on coral reefs. The parameters controlling aggression might in this case be the size of an invading fish and proximity to the nest; the behavior is once again described by a cusp catastrophe. A fish foraging far from
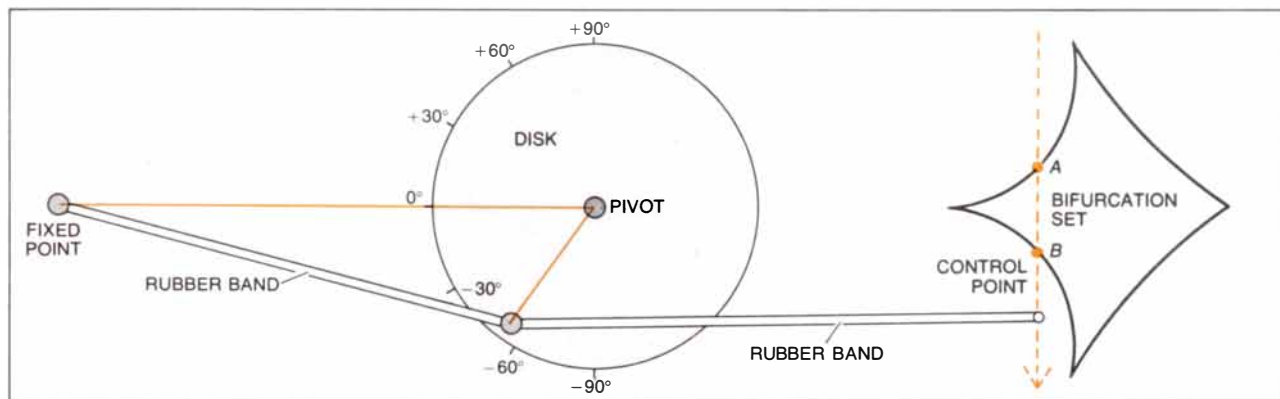
CATASTROPHE MACHINE invented by the author exhibits discontinuous behavior that can be described by a cusp catastrophe. The machine consists of a cardboard disk pivoted at its center, with two rubber bands attached at a point near the perimeter. The unstretched length of each rubber band is approximately equal to the diameter of the disk. The free end of one rubber band is fixed to the mounting board and the machine is operated by moving the other rubber band, the free end of which is designated the control point. The behavior measured is the angle formed by the fixed point, the pivot and the point at which the two rubber bands are attached to the disk. Many

movements of the control point cause only smooth rotation of the disk, but in some cases the disk swings suddenly from one side to the other. If all the positions of the control point at which such sudden movements take place are marked, a diamond-shaped curve is generated. The curve is made up of four cusps, each forming the bifurcation set of a cusp catastrophe. Moving the control point along the trajectory shown in color, there is no movement at point A, but the disk turns suddenly at B. If the path of the control point is reversed, it passes B uneventfully, but the disk moves when the control point reaches A. Inside the cusp there are two stable positions of the disk.

ENERGY FUNCTION governs the behavior of the catastrophe machine. The machine tends always to assume a position of minimum energy, that is, the disk rotates to minimize the tension on the rubber bands. When the control point is outside the bifurcation set, there is only one position of minimum energy, corresponding to the one stable position of the disk. As the control point is moved across the bifurca-

tion set a second local minimum develops at A, and eventually it becomes deeper than the original one. The machine cannot shift to the new local minimum, however, because it is separated from it by a local maximum. Only when the control point crosses the second line of the cusp at B is the local maximum eliminated; the equilibrium then breaks down and the machine moves suddenly to the new minimum.

its nest would flee on meeting a larger adversary; once it reached the "defense" perimeter of its own territory, however, its attitude would suddenly change and it would turn to defend its nest. Conversely, if the fish were threatened in its nest, it would chase the invader, but only until it reached the "attack" perimeter of its own territory, where it would abandon the chase and return to its nest. The distance from the nest at which the behavior would change would be determined by the cusp lines of the bifurcation set. Because of the shape of the cusp, the model makes the interesting prediction that the "defense" perimeter is smaller than the "attack" perimeter. Moreover, the size of both perimeters depends on the size of the adversary; a larger invader could approach the nest more closely before the fish would be provoked to fight. The model readily accounts for an observed feature of fish behavior: in mating pairs of territorial fish the partner that happens to be closer to the nest offers the more vigorous defense.

Models of Human Behavior

The cusp catastrophe also provides an interpretation of certain kinds of human behavior. For example, an argument often involves a display of aggression, and its progress is strongly influenced by anger and fear. A cusp catastrophe can be constructed with these emotions as the control parameters and with the intensity of the conflict as the behavior axis.

At the origin on the behavior surface is the most unemotional behavior: rational discussion. As anger and fear increase, the behavior point moves forward on the graph into the pleated region of the surface, where the behavior is bimodal. The opponents are then denied access to sober discourse, and they must either make stronger assertions or make concessions. With an additional escalation of emotion the available modes of behavior diverge further, and the alternatives are invective or apology; finally the opponents must choose between fury and tears. Once the argument has become a heated one a small increase in either anger or fear can cause an abrupt shift in behavior. An aggressive advocate who begins to waver in his opinion may abandon his position and apologize; a timid opponent forced to make repeated concessions may suddenly lose his temper and become truculent. The model even suggests a strategy for effective persuasion. If an argument is likely to induce both anger and fear, then it is best to state one's case and leave, allowing emotions to subside and enabling one's opponent to regain access to rational thought.

Another human behavioral pattern that can be described by the cusp model is self-pity and the catharsis that sometimes relieves it. In this case the controlling parameters, analogous to anger and fear, are the less extreme emotions frustration and anxiety. Moreover, the behavior axis measures not overt behavior, which in animals is the

CUSP MODEL of the catastrophe machine erects a pleated behavior surface over one segment of the bifurcation set, such as the cusp nearest the disk. Each point on the top and bottom sheets of the behavior surface gives the position of the disk having minimum energy for that position of the control point. Within the bifurcation set, where there are two stable positions of the disk, there are likewise two local minimums, one on the top sheet and the other on the bottom sheet. The middle sheet represents the local maximum in the energy function. Catastrophic changes in angular position are observed whenever the control point moves all the way across the cusp.

only kind that can be observed, but underlying moods, which in man can be identified directly. A typical range of moods might extend from anger and annoyance through neutral moods to dejection and self-pity.

