EXCERPTED FROM

STEPHEN WOLFRAM
A NEW KIND OF SCIENCE

SECTION 2.1

How Do Simple Programs Behave?

2

The Crucial Experiment

How Do Simple Programs Behave?

New directions in science have typically been initiated by certain central observations or experiments. And for the kind of science that I describe in this book these concerned the behavior of simple programs.

In our everyday experience with computers, the programs that we encounter are normally set up to perform very definite tasks. But the key idea that I had nearly twenty years ago—and that eventually led to the whole new kind of science in this book—was to ask what happens if one instead just looks at simple arbitrarily chosen programs, created without any specific task in mind. How do such programs typically behave?

The mathematical methods that have in the past dominated theoretical science do not help much with such a question. But with a computer it is straightforward to start doing experiments to investigate it. For all one need do is just set up a sequence of possible simple programs, and then run them and see how they behave.

Any program can at some level be thought of as consisting of a set of rules that specify what it should do at each step. There are many possible ways to set up these rules—and indeed we will study quite a few of them in the course of this book. But for now, I will consider a particular class of examples called cellular automata, that were the very first kinds of simple programs that I investigated in the early 1980s.

An important feature of cellular automata is that their behavior can readily be presented in a visual way. And so the picture below shows what one cellular automaton does over the course of ten steps.

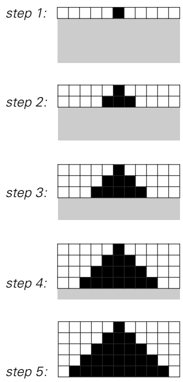

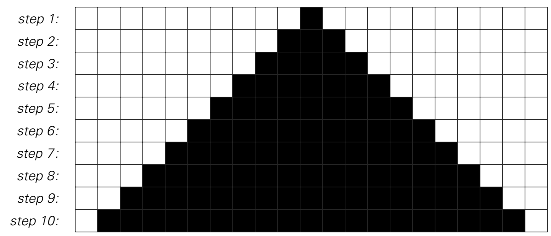

A visual representation of the behavior of a cellular automaton, with each row of cells corresponding to one step. At the first step the cell in the center is black and all other cells are white. Then on each successive step, a particular cell is made black whenever it or either of its neighbors were black on the step before. As the picture shows, this leads to a simple expanding pattern uniformly filled with black.

The cellular automaton consists of a line of cells, each colored either black or white. At every step there is then a definite rule that determines the color of a given cell from the color of that cell and its immediate left and right neighbors on the step before.

For the particular cellular automaton shown here the rule specifies—as in the picture below—that a cell should be black in all cases where it or either of its neighbors were black on the step before.

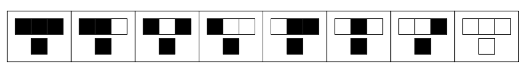

A representation of the rule for the cellular automaton shown above. The top row in each box gives one of the possible combinations of colors for a cell and its immediate neighbors. The bottom row then specifies what color the center cell should be on the next step in each of these cases. In the numbering scheme described in Chapter 3, this is cellular automaton rule 254.

And the picture at the top of the page shows that starting with a single black cell in the center this rule then leads to a simple growing pattern uniformly filled with black. But modifying the rule just slightly one can immediately get a different pattern.

As a first example, the picture at the top of the facing page shows what happens with a rule that makes a cell white whenever both of its neighbors were white on the step before—even if the cell itself was black before. And rather than producing a pattern that is uniformly filled with black, this rule now instead gives a pattern that repeatedly alternates between black and white like a checkerboard.

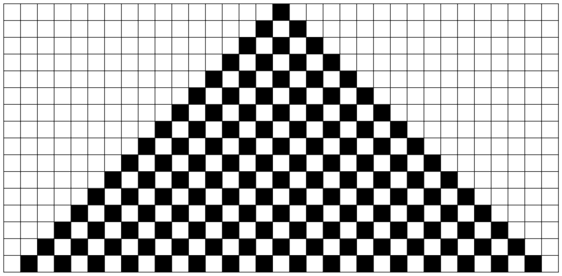

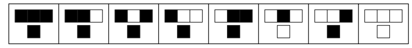

A cellular automaton with a slightly different rule. The rule makes a particular cell black if either of its neighbors was black on the step before, and makes the cell white if both its neighbors were white. Starting from a single black cell, this rule leads to a checkerboard pattern. In the numbering scheme of Chapter 3, this is cellular automaton rule 250.

This pattern is however again fairly simple. And we might assume that at least with the type of cellular automata that we are considering, any rule we might choose would always give a pattern that is quite simple. But now we are in for our first surprise.

The picture below shows the pattern produced by a cellular automaton of the same type as before, but with a slightly different rule.

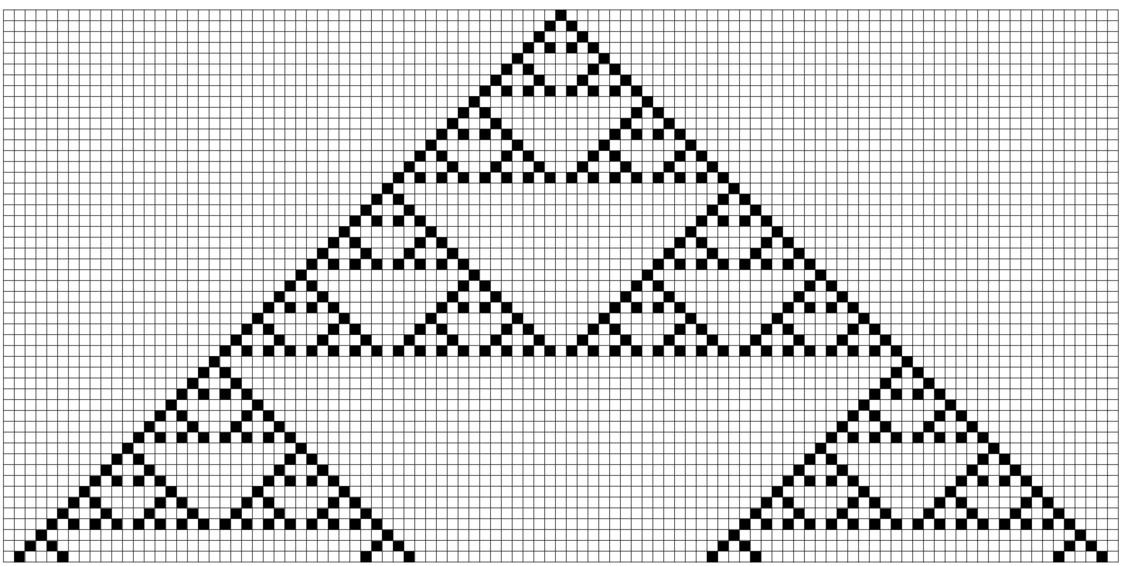

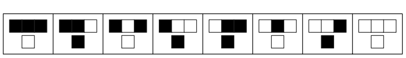

A cellular automaton that produces an intricate nested pattern. The rule in this case is that a cell should be black whenever one or the other, but not both, of its neighbors were black on the step before. Even though the rule is very simple, the picture shows that the overall pattern obtained over the course of 50 steps starting from a single black cell is not so simple. The particular rule used here can be described by the formula $a_i' = Mod[a_{i-1} + a_{i+1}, 2]$. In the numbering scheme of Chapter 3, it is cellular automaton rule 90.

This time the rule specifies that a cell should be black when either its left neighbor or its right neighbor—but not both—were black on the step before. And again this rule is undeniably quite simple. But now the picture shows that the pattern it produces is not so simple.

And if one runs the cellular automaton for more steps, as in the picture below, then a rather intricate pattern emerges. But one can now see that this pattern has very definite regularity. For even though it is intricate, one can see that it actually consists of many nested triangular pieces that all have exactly the same form. And as the picture shows, each of these pieces is essentially just a smaller copy of the whole pattern, with still smaller copies nested in a very regular way inside it.

