

I'm afraid you have not enough chaos in you to make a world.

GEORGE WILLIAM RUSSEL, DISCUSSION WITH JAMES JOYCE (UNKNOWN)

THE ESSENCE OF CHAOS

EDWARD LORENTZ

1993

Glimpses of Chaos

It only looks random

Words are not living creatures; they cannot breathe, nor walk, nor become fond of one another. Yet, like the human beings whom they are destined to serve, they can lead unique lives. A word may be born into a language with just one meaning, but, as it grows up, it may acquire new meanings that are related but nevertheless distinct.

Often these meanings are rather natural extensions of older ones. Early in our own lives we learn what “hot” and “cold” mean, but as we mature we discover that hot pursuit and cold comfort, or hot denials and cold receptions, are not substances or objects whose temperatures can be measured or estimated. In other instances the more recent meanings are specializations. We learn at an equally early age what “drink” means, but if later in life someone says to us, “You’ve been drinking,” we know that he is not suggesting that we have just downed a glass of orange juice. Indeed, if he tells someone else that we drink, he is probably implying not simply that we often consume alcoholic beverages, but that we drink enough to affect our health or behavior.

So it is with “chaos” – an ancient word originally denoting a complete lack of form or systematic arrangement, but now often used to imply the absence of some kind of order that ought to be present. Not withstanding

its age, this familiar word is not close to its deathbed, and it has recently outdone many other common words by acquiring several related but distinct *technical* meanings.

It is not surprising that, over the years, the term has often been used by various scientists to denote randomness of one sort or another. A recent example is provided by the penetrating book *Order Out of Chaos*, written by the Nobel Prize-winning physical chemist Ilya Prigogine and his colleague Isabelle Stengers. These authors deal with the manner in which many disorganized systems can spontaneously acquire organization, just as a shapeless liquid mass can, upon cooling, solidify into an exquisite crystal. A generation or two earlier, the mathematician Norbert Wiener would sometimes even pluralize the word, and would write about a chaos or several chaoes when referring to systems like the host of randomly located molecules that form a gas, or the haphazardly arranged collection of water droplets that make up a cloud.

This usage persists, but, since the middle 1970s, the term has also appeared more and more frequently in the scientific literature in one or another of its recently acquired senses; one might well say that there are several newly named kinds of chaos. In this volume we shall be looking closely at one of them. There are numerous processes, such as the swinging of a pendulum in a clock, the tumbling of a rock down a mountainside, or the breaking of waves on an ocean shore, in which variations of some sort take place as time advances. Among these processes are some, perhaps including the rock and the waves but omitting the pendulum, whose variations *are not random but look random*. I shall use the term *chaos* to refer collectively to processes of this sort – ones that appear to proceed according to chance even though their behavior is in fact determined by precise laws. This usage is argu-

ably the one most often encountered in technical works today, and scientists writing about chaos in this sense no longer feel the need to say so explicitly.

In reading present-day accounts, we must keep in mind that one of the other new usages may be intended. Sometimes the phenomena being described are things that appear to have random arrangements in space rather than random progressions in time, like wildflowers dotting a field. On other occasions, the arrangements or progressions are simply very intricate rather than seemingly random, like the pattern woven into an oriental rug. The situation is further complicated because several other terms, notably *nonlinearity*, *complexity*, and *fractality*, are often used more or less synonymously with *chaos* in one or several of its senses. In a later chapter I shall have a bit to say about these related expressions.

In his best-selling book *Chaos: Making a New Science*, which deals with chaos in several of its newer senses, James Gleick suggests that chaos theory may in time rival relativity and quantum mechanics in its influence on scientific thought. Whether or not such a prophecy comes true, the “new science” has without question jumped into the race with certain advantages. Systems that presumably qualify as examples of chaos can very often be seen and appreciated without telescopes or microscopes, and they can be recorded without time-lapse or high-speed cameras. Phenomena that are supposedly chaotic include simple everyday occurrences, like the falling of a leaf or the flapping of a flag, as well as much more involved processes, like the fluctuations of climate or even the course of life itself.

I have said “presumably” and “supposedly” because there is something about these phenomena that is not quite compatible with my description of chaos as something that

is random in appearance only. Tangible physical systems generally possess at least a small amount of true randomness. Even the seemingly regular swinging of the pendulum in a cuckoo clock may in reality be slightly disturbed by currents in the air or vibrations in the wall; these may in turn be produced by people moving about in a room or traffic passing down a nearby street. If chaos consists of things that are actually *not* random and only *seem* to be, must it exclude familiar everyday phenomena that have a bit of randomness, and be confined to mathematical abstractions? Might not such a restriction severely diminish its universal significance?

