

# A Brief Survey of Some Results on Mechanisms and Emergent Outcomes

Bruce Edmonds  
Centre for Policy Modelling  
Manchester Metropolitan University

**Abstract.** The mechanisms/abilities of agents compared to the emergent outcomes in three different scenarios from my past work is summarised: the El Farol Game; an Artificial Stock Market; and the Iterated Prisoner's Dilemma. Within each of these, the presence or absence of some different agent abilities was examined, the results being summarised here – some turning out to be necessary, some not. The ability in terms of the recognition of other agents, either by characteristics or by name is a recurring theme. Some of the difficulties in reaching a systematic understanding of these connections is briefly discussed.

## 1 INTRODUCTION

Put briefly the question of this workshop is: *what agent abilities are necessary for which emergent outcomes*. The importance of this can be seen in two ways: *firstly*, which of our abilities have evolved due to the advantage they give us in terms of social organisation (the *social intelligence hypothesis* [26]); and *secondly*, given the abilities we have, what kinds of social structure *could* emerge – i.e. what is the social *scope* of our species. To get a handle on this we need to understand some of the connections between agent abilities and what this means in terms of possible outcomes. Since this involves questions of scope and counterfactuals, a look at the evidence of one instance of a social species at one point in time is clearly not sufficient. Social simulation models provide one way of exploring this (examining a range of species and societies and gathering evidence is another).

Since this is a question I have thought about and investigated over at least 10 years, in a number of simulation models, I summarise some of these results here in brief. These models include both positive and negative results – that is, simulations in which certain abilities do seem to be necessary for certain outcomes and those in which they turn out not to be necessary. The paper is structured into the different games/environments in which the agents are placed. It must be noted that all these are very much simpler than those of a “complete” society and thus can be only taken as indicative of what might be the case there.

These are grouped together in terms of the kind of interaction the agents are put inside, with subsections for each mechanism investigated. I end with a short summary and a brief discussion of the difficulties of trying to work out what abilities cause which emergent outcomes and why.

## 2 EL FAROL BAR (OR *MINORITY*) GAME

In the El Farol Bar game [1] there are a fixed number of individuals (usually 100) have to decide each week whether they go to the El Farol Bar, or not. They want to go if it is not too crowded (more than 60) go but not otherwise. Each individual knows the past history of the attendance there and has to predict this week's attendance and decide accordingly. The result is that the attendance oscillates wildly (within a broad range of parameter settings) around the critical point. This was rechristened as the “minority game” in [3].

In [8] and [7] I extended this model to allow: sophisticated communication; a social network; open-ended learning; and strategies that allowed individuals to take their cue from specific others rather than just base their guesses on the total attendance numbers (as in the above version). This allows individuals to take their cue from identifiable others (I will go to the bar this week if agent-23 did last week, etc.).

### 2.1 Sophistication of Prediction Strategies

What Brian Arthur (and others who have investigated this model) found is that it does not matter which prediction strategies are available to the individuals [1] – the same outcomes seem to result as long as there is sufficient variety of strategies among the population of players]. Thus it does not matter whether the strategies for predicting patters in the attendance are sophisticate or simple, the structure of the game means that successful strategies are self-defeating after a short while. However this does seem to rely on a rough parity between cognitive power between the participants.

### 2.2 Open-endedness of learning ability

If the presence of certain kinds of open-ended learning, where strategies can be arbitrarily elaborated and individuals have a limited ability to search for new strategies this can lead to a spontaneous emergence of heterogeneity in terms of the kinds and styles of strategy developed by each participant. How different these strategy styles are depends on the range of strategy primitives available to an individual (the language that strategies belong to).

### 2.3 Ability to recognise individuals

The ability to explicitly recognise other individuals (i.e. by “name”) seemed to have several effects. I did allow the population as a whole to develop a better population of

strategies, in that it seemed able to facilitate the *discoordination* of decisions to attend between agents. At the same time it seemed to produce a greater variety of strategy fitnesses within agents [8]. Thus, as part of the expressiveness of the strategy space available to individual's this also had a similar effect to that described in section 2.2 above.

