

chapter 1

Scientific research can often seem obscure and even pointless to outsiders. This is not so much due to the intellectual difficulty involved in understanding such activity, for it is widely accepted that this will inevitably be the case. It is rather that many of the topics which are examined seem to be almost designed to incur the scorn and wrath of the layperson. Before sitting down this morning to write these very words, for example, my eye fell on a report in a serious British newspaper. An American psychologist had been visiting the country to carry out a study of rams in the English Lake District. His research was complete. ‘Ten per cent of all rams’, he proclaimed solemnly, ‘are homosexual’. At least, readers no doubt took consolation in the fact that this finding was obtained at the expense of the American taxpayer and not themselves.

Nor are such examples confined exclusively to the sciences. I have long admired Emily Brontë’s novel *Wuthering Heights*. The opening chapters, for example, in which Lockwood first encounters the ill-tempered Heathcliff and his assorted household, seem to me to be one of the finest pieces of comedy in the whole corpus of English literature. Realising that not everyone shares this opinion, and in order to improve my understanding, I recently opened a modern work of literary criticism on Brontë’s masterpiece. It was completely impenetrable. Many of the individual words were quite new to me, and whole sentences, indeed whole pages, appeared to lack any coherent meaning. I sought solace in the preface, where I learned that the density of the text was deliberate. ‘The analysis of literature and culture’, declared the author, ‘is a task no less difficult, and no less demanding of a specialized language, than the study of sub-atomic particles’. I hastened immediately to a textbook on orthodox economic theory in an effort to restore my sanity.

In the mid-1980s, entomologists carried out a series of experiments with ants which, at first sight, appear equally esoteric. Two identical food sources were placed equidistantly from a nest of ants, and were kept constantly replenished so that they always remained identical. In other words, every time an ant removed a grain from one of the sources, another was added to the pile. And the two piles were exactly the same distance from the nest. How would the ant colony divide itself between the two sources of food?

The experiments appear at first sight to be of little or no interest to anyone outside the world of biology. Even amongst biologists, ant behaviour seems a pretty specialised topic. Yet the results of the experiments turned out to be fiendishly difficult to explain. And a proper understanding of them has very widespread implications for behaviour far beyond that of the humble colony of ants, illuminating complex problems in human economies and societies, worlds living at the edge of chaos.

Thinking of the ant experiments themselves, there was, by design, absolutely no reason for the ants preferring one of the food sources to the other. So we might start by expecting that the ants would split evenly between them. A little reflection would lead us to think that, whilst this might very well be an outcome, any division would be possible. Suppose each ant emerges from the nest and visits one of the food piles at random. It is successful in obtaining food to bring back to the nest, and so on its next outing it has an incentive to re-visit the site of its previous success. The pile is always replenished, so it will always obtain food from this site.

If this theory were correct, the distribution of the ants between the two piles could be analysed in just the same way as an experiment in tossing a fair coin and observing the split between heads and tails. The first time an ant comes out of the nest to look for food, its destination is given by the equivalent of a toss of a coin, and the design of the experiment gives it a strong incentive to keep re-visiting its original choice. So, in theory, we could expect the colony to split in any proportion between the two piles. There would be a strong expectation that the split would be close to 50:50, because this is how a large number of tosses of a fair coin usually divide, but any distribution would be possible theoretically.

But the biologists had developed a more sophisticated version of this theory, based upon a known fact about ant behaviour. Once an ant has successfully found food - which it would by the construction of the experiment - it will re-visit the same site the next time and so on into the future. But when an ant which has found food returns to the nest, it physically stimulates another ant to follow it to the food source by chemical secretion. Some kinds go even further and recruit whole groups to follow them, by

laying a trail of secretions. So an ant emerging from the nest for the first time could be influenced in its decision by the trails of the ants it encounters in its journey.

Thinking in economic terms, the behaviour of agents - the ants in this example - is influenced directly by the behaviour of others. In this example, the interaction between ants takes place at what we term the *local* level. No ant can ever observe the overall division of the colony between the two food sources, and so this cannot influence the choice of destination. But each ant is open to recruitment by the limited number of other ants which pass by its immediate neighbourhood.

