The behaviour of the foreign exchange market

By Professor Alan Kirman.[1]

In this article, Alan Kirman considers what developments in economic theory have to contribute to an understanding of the recent evolution of the foreign exchange market. After outlining the standard, efficient-markets *model of the workings of the market, he looks at various reasons why that model has been questioned and examines the extent to which alternative models can offer a better explanation of the market's actual behaviour.*

Professor Kirman was a Houblon-Norman fellow at the Bank in August of last year.[2] The views expressed in this article are his, rather than those of the Bank.

The global foreign exchange (FX) market had a daily turnover of about $1.3 trillion in 1992, corresponding to a net daily turnover of $880 billion; this represents growth of 40% in three years.[3] 60% of this turnover is accounted for by the three main centres—the United Kingdom, the United States and Japan—of which the United Kingdom is the largest, with 60% more turnover than the next market, the United States. The market is becoming increasingly active; on the Reuters electronic dealing system, as many as 40,000 electronic 'conversations' occur per hour and there are 4,000 banks worldwide linked to this system with some 18,000 terminals. The market is also more and more global, with 60% of the daily transactions in 1992 being cross-border and with 80% of the aggregate FX turnover in London, for example, being done by foreign banks.[4]

Globally, 50% of gross foreign exchange transactions involved non-local currencies on both sides of transactions. The growth of the FX market has been much more rapid than that of foreign trade, giving weight to the idea that a growing percentage of the volume is accounted for by dealing for speculative purposes. This has led a number of commentators to argue that the market is becoming intrinsically unstable. The view is characterised by one economist, who says:[5]

> 'These [foreign] exchange transactions began as a means to smooth and facilitate the flows of traditional trade and investment. But this FX 'tail' has grown to be some hundred times larger than the original trade 'dog' . . . FX is a speculators' paradise.'

This does not take account of an alternative view, that the broadening and deepening of the FX market has allowed market participants to protect themselves against over-exposure by trading with other dealers. In this view, the increased volumes in the FX market simply reflect prudent behaviour on the part of dealers.

Whatever view is correct, the question remains whether the expansion of the market has been stabilising or destabilising. As movements in foreign exchange rates have become larger and more rapid—and with events such as those which followed 'Black Wednesday' (16 September 1992)—there have been calls from some authorities and a number of economists to impose some sort of control or restriction on the market. These have been reinforced by events such as the Swedish government's inability to maintain the level of the krona, even by raising overnight interest rates to 500%, and the depreciation of the lira.

Before considering such calls, however, it is worth examining what recent economic theory has to contribute to understanding these developments in the FX market.

There are a number of questions to be answered. Have the globalisation and increase in volumes in the FX market *of themselves* increased the volatility of exchange rate movements? Are exchange rates less linked to 'fundamentals' than they were, or have the fundamentals themselves changed? If there is less of a link to fundamentals, does that mean that dealers are behaving *irrationally*? In particular, are the large exchange rate movements evidence of irrationality or could they in fact reflect rational behaviour by participants? Would the introduction of 'frictions' into the market reduce the volatility of price movements?

I argue in this article that the increasing size and connectivity of the FX market may have led to increased

(1) Alan Kirman is Professor of Economics at the European University Institute in Florence.
(2) The Houblon-Norman Fund, established by the Bank in 1944, finances academic research into subjects relevant to central banking. More details of the Fund were given in an article in the August 1993 *Quarterly Bulletin.*
(3) The figures here are taken from 'Central bank survey of foreign exchange market activity', a document from the Bank for International Settlements' Monetary and Economic Department, Basle, February 1993, pages 1–42.
(4) Source: 'The foreign exchange market in London', *Quarterly Bulletin*, November 1992, pages 408–17.
(5) From Ohmae, K, *The borderless world*, Harper Business, New York 1990, quoted in Zaher, S, 'Market-makers: A study of the effects of global market integration in the currency trading industry', monograph, School of Management, University of Minnesota, 1994.

volatility. In addition, however, the structure and organisation of the market have to be understood as playing an important role. Recent economic theory, which takes account of the fact that individuals draw information from the actions of others in the market and may imitate the successful, suggests that 'herding' or 'informational cascades' may occur. Nevertheless, this sort of behaviour is by no means necessarily irrational and should not be attributed to some inexplicable market psychology.

Thus both those who argue that the foreign exchange market is now more efficient at revealing underlying imbalances in fundamentals and those who maintain that the intrinsic dynamics of the freer and more open markets have led to greater instability can find comfort from theory and empirical evidence.