An acceptable way to render the restriction unnecessary would be to stretch the definition of chaos to include phenomena that are slightly random, provided that their much greater apparent randomness is not a by-product of their slight true randomness. That is, real-world processes that appear to be behaving randomly – perhaps the falling leaf or the flapping flag – should be allowed to qualify as chaos, as long as they would continue to appear random even if any true randomness could somehow be eliminated.

In practice, it may be impossible to purge a real system of its actual randomness and observe the consequences, but often we can guess what these would be by turning to theory. Most theoretical studies of real phenomena are studies of approximations. A scientist attempting to explain the motion of a simple swinging pendulum, which incidentally is not a chaotic system, is likely to neglect any extraneous random vibrations and air currents, leaving such considerations to the more practical engineer.

Often he or she will even disregard the clockwork that keeps the pendulum swinging, and the internal friction that makes the clockwork necessary, along with anything else that is inconvenient. The resulting pencil-

and-paper system will be only a model, but one that is completely manageable. It seems appropriate to call a real physical system chaotic if a fairly realistic model, but one with the system's inherent suppressed, still appears to behave randomly.

Pinballs and Butterflies

[...] According to the narrower definition of randomness, a *random* sequence of events is one in which anything that can ever happen can happen next. Usually it is also understood that the probability that a given event will happen next is the same as the probability that a like event will happen at any later time. A familiar example, often serving as a paradigm for randomness, is the toss of a coin. Here either heads or tails, the only two things that can ever happen, can happen next. If the process is indeed random, the probability of throwing heads on the next toss of any particular coin, whether 50 percent or something else, is precisely the same as that of throwing heads on any other toss of the same coin, and it will remain the same unless we toss the coin so violently that it is bent or worn out of shape. If we already know the probability, knowing in addition the outcome of the last toss cannot improve our chances of guessing the outcome of the next one correctly.

It is true that knowing the results of enough tosses of the same coin can suggest to us what the probability of heads is, for that coin, if we do not know it already. If after many tosses of the coin we become aware that heads has come up 55 percent of the time, we may suspect that the coin is biased, and that the probability has been, is, and will be 55 percent, rather than the 50 percent that we might have presupposed.

The coin is an example of complete randomness. It is the sort of randomness that one commonly has in mind when thinking of random numbers, or deciding to use a random-number generator. According to the broader definition of randomness, a *random* sequence is simply one in which any one of several things can happen next, even though not necessarily *anything* that can ever happen can happen next. What actually is possible next will then depend upon what has just happened. An example, which, like tossing a coin, is intimately associated with games of chance, is the shuffling of a deck of cards. The process is presumably random, because even if the shuffler should wish otherwise – for example, if on each riffle he planned to cut the deck exactly in the middle, and then allow a single card to fall on the table from one pile, followed by a single card from the other pile, etc. – he probably could not control the muscles in his fingers with sufficient precision to do so, unless he happened to be a virtuoso shuffler from a gaming establishment. Yet the process is not completely random, if by what happens next we mean the outcome of the next single riffle, since one riffle cannot change any given order of the cards in the deck to any other given order. In particular, a single riffle cannot completely reverse the order of the cards, although a sufficient number of successive riffles, of course, can produce any order.

A *deterministic* sequence is one in which only one thing can happen next; that is, its evolution is governed by precise laws. Randomness in the broader sense is therefore identical with the absence of determinism. It is this sort of randomness that I have intended in my description of chaos as something that *looks* random.

Tossing a coin and shuffling a deck are processes that take place in discrete steps-successive tosses or riffles. For quantities that vary continuously, such as the speed of

a car on a highway, the concept of a next event appears to lose its meaning. Nevertheless, one can still define randomness in the broader sense, and say that it is present when more than one thing, such as more than one prespecified speed of a car, is possible at any specified future time. Here we may anticipate that the closer the future time is to the present, the narrower the range of possibilities – a car momentarily stopped in heavy traffic may be exceeding the speed limit ten seconds later, but not one second later. Mathematicians have found it advantageous to introduce the concept of a *completely* random continuous process, but it is hard to picture what such a process in nature might look like.

Systems that vary deterministically as time progresses, such as in mathematical models of the swinging pendulum, the rolling rock, and the breaking wave, and also systems that vary with an inconsequential amount of randomness – possibly a real pendulum, rock, or wave – are technically known as *dynamical systems*. At least in the case of the models, the state of the system may be specified by the numerical values of one or more *variables* – or the model pendulum, two variables – the position and speed of the bob will suffice; the speed is to be considered positive or negative, according to the direction in which the bob is currently moving. For the model rock, the position and velocity are again required, but, if the model is to be more realistic, additional variables that specify the orientation and spin are needed. A breaking wave is so intricate that a fairly realistic model would have to possess dozens, or more likely hundreds, of variables.