The situation is one in which, to introduce a technical term, positive feedback predominates. An ant goes out, finds food and encourages others to follow it back to its source. In this artificial experiment, the self-reinforcing mechanism is very strong, for each pile of food is constantly replenished. So the ants which are recruited find food with complete certainty, and return to recruit others. The more ants visit any particular site, the greater the chance that yet more of them will visit it in future.

In other words, the consequences of actions by individual ants are enhanced by their influence on the behaviour of others, hence the phrase 'positive feedback'. The term is purely descriptive, and does not carry any overtones of approval or desirability. It applies to any system, such as that of our ant colony, in which the initial impact of actions or events tends to be magnified over time. Its opposite is, unsurprisingly, 'negative feedback', used to describe systems in which initial effects are dampened and smoothed away. As we shall see later in the book, almost the whole of conventional economic theory can be thought of as describing systems of negative feedback. But in the real world of the economy and society, positive feedback generally rules.

The crucial trail laying quality of ants led to the more subtle theoretical expectations of the proportions which visit each of the sites. The signals left by the creatures mean that the random choices of the first few ants to leave the nest could exercise a decisive influence on the behaviour of the whole colony. If the choice of each ant were purely random each time it left the nest, because of the very large number of ants, there is a probability very close to one - in other words, almost complete certainty -

that the proportions will settle down very close to a 50:50 split.

But suppose half a dozen ants went out, foraged and returned with food. These then left trails for the next group to follow, and so on. But the random choices of a very small number of ants may not divide evenly between the two sites. Our fair coin tossed enough times *will* lead to an even split, but it is much less likely that a small number of tosses will give an equal number of heads and tails¹. The trails which the first returning ants leave potentially influence the decisions of those emerging for the first time and, precisely because the random choice of a *small* number can influence the subsequent decisions of the whole group, the eventual proportions visiting the two sites may differ quite markedly from a 50:50 split.

But a key feature of the biologists' theory was that the proportions in any given experiment would settle down to whatever they had been determined by in the early stages of the food foraging process. There would be some random fluctuations around this for a short time, but the eventual outcome would be stable.

This theoretical framework is an important one. It predicts that, once a few more ants start to visit one of the sites rather than the other, for whatever reason, there will be a strong tendency for that site to become more and more the favoured destination. Some of the early recruits to the other site might stay loyal, as it were, but we expect an unbalanced outcome to arise. And once this has arisen, the proportions will then remain fixed. Or, in the jargon, the system will stay locked in that particular solution.

In fact, what was observed to take place was a completely different outcome. Even when the experiment had been running for a long period of time - in ant terms - the proportion of the total ant population visiting any one site continued to fluctuate in an apparently random fashion. The average of this proportion was one-half, but this precise outcome was hardly ever observed. Indeed, the proportion was subject to constant change. Once a large majority of ants had piled up at one of the sites, the outcome tended to stay reasonably stable and exhibited small variations around that proportion for some considerable time. But inevitably the majority was always eroded and the ants

¹ With six tosses, for example, the odds are *against* an even split

switched to visiting the other site. Sometimes these shifts were not only very large, from, say an 80:20 division at one pile to the reverse outcome of 20:80, but also rapid.

The constant changes, often small but occasionally rapid and large, were entirely unexpected according to the biologists' theory. This conflict between the actual and theoretical outcomes led the experiment to be repeated in different ways. The exact recruitment mechanism which is used varies between species of ants, so different species were tried. The outcome was the same. Doubts then arose as to whether there was some subtle change in the food source which was the cause of the fluctuations, such as the piles not being replenished in an absolutely symmetrical way. So the experiment was tried with just one food source and two identical bridges, precisely the same distances away from the nest, and the proportion going over each of the bridges was observed. Again, the same pattern of behaviour was observed.

The economist Alan Kirman, then based at the European University Institute in Florence, turned his mind to the problem. By definition in circumstances such as the ant experiment, the idea that the system as a whole can be understood by the behaviour of a single, representative agent is a complete non-starter. For the overall outcome arises as a result of the interactions between individuals, and the changes in behaviour which they induce in one another. It is, quite literally, impossible to infer the behaviour of the group as a whole from an account of one of its individuals taken in isolation. Kirman has in fact been one of the world leaders in pioneering the development of interacting agent models in economics. But, to paraphrase the words of a popular song, what's ants got to do with it?