One associated argument is as to whether the volatility of markets could be reduced by the introduction of frictions, such as a tax on FX transactions. On this, I suggest that in the light of the evidence we would do better to use our improved understanding of the dynamics of foreign exchange markets, and of the role of their microstructure, to increase the efficiency of interventions to help maintain a certain orderliness, rather than impose restrictions on transactions.

In order to answer the questions posed above, it is necessary to outline the standard, 'efficient markets' model of the foreign exchange market, which suggests that prices should be linked rather strictly to fundamentals and therefore that the globalisation of financial markets should not radically affect prices. I then mention some empirical paradoxes and look at various reasons why this basic model has been questioned, and at whether alternative models offer better explanations of the facts.

The standard efficient-markets model

The behaviour of asset prices in general—and foreign exchange rates in particular—is typically explained by economists by *efficient-market* theories. This term, although used in different ways, basically reflects the following ideas.

The price of an asset should reflect underlying 'fundamentals'. Thus in the case of a stock or share, its price is supposed to reflect the discounted value of all dividends expected in the future and the value of the firm at the end of the life of the claim, if it is for a fixed term.

Secondly, the link between fundamentals and prices is such that all information available about fundamentals—both public and private—should be incorporated into prices. If this is the case, then the only reason for prices to change must be the arrival of completely new information which was not predictable. (If it had been predictable, it would have been predicted.) Hence asset prices must appear to fluctuate randomly; whether they do has been the subject of extensive debate.

In the case of exchange rates, the efficient-market hypothesis has to be rephrased. Suppose that assets are priced efficiently—in the above sense—in two different countries. Then changes in the exchange rate should effectively equalise the rates of return in both countries: this is so-called 'uncovered interest parity'. It can be objected to on the grounds that it does not take account of the risk involved in such transactions. However, using the forward market allows us to make a risk-free transaction and the forward exchange rate should equalise the rate of return. This is the 'covered interest parity' condition.

In either case, we are back to the original idea: fundamentals determine the rates of return in each country and exchange rates adapt to these. Thus modifications in exchange rates reflect changes in fundamentals. And changes in fundamentals will lead to compensating changes in exchange rates.

One important feature of this approach is that asset values or exchange rates reflect what is *expected* to happen to fundamentals. Since expectations are not measurable, it is clearly not possible to falsify the efficient-market hypothesis directly.

Problems for the efficient-markets model

Yet many of the facts about financial markets seem to be at odds with this kind of theory. Perhaps the most striking is the volatility of asset prices, compared with that of the underlying fundamentals. Despite numerous efforts by economists to explain it, the 'excess volatility' puzzle remains. Why should it be in the case of stocks, for example, that prices are so much more volatile than the associated dividend streams?(1) Once again, it can be argued that the relationship between fundamentals over time is highly 'non-linear' and that small changes in today's values may lead to large changes in the future, thus significantly changing the current price. It is difficult to believe, however, that there could be a sudden change in the fundamentals which would lead agents simultaneously within half a day to the view that returns in the future had gone down by over 20%. Yet this is what would have to be argued for the October 1987 episode on the New York Stock Exchange.(2)

The same is true for sudden and substantial changes in exchange rates. Do they really simply reflect modifications in expectations about future fundamentals? Why does the volatility of exchange rate changes vary over time? How does one reconcile the two ideas frequently expressed by traders, that on the one hand 'fundamentals matter in the long run' but on the other they do not drive exchange rates in the short run?(3)

(1) See Shiller, R J, *Market volatility*, Cambridge Mass: MIT Press, 1989.
(2) Indeed Miller has suggested that substantial changes in the future can result from very small changes in the present, and that such an explanation is not inconsistent with the Crash; see Miller, M H, *Financial innovation and market volatility*, Blackwell, 1991.
(3) See Goodhart, C A E and Figliuoli, L, 'Every minute counts in financial markets', *Journal of International Money and Finance, 10*, pages 23–52, 1991.

Two things have to be observed before these questions can be addressed. Firstly, changes in prices or exchange rates reflect changes in economic agents' *perceptions* of the future, and not necessarily what will actually occur. Secondly, a clear corollary of the efficient-markets view is that if asset prices change without any obvious change in the fundamentals, some agents must be acting irrationally.

For this reason the sort of empirical puzzles mentioned are frequently explained in terms of market psychology. The implication is that the movements involve some degree of irrationality on the part of those participating in the markets. Sudden changes or departures from fundamentals are taken not as evidence of the inapplicability of efficient-market theories but rather as evidence of a failure of investors or traders to act as rationally as those theories require.

