# CHAPTER 3
# Our Chaotic Weather

## *Prediction: A Tale of Two Fluids*

AN OLD JOKE about the supposedly inept weather forecaster never ceases making the rounds in one form or another. I recall one version that appeared shortly after Harry Truman, against all predictions except his own, had defeated Thomas Dewey for the presidency of the United States in 1948. It was a cartoon showing an applicant at a Weather Bureau employment office, and the interviewer was saying to him, "You worked for a public-opinion poll? I think we could use you."

Joking aside, however, the weather-forecasting community and the general public are both acutely aware that official forecasts, including those for later in the same day, are sometimes just plain wrong. To the often-heard question, "Why can't we make better weather forecasts?" I have been tempted to reply, "Well, why should we be able to make any forecasts at all?"

Why indeed should we expect to see the future, or at least a little part of it? First of all, we may believe that there is a set of physical laws governing the changes in the weather from one moment to the next, and that the very existence of these laws, whether or not we know them, ought to make prediction possible. Our faith may be fortified by the realization that other natural phenomena, governed by somewhat similar laws, have been regularly predicted with considerable success; consider the tides in the ocean, which we can predict rather accurately a few days ahead and almost as accurately many years ahead. Finally, our weather forecasts *are* correct far more often than they would be if they were pure guesses.

I must admit that my first encounter with the question of tidal prediction left me with an uneasy feeling. Apparently I had always looked at announcements of the times of coming high and low tides as statements of fact, as firmly established as the times of yesterday's tides, and not to be doubted. Yet it is apparent that they are predictions, and so, for that matter, are such "facts" as the times of sunrise and sunset that appear in almanacs. That the latter times are simply predictions with a high probability of being almost exactly right, rather than facts, becomes evident when we realize that an unforeseen cosmic catastrophe—perhaps a collision with an asteroid—could render them completely wrong. Even without a catastrophe they can be slightly in error; increases or decreases in the strength of the globe-encircling westerly-wind currents, which occur at irregular intervals, are compensated by small but measurable decreases or increases in the speed of rotation of the underlying earth, and the times of coming sunrises and sunsets may be delayed or advanced by a millisecond or so.

Returning to the ocean, let us for purposes of comparison with the atmosphere define the height of the tide as the height of the ocean surface above some fixed reference level, after individual waves have been averaged out. Let us compare the predictability of the tides, so defined, with the predictability of the temperature of the air—possibly the weather element whose prediction interests us most, although sometimes we may be more concerned with whether or not it is going to rain.

Both the atmosphere and the ocean are large fluid masses, and each envelops all or most of the earth. They obey rather similar sets of physical laws. They both possess fields of motion that tend to be damped or attenuated by internal processes, and both fields of motion are driven, at least indirectly, by periodically varying external influences. In short, each is a very complicated forced dissipative dynamical system. Perhaps it would be more appropriate to call them two components of a larger dynamical system, since each exerts a considerable influence on the other at the surface where they come into contact. The winds, which vary with the state of the system, produce most of the ocean's waves, and help to drive the great currents like the Gulf Stream. Evaporation from the ocean, which also varies with the state of the system, supplies the atmosphere with most of the moisture that subsequently condenses and still later falls as rain or snow. Why, with so many similarities, should we have had so much more success in tidal than in weather prediction? Are oceanographers more capable than meteorologists?

As a meteorologist whose close friends include a number of oceanographers, I would dispute any such hypothesis.

Take a look at the periodic external driving forces—primarily the heat emitted by the sun and the gravitational pull of the sun and the moon. The atmosphere and the ocean will respond to these forces by undergoing periodic oscillations, but, as with many dynamical systems, these oscillations will be accompanied by additional irregular behavior. Near the coast, the regular response of the ocean includes most of the tidal oscillations, while the irregular response includes the occasional anomalously high tides produced by unanticipated strong winds. The regular response of the atmosphere includes the normal excursions of temperature between summer and winter, or day and night, while the irregular response includes extended hot spells and cold spells, as well as the sudden temperature changes that often accompany the progression of large storms across the oceans and continents,

It appears, then, that in attempting to forecast the tides we are for the most part trying to predict the highly predictable regular response. We may also wish to predict the smaller irregular response, but even when we fail to do so we have usually made a fairly good forecast. In forecasting the weather, or, for definiteness, the temperature, we usually take the attitude that the regular response is already known—we know in advance that summer will be warmer than winter—and we regard our problem as that of predicting those things that we do not already know simply by virtue of knowing the climate. In short, when we compare tidal forecasting and weather forecasting, we are comparing prediction of predictable regularities and some lesser irregularities with prediction of irregularities alone. If oceanographers are smarter than meteorologists, it is in knowing enough to pick a readily solvable problem. I should hasten to add that most oceanographers are not tidal forecasters anyway, nor, for that matter, are most meteorologists weather forecasters. In most respects the oceans present just as many challenges as the atmosphere.

In the following pages I shall introduce the atmosphere as an example of an intricate dynamical system, and present the case for believing that its irregularities are manifestations of chaos. After a brief overview I shall enumerate various procedures through which the presence of chaos might be confirmed. Finally, I shall examine some of the consequences of the atmosphere's chaotic behavior.

## *Meteorology: Two Tales of One Fluid*

Chaotic dynamical systems come in many sizes. Our mathematical model of the sled on the ski slope has only three variables, yet it serves to illustrate many of the basic properties of chaos. It is by no means the simplest system with this capability; that honor, together with the honor of being the most intensively studied of all chaotic dynamical systems, goes to one with only one variable, which varies in accordance with a single quadratic difference equation—the kind of equation that defines a mapping.

The system is so simple that with a pocket calculator you can convince yourself in a few minutes that it can behave chaotically. Choose a fixed number; call it $c$. Choose a number between $-c$ and $c$ as the leading member of a sequence, and construct the remainder of the sequence by always squaring the most recent number and then subtracting $c$ to obtain the next number. Some but not all choices of $c$ between 1.4 and 2.0 will furnish you with sequences that lack periodicity. With one of these choices for $c$ you can also observe sensitive dependence directly, by repeating your calculations with a slightly different initial number. In a slightly altered form this system is known as the *logistic* equation, and it has been used in studies of population dynamics.

At the other extreme are systems with a large or even infinite number of variables. Among these we may expect to find the one whose states are simply the global weather patterns. Let us call it the *global weather system*.

Meteorologists have kept pace fairly well with their contemporaries in the art of creating esoteric terminology to describe esoteric concepts. They talk freely about potential pseudo-equivalent temperature, moist semigeostrophy, and dynamic anticyclogenesis, and they have even devised triple acronyms: GOCC stands for GATE Operations Control Center, GATE in turn stands for GARP Atlantic Tropical Experiment, and GARP is the Global Atmospheric Research Program, a multinational effort conceived in the 1960s and flourishing in the seventies and eighties. Nevertheless, the variables of the global weather system include no mysterious quantities. They are the familiar weather elements that we have always known—the ones that make us acutely aware of their presence, and whose values we can often estimate with fair accuracy, whenever we step outdoors. They are the temperature, the wind, the humidity, and some representation of the clouds that may be enveloping us and the rain or snow that may be falling on us. To these we must add pressure—a familiar item in many weather reports even though it bears less directly on our comfort. We can easily detect the rise in pressure that we experience as we drive down a long steep hill, but most of us would find it difficult, on waking in the morning, to say whether the pressure was higher or lower than it had been when we fell asleep. The humidity can be expressed as relative humidity, wet-bulb temperature, dew point, or water-vapor concentration; any one of these, in combination with temperature and pressure, determines the others.

If we lived on a planet whose atmosphere consisted of a pure gas of uniform composition, we would have only temperature, wind, and pressure to worry about. By the wind I mean the three-dimensional wind, with the strength of an updraft or downdraft appearing as one velocity component. Our own atmosphere, not to mention some other atmospheres in our solar system, is more complicated, in that one of its most important constituents, in our case water vapor, occurs in highly variable concentrations, ordinarily comprising more than two percent of the mass of the atmosphere in the humid tropics, but less than one-tenth of one percent in the colder air at high latitudes or high elevations. Water also occurs as suspended or falling liquid drops and solid particles, so that in reality the atmosphere is not wholly a gas. The variables of the system must therefore include the concentrations of water vapor, which we perceive as humidity, and liquid and solid water, which we observe as denseness of clouds and intensity of rain or snow. It should probably also include the concentrations of such pollutants as dust and smoke.

