CHAPTER 5

Mechanistic Modelling

Seventeenth-century Europe was a place of great religious ferment. Theological differences between the various Christian churches and sects were often magnified into issues involving life and death. A belief common to many of the protagonists was that their opponents either were, or had the potential to become, the Antichrist. Some very considerable intellects, along with many others of lesser calibre, applied themselves to the problem of identifying either whether Antichrist had actually appeared already, or was about to appear.

A technique which was used widely in this investigation was the so-called 'science' of numerology. Scholars would look carefully for coincidences in the appearance of sequences of numbers, searching for clues to identify Antichrist's appearance on Earth. Many predictions were made that he had in fact appeared. But since his appearance is prophesied to bring about a reign of dreadful terror and chaos, more thoughtful minds began to question whether this was in fact the case. Indeed, they began to wonder whether the discipline of numerology had any value whatsoever.

In 1690, John Owen, Chancellor of Oxford University, went so far as to state: 'Take heed of computation! How woefully and wretchedly we have been misled by it!'

This could serve as the epitaph for the discipline of economics in its current state. And an epitaph in particular for macro-economic modelling and forecasting, the theme of this chapter.

So far, most of our attention has been devoted to the micro-level, to the model of competitive equilibrium, because it is the core model of economic theory. Whatever nuances and qualifications they might make, orthodox economists ultimately subscribe to this model as a reasonable approximation of how the world operates. Its workings rarely, if ever, catch the attention of the media. And with good reason. For within standard economics the model is about as controversial as the fact that the Pope believes in God.

It is much more exciting, not to say amusing, for the media to feature disputes among economists as to how the economy operates at the

aggregate, macro-level, and to ridicule the accuracy of macro-economic forecasts.

Almost every day, economists, particularly those employed by the publicity-conscious financial institutions, can be heard and seen pontificating and arguing with each other about the significance of the latest economic statistic to emerge, or pronouncing with great confidence on what the government should be doing about this or that problem. The remark of the Russian physicist Lev Landau, made in a different context, might almost have been designed as a perfect description of the gurus of Wall Street and the City of London: 'Cosmologists,' he wrote, 'are often in error but never in doubt.'

Whatever solution is offered, whatever prognostication is made, one quality above all is lacking from the various utterances. Namely, that of humility. For the record of economists in understanding and forecasting the economy at the macro-level is not especially impressive. Indeed, uncharitable writers might be inclined to describe it as appalling.

To be fair, economics is not a completely empty box in the way in which numerology was so clearly shown to be. Its status is similar to that of science before Newton. At the micro-level in particular, in certain well-defined circumstances, the discipline can offer useful insights into behaviour. But the contrast between the actual scientific achievements of the discipline and the confidence with which claims are made for it by its protagonists is striking.

A long-standing proposition in economics, for example, is the quantity theory of money. This asserts that inflation is always and everywhere caused by increases in the supply of money. At one level, the proposition is simply an accounting identity. It states that the amount of goods and services produced in an economy in any one period multiplied by the average price of the same goods and services by definition equals the amount of money in circulation multiplied by the 'velocity' of circulation of money.

It will be helpful to the subsequent discussion to write down this somewhat wordy proposition in the form of a simple equation. Using the symbol 'M' to denote the money supply, 'V' the velocity of circulation, 'P' the average price level and 'Q' the amount of goods and services produced in an economy, the quantity theory of money states that: $M \times V = P \times Q$. In other words, the amount of money in an economy, multiplied by the speed with which it circulates, equals the amount of goods and services produced multiplied by the price level.

One way in which this equation can be used is to measure V. The central bank of an economy knows how much money is in circulation, and governments make separate estimates of P and Q. So, given values for M, P and Q, V is measured as being that value which, when multiplied by M, is equal to P times Q. In fact, this is the only way in which V can be measured, for there is no independent way of doing so, unlike M, P and Q, for which well-defined and well-used procedures exist to estimate their values.

But it is not terribly interesting to use this equation simply to measure V, a rather esoteric factor which measures how fast money is circulating round any particular economy. And, indeed, the quantity theory of money does make much more imaginative use of it.

Suppose that the government or the central bank of an economy can control the amount of money in circulation, M. Suppose further that, ultimately, changes in the amount of money do not affect the amount of goods and services produced in an economy, Q. Temporarily, perhaps for as long a period as two or three years, increasing or decreasing the amount of money in an economy might lead to changes in how much is produced. But suppose that, eventually, changes in M have no effect on Q. Suppose, finally, that V does not change over time, but always keeps the same value. Rather more subtly, V can be allowed to change, but suppose that in some way these changes can always be predicted accurately.