To take an example, in his classic book on market volatility Shiller refers to alternative explanations of price movements as being associated with 'capricious' behaviour or as being 'for no good reason'.(1) His argument to explain 'excess volatility' is that investors may take actions as a result of a movement in asset prices which result in a further change in prices; a sequence of such events may lead to a 'price bubble' which detaches prices from fundamentals. Despite his insistence on the importance of this behaviour, Shiller seems to emphasise the irrationality of such self-induced price movements.

Before considering the argument about irrationality, however, a basic point has to be made. There is no clear consensus concerning the nature of the relationship between fundamentals and prices in many financial markets. In the foreign exchange market, most participants believe in the existence of a relationship between exchange rates and certain macroeconomic fundamentals. But such a relationship is difficult to estimate and may well not be invariant over time, even if one knew which macro variables were important. Thus even without any shock, it is easy to see that there is a potential source of instability. If traders or participants in the market change their view about the nature of the relationship—or about the particular macroeconomic variables which are important—they may, by their very actions, modify the relationship and make it self-fulfilling.

The well known literature on *sunspots* makes this point in an even more striking way.(2) If market agents believe that prices are correlated with sunspots, they will buy and sell accordingly and, as a result, prices will indeed become coordinated with sunspots. Yet one could ask how would this come about, and it can be shown that sensible agents using sensible learning rules may come to believe in the importance of sunspots and that their beliefs will be self-confirming.(3) Thus movements of exchange rates which are not directly correlated to movements in fundamentals are not necessarily the result of irrationality.

An alternative kind of model

What Shiller does show, despite his comments on the irrationality of such behaviour, is the importance of agents' reactions to one another's behaviour. The important point to make is that in financial markets agents do indeed interact directly with one another and not only indirectly through market prices. This apparently innocent remark has significant consequences for the aggregate behaviour of markets. Instead of thinking of a single 'typical' economic agent's response to events, one should consider individuals in a market as not necessarily being homogeneous, and as observing and anticipating the behaviour of other participants. Once one does this, it is much easier to see how a common view can take over a market temporarily and then be replaced by another view.

Once we consider the market as a complex interactive system in which heterogeneous agents with different horizons and different attitudes to risk participate, the price dynamics can be very different from those of more conventional models. Although this sort of idea is familiar to mathematicians and physicists—a number of whom are now employed by major financial institutions—it has only recently influenced the development of economic models of financial markets.

Such ideas are, on an informal level, far from new. In economics, Keynes' beauty contest example has been frequently discussed. Keynes' point was that in deciding which contestant would win a beauty contest one should not take into account one's own judgment, but rather should try to assess which of the candidates was likely to be most pleasing to the judges. The argument can be extended to a situation in which there is a popular vote to decide the winner. And the reasoning can be used in the case of financial markets. Keynes made the point that trying to act with the majority was important for a manager of funds: he or she is less likely to be criticised for making an investment which turns out to be unprofitable if many other market participants made similar investments than if it was purely the result of his or her own judgment.

Arguments such as those of Keynes, while frequently evoked and suggesting that it is not necessarily irrational to 'follow the herd', have until recently been largely anecdotal. However, the particular argument underlying the beauty contest example has now been developed formally,(4) and in the context of a principal-agent relationship it can be shown that agents who invest on the part of others may have strong incentives to imitate the actions of the market participants that they observe. The idea is a simple formalisation of Keynes' notion that sanctions are asymmetric and in this case the agent will have every interest to conform. Yet this may well result in conformity of a sort which is not efficient from a welfare point of view.(5) However, conformism as a

(1) See Shiller (*op. cit.*).
(2) See Cass, D, and Shell, K, 'Do sunspots matter?', *Journal of Political Economy, 91*, pages 193–227, 1983.
(3) See Woodford, M, 'Learning to believe in sunspots', *Econometrica, 58*, pages 277–307, 1990.
(4) See Sharfstein, D S, and Stein, J C, 'Herd behaviour and investment', *American Economic Review, 80, 3*, pages 465–79, 1990.
(5) The same sort of thing can occur when firms who adopt a technology provide a positive effect for other users of the same technology; a whole industry can get locked into an inferior technology. See Arthur, W B, 'Competing technologies, increasing returns and lock-in by historical events', *Economic Journal, IC*, pages 116–31, 1989.

form of risk aversion is but one of several explanations of why individuals may be influenced by the actions or opinions of others.

It is important to take account not only of the interaction between agents but also of *how* that interaction takes place, and which individuals interact with—or react to—which others. That will depend on the way in which the market is organised; and there is a growing interest in how market microstructure affects the evolution of prices.(1)

Of course, those who maintain the market-efficiency point of view could argue that if different structures give rise to different prices, some of these structures must be unsatisfactory from the point of view of economic efficiency. Empirically though, it may be very difficult to establish this.

