What can Agents Know? The Feasibility of Advanced Cognition in Social and Economic Systems

Paul Ormerod[1]

Abstract The purpose of this paper is to suggest that in many social and economic contexts, self-awareness of agents is of little consequence. The complexity of many such systems is very high. No matter how advanced the cognitive abilities of agents in abstract intellectual terms, it is as if they operate with relatively low cognitive ability within the system. This can be the case even when the emergent properties of the system are known to individual agents. Examples are given from macro-economics, the evolution of firms, financial markets and games.

1 INTRODUCTION

The purpose of this short paper is to suggest that in many social and economic contexts, self-awareness of agents is of little consequence. The complexity of many such systems is very high. No matter how advanced the cognitive abilities of agents in abstract intellectual terms, it is as if they operate with relatively low cognitive ability within the system. This can be the case even when the emergent properties of the system are known to individual agents.

This is not to say that the arguments of Gilbert [1] and of Goldspink and Kay [2] are not valid in some empirical settings. But I want to suggest that the situations in which they are valid might be rather circumscribed. The Turing rule that the vast majority of real life problems have no algorithmic solution limits the empirical usefulness of the assumption that agents operate with advanced cognitive ability. In many real life situations, the dimension of the problem scales super-exponentially, even when considering situations in which interactions between agents and emergent properties of the system are absent. Keen [3], for example, provides an illustration of the dimensions involved in even quite simple consumer choice decisions.

I am not pretending to offer in any way a proof of my main point. But I want to provide empirical examples. I am suggesting that a key issue in the context of the theme of this session of the conference is prior consideration of the feasibility of agents exhibiting advanced cognition in any particular context, regardless of their inherent intellectual abilities.

2 THE MACRO-ECONOMY

The problems of assigning advanced cognition to agents, even with no emergence, can be illustrated first of all by reference to the macro-economy.

The economy is undoubtedly a complex system with emergent properties. The decisions of millions of individual consumers and firms interact to produce the movements we observe at the aggregate, macro-economic level in variables such as total output (GDP), inflation and unemployment.

Policy makers have a strong incentive to be in possession of forecasts which are systematically accurate over time. If they have little idea of where the economy is likely to be in a year's time, say, the ability to carry out a successful policy intervention is obviously limited.

Policy makers and their advisers are very much aware of the emergent phenomena of a macro-economy. Indeed, it is precisely this data, which emerges from the interactions of the agents at the micro level, which they are trying to influence.

There is a large literature which shows that even on one-year ahead predictions, the forecast errors are large relative to the size of the data. Forecasters seem to do well when the economy is pretty stable, but are quite unable to capture turning points. The evidence from the forecasting record suggests that we can do ever so slightly better than the naïve rule which says next year's growth is the same as this, but there is not much in it. The only economy which seems to deviate from this finding to any significant extent is that of the United States, yet even here the level of predictability is very low by scientific standards.

There is a whole branch of economic theory, real business cycle theory, which argues precisely that cycles arise from random exogenous shocks. As it happens, I think the cycle is mainly endogenous and not exogenous. But here is a serious part of mainstream theory which hypothesises that the short-term growth rate of the economy is unpredictable. Of course, predictions can always be carried out, but 'unpredictable' here means that it is not possible to make systematically accurate forecasts. So our ability to learn to make short-term macro forecasts is very severely constrained. We do not appear to be able to process successfully the available information.

I started off many years ago as a forecaster, and it soon became clear to me that it didn't work. I've been interested for a long time in why this should be the case, and a few years ago I finally worked it out. Physicists and mathematicians have developed a technique, random matrix theory, which enables us to decompose time series data into what we might usefully think of as signal and noise. Signal is the bit that contains genuine information, and noise is, well, noise. I used this technique and published an article in *Physica A* on the failure of macro-economic forecasting [4]. Essentially, the data is dominated by noise rather than signal.

The lack of predictability of the cycle does not mean that the agents taking the decisions which generate the cycle are acting at random. It is a question of dimensionality. The dimension of the problem leads to the data appearing 'as if', a favourite phrase of economists, it is close to random.