More importantly this greatly increased the amount of social embedding (in the sense of [19]) that seemed to occur – that is the extent to which different agent's strategies become interdependent on the results of others' strategies etc. Such embedding has positives and negatives from the point of view of the individual – it means that information is efficiently used within the “society” of individuals, so that newcomers/outsideers are usually at a disadvantage since they have not learnt to exploit this resource. However occasionally such outsiders might do much better than any because they can more easily discover totally new strategies and others will take a while to learn how to use its decisions [7].

## 2.4 Imitation

The ability to imitate a strategy is different from observing someone else's behaviour and using it as an input to one's own decision making strategy. It is an ability to copy the strategy behind the observed behaviour and thus reproduce that behaviour. In this scenario, imitation did not seem to facilitate either the general (dis)coordination of behaviour or the social embedding. Maybe this is due to the nature of the task which favours anti-coordination and a dissimilarity of strategies between individuals.

## 2.5 Communication

Similarly to imitation, described above, communication between agents in this model does not seem to have a crucial impact on the outcomes in this model. Since communicating in this model gives another individual additional information that it could use to its advantage (and the communicating agent's disadvantage) *if* it does contain a message that can be correlated with the sender's actions then it is so the advantage of the sender to “fool” the receiver, by constantly changing that message. Thus in some of the simulation runs agents were developing utterances of the form “*not not not not ... not not what I did last time*” as they appended “not”s in an effort to reverse the meaning of their communication. In summary the situation here is not conducive to the emergence of communication since it is to an agent's disadvantage to say anything that is meaningful in terms of its own actions.

## 3 ARTIFICIAL STOCK MARKET

In this situation there are a number of traders and a single market maker that sets prices and buys and sells the various stocks. Thus each trader can hold different mixes of cash the the various stocks available. Each trader has a small regular external income and aims to make money by investing or speculating. Each stock has a fundamental in terms of an dividend rate which follows a slow random walk – the trader gains interest from holding the stock in terms of this fundamental. There is a trading cost, so that it is not a good idea to be over-trading. The market maker is under obligation to buy stocks from traders at the current price or sell what stock it has, but it also sets the

prices. The prices are set according to a simple rule which compares supply and demand – if supply is greater than demand (more want to sell than buy) the price goes down, and in the opposite case upwards. This is basically an extension of [27]

This situation is similar to that of the El Farol bar game in that traders are competing to “out-guess” each other. It is a good idea to start buying a stock after and when others have been selling them and to sell after others have been buying into it. If one just follows others, one step behind, buying when others have bought and selling when others have sold last time you just lose money. If the market prices are fairly stable, it is best to buy and hold the stock with the highest dividend rate, but if dividends are low and prices volatile, it might be possible to make more money by speculating. Thus this is a fundamentally competitive situation, and chances for cooperation are minimal.

## 3.1 Type of Learning Mechanism

In [16] I investigated what happens if I swapped two different learning algorithms, but ones with the same inputs, outputs and rough expressive power. One was a neural network and one a genetic-programming algorithm. They both had a similar theoretical ability in terms of the functions between inputs and outputs that they could learn. However they did have different characteristics. The NN learns more smoothly, sampling intermediate functions and values from one strategy to another. The GP algorithm encodes rougher approximations and is able to change more sharply from one strategy to another in its memory.

Since a stock market model like this one acts to “amplify” and “react” to sudden changes running the model with each of these as learning mechanisms makes a lot of difference to the outcomes. The version with NN learning reacted smoothly and gradually, showing gradual adaption of prices and long-term price “waves”, whilst the GP version showed sudden, clustered volatility and characteristic speculative bubbles. Thus we can conclude (unlike others) that in some situations the exact cognitive model can make a crucial difference, even if they have the same “abilities”.