We could make a case for adding still more quantities, but what makes the atmosphere so complicated as a dynamical system is not so much the proliferation of physical variables as the fact that their values vary from one point to another and not merely from one time to another. To know a single state of the global weather system, we must therefore know the value of each variable at every point. Since there are plainly an infinite number of points in the atmosphere, the system would seem to have an infinite number of variables.

Actually the situation is not quite so bad. On a fine enough spatial scale the weather elements vary rather smoothly, and if two states are nearly alike at each of a sufficiently dense network of well-spaced points, they will be nearly alike at the intervening locations. It is therefore legitimate to treat the atmosphere as a system with a finite number of variables, and to conclude that it is compact. What is not so legitimate is to treat it as a system with a *small* finite number of variables; the number is truly enormous.

How ought we to approach a system of such complexity? Let me present two answers that we might have received at the midpoint of the twentieth century, before the computer had come along and changed everything. Each point of view had its ardent supporters.

Consider first the methods of a subdiscipline that has been known for a century or so as dynamic meteorology, although it might more accurately be called meteorological dynamics. To dynamic meteorologists, the state of the atmosphere consists of the spatial distributions of the temperature, wind, and other weather elements. The dynamicist starts out with the physical laws that govern the behavior of the atmosphere, and usually expresses these laws as mathematical equations. Among them is one of Isaac Newton's laws of motion, familiar to many as "Force equals mass times acceleration," but rearranged for meteorological use as "Acceleration equals force divided by mass," or "Rate of change of velocity equals force per mass." With a knowledge of the state of the atmosphere one can in principle evaluate the force at any point, and thus learn how the velocity of the air passing that point will change as time progresses. The laws of thermodynamics will tell us how the temperature will behave, and other laws will allow us to handle the remaining variables. In short, there is a dynamical basis for forecasting the weather as it evolves, and more generally for treating the atmosphere as a dynamical system.

The synoptic meteorologist would tell a far different tale, regarding the dynamicist's description as grossly incomplete and perhaps irrelevant. Synoptic meteorology is the study of the characteristic structures into which states of the atmosphere can be analyzed. These include meandering jet streams that may encircle the globe in middle latitudes; vortices of subcontinental size, also known as high- and low-pressure systems or simply highs and lows, that travel across the oceans and continents in middle and higher latitudes and bring many of our day-to-day weather changes; smaller and more intense vortices at lower latitudes, known as hurricanes, typhoons, or cyclones according to the ocean over which they originate; towering cumulonimbus clouds with their accompanying thunderstorms and occasional tornados; and small innocuous clouds scattered through an otherwise clear sky. Typical horizontal extents of the structures mentioned are respectively 10, 000, 1, 000, 100, 10, and 1 kilometers, or a bit more; the list is only a sample. Structures of each type can be counted on to appear again and again, but each individual structure will have its own peculiarities, just as the human race continues, but different people inevitably have their distinguishing personalities. The synoptician's principal tool for identifying and studying the larger structures is the weather map.

The practicing forecaster in precomputer days was effectively an applied synoptic meteorologist. Individual forecasters would learn through their own experience and that of their predecessors how each structure typically develops and moves. They would find that certain peculiarities in a structure signal certain unusual happenings, and they would recognize the telltale signs for the appearance of new structures and the demise of older ones. They would discover, for example, that a high and a low, seemingly heading for the same spot, will retain their identities and simply deflect each other, rather than annihilating each other. In preparing a forecast, a forecaster would most likely construct a prognostic chart, which would be a personal estimate of what the next day's weather map would look like, and he or she would use the chart to infer the coming local weather conditions.

Effectively the forecasting rules are to the synoptician what the physical laws are to the dynamicist. If they could be formulated in such a way as to give a unique prediction in any conceivable situation, they would define an alternative mathematical model, which would constitute another dynamical system.

Why, aside from tradition, didn't some midcentury practicing forecasters opt for the methods of dynamic meteorology? The most likely reason is a practical one; no acceptable weather forecast based primarily on the dynamic equations had ever been produced.

What, then, did dynamic meteorologists have to show for their many years of efforts? As scientists rather than technicians, their interest was directed toward a true understanding of the atmosphere in terms of the physical laws that govern it. They would have been happy to discover why a particular weather pattern with its inevitable peculiarities would evolve as it did, but they were far more interested in why weather patterns vary at all from day to day, or why they are inevitably filled with the large-scale vortices that synoptic meteorologists take for granted. They would seek to learn what processes would allow these vortices to develop and persist for a while, in the face of the ubiquitous dissipative processes that by themselves would always act to destroy them. By midcentury, they had found many of the answers.

I am not trying to imply that synoptic meteorologists are technicians rather than scientists, nor, for that matter, that technicians are somehow inferior to scientists. It is seldom that a single approach to a problem proves to be the only fruitful one. Synopticians have a keen scientific interest in documenting the properties of the structures that they observe, and in establishing regular relationships between neighboring structures. They are not averse to examining their findings for consistency with the physical laws, but their final conclusions are based more on careful analyses of extensive sequences of weather patterns.

Neither do I wish to suggest that someone who knows all about both dynamic and synoptic meteorology knows all about meteorology. There are numerous other subdisciplines, each with its own group of experts. To name just a couple, one is *cloud physics,* where a fundamental concept is the distribution of the sizes of the drops in a cloud, and where one studies the processes by which tiny suspended droplets and ice crystals become converted into larger drops and particles, which will then fall out as rain or snow. Another is *instrumentation,* where one investigates the strong and weak points of the various instruments via which we have discovered much of what we believe we know about the weather, and where one also designs new instruments in the hopes of gathering hitherto inaccessible information.

Dynamic and synoptic meteorology are not wholly divorced. There have always been some meteorologists who have been outstanding in both subdisciplines. In strong academic meteorology departments, the programs in dynamic and synoptic meteorology tend to be well coordinated. There appear to be other institutions, however, where excellent work in dynamic meteorology may take place in an applied mathematics department, while correspondingly good work in synoptic meteorology may be found in a geography department, but where any communication between the departments is hard to detect. At midcentury, the history of meteorology was marked by both cooperation and contention between the two methodologies.

### *The Unperformable Experiment*

The most direct way to look for chaos in a concrete system, whether it is a simple object sliding down a slope or an atmosphere with its multitude of interdependent structures, is to work with the system itself. If we have released a board and watched or perhaps photographed it on its downward trip, we can easily retrieve it and release it from nearly the same point, to see whether it will follow nearly the same path. Unfortunately for scientific experimentation, but perhaps fortunately for humanity as a whole, we cannot stop the advance of the weather and then reestablish a pattern that has previously been observed, in order to disturb it slightly and then see how rapidly the resulting weather will diverge from the weather that occurred earlier. We can readily disturb the existing weather, perhaps violently by setting off an explosion or starting a fire, or more gently by dropping crystals of dry ice into a cloud—or perhaps even by releasing a butterfly—and we can observe what will happen, but then we shall never know what would have happened if we had left things alone.

What about comparing what happens after we disturb the weather with a forecast of what would have happened if we had not interfered? The forecast is based upon extrapolation from incompletely observed conditions; at best, it can tell us what would have happened if someone had introduced a disturbance similar in magnitude and structure to the observational error—the error in estimating the initial state. If the disturbance that *we* introduce is to tell us anything in addition, it must be large enough not to be swamped by the observational error. However, a disturbance of this magnitude seems hard to produce, when we note that entire thunderstorms may go undetected between observing sites.

Lacking the ability to change the weather to suit our needs, we can wait for what meteorologists call an analogue—a weather pattern that closely resembles one that has previously been observed —in order to see how closely the behavior following the second occurrence resembles that following the first. This method also fails; even though the atmosphere seems to be a compact system—one in which pairs of analogues must eventually occur—good analogues on a global scale have not been found within the few decades that global weather conditions have been recorded. Patterns that are much alike over regions of continental size are sometimes observed, but, when these fail to develop similarly, they may do so because dissimilar weather structures have moved in from distant regions, rather than because of any sensitivity to small local differences.

