EXCERPTED FROM

STEPHEN WOLFRAM

A NEW KIND OF SCIENCE

SECTION 8.8

Financial Systems

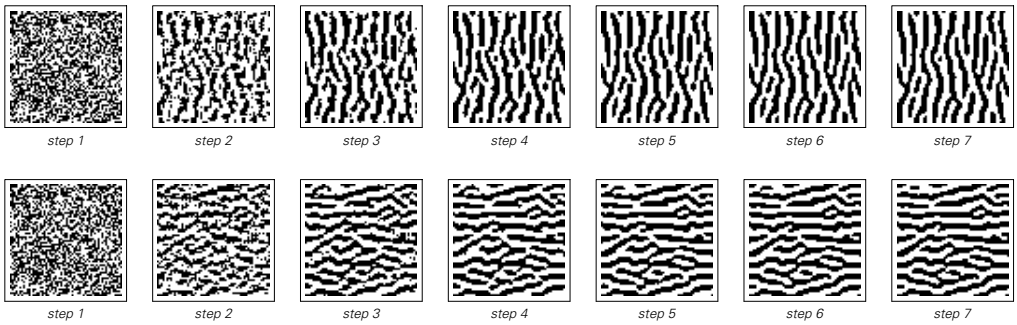

Examples of rules in which cells in the horizontal and vertical directions are weighted differently. In the first case, cells at distances 2 and 3 only have an effect in the vertical direction; in the second case, they only have an effect in the horizontal direction. The result is the formation of either vertical or horizontal stripes.

Financial Systems

During the development of the ideas in this book I have been asked many times whether they might apply to financial systems. There is no doubt that they do, and as one example I will briefly discuss here what is probably the most obvious feature of essentially all financial markets: the apparent randomness with which prices tend to fluctuate.

Whether one looks at stocks, bonds, commodities, currencies, derivatives or essentially any other kind of financial instrument, the sequences of prices that one sees at successive times show some overall trends, but also exhibit varying amounts of apparent randomness.

So what is the origin of this randomness?

In the most naive economic theory, price is a reflection of value, and the value of an asset is equal to the total of all future earnings—such as dividends—which will be obtained from it, discounted for the interest that will be lost from having to wait to get these earnings.

With this view, however, it seems hard to understand why there should be any significant fluctuations in prices at all. What is usually said is that prices are in fact determined not by true value, but rather by the best estimates of that value that can be obtained at any given time. And it is then assumed that these estimates are ultimately affected by all sorts of events that go on in the world, making random movements

in prices in a sense just reflections of random changes going on in the outside environment.

But while this may be a dominant effect on timescales of order weeks or months—and in some cases perhaps even hours or days—it is difficult to believe that it can account for the apparent randomness that is often seen on timescales as short as minutes or even seconds.

In addition, occasionally one can identify situations of seemingly pure speculation in which trading occurs without the possibility of any significant external input—and in such situations prices tend to show more, rather than less, seemingly random fluctuations.

And knowing this, one might then think that perhaps random fluctuations are just an inevitable feature of the way that prices adjust to their correct values. But in negotiations between two parties, it is common to see fairly smooth convergence to a final price. And certainly one can construct algorithms that operate between larger numbers of parties that would also lead to fairly smooth behavior.

So in actual markets there is presumably something else going on. And no doubt part of it is just that the sequence of trades whose prices are recorded are typically executed by a sequence of different entities—whether they be humans, organizations or programs—each of which has its own detailed ways of deciding on an appropriate price.

But just as in so many other systems that we have studied in this book, once there are sufficiently many separate elements in a system, it is reasonable to expect that the overall collective behavior that one sees will go beyond the details of individual elements.

It is sometimes claimed that it is somehow inevitable that markets must be random, since otherwise money could be made by predicting them. Yet many people believe that they make money in just this way every day. And beyond certain simple situations, it is difficult to see how feedback mechanisms could exist that would systematically remove predictable elements whenever they were used.

No doubt randomness helps in maintaining some degree of stability in markets—just as it helps in maintaining stability in many other kinds of systems that we have discussed in this book. Indeed, most markets are set up so that extreme instabilities associated with

certain kinds of loss of randomness are prevented—sometimes by explicit suspension of trading.

But why is there randomness in markets in the first place?

Practical experience suggests that particularly on short timescales much of the randomness that one sees is purely a consequence of internal dynamics in the market, and has little if anything to do with the nature or value of what is being traded.

So how can one understand what is going on? One needs a basic model for the operation and interaction of a large number of entities in a market. But traditional mathematics, with its emphasis on reducing everything to a small number of continuous numerical functions, has rather little to offer along these lines.

The idea of thinking in terms of programs seems, however, much more promising. Indeed, as a first approximation one can imagine that much as in a cellular automaton entities in a market could follow simple rules based on the behavior of other entities.

To be at all realistic one would have to set up an elaborate network to represent the flow of information between different entities. And one would have to assign fairly complicated rules to each entity—certainly as complicated as the rules in a typical programmed trading system. But from what we have learned in this book it seems likely that this kind of complexity in the underlying structure of the system will not have a crucial effect on its overall behavior.

And so as a minimal idealization one can for example try viewing a market as being like a simple one-dimensional cellular automaton. Each cell then corresponds to a single trading entity, and the color of the cell at a particular step specifies whether that entity chooses to buy or sell at that step. One can imagine all sorts of schemes by which such colors could be updated. But as a very simple idealization of the way that information flows in a market, one can, for example, take each color to be given by a fixed rule that is based on each entity looking at the actions of its neighbors on the previous step.

With traditional intuition one would assume that such a simple model must have extremely simple behavior, and certainly nothing like what is seen in a real market. But as we have discovered in this book,

simple models do not necessarily have simple behavior. And indeed the picture below shows an example of the behavior that can occur.

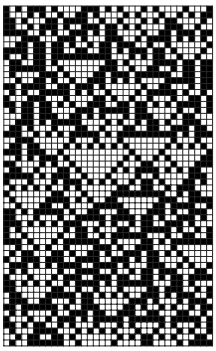

An example of a very simple idealized model of a market. Each cell corresponds to an entity that either buys or sells on each step. The behavior of a given cell is determined by looking at the behavior of its two neighbors on the step before according to the rule shown. The plot below gives as a rough analog of a market price the running difference of the total numbers of black and white cells at successive steps. And although there are patches of predictability that can be seen in the complete behavior of the system the plot on the right looks in many respects random.

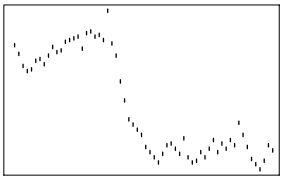

In real markets, it is usually impossible to see in detail what each entity is doing. Indeed, often all that one knows is the sequence of prices at which trades are executed. And in a simple cellular automaton the rough analog of this is the running difference of the total numbers of black and white cells obtained on successive steps.

And as soon as the underlying rule for the cellular automaton is such that information will eventually propagate from one entity to all others—in effect a minimal version of an efficient market hypothesis—it is essentially inevitable that running totals of numbers of cells will exhibit significant randomness.

One can always make the underlying system more complicated—say by having a network of cells, or by allowing different cells to have different and perhaps changing rules. But although this will make it more difficult to recognize definite rules even if one looks at the complete behavior of every element in the system, it does not affect the basic point that there is randomness that can intrinsically be generated by the evolution of the system.