Returning to chaos, we may describe it as behavior that is deterministic, or is nearly so if it occurs in a tangible system that possesses a slight amount of randomness, but does not *look* deterministic. This means that the present

state completely or almost completely determines the future, but does not appear to do so. How can deterministic behavior look random? If truly identical states do occur on two or more occasions, it is unlikely that the identical states that will necessarily follow will be perceived as being appreciably different. What can readily happen instead is that almost, but not quite, identical states occurring on two occasions will *appear* to be just alike, while the states that follow, which need not be even nearly alike, will be observably different. In fact, in some dynamical systems it is normal for two almost identical states to be followed, after a sufficient time lapse, by two states bearing no more resemblance than two states chosen at random from a long sequence. Systems in which this is the case are said to be *sensitively dependent on initial conditions*. With a few more qualifications, to be considered presently, sensitive dependence can serve as an acceptable definition of chaos, and it is the one that I shall choose.

“Initial conditions” need not be the ones that existed when a system was created. Often they are the conditions at the beginning of an experiment or a computation, but they may also be the ones at the beginning of any stretch of time that interests an investigator, so that one person’s initial conditions may be another’s midstream or final conditions.

Sensitive dependence implies more than a mere increase in the difference between two states as each evolves with time. For example, there are deterministic systems in which an initial difference of one unit between two states will eventually increase to a hundred units, while an initial difference of a hundredth of a unit, or even a millionth of a unit, will eventually increase to a hundred units also, even though the latter increase will inevitably consume more time. There are other deterministic systems

in which an initial difference of one unit will increase to a hundred units, but an initial difference of a hundredth of a unit will increase only to one unit. Systems of the former sort are regarded as chaotic, while those of the latter sort are not considered to constitute chaos, even though they share some of its properties.

Because chaos is deterministic, or nearly so, games of chance should not be expected to provide us with simple examples, but games that *appear* to involve chance ought to be able to take their place. Among the devices that can produce chaos, the one that is nearest of kin to the coin or the deck of cards may well be the pinball machine. It should be an old-fashioned one, with no flippers or flashing lights, and with nothing but simple pins to disturb the free roll of the ball until it scores or becomes dead.

One spring in the thirties, during my undergraduate years at Dartmouth, a few pinball machines suddenly appeared in the local drugstores and eating places. Soon many students were occasionally winning, but more often losing, considerable numbers of nickels. Before long the town authorities decided that the machines violated the gambling laws and would have to be removed, but they were eventually persuaded, temporarily at least, that the machines were contests of skill rather than games of chance, and were therefore perfectly legal.

If this was indeed so, why didn’t the students perfect their skill and become regular winners? The reason was chaos. As counterparts of successive tosses of a coin or riffles of a deck, let the “events” be successive strikes on a pin. Let the outcome of an event consist of the particular pin that is struck, together with the direction from the pin to the center of the ball, and the velocity of the ball as it leaves the pin. Note that I am using *velocity* in the technical sense, to denote speed together with direction of motion,

just as position with respect to some reference point implies distance together with direction of displacement.

Suppose that two balls depart one after the other from the same pin in slightly different directions. When the balls arrive at the next pin, their positions will be close together, compared to the distance between the pins, but not necessarily close, compared to the diameter of a ball. Thus, if one ball hits the pin squarely and rebounds in the direction from which it came, the other can strike it obliquely and rebound at right angles. This is approximately what happens in Figure 1, which shows the paths of the centers of two balls that have left the plunger of a pinball machine at nearly equal speeds. We see that the angle between two paths can easily increase tenfold whenever a pin is struck, until soon one ball will completely miss a pin that the other one hits. Thus a player will need to increase his or her control tenfold in order to strike one more pin along an intended pathway.

Of course, the pinball machine in Figure 1 is really a mathematical model, and the paths of the balls have been computed. The model has incorporated the decelerating effect of friction, along with a further loss of energy whenever a ball bounces from a pin or a side wall, but, in a real machine, a ball will generally acquire some side spin as it hits a pin, and this will alter the manner in which it will rebound from the next pin. It should not alter the conclusion that the behavior is chaotic – that the path is sensitively dependent on the initial speed.

