




EXCERPTED FROM

STEPHEN  
WOLFRAM  
A NEW  
KIND OF  
SCIENCE

---

CHAPTER 1

*The Foundations for a  
New Kind of Science*



---

# The Foundations for a New Kind of Science

## **An Outline of Basic Ideas**

Three centuries ago science was transformed by the dramatic new idea that rules based on mathematical equations could be used to describe the natural world. My purpose in this book is to initiate another such transformation, and to introduce a new kind of science that is based on the much more general types of rules that can be embodied in simple computer programs.

It has taken me the better part of twenty years to build the intellectual structure that is needed, but I have been amazed by its results. For what I have found is that with the new kind of science I have developed it suddenly becomes possible to make progress on a remarkable range of fundamental issues that have never successfully been addressed by any of the existing sciences before.

If theoretical science is to be possible at all, then at some level the systems it studies must follow definite rules. Yet in the past throughout the exact sciences it has usually been assumed that these rules must be ones based on traditional mathematics. But the crucial realization that led me to develop the new kind of science in this book is that there is in fact no reason to think that systems like those we see in nature should follow only such traditional mathematical rules.

Earlier in history it might have been difficult to imagine what more general types of rules could be like. But today we are surrounded

by computers whose programs in effect implement a huge variety of rules. The programs we use in practice are mostly based on extremely complicated rules specifically designed to perform particular tasks. But a program can in principle follow essentially any definite set of rules. And at the core of the new kind of science that I describe in this book are discoveries I have made about programs with some of the very simplest rules that are possible.

One might have thought—as at first I certainly did—that if the rules for a program were simple then this would mean that its behavior must also be correspondingly simple. For our everyday experience in building things tends to give us the intuition that creating complexity is somehow difficult, and requires rules or plans that are themselves complex. But the pivotal discovery that I made some eighteen years ago is that in the world of programs such intuition is not even close to correct.

I did what is in a sense one of the most elementary imaginable computer experiments: I took a sequence of simple programs and then systematically ran them to see how they behaved. And what I found—to my great surprise—was that despite the simplicity of their rules, the behavior of the programs was often far from simple. Indeed, even some of the very simplest programs that I looked at had behavior that was as complex as anything I had ever seen.

It took me more than a decade to come to terms with this result, and to realize just how fundamental and far-reaching its consequences are. In retrospect there is no reason the result could not have been found centuries ago, but increasingly I have come to view it as one of the more important single discoveries in the whole history of theoretical science. For in addition to opening up vast new domains of exploration, it implies a radical rethinking of how processes in nature and elsewhere work.

Perhaps immediately most dramatic is that it yields a resolution to what has long been considered the single greatest mystery of the natural world: what secret it is that allows nature seemingly so effortlessly to produce so much that appears to us so complex.

It could have been, after all, that in the natural world we would mostly see forms like squares and circles that we consider simple. But in fact one of the most striking features of the natural world is that

across a vast range of physical, biological and other systems we are continually confronted with what seems to be immense complexity. And indeed throughout most of history it has been taken almost for granted that such complexity—being so vastly greater than in the works of humans—could only be the work of a supernatural being.

But my discovery that many very simple programs produce great complexity immediately suggests a rather different explanation. For all it takes is that systems in nature operate like typical programs and then it follows that their behavior will often be complex. And the reason that such complexity is not usually seen in human artifacts is just that in building these we tend in effect to use programs that are specially chosen to give only behavior simple enough for us to be able to see that it will achieve the purposes we want.

One might have thought that with all their successes over the past few centuries the existing sciences would long ago have managed to address the issue of complexity. But in fact they have not. And indeed for the most part they have specifically defined their scope in order to avoid direct contact with it. For while their basic idea of describing behavior in terms of mathematical equations works well in cases like planetary motion where the behavior is fairly simple, it almost inevitably fails whenever the behavior is more complex. And more or less the same is true of descriptions based on ideas like natural selection in biology. But by thinking in terms of programs the new kind of science that I develop in this book is for the first time able to make meaningful statements about even immensely complex behavior.

In the existing sciences much of the emphasis over the past century or so has been on breaking systems down to find their underlying parts, then trying to analyze these parts in as much detail as possible. And particularly in physics this approach has been sufficiently successful that the basic components of everyday systems are by now completely known. But just how these components act together to produce even some of the most obvious features of the overall behavior we see has in the past remained an almost complete mystery. Within the framework of the new kind of science that I develop in this book, however, it is finally possible to address such a question.

From the tradition of the existing sciences one might expect that its answer would depend on all sorts of details, and be quite different for different types of physical, biological and other systems. But in the world of simple programs I have discovered that the same basic forms of behavior occur over and over again almost independent of underlying details. And what this suggests is that there are quite universal principles that determine overall behavior and that can be expected to apply not only to simple programs but also to systems throughout the natural world and elsewhere.

In the existing sciences whenever a phenomenon is encountered that seems complex it is taken almost for granted that the phenomenon must be the result of some underlying mechanism that is itself complex. But my discovery that simple programs can produce great complexity makes it clear that this is not in fact correct. And indeed in the later parts of this book I will show that even remarkably simple programs seem to capture the essential mechanisms responsible for all sorts of important phenomena that in the past have always seemed far too complex to allow any simple explanation.

It is not uncommon in the history of science that new ways of thinking are what finally allow longstanding issues to be addressed. But I have been amazed at just how many issues central to the foundations of the existing sciences I have been able to address by using the idea of thinking in terms of simple programs. For more than a century, for example, there has been confusion about how thermodynamic behavior arises in physics. Yet from my discoveries about simple programs I have developed a quite straightforward explanation. And in biology, my discoveries provide for the first time an explicit way to understand just how it is that so many organisms exhibit such great complexity. Indeed, I even have increasing evidence that thinking in terms of simple programs will make it possible to construct a single truly fundamental theory of physics, from which space, time, quantum mechanics and all the other known features of our universe will emerge.

When mathematics was introduced into science it provided for the first time an abstract framework in which scientific conclusions could be drawn without direct reference to physical reality. Yet despite

all its development over the past few thousand years mathematics itself has continued to concentrate only on rather specific types of abstract systems—most often ones somehow derived from arithmetic or geometry. But the new kind of science that I describe in this book introduces what are in a sense much more general abstract systems, based on rules of essentially any type whatsoever.

One might have thought that such systems would be too diverse for meaningful general statements to be made about them. But the crucial idea that has allowed me to build a unified framework for the new kind of science that I describe in this book is that just as the rules for any system can be viewed as corresponding to a program, so also its behavior can be viewed as corresponding to a computation.

Traditional intuition might suggest that to do more sophisticated computations would always require more sophisticated underlying rules. But what launched the whole computer revolution is the remarkable fact that universal systems with fixed underlying rules can be built that can in effect perform any possible computation.

The threshold for such universality has however generally been assumed to be high, and to be reached only by elaborate and special systems like typical electronic computers. But one of the surprising discoveries in this book is that in fact there are systems whose rules are simple enough to describe in just one sentence that are nevertheless universal. And this immediately suggests that the phenomenon of universality is vastly more common and important—in both abstract systems and nature—than has ever been imagined before.