There remains the reasonably well established observation that weather variations are not periodic. Of course they have periodic components, the most obvious ones being the warming and cooling that occur with the passage of the seasons of the year or the hours of the day. Careful measurements have also detected weak signals with a lunar period, probably gravitational effects, and there is virtually no limit to the number of periods that investigators have *claimed* to have discovered. Some of these have been stated to several decimal places. Nevertheless, if we take an extended record of temperature or some other weather variable and subtract out all verified or suspected periodic components, we are left with a strong irregular signal. Migratory storms that cross the oceans and continents are still present in full force. These are presumably manifestations of chaos.

Since good analogues of global extent have yet to be discovered, we cannot with certainty rule out the possibility that, when one finally does appear, the subsequent weather will repeat the earlier sequence. That is, the atmosphere may really be behaving periodically, with a period whose length exceeds that of any weather records. We are left with the strong impression that the atmosphere is chaotic, but we would like additional evidence.

## *Voices from Dishpans*

It is fairly easy to construct a scale model of a bumpy slope and observe the descent of a ball or some other object. Modeling a planetary atmosphere in the laboratory is another matter. We might think of letting a fluid fill the space between two concentric spheres. The inner sphere could represent the planet, while the outer one could take the place of gravity to the extent of preventing the fluid from leaving the planet, but how could we then introduce a force that would simulate gravity within the fluid by always being directed toward the planet's surface?

The pioneers in laboratory modeling were already well versed in dynamics. Dynamic meteorologists have long been accustomed to simplifying their equations before putting them to use. Sometimes the simplifications are merely deletions of terms that appear to be inconsequential, but equally often they consist of omitting or significantly altering certain physical features or processes. Thus, effectively, they replace the atmosphere by a different atmosphere, which Napier Shaw described in the early twentieth century, in his four-volume treatise *Manual of Meteorology,* as a fairy tale, but which today we would call a model. The dynamicist assumes, or at least hopes, that the weather in a make-believe atmosphere will be more or less like the real weather in its gross features, and will differ mainly in minor details.

The equations expressing the laws that govern evaporation and condensation of water in the atmosphere are rather awkward, while those governing the transformation of a cloud composed of tiny suspended water droplets into a rain cloud are even more forbidding, and dynamicists often work with model atmospheres that are devoid of water in any form. Likewise, many of them have an undisguised aversion to spherical geometry, and their atmospheres may flow over a flat rotating earth instead of one whose surface is curved. Any dynamic meteorologist who could explain the development and persistence of large-scale vortices in the fairyland of dry atmospheres and flat earths would feel that he had completed the major part of the work of solving the real-world problem.

Laboratory modelers found it quite acceptable to build into their apparatus the same distortions that dynamic meteorologists traditionally built into their equations. The first experiments to bear fruit were designed by Dave Fultz at the University of Chicago. After a number of trials, he settled upon a cylindrical vessel partly filled with water, placed on a rotating turntable, and subjected to heating near the periphery and cooling near the center. Figure 28 is a schematic oblique view of his apparatus. The bottom of the container is intended to simulate one hemisphere of the earth's surface, the water is intended to simulate the air above this hemisphere, the rotation of the turntable simulates the earth's rotation, the heating and cooling simulate the excess external heating of the atmosphere in low latitudes and the excess cooling in high latitudes, and gravity simulates itself. Fultz had hoped that the motions that developed in the water might resemble the large-scale wind patterns in the atmosphere.

Originally the edge of the container extended beyond the rim of the turntable, and the heating was accomplished by a fixed Bunsen burner, while exposure to room temperature was supposed to take care of the cooling. This setup was soon supplanted by more easily controlled heating coils arranged around the periphery of the

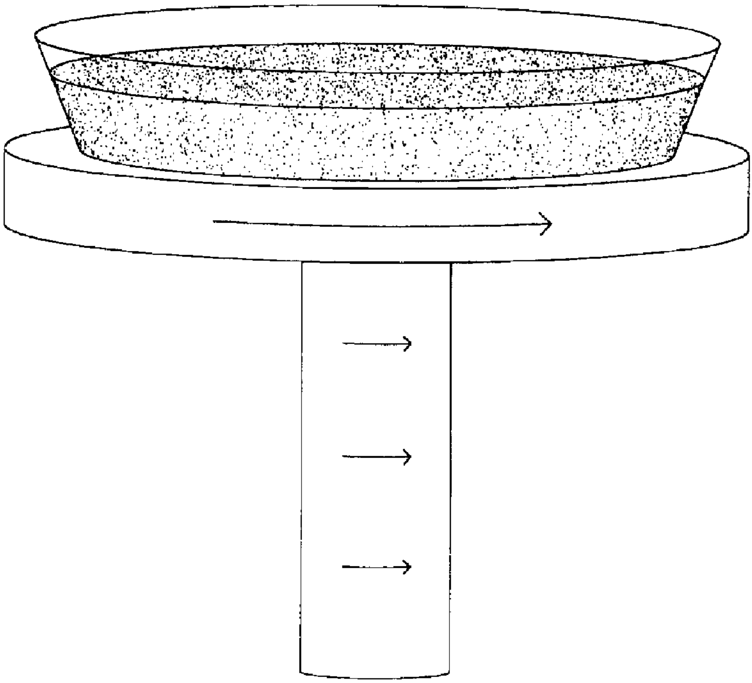

**Figure 28.** A schematic oblique view of the apparatus used by Fultz in the dishpan experiments. The arrows indicate the direction of rotation of the turntable. Heating is applied at the rim of the container and cooling is applied at the center.

container, while the cooling was sometimes accomplished by an upward jet of cold water through a hole in the turntable. Flow at the upper surface, which was intended to simulate atmospheric motion at high elevations, was made visible by a sprinkling of aluminum powder. A special camera that effectively rotated with the turntable took time exposures, so that a moving aluminum particle would appear as a streak, and sometimes each exposure ended with a flash, which would add an arrowhead to the forward end of each streak. Flow deeper within the fluid could be detected by injecting a dye, and thermometers were often inserted to record the temperature fluctuations that were expected to accompany the passage of the simulated weather structures. The turntable generally rotated counterclockwise, as does the earth when viewed from above the north pole. The collection of components cost about forty thousand dollars—a fair sum for 1950—but the central component —the container—was an ordinary dishpan purchased for a dollar or two, and the work became known as the dishpan experiments.

Fultz had assumed that it would make little difference whether the working fluid was a gas or a liquid, and water was certainly the simplest choice. The water possessed no impurities to simulate real atmospheric water vapor and clouds, and the bottom of the dishpan was essentially flat, with nothing to distinguish between oceans and continents. Dynamicists who might have been criticized for omitting the water in the atmosphere and the curvature of the earth could have claimed that they were really trying to model the dishpan experiments.

Because everything in the experiments was arranged with perfect symmetry about the axis of rotation, at least within the limits of experimental control, one might have anticipated that the resulting flow patterns would also be symmetric, looking perhaps like the one shown schematically in Figure 29. Fultz was hoping for something more like Figure 30, with a meandering jet stream and an assemblage of vortices typical of the atmosphere.

He got more than he bargained for. Both flow patterns appeared, the choice depending upon the speed of the turntable's rotation and the intensity of the heating. In brief, with fixed heating, a transition from circular symmetry would take place as the rotation increased past a critical rate. With fixed, sufficiently rapid rotation, a similar transition would occur when the heating reached a critical strength, while another transition back to symmetry would occur when the heating reached a still higher critical strength. The experiments proved to be repeatable, producing transitions at the same combinations of values of rotation and heating when run again.

In the early experiments, the flow that was asymmetric appeared to be irregular also, changing continuously from one pattern to another, much as the real atmosphere changes. We now recognize Fultz's transitions as classical bifurcations, between a steady system, whose attractor consists of a single point in phase space, and an unsteady, apparently chaotic one, whose attractor should be composed of an infinite complex of multidimensional manifolds.

