Gossip, Sexual Recombination and the El Farol Bar: modelling the emergence of heterogeneity

Bruce Edmonds

Centre for Policy Modelling
Manchester Metropolitan University
http://www.cpm.mmu.ac.uk/~bruce

Abstract

Brian Arthur's 'El Farol Bar' model is extended so that the agents also learn and communicate. The learning and communication is implemented using an evolutionary process acting upon a population of mental models inside each agent. The evolutionary process is based on a Genetic Programming algorithm. Each gene is composed of two tree-structures: one to control its action and one to determine its communication.

A detailed case-study from the simulations show how the agents have differentiated so that by the end of the run they had taken on very different roles. Thus the introduction of a flexible learning process and an expressive internal representation has allowed the emergence of heterogeneity.

1. Introduction - the El Farol Bar

In 1994, Brian Arthur introduced the 'El Farol Bar ' problem as a paradigm of complex economic systems. In this model a population of agents have to decide whether to go to the bar each thursday night. All agents like to go to the bar unless it is too crowded (i.e. when more that 60% of the agents go). So in order to epitomise its own utility each agent has to try and predict what everybody else will do. The problem is set up so that any model of the problem that is shared by most of the agents is self-defeating. For if most agents predict that the bar will *not* be too crowded then they will all go and it *will* be too crowded, and *vice versa*.

Brian Arthur modelled this by randomly giving each agent a fixed menu of potentially suitable models to predict the number who will go given past data (e.g. the same as two weeks ago, the average of the last 3 week, or 90 minus the number who went last time). Each week each agent evaluates these models against the past data and chooses the one that was the best predictor on this data and then uses this to predict the number who will go this time. It will go if this prediction is less than 60 and not if it is more than 60.

As a result the number who go to the bar oscillates in an apparently random manner around the critical 60% mark, but this is not due to any single pattern of behaviour - different groups of agents swap their preferred model of the process all the time. Although each agent is applying a different model at any one time chosen from a different menu of models, with varying degrees of success, when viewed globally they seem pretty indistinguishable, in that they all regularly swap their preferred model and join with different sets of other agents in going or not. None takes up any particular strategy for any length of time or adopts any identifiably characteristic role. Viewed globally they seem to be acting stochastically and homogeneously, despite the fact that the whole system is completely deterministic and each agent is initialised with a different repetoire of models [12].

The purpose of this paper is to report on the difference in their behaviour when these agents are given a suitably powerful learning and communicative mechanisms and the whole system is allowed to co-evolve. It can thus be seen as an extension of the work in [2]

The approach taken is to endow each agent with a form of bounded rationality in the form of an evolutionary process among a population of competing mental models inside each agent. This is described in section 2. Then in section 3 I describe how this is applied to the El Farol Bar problem in a way which will allow social relations to emerge among the agents.

The results are considered in section 4 at the macroscopic level as well as in detail in the form of a case study of the interactions in the model at the last date of a particular (but representative) run. The heterogeneity which emerges in the discussed in section 5

2. Modelling Boundedly Rational Agents using the Evolution of Mental Models

The purpose of the model is to explore some of the possible ways that emer gent social structures in an agent community might effect the overall behaviour of that collection of agents.

Since we primarily have humans in mind in this exercise we wish for our software agents to at least capture some of the known qualitative characteristics of humans. In particular we are interested in agents:

- who incrementally develop models of their environment;
- who develop their models in a parallel manner so that different (and even contrary) models can be brought to bear in different circumstances;
- whose mechanisms of learning dominate those of inference;
- who are able to identify other agents individually and develop models specifically concerned with those agents;
- who have a flexible and expressive internal system of representation, so that they are as little constrained as possible in what model the can develop;
- who are able to develop any connection between their communication and their action that is appropriate;
- who are able to deal with received communication in what-ever way is best for it;

The agent modelling approach adopted broadly follows [5]. *Each* agent has a population of mental models, which broadly correspond to alternative models of its world. This population develops in a slow evolutionary manner based on what its past success at gaining utility might be.

Each notional week, the new population of models is produced using a genetic programming (GP) algorithm (Koza 1992). In GP each 'gene' is a tree structure, representing a program or other formal expression of arbitrary complexity. A population of such genes is evolved using a version of crossover that swaps randomly selected sub-trees and propagation. Selection of genes for crossover and propagation is done probilistically with a likelihood of selection in proportion to its fitness.

I have slightly modified this here by only using *some* tree crossover but with a high degree of propagation and also some new random genes introduced each time. Then the best model is selected and used to determine first its communicative action and subsequently whether to go to El Farol's or not. Thus the evolution of mental models is a rough representation of learning.