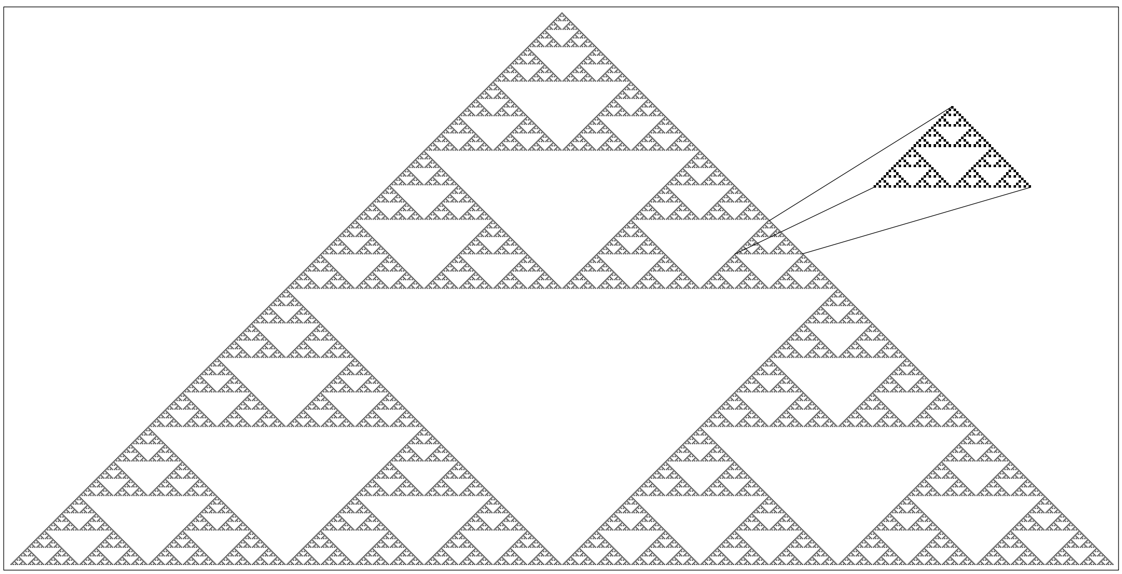

A larger version of the pattern from the previous page, now shown without a grid explicitly indicating each cell. The picture shows five hundred steps of cellular automaton evolution. The pattern obtained is intricate, but has a definite nested structure. Indeed, as the picture illustrates, each triangular section is essentially just a smaller copy of the whole pattern, with still smaller copies nested inside it. Patterns with nested structure of this kind are often called "fractal" or "self-similar".

So of the three cellular automata that we have seen so far, all ultimately yield patterns that are highly regular: the first a simple uniform pattern, the second a repetitive pattern, and the third an intricate but still nested pattern. And we might assume that at least for cellular automata with rules as simple as the ones we have been using these three forms of behavior would be all that we could ever get.

But the remarkable fact is that this turns out to be wrong.

And the picture below shows an example of this. The rule used—that I call rule 30—is of exactly the same kind as before, and can be described as follows. First, look at each cell and its right-hand neighbor. If both of these were white on the previous step, then take the new color of the cell to be whatever the previous color of its left-hand neighbor was. Otherwise, take the new color to be the opposite of that.

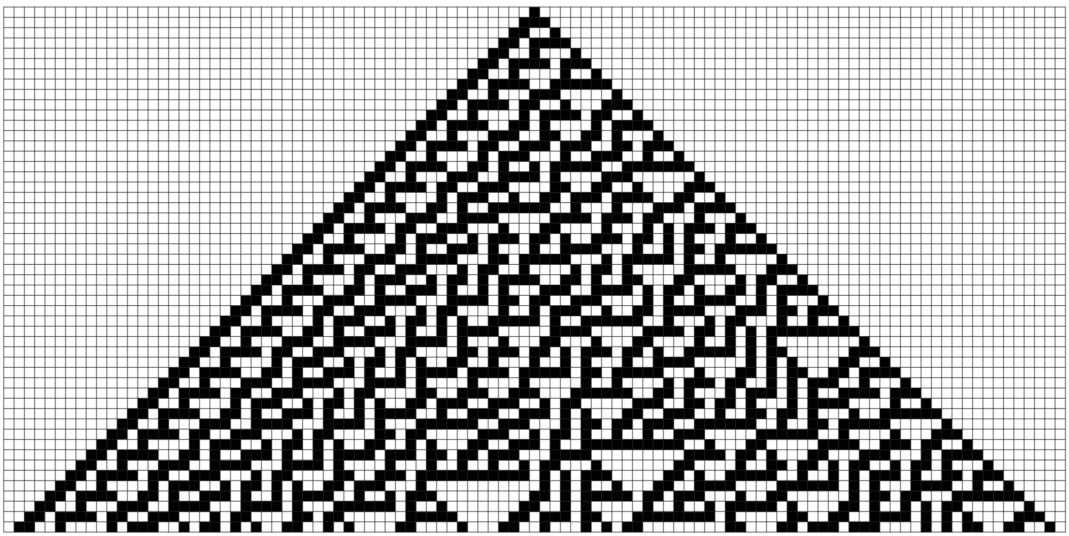

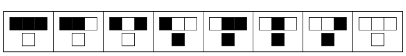

A cellular automaton with a simple rule that generates a pattern which seems in many respects random. The rule used is of the same type as in the previous examples, and the cellular automaton is again started from a single black cell. But now the pattern that is obtained is highly complex, and shows almost no overall regularity. This picture is our first example of the fundamental phenomenon that even with simple underlying rules and simple initial conditions, it is possible to produce behavior of great complexity. In the numbering scheme of Chapter 3, the cellular automaton shown here is rule 30.

The picture shows what happens when one starts with just one black cell and then applies this rule over and over again. And what one sees is something quite startling—and probably the single most surprising scientific discovery I have ever made. Rather than getting a simple regular pattern as we might expect, the cellular automaton instead produces a pattern that seems extremely irregular and complex.

But where does this complexity come from? We certainly did not put it into the system in any direct way when we set it up. For we just used a simple cellular automaton rule, and just started from a simple initial condition containing a single black cell.

Yet the picture shows that despite this, there is great complexity in the behavior that emerges. And indeed what we have seen here is a first example of an extremely general and fundamental phenomenon that is at the very core of the new kind of science that I develop in this book. Over and over again we will see the same kind of thing: that even though the underlying rules for a system are simple, and even though the system is started from simple initial conditions, the behavior that the system shows can nevertheless be highly complex. And I will argue that it is this basic phenomenon that is ultimately responsible for most of the complexity that we see in nature.

The next two pages show progressively more steps in the evolution of the rule 30 cellular automaton from the previous page. One might have thought that after maybe a thousand steps the behavior would eventually resolve into something simple. But the pictures on the next two pages show that nothing of the sort happens.

Some regularities can nevertheless be seen. On the left-hand side, for example, there are obvious diagonal bands. And dotted throughout there are various white triangles and other small structures. Yet given the simplicity of the underlying rule, one would expect vastly more regularities. And perhaps one might imagine that our failure to see any in the pictures on the next two pages is just a reflection of some kind of inadequacy in the human visual system.

But it turns out that even the most sophisticated mathematical and statistical methods of analysis seem to do no better. For example, one can look at the sequence of colors directly below the initial black cell. And in the first million steps in this sequence, for example, it never repeats, and indeed none of the tests I have ever done on it show any meaningful deviation at all from perfect randomness.