## 3.2 Anticipation

There are two different kinds of feedback that can be observed when one has taken an action or implemented a strategy – the success of that strategy (which explicitly or implicitly requires comparison with a goal and may be given a numerical measure such as a *fitness*) and the how well it results in the outcomes that were predicted (the *accuracy*). This corresponds to *instrumentalist* and *realist* strategies of learning, respectively: the former simply using raw feedback according to success to learn what to do; and the later which learns to make good predictions about the effects of actions (or causes) using feedback as to their accuracy and then uses these to reason about which action or strategy should be chosen assuming these that will best achieve their goal. It is generally supposed that the later kind, labelled anticipatory, will be more sophisticated but require more computational resources (although there are leaner versions of this, e.g. [28]). It is clear that sometimes such anticipation is needed, but that at other times it is not.

Thus I mixed traders which uses an instrumental style of trading (the fitness of just what worked recently and the content just being the immediate buy/sell decision) with one which evaluated GP expressions as to how accurately they predicted

prices (recently), the currently best being used to predict the next price changes and hence determine its buy/sell decisions. What I found in this model [10] is that anticipation did not help a trader earn more money in the long term, but have a different pattern of trading. The anticipatory traders would detect patterns in the short term and for a while do better than the instrumental traders, however the market would then change unexpectedly due to the instrumental traders discovering better strategies (due to the drain on them extracted by the anticipatory traders) and the anticipatory traders would then lose a lot of money, until they recognised their predictions had failed. Thus anticipation did not make them better traders in the long run.

### 3.3 Context Awareness

Learning in a context-aware manner means that there is a two stage process to learning: *firstly*, recognising the appropriate type of context; and *secondly*, learning assuming that context. This *does* depend upon the assumption that the situations being learned about *do* divide usefully into a series of recognisable contexts [6]. The folk theory of traders does suggest that they recognise and respond to market “moods”, so it seems plausible to suggest that context recognition might be useful for a trader in such a market.

[9] suggests an algorithm for simultaneously learning the appropriate contexts and the knowledge within these contexts. This was tried in a stock market model in [11]. Agents that used context-dependent learning was compared with those that used a similar style of learning but assumed there was only one context to learn within (effectively being context-free learning).

Preliminary results did seem to indicate that context-aware traders did better than context-free traders.

## 4 PRISONER’S DILLEMA GAME

Social dilemmas are when there is an outcome that is desirable for all (usually couched as when everyone *cooperates*), where individuals can do even better for themselves if they act selfishly (usually *defection*) but if everyone does so everyone does very badly. The situations characterised as the “Tragedy of the Commons” [24] are a classic example of this. Clearly if everyone just acts myopically in their self interest, then the worst outcome is inevitable.

In order for this to be avoided some additional mechanisms or social structure is necessary. Over the years quite a number of different ways of achieving this basic type of cooperation have been developed, including: kin selection, iterative play with memory of past interactions with each player, enforceable contracts, and group formation.

The version of this situation studied here is when each interaction individuals play a number of rounds of the prisoner’s dilemma game, and the individuals are propagated into the next iteration with a probability that depends upon their total score at playing the game. Some of the work in this area is summarised in [18].

### 4.1 Tags

The ability to recognise whether the person you are playing with will be cooperative is obviously advantageous to an individual in this game. It would also be useful to be able to recognise and remember every individual so that if you met them again you

would have information about their behaviour. The former of these is impossible in most situations and the later expensive in terms of cognitive resources and impractical in large populations.

However it turns out that simply being able to recognise some observable features of individuals and preferentially chose those similar to oneself can help establish cooperative groups, even where these features have no necessary link to behaviour. This was identified and called “tags” in [25]. In other words a high over-all level of cooperation is maintained even where this is not an evolutionary stable strategy. What seem to happen is the following:

1. A small group of co-operators with the same (or similar) tags happen to form;
2. Since individuals in this group preferentially play each other, they outperform those in the general population and are preferentially populated into the next iteration, so the group grows;
3. Eventually a defector appears in the group (due to an invading defector or a mutation) and does even better than the others in that group at their expense and hence multiples faster than the others;
4. The group becomes dominated by defectors and the fitness of the members rapidly decreases, leading to the death of its members since they now score less than those in the surrounding population.