[1] Volterra Consulting, London and Institute of Advanced Study, University of Durham Email: pormerod@volterra.co.uk

Theoretical approaches which contain sophisticated autonomous agents following maximising rules, such as real business cycle theory, are unable to capture key emergent features of the system, such as the (weak) time and frequency domain regularities, the distribution of the duration and size of recessions, the overall distribution of GDP growth rates and the distribution of individual firm growth rates. In contrast,. models which have simple agent rules, but in which agents are connected on networks, can.[for example, 5]

A series which, at least for the UK, is not possible to distinguish from a random one is the *change* in the rate of inflation [6] The Monetary Policy Committee of the Bank of England have as their objective the maintenance of a particular rate of inflation. All the members are thoughtful and intelligent. They have large teams of highly qualified researchers employed to discover the emergent properties of the system, in other words how the macro-economy operates.

Yet, given that the change in inflation is indistinguishable from a random series, they do not know what the rate of inflation is going to be in, say, one year's time. Specifically, they do not know whether it will be higher or lower than it is at present. So their ability to control the rate of inflation, to meet the target, is very seriously constrained[2].

The macro-economy is an important example of our inability to learn in any meaningful sense because of limits to the cognitive ability of agents. It is not that the agents – policy makers – do not have advanced cognition in general, it is as 'as if' in specific contexts they do not.

3 FIRM EVOLUTION AND DEATH

Firms have a very strong incentive to survive. Management spends a large amount of both time and money in trying to understand the properties of the system which their particular firm inhabits. Yet even amongst the very largest firms in the world, success tends not to persists. Batty [7] notes that over the 1955-1994 period 'From the 100 firms making up the Fortune 500list in 1955, 39 (percent) remain in 1994, and if the changes in each year from 1955 are examined, this reveals a more dramatic micro-dynamics with firms entering and leaving the list with great rapidity'.

Even more dramatically, firms both large and small actually disappear all the time. In both America and Europe, for example, more than 10 per cent of firms become extinct every year. A stylised fact that is established in the literature relates to the relationship between the age of the firm and the probability of extinction, or survival, looking at it from the opposite perspective. The probability of extinction is high in the early years of a firm's life. It falls rapidly, and becomes essentially flat. The finding seems to apply to firms regardless of size.

The second relates to the concept of the size of an extinction event. In other words, during a given period of time, we identify what proportion of firms become extinct. This gives us the size of the event. The bigger the proportion, the bigger the size. We then relate the size to the frequency with which it is observed over the whole of the data.

[2] Even assuming economists understand the connection between changes in interest rates now (the policy instrument available to the Bank) and changes in the rate of inflation in the future, which they do not. For example, the impact of interest rate changes on the exchange rate, which is a determinant of inflation, is theoretically indeterminate

This approach is well known in the biological fossil record. The frequency with which any given size of extinction event is observed – in this case the proportion of all species becoming extinct during given time interval – is inversely proportional to the square of the size. In other words, large extinction events are much less frequent than small ones.

It turns out that a very similar relationship exists for the frequency-size relationship observed for the extinction of firms [8]. So we have two stylised facts to explain.

I developed a theoretical model to account for them [9]. The actions of any given firm can have either positive or negative net impacts on the fitness for survival of any other firm. In other words, firms might produce complimentary products, or they might compete. The impacts are expressed through a matrix, each cell of which specifies the impact of firm i on the fitness for survival of firm j. There is a rule which specifies what level of fitness a firm needs to survive, and how extinct firms are replaced.

The key feature of the model is that, in each step of the model, the net impact of the actions of firm i on the fitness for survival of firm j is updated at random. In other words, it is as if firms had no knowledge of the impact of their strategies. It is as if they are unable to learn from their past experience.

This model replicates very accurately the two key stylised facts on firm extinction. I go on to allow firms to have a certain amount of knowledge of the consequences of their actions. In other words, it is as if they are able to learn. I vary both the proportion of all firms in the model with this ability, and the amount which they are able to learn.

There are very considerable gains to being able to learn. The average life of agents able to learn, even small amounts, is considerably larger than agents who cannot learn. In the limit, of course, as the proportion of firms able to learn approaches unity, and as the amount of knowledge which they learn increases, the firms approach infinite lives and never die.

But firms are only able to know very small amounts before the model ceases to replicate the stylised facts on extinction. It appears that a firm cannot learn very much at all either about the impact of their strategies on the ability of other firms to survive, or about the impact of the strategies of other firms on its own ability to survive.

As with business cycle theory, the key to the ability of the theory to account for emergent behaviour is not the sophisticated rules of the agents, but the fact that they are connected on a network. This structure of the system is more important than specific agent rules. Knowledge of the emergent properties of the system would essentially be of no use to an individual firm.