In a sense, however, there is a certain simplicity to such perfect randomness. For even though it may be impossible to predict what

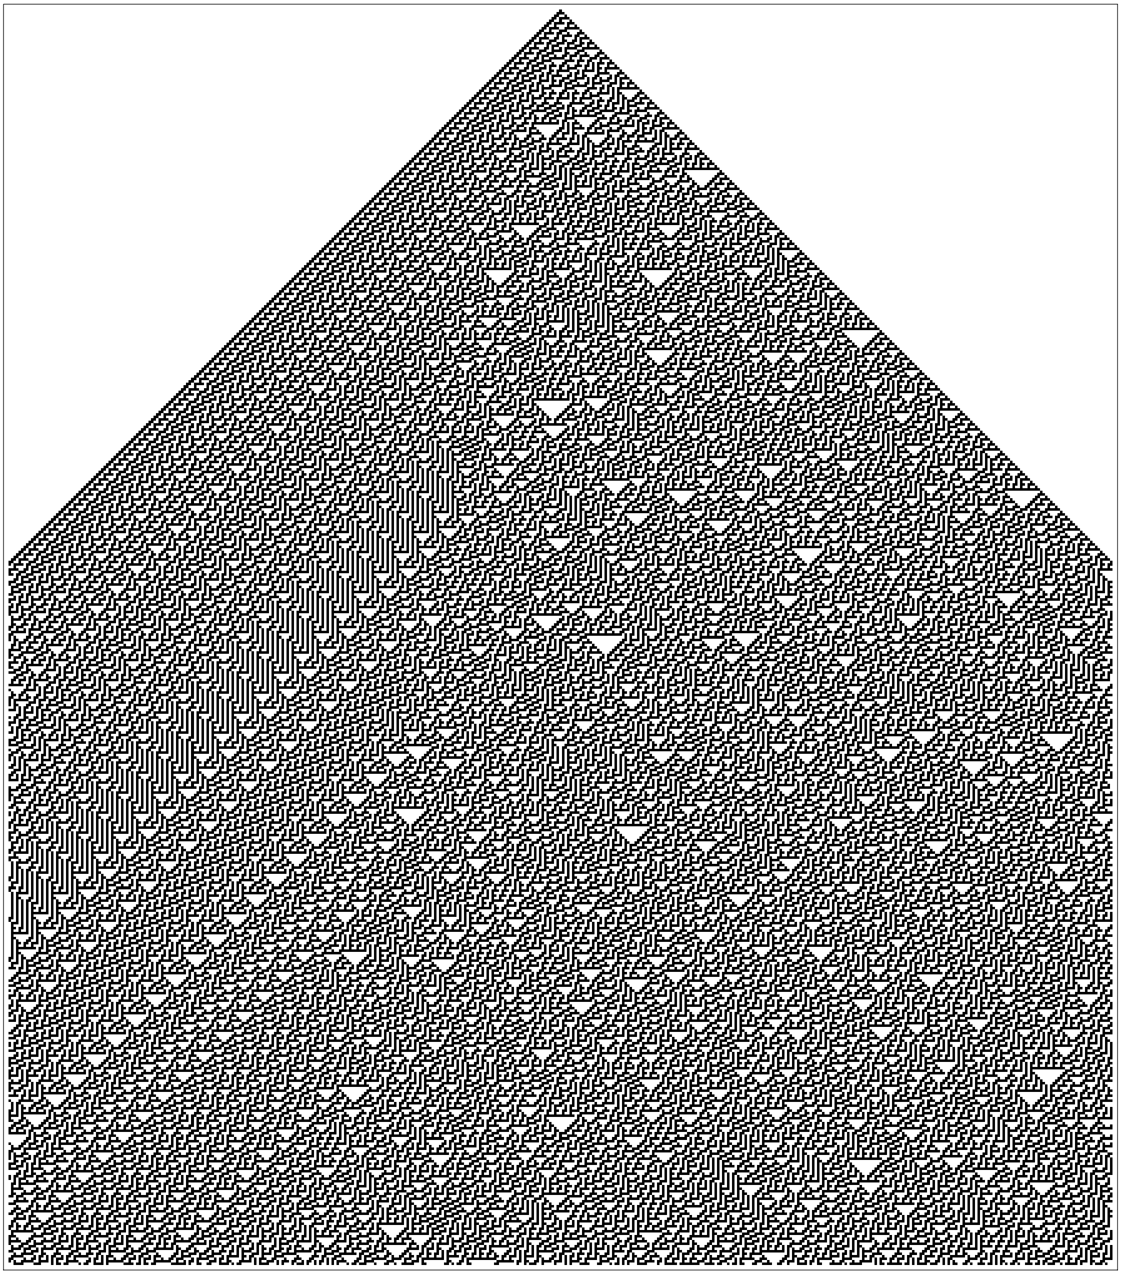

Five hundred steps in the evolution of the rule 30 cellular automaton from page 27. The pattern produced continues to expand on both left and right, but only the part that fits across the page is shown here. The asymmetry between the left and right-hand sides is a direct consequence of asymmetry that exists in the particular underlying cellular automaton rule used.