But on the basis of many discoveries I have been led to a still more sweeping conclusion, summarized in what I call the Principle of Computational Equivalence: that whenever one sees behavior that is not obviously simple—in essentially any system—it can be thought of as corresponding to a computation of equivalent sophistication. And this one very basic principle has a quite unprecedented array of implications for science and scientific thinking.

For a start, it immediately gives a fundamental explanation for why simple programs can show behavior that seems to us complex. For like other processes our own processes of perception and analysis can be

thought of as computations. But though we might have imagined that such computations would always be vastly more sophisticated than those performed by simple programs, the Principle of Computational Equivalence implies that they are not. And it is this equivalence between us as observers and the systems that we observe that makes the behavior of such systems seem to us complex.

One can always in principle find out how a particular system will behave just by running an experiment and watching what happens. But the great historical successes of theoretical science have typically revolved around finding mathematical formulas that instead directly allow one to predict the outcome. Yet in effect this relies on being able to shortcut the computational work that the system itself performs.

And the Principle of Computational Equivalence now implies that this will normally be possible only for rather special systems with simple behavior. For other systems will tend to perform computations that are just as sophisticated as those we can do, even with all our mathematics and computers. And this means that such systems are computationally irreducible—so that in effect the only way to find their behavior is to trace each of their steps, spending about as much computational effort as the systems themselves.

So this implies that there is in a sense a fundamental limitation to theoretical science. But it also shows that there is something irreducible that can be achieved by the passage of time. And it leads to an explanation of how we as humans—even though we may follow definite underlying rules—can still in a meaningful way show free will.

One feature of many of the most important advances in science throughout history is that they show new ways in which we as humans are not special. And at some level the Principle of Computational Equivalence does this as well. For it implies that when it comes to computation—or intelligence—we are in the end no more sophisticated than all sorts of simple programs, and all sorts of systems in nature.

But from the Principle of Computational Equivalence there also emerges a new kind of unity: for across a vast range of systems, from

simple programs to brains to our whole universe, the principle implies that there is a basic equivalence that makes the same fundamental phenomena occur, and allows the same basic scientific ideas and methods to be used. And it is this that is ultimately responsible for the great power of the new kind of science that I describe in this book.

### **Relations to Other Areas**

**Mathematics.** It is usually assumed that mathematics concerns itself with the study of arbitrarily general abstract systems. But this book shows that there are actually a vast range of abstract systems based on simple programs that traditional mathematics has never considered. And because these systems are in many ways simpler in construction than most traditional systems in mathematics it is possible with appropriate methods in effect to go further in investigating them.

Some of what one finds are then just unprecedentedly clear examples of phenomena already known in modern mathematics. But one also finds some dramatic new phenomena. Most immediately obvious is a very high level of complexity in the behavior of many systems whose underlying rules are much simpler than those of most systems in standard mathematics textbooks.

And one of the consequences of this complexity is that it leads to fundamental limitations on the idea of proof that has been central to traditional mathematics. Already in the 1930s Gödel's Theorem gave some indications of such limitations. But in the past they have always seemed irrelevant to most of mathematics as it is actually practiced.

Yet what the discoveries in this book show is that this is largely just a reflection of how small the scope is of what is now considered mathematics. And indeed the core of this book can be viewed as introducing a major generalization of mathematics—with new ideas and methods, and vast new areas to be explored.

The framework I develop in this book also shows that by viewing the process of doing mathematics in fundamentally computational terms it becomes possible to address important issues about the foundations even of existing mathematics.



**Physics.** The traditional mathematical approach to science has historically had its great success in physics—and by now it has become almost universally assumed that any serious physical theory must be based on mathematical equations. Yet with this approach there are still many common physical phenomena about which physics has had remarkably little to say. But with the approach of thinking in terms of simple programs that I develop in this book it finally seems possible to make some dramatic progress. And indeed in the course of the book we will see that some extremely simple programs seem able to capture the essential mechanisms for a great many physical phenomena that have previously seemed completely mysterious.

Existing methods in theoretical physics tend to revolve around ideas of continuous numbers and calculus—or sometimes probability. Yet most of the systems in this book involve just simple discrete elements with definite rules. And in many ways it is the greater simplicity of this underlying structure that ultimately makes it possible to identify so many fundamentally new phenomena.

Ordinary models for physical systems are idealizations that capture some features and ignore others. And in the past what was most common was to capture certain simple numerical relationships—that could for example be represented by smooth curves. But with the new kinds of models based on simple programs that I explore in this book it becomes possible to capture all sorts of much more complex features that can only really be seen in explicit images of behavior.

In the future of physics the greatest triumph would undoubtedly be to find a truly fundamental theory for our whole universe. Yet despite occasional optimism, traditional approaches do not make this seem close at hand. But with the methods and intuition that I develop in this book there is I believe finally a serious possibility that such a theory can actually be found.

**Biology.** Vast amounts are now known about the details of biological organisms, but very little in the way of general theory has ever emerged. Classical areas of biology tend to treat evolution by natural selection as

a foundation—leading to the notion that general observations about living systems should normally be analyzed on the basis of evolutionary history rather than abstract theories. And part of the reason for this is that traditional mathematical models have never seemed to come even close to capturing the kind of complexity we see in biology. But the discoveries in this book show that simple programs can produce a high level of complexity. And in fact it turns out that such programs can reproduce many features of biological organisms—and for example seem to capture some of the essential mechanisms through which genetic programs manage to generate the actual biological forms we see. So this means that it becomes possible to make a wide range of new models for biological systems—and potentially to see how to emulate the essence of their operation, say for medical purposes. And insofar as there are general principles for simple programs, these principles should also apply to biological organisms—making it possible to imagine constructing new kinds of general abstract theories in biology.

**Social Sciences.** From economics to psychology there has been a widespread if controversial assumption—no doubt from the success of the physical sciences—that solid theories must always be formulated in terms of numbers, equations and traditional mathematics. But I suspect that one will often have a much better chance of capturing fundamental mechanisms for phenomena in the social sciences by using instead the new kind of science that I develop in this book based on simple programs. No doubt there will quite quickly be all sorts of claims about applications of my ideas to the social sciences. And indeed the new intuition that emerges from this book may well almost immediately explain phenomena that have in the past seemed quite mysterious. But the very results of the book show that there will inevitably be fundamental limits to the application of scientific methods. There will be new questions formulated, but it will take time before it becomes clear when general theories are possible, and when one must instead inevitably rely on the details of judgement for specific cases.

**Computer Science.** Throughout its brief history computer science has focused almost exclusively on studying specific computational systems set up to perform particular tasks. But one of the core ideas of this book is to consider the more general scientific question of what arbitrary computational systems do. And much of what I have found is vastly different from what one might expect on the basis of existing computer science. For the systems traditionally studied in computer science tend to be fairly complicated in their construction—yet yield fairly simple behavior that recognizably fulfills some particular purpose. But in this book what I show is that even systems with extremely simple construction can yield behavior of immense complexity. And by thinking about this in computational terms one develops a new intuition about the very nature of computation.