In the model a large increase in anxiety induces a persistent mood of self-pity; the point representing the mood is trapped on the bottom sheet of the pleated behavior surface [see top illustration on page 75]. Self-pity is a defensive attitude commonly adopted by children, and it often seems that sympathy is powerless to alleviate it. A sarcastic remark, on the other hand, may provoke a sudden loss of temper and by releasing tension may open a pathway back to a less emotional state. It is unfortunate that sarcasm should succeed where sympathy fails, but the cause of that irony is apparent in the model. The sarcasm brings an increase in frustration, and as a result the point representing mood travels across the behavior surface as far as the fold curve; having reached the extremity of the bottom sheet, it is forced to make a catastrophic jump to the top sheet, and self-pity is transformed into anger.

These examples of the cusp catastrophe offer an interesting and apparently successful model of certain modes of animal and human behavior, but it is a phenomenological model only; it cannot yet be said to explain the behavior. The question of why a dog or a fish behaves as it does has not been answered but merely recast at a level of greater abstraction. We must now ask why the model works. In particular, why does the behavior point follow the surface to the edge of a pleat, then catastrophically jump to another sheet? Why does it not tunnel smoothly from one surface to another so that there is a gradual transition? What mechanism holds the state of the system on the behavior surface? The answers to these questions can best be approached through another example of the cusp catastrophe, one dealing with the behavior of a much simpler system than a dog, a fish or a man.

A Catastrophe Machine

Elementary catastrophes can be generated at will with a simple device made from stiff cardboard, rubber bands and a few other materials. The heart of the machine is a disk of cardboard pivoted at its center and with two rubber bands attached at a point near the perimeter. The free end of one rubber band is attached at a point outside the disk; the other rubber band serves to control the motion of the disk, and the position of its free end is designated the control point [see upper illustration on opposite page].

The catastrophe machine is operated by

69

moving the control point in the plane of the disk. In many cases the result is smooth rotation of the disk; in one region, however, in the vicinity of a point diametrically opposite the anchorage of the fixed rubber band, smooth movement of the control point can cause abrupt motion of the disk. If the position of the control point is marked each time the disk jumps, a concave, diamond-shaped curve is generated. This curve is made up of four connected cusps, the bifurcation sets of four cusp catastrophes.

If we consider only one of the four cusps, such as the one closest to the disk, the corresponding behavior surface can be constructed by arranging the fold curve so that it lies directly over the cusp. For any position of the control point outside the cusp the behavior surface has only one sheet and the disk has only one stable position. If the control point is inside the cusp, the behavior surface has three sheets, but again the middle one is to be excluded because it corresponds to an unstable equilibrium. As a result there are two stable positions for the disk. That the behavior of the machine conforms to this model can be verified by moving the control point from left to right across the graph. The disk moves smoothly and only slightly until the control point reaches the right edge of the cusp; the disk then suddenly turns as the behavior point falls off the extremity of the bottom sheet and jumps to the top one. When the path of the control point is reversed, the point crosses the right edge of the cusp uneventfully and the disk continues to move smoothly until the left edge of the cusp is reached; this time the behavior point jumps from the top sheet to the bottom one.

In the catastrophe machine the cause of this behavior is readily discovered: it is the tendency of all physical systems in which friction is important to assume a state of minimum energy. The energy to be minimized is the potential energy stored in the rubber bands, and the disk therefore rotates until the tension on the two rubber bands is at a minimum. At that position the machine is in stable equilibrium. Unless energy is added to the system the machine must remain at the equilibrium point; the process that keeps it there is called the dynamic.

The operation of the dynamic can be demonstrated by a series of graphs, each one showing for a single position of the control point the energy of the machine for all possible rotations of the disk [see lower illustration on page 68]. As long as the control point is outside the cusp the graph is a smooth curve with a single trough, or minimum, and the state of the machine will always move swiftly to the state of minimum energy at the bottom of the trough. As the control point enters the interior of the cusp a second trough, or local energy minimum, develops next to the original one. This second trough gradually grows deeper, but the machine cannot enter it because the two troughs are separated by a small peak, a local energy maximum. The state of the machine does not change until the second trough coalesces with the local maximum. This takes place as the control point crosses the second cusp line, and the state of the machine is then carried swiftly by the dynamic to a new unique position of minimum energy.

The significance of the dynamic becomes apparent when it is revealed that the behav-ior surface of the cusp catastrophe is the graph of all the minimums and maximums of the energy function. Outside the cusp there is a single energy minimum and there are no maximums; the behavior surface therefore has a single sheet. At the cusp line a new local minimum and maximum are created, and consequently two new sheets form in the behavior surface. The state of the machine can never lie stably on the middle sheet because that is a position of maximum energy.

The mathematical procedure for drawing the behavior surface comes from elementary calculus: the behavior surface is a graph of all the points where the first derivative of the energy function is equal to zero. It is not necessary to understand how this operation is performed; it is sufficient to know that the first derivative is zero wherever the graph of the energy function is horizontal (where its slope is zero). It is horizontal only at the minimums and maximums and at inflection points. The minimums form the stable top and bottom sheets, the maximums form the unstable middle sheet and the inflection points form the fold curve that marks the boundaries of the sheets.

The behavior associated with the cusp farthest from the disk could be analyzed in the same way that we have explained the nearer cusp. The two cusps at the sides, however, differ in a crucial respect: for them the energy function is inverted, so that inside the cusp there are two points—corresponding to two positions of the disk—with maximum energy and only one point with minimum energy. The dynamic compels the machine to remain in the single stable position of minimum energy. On the behavior surface also the positions of minimums and maximums are reversed: the middle sheet represents stable energy minimums, the top and bottom sheets represent unstable maximums. The behavior point therefore can lie only on the middle sheet. The graph is called a dual cusp catastrophe.