Even so, the model as it stands fails in one respect to provide a perfect example of chaos, since the chaotic behavior ceases after the last pin is struck. If, for example, a particular ball hits only seven pins on its downward journey, a change of a millionth of a degree in its initial direction would amplify to ten degrees, but a change of a

ten-millionth of a degree would reach only one degree. To satisfy all of the requirements for chaos, the machine would have to be infinitely long – a possibility in a model if not in reality – or else there would have to be some other means of keeping the ball in play forever. Any change in direction, even a millionth of a millionth of a degree, would then have the opportunity to amplify beyond ten degrees.

An immediate consequence of sensitive dependence in any system is the impossibility of making perfect predictions, or even mediocre predictions sufficiently far into the future. This assertion presupposes that we cannot make measurements that are completely free of uncertainty. We cannot estimate by eye, to the nearest tenth of a degree nor probably to the nearest degree, the direction in which a pinball is moving. This means that we cannot predict, to the nearest ten degrees, the ball's direction after one or two strikes on a pin, so that we cannot even predict which pin will be the third or fourth to be struck. Sophisticated electronic equipment might measure the direction to the nearest thousandth of a degree, but this would merely increase the range of predictability by two or three pins. As we shall see in a later chapter, sensitive dependence is also the chief cause of our well-known failure to make nearly perfect weather forecasts.

I have mentioned two types of processes – those that advance step by step, like the arrangements of cards in a deck, and those that vary continuously, like the positions or speeds of cars on a highway. As dynamical systems, these types are by no means unrelated. The pinball game can serve to illustrate a fundamental connection between them.

Suppose that we observe 300 balls as they travel one by one through the machine. Let us construct a diagram containing 300 points. Each point will indicate the position of the centre of one ball when that ball strikes




FIGURE 1. THE PINBALL MACHINE

its first pin. Let us subsequently construct a similar diagram for the second strike. The latter diagram may then be treated as a full-scale map of the former, although certainly a rather distorted map. A very closely spaced cluster of points in the first diagram may appear as a recognizable cluster in the second. Dynamical systems that vary in discrete steps, like the pinball machine whose “events” are strikes on a pin, are technically known as *mappings*. The mathematical tool for handling a mapping is the *difference equation*. A system of difference equations amounts to a set of formulas that together express the values of all of the variables at the next step in terms of the values at the current step.

I have been treating the pinball game as a sequence of events, but of course the motion of a ball between strikes is as precisely governed by physical laws as are the rebounds when the strikes occur. So, for that matter, is the motion of a coin while it is in the air. Why should the latter process be randomness, while the former one is chaos? Between any two coin tosses there is human intervention, so that the outcome of one toss fails to determine the outcome of the next. As for the ball, the only human influence on its path occurs before the first pin is struck, unless the player has mastered the art of jiggling the machine without activating the tilt sign.

Since we can observe a ball between strikes, we have the option of plotting diagrams that show the positions of the centres of the 300 balls at a sequence of closely spaced times, say every fiftieth of a second, instead of only at moments of impact. Again each diagram will be a full-scale map of the preceding one. Now, however, the prominent features will be only slightly changed from one diagram to the next, and will appear to flow through the sequence. Dynamical systems that vary continuously, like

the pendulum and the rolling rock, and evidently the pinball machine when a ball's complete motion is considered, are technically known as *flows*. The mathematical tool for handling a flow is the *differential equation*. A system of differential equations amounts to a set of formulas that together express the rates at which all of the variables are currently changing, in terms of the current values of the variables.

When the pinball game is treated as a flow instead of a mapping, and a simple enough system of differential equations is used as a model, it may be possible to solve the equations. A complete solution will contain expressions that give the values of the variables at any given time in terms of the values at any previous time. When the times are those of consecutive strikes on a pin, the expressions will amount to nothing more than a system of difference equations, which in this case will have been derived by solving the differential equations. Thus a mapping will have been derived from a flow.

Indeed, we can create a mapping from *any* flow simply by observing the flow only at selected times. If there are no special events, like strikes on a pin, we can select the times as we wish – for instance, every hour on the hour. Very often, when the flow is defined by a set of differential equations, we lack a suitable means for solving them – some differential equations are intrinsically unsolvable. In this event, even though the difference equations of the associated mapping must exist as relationships, we cannot find out what they look like. For some real-world systems we even lack the knowledge needed to formulate the differential equations; can we honestly expect to write any equations that realistically describe surging waves, with all their bubbles and spray, being driven by a gusty wind against a rocky shore?

If the pinball game is to chaos what the coin toss is to complete randomness, it has certainly not gained the popularity as a symbol for chaos that the coin has enjoyed as a symbol for randomness. That distinction at present seems to be going to the butterfly, which has easily outdistanced any potential competitors since the appearance of James Gleick's book, whose leading chapter is entitled "The Butterfly Effect."