One consequence is a dramatic broadening of the domain to which computational ideas can be applied—in particular to include all sorts of fundamental questions about nature and about mathematics. Another consequence is a new perspective on existing questions in computer science—particularly ones related to what ultimate resources are needed to perform general types of computational tasks.

**Philosophy.** At any period in history there are issues about the universe and our role in it that seem accessible only to the general arguments of philosophy. But often progress in science eventually provides a more definite context. And I believe that the new kind of science in this book will do this for a variety of issues that have been considered fundamental even since antiquity. Among them are questions about ultimate limits to knowledge, free will, the uniqueness of the human condition and the inevitability of mathematics. Much has been said over the course of philosophical history about each of these. Yet inevitably it has been informed only by current intuition about how things are supposed to work. But my discoveries in this book lead to radically new intuition. And with this intuition it turns out that one can for the first time begin to see resolutions to many longstanding issues—typically along rather different lines from those expected on the basis of traditional general arguments in philosophy.

**Art.** It seems so easy for nature to produce forms of great beauty. Yet in the past art has mostly just had to be content to imitate such forms. But now, with the discovery that simple programs can capture the essential mechanisms for all sorts of complex behavior in nature, one can imagine just sampling such programs to explore generalizations of the forms we see in nature. Traditional scientific intuition—and early computer art—might lead one to assume that simple programs would always produce pictures too simple and rigid to be of artistic interest. But looking through this book it becomes clear that even a program that may have extremely simple rules will often be able to generate pictures that have striking aesthetic qualities—sometimes reminiscent of nature, but often unlike anything ever seen before.

**Technology.** Despite all its success, there is still much that goes on in nature that seems more complex and sophisticated than anything technology has ever been able to produce. But what the discoveries in this book now show is that by using the types of rules embodied in simple programs one can capture many of the essential mechanisms of nature. And from this it becomes possible to imagine a whole new kind of technology that in effect achieves the same sophistication as nature. Experience with traditional engineering has led to the general assumption that to perform a sophisticated task requires constructing a system whose basic rules are somehow correspondingly complicated. But the discoveries in this book show that this is not the case, and that in fact extremely simple underlying rules—that might for example potentially be implemented directly at the level of atoms—are often all that is needed. My main focus in this book is on matters of basic science. But I have little doubt that within a matter of a few decades what I have done will have led to some dramatic changes in the foundations of technology—and in our basic ability to take what the universe provides and apply it for our own human purposes.

## Some Past Initiatives

My goals in this book are sufficiently broad and fundamental that there have inevitably been previous attempts to achieve at least some of them. But without the ideas and methods of this book there have been basic issues that have eventually ended up presenting almost insuperable barriers to every major approach that has been tried.

**Artificial Intelligence.** When electronic computers were first invented, it was widely believed that it would not be long before they would be capable of human-like thinking. And in the 1960s the field of artificial intelligence grew up with the goal of understanding processes of human thinking and implementing them on computers. But doing this turned out to be much more difficult than expected, and after some spin-offs, little fundamental progress was made. At some level, however, the basic problem has always been to understand how the seemingly simple components in a brain can lead to all the complexities of thinking. But now finally with the framework developed in this book one potentially has a meaningful foundation for doing this. And indeed building on both theoretical and practical ideas in the book I suspect that dramatic progress will eventually be possible in creating technological systems that are capable of human-like thinking.

**Artificial Life.** Ever since machines have existed, people have wondered to what extent they might be able to imitate living systems. Most active from the mid-1980s to the mid-1990s, the field of artificial life concerned itself mainly with showing that computer programs could be made to emulate various features of biological systems. But normally it was assumed that the necessary programs would have to be quite complex. What the discoveries in this book show, however, is that in fact very simple programs can be sufficient. And such programs make the fundamental mechanisms for behavior clearer—and probably come much closer to what is actually happening in real biological systems.

**Catastrophe Theory.** Traditional mathematical models are normally based on quantities that vary continuously. Yet in nature discrete changes are often seen. Popular in the 1970s, catastrophe theory was

concerned with showing that even in traditional mathematical models, certain simple discrete changes could still occur. In this book I do not start from any assumption of continuity—and the types of behavior I study tend to be vastly more complex than those in catastrophe theory.

**Chaos Theory.** The field of chaos theory is based on the observation that certain mathematical systems behave in a way that depends arbitrarily sensitively on the details of their initial conditions. First noticed at the end of the 1800s, this came into prominence after computer simulations in the 1960s and 1970s. Its main significance is that it implies that if any detail of the initial conditions is uncertain, then it will eventually become impossible to predict the behavior of the system. But despite some claims to the contrary in popular accounts, this fact alone does not imply that the behavior will necessarily be complex. Indeed, all that it shows is that if there is complexity in the details of the initial conditions, then this complexity will eventually appear in the large-scale behavior of the system. But if the initial conditions are simple, then there is no reason for the behavior not to be correspondingly simple. What I show in this book, however, is that even when their initial conditions are very simple there are many systems that still produce highly complex behavior. And I argue that it is this phenomenon that is for example responsible for most of the obvious complexity we see in nature.

**Complexity Theory.** My discoveries in the early 1980s led me to the idea that complexity could be studied as a fundamental independent phenomenon. And gradually this became quite popular. But most of the scientific work that was done ended up being based only on my earliest discoveries, and being very much within the framework of one or another of the existing sciences—with the result that it managed to make very little progress on any general and fundamental issues. One feature of the new kind of science that I describe in this book is that it finally makes possible the development of a basic understanding of the general phenomenon of complexity, and its origins.

**Computational Complexity Theory.** Developed mostly in the 1970s, computational complexity theory attempts to characterize how difficult certain computational tasks are to perform. Its concrete results have tended to be based on fairly specific programs with complicated structure yet rather simple behavior. The new kind of science in this book, however, explores much more general classes of programs—and in doing so begins to shed new light on various longstanding questions in computational complexity theory.

**Cybernetics.** In the 1940s it was thought that it might be possible to understand biological systems on the basis of analogies with electrical machines. But since essentially the only methods of analysis available were ones from traditional mathematics, very little of the complex behavior of typical biological systems was successfully captured.

**Dynamical Systems Theory.** A branch of mathematics that began roughly a century ago, the field of dynamical systems theory has been concerned with studying systems that evolve in time according to certain kinds of mathematical equations—and in using traditional geometrical and other mathematical methods to characterize the possible forms of behavior that such systems can produce. But what I argue in this book is that in fact the behavior of many systems is fundamentally too complex to be usefully captured in any such way.

**Evolution Theory.** The Darwinian theory of evolution by natural selection is often assumed to explain the complexity we see in biological systems—and in fact in recent years the theory has also increasingly been applied outside of biology. But it has never been at all clear just why this theory should imply that complexity is generated. And indeed I will argue in this book that in many respects it tends to oppose complexity. But the discoveries in the book suggest a new and quite different mechanism that I believe is in fact responsible for most of the examples of great complexity that we see in biology.