If all these suppositions are correct, the quantity theory of money is valid, and there is a direct and predictable link between increases in the money supply and inflation. If M is increased by the government, and either V is constant or we know how it is going to change, we can compute how much the left-hand side of our equation, M times V, will change. And since eventually, on the right-hand side, Q will not change as a result of changes to M or V, P must rise by an amount equal to the change in M times V. The price level, P, rises, or in other words increasing the money supply has caused inflation.

The reader might reflect that an awful lot of supposing has to take place in order for the quantity theory of money to be true. But this has not stopped it from being of great practical importance, particularly over the past twenty years or so. The proselytising work of Milton Friedman, Nobel Prize winner in economics and a firm believer in the quantity theory, did much to push the theory forward into the policy-making domain. The conduct of economic policy in many countries has been

strongly influenced by the view that inflation is a purely monetary phenomenon. If the money supply is controlled, so is inflation.

For example, the Bundesbank monitors changes in the German money supply very closely, believing that if it can be kept under control, so can inflation. And given the dominant position of the Deutschmark in the European ERM in the 1980s and early 1990s, policies followed by the Bundesbank have had a powerful influence on policies in the other countries of the European Community. Latin American countries are repeatedly urged by the IMF to exercise monetary 'discipline' and control their money supplies in order to control the endemic inflation of that continent.

Staying with the theme of discipline, Mrs Thatcher, in her early years in power, was a great believer that inflation could only be caused by changes in the money supply. The sharp rise in inflation which took place in Britain in 1980 came as a complete surprise to her newly elected government, which had set strict new targets for the growth of the money supply. The fact that the government had also conceded enormous wage increases across the public sector and at the same time virtually doubled the rate of VAT – the tax rate on sales of goods and services – early in the same year could not, on the logic of the government and its advisers, affect inflation, for these policies did not involve an increase in the money supply. Unfortunately for the government and the theory, workers in the private sector demanded wage increases of a similar size, pushing up industry's costs, and companies simply raised their prices in line both with these higher costs and with the new spending-tax rate. Inflation soared towards 20 per cent.

There are circumstances in which increases in the money supply do trigger increases in inflation. In the middle of the third century AD, for example, the Roman Empire experienced a tremendous crisis which almost destroyed it. In the space of fifty years, there were no fewer than twenty emperors: eighteen of them died violent deaths, one was held captive abroad, and the other was a victim of the plague. There was constant debasement of the currency, and the price of corn rose to well over a hundred times its previous levels. But such inflation was modest in comparison to the dramatic hyper-inflations of our own century, as, for example, in Germany in the 1920s; in Central European countries, most notably Hungary, at the end of the Second World War; and at present in the newly formed countries of the former Soviet bloc. Inflation in some of these examples far exceeded an annual rate of 1 million per cent. In

the single month of July 1946, prices in Hungary rose at an annualised percentage rate of more than 1 followed by sixteen zeros, and in August 1993 a similar calculation for Serbia produced a percentage rate of more than 3 followed by seventeen zeros.

But the claim that inflation is always and everywhere purely caused by increases in the money supply, and that the rate of inflation bears a stable, predictable relationship to increases in the money supply is ridiculous.

At a purely practical level, there is no unique definition of what constitutes the money supply. A range of indicators is used and appears in the financial press. The various definitions are, by convention, denoted by the capital letter 'M' followed by a number. So we have, for example, Mo, M1 and so on, usually up to M5, although there can be subtle variants such as M1A. The rule is, the bigger the number, the more factors are included in the definition of money.

At a narrow level, money might be defined simply as the total value of notes and coins circulating in the economy. But one could argue that money which is held in bank accounts and which can be withdrawn on demand is virtually the same thing as cash, so this, too, should be included in the definition. A case can be made for including the money held in other bank accounts. For example, a deposit account which requires one month's notice before it can be withdrawn is not as accessible as a pile of banknotes in the back pocket, but it is money which the owner can obtain.

The differences would not matter if all definitions of money tended to move together, but throughout the Western world at any point in time there are often large differences between the rates of growth of the various definitions of M. Simply by looking at the statistics published by the IMF for the growth in M1 and M2 – two measures which are close together in terms of what is and what is not included in their definitions, and so might be expected to move together quite closely – it is easy to pick out examples. Between 1985 and 1987, money as measured by M1 in the United States grew by 15 per cent, but by 25 per cent on its neighbouring definition given by M2. Between 1987 and 1990, M1 grew by 14 per cent, and M2 by only 8 per cent. In Britain, in the two years from 1988, M2 money grew by 23 per cent, but M1 grew by no less than 37 per cent. The list of examples could go on for a long time.

Virulent debates rage over the exact definition of M which should be used for the monetary theory to work. But the problems with the quantity

theory of money run even deeper. The assumption of the theory that the velocity of circulation of money is constant simply does not hold. In other words, increases or decreases in the money supply are often accompanied, even over a period of years, by changes in the speed with which money circulates. Since the right-hand side of the equation in the quantity theory of money is the amount of money in circulation multiplied by the speed with which it circulates, it follows that if both of these factors change, the direct link between money and prices is broken.