The cross-over operation is not very realistic but does as a first approximation. For a critique of cross-over and further discussion of the philosophy of agent design for the purposes of the credible modelling of human agents, see [5]. This model of learning fits into the wider framework of modelling economic learning as modelling described in [9].

3. Extending the El Farol Bar Model with Learning and Communication

In this extension of the model agents have a chance to communicate with other agents before making their decision whether to go to El Farol' s Bar. Each of the agents' models of their environment is composed of a *pair* of expressions: one to determine the action (whether to go or not) and a one second to determine their communication with other agents. The action can be dependent upon both the content and the source of communications received from other agents. Although the beliefs and goals of other named agents are not explicitly represented, they emerge implicitly in the effects of the agents' models.

The two parts of each model are expressions from a two-typed language specified (by the programmer) at the start [1]. A simple but real example is shown in figure 1 below. Translated this example means: that it will say that it will go to El Farol's if the trend predicted over observed number going over the last two weeks is greater than 5/3 (the total population was 5 in this example); but it will only *actually* go if it said it would go or if barGoer-3 said it will go.

```
talk:     [greaterThan [trendOverLast [2]]
                       [divide [5] [3]]]

action:   [OR [saidBy ['barGoer-3']]
              [ISaid]]
```

Figure 1: A simple example model

The agent gains utility by going to the El Farol Bar when it is not too crowded. Thus each agent is competitively developing its models of what the other agents are going to do.

The model was implemented in a declarative forward-chaining programming language called SDML [4, 10] which has been written specifically for agent-based modelling in the fields of business, management, organisation theory and economics. SDML is particularly suited to this model because is provides facilities for the easy programming of multi-layered object-orientated structures (so the populations of genes within a population of agents is easy) with several levels of time (in this case weeks and days)[2].

4. Results and A Case Study From the Model

In the output of the model the attendance at the bar fluctuates chaotically about the critical number of patrons (see the example plot in figure 2).

Figure 2: Number of people going to El Farol's each week in a typical run

The average fitness of the agents' models fluctuate wildly at the beginning but as the simulation progresses they settle down somewhat but the fluctuations do not damp down to zero. The deviance between different models of the same agent reduces only slightly (figure 3).

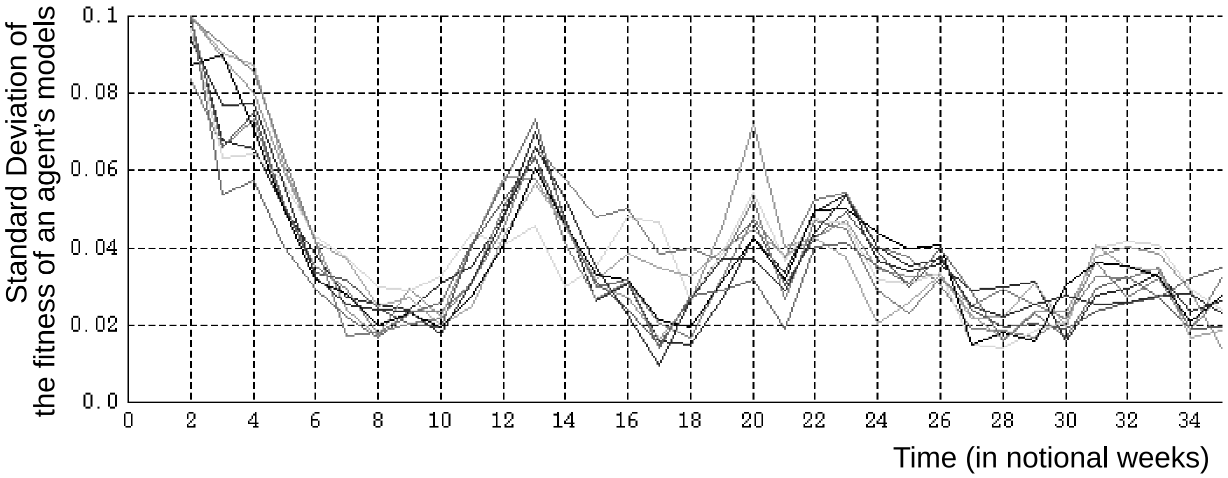

Figure 3: The change in variance (in standard deviations) of the Agents' population of models over time in (another) typical run

The graph of the utilities gained shows that different agents predominate at different times during the simulation with no one agent permanently dominating the others (figure 4).