Thus although cooperation in these groups only lasts in the medium term, there is continual arising of new groups to replace those that become “infected” by defectors.

### 4.2 Link Imitation

In a social network where a link indicates whether the nodes will play together or not, any “groups” are not indicated by similarity of tags, but rather by the structure of the links themselves. An effect similar to the one in 4.1 above can be observed where nodes compare their success with other nodes and their links and strategy copied if they are doing better than themselves. This is the SLAC algorithm of [21].

As above this can allow the network to structure itself, by breaking into groups or otherwise partially isolating defectors, and thus maintaining cooperative groups or regions of the network in a dynamic fashion.

## 5 OPINION DYNAMICS

Opinion dynamics refers to a wide range of models, where there are a set of individuals each of which has a particular opinion (from a given range) at any one time. Over the simulation these individuals interact in a pair wise fashion (either randomly chosen pairs or restricted by a given social network) in which the opinion of one node affects that of the other to make it more similar to its own and to have a greater impact the closer or more coherent are the two individuals’ opinions. The overriding dynamic in such models is the clustering of individuals into “groups” with similar opinions – the key questions being how many clusters form, how long does it take and how stable they are. The original and best know of these is the family of models starting with [4].

Departing from the spirit of the [5] model, is a new, simpler model. This has a fixed number of nodes and directed arcs

between these. Each node has a numeric value representing the strength of its belief on a certain issue as well as a value representing its susceptibility to influence by another. Each iteration of the simulation a random arc is selected and the opinion of the node at the origin of this is copied into that of the destination with a probability of its susceptibility. Thus eventually (without noise and give the network is connected) all nodes will have the same opinion and change will cease.

An elaboration of this model is designed to model the process of consensus formation among agents. Here the opinion of each agent can be thought of as a binary string, where each bit indicates the belief (or not) of each of a sequence of possible beliefs. There is a consistency function from possible bit strings to the interval  $[-1, 1]$  that indicates the consistency of this set of beliefs. The copy process involves the copying of a single bit from origin to destination according to a probability determined by the change in consistency that would result in the destination node. There is also a "drop" process done by single nodes which may drop a belief according to a probability related to the change in consistency that would result from this. Thus it is much more likely that the consistency of the beliefs in each node will increase over time, and also that nodes with similar beliefs will be clustered together. This model is described more in [14].

## 5.1 Topology

One surprising and quite general result is that, for a broad range of topologies (i.e. those that naturally occur in social networks – that are connected and with a relatively small diameter) the topology does not seem to make much of a difference to the consensus formation process. This has, in fact, been proved in a slightly abstracted general model.

## 5.2 Belief rejection

In the above model the effect of the presence of the ability of agents to drop beliefs is obviously important, otherwise eventually all nodes will end up with all possible beliefs (depending a little on there always being a positive probability of the copy process occurring regardless of the resulting change in consistency).

## 6 CONCLUDING DISCUSSION

Clearly whether the abilities of agents are critical for the emergence of certain emergent outcomes depends upon the situation that the agents are within as well as the way in which they interact. Thus a general pattern of necessary abilities is not clear. However, in the above results the ability to recognise social clues and identities is a recurring theme, and is clearly necessary in some instances.

There is a general tendency to always argue that more abilities or more sophisticated abilities are necessary, on the grounds that one can almost always conceive of a situation where it would be required. This is plausible in that we know of the wide range of abilities that humans possesses, so it is natural to suppose that these are necessary.

However this need not be the case. *Firstly*, it is probably that nature evolves multiple and overlapping mechanisms to achieve any particular purpose, since this makes for a more robust and certain result in an uncertain world, so just because humans possess an ability does not mean that it is necessary (though

might indicate it is helpful). *Secondly*, it assumes that human intelligence is a general ability, which requires a set of features to be obtained but is then sufficient. However, as I argue elsewhere [12], any intelligence will have different pros and cons, be better for some tasks than others – i.e. a *general* intelligence is impossible (unlike a general computation device, which clearly *is* possible [29]).