In most Western countries during the past decade, the velocity of circulation of money, on all definitions, has altered with the impact of financial deregulation. Just taking evidence for the velocity of M2 published by the IMF illustrates the point. In America, velocity fell by 12 per cent between 1981 and 1986, and then rose by some 6 per cent by the early 1990s. In Australia, velocity fell 13 per cent between 1981 and 1988, and has since risen by no less than 30 per cent. In France, from the early 1980s to the early 1990s, it rose 14 per cent, but in neighbouring Germany it fell 13 per cent. The most conspicuous fall was in Britain, where over the decade the velocity of M2 fell by 48 per cent. In Britain, in fact, despite the fact that during this period the money supply rose fourfold, prices barely doubled, the potential link between monetary growth and inflation being broken by the collapse in the measure of velocity.

Of course, as we noted above, the velocity of circulation does not have to be strictly constant for the theory to have any chance of success. It is rather that any changes in it have to be predictable. But the deregulation of financial markets in many countries in the 1980s is just one example of a range of factors which has prevented velocity from being predicted with any degree of accuracy. Such events change the rules of the game, as it were, and disrupt any established relationships which might exist.

Perhaps the only example in history of completely successful predictions being made with M occurred in Thomas Middleton's Jacobean play *The Changeling*. One of the characters possesses a magic elixir in a phial marked, quite simply, 'M'. He is able to predict repeatedly and with great accuracy the impact on others of consuming the potion. How unlike the economic forecasts of today! The forecasting problems encountered with the quantity theory of money extend to other aspects of macroeconomics. The media highlight the differences between monetarist and Keynesian economists, and this is not unreasonable, for they can make a lively story. But in many ways, the juxtaposition of monetarist and

Keynesian positions is a false one. Importantly, both approaches share the same Victorian, mechanistic view of the world.

Many of the differences can in fact be understood from the quantity theory of money equation discussed above. In essence, and with some but not a great deal of risk of simplification, the key distinction between a monetarist position and a Keynesian one is that the former holds that changes in M, however defined, can ultimately have no impact on Q, while the latter believes that they can.

This is a very important practical difference for macro-economic policy. Either such policy has no influence on how much is produced in the economy as a whole, or it does have some. The reader may feel, in the words of a catch-phrase fashionable in Britain: 'I think we should be told.' But macro-economics cannot give a clear-cut answer.

Governments have a whole range of macro-economic policy measures which they can take, and changes in many of these measures will have implications for the growth of the money supply in the economy. For example, in the United States there is a gap of hundreds of millions of dollars each year between the amount which the federal government spends and the amount of income it receives from taxes. In Japan, in response to the deepest recession since the war, the government has recently decided to increase the level of its expenditure without raising taxes to pay for the increase. Increased spending on public transport and on housing, for example, is part of the Japanese measures.

But these differences between government spending and its income from taxation have to be financed in some way. At its simplest, unlike the rest of us when we go into debt, the government can simply print some more money to pay for its extra spending. Or it can try to persuade – by, for example, offering high interest rates – both individuals and institutions such as pension funds to lend it the money instead.

In practice, the links between changes in the amount which the government needs to pay for its spending and changes in the money supply can be considerably more complex and circuitous than the basic textbook examples mentioned directly above, and a substantial proportion of the differences between the various schools of thought in macro-economics arises from differences in their technical accounts of these imperfectly understood mechanisms.

But whether governments increase spending without raising taxes, or cut taxes without reducing spending, the resulting need to pay for the gap in the government's budget has implications for the money supply.

Monetarists maintain that such policies, however they are financed, cannot influence the total amount of goods and services which are produced, while Keynesians believe that they can. But just as the monetarist camp is split as to the precise measure of M which is relevant for their theory of inflation, Keynesians are divided on the exact strength of the impact of government macro-economic policy on output as a whole.

Theoretical macro-economics has developed in recent years to produce what effectively amounts to a synthesis of the various schools of thought. And it is this synthesis which is taught the world over on courses in macro-economics. What might be termed pure monetarist or pure Keynesian positions can be regarded simply as special cases of this more general, synthesised model of macro-economics. If certain links in the model are activated with a certain strength, results which are more inclined to the monetarist position are obtained, while the activation of other links gives the model varying degrees of Keynesianism.