**Experimental Mathematics.** The idea of exploring mathematical systems by looking at data from calculations has a long history, and has gradually become more widespread with the advent of computers and

*Mathematica*. But almost without exception, it has in the past only been applied to systems and questions that have already been investigated by other mathematical means—and that lie very much within the normal tradition of mathematics. My approach in this book, however, is to use computer experiments as a basic way to explore much more general systems—that have never arisen in traditional mathematics, and that are usually far from being accessible by existing mathematical means.

**Fractal Geometry.** Until recently, the only kinds of shapes widely discussed in science and mathematics were ones that are regular or smooth. But starting in the late 1970s, the field of fractal geometry emphasized the importance of nested shapes that contain arbitrarily intricate pieces, and argued that such shapes are common in nature. In this book we will encounter a fair number of systems that produce such nested shapes. But we will also find many systems that produce shapes which are much more complex, and have no nested structure.

**General Systems Theory.** Popular especially in the 1960s, general systems theory was concerned mainly with studying large networks of elements—often idealizing human organizations. But a complete lack of anything like the kinds of methods I use in this book made it almost impossible for any definite conclusions to emerge.

**Nanotechnology.** Growing rapidly since the early 1990s, the goal of nanotechnology is to implement technological systems on atomic scales. But so far nanotechnology has mostly been concerned with shrinking quite familiar mechanical and other devices. Yet what the discoveries in this book now show is that there are all sorts of systems that have much simpler structures, but that can nevertheless perform very sophisticated tasks. And some of these systems seem in many ways much more suitable for direct implementation on an atomic scale.

**Nonlinear Dynamics.** Mathematical equations that have the property of linearity are usually fairly easy to solve, and so have been used extensively in pure and applied science. The field of nonlinear dynamics is concerned with analyzing more complicated equations. Its greatest success has been with so-called soliton equations for which



careful manipulation leads to a property similar to linearity. But the kinds of systems that I discuss in this book typically show much more complex behavior, and have no such simplifying properties.

**Scientific Computing.** The field of scientific computing has usually been concerned with taking traditional mathematical models—most often for various kinds of fluids and solids—and trying to implement them on computers using numerical approximation schemes. Typically it has been difficult to disentangle anything but fairly simple phenomena from effects associated with the approximations used. The kinds of models that I introduce in this book involve no approximations when implemented on computers, and thus readily allow one to recognize much more complex phenomena.

**Self-Organization.** In nature it is quite common to see systems that start disordered and featureless, but then spontaneously organize themselves to produce definite structures. The loosely knit field of self-organization has been concerned with understanding this phenomenon. But for the most part it has used traditional mathematical methods, and as a result has only been able to investigate the formation of fairly simple structures. With the ideas in this book, however, it becomes possible to understand how vastly more complex structures can be formed.

**Statistical Mechanics.** Since its development about a century ago, the branch of physics known as statistical mechanics has mostly concerned itself with understanding the average behavior of systems that consist of large numbers of gas molecules or other components. In any specific instance, such systems often behave in a complex way. But by looking at averages over many instances, statistical mechanics has usually managed to avoid such complexity. To make contact with real situations, however, it has often had to use the so-called Second Law of Thermodynamics, or Principle of Entropy Increase. But for more than a century there have been nagging difficulties in understanding the basis for this principle. With the ideas in this book, however, I believe that there is now a framework in which these can finally be resolved.

## The Personal Story of the Science in This Book


I can trace the beginning of my serious interest in the kinds of scientific issues discussed in this book rather accurately to the summer of 1972, when I was twelve years old. I had bought a copy of the physics textbook on the right, and had become very curious about the process of randomization illustrated on its cover. But being far from convinced by the mathematical explanation given in the book, I decided to try to simulate the process for myself on a computer.

The computer to which I had access at that time was by modern standards a very primitive one. And as a result, I had no choice but to study a very simplified version of the process in the book. I suspected from the start that the system I constructed might be too simple to show any of the phenomena I wanted. And after much programming effort I managed to convince myself that these suspicions were correct.

Yet as it turns out, what I looked at was a particular case of one of the main kinds of systems—cellular automata—that I consider in this book. And had it not been for a largely technical point that arose from my desire to make my simulations as physically realistic as possible, it is quite possible that by 1974 I would already have discovered some of the principal phenomena that I now describe in this book.

As it was, however, I decided at that time to devote my energies to what then seemed to be the most fundamental area of science: theoretical particle physics. And over the next several years I did indeed manage to make significant progress in a few areas of particle physics and cosmology. But after a while I began to suspect that many of the most important and fundamental questions that I was encountering were quite independent of the abstruse details of these fields.

And in fact I realized that there were many related questions even about common everyday phenomena that were still completely unanswered. What for example is the fundamental origin of the complicated patterns that one sees in turbulent fluids? How are the intricate patterns of snowflakes produced? What is the basic mechanism that allows plants and animals to grow in such complex ways?



The book cover that originally sparked my interest in some of the issues discussed in this book.

To my surprise, very little seemed to have been done on these kinds of questions. At first I thought it might be possible to make progress just by applying some of the sophisticated mathematical techniques that I had used in theoretical physics. But it soon became clear that for the phenomena I was studying, traditional mathematical results would be very difficult, if not impossible, to find.

So what could I do? It so happened that as an outgrowth of my work in physics I had in 1981 just finished developing a large software system that was in some respects a forerunner to parts of *Mathematica*. And at least at an intellectual level the most difficult part of the project had been designing the symbolic language on which the system was based. But in the development of this language I had seen rather clearly how just a few primitive operations that I had come up with could end up successfully covering a vast range of sophisticated computational tasks.

So I thought that perhaps I could do something similar in natural science: that there might be some appropriate primitives that I could find that would successfully capture a vast range of natural phenomena. My ideas were not so clearly formed at the time, but I believe I implicitly imagined that the way this would work is that such primitives could be used to build up computer programs that would simulate the various natural systems in which I was interested.

There were in many cases well-established mathematical models for the individual components of such systems. But two practical issues stood in the way of using these as a basis for simulations. First, the models were usually quite complicated, so that with realistic computer resources it was very difficult to include enough components for interesting phenomena to occur. And second, even if one did see such phenomena, it was almost impossible to tell whether in fact they were genuine consequences of the underlying models or were just the result of approximations made in implementing the models on a computer.

But what I realized was that at least for many of the phenomena I wanted to study, it was not crucial to use the most accurate possible models for individual components. For among other things there was evidence from nature that in many cases the details of the components did not matter much—so that for example the same complex patterns

of flow occur in both air and water. And with this in mind, what I decided was that rather than starting from detailed realistic models, I would instead start from models that were somehow as simple as possible—and were easy to set up as programs on a computer.

At the outset, I did not know how this would work, and how complicated the programs I would need would have to be. And indeed when I looked at various simple programs they always seemed to yield behavior vastly simpler than any of the systems I wanted to study.