The Role of the Dynamic

The success of catastrophe theory in accounting for the behavior of the catastrophe machine can now be explained. The crucial concept is that of the dynamic, which has two functions. First, it holds the behavior point firmly on the top or bottom sheet of the behavior surface. If the disk is turned by hand against the force of the rubber bands and is then released, the dynamic brings it sharply back to equilibrium, that is, it brings the behavior point back onto the surface. Second, when the behavior point crosses the fold curve, it is the dynamic that causes the catastrophic jump from one sheet to the other.

The same principles can be applied to the psychological models considered above. The likelihood functions of these models are analogous to the energy function in the catastrophe machine, except that the roles of minimums and maximums are reversed. The top and bottom sheets of the behavior
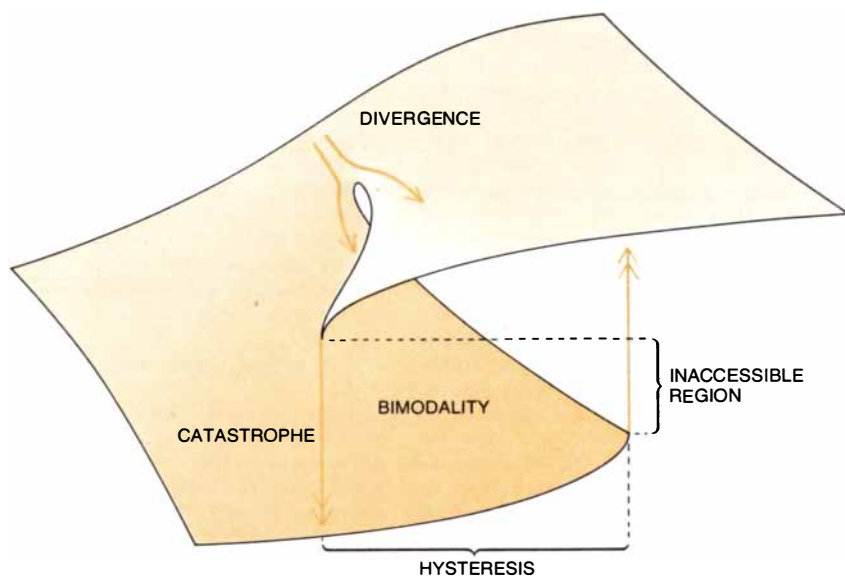
FIVE PROPERTIES characterize phenomena that can be described by the cusp catastrophe. The behavior is always bimodal in some part of its range, and sudden jumps are observed between one mode of behavior and the other. The jump from the top sheet of the behavior surface to the bottom sheet does not take place at the same position as the jump from the bottom sheet to the top one, an effect called hysteresis. Between the top and bottom sheets there is an inaccessible zone on the behavior axis; the middle sheet, representing least likely behavior, has been omitted for clarity. Finally, the cusp catastrophe implies the possibility of divergent behavior.

surface are made up of all the points representing maximum likelihood and the middle sheet is made up of those representing minimum likelihood. An important question remains: What is the dynamic? In the model of aggression what compels the dog to express the most likely behavior, and in the model of self-pity why is the "most likely mood" the one that is adopted?

An energy minimum in a physical system such as the catastrophe machine is a special instance of a concept called an attractor. In this case it is the simplest kind of attractor, a single stable state, and its effect is like that of a magnet: everything within its range of influence is drawn toward it. Under the influence of the attractor the system assumes a state of static equilibrium.

In the psychological models there must also be attractors, although they need not be as simple as this one. The attractor of a system that is in dynamical equilibrium consists of the entire stable cycle of states through which the system passes. For example, a bowed violin string repeats the same cycle of positions over and over at its resonant frequency, and that cycle of positions represents an attractor of the bowed string.

In the psychological models the obvious place to seek attractors is in the neural mechanisms of the brain. Of course the brain is far more complicated and less well understood than a violin string, but its billions of neurons are known to be organized in large, interconnected networks that form a dynamical system; the equilibrium states of any dynamical system can be represented by attractors. Some of these attractors may be single states, but others are more likely to be stable cycles of states, or higher-dimensional analogues of stable cycles. As various parts of the brain influence one another, the attractors appear and disappear, in some cases rapidly and in others slowly. As one attractor gives way to another the stability of the system may be preserved, but often it is not; then there is a catastrophic jump in the state of the brain.