An analogy with a railway signalbox which controls a complicated and busy junction might help to illustrate the point. Economists and politicians often speak of the 'levers' of economic policy, even if in the high-tech world of today it is buttons which are pressed rather than levers which are pulled. All the people who draw up the plans for which train is to be routed where share a great deal in common about the operation of the railway. They all agree that the railway has engines and carriages, and that they run on rails. This might seem obvious, but if the signalbox were controlled by someone who imagined he was supervising a horse race such as the Grand National or Kentucky Derby instead, all sorts of problems might ensue. With this common background, the task of the particular signalman on duty is to pull various levers, which activate certain signals and certain patterns on the track, and which in turn choose the direction the train will take. One set of levers will activate the equipment to send the train in one direction, and a different set will ensure that the train ends up on a different track. Similarly, the choice of which connections are activated will decide the direction in which a particular theoretical specification of a macro-economic model goes, towards monetarism or Keynesianism.

The role of the signalman in applied macro-economics is played by people with statistical training, known as econometricians. Their job is to analyse macro-economic data using statistical techniques, to try to decide which links in the system are the appropriate ones to activate. But this is a rather humdrum description of what can be an intricate and intellectually

challenging task. Econometricians working with macro-data may often see themselves in the much grander role of arbitrating the disputes between the various schools of macro-theorists.

I should know, for I spent most of the 1970s and early 1980s as an applied macro-economist, before the gradual accumulation of evidence forced me to the conclusion that, despite intensive effort, econometrics had not succeeded in resolving the various theoretical disputes – nor has it done in the succeeding decade. The weight of evidence may sometimes point very clearly in one particular direction, but it is never quite convincing enough to persuade someone who really does not want to believe the result.

Despite the problems, there is a large number of macro-economic models in regular use around the world, in treasuries, central banks and in large commercial companies, almost all of which are based upon the general, shared theoretical framework of macro-economic behaviour. The models differ in the strengths of their various linkages, which are still the source of endless discussion among applied econometricians, but the underlying approach is common to all.

It is in these models that the theory of macro-economics is applied and confronted with empirical data on a regular basis. Despite the potential importance of this in adjudicating between theoretical arguments, the activity of building and using these macro-models is rather looked down on by the academic profession. An important reason for this is that the interest of most academics is in micro-economics, in the implications of the model of competitive equilibrium.

A secondary reason, perhaps, is that those who get involved in using data are thought by some to be not quite gentlemen, almost as if they were involved in trade, or as if, Heaven forbid, they got their hands dirty. An experience early in my career illustrated the point to me clearly. An aspiring economist and close contemporary of mine, while still a very young man, was highly regarded by the mainstream academic theorists in Britain. Within just a couple of years of taking his first degree, before obtaining a doctorate, he was elected to the economics faculty in the University of Oxford. He proceeded to consolidate his reputation, but after only a few years he decided to abandon academic life and to try and make money as a macro-economic forecaster in business. Great was the consternation among Oxford theorists, and eloquent were the pleas for him to see the error of his ways. The phrase 'prostituting himself', if not

used openly, was hinted at strongly. But to no avail.

As it happens, there is an amusing sequel to this story. Several years later the young economist, then prospering, was approached by a drinks company to see if he could do some lucrative consultancy work for them in a dispute they were having with the European Commission on sales of one of their products. They had, they said, already commissioned another economist to work on this, but they would bring him along. An appointment was made and, on the due day, the management team from the drinks company entered the economist's office. Following them, at least having the grace to look sheepish, was the other economist, the author of their econometric study of the drinks market in the European Community, at last identified as being none other than one of the highest of Oxford's theoretical high priests!

The main practical uses of macro-economic models are not so much to adjudicate between rival theories as to act as a source of advice to governments and to business. A substantial amount of resources has been devoted in the past twenty years to developing and refining these models. In Europe, this has been mainly at the expense of the taxpayer, although in the United States commercial funding of the models has become the norm. But despite this effort and attention, the performance of the models is sadly lacking.

The nature of the advice the models are supposed to be able to provide is twofold. First, to assess the likely consequences, across the economy as a whole, of changes in government policy. The various linkages within the models discussed above trace and monitor, in complicated ways, the implications of changes in one part of the model as they feed through into other parts of the model, often taking time for the full consequences of the various connections in the model to take effect. Second, the models are used to assist in the production of forecasts of factors such as the growth of national output, unemployment, interest rates, inflation, and so on.

As an example of their use to assess the impact of policy changes, VAT is levied on most items of consumer spending in the countries of the European Community. Macro-models can be used to supply answers to the question: what would happen if the rate of VAT were changed? This very question was actually asked of the six leading macro-models in Britain in an exercise to compare their structures carried out in the summer of 1993. For the purposes of the exercise, VAT was assumed to be reduced by one percentage point.

In the first instance, it seems logical that a reduction in the rate of tax on spending would lead to some reduction in the average price at which goods and services across the economy are sold. Indeed, all six models agreed that initially this would happen, but in varying degrees. A couple of models thought that average prices would fall at once by 0.6 per cent, while at the other extreme another model considered that prices would hardly fall at all, by just 0.1 per cent. One reason for the variations is the different ways in which the models represent the links between taxes on consumer spending and the way in which companies set their prices. If VAT changes, companies have a choice of either passing the change on to the consumer through changes in prices, or absorbing the change themselves through their profit margins, leaving prices to the consumer unchanged. The more companies pass changes in tax on to consumers, the more will prices change in response to a change in tax rates.