All of these arguments suggest that a satisfactory model of financial markets should include the following features:

- Agents should react directly to one another's behaviour.
- The heterogeneity of agents—in terms both of expectations and horizons—should be included.
- The market microstructure, in particular the network of communications within which agents operate, should be considered.

The remainder of this section looks at each of these points and presents a simple example of an economic model which incorporates such features to see how its behaviour compares with the empirical data.

Inferences from the behaviour of others

One of the most important features of markets is that the actions of individuals reveal something about the information that they possess. This feature is poorly incorporated in most economic models and difficult to include in the efficient-markets framework.

To take a few examples that economists have considered, let me look first at the case in which agents receive private signals but also observe the choices made by others. So in the foreign exchange market, in addition to any information acquired from a private source, a trader observes what other participants are doing—or at least proposing to do. If the agent changes his action in the light of this information, a so-called 'information cascade' arises.(2) As more and more individuals act in this way, a trader would have to have almost unbounded confidence in his own information not to conform, particularly if such cascades lead to self-fulfilling outcomes.

In the FX market, a trader's goal is to anticipate the direction of movements in market prices, so he or she gains a great deal of information by listening to the brokers, watching the bid and ask prices on the screens, and telephoning other traders to ask for a quote. Each piece of information modifies the individual information set, but since there is no central equilibrium price this information cannot be incorporated and become public through that price, but only through the observable actions of the individual. The problem with information cascades is that as the number of people involved increases, the cascade reinforces itself. Although quite fragile to start with, cascades later become almost immune to relevant information.

There is a significant loss of efficiency here. The information acquired by early agents would be of use to their later counterparts, but if they choose to follow what others do this information is not made available. In this way, possibly relevant information about fundamentals, for example, might never be used and prices could get detached from these fundamentals. A conclusion that can be drawn from work by a number of economists is that the information obtained by observing the actions of others can outweigh the information obtained by the individuals themselves. It is also clear that as the behaviour of market participants becomes more and more instantly observable—with the development of modern communication technology—the probability of cascades is increased.

Imitation

A second source of herd behaviour is the tendency to imitate those who are successful. This can occur in two ways. Either individuals may be converted to the beliefs held by their successful counterparts, or they may simply imitate directly the choices of the successful. This, in itself, might merely imply a learning process which would lead less successful participants to improve their performance. However two things can happen. An individual may become successful as a result of some chance event, or series of chance events. The fact that he is then imitated may lead to the market moving in the direction he predicts, or it may end in a collapse of what will become apparent was a bubble. It is possible indeed that imitation of success will lead to perpetually changing patterns of behaviour in the market.(3)

Market microstructure: network effects

The way in which a market is organised can have important consequences for the way in which prices evolve. In a market in which there is no centralised price determination, agents will trade with and observe other traders. But traders do not pay equal attention to all the other traders operating in their currencies. Typically, they operate with a limited subset of partners and there are clear reasons for this, in terms of the time cost of monitoring and communicating with others. So the market may be viewed as a complex

(1) See O'Hara, M, *Market microstructure theory*, Cambridge Mass: Blackwell, 1995.
(2) See Hirschleifer, D, 'The blind leading the blind: social influence, fads and informational cascades', *Finance Working Paper No 24–93*, School of Management, UCLA, 1993; Bikhchandani, S, Hirschleifer, D, and Welch, I, 'A theory of fads, fashion, custom and cultural change as informational cascades', *Journal of Political Economy, 100*, pages 992–1,026, 1992; Banerjee A, 'A simple model of herd behaviour', *Quarterly Journal of Economics, 108*, pages 797–817, 1992; Welch, I, 'Sequential sales, learning and cascades', *The Journal of Finance, 47*, 1992; and Kirman, A P, 'Communication in markets: a suggested approach', *Economics Letters, 12, No 1*, pages 101–8, 1983.
(3) See Ellison, G, and Fudenberg, D, 'Rules of thumb for social learning', *Journal of Political Economy, 101, No 41*, pages 612–43, 1993.

A simple example

This example illustrates how some of the features discussed can be incorporated into a model of a financial market. Consider a simple situation in which opinions in the foreign exchange market are divided between those who are *chartists*—ie who extrapolate prices in a more or less sophisticated way—and those who are *fundamentalists*, who believe that prices are essentially determined by underlying fundamental values.(1) As people meet each other in the market (where 'meet' may mean observing an action, making a telephone call or receiving a signal), they will be influenced by the expectations of those they meet, and will be recruited to another's opinion—if different from their own—with a certain probability. Clearly if they keep meeting or observing individuals who share their own opinions, these can only be reinforced. Thus the proportion of people in the market holding a certain opinion will fluctuate as a result of the sequence of meetings that take place.