Thom's theory states that all possible sudden jumps between the simplest attractors—points of static equilibrium—are determined by the elementary catastrophes. Thus if the brain dynamic had only point attractors, it could exhibit only the elementary catastrophes. That is not the case; one obvious item of evidence for more complicated attractors is the alpha rhythm of brain waves, a cyclic attractor. The rules governing jumps between cyclic attractors and higher-dimensional ones are not yet known; they must include not only elementary catastrophes but also generalized catastrophes, and their study is today an active area of research in mathematics. Hence as yet there is no complete theory for the description of all brain dynamics. Nevertheless, the elementary catastrophes provide meaningful models of some brain activities. The models are explicit and sometimes disarmingly simple, but the powerful mathematical theory on which they are based implicit-
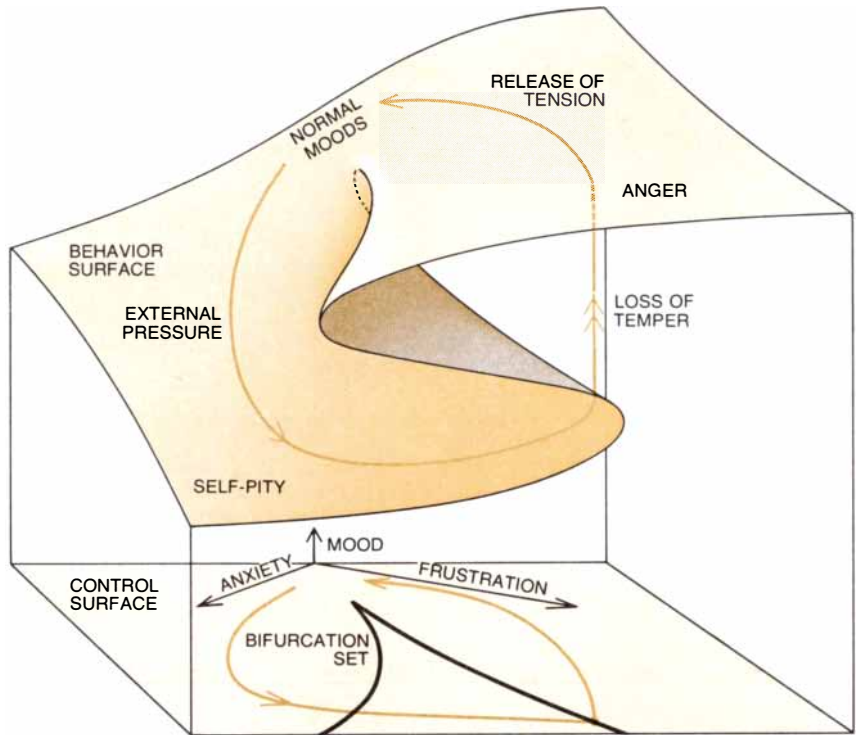
CATHARTIC RELEASE FROM SELF-PITY is described by a cusp catastrophe in which anxiety and frustration are conflicting factors influencing mood. Self-pity is induced by an increase in anxiety; it can be relieved by some event, such as a sarcastic remark, that causes an increase in frustration. As the control point crosses the cusp the mood changes catastrophically from self-pity to anger; the resulting release of tension gives access to calmer emotional states.

BEHAVIOR OF THE STOCK MARKET is described by a model in which the controlling parameters are excess demand for stock and the proportion of the market held by speculators as opposed to that held by investors. The behavior itself is measured by the rate at which the index of stock prices is rising or falling. The control factors are oriented not as conflicting factors but as normal and splitting factors. A fall from the top sheet to the bottom represents a crash; the slow recovery is effected through feedback of the price index on the control parameters.

75

ly allows for the complexity of the underlying neural network.

The concept of an attractor of the brain dynamic provides what is needed in our models of human and animal behavior. The neural mechanism responsible for a mood such as self-pity is not known, but the existence of the mood as a stable state implies that that mechanism is an attractor. Indeed, in the model of self-pity every point on the behavior surface corresponds to an attractor for the system in the brain that determines mood. If that neural system is disturbed in some way, it is quickly returned, under the influence of an attractor, to the behavior surface, just as the catastrophe machine returns to equilibrium. Abrupt changes in mood are encountered when the stability of an attractor breaks down, allowing the mood-determining system to come under the influence of another attractor, toward which it immediately moves.

Through this hypothetical mechanism catastrophe theory provides a model not only of expressed behavior but also of the activity of the brain that directs the behavior. The model is probably most appropriate to primitive regions such as the midbrain, where the networks are highly interconnected and therefore may act as a unit. (In the phylogenetically younger cerebral cortex, patterns of activity are much more complex.) The psychological models we have considered are largely concerned with emotion or mood, and it is thought that the part of the midbrain called the limbic system is primarily responsible for generating mood.

Features of the Cusp Model

The preceding examples and analysis suggest several features common to all cusp catastrophes. One invariant characteristic is that the behavior is bimodal over part of its range, and sudden changes are observed from one mode of behavior to the other. In addition, the pattern of the sudden changes exhibits the effect called hysteresis, that is, the transition from the top sheet to the bottom one does not take place at the same point as the transition from the bottom sheet to the top one. The sudden change does not come at the middle of the cusp but is delayed until the bifurcation set is reached. Another characteristic is that inside the cusp, where the behavior becomes bimodal, the middle zone on the behavior axis becomes inaccessible. Finally, the model implies the possibility of divergence, so that a small perturbation in the initial state of the system can result in a large difference in its final state. These five qualities—bimodality, sudden transitions, hysteresis, inaccessibility and divergence—are related to one another by the model itself. If any one of them is apparent in a process, the other four should be looked for, and if more than one is found, then the process should be considered a candidate for description as a cusp catastrophe.

One process where the five qualities can be detected is the development of hostilities between nations, a situation with obvious analogies to the models of aggression and argument. The control parameters in those models were rage and fear; here we substitute threat and cost. The behavior axis describes the possible actions of the nation, ranging from full-scale attack to lesser military responses such as blockade through neutrality to appeasement and surrender. In a situation where threat and cost are both high public opinion often becomes bimodal as the nation is divided into "doves" advocating surrender and "hawks" advocating attack. The dynamic in the model is the sensitivity of the government to the will of its constituency; the government continuously adapts its policy in order to increase its support, and therefore it remains on the behavior surface. From the model we can deduce and recognize the possible catastrophes. A threatened nation may make repeated concessions, but there is a limit beyond which further threat elicits a sudden declaration of war. Conversely, as costs rise a nation may escalate a war, but there is a limit beyond which further costs may result in a sudden surrender. Hysteresis is recognizable in the delays observed before declaring war or surrendering. The inaccessible region of the behavior axis is the middle zone representing negotiation or compromise. Finally, divergence can be observed in a conflict between two equally strong nations, in which the distribution of public opinion is similar, but the response of the governments is quite different, one becoming increasingly aggressive as the other grows more submissive.

Another candidate for analysis as a cusp catastrophe is the behavior of the stock market, where the terms "bull market" and "bear market" suggest an obvious bimodality. Moreover, a crash, or collapse, of the market is readily explained as a catastroph-

BUCKLING OF AN ELASTIC BEAM is controlled by load and compression, which are respectively normal and splitting factors in a cusp catastrophe. If the beam is flat, an increase in compression forces it to buckle upward or downward. If it buckles upward, a subsequent increase in load drives the control point across the cusp, causing a catastrophic downward motion.