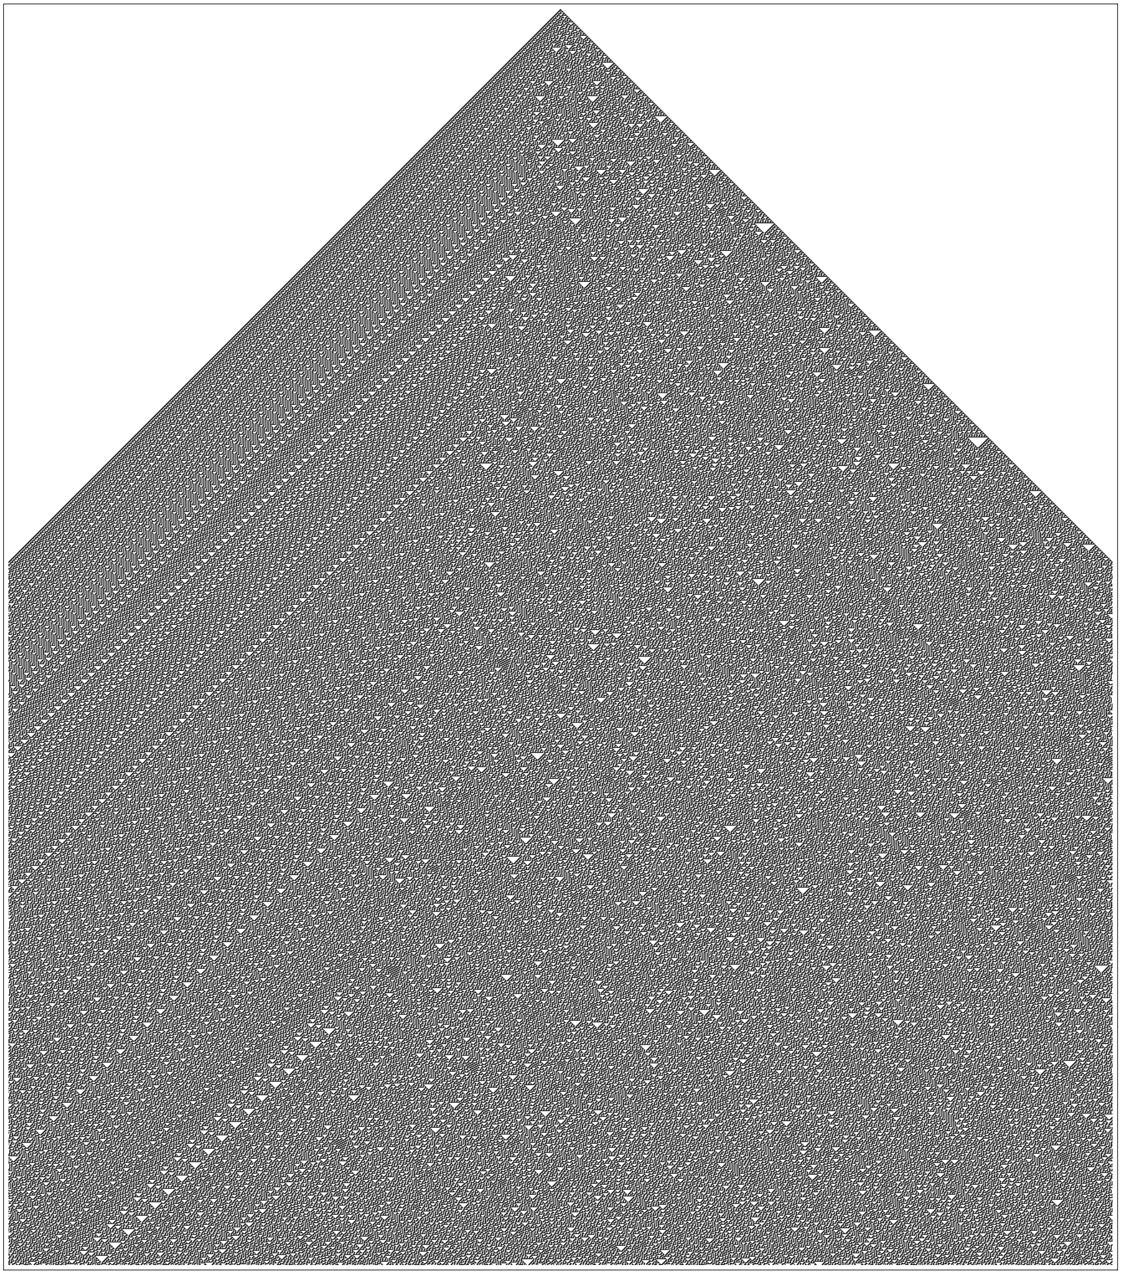

Fifteen hundred steps of rule 30 evolution. Some regularities are evident, particularly on the left. But even after all these steps there are no signs of overall regularity—and indeed even continuing for a million steps many aspects of the pattern obtained seem perfectly random according to standard mathematical and statistical tests. The picture here shows a total of just under two million individual cells.

color will occur at any specific step, one still knows for example that black and white will on average always occur equally often.

But it turns out that there are cellular automata whose behavior is in effect still more complex—and in which even such averages become very difficult to predict. The pictures on the next several pages give a rather dramatic example. The basic form of the rule is just the same as before. But now the specific rule used—that I call rule 110—takes the new color of a cell to be black in every case except when the previous colors of the cell and its two neighbors were all the same, or when the left neighbor was black and the cell and its right neighbor were both white.

The pattern obtained with this rule shows a remarkable mixture of regularity and irregularity. More or less throughout, there is a very regular background texture that consists of an array of small white triangles repeating every 7 steps. And beginning near the left-hand edge, there are diagonal stripes that occur at intervals of exactly 80 steps.

But on the right-hand side, the pattern is much less regular. Indeed, for the first few hundred steps there is a region that seems essentially random. But by the bottom of the first page, all that remains of this region is three copies of a rather simple repetitive structure.

Yet at the top of the next page the arrival of a diagonal stripe from the left sets off more complicated behavior again. And as the system progresses, a variety of definite localized structures are produced.

Some of these structures remain stationary, like those at the bottom of the first page, while others move steadily to the right or left at various speeds. And on their own, each of these structures works in a fairly simple way. But as the pictures illustrate, their various interactions can have very complicated effects.

And as a result it becomes almost impossible to predict—even approximately—what the cellular automaton will do.

Will all the structures that are produced eventually annihilate each other, leaving only a very regular pattern? Or will more and more structures appear until the whole pattern becomes quite random?

The only sure way to answer these questions, it seems, is just to run the cellular automaton for as many steps as are needed, and to

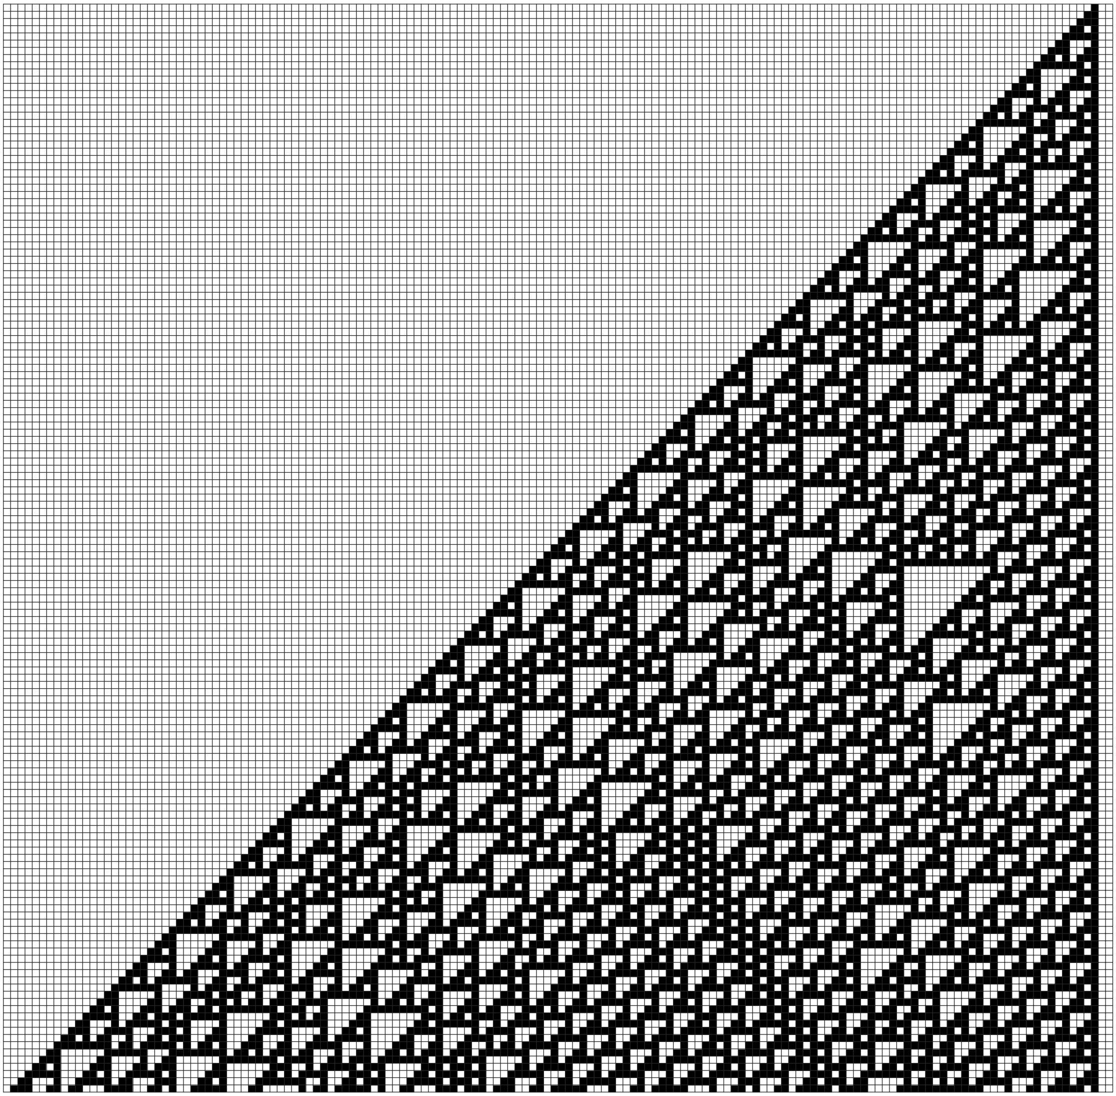

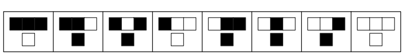

A cellular automaton whose behavior seems neither highly regular nor completely random. The picture is obtained by applying the simple rule shown for a total of 150 steps, starting with a single black cell. Note that the particular rule used here yields a pattern that expands on the left but not on the right. In the scheme defined in Chapter 3, the rule is number 110.

More steps in the pattern shown above. Each successive page shows a total of 700 steps. The pattern continues to expand on the left forever, but only the part that fits across each page is shown. For a long time it is not clear how the right-hand part of the pattern will eventually look. But after 2780 steps, a fairly simple repetitive structure emerges. Note that to generate the pictures that follow requires applying the underlying cellular automaton rule for individual cells a total of about 12 million times.