Individuals then try to assess what the majority expectation is and make their bids accordingly. And the exchange rate is then set to clear the market.

What can be shown in such a simple model is:(2)

- that the proportion of opinions will never settle down, but will continually change;
- at any time, the individuals will be nearly all chartists or nearly all fundamentalists, but periodically the market will switch from being dominated by one to being dominated by the other; and
- that although a market currently dominated by chartists will always return to fundamentals, the time at which it will do so is indeterminate.

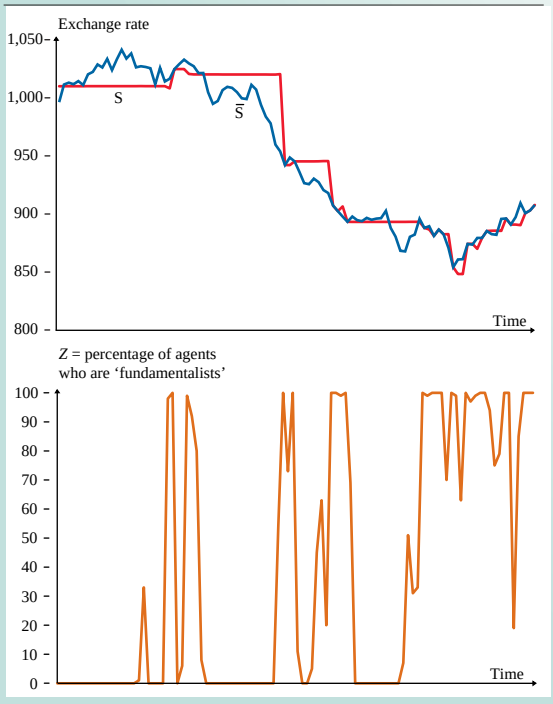

A simulation of such a model is shown in the chart. Z is the percentage of agents who act as fundamentalists, S is the exchange rate and $\overline{S}$ is the exchange rate were it to be completely determined by fundamentals. What can clearly be seen is that the exchange rate S moves away from $\overline{S}$ for a while and then returns sharply. The example has been constructed so that periods of chartist domination maintain the exchange rate constant. This has been done in order to make the figure more readily interpretable, but is just an artefact of the parameters.

Three things are worth noting:

(i) The equilibrium of such a market should be thought of in terms of the so-called 'limit' distribution—the proportion of time the system spends in each state.

(ii) In the model, 'herd behaviour' is rational—in that it is more profitable to act with the majority, particularly if it is known that others do so.

(iii) Individuals do not systematically make mistakes but switch opinions and are justified in doing so, since most of the time their expectations are self-fulfilling.

These features are clearly heavily dependent on the assumption that decisions take place—and expectations modified—in a sequential way. This, however, seems to be an appropriate assumption, given the way in which the FX market works. The rapidity with which meetings, conversations and observations occur in the modern FX market would, of course, have a significant impact on the time it takes to switch from one extreme to the other. *Ceteris paribus*, the more frequent these meetings, the more frequent the change of 'regime'.

Such a model enables one to formalise the idea that in the short run a movement may persist, even though it seems to run counter to what the fundamentals would indicate. Since no information is to be gathered from how long one has been in one state as to when opinions in the market might shift, there is little to be gained from taking a position on the basis of a return to fundamentals at some indeterminate time in the future. This is particularly important since 90% of FX volume is accounted for by intra-day, ie short-term trading.

It is also worth noting that econometric standard tests as to whether the movement of the exchange rate is a random walk—which is what it would have been had it always followed fundamentals and thus satisfied the 'efficient-market' criterion—failed to detect the presence of strong deviations from the fundamentals. So these tests do not seem capable of detecting the presence of bubbles in a series which for part of the time follows a random walk. Furthermore well-known tests for time-dependency in the volatility of the series (ARCH, GARCH, etc) did not reject the presence of such an effect, despite the fact that—by construction—the sort of time structure these tests are supposed to detect was not present in this model.

(1) For a similar model, see Frankel, J A, and Froot, K A, 'The dollar as an irrational speculative bubble: the tale of fundamentalists and chartists', *Marcus Wallenberg Papers on International Finance*, *1*, pages 27–55, 1986.
(2) For details, see Kirman, A P, 'Ants, rationality and recruitment', *The Quarterly Journal of Economics*, *108*, pages 137–56, February 1993; and 'Testing for bubbles', *mimeo*, European University Institute, Florence, 1994.

network, with each participant being linked to a subset of others. Not all of these links will be in use at any one time and indeed they will be used with a certain probability, depending on the terms offered and the positions held by the various partners. Thus the market can be viewed as a 'stochastic or random graph'—*graph*, since traders are linked with certain other traders; and *random*, since these links are only used with a certain probability.(1)

The connectivity of this network will be of considerable importance for the transmission of information and for the speed with which a particular view takes over in the market. One thing that is known is that if the probability that any two agents in the market are in contact with each other is not too small, the larger the market the faster information will disseminate. Why is this true?