76

ic jump from one sheet of the behavior surface to the other.

In the construction of the model a small modification is necessary. In this case the control axes do not diverge on each side of the cusp, as they do in all the previous examples; instead one axis comes straight forward on the graph, bisecting the cusp, and the other is perpendicular to the cusp [see bottom illustration on page 75]. The parameter that bisects the cusp is called a splitting factor, since increasing it causes a progressively larger divergence between the top and bottom sheets. The other factor is the normal factor, since at the back of the behavior surface the behavior increases continuously with it.

In the stock-market model the normal factor is excess demand for stock; the splitting factor is perhaps more difficult to identify, but it might be related to the amount of stock held by speculators compared with that held by long-term investors. The behavior axis is best measured by the rate of change of the index of stock prices. A market with a rising index is a bull market, and its behavior point is on the top sheet; a falling index, or bear market, places the behavior point on the bottom sheet.

The mechanism of the crash can now be understood. A market with some excess demand and a high proportion of speculators is a bull market, on the top sheet of the behavior surface. A crash can be precipitated by any event that reduces demand enough to push the behavior point over the fold curve. The larger the share of the market held by speculators, the severer the crash. One is immediately prompted to ask why the subsequent recovery is usually slow, and why there is no "upward crash" from a bear market to a bull one. The probable answer is that the behavior axis (the rate of change of the index) has an influence on the control parameters through a feedback loop. A falling, bear market discourages speculation, but after a while the resulting undervaluation encourages long-term investment; as a consequence after a crash the splitting factor is reduced and the market moves back on the graph to the region where the behavior surface is no longer bimodal. As confidence grows and produces excess demand the index rises, but slowly and smoothly, without catastrophes. Speculation is now encouraged and investment discouraged, and the stage is set for another cycle of boom and bust.

As was pointed out above, most applications of catastrophe theory have been in biology and the human sciences, where other modeling techniques are often uninformative, but there are many situations in physics (a science with a highly developed mathematical language) where the theory can contribute to understanding. One such instance is the transition between the liquid phase and the gaseous phase of matter. We can rewrite the equations of J. D. van der Waals as a cusp catastrophe where temperature and pressure are conflicting control

PHASE TRANSITIONS between the liquid and the gaseous state of matter conform to a modified cusp model in which temperature and pressure are the control factors. Ordinarily both boiling and condensation take place at the same values of temperature and pressure. Thus there are catastrophic changes, but there is no hysteresis. Under special circumstances, however, a vapor can be cooled below its dew point and a liquid can be heated above its boiling point, so that the behavior surface is followed all the way to the fold curve. The critical point, where liquid and gas exist simultaneously, is represented by the singularity where the pleat disappears.

factors and density is the behavior axis. The top sheet is then the liquid phase and the bottom sheet the gaseous phase; the two catastrophes represent boiling and condensation. The vertex of the cusp is the critical point, where liquid and gas exist simultaneously. By going around the back of the cusp a liquid can be converted into a gas without boiling.

Under exceptional circumstances the physical system can be made to follow the behavior surface all the way to the edge of each sheet. With due care, for example, liquid water can be heated well beyond its normal boiling point and water vapor can be cooled below its dew point before the phase transitions take place. Such superheating and supercooling are employed in the bubble chambers and cloud chambers used to detect subatomic particles. Ordinarily, however, a substance boils and condenses at the same temperature and pressure, so that a "cliff" forms in the behavior surface, cutting across the middle of the pleat [see illustration above]. The formation of the cliff is explained by a rule called Maxwell's convention, and it reflects the fact that the model is a statistical one, averaging the behavior of many particles.

Another cusp catastrophe in physics, which is derived from the work of Leonhard Euler in the 18th century, is the buck-

ling of an elastic beam under horizontal compression and a vertical load. The compression is a splitting factor, the load a normal factor. An increase in compression causes the behavior point to move forward on the graph, into the region of the cusp, where the beam has two stable states, one buckled upward and the other downward. If the initial buckling is upward, increasing the load can move the behavior point across the cusp, causing the beam to snap suddenly downward. The effect can be observed with a piece of cardboard held between the fingers. When it happens to a girder supporting a bridge, the result is a catastrophe in both the mathematical and the mundane sense.

Another beautiful example in physics is provided by the bright geometric patterns called light caustics, which are created when light is reflected or refracted by a curved surface. A familiar caustic is the cusp-shaped curve that sometimes appears on the surface of a cup of coffee in bright sunlight; it is caused by the reflection of the sun's rays from the inside of the cup.

Another familiar caustic, which exhibits temporal as well as spatial discontinuities of brightness, is the changing pattern on the bottom of a swimming pool in sunlight. The rainbow is a family of colored caustic cones. More complex caustics can be produced by

	CATASTROPHE	CONTROL DIMENSIONS	BEHAVIOR DIMENSIONS	FUNCTION	FIRST DERIVATIVE
CUSPOIDS	FOLD	1	1	$\frac{1}{3}x^3 - ax$	$x^2 - a$
	CUSP	2	1	$\frac{1}{4}x^4 - ax - \frac{1}{2}bx^2$	$x^3 - a - bx$
	SWALLOWTAIL	3	1	$\frac{1}{5}x^5 - ax - \frac{1}{2}bx^2 - \frac{1}{3}cx^3$	$x^4 - a - bx - cx^2$
	BUTTERFLY	4	1	$\frac{1}{6}x^6 - ax - \frac{1}{2}bx^2 - \frac{1}{3}cx^3 - \frac{1}{4}dx^4$	$x^5 - a - bx - cx^2 - dx^3$
UMBILICS	HYPERBOLIC	3	2	$x^3 + y^3 + ax + by + cxy$	$3x^2 + a + cy$ $3y^2 + b + cx$
	ELLIPTIC	3	2	$x^3 - xy^2 + ax + by + cx^2 + cy^2$	$3x^2 - y^2 + a + 2cx$ $-2xy + b + 2cy$
	PARABOLIC	4	2	$x^2y + y^4 + ax + by + cx^2 + dy^2$	$2xy + a + 2cx$ $x^2 + 4y^3 + b + 2dy$