Kirman set up a theoretical model which gives an excellent account of the observed behaviour of the seemingly perverse ants. And it can also be stated quite simply. An ant coming out of the nest follows one of three possibilities. It either visits the food pile it previously visited; it is persuaded by a returning ant to visit the other source; or, of its own volition, it decides to try the other pile itself. And this is almost all that is required to explain the complex and seemingly baffling phenomenon of the fluctuations in the proportions of ants visiting the respective piles.

I use these simple basic principles throughout the book to explain many economic and social problems. At any point in time an individual agent - whether an ant, a person, a company or whatever - can follow one of three choices: to stay with its previous decision; to select an alternative of its own accord; or to be persuaded to switch to the alternative by the actions of others.

In such circumstances, no single outcome of an experiment will ever be identical to another, for the choices of individual ants are not fixed, but can be altered each time with given probabilities. It is this random element to the whole process which means that each solution of Kirman's theoretical model, and each outcome of each practical experiment, is unique. But a typical simulation, or outcome, of Kirman's model is plotted in Figure 1.1, which shows the proportion of ants visiting one of the food sources at any one time. The chart illustrates the typical patterns of constant small changes and occasional large shifts which are observed.

Typical solution of ants model

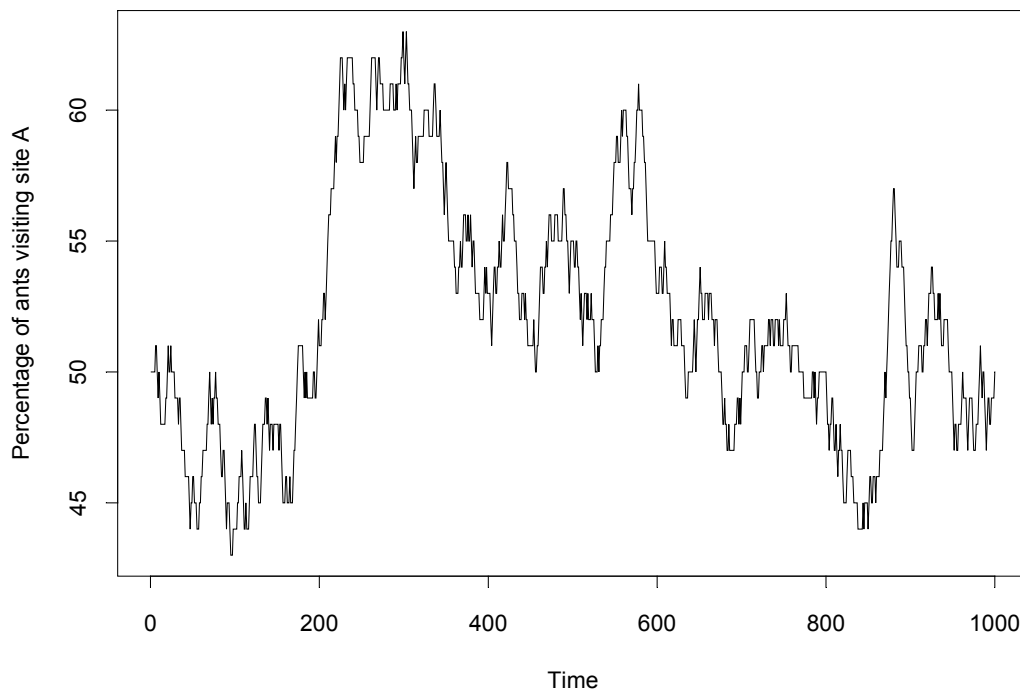


Figure 1.1

0

When its properties are examined more deeply, such simulated data exhibits characteristics which are entirely typical of situations in which the behaviour of any individual agent is influenced directly by the behaviour of others. In the short-term, movements in the series are quite unpredictable. Even with completely accurate

knowledge of the equations which describe the behaviour of the individual ants, it is not possible to predict with any degree of accuracy the direction of change of the proportion of ants which visit either of the food sources.

Indeed, in this particular system, non-predictability appears in its most extreme form. We can work out the probability of the very next ant about to collect food visiting a particular site, but we can never do any better than this. In other words, all we can ever say is that the next ant has a certain probability of visiting one site, and a certain probability of visiting the other. In the same way, with the toss of a fair coin, we can never do better than say that there is a probability that a head will appear, and one that a tail will appear. Any 'prediction' can be no better than a pure guess.