4 FINANCIAL MARKETS

A similar result on structure is obtained by Farmer et. al. in their study of share prices on the London Stock Exchange [10] Agents place orders to buy and sell at random, subject to constraints imposed by current prices (which limit the size of order which can be placed by an individual). The model explains 96 per cent of the variance of the gap between the best buying and selling prices (the spread) using a sample of 11 stocks, and 76 per cent of the variance of the price diffusion rate, which determines the size and frequency of changes to prices.

The price setting mechanism is a continuous double auction, the process actually used on the Stock Exchange. The auction is called "double," because traders can submit orders to both buy and sell, and "continuous," because they can do so at any time. The results of the Farmer et. al. paper arise, for reasons which no-one yet understands, because of this particular price setting mechanism which is used.

In real life, of course, agents engage in what they believe to be clever strategic behaviour, yet a model which neglects this entirely performs impressively from a scientific perspective. So, again, the *structure* of what we might think of as the game, the price setting mechanism, is very important in determining outcomes, more so than accurate modelling of agent behaviour.

Further, knowledge of the emergent properties of the system, the subtle properties of price changes, is of no use to individual agents in the system. They cannot exploit knowledge of these properties in their individual decision rules.

5 GAMES

The more abstract world of game theory offers further illustrations both of the difficulties of assigning high cognition to agents, and of the importance of the overall structure of the system, the institutional rules. But, again, knowledge of the 'emergent' properties of the game would not be of much help to individuals.

A trivially easy game is noughts and crosses (British English) or tic-tac-toe (American English). There are multiple Nash equilibria in the game, so many that almost any move in any strategy leads to the optimal outcome, a draw. Knowledge of this emergent property might be useful, but it is easy to discover by experiment. Even young children rapidly learn this. But we do not need to consider games which are very much more complicated before things become less clear cut.

The rules of chess can be stated very readily, and a reasonably intelligent person can remember them quickly. But the computational power required to analyse most position in a game scales super-exponentially. In the vast majority of positions which can exist, we are completely unable to determine which is the best move.

Computers have essentially made progress in chess by pure number crunching. In other words, by the exhaustive examination of permutations of moves in a given situation. The world's leading player for two decades after the Second World War, Botvinnik, believed that computers would eventually beat humans if they could in some way be programmed to understand the nuances of positional play in chess, rather than by exhaustive examination of the possibilities. He led a Soviet research programme on this, but essentially got nowhere. The gains have not been made by computers exhibiting advanced cognition in understanding the subtleties of positional play in chess – the emergent properties, as it were - but by grinding out tactical calculations.

So chess is an example of a game which can be described very simply, but where the dimension of the problem of solving it scales in a super-exponential way. Even very powerful modern computers can only solve a limited proportion of all possible 6 piece combinations yet the game itself involves 32 pieces.

Chess of course is a recreational game without wider applications. A game which is often held to have many practical application is that of the Prisoner's Dilemma. The rules are very simple and are time invariant. Agents are assumed to have a great deal of information. In particular, in its simplest form, each agent is assumed to know the payoff values of his or her opponent. This is a pretty strong assumption to make when you think about it. Yet do we know the best strategy? Well, we do when we make the very specific assumption that the game will end in a fixed number of moves, and that both players know this. We might think of this as removing any uncertainty which the existence of the future might bring. In other words, it limits the dimension of the problem.

There is a vast literature on the Prisoner's Dilemma when it ends at random. But the optimal strategy remains unknown. Agents do not have the cognitive ability to compute it. The scientific community has invested a great deal of effort in trying to discover, to learn the best strategy, but still we do not know.

The Beauty Contest game is based on Keynes' famous comment on the stock market, which he likened to a newspaper game popular in the UK in the 1930s. Newspapers published picture of 100 women, and to win it was necessary to guess the 6 which the most participants would select as the most beautiful. As Keynes wrote [11] 'It is not a case of choosing those [faces] which, to the best of one's judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees'

In the modern version, a group of individuals is asked to select a number between (and including) 0 and 100. The winner is the person whose guess is nearest to a specified fraction of the average of all the guesses. If all players practice a 'high degree' of reasoning, the winning number will be close to the Nash equilibrium of 0.

But experiments designed to elicit the degree of reasoning which agents use all show that it is low, typically between 1 and 3. For example, Duffy and Nagel [12] set up a game in which the winner is the person(s) whose guess was closest to half of either the median, the mean, or the maximum number chosen by all players. They found that players used a very low order degree of reasoning when forming expectations on other players' expectations. If the winning number were announced and the game repeated with the same players, they found that the winning number did approach zero, but even after repeated plays of the game the degree of reasoning remained low.