SEVEN ELEMENTARY CATASTROPHES describe all possible discontinuities in phenomena controlled by no more than four factors. Each of the catastrophes is associated with a potential function in which the control parameters are represented as coefficients (a, b, c, d) and the behavior of the system is determined by the variables (x, y). The behavior surface in each catastrophe model is the graph of all the points where the first derivative of this function is equal to zero or, when there are two first derivatives, where both are equal to zero.

shining a beam of light in a concave mirror or through spherical or cylindrical lenses (such as a light bulb or a beaker filled with water). In this application catastrophe theory has led to a better understanding of the phenomenon; Thom has shown that stationary caustics can have only three kinds of singular point. A mathematical

subtlety of the catastrophe-theory analysis of light caustics is that there is no dynamic; instead there is a variational principle that gives equal importance to both minimums and maximums.

The cusp catastrophe is a three-dimensional figure: two dimensions are required for the two control parameters and one

more is required for the behavior axis. Actually the behavior axis need not represent a single behavior variable; in models of brain function, for example, it may represent the states of billions of neurons, all varying at the same time. Nevertheless, catastrophe theory shows that it is always possible to select a single behavior variable and to plot

FOLD

CUSP

SWALLOWTAIL

HYPERBOLIC UMBILIC

ELLIPTIC UMBILIC

GRAPHS of five of the elementary catastrophes suggest the nature of their geometry. The fold catastrophe is a transverse section of a fold curve of the cusp catastrophe, and its bifurcation set consists of a single point. The cusp is the highest-dimensional catastrophe that can be drawn in its entirety. The swallowtail is a four-dimensional catastrophe and the hyperbolic umbilic and the elliptic umbilic catastrophes are five-dimensional. For these graphs only the three-dimensional bifurcation sets can be drawn; the behavior surfaces are not shown.

the behavior surface with respect to that axis only, and so to obtain the familiar three-dimensional graph.

If the graph is reduced to two dimensions, an even simpler model results: the fold catastrophe. In the fold catastrophe there is only one control parameter; the control space is a straight line and the bifurcation set is a single point on that line. The behavior space is a parabola, half of which represents stable states, the other half unstable. The two regions are separated by a fold point directly above the bifurcation point.

The Classification Theorem

The fold catastrophe can be regarded as a transverse section of the fold curve of the cusp catastrophe. The cusp in turn can be regarded as a stack of many fold catastrophes, together with one new singular point at the origin. More complicated catastrophes of higher dimension can be constructed on the same plan: each consists of all the lower-order catastrophes together with one new singularity at the origin.

If the control space is made three-dimensional while the behavior space remains one-dimensional, a unique four-dimensional catastrophe can be constructed. The behavior surface becomes a three-dimensional hypersurface, and instead of being folded along curves as in the cusp catastrophe it is folded along entire surfaces, a configuration that cannot easily be visualized. The bifurcation set no longer consists of curves with cusp points in two dimensions but is made up of surfaces in three dimensions that meet in cusps at their edges. A new singularity appears at the origin, called the swallowtail catastrophe. It is impossible to draw the complete swallowtail catastrophe because we cannot draw four-dimensional pictures. We can, however, draw its bifurcation set, which is three-dimensional, and from that drawing it is possible to derive some geometrical intuition about the swallowtail, just as it would be possible to describe the cusp catastrophe by drawing its bifurcation set (the cusp) in two dimensions and bearing in mind that the behavior surface is bimodal over the inside of the cusp. The catastrophe is called a swallowtail because the bifurcation set looks somewhat like one; the name was suggested by Bernard Morin, a blind French mathematician.

If yet another control parameter is added, a five-dimensional catastrophe is created. The fold, the cusp and the swallowtail again appear as sections, and a new singularity is associated with a "pocket" formed by the interpenetration of several surfaces. The shape of this pocket, or of sections of it, has suggested the name butterfly catastrophe. In the butterfly even the bifurcation set is four-dimensional and therefore cannot be drawn. It can be illustrated only through two- or three-dimensional sections [see illustration on this page].

There are two more five-dimensional ca-

BUTTERFLY

PARABOLIC UMBILIC

SECTIONS are the only recourse for illustrating the remaining two catastrophes, since even their bifurcation sets have more than three dimensions. The four-dimensional bifurcation set of the butterfly catastrophe is shown in three-dimensional sections; the fourth dimension is the butterfly factor, and if it happens to be time, then one configuration evolves into the other. Moving from left to right in each drawing is equivalent to changing the bias factor; two-dimensional "slices" reveal the effect of this factor more clearly. The four-dimensional bifurcation set of the parabolic umbilic catastrophe is also shown in a three-dimensional section. It is based on a drawing prepared with a computer by A. N. Godwin of Lanchester Polytechnic in England.

tastrophes, formed when the control space has three dimensions and the behavior space has two dimensions. They are called the hyperbolic umbilic and the elliptic umbilic catastrophes. As in the case of the swallowtail, their bifurcation sets consist of surfaces with cusped edges, and since they are three-dimensional they can be drawn. Finally, the six-dimensional catastrophe generated by a four-dimensional control space and a two-dimensional behavior space is called the parabolic umbilic. Its geometry is complex, and again only sections of its bifurcation set can be drawn.

By increasing the dimensions of the control space and the behavior space an infinite list of catastrophes can be constructed. The Russian mathematician V. I. Arnold has

classified them up to at least 25 dimensions. For models of phenomena in the real world, however, the seven described above are probably the most important because they are the only ones with a control space having no more than four dimensions. One particularly common class of processes, those determined by position in space and by time, cannot require a control space with more than four dimensions, since our world has only three spatial dimensions and one time dimension.