One way of looking at this is to see if we can draw any conclusions about the way the system will move from any given split of the colony between the sites. Look, for example, at what happens when the split is 55:45. Reading across from the point marked '55' on the left hand axis, we can see a number of occasions on which this split occurred in this particular simulation of the model. The first time, the proportion of ants visiting site A then rose rapidly to over 60 per cent. The next time the 55:45 split happened, the proportion visiting this site subsequently fell by a small amount. Moving across to the peak at the far right of the chart, the proportion visiting site A rose by a small amount for a short time. But then, as it fell back through 55 it continued to fall quite sharply. In other words, the proportions we observe at any point in time give us no information about what will happen to the proportions in the immediate future.

But the system does have a very distinct pattern in the longer term. Figure 1.2 sets out for the ants model how much time the system will spend at any given distribution of the ant colony between the food sources, whenever the experiment is run for a reasonable length of time. The precise shape of this distribution will vary according to the persuasiveness with which ants can convert others, and on the propensity of individuals to change their own minds.

Figure 1.2 shows the relative amounts of time which the ant population spends at different proportions between the two sites, when the propensities to switch behaviour

are low. The values along the bottom axis of the chart move from zero to 100. This shows the percentage visiting site A, so when the value is close to zero, by implication almost 100 per cent of the ants are visiting site B, and vice versa. The left hand axis of the chart shows the relative proportion of time when the various proportions of the ants visiting site A are observed. The U-shape of the curve tells us that the ants spend much more time at extremes of the split between the two sites than they do at reasonably equal distributions. In other words, the colony spends a large proportion of its time in situations in which almost every ant visits site A and very few site B, and an equally large proportion when almost every ant visits site B and very few site A. In contrast, the occasions on which a split close to 50/50 is observed are relatively few and far between.

Relative amounts of time for different percentages of ants at site A
low propensity to switch behaviour

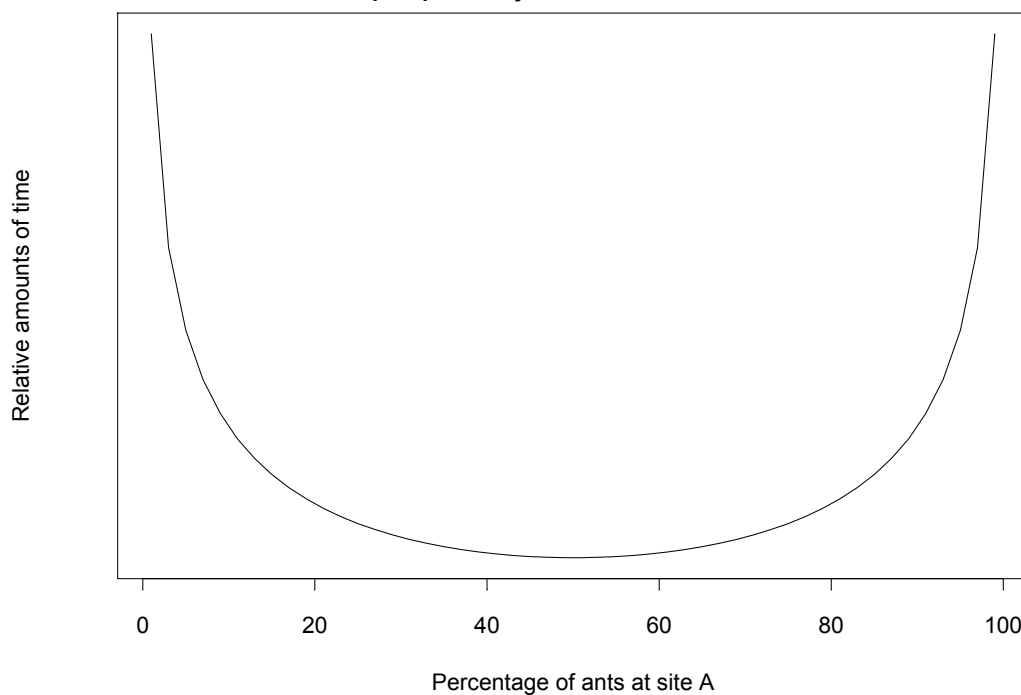


Figure 1.2

Figure 1.3 sets out the same kind of plot as Figure 1.2, but one feature has changed. In this case the propensity of the ants to switch behaviour is high.