So the six models differed in their account of what would happen to prices as soon as the rate of VAT was changed. But at least they all agreed that prices would fall. An even bigger disagreement arises when the models trace through the consequences of a change in VAT over a period of three or four years. After four years, two of the models continued to give the answer that prices would fall, and by amounts greater than the initial impact. But one model said that, by then, prices would not have altered at all, while the other three answered that a reduction in VAT now would actually lead to higher prices in four years' time!

In other words, a finance minister trying to decide whether or not actually to change VAT, or a managing director trying to understand the consequences for his or her business of such a change, would be given quite different answers depending on which particular model was selected to tackle the question. In the immortal words of the salesman: 'You pays your money, and you takes your choice.'

Such disagreements between models about the empirical consequences of practical policy changes are widespread, and other examples could readily be supplied, from models both of the British and of other economies. The differences in the answers arise from the cumulative effect of what are often small and apparently insignificant differences in the various linkages within the models. Providing a full account of the reasons for such differences can be a challenging job for the model operators, in the same way that, for example, tracing connections on a complicated electrical switchboard requires skill.

But a flavour can be given of the connections which are activated when the impact of even such a simple policy as changing the rate of VAT is examined. Prices change, and so then do wages, for wage bargainers, whether in trade unions or not, do take account of what is happening to prices. Changes in wages imply changes in industry's costs, which in turn have further implications for prices. A reduction in VAT reduces the amount of income which the government receives from tax, so that the gap between government expenditure and income widens; this gap has to be financed in some way or other, so the money supply has to expand and/or interest rates increase. As prices and wages change, the purchasing power of consumers changes. If the net impact of the changes in prices and wages is to make consumers better off, for example, they will spend more, and some of this spending goes on imports. So the country's trade balance worsens, which again has to be financed in some way, which may have implications for the exchange rate of the country's currency, changes that may feed through into changes in the costs of imported materials, which in turn has implications for prices.

The above paragraph may seem complicated. And so it is, for these models have grown and grown in size. A single model may now be made up of literally hundreds of equations, each purporting to represent some aspect of economic behaviour, and each connected to other equations in the system.

The models may seem impressive and intimidating when their mathematical specifications are set down on paper, but when it comes to their use in forecasting they are so unreliable that virtually no model operator in the world dares allow his or her (although almost invariably his) model loose on its own.

Just after the last world war, the mysterious Howard Hughes had an aeroplane built of truly enormous size. By far the largest in the world, it looked wonderful on the drawing board, and even more marvellous once it was built. The only slight problem was that it could barely get off the ground. It had to be scrapped.

In the same way, the macro-economic models are unable to produce forecasts on their own. The proprietors of the models interfere with their output before it is allowed to see the light of day. These 'judgmental adjustments' can be, and often are, extensive. Every model builder and model operator knows about the process of altering the output of a model, but this remains something of a twilight world, and is not well documented in the literature. One of the few academics to take an

interest is Mike Artis of Manchester University, a former forecaster himself, and his study carried out for the Bank of England in 1982 showed definitively that the forecasting record of models, without such human intervention, would have been distinctly worse than it has been with the help of the adjustments, a finding which has been confirmed by subsequent studies.

The importance of adjusting the output of such models was brought home to me at an early stage. At the end of 1975 I was employed at the National Institute of Economic and Social Research in London, then as now one of Britain's most prestigious macro-economic modelling and forecasting institutions. Purely by coincidence, a number of senior people left more or less at the same time, and I found myself, at the age of twenty-five, in charge of producing the forecast with the Institute's model of the British economy.

At the time, the British economy was in dire straits. Inflation was high, the pound was falling and the balance of payments was in deficit. But there was a glimmer of hope, for the extraction of oil from the North Sea was beginning to come on stream. By the simple process of adding several billion pounds of oil exports to the model's forecast of total British exports and subtracting a similar figure from the forecast of Britain's import bill to represent savings on imports of oil, I was able to produce a forecast for Britain in 1976 and 1977 which, by the standards of the time, was very optimistic. To my amazement and trepidation, the media seized on this sliver of potential good news, and the Institute's forecasts, for the first and probably the last time, were headline news in the mass-circulation, tabloid press. Fortunately for me, this was one adjustment of a model's output which turned out well, although every forecaster will have his or her own personal horror story to tell, over which a veil will be discreetly drawn.

In many ways, the models are like tarantula spiders. Careful handling of both is required, but in the case of the spiders this is not just for the obvious reason. For if their handlers lose control and drop them, they explode. These interesting creatures use hydraulic pressure to extend their legs, and the fluid splatters on impact. In the same way, macro-models have a particular fascination of their own, and a tendency to explode.