Another difficulty with this kind of exploration is that the advantage (or even necessity) of an ability may only become apparent in the presence of certain emergent social structures. Thus there may well be a processing of social bootstrapping that must occur if certain other social structures or other phenomena will be developed.

If one attempts to evolve whole societies from scratch and one is successful in this, one learns one set of abilities that are necessary, but one still does not know if there are other (e.g. simpler) ways of achieving the same result or whether *all* the abilities are indeed necessary.

The fact that some of these social structures do seem to be emergent makes the prediction of when they will result and when not almost impossible to tell without looking at case studies or doing extensive simulation exploration. Thus the problem of studying the connection between abilities and emergent outcomes is an extremely hard one.

One thing that *is* needed is the systematic recording of simulations and results, so that some of the conditions and connections can be started to be mapped. There *are* models that share some almost standard: kinds of interaction, topology of interaction, abilities of agents, temporal structure (e.g. evolutionary propagation), system goals, etc. or at least are various of these relative to a common ancestor (e.g. a PD game). A website that mapped versions of these things (e.g. different versions of a PD game, or topology) and then linked these together to a simulation that brought them together (with the various results) would start to put this jigsaw together.

## 7 REFERENCES

I apologise about the number of citations to my own work, however this seems unavoidable in a summary of my own results. Versions of almost all my papers are accessible from the URL: <http://bruce.edmonds.name/pubs.html>

[1] Arthur, B. (1994). Inductive Reasoning and Bounded Rationality. American Economic Association Papers, 84: 406-411.

[2] Axelrod, R. (1984) The Evolution of Cooperation, Basic Books, New York.

[3] Challet, D. and Y.-C. Zhang, Emergence of Cooperation and Organization in an Evolutionary Game. Physica A, 1997. 246: p. 407.

<http://xxx.lanl.gov/abs/adaporg/9708006>.

[4] Deffuant, G, Neau D, Amblard F, and Weisbuch G (2000) Mixing beliefs among interacting agents. Advances in Complex Systems 3:87-98.

[5] Deffuant, G., Amblard, F., Weisbuch, G. and Faure, T. (2002), How can extremism prevail? A study based on the relative agreement interaction model. Journal of Artificial Societies and Social Simulation, 5(4) <http://jasss.soc.surrey.ac.uk/5/4/1.html>

[6] Edmonds, B. (1999) *The Pragmatic Roots of Context*. CONTEXT'99, Trento, Italy, September 1999. Lecture Notes in Artificial Intelligence, 1688:119-132.

[7] Edmonds, B. (1999). Capturing Social Embeddedness: a Constructivist Approach. Adaptive Behavior, 7:323-348.

[8] Edmonds, B. (1999). Gossip, Sexual Recombination and the El Farol Bar: modelling the emergence of heterogeneity. Journal of