Even the catastrophes that cannot be drawn can be employed in modeling phenomena. Their geometry is completely determined, and the movement of a point over the behavior surface can be studied analytically if it cannot be seen graphically. Each

79

catastrophe is defined by a potential function, and in each case the behavior surface is the graph of all the points where the first derivatives of that function are zero.

The power of Thom's theory lies in its generality and its completeness. It states that if a process is determined by minimizing or maximizing some function, and if it is controlled by no more than four factors, then any singularity of the resulting behavior surface must be similar to one of the seven catastrophes I have described. If the process is governed by only two control factors, then the behavior surface can have only folds and cusps. The theorem states in essence that in any process involving two causes the cusp catastrophe is the most complicated thing that can happen to the graph. The proof of the theorem is too technical and too long to be presented here, but its consequences are straightforward: whenever a continuously changing force has an abruptly changing effect, the process must be described by a catastrophe.

After the cusp, the catastrophe with the richest spectrum of applications is the butterfly. Just as bimodal behavior determines the cusp model, so trimodal behavior determines the butterfly. In the cusp model of war policy, for example, where public opinion is divided between "doves" and "hawks," the butterfly model provides for the emergence of a compromise opinion favoring negotiation. The new mode of behavior arises as a new sheet of the behavior surface, growing smoothly out of the back of the pleat [see illustration below].

The geometry of the butterfly is controlled by four parameters. Two of them are familiar from the cusp models: the normal factor and the splitting factor. The remaining two are new: the bias factor and the butterfly factor. The effect of the bias factor is to alter the position and shape of the cusp; it swings the main part of the cusp left or right, but the vertex of the cusp bends the opposite way. At the same time the bias factor moves the behavior surface up and down.

The effect of the butterfly factor is to create the third stable mode of behavior. As the butterfly factor increases, the cusp on the control surface evolves into three cusps, which form a triangular "pocket." Above the pocket is the new, triangular sheet on the behavior surface, between the top and bottom sheets.

In order to draw the butterfly catastrophe two of the four parameters must be suppressed, and ordinarily the bias and butterfly factors are chosen. Their influence on the graph cannot be ignored, however. One effect of the bias factor is to reduce one side of the pocket until it disappears in a swal-lowtail catastrophe; the bias factor therefore tends to destroy a compromise. Since the butterfly factor controls the growth of the intermediate behavior sheet it enhances the stability of a compromise.

Anorexia Nervosa

A second application of the butterfly catastrophe, and an exceptionally fertile one, is to anorexia nervosa, a nervous disorder suffered mainly by adolescent girls and young women in whom dieting has degenerated into obsessive fasting. The model was developed by me in collaboration with J. Hevesi, a British psychotherapist who has introduced trance therapy in the treatment of anorexia. In a recent survey of 1,000 anorexic patients his were the only ones to state they had been completely cured.

In the initial phase of anorexia the obsessive fasting can lead to starvation, and sometimes to death. With the passage of time the patient's attitudes toward food and her behavior become progressively more abnormal. After about two years a second phase, called bulimia, usually develops, in which the victim alternately fasts and gorges herself. The bimodal behavior of this second phase immediately suggests a cusp catastrophe. The anorexic is caught in a hysteresis cycle, jumping catastrophically between two extremes, and she is denied access to the normal behavior in between. Catastrophe theory also suggests a theoretical cure: if a further bifurcation could be induced according to the butterfly catastrophe, a new pathway back to normality might be created.

The behavior surface in this model represents the overt behavior of the patient, ranging from uncontrolled gorging through normal eating to satiety and obsessive fasting. It also provides some indication of the underlying states of the brain; as in the models of aggression, we are concerned with emotional states that probably originate in the limbic system. Psychological evidence suggests that the behavior variable may actually be a measure of the relative weight given by the limbic system to inputs from the body as opposed to inputs from the cerebral cortex. In a normal person these inputs may be in some sense balanced, but in the anorexic one or the other may tend to dominate.

Among the control parameters the normal factor is hunger, which in normal people governs the rhythmic cycle between eating and satiety. The splitting factor is the degree of abnormality of the anorexic's attitudes toward food; the abnormality steadily increases as her condition deteriorates. Diets become more severe, entire classes of foods are eliminated; carbohydrates are at first avoided and later actively feared.

The bias factor in the butterfly graph is loss of self-control, which can be measured by loss of weight. In the first phase of the disorder the anorexic's attitudes are already abnormal, but she has control of herself. As a result she is trapped on the bottom sheet

BUTTERFLY CATASTROPHE provides for the emergence of compromise opinion in a model of the development of war policy. In the butterfly four controlling parameters are required, but here only two are shown (threat and cost) and the other two are assumed to remain constant. The bifurcation set is one of the sections on the preceding page; it is a complex curve with three cusps and a "pocket" in the middle. On the behavior surface above the pocket is a new sheet that provides for a new intermediate mode of behavior. If both the threat and the cost of war are high, the cusp model would allow for only the extreme positions advocating attack or surrender. The new sheet in the butterfly model represents a compromise opinion, advocating negotiation.

80

of the behavior surface; all the time she is awake the limbic system remains in states corresponding to a fasting frame of mind, even when she is eating her minimal meals.

As the anorexic loses weight she also loses control, and the bias factor gradually increases. As a result the cusp swings to the left on the graph [see illustration at right]; if it moves far enough, the right-hand side of the cusp intersects the anorexic cycle, bringing the sudden onset of the second phase. Now the anorexic is no longer trapped in a cycle of constant fasting but is caught in a hysteresis cycle, jumping from the bottom sheet to the top one and back again. In the words of a typical anorexic, the catastrophic jump from fasting to gorging takes place when she "lets go" and watches helplessly as the "monster within her" devours food for several hours, sometimes vomiting as well. The catastrophic return to fasting comes when exhaustion, disgust and humiliation sweep over her, an experience that many anorexics call the "knockout."

