EXCERPTED FROM

STEPHEN WOLFRAM

A NEW KIND OF SCIENCE

SECTION 1.4

The Personal Story of the Science in This Book

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.

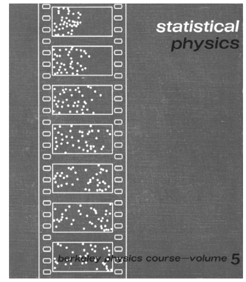

The book cover that originally sparked my interest in some of the issues discussed in this book.

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?

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.

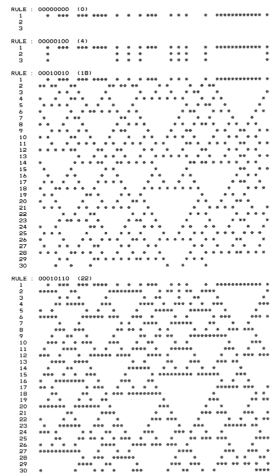

A reproduction of the computer printout that first gave me a hint of some of the central phenomena in this book.

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 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.