Knowledge of the Nash equilibrium solution would, except possibly in the final stages of successive plays of the game by the same set of players, be of no use to an individual. Indeed, anyone using this as his or her rule would in general lose. The key to success is not knowledge of the emergent equilibrium, but guessing the degree of reasoning which is being used by other players.

Interestingly, Duffy and Nagel found that the structure of the game, in this case the statistic which determines the winning number, had an important influence on the speed with which players converged towards zero as the winning number.

The Ultimatum Game [13] offers an example of how the wider setting in which a game is played can be of crucial importance to the outcome. Two players interact once only. The first player proposes how to divide a sum of money between themselves, and the second player can either accept or reject this proposal. If the second player rejects, neither player receives anything. If the second player accepts, the money is split according to the proposal.

There is now a vast literature on this game, almost rivalling that of the Prisoner's Dilemma. The reason for this is well known. Theoretically, the first player should offer the smallest non-zero amount possible, since from the point of view of the second player anything is better than nothing[3] . Yet the evidence that in general people do not follow this strategy is very strong. The game has been played in many settings, including ones in which the amount of money on offer is large to the participants, and this result appears to hold.

I am not suggesting that we know for certain why this is the case. But broader concepts of equity and fairness are surely important, 'broader' here in the sense of existing quite independently of the game itself. How agents play the game does not influence these broader sets of values, but the values influence the outcome of the game.

So, clearly, there are 'higher levels' of emergence which are important for outcomes. With the Beauty Contest game, it is the policy maker (the experimenter) who selects the statistic which is to be used. And with the Ultimatum Game, the outcomes appear to be determined by the broader values of society.

6 BRIEF COMMENTS

- In many social and economic contexts, self-awareness of agents is of little consequence. The complexity of many such systems is very high. No matter how advanced the cognitive abilities of agents in abstract intellectual terms, it is as if they operate with relatively low cognitive ability within the system. This can be the case even when the emergent properties of the system are known to individual agents.
- The Turing rule that the vast majority of real life problems have no algorithmic solution limits the empirical usefulness of the assumption that agents operate with advanced cognitive ability.
- The more useful 'null model' in social science agent modelling is one close to zero intelligence. It is only when this fails that more advanced cognition of agents should be considered.

REFERENCES

[1] N. Gilbert, *Varieties of Emergence.* Paper presented at the Social Agents: Ecology, Exchange, and Evolution Conference Chicago, (2002).

[2] C. Goldspink and R. Kay, *Social Emergence: Distinguishing Reflexive and Non-reflexive Modes* Paper presented at the AAAI Fall Symposium Washington (2007).

[3] S.Keen, *Debunking Economics: The Naked Emperor of the Social Sciences,* Sydney: Pluto Press (2001).

[4] P Ormerod and C Mounfield, 'Random Matrix Theory and the Failure of Macro-economic Forecasting,', *Physica A*, (2000).

[5] H Li and T Gao, 'A GDP fluctuation model based on interacting firms'. *Physica A*, forthcoming (2008).

[6] P Ormerod, 'Why Membership of the Monetary Policy Committee is a Sinecure', *Manchester Statistical Society,* (2001).

[7] M Batty, 'Visualising Creative Destruction', *Centre for Advanced Spatial Analysis, University College London, Working Paper 112,* (2007)

[8] C Di Guilmi, M Gallegati and P Ormerod, 'Scaling invariant distributions of firms' exit in OECD countries' *Physica A*, **334**, 267-273 (2004).

[9] P Ormerod and B Rosewell, 'What Can Firms Know?', *Proc. NACCSOS Conference,* Pittsburgh, (2003).

[10] J.D. Farmer, P. Patelli and I.I. Zovko, 'The Predictive Power of Zero Intelligence in Financial Markets', *Proc. Nat. Academy of Sciences*, **102**, 2254-59, (2005).

[11] JM Keynes, The General Theory of Employment, Interest and Money, *chapter 12, Macmillan,* (1936).

[12] J Duffy and R Nagel, 'On the Robustness of Behaviour in Experimental "Beauty Contest" Games', *Economic Journal*, **107**, 1684-1700, (1997).

[13] W. Guth, R. Schmittberger, and B. Schwarze , 'An Experimental Analysis of Ultimatum Bargaining', *Journal of Economic Behavior and Organization*, 3, 367-88, (1982).

[3] There is an argument as to whether offering zero itself to the second player also constitutes a Nash equilibrium