- Artificial Societies and Social Simulation, 2(4). (<http://www.soc.surrey.ac.uk/JASSS/2/3/2.html>)
- [9] Edmonds, B. (2001) *Learning Appropriate Contexts*. In: Akman, V. et. al (eds.) *Modelling and Using Context - CONTEXT 2001*, Dundee, July, 2001. Lecture Notes in Artificial Intelligence, 2116:143-155.
- [10] Edmonds, B. (2002) Exploring the Value of Prediction in an Artificial Stock Market. Workshop on Adaptive Behavior in Anticipatory Learning Systems 2002, Edinburgh, Scotland, August, 2002. (ABiALs 2002). Butz V. M., Sigaud, O. and Gérard, P. (eds.) *Anticipatory Behavior in Adaptive Learning Systems*. Springer, Lecture Notes in Artificial Intelligence, 2684:262-281.
- [11] Edmonds, B. (2002) Learning and Exploiting Context in Agents. Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), Bologna, Italy, July 2002. ACM Press, 1231-1238.
- [12] Edmonds, B. (2002) The Social Embedding of Intelligence - Towards producing a machine that could pass the Turing Test. CPM Report 02-95, MMU.
- [13] Edmonds, B. (2006) The Emergence of Symbiotic Groups Resulting From Skill-Differentiation and Tags. *Journal of Artificial Societies and Social Simulation*, 9(1). (<http://jasss.soc.surrey.ac.uk/9/1/10.html>).
- [14] Edmonds, B. (2008). Achieving Consensus Among Agents – an opinion-dynamics model. CPM Report CPM-08-185, MMU, Manchester, UK.
- [15] Edmonds, B. and Hales, D. (2003) Replication, Replication and Replication - Some Hard Lessons from Model Alignment. *Journal of Artificial Societies and Social Simulation* 6(4) (<http://jasss.soc.surrey.ac.uk/6/4/11.html>)
- [16] Edmonds, B. and Moss, S. (2001) The Importance of Representing Cognitive Processes in Multi-Agent Models, Invited paper at Artificial Neural Networks - ICANN'2001, Aug 21-25 2001, Vienna, Austria. Published in: Dorffner, G., Bischof, H. and Hornik, K. (eds.), *Lecture Notes in Computer Science*, 2130:759-766.
- [17] Edmonds, B. and Norling, E. (2007) Integrating Learning and Inference in Multi-Agent Systems Using Cognitive Context. In Antunes, L. and Takadama, K. (Eds.) *Multi-Agent-Based Simulation VII*, 4442:142-155.
- [18] Edmonds, B., Norling, E. and Hales, D. (2008, in press) Towards the Emergence of Social Structure. *Computational and Mathematical Organization Theory*.
- [19] Granovetter, M. (1985) Economic-Action and Social-Structure – The Problem of Embeddedness. *American Journal Of Sociology* 91:481-510.
- [20] Hales, D. (2001) Tag Based Co-operation in Artificial Societies. Ph.D. Thesis, Department of Computer Science, University of Essex, UK.
- [21] Hales, D. (2004) From Selfish Nodes to Cooperative Networks – Emergent Link-based Incentives in Peer-to-Peer Networks. In proceedings of The Fourth IEEE International Conference on Peer-to-Peer Computing (p2p2004), 25-27 August 2004, Zurich, Switzerland. IEEE Computer Society Press.
- [22] Hales, D. and Edmonds, B. (2003) Evolving Social Rationality for MAS using “Tags”, In Rosenschein, J. S., et al. (eds.) *Proceedings of the 2nd International Conference on Autonomous Agents and Multiagent Systems*, Melbourne, July 2003 (AAMAS03), ACM Press, 497-503.
- [23] Hales, D. and Edmonds, B. (2005) Applying a socially-inspired technique (tags) to improve cooperation in P2P Networks. *IEEE Transactions in Systems, Man and Cybernetics*, 35:385-395.
- [24] Hardin, G. (1968) *The Tragedy of the Commons*, Science, 162:1243-1248.
- [25] Holland, J. *The Effect of Labels (Tags) on Social Interactions*. SFI Working Paper 93-10-064. Santa Fe Institute, Santa Fe, NM. 1993.
- [26] Kummer, H., Daston, L., Gigerenzer, G. and Silk, J. (1997). The social intelligence hypothesis. In Weingart et. al (eds.), *Human by Nature: between biology and the social sciences*. Hillsdale, NJ: Lawrence Erlbaum Associates, 157-179.
- [27] Palmer, R., Arthur, W.B., Holland, J.H., LeBaron, B., Taylor, P. (1994) *Artificial economic life – a simple model of a stock market*. *Physica D* 75:264–274
- [28] Stolzmann, W. *Anticipatory Classifier Systems*. in *Genetic Programming*. 1998. University of Wisconsin, Madison, Wisconsin: Morgan Kaufmann.
- [29] Turing, A.. M. (1936) On computable numbers, with an application to the Entscheidungsproblem. *Proc. Lond. Math. Soc.* 42:230-65; 43:544-6.