In recent years, even the combination of models supported by the judgments of their operators have produced forecasts which have been seriously misleading. During 1992 and early 1993 alone, big errors have been made in forecasts. We have seen that the Japanese recession, by far

the deepest since the war, was not predicted. Neither the strength of the recovery in America in the second half of 1992 nor the slow-down in the first half of 1993 was really anticipated. And in Europe, neither the turmoil in the ERM nor the depth of the recession in Germany was foreseen by the models.

A survey published by the Paris-based international body, the Organisation for Economic Co-operation and Development (OECD) in June 1993, illustrates the problem quite clearly.

The forecasting records of the two major publicly funded international bodies, the forecasts of the OECD itself and the IMF, and of the national governments were compared. For the major seven world economies, the forecasts for the next year ahead for output growth and inflation were examined.

The benchmark used by the survey for comparison was a naive projection that next year's growth of output or inflation would simply be equal to this year's. In other words, this benchmark required no knowledge of economics to produce, and a forecast could be made with it by anyone who understood the elementary arithmetic of percentage changes.

Over the 1987–92 period, this extremely simple rule performed at least as well as the professional forecasters in projecting next year's economic growth rate. And in terms of inflation, the rule performed as well as the OECD and IMF, and slightly better than the national governments. In other words, the combined might of the macro-economic models and the intellectual power of their operators, whether based in national governments or installed in tax-free splendour at public expense in Paris or in the IMF in Washington, could not perform any better than the simplest possible rule which could be used to make a forecast.

The record of forecasting is poor whatever the theoretical nuance of the model concerned, whether it leans towards monetarist or towards Keynesian properties. A survey of the accuracy of British economic forecasts, for example, carried out by the London Business School in 1993, concluded that differences over time between the predictions from the various schools of thought are very small. But the most striking fact to emerge from this study is that errors in forecasts are much greater than differences between apparently contending schools of thought. This is by no means a new discovery, but it represents valuable confirmation of previous studies over the years which have come to the same conclusion.

The best recipe for forecasting success, conclude the London Business School researchers, displaying a degree of irony as welcome as it is rare among economists, is to 'forecast often and forecast late'. In other words, the more forecasts made during the course of the year, the greater the chance that, purely at random, one of them will prove to be reasonably accurate. And by forecasting as close as possible in time to the actual period being forecast, much more information becomes available about what is likely to happen. Of course, this information is not confined by some secret code to economic forecasters. It is information in the public domain, available to anyone wishing to make an informed guess about the prospects in the immediate future.

Occasionally a forecaster will get things broadly right for a couple of years, and be lionised as a hero. But there is no guarantee that next year's forecast will not be the worst of the lot. Indeed, this is not merely a supposition. It has actually happened a number of times, as I once learned to my cost, back in the late 1970s.

All the problems of macro-economic models – the contradictory answers different models give to the same question, the poor forecasting record, the inability to trust a model on its own – exist despite the effort devoted to their maintenance and construction. This is despite the fact that model builders and operators, particularly in Europe, where a greater proportion of their work is funded by grants from the taxpayer than it is in America, pride themselves on incorporating the latest nuances of macro-economic theory into the specification of their models.

The overwhelmingly dominant fashion in macro-economic theory in the past fifteen years or so has been the concept of 'rational expectations'. Such is its hold on the profession that it has been scarcely possible in recent years to publish articles in many journals, either on theoretical or applied macro-economics, without using it.

The idea, which has been incorporated extensively in the applied macro-economic models discussed above, has a powerful attraction for economists. As we have seen in previous chapters, economic theory at the micro-level is based upon the concept of the rational, self-interested individual, and the properties of the model as a whole are built up by adding together the equations which describe the behaviour of people and companies.

Macro-economists had a long-standing concern that their own theoretical models, intended to describe behaviour of the economy at the

aggregate level, were not founded so clearly on the principle of individual rationality. Spending by individual consumers, for example, is perhaps the single most important feature of any developed economy, accounting for well over half the total amount of spending in every Western country. From the very beginnings of macro-economic models in the late 1940s and early 1950s, equations have been built to try to describe the behaviour of consumer spending. Various factors are thought to influence the total amount consumers spend. Their incomes, for example, are obviously important, and possibly their wealth, the amount of debt they hold, the level of unemployment, interest rates and so on. Econometricians, using statistical techniques, have worked hard to discover both which factors really do influence consumer spending, and how strong each influence actually is.

But leaving aside the fact that there is still no agreement on the best set of factors to use when trying to account for the total amount which consumers spend in this way, the process of simply searching the data with econometric techniques leaves many economists uneasy, for such results could be seen as mere descriptions of fluctuations in economic series at the aggregate level, giving no account of why individuals behaved as they did, in ways which led to these movements in the data series for spending as a whole.