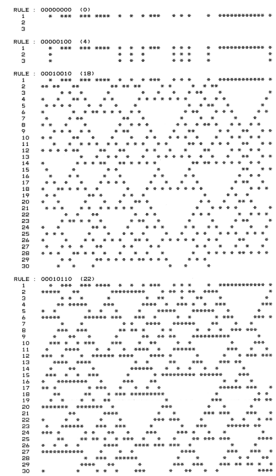
But in the summer of 1981 I did what I considered to be a fairly straightforward computer experiment to see how all programs of a particular type behaved. I had not really expected too much from this experiment. But in fact its results were so surprising and dramatic that as I gradually came to understand them, they forced me to change my whole view of science, and in the end to develop the whole intellectual structure of the new kind of science that I now describe in this book.

The picture on the right shows a reproduction of typical output from my original experiment. The graphics are primitive, but the elaborate patterns they contain were like nothing I had ever seen before. At first I did not believe that they could possibly be correct. But after a while I became convinced that they were—and I realized that I had seen a sign of a quite remarkable and unexpected phenomenon: that even from very simple programs behavior of great complexity could emerge.

But how could something as fundamental as this never have been noticed before? I searched the scientific literature, talked to many people, and found out that systems similar to the ones I was studying had been named “cellular automata” some thirty years earlier. But despite a few close approaches, nobody had ever actually tried anything quite like the type of experiment I had.

Yet I still suspected that the basic phenomenon I had seen must somehow be an obvious consequence of some known scientific principle. But while I did find that ideas from areas like chaos theory and fractal geometry helped in explaining some specific features, nothing even close to the phenomenon as a whole seemed to have ever been studied before.

My early discoveries about the behavior of cellular automata stimulated a fair amount of activity in the scientific community. And



A reproduction of the computer printout that first gave me a hint of some of the central phenomena in this book.

by the mid-1980s, many applications had been found in physics, biology, computer science, mathematics and elsewhere. And indeed some of the phenomena I had discovered were starting to be used as the basis for a new area of research that I called complex systems theory.

Throughout all this, however, I had continued to investigate more basic questions, and by around 1985 I was beginning to realize that what I had seen before was just a hint of something still much more dramatic and fundamental. But to understand what I was discovering was difficult, and required a major shift in intuition.

Yet I could see that there were some remarkable intellectual opportunities ahead. And my first idea was to try to organize the academic community to take advantage of them. So I started a research center and a journal, published a list of problems to attack, and worked hard to communicate the importance of the direction I was defining.

But despite growing excitement—particularly about some of the potential applications—there seemed to be very little success in breaking away from traditional methods and intuition. And after a while I realized that if there was going to be any dramatic progress made, I was the one who was going to have to make it. So I resolved to set up the best tools and infrastructure I could, and then just myself pursue as efficiently as possible the research that I thought should be done.

In the early 1980s my single greatest impediment had been the practical difficulty of doing computer experiments using the various rather low-level tools that were available. But by 1986 I had realized that with a number of new ideas I had it would be possible to build a single coherent system for doing all kinds of technical computing. And since nothing like this seemed likely to exist otherwise, I decided to build it.

The result was *Mathematica*.

For five years the process of building *Mathematica* and the company around it absorbed me. But in 1991—now no longer an academic, but instead the CEO of a successful company—I was able to return to studying the kinds of questions addressed in this book.

And equipped with *Mathematica* I began to try all sorts of new experiments. The results were spectacular—and within the space of a few months I had already made more new discoveries about what

simple programs do than in all the previous ten years put together. My earlier work had shown me the beginnings of some unexpected and very remarkable phenomena. But now from my new experiments I began to see the full force and generality of these phenomena.

As my methodology and intuition improved, the pace of my discoveries increased still more, and within just a couple of years I had managed to take my explorations of the world of simple programs to the point where the sheer volume of factual information I had accumulated would be the envy of many long-established fields of science.

Quite early in the process I had begun to formulate several rather general principles. And the further I went, the more these principles were confirmed, and the more I realized just how strong and general they were.

When I first started at the beginning of the 1980s, my goal was mostly just to understand the phenomenon of complexity. But by the mid-1990s I had built up a whole intellectual structure that was capable of much more, and that in fact provided the foundations for what could only be considered a fundamentally new kind of science.

It was for me a most exciting time. For everywhere I turned there were huge untouched new areas that I was able to explore for the first time. Each had its own particular features. But with the overall framework I had developed I was gradually able to answer essentially all of what seemed to be the most obvious questions that I had raised.

At first I was mostly concerned with new questions that had never been particularly central to any existing areas of science. But gradually I realized that the new kind of science I was building should also provide a fundamentally new way to address basic issues in existing areas.

So around 1994 I began systematically investigating each of the various major traditional areas of science. I had long been interested in fundamental questions in many of these areas. But usually I had tended to believe most of the conventional wisdom about them. Yet when I began to study them in the context of my new kind of science I kept on seeing signs that large parts of this conventional wisdom could not be correct.

The typical issue was that there was some core problem that traditional methods or intuition had never successfully been able to address—and which the field had somehow grown to avoid. Yet over

and over again I was excited to find that with my new kind of science I could suddenly begin to make great progress—even on problems that in some cases had remained unanswered for centuries.

Given the whole framework I had built, many of the things I discovered seemed in the end disarmingly simple. But to get to them often involved a remarkable amount of scientific work. For it was not enough just to be able to take a few specific technical steps. Rather, in each field, it was necessary to develop a sufficiently broad and deep understanding to be able to identify the truly essential features—that could then be rethought on the basis of my new kind of science.

Doing this certainly required experience in all sorts of different areas of science. But perhaps most crucial for me was that the process was a bit like what I have ended up doing countless times in designing *Mathematica*: start from elaborate technical ideas, then gradually see how to capture their essential features in something amazingly simple. And the fact that I had managed to make this work so many times in *Mathematica* was part of what gave me the confidence to try doing something similar in all sorts of areas of science.

Often it seemed in retrospect almost bizarre that the conclusions I ended up reaching had never been reached before. But studying the history of each field I could in many cases see how it had been led astray by the lack of some crucial piece of methodology or intuition that had now emerged in the new kind of science I had developed.

When I made my first discoveries about cellular automata in the early 1980s I suspected that I had seen the beginning of something important. But I had no idea just how important it would all ultimately turn out to be. And indeed over the past twenty years I have made more discoveries than I ever thought possible. And the new kind of science that I have spent so much effort building has seemed an ever more central and critical direction for future intellectual development.