In England, meanwhile, Raymond Hide was experimenting at Cambridge University with a somewhat similar apparatus. It differed mainly in that the fluid occupied a ring-shaped region between two concentric cylinders instead of the interior of a single cylinder. Hide found similar transitions between symmetric and

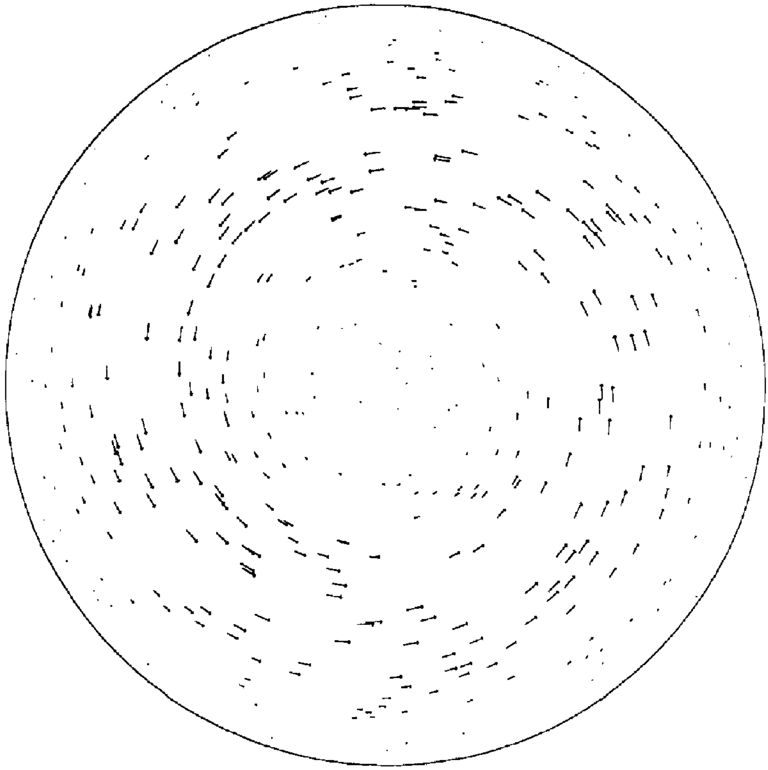

**Figure 29.** A schematic view of symmetric flow at the upper surface of the water in the dishpan.

asymmetric flow, but, possibly because of the restrictive effect of the inner cylinder, the asymmetric flow was often regular, and would consist of a chain of apparently identical waves, which would travel around the cylinder without changing their shape. Here was a dynamical system with a one-dimensional attractor—a circle in a suitably chosen phase space—the separate states being distinguished only by the longitudes of the waves.

Hide also discovered a remarkable regular phenomenon, which he called *vacillation*. Here also a chain of identical waves would appear, but, as they traveled along, they would alter their shape in unison in a regular periodic fashion, and, after many "days"—many rotations of the turntable—they would regain their original shape and then repeat the cycle. Here the system had a two-dimensional

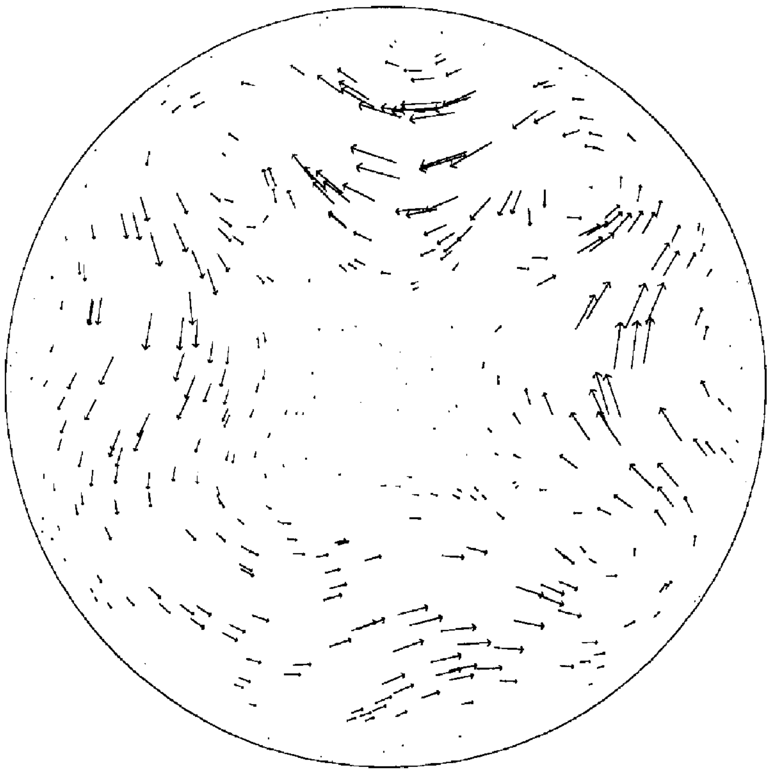

**Figure 30.** A schematic view of asymmetric flow at the upper surface of the water in the dishpan.

attractor, the two varying quantities being the longitudinal phase of the waves and the phase of the vacillation cycle.

Hide was not a meteorologist, although he has since become one of the leading dynamicists in the meteorological community, and he was actually attempting to model the motions in the earth's magnetic core, but as an all-arounder he quickly saw the relevance of his experiments for the atmosphere, and noted the resemblance between his vacillation cycles and the frequently seen fluctuations between intervals of strong and weak westerly winds. He and Fultz soon traded information. In due time Fultz produced both uniform wave motion and vacillation in the dishpan, and Hide found that his own apparatus would support irregular behavior.

Other scientists soon took up the laboratory modeling, although the number remained small compared, for example, with those who continued to favor equations. By using such components as a discarded phonograph turntable, Alan Faller managed to reproduce the essence of Fultz's early experiments at a cost of four dollars. Subsequently, at the Woods Hole Oceanographic Institution, he built a "dishpan" eight feet in diameter, and was able to produce cold fronts and warm fronts—narrow zones separating extensive air masses, or, in the experiments, water masses, with contrasting temperatures. Fronts are ubiquitous features of sea-level weather maps.

Figure 31 shows a vacillation cycle as captured photographically by Fultz. The separate pictures were taken at intervals of four "Zdishpan days." In the first picture, a meandering circumpolar jet stream, identifiable by the longest bright streaks, appears to be made up of five virtually identical waves. The waves proceed to change their shape, and after eight days they have become transformed into nearly circular vortices. The vortices subsequently elongate, and by sixteen days the pattern, not shown, has become virtually indistinguishable from the initial one.

Figure 32 shows two photographs of the dishpan, one "Zday" apart, during an irregular and presumably chaotic regime. Accompanying them are Fultz's streamline analyses, based on the photographs. A nearly circular vortex below the center may be seen to elongate considerably as it propagates "eastward." The vortex rotates counterclockwise, and, like its counterparts in the northern hemisphere of the real atmosphere, it is a true low-pressure center.

In Figure 33 the first streamline analysis has been inverted so as to simulate southern-hemisphere flow, and it is compared with an actual upper-level southern-hemisphere weather map, containing approximate streamlines. The patterns are not superposable, but qualitatively they are so much alike that they might almost have been selected from the same attractor.

The implications of the laboratory experiments are profound. Structures such as jet streams, traveling vortices, and fronts appear to be basic features of rotating heated fluids, and are not peculiar to atmospheres. Efforts in dynamic meteorology had not always been success stories, and it had been proposed at times that the failures might result from using the wrong dynamics—possibly from being unaware of some strange force. The similarities and differences between the atmosphere and the experiments strongly