Relative amounts of time for different percentages of ants at site A
high propensity to switch behaviour

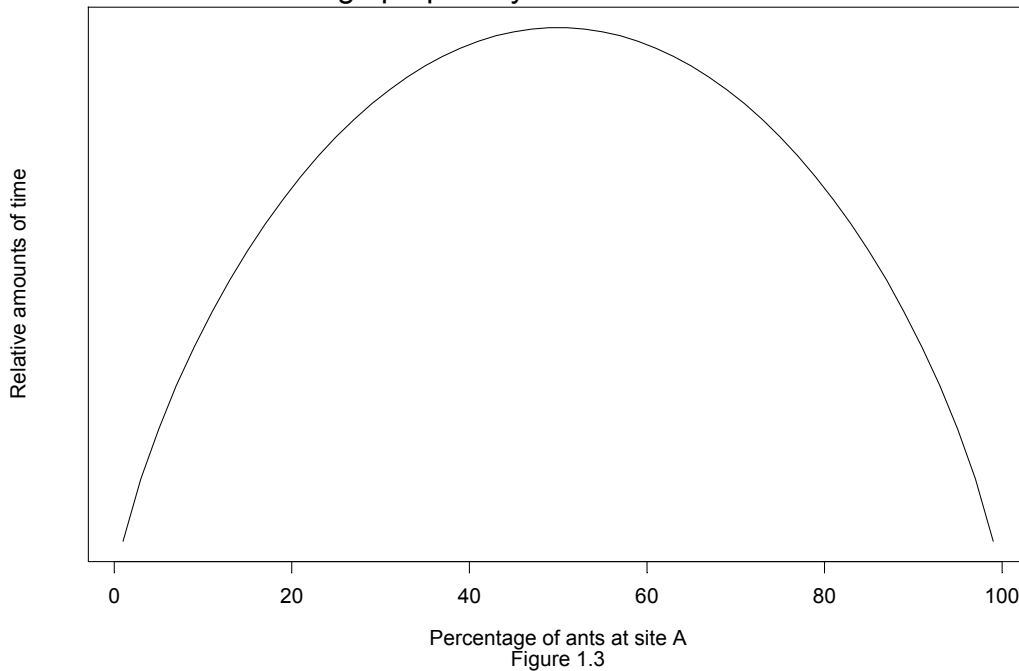


Figure 1.3

Comparing Figures 1.2 and 1.3, a potential paradox appears to arise. In the first figure, ants have only a low propensity to change their behaviour and visit a different site, and in the second they are much more likely to switch. Yet in the former case, a much larger percentage of the total time is spent with most of the ants visiting either site A or B than is the case in Figure 1.3. In this latter instance, the ants spend much more time split closer to 50/50 between the sites.

A first impression might suggest that a high likelihood of changing behaviour would drive the system to the extremes, rather than a low one. But, on the contrary, if ants often change the site they visit, the chances of most of them ending up at one or other of the sites is very low, for the very reason that lots of them change their mind each period. In contrast, if changes are only occasional, once the proportion has wandered off, as it were, to an extreme split, it will take a very long time to change. It may take a long time to ever get into such a situation, but once there, the proportion will take even longer to be altered significantly.

The behaviour of individual ants, their direct influence on the behaviour of

others, and the consequences of this interaction between individuals for the colony as a whole offer a very general description, or model, of a wide range of economic and social phenomena. For the principles which govern the behaviour of ants also apply to humans. Much of the time, individuals face a limited number of choices in any particular situation.

Often, there are more than two choices, but this is just an extension of the fundamental ideas which can be readily incorporated. There are other extensions, complications and simplifications which we will come across in the course of the book as we consider different circumstances and different problems.

But the essential principles of the ants model remain. In most circumstances, a person can either stay with the pattern of behaviour he or she previously followed (an ant visiting its previous site), can decide to switch of his or her own volition, or can be influenced into switching by the observed behaviour of others.

And, as we shall see, the consequences of this description of individual behaviour has deep implications for the outcome for the human colony as a whole. Many important social and economic issues share the key characteristics of ant behaviour, of unpredictability in the short-run merging imperceptibly over time into a form of regularity, of complex systems living at the edge of chaos.