# The Foundations for a New Kind of Science

## An Outline of Basic Ideas

■ **Mathematics in science.** The main event usually viewed as marking the beginning of the modern mathematical approach to science was the publication of Isaac Newton's 1687 book *Mathematical Principles of Natural Philosophy* (the *Principia*). The idea that mathematics might be relevant to science nevertheless had long precursors in both practical and philosophical traditions. Before 500 BC the Babylonians were using arithmetic to describe and predict astronomical data. And by 500 BC the Pythagoreans had come to believe that all natural phenomena should somehow be reducible to relationships between numbers. Many Greek philosophers then discussed the general concept that nature should be amenable to abstract reasoning of the kind used in mathematics. And at a more practical level, the results and methodology of Euclid's work on geometry from around 300 BC became the basis for studies in astronomy, optics and mechanics, notably by Archimedes and Ptolemy. In medieval times there were some doubts about the utility of mathematics in science, and in the late 1200s, for example, Albertus Magnus made the statement that "many of the geometer's figures are not found in natural bodies, and many natural figures, particularly those of animals and plants, are not determinable by the art of geometry". Roger Bacon nevertheless wrote in 1267 that "mathematics is the door and key to the sciences", and by the 1500s it was often believed that for science to be meaningful it must somehow follow the systematic character of mathematics. (Typical of the time was the statement of Leonardo da Vinci that "no human inquiry can be called science unless it pursues its path through mathematical exposition and demonstration".) Around the end of the 1500s Galileo began to develop more explicit connections between concepts in mathematics and in physics, and concluded that the universe could be understood only in the "language of mathematics", whose "characters are triangles, circles and other geometric figures".

What Isaac Newton then did was in effect to suggest that natural systems are at some fundamental level actually governed by purely abstract laws that can be specified in terms of mathematical equations. This idea has met with its greatest success in physics, where for the past three centuries essentially every major theory has been formulated in terms of mathematical equations. Starting in the mid-1800s, it has also had increasing success in chemistry. And in the past century, it has had a few scattered successes in dealing with simpler phenomena in fields like biology and economics. But despite the vast range of phenomena in nature that have never successfully been described in mathematical terms, it has become quite universally assumed that, as David Hilbert put it in 1900, "mathematics is the foundation of all exact knowledge of natural phenomena". There continue to be theories in science that are not explicitly mathematical—examples being continental drift and evolution by natural selection—but, as for example Alfred Whitehead stated in 1911, it is generally believed that "all science as it grows toward perfection becomes mathematical in its ideas".

■ **Definition of mathematics.** When I use the term "mathematics" in this book what I mean is that field of human endeavor that has in practice traditionally been called mathematics. One could in principle imagine defining mathematics to encompass all studies of abstract systems, and indeed this was in essence the definition that I had in mind when I chose the name *Mathematica*. But in practice mathematics has defined itself to be vastly narrower, and to include, for example, nothing like the majority of the programs that I discuss in this book. Indeed, in many respects, what is called mathematics today can be seen as a direct extension of the particular notions of arithmetic and geometry that apparently arose in Babylonian times. Typical dictionary definitions reflect this by describing mathematics as the study of number and space, together with their abstractions and generalizations. And even logic—an abstract system that dates from antiquity—is not normally



considered part of mainstream mathematics. Particularly over the past century the defining characteristic of research in mathematics has increasingly been the use of theorem and proof methodology. And while some generalization has occurred in the types of systems being studied, it has usually been much limited by the desire to maintain the validity of some set of theorems (see page 793). This emphasis on theorems has also led to a focus on equations that statically state facts rather than on rules that define actions, as in most of the systems in this book. But despite all these issues, many mathematicians implicitly tend to assume that somehow mathematics as it is practiced is universal, and that any possible abstract system will be covered by some area of mathematics or another. The results of this book, however, make it quite clear that this is not the case, and that in fact traditional mathematics has reached only a tiny fraction of all the kinds of abstract systems that can in principle be studied.

■ **Reasons for mathematics in science.** It is not surprising that there should be issues in science to which mathematics is relevant, since until about a century ago the whole purpose of mathematics was at some level thought of as being to provide abstract idealizations of aspects of physical reality (with the consequence that concepts like dimensions above 3 and transfinite numbers were not readily accepted as meaningful even in mathematics). But there is absolutely no reason to think that the specific concepts that have arisen so far in the history of mathematics should cover all of science, and indeed in this book I give extensive evidence that they do not. At times the role of mathematics in science has been used in philosophy as an indicator of the ultimate power of human thinking. In the mid-1900s, especially among physicists, there was occasionally some surprise expressed about the effectiveness of mathematics in the natural sciences. One explanation advanced by Albert Einstein was that the only physical laws we can recognize are ones that are easy to express in our system of mathematics.

■ **History of programs and nature.** Given the idea of using programs as a basis for describing nature, one can go back in history and find at least a few rough precursors of this idea. Around 100 AD, for example, following earlier Greek thinking, Lucretius made the somewhat vague suggestion that the universe might consist of atoms assembled according to grammatical rules like letters and words in human language. From the Pythagoreans around 500 BC through Ptolemy around 150 AD to the early work of Johannes Kepler around 1595 there was the notion that the planets might follow definite geometrical rules like the elements of a mechanical clock. But following the work of Isaac Newton in the late 1600s it increasingly came to be believed that systems

could only meaningfully be described by the mathematical equations they satisfy, and not by any explicit mechanism or rules. The failure of the concept of ether and the rise of quantum mechanics in the early 1900s strengthened this view to the point where at least in physics mechanistic explanations of any kind became largely disreputable. (Starting in the 1800s systems based on very simple rules were nevertheless used in studies of genetics and heredity.) With the advent of electronics and computers in the 1940s and 1950s, models like neural networks and cellular automata began to be introduced, primarily in biology (see pages 876 and 1099). But in essentially all cases they were viewed just as approximations to models based on traditional mathematical equations. In the 1960s and 1970s there arose in the early computer hacker community the general idea that the universe might somehow operate like a program. But attempts to engineer explicit features of our universe using constructs from practical programming were unsuccessful, and the idea largely fell into disrepute (see page 1026). Nevertheless, starting in the 1970s many programs were written to simulate all sorts of scientific and technological systems, and often these programs in effect defined the models used. But in almost all cases the elements of the models were firmly based on traditional mathematical equations, and the programs themselves were highly complex, and not much like the simple programs I discuss in this book. (See also pages 363 and 992.)

■ **Extensions of mathematics.** See page 793.

■ **The role of logic.** In addition to standard mathematics, the formal system most widely discussed since antiquity is logic (see page 1099). And starting with Aristotle there was in fact a long tradition of trying to use logic as a framework for drawing conclusions about nature. In the early 1600s the experimental method was suggested as a better alternative. And after mathematics began to show extensive success in describing nature in the late 1600s no further large-scale efforts to do this on the basis of logic appear to have been made. It is conceivable that Gottfried Leibniz might have tried in the late 1600s, but when his work was followed up in the late 1800s by Gottlob Frege and others the emphasis was on building up mathematics, not natural science, from logic (see page 1149). And indeed by this point logic was viewed mostly as a possible representation of human thought—and not as a formal system relevant to nature. So when computers arose it was their numerical and mathematical rather than logical capabilities that were normally assumed relevant for natural science. But in the early 1980s the cellular automata that I studied I often characterized as being based on logical rules, rather than traditional mathematical ones. However, as

we will see on page 806, traditional logic is in fact in many ways very narrow compared to the whole range of rules based on simple programs that I actually consider in this book.