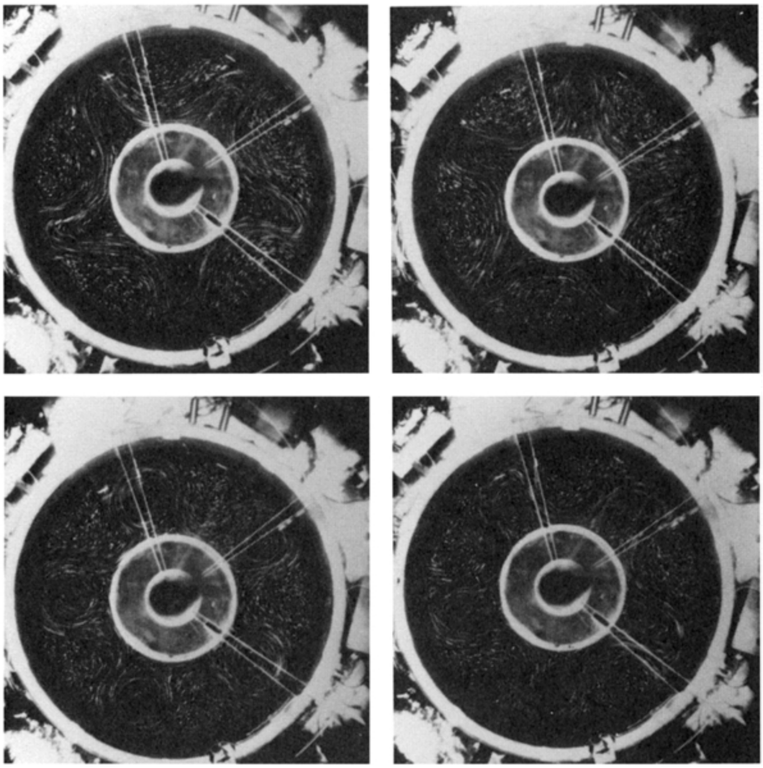

**Figure 31.** Streak photographs of the flow at the upper surface of the dishpan, at four phases of a vacillation cycle. The upper right, lower left, and lower right patterns follow the upper left by four, eight, and twelve rotations. The pattern after four more rotations, not shown, is almost indistinguishable from the first. Photographs by Dave Fultz.

suggest that the principal causes of the gross behaviors are to be found in forces and processes that influence both systems—gravity, rotation, and differential heating—while such properties as compressibility, which air possesses but water does not, are secondary. If a mysterious force is at work in the atmosphere, it would have to be at work in the dishpan also.

Finally, it may be a prehistorical accident that our day is about twenty-four hours long, instead of several times longer. If so, it may be an accident that our atmosphere behaves like a rapidly rotating dishpan instead of a slow one, fluctuating chaotically instead of regularly, and that our weather is rather unpredictable instead of

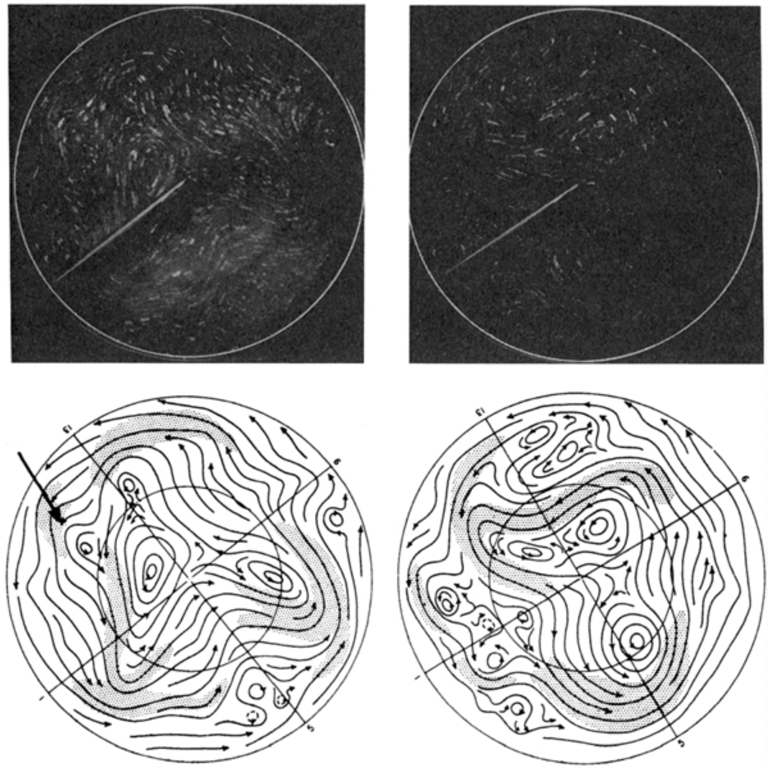

**Figure 32.** Streak photographs of the upper surface of the dishpan in an experiment revealing chaotic behavior, accompanied by streamline analyses based on the photographs. The right-hand patterns follow those on the left by one rotation. Photographs and analyses by Dave Fultz, reproduced by permission of the American Meteorological Society.

continually executing a monotonous cycle, or perhaps not changing at all except for the slow advance of the seasons.

*The Five-Million-Variable Dynamical System*

By far the strongest evidence for a chaotic atmosphere has come from mathematical models. Strictly speaking, these models are what dynamicists have been using since the dawn of dynamic meteorology, but more recently a "model" has generally meant a system of equations arranged for numerical solution on a computer. The history of the more realistic models of this sort is essentially

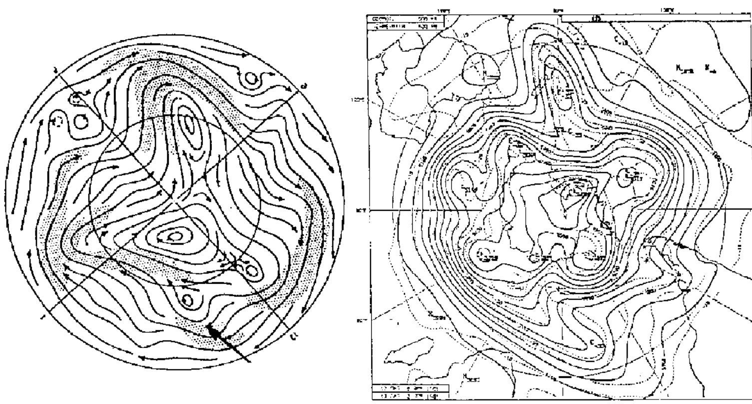

**Figure 33.** The left-hand streamline analysis from Figure 32, inverted so as to simulate southern hemisphere flow, compared with a real-world height-contour analysis of the 500-millibar surface over a portion of the southern hemisphere. At this surface, whose average height is close to 5.5 kilometers, the height contours closely resemble streamlines.

the history of the use of the dynamic equations for weather forecasting.

The story opens in Norway in the early twentieth century, when Vilhelm Bjerknes, considered by some to have been the all-time great meteorologist, proposed that the weather-forecasting problem was simply the problem of solving the system of equations representing the basic laws, using a set of simultaneous weather observations as the initial state. He maintained that the equations were known, but recognized that there was no practical method of solving them. It was Bjerknes who, many years later, championed the idea that the reason that vortices and other structures of continental or subcontinental size must be present in the atmosphere is not the dynamic impossibility of a flow pattern without them, which would look like symmetric flow in the dishpan, and would constitute a state of equilibrium. Rather, it is the instability of such a pattern with respect to inevitable disturbances of large horizontal extent but small amplitude. These disturbances would proceed to inten sify and then persist as part of the complete pattern.

The next chapter begins in England during World War I, when the versatile scientist Lewis Richardson, who was undaunted by the formidable nature of the equations, attempted to solve them numerically. In his procedure, he replaced the continuous distributions of pressure, wind, and other quantities, which in any event could be estimated only by interpolating between reports at weather stations, by the values of these quantities at a regular grid of points. The gradients of these quantities—the rates at which they varied horizontally—were then approximated by differences between values at adjacent grid points.

As a Quaker, Richardson objected to armed combat, but he had no fear of the front, and was happy to serve during the war as an ambulance driver. He brought with him an extensive set of weather data for one selected day, and between shifts he would perform the thousands of additions, subtractions, and multiplications needed to produce a single six-hour forecast for an area no larger than Europe. His predicted weather pattern not only turned out to be wrong, but did not even resemble any pattern that had been seen before. Richardson correctly attributed his failure to inadequacies in the initial wind measurements, although subsequent analysis has shown that his procedure would have produced serious although less drastic errors even with perfect initial data.

Imagine an enormous creature from outer space that swoops down close to the earth, reaches out with a giant paddle, and stirs the atmosphere for a short while before disappearing. Wholly aside from the possibly disastrous effect upon the living beings of the earth, what will be the likely effect on the weather?

Air that has simply been moving around a low pressure system, for example, as it normally does, may be left moving predominantly into it. The low will rapidly fill, soon becoming a high, after which the now piled-up air will surge outward, leaving a deep low, into which air will rush a second time before rushing out again. The precise chain of events will be further complicated by the ever-present deflecting effect of the earth's rotation. Rather similar events will take place at locations where the creature has left the air moving out of a low, or into or out of a high. In short, there will be violent fluctuations of pressure and accompanying fluctuations in the wind. Theory indicates that the period of an oscillation—a change from low to high to low again—will be comparable to one day.