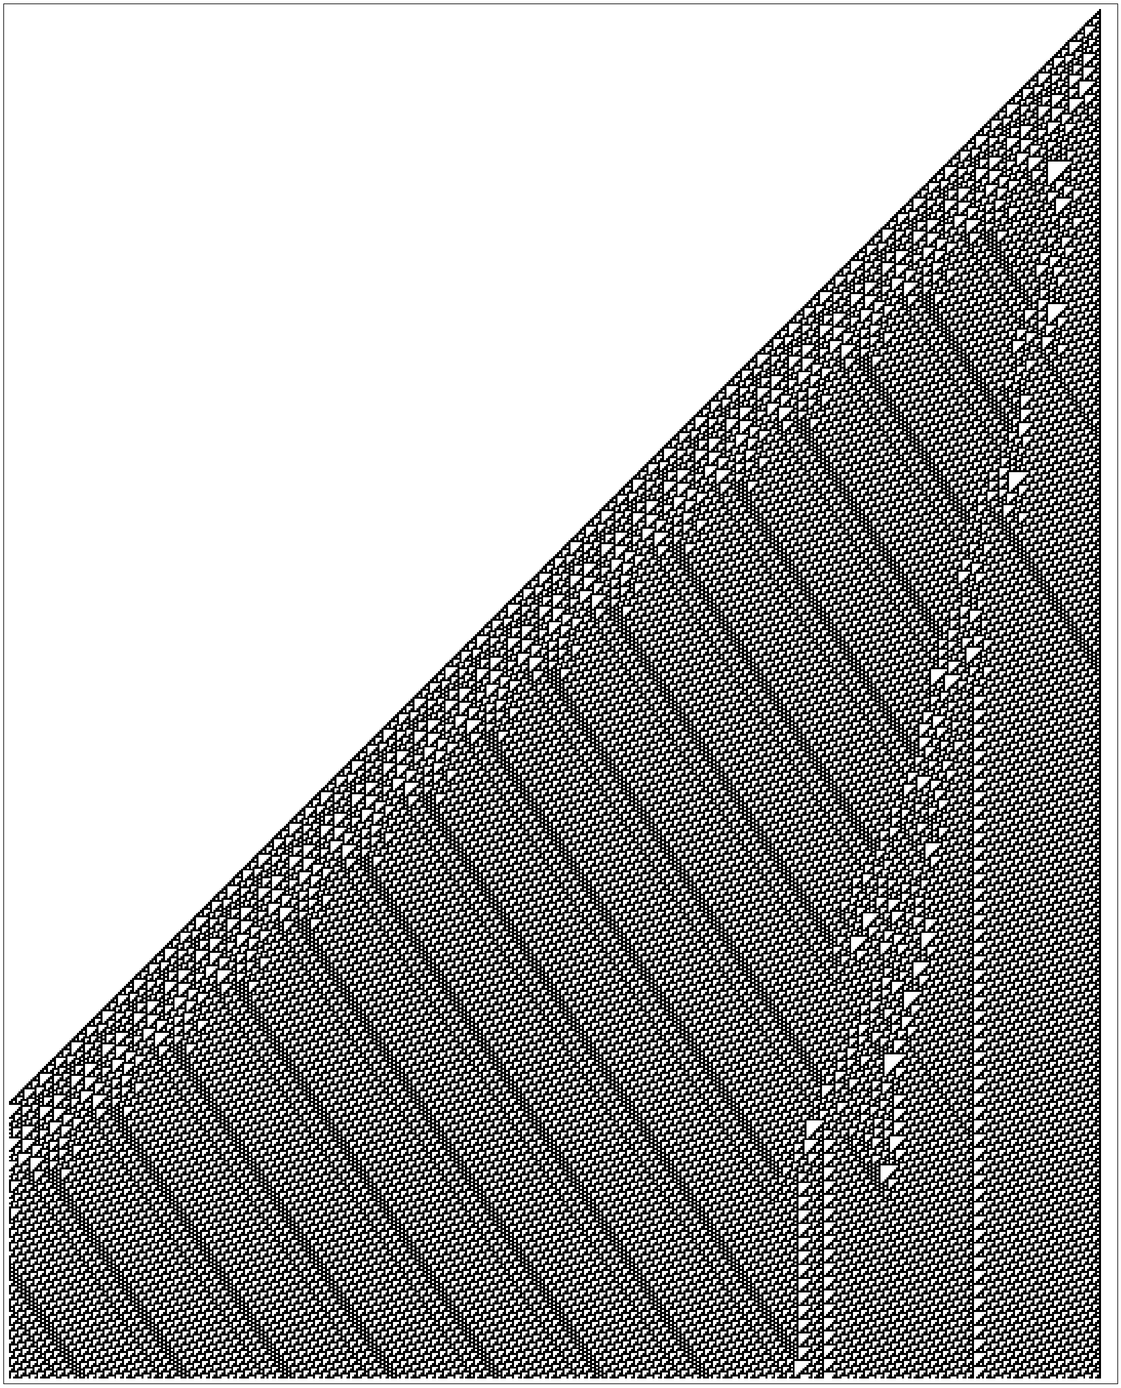

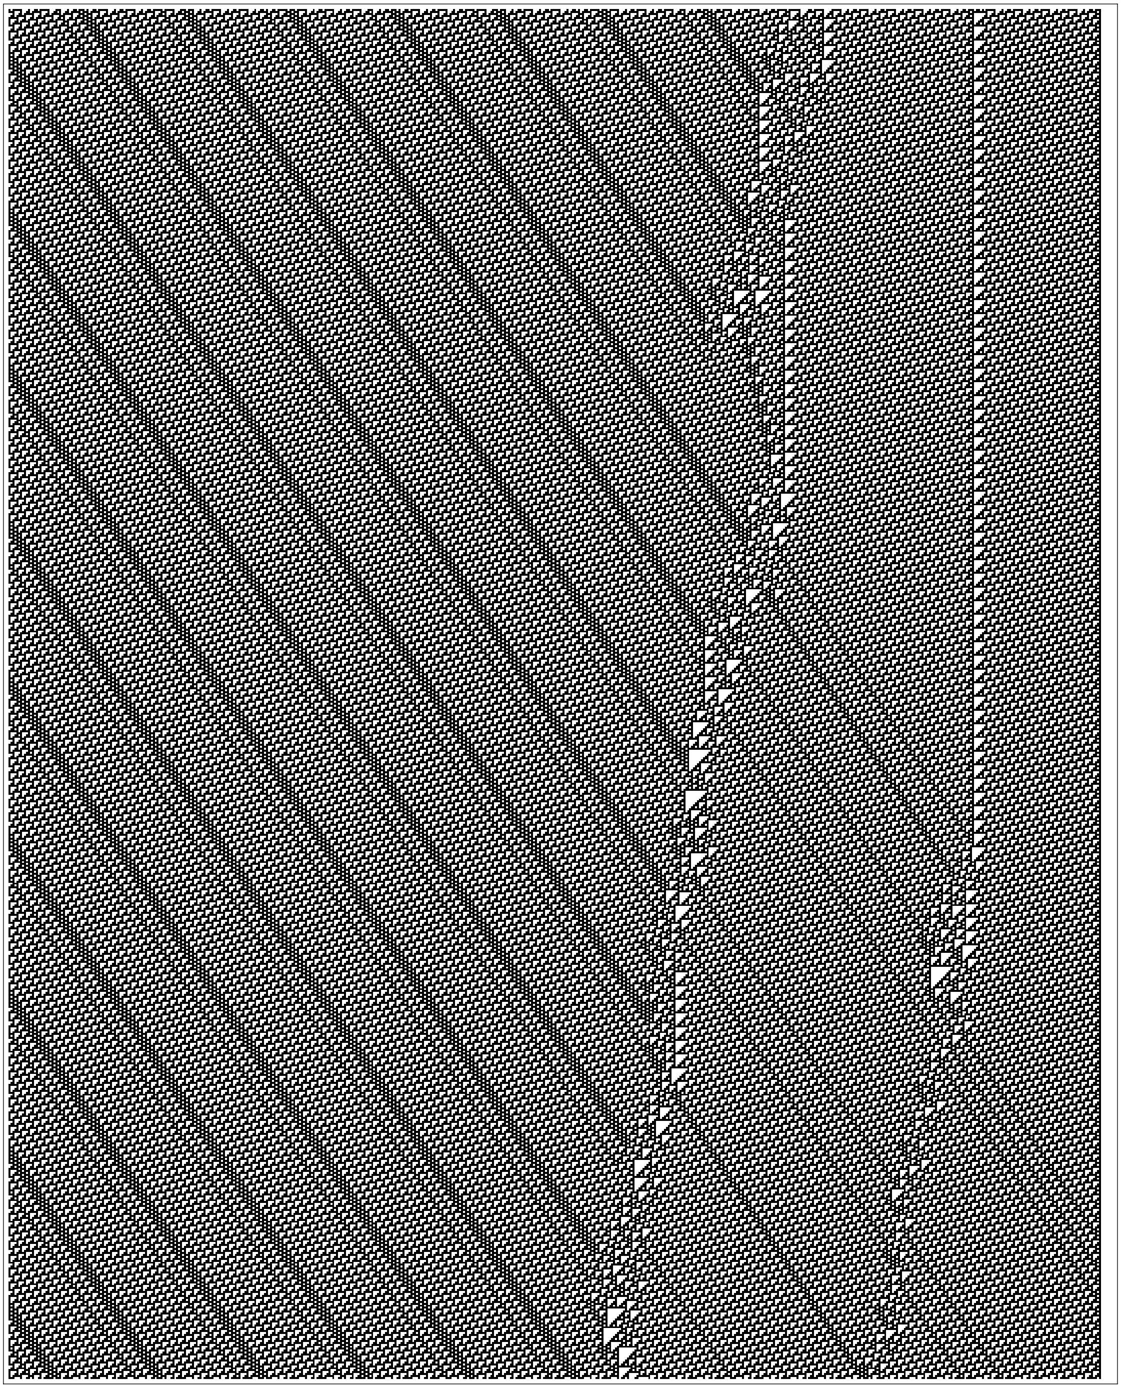

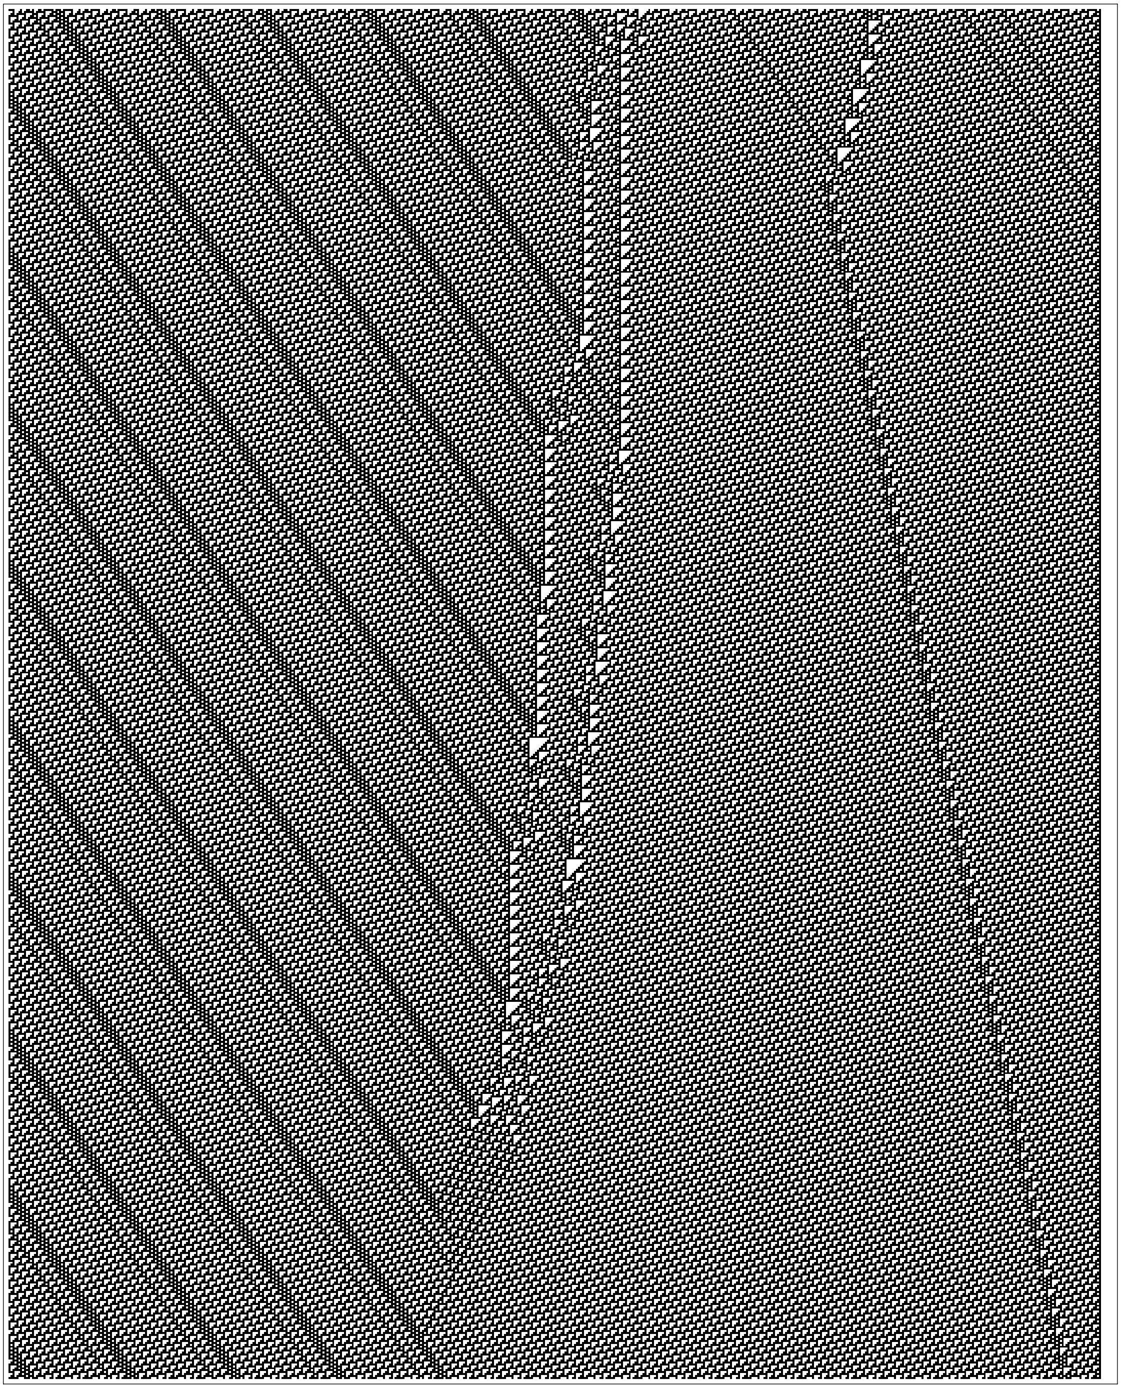

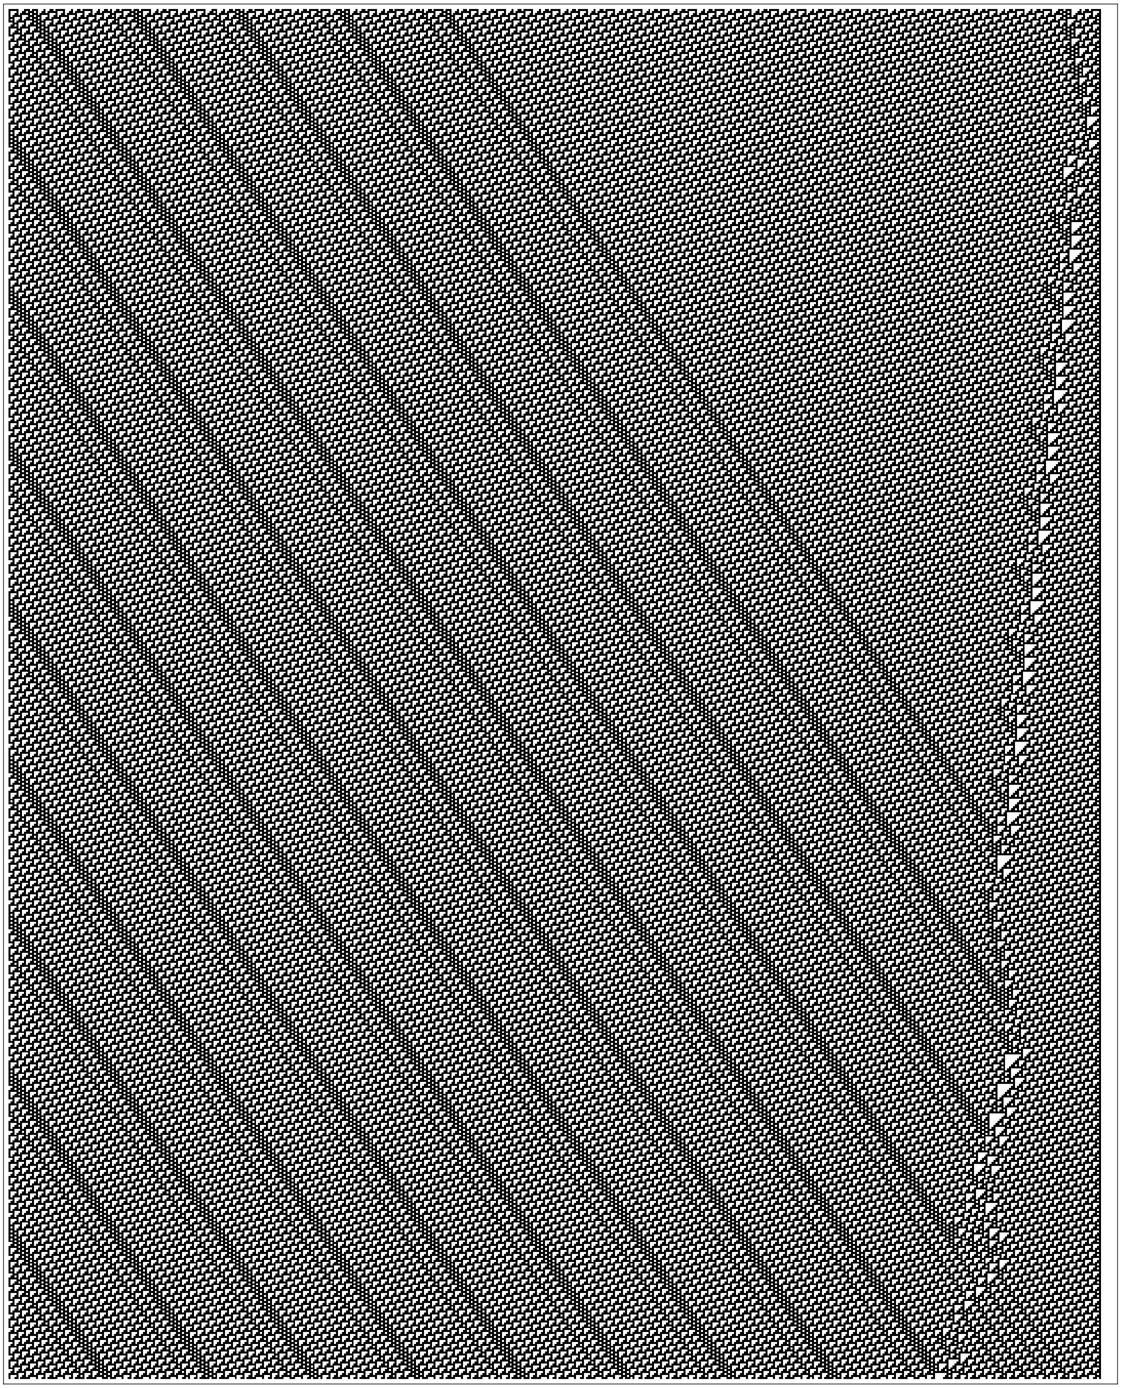

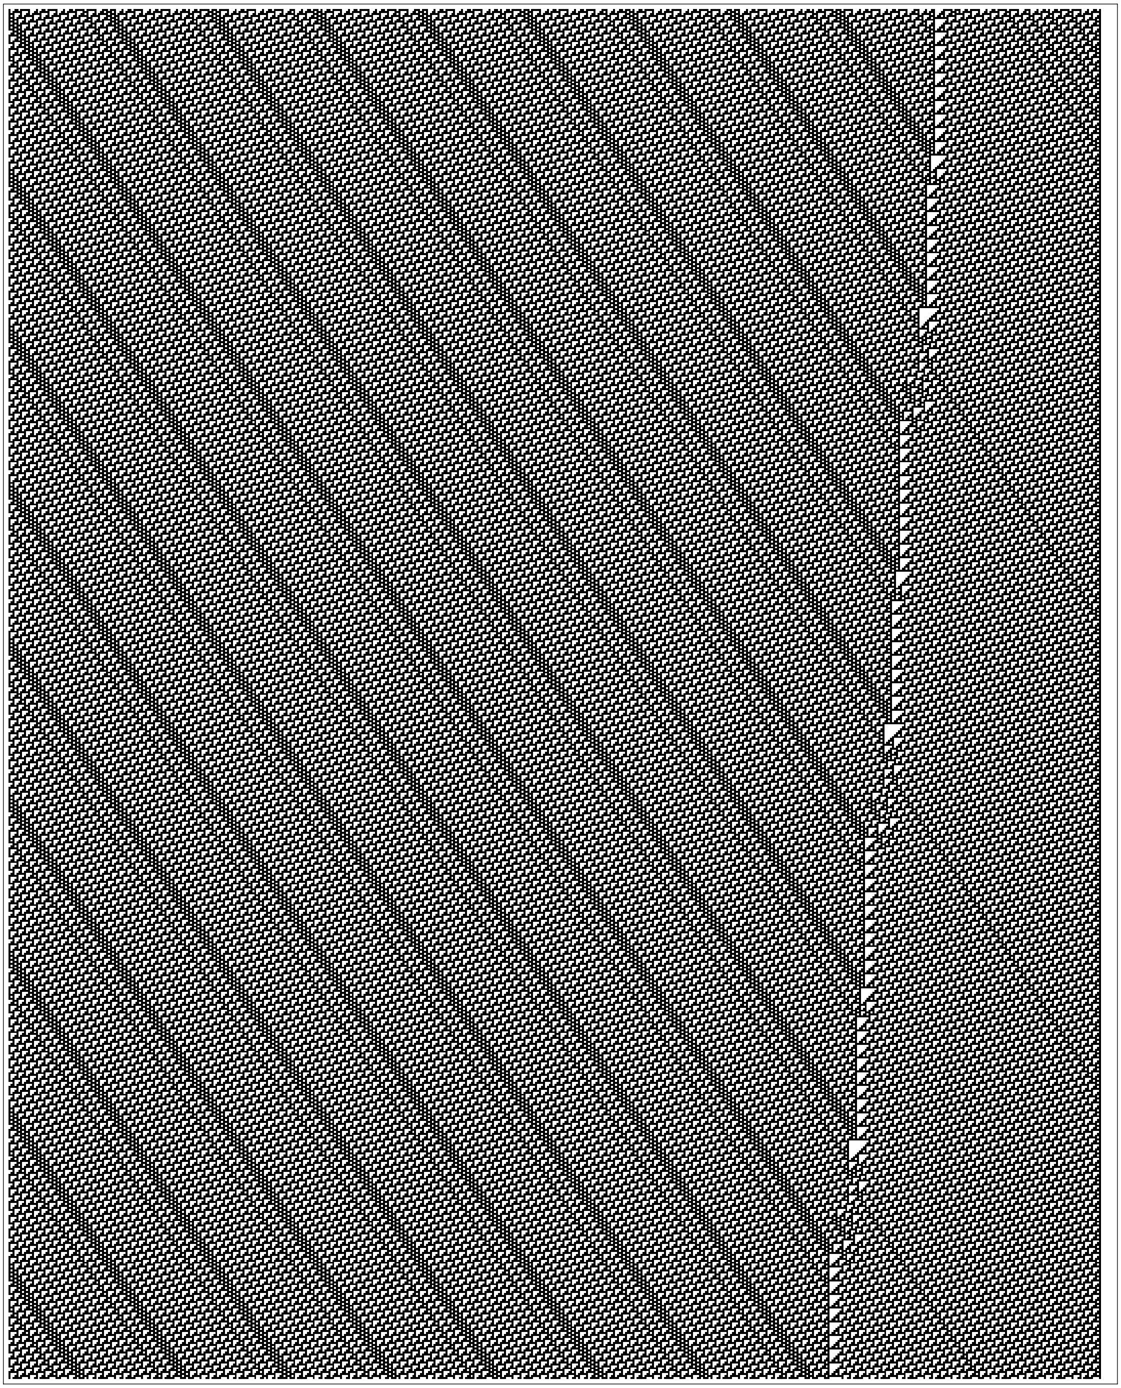

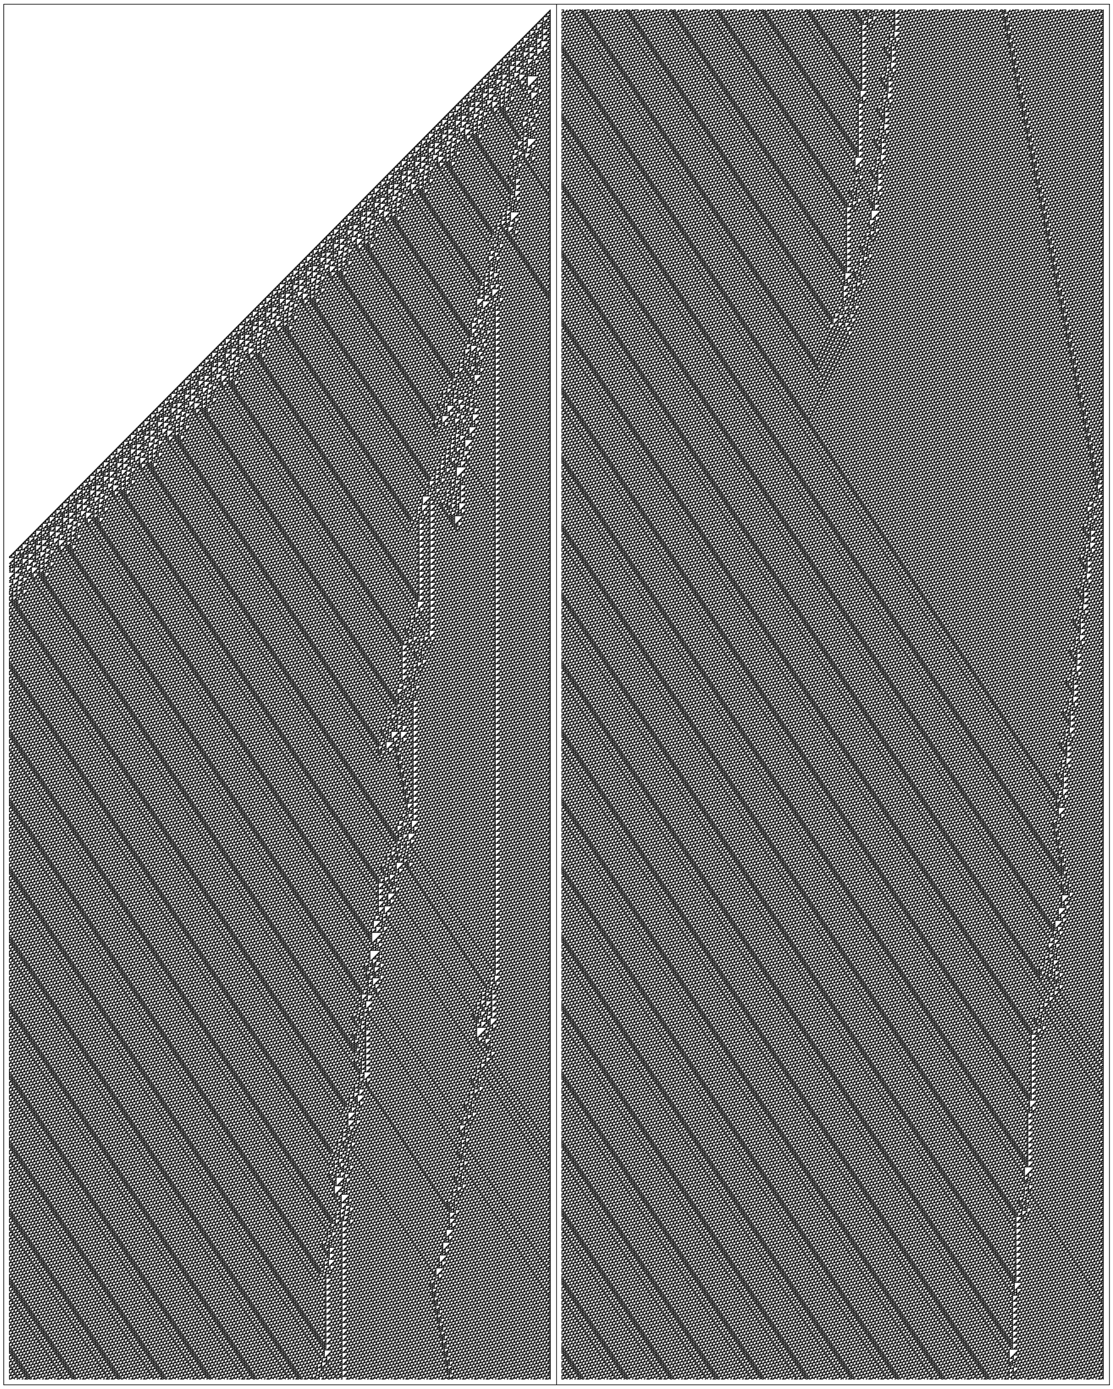

watch what happens. And as it turns out, in the particular case shown here, the outcome is finally clear after about 2780 steps: one structure survives, and that structure interacts with the periodic stripes coming from the left to produce behavior that repeats every 240 steps.

However certain one might be that simple programs could never do more than produce simple behavior, the pictures on the past few pages should forever disabuse one of that notion. And indeed, what is perhaps most bizarre about the pictures is just how little trace they ultimately show of the simplicity of the underlying cellular automaton rule that was used to produce them.

One might think, for example, that the fact that all the cells in a cellular automaton follow exactly the same rule would mean that in pictures like the last few pages all cells would somehow obviously be doing the same thing. But instead, they seem to be doing quite different things. Some of them, for example, are part of the regular background, while others are part of one or another localized structure. And what makes this possible is that even though individual cells follow the same rule, different configurations of cells with different sequences of colors can together produce all sorts of different kinds of behavior.

Looking just at the original cellular automaton rule one would have no realistic way to foresee all of this. But by doing the appropriate computer experiments one can easily find out what actually happens—and in effect begin the process of exploring a whole new world of remarkable phenomena associated with simple programs.

The Need for a New Intuition

The pictures in the previous section plainly show that it takes only very simple rules to produce highly complex behavior. Yet at first this may seem almost impossible to believe. For it goes against some of our most basic intuition about the way things normally work.

◀ A single picture of the behavior from the previous five pages. A total of 3200 steps are shown. Note that this is more than twice as many as in the picture on page 30.