Suppose that there are N individuals in the market. These individuals trade with or observe only a limited number of others. Thus the reaction by other agents to an action of one will not be instantaneous, but will take place only when one of the agents with whom they are in contact reacts. The reaction will 'percolate' through the system. Assume, for example, that each of the agents observes $\sqrt{N}$ others, ie in a market with 900 participants each individual observes 30 of his counterparts, which does not seem unreasonable. One can show that, if agents do this and N is large, it will take only two steps before every agent is alerted to the fact that an action has taken place.(2)

So, perhaps counterintuitively, epidemics or herd behaviour are more likely to develop rapidly when there are many agents in the market. Two things offset this, however. Traders are not all linked with equal probability to others. The global market is, in fact, still quite strongly segregated into three regions: Asia, Europe and North America. There are troughs of activity at around 4.00 am GMT and between 7.00 pm and 11.00 pm GMT as regional markets open, close or diminish activity. This, together with the second observation—that currencies tend to be traded more specifically in their own markets—probably slows down the transmission of reactions, and may diminish the effect of any particular local movement. (Globalisation may, in this sense, be destabilising if it leads to more integrated markets.)

Although it is clear that the network of communications that traders use is important, what is more difficult to analyse is how the structure of the network develops in the first place. How do traders choose their partners? Why do the probabilities of trading with others evolve away from the uniform situation? Although economists are now paying some attention to this sort of question,(3) little formal analysis has been done. A typical feature that has to be explained is the advent of traders or clients who become the focus of attention of many members of the networks. Their actions are closely monitored and often imitated for a period, and then the links to them become less important as attention switches elsewhere.

Market microstructure: organisation and prices

In the standard efficient-markets view, little attention is paid to precisely how the market is organised. Thus the particular microstructure of a market is assumed not to have an impact on the evolution of prices. Yet a number of empirical observations bely this. Markets which are organised on an auction basis do not, in general, exhibit the same price behaviour as those for the same product organised on a posted-price basis, or with bilateral deals. Indeed considerable attention is paid by governments, for example, as to which mechanism to use when selling government bonds and privatising public enterprises.(4) This is precisely because the mechanism chosen will affect the prices obtained. So the prices obtained do not simply reflect the underlying 'fundamental' value of the assets or goods being sold, but also the choice of mechanism used to sell them.

In the FX market, there are dozens of market-makers all simultaneously announcing bid and offer prices at which they are prepared to trade in particular pairs of currencies. Even though these prices are posted on screens, at any point in time there will be a dispersion of prices available and transactions will often take place at the same moment in the same currencies at different prices. The explanations for this are clear. It is not possible for traders to keep track of all prices simultaneously, the prices announced are indicative, actual transactions will often take place within and not necessarily at the announced spread, and the prices offered by market-makers, announced by brokers and revealed by electronic broking systems are constantly shifting.

The way in which prices evolve is not the same as it would if there were a central auctioneer who periodically set prices to clear existing bids and offers.(5) Each market-maker's action will depend crucially on his time horizon. If he is 'hit' too often on one side, he will start to acquire a short or long position; this has two consequences. Firstly there is now a substantial element of risk involved and secondly he is acquiring a position which—for most market-makers—has to be closed by the end of the day. Both features will cause the market-maker to modify his prices and possibly his spread. If all trades were between market-makers, for every such movement there would be a countermovement. Even this would not necessarily eliminate any change, since the reaction of the two agents to the trades they had just effected might not be symmetric. In fact, of course, a number of the trades are motivated by outside orders.

The standard argument is that all this is irrelevant to the economist or actor such as a central bank, who is interested in observing and predicting the way in which prices move.

(1) For a simple example of the application of this tool to economics, see Kirman (*op. cit.*).
(2) This rather simplified explanation is based on a mathematical result of Bollobas.
(3) See, for example, Stanley, E A, Ashlock, D, and Tesfatsion, L, 'Iterated prisoner's dilemma with choice and refusal of partners' in *Artificial Life III*, Ed Langton, C G, Santa Fe Institute Studies in the Sciences of Complexity, proc. vol. XVII, Addison-Wesley, 1994.
(4) In general, the simple aim is to maximise revenue subject to certain distributional or 'fairness' constraints.
(5) In fact in some countries there are daily 'fixings' but these only involve a small volume of trade.