The atmosphere has its own method of getting rid of any such fluctuations; otherwise they might be a part of our normal weather. Mechanical and thermal damping play an essential role in their removal. After a few weeks, the oscillations will be hardly detectable and the weather will be back to normal, although the particular sequence of weather patterns will undoubtedly not be the one that would have developed without the disturbance. Stated otherwise, the normal weather patterns that occur day after day belong to the attractor of the global weather system. The alien creature will produce a new state—think of it as an initial state—that is well removed from the attractor, but, as in any dissipative dynamical system, the transient effects will ultimately damp out, and normal behavior will resume.

Now imagine that Richardson had wished to use his numerical procedure to discover what would happen if such a creature should pay a visit. He could have done no better than to do what he actually did. Observations of the weather such as those that he used, and interpolations to standard geographical locations, are fraught with small errors. The true state of the atmosphere, and the state as Richardson could best estimate it, differed in much the same way as the states of an atmosphere before and after being stirred. The true state belonged to the attractor, and the estimated state did not. Inevitably Richardson predicted the violent oscillations that his assumed initial state demanded.

In his monumental book *Weather Prediction by Numerical Process,* completed in 1922, Richardson presented his procedure in full detail, and discussed his forecast. He ended by envisioning a weather center where sixty-four thousand people working in shifts could produce a weather forecast more rapidly than the weather itself could advance. The one feature that he failed to envision was the device that within half a century would be working as rapidly as sixty-four thousand people.

Following Richardson's efforts, the general attitude toward numerical weather prediction became pessimistic. Many prominent meteorologists seriously doubted that wind observations could ever become accurate enough to suppress the spurious oscillations. Those who felt otherwise tended to be discouraged by the sheer magnitude of the needed computations.

Fortunately one of the optimists was the Swedish scientist Carl-Gustaf Rossby, a dynamic meteorologist in the literal as well as the technical sense, and certainly another candidate for the title of all-time great meteorologist. One of his contributions was the suggestion that the key to understanding the atmosphere was to be found in the wind instead of the pressure. A low-pressure system is also a spinning vortex, and, although the barometer provides the easiest and most accurate means of detecting and mapping the structure, the wind pattern may exert the greater influence on its behavior, with the pressure serving largely as an indicator.

As the middle of the century approached, the renowned mathematician John von Neumann became interested in designing automatic computers and applying them to involved problems. Although not a meteorologist, he recognized the weather-forecasting problem as ideal for his needs. Soon he went about assembling a group of meteorologists and other scientists to work on the numerical forecasting problem, at the Institute for Advanced Study in Princeton, New Jersey. In addition to the largely computational matters to be faced, there remained the problem of the spurious fluctuations.

The states that the atmosphere assumes as the weather progresses are all supposed to be restricted to the attractor of the global weather system; any transient effects should have disappeared long ago. If some system of equations is to be used to produce short-range forecasts, say one or two days in advance, it must handle the states that are on the attractor, or the best approximations to these states that it is able to depict, in approximately the way that the atmosphere handles them. On the other hand, there is no need for it to be able to deal properly with states that are not on the attractor, since these will never arise.

One member of von Neumann's group who recognized this situation was the then-young meteorologist Jule Charney, later to be recognized as still another possible all-time great. Before arriving in Princeton, Charney had become acquainted with Rossby. Starting from Rossby's ideas as to the importance of the wind, he had managed to construct a system of equations that effectively failed to distinguish between unrealistic weather patterns in which strong oscillations would have been expected to develop, and slightly different but more realistic ones in which they would not, and, with either type of pattern as an initial state, would predict that the oscillations would not arise. His system could not have detected a visit from the creature from outer space. Effectively the new equations filtered out the oscillations, and later were sometimes called the filtered equations, while the more nearly exact equations that Richardson had used became known as the primitive equations. With various modifications that rendered them more adaptable to computation, the filtered equations became the basis for the first experimental series of numerical weather predictions. The story of these early efforts has been aptly recounted by both Philip Thompson and George Platzman, two of the original participants.

By the middle fifties, the moderate success of the forecasts, even though they did not match up to the ones turned out by experienced synoptic forecasters, led to the introduction of numerical forecasting as a part of the operational procedure of various national weather services. At the very least, synoptic forecasters now had, in addition to everything that they had formerly used, the information that "this is what the computer says will happen". They could use or reject this information as they saw fit. As the years advanced, forecasters came to rely more and more on the numerical product.

As dynamical systems, the new models were rather peculiar. They did not possess attractors that resembled the attractor of the real atmospheric system. If they had been used to make long-range forecasts, say a month in advance, they would have produced weather patterns quite unlike anything seen in nature. Indeed, in the earliest forecasts the external forcing and internal dissipation were completely omitted from the equations, on the grounds that, no matter how important they might have been in bringing about an initial state, their added influence during the next day or two would be minor. Thus, aside from any changes that the numerical scheme might have introduced, the models conserved total energy, and, like other Hamiltonian systems, possessed no attractors at all.

An outgrowth of numerical weather prediction that recognized this shortcoming was global circulation modeling. The equations used were much like those already used in short-range prediction, but, as the name suggests, the area to which they were applied covered the whole globe, or at least most of one hemisphere, rather than a more restricted region. The purpose of the new models was to produce simulated weather whose long-term behavior was realistic in as many respects as possible, rather than to make forecasts, and the initial conditions were often purposely chosen not to look like real weather patterns, in the hopes that reasonable patterns would soon develop. Stated otherwise, it was hoped that the new models would possess realistic attractors. Needless to say, external forcing and internal dissipation had to be included.

The prototype global circulation model was constructed by Norman Phillips, who had been working closely with Charney at the Institute for Advanced Study. His dynamical system had 450 variables. He extended his solution for one month, and produced a meandering jet stream and traveling vortices before he encountered computational problems. Subsequently he succeeded in diagnosing the computational difficulty and prescribing a cure, thus paving the way for the countless models that were to follow.

During the sixties it became apparent that the filtered equations, which had made numerical forecasting possible with the early computers, were not going to produce forecasts of the quality that some had hoped for. With the advent of more powerful computers, some meteorologists turned their attention back to the primitive equations. The solution to the problem of the violent oscillations, which had led to the rejection of the primitive equations a decade earlier, turned out to be surprisingly simple in concept, although not so easy to implement.

The initial patterns of wind and perhaps pressure, as interpolated from observations, are inaccurate in any case; otherwise they would not give rise to oscillations so much stronger than those observed in nature. Why then shouldn't we tamper with these patterns a bit, at the risk of making them slightly more inaccurate, to produce *new initial states* from which oscillations *cannot* develop, as an alternative to using *equations* that will not *predict* that oscillations will develop? After extensive research, several methods of making the needed adjustments were devised; improvements are still being introduced. The tampering or adjusting process, known as *initialization,* is now a standard part of every routine forecasting procedure that uses the primitive equations. Let us note that initialization need not produce the correct initial state; it simply produces one that might be correct instead of one that cannot be. Ideally, it will move the observed state to some nearby state on the attractor.

By the seventies, global circulation modelers were also turning to the primitive equations. As the years advanced and computers became ever more powerful, the distinction between global circulation models and numerical forecasting models tended to disappear. Operational forecasting centers could now afford to use models that covered the globe, or at least a hemisphere, and, with increasing interest in predicting several days ahead, during which time storms could move in from distant areas, there was good reason to do so.

The big model with which I am personally most familiar is the operational model of the European Centre for Medium Range Weather Forecasts in Reading, England. The Centre is a joint venture of the weather services of more than a dozen European nations. As its name implies, its mission is to produce forecasts at medium range, which in this case has meant up to ten days ahead. The scientists who have worked there, either directly with the model or on problems relevant to its construction and performance, have included not only representatives from the member nations but visitors from around the world.

The model itself is global, and, like most large models today, is based on the primitive equations. I should probably not call it *the* model, because, during the ten years or so that I have intermittently worked with it, it has been frequently subjected to modifications aimed at improving its performance. As of 1985, it possessed three physical quantities—temperature and two wind components—defined at each of nineteen elevations, and a fourth quantity—water-vapor content—defined at all but the high elevations. Pressure was explicitly defined only at the lowest level, since pressures at higher levels could be inferred from those at lower levels when temperatures and humidities at intervening levels were known. Other auxiliary variables such as soil moisture were defined where appropriate. Each physical quantity at each level was for practical purposes specified at a grid of more than 11,000 points, spanning the globe. This produced a total of some 800,000 variables.