The expression has a somewhat cloudy history. It appears to have arisen following a paper that I presented at a meeting in Washington in 1972, entitled "Does the Flap of a Butterfly's Wings in Brazil Set Off a Tornado in Texas?" I avoided answering the question, but noted that if a single flap could lead to a tornado that would not otherwise have formed, it could equally well prevent a tornado that would otherwise have formed. I noted also that a single flap would have no more effect on the weather than any flap of any other butterfly's wings, not to mention the activities of other species, including our own. The paper is reproduced in its original form as Appendix 1.

The thing that has made the origin of the phrase a bit uncertain is a peculiarity of the first chaotic system that I studied in detail. Here an abbreviated graphical representation of a special collection of states known as a "strange attractor" was subsequently found to resemble a butterfly, and soon became known as the butterfly. In Figure 2 we see one butterfly; a representative of a closely related species appears on the inside cover of Gleick's book. A number of people with whom I have talked have assumed that the butterfly effect was named after this attractor. Perhaps it was.

Some correspondents have also called my attention to Ray Bradbury's intriguing short story "A Sound of Thunder," written long before the Washington meeting.

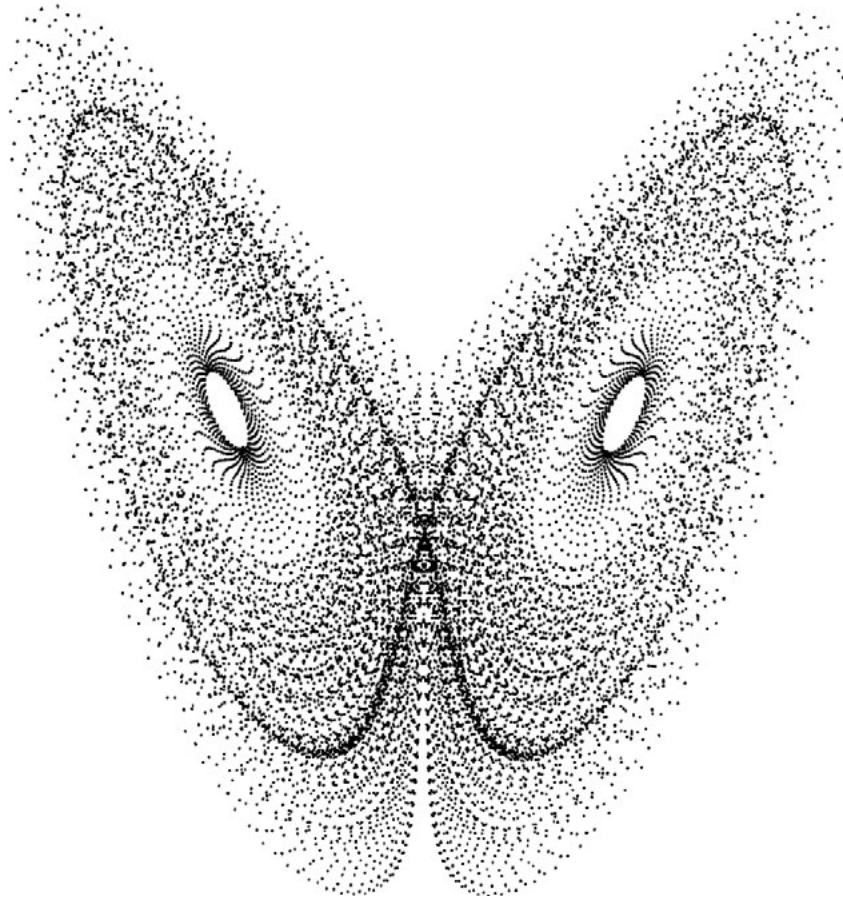


FIGURE 2, THE BUTTERFLY

Here the death of a prehistoric butterfly, and its consequent failure to reproduce, change the outcome of a present-day presidential election.

Before the Washington meeting I had sometimes used a seagull as a symbol for sensitive dependence. The switch to a butterfly was actually made by the session convenor, the meteorologist Philip Merilees, who was unable to check with me when he had to submit the program titles. Phil has recently assured me that he was not familiar with Bradbury's story. Perhaps the butterfly, with its seeming frailty and lack of power, is a natural choice for a symbol of the small that can produce the great.

Other symbols have preceded the seagull. In George R. Stewart's novel *Storm*, a copy of which my sister gave me for Christmas when she first learned that I was to become a meteorology student, a meteorologist recalls his professor's remark that a man sneezing in China may set people to shovelling snow in New York. Stewart's professor was simply echoing what some real-world meteorologists had been saying for many years, sometimes facetiously, sometimes seriously.