The process is analogous, it is argued, to that of a Walrasian auctioneer. If one currency is being sold then those with the highest bids will be hit first, their bids will fall and other participants with lower bids will be hit until the market is back in equilibrium.

But the essential point is that the dynamics of price movements are more complicated than this, for two simple reasons. Firstly if a transaction is observed to take place, it provides information to other market-makers, and this fact alone may cause them to modify their own bids and offers. Furthermore it also provides information to those who place orders with market-makers and their demand may be affected by this information. The identity of the person making the transaction (or on behalf of whom the transaction is being made) may also convey information.

So depending on the way in which a transaction is carried out, the information revealed will be different. Central banks are, of course, well aware of this and can choose, in consequence, whether to intervene openly or to do so in a less detectable way. The significance of this is that their impact on prices will be different depending on their behaviour and not on the fundamentals.

Lastly, as has been mentioned previously, the particular structure of the network of communications within the market and its connectivity will have a significant impact on the way information is transmitted and thus on the evolution of prices.

For all these reasons, the microstructure of the market cannot be ignored when trying to understand the nature of price formation in the FX market. Two aspects of the microstructure must be emphasised. Firstly it will determine how the heterogeneity of opinions or positions of agents are translated into prices and transactions, and secondly it will determine how—and how quickly—information will be transmitted.

Implications of interaction for price dynamics

Before considering the impact of the sort of phenomena discussed above, it is perhaps worth mentioning what is meant by price here. A lot of interest has been focused on high-frequency, or 'tick by tick', data. Each observation is the average of the bid and ask of the indicative quote in question. If the market is adjusting in one direction, the level of the prices will be misleading but the direction of the change will not.

If the market were operating as efficient-market theory says it should and each quote corresponded to the equilibrium price, then there should be no *auto-correlation* between price changes, ie no correlation between successive changes.

However if—as seems more likely—the market adjusts through a series of transactions in response to quotes, then one would expect positive auto-correlation, that is the change in one period is likely to be in the same direction as that in the next. Such a pattern would also be consistent with the sort of herd behaviour I have mentioned.

At the highest frequency, however, auto-correlation is actually negative.(1) This may be a result of looking in *too* great detail at the price series. In fact agents have heterogeneous horizons and may have different expectations. For example, some traders may not be allowed to hold open positions overnight while others will be taking positions on a much longer-term basis. When the market is unsettled, opinions may vary as to the direction of price changes even in the short run. As a result, successive trades will not necessarily be in the same direction. It has also been observed(2) that since different banks offer different spreads, prices may bounce back and forth, for example at the start of a movement from one expectations regime to another.

If we couple together two features of the sort of models discussed above—the different horizons of the agents and the emergence of speculative bubbles—we would have a situation in which volatility would be time-varying, but shorter-term and longer-term volatility would be linked. An initial switch would be transmitted to those with different horizons and would trigger off actions by a large number of market participants. In fact this is precisely what happens in the foreign exchange market: there is correlation between longer-term and shorter-term volatility, but not, for example, between successive observations of short-term volatility.(3)

There is by now a substantial literature on the detection of the time structure of the volatility of foreign exchange rates. Some success has been reported using ARCH, GARCH and more sophisticated versions of these tests.(4) The evidence is far from conclusive and indeed in the simple example discussed in the box on page 290 these tests failed. Nevertheless the total absence of structure in the time series of exchange rates—which the efficient-markets model would suggest—is not borne out by the evidence, and the effort being put into developing and testing trading rules based on the structure of the stochastic process generating exchange rate movements suggests that many market participants realise this.

Although trading rules based on the extrapolation (however sophisticated) of past prices have not been accorded much interest by academics with the exception of two centres,(5) the same is not true for major financial institutions who, according to a recent survey, almost all use 'technical analysis' in their forecasting. Indeed Brock *et al* have shown that even relatively simple 'technical' rules based on the

(1) See Guillaume, D M, Dacorogna, A M, Davé, R R, Müller, U A, Olsen, R B, and Pictet, O V, 'From the bird's eye to the microscope: a survey of new stylized facts of the intra-daily foreign exchange markets', *O&A Research Group Discussion Paper*, 1994.
(2) See Bollerslev, T, and Domowitz, I, 'Trading patterns and prices in the interbank foreign exchange market', *The Journal of Finance, 48*, pages 1,421–43, 1993.
(3) See Guillaume *et al (op. cit.*).
(4) See for example Chiandotto, B, and Gallo, G, Eds, *In Quest of the Philosopher's Stone*, Società Italiana di Statistica, Florence, 1994.
(5) The University of Wisconsin and the Santa Fe Institute.

differences between short-term and long-term moving averages do have some predictive power and therefore are profitable.(1)

Such evidence would seem to be consistent with the sort of stochastic model I have mentioned; the rules would be picking up the changes in prevailing expectations. All of this suggests that the data from the FX market share some of the features of a stochastic model of changing expectations regimes—periods of calm interspersed with periods of higher volatility, and periodic switches from one type of expectation to another.