The period of fasting that follows the "knockout" in the hysteresis cycle is different from the constant fasting of the first phase. It lies at a different position on the behavior axis and might better be called purging. The limbic state associated with the earlier fasting is dominated by inputs from the cerebral cortex and is directed toward forbidding food to enter. During gorging the limbic system is dominated by inputs from the body. The limbic state underlying purging is again dominated by cerebral inputs, but it has a bodily component directed toward ridding the body of contamination.

The trance therapy employed by Hevesi reassures the patient, reduces her insecurity and thereby enables her to regain access to normal behavior. Anorexics tend to sleep fitfully, and when they are awake, they experience naturally occurring trancelike periods; it is on these periods that the therapist builds. The trance may represent a third state of the limbic system, in the otherwise inaccessible zone between the gorging and the purging states. When the patient is fasting she views the outer world with anxiety and when she is gorging she is overwhelmed by that world, but during the trance she is isolated, her mind free both of food and of scheming to avoid food. It is only then that reassurance is possible.

Reassurance becomes the butterfly factor in the model. It creates the new sheet of the behavior surface, which lies between the other two sheets and which eventually gives access to the stable, normal region behind the cusp. Because therapy usually takes place during the fasting part of the cycle, entering the trance is a catastrophic jump from the bottom sheet onto the intermediate sheet. Coming out of the trance is another catastrophe, which can take the patient either to the bottom sheet or to the top one.

After about two weeks of therapy and in about the seventh session of trance the patient's abnormal attitudes usually break

ANOREXIA NERVOSA, a nervous disorder of adolescent girls and young women that involves obsessive fasting, can be described as a butterfly catastrophe. Two of the controlling factors are hunger and the abnormality of attitudes toward food. In normal people hunger leads to a cycle of behavior that oscillates between eating and satiety; in the anorexic person, with abnormal attitudes, the same hunger cycle leads to quite different behavior. In the first phase of the disease (top) the cycle is trapped on the bottom sheet of the behavior surface, and the anorexic remains constantly in a fasting frame of mind. The second phase (bottom) is induced by a change in a third controlling factor, self-control. As the patient loses control of herself over a period of two years or more, the bifurcation set is gradually skewed to the left until the hunger cycle crosses the right-hand side of the cusp. She then enters a hysteresis cycle: she fasts until hunger causes her to catastrophically "let go," then she gorges herself until, after a "knockout" catastrophe, she reverts to fasting and purges herself of what she perceives as contamination.

81

down catastrophically and the personality is fused into a complete whole again. When the patient awakens from this trance, she may speak of it as a "moment of rebirth," and she finds that she can eat again without fear of gorging. The trance has seemingly opened a pathway in the brain back to the more balanced limbic states, so that the patient regains access to normal behavior. Subsequent trance sessions reinforce the experience.

The model of anorexia presented here is incomplete in several respects. I have omitted an additional control factor, drowsiness, which governs the behavioral distinction between waking and sleeping and the associated catastrophes of falling asleep and waking up. As a consequence the path from trance to normality in the model is misleading in that it omits the catastrophe of awakening. I have also not discussed the other half of the model, which concerns personality as opposed to behavior and which explains the escalation of the disorder, its rigidity, the stability of the abnormal attitudes and the breakdown of that stability at the moment of cure.

One of the strengths of the catastrophe-theory model of anorexia is that it explains the patient's own description of herself. The seemingly incomprehensible terms in which some anorexics describe their illness turn out to be quite logical when viewed in the framework of the catastrophe surfaces. The

advantage of a mathematical language in such applications is that it is psychologically neutral. It allows a coherent synthesis of observations that would otherwise appear to be disconnected.

The Future of Catastrophe Theory

Catastrophe theory is a young science: Thom published the first paper on it in 1968. So far its greatest impact has been on mathematics itself; in particular it has stimulated developments in other branches of mathematics that were required for the proofs of its theorems. The most important outstanding problems in the development of the theory concern the understanding and classification of generalized catastrophes and the more subtle catastrophes that arise when symmetry conditions are imposed. In addition there are problems associated with how catastrophe theory can be employed in conjunction with other mathematical methods and concepts, such as differential equations, feedback, noise, statistics and diffusion.

New applications of the theory are being explored in many fields. In physics and engineering, models have been developed for the propagation of shock waves, the minimum area of surfaces, nonlinear oscillations, scattering and elasticity. Michael V. Berry of the University of Bristol has recently employed the umbilic catastrophes

TREATMENT OF ANOREXIA relies on creating a third, intermediate mode of behavior. The new behavior is made possible by increasing the fourth control parameter of the butterfly catastrophe: reassurance. The effect of this fourth parameter is to create the pocket in the bifurcation set and thus to create the intermediate sheet in the behavior surface. In a system of therapy developed by J. Hevesi, a British psychotherapist, reassurance is encouraged by putting the patient in a trance. Initially the patient enters and leaves the trance through catastrophic

82

to predict new results in the physics of caustics and fluid flow, and he has confirmed those results by experiment.

Thom's own *Structural Stability and Morphogenesis,* inspired by the work of D'Arcy Wentworth Thompson and C. H. Waddington, was largely concerned with embryology, but as yet few biologists have pursued his ideas in the laboratory. I have constructed catastrophe models of the heartbeat, the propagation of nerve impulses and the formation of the gastrula and of somites in the embryo. Recent experiments conducted by J. Cooke of the Medical Research Council laboratories in London and by T. Elsdale of the Medical Research Council laboratories in Edinburgh appear to confirm some of my predictions.

Most of my own work, however, has been in the human sciences, as is suggested by the models described in this article. An increasing number of investigators are now suggesting models derived from catastrophe theory, and in the coming decade I look forward to seeing those models tested by experiment. Only then can we judge the true worth of the method.

Thom has employed the theory in an endeavor to understand how language is generated. It is an intriguing thought that the same mathematics may underlie not only the way the genetic code causes the embryo to unfold but also the way the printed word causes our imagination to unfold.

jumps from the intermediate sheet to the top and bottom sheets, as in the section at right. When the course of therapy has been successful, however, the patient makes a smooth transition from the intermediate sheet to the normal modes of behavior behind the pocket.