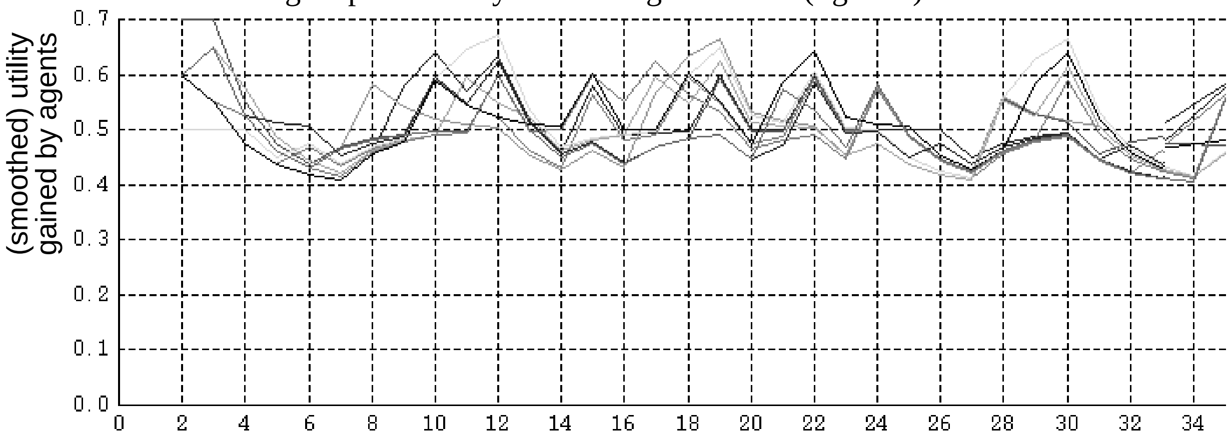

Figure 4: (smoothed) utility gained by agents over time

What is perhaps more revealing is the detail of what is going on, so I will exhibit here a case study of the agents at the end of a typical simulation.

Here I have chosen a 5-agent simulation at date 100. In this simulation the agents judge their internal models by the utility they would have resulted in over the past 5 time periods. Each agent had 40 mental models of average depth of 5 generated from the language of nodes and terminal specified in figure 2.

possible nodes for talk gene:
greaterThan lessThan previous times plus minus divide averageOverLast boundedByPopulation wentLag trendOverLast randomIntegerUpTo

possible terminals for talk gene:
wentLastTime maxPopulation IPredictedLastWeek 1 2 3 4 5

possible nodes for action gene:
AND OR NOT saidBy

possible terminals for action gene:
randomDecision ISaidYesterday IWentLastWeek T F 'barGoer-1' 'barGoer-2' 'barGoer-3' 'barGoer-4' 'barGoer-5' 1 2 3 4 5

Figure 5: Possible nodes and terminals of the tree-structured genes

The formal languages indicated in figure 2 allow for a great variety of possible models, including arithmetic projections, stochastic models, models based on an agents own past actions, or the actions of other agents, logical expressions and simple trend projections.

The utility that agents get is 0.4 if they go when it is two crowded, 0.5 if they stay at home and 0.6 if they go when it is not too crowded (where too crowded means greater than 60% of the total population).

The best (and hence active) genes of each agent are summarised above in figure 3. I have simplified each so as to indicate is *logical* effect only. The actual genes contain much logically redundant material which may put in an appearance in later populations due to the activity of cross-over in producing later models. Also it must be remembered that other alternative models may well be selected in subsequent weeks, so that the behaviour of each agent may ‘flip’ between different modes (represented by different models) depending on the context of the other agent’s recent behaviour.

talk-1: averageOverLast(numWentLast) > previous(trendOverLast(numWentLast))
action-1: wentLastTime

talk-2: trendOverLast(numWentLast) - 2 * numWentLag(2) > numWentLag(numWentLast)
action-2: NOT Isaid

talk-3: randomNumberUpTo(8) < 8/3
action-3: True

talk-4: averageOverLast(4) / averageOverLast(5) < numWentLag(15)
action-4: (Isaid AND randomDecision) OR (saidBy agent-2)

talk-5: trandOverLast(20) < numWentLag(2) - averageOverLast(numWentLast)
action-5: randomDecision OR (saidBy agent-4)

Figure 6: Simplified talk and action genes for the five agents at date 100

The effect of the genes is tricky to analyse even in its simplified form. For example agent-1 will tell its friends it will go to El Farol’s if the average attendance over a previous number of time periods (equal to the number who went last time) is greater than the predicted number indicated by the trend estimated over the same number of time periods but evaluated as from the previous week! However its rule for whether it goes is simpler - it goes if it went last week[1].

You can see that for only one agent what it says indicates what it does in a positive way (agent 4) and one which will do the exactly the opposite of what it says (agent 2). It may seem that agents 1 and 3 are both static but this is not so because figure 3 only shows the fittest genes for each agent at the moment in terms of the utility they would have gained in previous weeks. During the next week another gene may be selected as the best.