Would the introduction of friction stabilise the market?

A number of economists have called for the imposition of some sort of globally applied tax on trading. Apart from the difficulty of implementing such a measure, two recent arguments related to the alternative models discussed above suggest that such a measure might be counterproductive.

The first is that since individuals convey information when they act, the introduction of a tax would mean that the signal given by an action might, when individuals act less frequently, be regarded by others as more important and more informative. This is simply because, since the cost of taking any action has been increased, the profit that an agent expects to make before taking an action must be higher. As a result, those observing will interpret an action as having more predictive value than previously.(2) So although the market may be quieter for longer, it will be more susceptible to large and sudden movements.

A rather different argument is that when the market is quiet and exchange rates are less volatile simple predictive rules are quite effective and will start to take over from more sophisticated and time-consuming predictive methods. But it may be the case that the simple rules are unstable, in the sense that, once perturbed, the exchange rate will not return quickly to some stable value. Suppose for the sake of argument that in equilibrium both methods would predict the same rate, then when sophisticated prediction methods are being used the rate will be robust to perturbations. However, precisely because of the market's stability the simple prediction method will reappear and will, sooner or later, lead to a period of high volatility. With more costly transactions the quiet period may be prolonged, but at the possible expense of experiencing greater volatility when a shock occurs.(3)

Conclusion

What I have suggested is a view of the FX market rather different in nature to that underlying the standard efficient-markets model.

Interaction between agents, in terms of the information that is passed and inferred, plays an important role in determining the dynamics of exchange rates. If the view is correct, the greater flow of information as communication technology develops will lead to more frequent changes in 'market opinion', and the increasing number of participants in the market will speed up the transmission of information between them. This in turn will increase the speed with which different price expectations can come to prevail. These tendencies have led—and will no doubt in future lead—to demands for the introduction of some sort of friction into the market, such as a tax on trading to diminish the volatility of exchange rate movements. In statistical terms, the market characteristics could be modified by the introduction of such frictions into the system. It is true that they would change the distribution of price changes and reduce average volatility. But this would be at the expense of a much larger probability of extreme events.(4)

A particularly interesting feature of the sort of model described is that it allows us to explain two important features of the FX market. Firstly, if there is a view on the part of some market participants that fundamentals are important, eventually this view will come to prevail—at least for a period. But the time at which this will occur is unpredictable. This makes the view that fundamentals are of little importance in the short term but matter in the long run perfectly consistent with the facts.

Secondly, the stochastic element in the communication between agents weakens any direct deterministic link between exchange rates and fundamentals in the short run, and explains why in two apparently similar situations there can be very different exchange rate movements.

The relationship between fundamentals and exchange rates is not well understood and seems to vary considerably over time. Furthermore, exchange rate movements depend on what market participants believe that relationship to be, and this also adds weight to the importance of the role of communication between agents in the market. The tendency of commentators to attribute a particular change to a particular piece of news about fundamentals is not justified by the behaviour of the system. Such explanations may give a plausible account of what happened, but do not enable us to make good predictions about future changes: one can make reasonable 'in-sample fits' but can not make good 'out-of-sample predictions'.

Finally, the view of the foreign exchange market as a complex interactive system with many heterogeneous agents, which undergoes periodic shifts in its state, seems not only to be able to explain some of the characteristics of the actual price dynamics of the market, but also to cohere with some of the stylised facts about the way in which the market actually works.

(1) See Brock, W A, Lakonishok, J, and Le Baron, B, 'Simple technical trading rules and the stochastic properties of stock returns', *Discussion Paper, University of Wisconsin*, Madison, 1991.
(2) This argument has been put in Caplin, A, and Leahy, J, 'Business as usual, market crashes and wisdom after the fact', *Harvard University Economics Department, Discussion Paper No 602*, 1992.
(3) This argument is developed in Brock, W A and Hommes, C, 'Rational paths to randomness', *mimeo*, University of Wisconsin, 1994.
(4) See Guillaume *et al* (*op. cit.*), who observed that this characterised the EMS record when bands were narrow.