■ **Complexity and theology.** Both complexity and order in the natural world have been cited as evidence for an intelligent creator (compare page 1195). Early mythologies most often assume that the universe started in chaos, with a supernatural being adding order, then creating a series of specific complex natural systems. In Greek philosophy it was commonly thought that the regularities seen in astronomy and elsewhere (such as the obvious circular shapes of the Sun and Moon) were reflections of perfect mathematical forms associated with divine beings. About complexity Aristotle did note that what nature makes is “finer than art”, though this was not central to his arguments about causes of natural phenomena. By the beginning of the Christian era, however, there is evidence of a general belief that the complexity of nature must be the work of a supernatural being—and for example there are statements in the Bible that can be read in this way. Around 1270 Thomas Aquinas gave as an argument for the existence of God the fact that things in nature seem to “act for an end” (as revealed for example by always acting in the same way), and thus must have been specifically designed with that end in mind. In astronomy, as specific natural laws began to be discovered, the role of God began to recede somewhat, with Isaac Newton claiming, for example, that God must have first set the planets on their courses, but then mathematical laws took over to govern their subsequent behavior. Particularly in biology, however, the so-called “argument by design” became ever more popular. Typical was John Ray’s 1691 book *The Wisdom of God Manifested in the Works of the Creation*, which gave a long series of examples from biology that it claimed were so complex that they must be the work of a supernatural being. By the early 1800s, such ideas had led to the field of natural theology, and William Paley gave the much quoted argument that if it took a sophisticated human watchmaker to construct a watch, then the only plausible explanation for the vastly greater complexity of biological systems was that they must have been created by a supernatural being. Following the publication of Charles Darwin’s *Origin of Species* in 1859 many scientists began to argue that natural selection could explain all the basic phenomena of biology, and although some religious groups maintained strong resistance, it was widely assumed by the mid-1900s that no other explanation was needed. In fact, however, just how complexity arises was never really resolved, and in the end I believe that it is only with the ideas of this book that this can successfully be done.

■ **Artifacts and natural systems.** See page 828.

■ **Complexity and science.** Ever since antiquity science has tended to see its main purpose as being the study of regularities—and this has meant that insofar as complexity is viewed as an absence of regularities, it has tended to be ignored or avoided. There have however been occasional discussions of various general aspects of complexity and what can account for them. Thus, for example, by 200 BC the Epicureans were discussing the idea that varied and complex forms in nature could be made up from arrangements of small numbers of types of elementary atoms in much the same way as varied and complex written texts are made up from small numbers of types of letters. And although its consequences were remarkably confused, the notion of a single underlying substance that could be transmuted into anything—living or not—was also a centerpiece of alchemy. Starting in the 1600s successes in physics and discoveries like the circulation of blood led to the idea that it should be possible to explain the operation of almost any natural system in essentially mechanical terms—leading for example René Descartes to claim in 1637 that we should one day be able to explain the operation of a tree just like we do a clock. But as mathematical methods developed, they seemed to apply mainly to physical systems, and not for example to biological ones. And indeed Immanuel Kant wrote in 1790 that “it is absurd to hope that another Newton will arise in the future who will make comprehensible to us the production of a blade of grass according to natural laws”. In the late 1700s and early 1800s mathematical methods began to be used in economics and later in studying populations. And partly influenced by results from this, Charles Darwin in 1859 suggested natural selection as the basis for many phenomena in biology, including complexity. By the late 1800s advances in chemistry had established that biological systems were made of the same basic components as physical ones. But biology still continued to concentrate on very specific observations—with no serious theoretical discussion of anything as general as the phenomenon of complexity. In the 1800s statistics was increasingly viewed as providing a scientific approach to complex processes in practical social systems. And in the late 1800s statistical mechanics was then used as a basis for analyzing complex microscopic processes in physics. Most of the advances in physics in the late 1800s and early 1900s in effect avoided complexity by concentrating on properties and systems simple enough to be described by explicit mathematical formulas. And when other fields tried in the early and mid-1900s to imitate successes in physics, they too generally tended to concentrate on issues that seemed amenable to explicit mathematical

formulas. Within mathematics itself—especially in number theory and the three-body problem—there were calculations that yielded results that seemed complex. But normally this complexity was viewed just as something to be overcome—either by looking at things in a different way, or by proving more powerful theorems—and not as something to be studied or even much commented on in its own right.

In the 1940s, however, successes in the analysis of logistical and electronic systems led to discussion of the idea that it might be possible to set up some sort of general approach to complex systems—especially biological and social ones. And by the late 1940s the cybernetics movement was becoming increasingly popular—with Norbert Wiener emphasizing feedback control and stochastic differential equations, and John von Neumann and others emphasizing systems based on networks of elements often modelled after neurons. There were spinoffs such as control theory and game theory, but little progress was made on core issues of complexity, and already by the mid-1950s what began to dominate were vague discussions involving fashionable issues in areas such as psychiatry and anthropology. There also emerged a tradition of robotics and artificial intelligence, and a few of the systems that were built or simulated did show some complexity of behavior (see page 879). But in most cases this was viewed just as something to be overcome in order to achieve the engineering objectives sought. Particularly in the 1960s there was discussion of complexity in large human organizations—especially in connection with the development of management science and the features of various forms of hierarchy—and there emerged what was called systems theory, which in practice typically involved simulating networks of differential equations, often representing relationships in flowcharts. Attempts were for example made at worldwide models, but by the 1970s their results—especially in economics—were being discredited. (Similar methods are nevertheless used today, especially in environmental modelling.)

With its strong emphasis on simple laws and measurements of numbers, physics has normally tended to define itself to avoid complexity. But from at least the 1940s, issues of complexity were nevertheless occasionally mentioned by physicists as important, most often in connection with fluid turbulence or features of nonlinear differential equations. Questions about pattern formation, particularly in biology and in relation to thermodynamics, led to a sequence of studies of reaction-diffusion equations, which by the 1970s were being presented as relevant to general issues of complexity, under names like self-organization, synergetics and dissipative structures. By the late 1970s the work of

Benoit Mandelbrot on fractals provided an important example of a general approach to addressing a certain kind of complexity. And chaos theory—with its basis in the mathematics of dynamical systems theory—also began to become popular in the late 1970s, being discussed particularly in connection with fluid turbulence. In essentially all cases, however, the emphasis remained on trying to find some aspect of complex behavior that could be summarized by a single number or a traditional mathematical equation.

As discussed on pages 44–50, there were by the beginning of the 1980s various kinds of abstract systems whose rules were simple but which had nevertheless shown complex behavior, particularly in computer simulations. But usually this was considered largely a curiosity, and there was no particular sense that there might be a general phenomenon of complexity that could be of central interest, say in natural science. And indeed there remained an almost universal belief that to capture any complexity of real scientific relevance one must have a complex underlying model. My work on cellular automata in the early 1980s provided strong evidence, however, that complex behavior very much like what was seen in nature could in fact arise in a very general way from remarkably simple underlying rules. And starting around the mid-1980s it began to be not uncommon to hear the statement that complex behavior can arise from simple rules—though often there was great confusion about just what this was actually saying, and what, for example, should be considered complex behavior, or a simple rule.