Rational expectations appeared to offer a solution which would place macro-economic theory on what, for economists, was a much more secure foundation. Expectations are obviously important in a range of economic decisions, such as the income which an individual expects to receive over the next few years, or the rate of return expected on an investment. But it is not usually possible to observe or measure expectations directly, apart from in certain specialised areas of financial markets where, as we mentioned in the previous chapter, deals can be struck about prices which are expected to obtain in the future. There are also surveys of spending plans by companies, or of the degree of optimism felt by consumers, but in general data on expectations are fairly sparse.

The idea of rational expectations is that 'rational' economic agents – in other words, people and companies – will not just learn from their mistakes in forming expectations, but that the learning process will enable them to identify the true model determining the behaviour of the economy. They will then use this model in order to form expectations. It is as if (that favourite phrase of economists) everyone not only has his or her own personal macro-model of the economy, but everyone has the

same model and, to stretch credibility even further, a model which happens to be a true and correct representation of how the economy behaves. This model will then be used to form expectations. By definition, economic agents are 'rational', so they should use this rational method of forming expectations.

Apart from its logical attraction for economists, the concept had the added bonus, at both the theoretical and applied levels, of creating lots of opportunities to indulge in fascinating intellectual exercises, replete with mathematical manipulation in the differential calculus. But even a moment's reflection might lead one to believe that many individuals have very little grasp of many of the true models which govern their everyday lives, let alone a personal copy of the correct macro-model of the economy.

In the present day, for example, at the very end of the twentieth century, surveys appear from time to time in all Western countries which show that many adults think the Sun revolves around the Earth. On a more trivial level, a survey published in Britain in the spring of 1993 revealed that dog owners attributed a level of intelligence to their pets which was on average that of a nine-year-old child. A substantial minority believed their animals to be more intelligent than university students.

An orthodox economist would dismiss these examples as frivolous, since it is not necessary to know either how the solar system operates or how intelligence in different species can or cannot be compared in order to maximise one's individual economic welfare. They equally dismiss another serious problem for rational expectations – the fact that, while ordinary consumers are presumed to know the true model of the economy as they go about their daily business, economists themselves in their professional capacities clearly do not know the 'true' model which governs the economy. This particular problem is again dismissed by invoking the trusty phrase 'as if'.

Baseball players or cricketers do not need to be able to solve explicitly the non-linear differential equations which govern the flight of the ball. They just catch it. So it is always possible to argue that people act 'as if' they knew the true structure of any particular system, without necessarily being able to articulate it.

However, a curious form of schizophrenia is invoked by this argument. A team of applied mathematicians playing baseball or cricket could solve the equations governing the flight of the ball. So a group of economists

who, in their roles as consumers and employees are assumed to act rationally and act as if they know the true structure of the economy, ought to be able to articulate this in their professional capacity. This, as we have seen, they are singularly unable to do.

It is not surprising that the economist's concept of rationality does not bear scrutiny at the aggregate level. Given its crucial importance, the amount of empirical work which has been done to test its validity is ridiculously small, but some interesting work has at last been done in the past decade.

In the mid-1980s, several studies in Britain and the United States examined data from surveys of economic expectations to see if they stood up to the tenets of rationality. As mentioned above, such surveys cover factors such as the expectations of business about investment and inflation, and the overall confidence of consumers. Careful examination of the evidence by researchers such as Lovell in America and Ken Holden and David Peel of Liverpool University in Britain concluded that, in the detached language of science, 'the hypothesis of rational expectations did not appear to be consistent with the data'. In other words, it was wrong.

These studies seem conclusive. But this is to reckon without the intellectual dexterity of the orthodox, who with an incisiveness rivalling that of Perry Mason or Columbo, argued that surveys of business opinion could not be used to provide evidence either for or against the hypothesis of rationality. Companies which are asked to complete a survey questionnaire have no economic incentive to make accurate forecasts in it. Since they are not paid, they do not need to make accurate responses.

To argue that, because no money is at stake, companies might not transcribe their genuine views into surveys carried out by government bodies and by their own business and trade associations is to ignore the fact that companies have every incentive to construct accurate forecasts for their own business decisions.

More recent work in the United States has examined data generated by groups who *do* have a direct incentive to get things right – namely, economic forecasters. As we have noted already, their forecasts are financed almost entirely on a commercial basis, and yet several studies using this data conclude that even forecasts generated by professional economic forecasters do not support the hypothesis of rational expectations. For example, in 1993 Baghestani and Kianian of Colorado published a paper which examined American macro-economic forecasts over the 1981–91 period. They concluded that, out of seventy-five forecast

series examined, only two satisfied the conditions of rationality. In other words, no fewer than seventy-three of them failed to support the concept of rational expectations.