As if these were not enough, as of late 1991 the resolution was doubled in both the latitudinal and longitudinal directions, producing effectively 45,000 grid points, while the number of standard elevations was increased to 31. This produced a model with five million variables. Of such stuff are modern global circulation models made.

Lest a system of 5,000, 000 simultaneous equations in as many variables appear extravagant, let us note that, with a horizontal grid of less than 50,000 points, each point must account for more than 10,000 square kilometers. Such an area is large enough to hide a thunderstorm in its interior. I have heard speculations at the Centre that another enlargement of the model is unlikely to occur soon, and this evidently means that not all significant weather structures will soon be resolved.

What about chaos? Almost all global models, aside from the very earliest, have been used for predictability experiments, in which two or more solutions originating from slightly different initial states have been examined for the presence of sensitive dependence.

Interest has not been so much in chaos itself as in the feasibility of producing extended-range forecasts, particularly at the two-week range.

Almost without exception, the models have indicated that small initial differences will amplify until they are no longer small. There is even good quantitative agreement as to the rate of amplification. Unless we wish to maintain that the state-of-the-art model at the European Centre, and competitive models at the National Meteorological Center in Washington and other centers, do not really behave like the atmosphere, in spite of the rather good forecasts that they produce at short range, we are more or less forced to conclude that the atmosphere itself is chaotic.

## *The Consequences*

The possibility that an object may slide chaotically down a slope is largely a matter of academic interest. Chaos in the atmosphere has farther-reaching consequences.

The most obvious effect is the limitation that it imposes upon our ability to forecast. Imagine that you are a computer-age weather forecaster, and that instead of making just one extended-range forecast you have decided to make a dozen or so. You take a dozen estimates of the initial state that are more or less alike, differing from each other by no more than the typical uncertainty in estimating the true initial state: temperatures at some locations might differ by a degree or so, while wind speeds might vary by two or three knots. To fend off anticipated spurious oscillations, you apply the initialization procedure to each state. When you make a two-week forecast from each state, using the best approximation to the true physical laws that your computer can handle, you will find that, because of sensitive dependence, the dozen predictions are not much alike. If you have no basis for saying which, if any, of the dozen initial states is correct, you will have no basis for saying which of the predictions should become the official forecast.

The process that you will have carried out is not somebody's wild fantasy. It shows signs of becoming the forecasting procedure of the future, when computers have become still more powerful. It is known as *Monte Carlo* forecasting. It takes its name from the famous gambling resort because the original idea was that, out of a virtually infinite collection of eligible initial states, a few should be chosen at random, although it now appears that, if the procedure becomes operational, the states may be chosen systematically. Monte Carlo methods have in fact been used in numerous fields almost since the advent of computers as a means of generating statistical distributions.

The Monte Carlo procedure can give some idea of the degree of confidence to be put in a particular day's forecast. If the separate forecasts show little resemblance to each other, the confidence will be low, whether one of the forecasts is selected arbitrarily as the official one, or whether some average is used. If the forecasts are much alike, any one of them is likely to be fairly good.

What is the basis for choosing two weeks as a time after which the forecasts might differ significantly? That story goes back to the early 1960s, when preparations were under way for the Global Atmospheric Research Program, the international effort to obtain world-wide observations of a higher quality than previously known, and to extend our knowledge of atmospheric dynamics so that the new observations might be put to optimum use. Among the original aims of the program was the production of good two-week forecasts. Such a vast program obviously would require vast funding, and "aims" should perhaps be viewed as a euphemism for "selling points."

It was just at that time that the possibility of sensitive dependence in the atmosphere's behavior was beginning to be recognized by meteorologists. Jule Charney, who was one of the leaders in organizing the program, became concerned that two-week forecasting might be proven impossible even before the first two-week forecast could be produced, and he managed to replace the aim of making these forecasts with the more modest aim of determining whether such forecasts were feasible. In 1964, a special conference held in Boulder, Colorado, was attended by a wide assortment of dynamicists, synopticians, and other meteorologists, including all the then-active global circulation modelers. The agenda included scientific papers presented by representatives of ten nations, and Charney talked about the possibility of chaotic behavior. Between sessions, however, when the real work of such conferences generally takes place, he managed to persuade all of the global-circulation modelers to use their models to perform numerical experiments in which pairs of forecasts originating from slightly different conditions would be examined for sensitive dependence. From these experiments, performed in the ensuing months, Charney's committee concluded that a reasonable estimate of the average doubling time for small errors in the temperature or wind pattern was five days.

The doubling time soon acquired the status of a standard measure of predictability. If we have a fair idea of the size of typical errors in estimating the initial state, and if we have decided how great an error we can tolerate in the forecast, we will know how many doublings are acceptable, and, if we also know the typical doubling time, we can calculate the range of acceptable predictability. This range should then be adjusted upward, since errors typically grow less rapidly after they have become moderately large. The five-day doubling time seemed to offer considerable promise for one-week forecasts, but very little hope for one-month forecasts, while two-week forecasts seemed to be near the borderline.

What typically happens when a more powerful computer becomes available to the meteorological community is that larger mathematical models are built, so that a one-day or a one-week forecast takes about as long to produce as it did before. The enlargements generally entail increases in horizontal and vertical resolution, but they may also involve more realistic formulations of certain physical processes, such as the absorption and emission of radiation by the atmosphere, or the flow of air over mountainous terrain. One specific enlargement in the sixties was the change from filtered to primitive equations. With the new models came new predictability experiments, and by 1970 the typical doubling time seemed to be closer to three days than five. By the early eighties, the European Centre model and other models had reduced the time to about two days; this estimate still stood in 1990. Thus, although it had become fairly well established that two-week forecasts showed a slight edge over pure guesswork, scenarios in which the locations and intensities of migratory storms were predicted two weeks ahead appeared less and less realistic.

Some promise for further improvement in forecasting has come from the observation that, with the European Centre model, differences between two forecasts that start from different states regularly amplify more slowly than differences between either forecast and the weather that actually transpires. If the model perfectly represented the physical laws, the rates of amplification should all be the same. Improvements must therefore still be possible. Computations indicate that, if the present model is correctly estimating the atmosphere's doubling time, a perfect model should produce three-day forecasts as good as today's one-day forecasts, which generally are good; one-week forecasts as good as today's three-day forecasts, which occasionally are good; and two-week forecasts comparable to today's one-week forecasts, which, although not very good, may contain some useful information. This is the optimistic view; one can always take the alternative view that, since the present models are not perfect, the appropriate doubling time may be even less than the estimated two days.

Let us take another look at the calculated doubling times. First of all, they are properties of the models that have been used to compute them. As a property of the real-world system, a two-day doubling time can at most refer to a doubling time for structures that are resolved by the models—jet streams, temperate latitude vortices, and perhaps tropical storms. Structures that are not regularly resolved either by a global observational network or by the computational grid of a global model have an important influence on the resolved structures; thunderstorms, and to some extent less intense showers, are effective in altering the global temperature and humidity patterns by carrying heat and water from low to high elevations. Failure to include these effects in a model leads to inferior forecasts. The larger-scale pattern tends to be indicative of the presence or absence of significant smaller-scale structures, and the standard procedure is to estimate, at each point, the most probable effect of the smaller scales. This procedure, know as *parameterization*, has been the subject of entire conferences. It is still one of the less realistically formulated aspects of large models, and amendments or alternative schemes are continually being introduced.

If the models could ever include so many variables that individual thunderstorms and other smaller-scale structures would be properly represented, and parameterization would no longer be needed, it would be totally unreasonable to expect that errors in the details of these structures would require two days to double. Individual thunderstorms typically last only a few hours, and, with an assumed two-day doubling time, the error growth during those hours would be nearly imperceptible. Since a thunderstorm can in reality easily double its severity in less than one hour, we should expect that the difference between two rather similar thunderstorms would double just as rapidly.