That complexity could be identified as a coherent phenomenon that could be studied scientifically in its own right was something I began to emphasize around 1984. And having created the beginnings of what I considered to be the necessary intellectual structure, I started to try to develop an organizational structure to allow what I called complex systems research to spread. Some of what I did had fairly immediate effects, but much did not, and by late 1986 I had started building *Mathematica* and decided to pursue my own scientific interests in a more independent way (see page 20). By the late 1980s, however, there was widespread discussion of what was by then being called complexity theory. (I had avoided this name to prevent confusion with the largely unrelated field of computational complexity theory). And indeed many of the points I had made about the promise of the field were being enthusiastically repeated in popular accounts—and there were starting to be quite a number of new institutions devoted to the field. (A notable example was the Santa Fe Institute, whose orientation towards complexity seems to have been a quite direct consequence of my efforts.)

But despite all this, no major new scientific developments were forthcoming—not least because there was a tremendous tendency to ignore the idea of simple underlying rules and of what I had discovered in cellular automata, and instead to set up computer simulations with rules far too complicated to allow them to be used in studying fundamental questions. And combined with a predilection for considering issues in the social and biological sciences that seem hard to pin down, this led to considerable skepticism among many scientists—with the result that by the mid-1990s the field was to some extent in retreat—though the statement that complexity is somehow an important and fundamental issue has continued to be emphasized especially in studies of ecological and business systems.

Watching the history of the field of complexity theory has made it particularly clear to me that without a major new intellectual structure complexity cannot realistically be studied in a meaningful scientific way. But it is now just such a structure that I believe I have finally been able to set up in this book.

### Relations to Other Areas

■ **Page 7 • Mathematics.** I discuss the implications of this book for the foundations of mathematics mainly on pages 772–821 and in the rather extensive corresponding notes. With a sufficiently general definition of mathematics, however, the whole core of the book can in fact be viewed as a work of experimental mathematics. And even with a more traditional definition, this is at least true of much of my discussion of systems based on numbers in Chapter 4. The notes to almost all chapters of the book contain a great many new mathematical results, mostly emerging from my analysis of some of the simpler behavior considered in the book. Pages 606–620 and 737–750 discuss in general the capabilities of mathematical analysis, while pages 588–597 address the foundations of statistics. Note that some ideas and results highly relevant to current frontiers in mathematics appear in some rather unexpected places in the book. Specific examples include the parameter space sets that I discuss in connection with shapes of plant leaves on page 407, and the minimal axioms for logic that I discuss on page 810. A more general example is the issue of smooth objects arising from combinatorial data that I discuss in Chapter 9 in connection with the nature of space in fundamental physics.

■ **Page 8 • Physics.** I discuss general mechanisms and models relevant for physical systems in Chapter 7, specific types of everyday physical systems in Chapter 8, and applications to basic foundational problems in physics in Chapter 9. I

mention some further fundamental issues in physics around page 730 and in chemistry on page 1193.

■ **Page 8 • Biology.** The main place I discuss applications to biology is on pages 383–429 of Chapter 8, where I consider first general questions about biology and evolution, and then more specific issues about growth and pattern in biological organisms. I consider visual and auditory perception on pages 577–588, and the operation of brains on pages 620–631. I also discuss the definition of life on pages 823 and 1178, as well as mentioning protein folding and structure on pages 1003 and 1184.

■ **Page 9 • Social and related sciences.** I discuss the particular example of financial systems on pages 429–432, and make some general comments on page 1014. The end of Chapter 10, as well as some parts of Chapter 12, also discuss various issues that can be viewed as foundational questions.

■ **Page 10 • Computer science.** Chapter 11 as well as parts of Chapter 12 (especially pages 753–771) address foundational issues in computer science. Chapter 3 uses standard computer science models such as Turing machines and register machines as examples of simple programs. In many places in the book—especially these notes—I discuss all sorts of specific problems and issues of direct relevance to current computer science. Examples include cryptography (pages 598–606), Boolean functions (pages 616–619 and 806–814), user interfaces (page 1102) and quantum computing (page 1147).

■ **Page 10 • Philosophy.** Chapter 12 is the main place I address traditional philosophical issues. On pages 363–369 of Chapter 8, however, I discuss some general issues of modelling, and in Chapter 10 I consider at length not only practical but also foundational questions about perception and to some extent general thinking and consciousness. (See page 1196.)

■ **Page 11 • Technology.** The notes to this book mention many specific technological connections, and I expect that many of the models and methods of analysis that I use in the book can be applied quite directly for technological purposes. I discuss foundational questions about technology mainly on pages 840–843.

■ **Scope of existing sciences.** One might imagine that physics would for example concern itself with all aspects of physical systems, biology with all aspects of biological systems, and so on. But in fact as they are actually practiced most of the traditional sciences are much narrower in scope. Historically what has typically happened is that in each science a certain way of thinking has emerged as the most successful. And then over the course of time, the scope of the science itself has come to be defined to encompass just those issues that this

way of thinking is able to address. So when a new phenomenon is observed, a particular science will typically tend to focus on just those aspects of the phenomenon that can be studied by whatever way of thinking has been adopted in that science. And when the phenomenon involves substantial complexity, what has in the past usually happened is that simpler and simpler aspects are investigated until one is found that is simple enough to analyze using the chosen way of thinking.

### The Personal Story of the Science in This Book

■ **Page 17 · Statistical physics cover.** The pictures show disks representing idealized molecules bouncing around in a box, and the book claims that as time goes on there is almost inevitably increasing randomization. The pictures were made in about 1964 by Berni Alder and Frederick Reif from oscilloscope output from the LARC computer at what was then Lawrence Radiation Laboratory. A total of 40 disks were started with positions and velocities determined by a middle-square random number generator (see page 975), and their motion was followed for about 10 collision times—after which roundoff errors in the 64-bit numbers used had grown too big. From the point of view of this book the randomization seen in these pictures is in large part just a reflection of the fact that a random sequence of digits were used in the initial conditions. But what the discoveries in this book show is that such randomness can also be generated

without any such random input—finally clarifying some very basic issues in statistical physics. (See page 441.)

■ **Page 17 · My 1973 computer experiments.** I used a British Elliott 903 computer with 8 kilowords of 18-bit ferrite core memory. The assembly language program that I wrote filled up a fair fraction of the memory. The system that I looked at was a 2D cellular automaton with discrete particles colliding on a square grid. Had I not been concerned with physics-like conservation laws, or had I used something other than a square grid, the teleprinter output that I generated would have shown randomization. (See page 999.)

■ **Page 19 · Computer printouts.** The printouts show a series of elementary cellular automata started from random initial conditions (see page 232). I generated them in 1981 using a C program running on a VAX 11/780 computer with an early version of the Unix operating system. (See also page 880.)

■ **Timeline.** Major periods in my work have been:

- 1974–1980: particle physics and cosmology
- 1979–1981: developing SMP computer algebra system
- 1981–1986: cellular automata etc.
- 1986–1991: intensive *Mathematica* development
- 1991–2001: writing this book

(Wolfram Research, Inc. was founded in 1987; *Mathematica* 1.0 was released June 23, 1988; the company and successive versions of *Mathematica* continue to be major parts of my life.)

■ **Detailed history.** See pages 880–882.