The techniques of behavioural science have also begun to be used to test the concept of rationality from a number of perspectives, with some devastating results. We saw in Chapter 2 that a recent study from Cornell University, the latest of an increasing number of such papers, showed quite clearly that people, except those whose minds had been scrambled by the study of micro-economic theory, are far more co-operative and less competitive than the postulates of economic theory assert rational individuals should be.

A paper published as long ago as 1982 by Joël Huber illustrated yet another dimension of apparent irrationality in consumer behaviour. A group of students was asked to choose between two apartments, one close to college but expensive, the other further away but cheap. The distance and rents were calibrated so that each apartment was chosen by roughly half the students in the experiment. A second group was then asked to compare the two, but this time a third apartment was on offer. This was both the furthest away from college of all three, and more expensive than the one further away in the original experiment. The addition of this third choice increased dramatically the proportion of students choosing the second apartment. In other words, instead of dividing equally between the first two apartments on offer, many more now chose, for no apparent logical reason, the second one. According to the economic concept of rationality, this should not happen.

To take yet another example, theory requires that because, by definition, individuals are rational, the preferences of any individual must be internally consistent. For example, if a consumer prefers Coca-Cola to Pepsi-Cola, and Pepsi-Cola to lemonade, then he or she is deemed to prefer Coca-Cola to lemonade. This characteristic of preferences is known technically as transitivity, and must hold if an individual is to have consistent preferences. If product A is preferred to product B, and product B to product C, then, for consistency, A must be preferred to C.

A paper by Graham Loomes and colleagues at the Centre for Experimental Economics at the University of York was published in 1991, to the credit of the journal, in *Econometrica*, one of the world's leading technical economic journals. The paper provides strong evidence against the view that consumer preferences are transitive. Loomes approached the question using the methodology of behavioural science, constructing

games which people played under controlled conditions, and analysing their behaviour. A number of papers in the 1980s had used the same methodology and had also rejected the concept of transitive preferences, but their exact approaches had been the subject of detailed criticism. Loomes and his collaborators designed their experiments to meet these specific objections, and still their results rejected transitivity.

The implications of the non-transitivity of individual consumer preferences are profound.

Economists believe that consumers reveal their true preferences through their actual purchasing decisions. Indeed, the phrase 'revealed preferences' is used in theoretical models of consumer behaviour to describe this activity. Firms observe the revealed preferences of consumers, and the market mechanism, the Invisible Hand of Adam Smith, comes into play to ensure that these preferences are met. The concept of revealed preference is essential to the view that 'the market knows best'. By acting in their individual self-interest in a competitive environment, companies are able to satisfy the true preferences of consumers.

The axiomatic claim that the market mechanism is satisfying the true preferences of consumers is undermined if these preferences themselves are non-transitive. For the signals sent by revealed preference are then potentially contradictory. Companies receive conflicting messages about the range of goods and services which should be provided to satisfy consumer needs.

The soft drinks example used above illustrates the point. Suppose a small neighbourhood shop has space on its shelves to stock only two types of soft drink, Pepsi and lemonade. Sales of Pepsi heavily outsell those of lemonade. The proprietor, a firm believer in the free market, is anxious to please his customers by offering them a choice, so he stops buying lemonade and puts Coca-Cola on his shelves instead. He notices that most people switch to this, and it is the sales of Pepsi which languish. But one week the delivery van has no more Pepsi available, so the shop-owner has some lemonade put back on his shelves, alongside the Coca-Cola. To his bewilderment, lemonade sales are much higher than those of Coca-Cola. When the delivery van calls next week, what should the poor man do to satisfy his non-transitive customers? The market has given him contradictory signals.

The appropriation of the word 'rational' to describe the basic postulates of orthodox economic theory was a propaganda coup of the highest order. The world's most expensive public-relations firms could not have

done better. It carries the implication that any criticisms of it, or any alternatives put forward, are by definition irrational, and hence not worthy of serious contemplation.

But those economists bold enough, along with colleagues from other disciplines, to test and challenge the traditional view of rationality find it seriously wanting.

And the large-scale macro-economic models, the meeting place of macro-theory and empirical data, fully equipped as they have been with the economist's concept of rationality, have made little or no progress in improving our ability to understand and predict the world.

By any reasonable criteria, the discipline of economics as a whole, in its present state, is sadly lacking. Encouragingly, more and more attention is being given to alternative approaches which offer hope of escape from the present impasse. In the words of St Paul's Epistle to the Corinthians: 'When I was a child, I understood as a child, I thought as a child. But when I became a man I put away childish things. For now we see through a glass darkly; but then face to face.' The exciting possibility is opening up that one day we will be able to see the workings of the economy 'face to face', and the focus of this book now moves to the articulation of ways by which we might indeed obtain a better understanding of the behaviour of economies at the macro-level.

PART II

TOWARDS THE FUTURE OF ECONOMICS