If this is the case, the outcome would be that, after several hours, forecasts of the details of small-scale structures would be no better than guesswork, and subsequent representations of their effects on the larger scales would be no better than parameterization. In other words, if we could use such a model with its unbelievably high resolution for perhaps the first half day, we might as well return to one of today's models for the remainder of the forecast. The implication is that introducing such impossibly high resolution would increase the range of practical predictability by only a few hours. As a corollary, it appears that coming improvements in forecasting may have to come from better numerical representations of the structures that are supposedly already resolved, or better formulations of some of the physical processes. The apparent drop in returns with continued increases in resolution has led some forecasters to propose that the anticipated additional computer power in the middle nineties can be more advantageously used to carry out some Monte Carlo procedure.

With all these obstacles around, it may surprise us to learn that within our chaotic atmosphere there are certain weather elements at a few locations that can be rather accurately predicted not just two weeks but two months or even two years ahead. The most spectacularly predictable of these are the winds at high levels in equatorial regions, which are dominated by the so-called quasi-biennial oscillation, first recognized by Richard Reed of the University of Washington. Since the cataclysmic eruption of Krakatau west of Java in 1883, it had been common "knowledge" that the winds at 20 or 25 kilometers above the equator blew from the east; a cloud of volcanic dust had even been observed to circle four times about the globe.

In the 1950s, when sporadic equatorial balloon soundings first reached high enough elevations, a few meteorologists noted that the "normal" so-called Krakatoa easterlies were sometimes missing. I was fortunate enough to be present at the meeting in 1960 when Reed announced his findings, and I could see members of the audience shaking their heads as he maintained that at these heights the equatorial winds would blow continually from the east for about a year, and then from the west for a year, and then from the east again for another year, and that, if Krakatau had blown up a year earlier or later, the meteorological language would have had the term "Krakatoa westerlies."

The subsequent years have fully confirmed his claims. In Figure 34, the plotted points show daily observed values of the eastward component of the wind above Singapore, one degree north

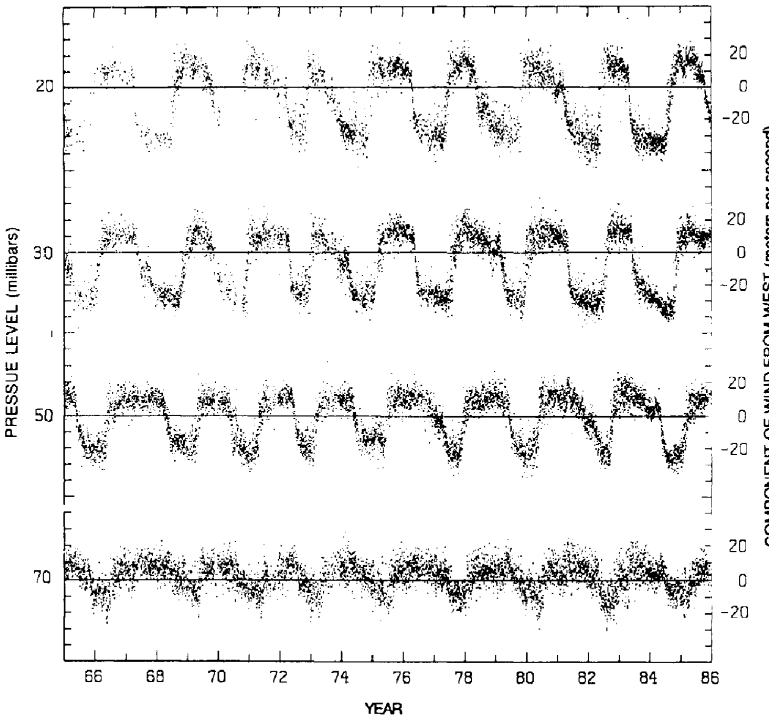


**Figure 34.** The points, some of which are too closely packed to be individually distinguishable, show daily values of the eastward component of the wind at the 70-, 50-, 30-, and 20-millibar surfaces over Singapore, from 1965 through 1985 as indicated by the scale at the base. The solid lines are zero-lines. Values above zero indicate winds from the west. The approximate two-year periodicity is evident.

of the equator, at the four standard pressure levels of 70, 50, 30, and 20 millibars—the pressure is about 1000 millibars at sea level. Above Singapore, these pressures are reached at elevations that fluctuate a few hundred meters about averages of 18.6, 20.6, 23.8, and 26.3 kilometers—roughly twice the height at which commercial jets typically fly. The sequences extend from the beginning of 1965, when the sounding balloons released at Singapore first regularly reached the high levels, to the end of 1985, and they cover nine complete cycles. No smoothing or averaging has been performed, so that the points represent the values that one would attempt to forecast when forecasting for particular moments. Although there are always some day-to-day fluctuations, by far the stronger part of the signal at the upper three levels consists of the oscillation itself, and it is apparent that forecasts with reasonably small errors, for the winds on most of the individual days a cycle or two in advance, can be produced by subjectively extrapolating the phase, and predicting conditions that are average for that phase. As with any other forecasts, these ones will sometimes fail completely, particularly for the times when the rather sharp transitions between westerlies and easterlies will be occurring. Note that the phases propagate downward, taking about a year to descend from 20 to 70 millibars.

The period of oscillation is not exactly two years, as had been conjectured when fewer cycles had been observed, and the separate cycles are not of identical length, so that the oscillation is presumably chaotic, and its phase cannot be predicted decades in advance. Yet the chaos is characterized by an entirely different time scale from that of storms of continental size, just as these storms have a different time scale from thunderstorms. Perhaps the principal lesson is that we still have much to learn about what can happen in chaotic dynamical systems with many interconnected parts.

Other conclusions as to the consequences of chaos in the atmosphere are more speculative, and result from comparing the real world with guesses as to what the world would be like if the weather were not chaotic. In the dishpan we have seen transitions between regimes of symmetric flow, steadily progressing waves, vacillation, and chaos, but I know of no cases in which the flow has assumed an extremely wide variety of patterns during an extended interval before regularly repeating itself. This strongly suggests, although it provides no proof, that if the atmosphere were not in a chaotic regime it would undergo rather simple periodic oscillations not appreciably more complicated than vacillation, with a period of perhaps a few weeks, although the quasi-biennial oscillation, if it could still exist in a nonchaotic regime, could upset things. Any simple behavior would also have to be modulated by the advance of the seasons, so that true repetition would occur only after a year, but each year could be a repetition of all of the previous ones.

Large migratory storms, which are features of both the dishpan and the atmosphere, would undoubtedly be found in our make-believe periodic atmosphere, and, because of the seasonal modulation, successive storms would travel on somewhat different paths. In the course of a year, a considerable portion of the earth's surface might then receive abundant rain, sufficient for agriculture, falling at each location on a particular set of dates. Without the seasons, rainfall would perhaps be confined to a few narrow belts.

At the other extreme in scale, thunderstorms and showers should be abundant enough to strike much of the earth in the course of a year. Hurricanes would be another matter, if they still occurred with a frequency characteristic of the real atmosphere. This could well be the case if the ocean-surface temperatures were comparable to those of the real oceans, since the formation and maintenance of hurricanes is strongly influenced by the temperature of the oceans beneath them. World-wide, a few dozen hurricanes might form during any year. Each might acquire a name, and the same name might be given to the identical hurricane arriving a year later, since the storm would be perceived as an annual event, just as El Nino, the sporadic warm current off the South American coast, is called El Niño whenever it returns. Thus there might have been a Hurricane Amy 1964, or a Hurricane Ben 1977.

Since every named hurricane would be following a track that hurricanes had been following for countless years, there would be little reason to expect much damage. Builders could avoid the paths of the stronger hurricanes, which together would not occupy too much real estate, but it would not be surprising if they built there anyway, presumably taking into account the known maximum wind speeds and the depths of any flash floods. In many respects, planning ahead without the vicissitudes of chaotic behavior would be a much simpler process. The greater difficulty in planning things in the real world, and the occasional disastrous effects of hurricanes and other storms, must therefore be attributed to chaos.

Weather forecasters using twentieth-century methods would not be needed, since this year's weather would be last year's. Meteorologists would still be active, just as tidal theorists are active in the real world, and they would seek explanations for the phenomena that they would be observing with such monotonous regularity. With the global circulation models that might be the fruits of a Global Atmospheric Research Program, modelers might succeed in simulating the significant weather structures, but it is doubtful that such a program would ever be initiated; with no need to improve the process of weather forecasting, who would supply the funds?