The interactions are summarised in figure 4, which shows the five agents as numbered circles. It has simple arrows to indicate a positive influence (i.e. if agent-2 says she is going this makes it more likely that agent-4 would go) and crossed arrows for negative influences (e.g. if agent-2 says she will go this makes it *less* likely she will go). The circles with an “R” represent a random input.

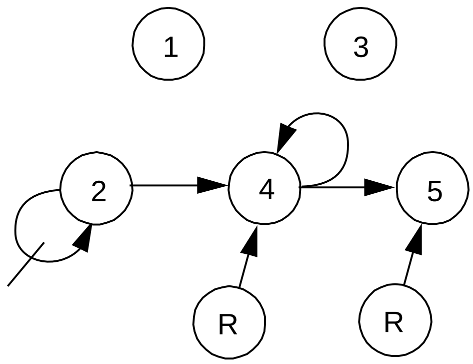

Figure 7: Talk to action causation

It is not obvious from the above, but agent-2 has developed its action gene so as to gradually increase the number of ‘NOT’s. By date 100 it had accumulated 9 such ‘NOT’s (so that it actually read NOT [NOT [... NOT [Isaid]...]]). In this way it appears that it has been able to ‘fool’ agent-4 by sometimes lying and sometimes not.

5. The emergence of heterogeneity

Unlike the Brian Arthur’s original El Farol model, this model shows the clear development of different roles[2].

By the end of the run described above agent-3 and agent-1 had developed a stand-alone repetoire of strategies which largely ignored what other agent said. Agent-3 had settled on what is called a mixed strategy in game theory, namely that it would go about two-thirds of the time in a randomly determined way, while agent-1 relied on largely deterministic forecasting strategies.

The other three agents had developed what might be called social strategies. Agent-2 seemed to have come to rely on ‘tricking’ agent-4 into going when it was not, which explains the gradual accumulation of ‘NOT’s in the example gene described above. Agent-4 has come to rely (at least somewhat) on what agent-2 says and likewise agent-5 uses what agent-4 says (although both mix this with other methods including a degree of randomness).

Thus although all agents were indistinguishable at the start of the run in terms of their resources and computational structure, they evolved not only different models but also very distinct strategies and roles.

The conclusion of the paper is that if one only allows global communicative mechanisms, and internal models of limited expressiveness then one might well be preventing the emergence of heterogeneity in your model. Or, to put it another way, endowing ones agents with the ability to make real social distinctions and (implicit or explicit) models of each other may allow the emergence of social situated behaviour that might be qualitatively different than a model without this capacity.

Such a conclusion marries well with other models which enable local and specific communication between its agents (e.g. [11]) and goes some way to addressing the criticisms in [6].

Acknowledgements

SDML is developed in VisualWorks 2.5.1. the Smalltalk-80 environment produced by ParcPlace-Digitalk. Free distribution of SDML for use in academic research is made possible by the sponsorship of ParcPlace-Digitalk (UK) Ltd. The research reported here was funded by the Economic and Social Research Council of the United Kingdom under contract number R000236179 and by the Faculty of Management and Business, Manchester Metropolitan University.

References

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

[2] Akiyama, E. and Kaneko, K. 1996. Evolution of Cooperation, Differentiation, Complexity, and Diversity in an Iterated Three-person Game, *Artificial Life*, 2, 293-304.

[3] Edmonds, B. 1997. Modelling Socially Intelligent Agents in Organisations. *AAAI Fall Symposium on Socially Intelligent Agents*, Boston, MA, November 1997. (Also available at http://www.cpm.mmu.ac.uk/cpmrep26.html)

[4] Edmonds, B., Moss, S. J. and Wallis, S. 1996. Logic, Reasoning and a Programming Language for Simulating Economic and Business Processes with Artificial Intelligent Agents. In Ein-Dor, P. (ed.): *Artificial Intelligence in Economics and Management*. Boston: Kluwer Academic, 221-230.

[5] Edmonds, B. and Moss, S. J. 1997. Modelling Bounded Rationality using Evolutionary Techniques. *AISB'97 workshop on Evolutionary Computation*, Lecture Notes in Computer Science, 1305, 31-42. (Also available at http://www.cpm.mmu.ac.uk/cpmrep10.html)

[6] Gaylard, H. A Cognitive Approach to Modelling Structural Change. *Workshop on Modelling Structural Change*, Manchester Metropolitan University, May, 1996. (Also available at http://www.cpm.mmu.ac.uk/cpmrep20.html)

[7] Koza, J. R. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press.

[8] Montana, D. J. 1995. Strongly Typed Genetic Programming, *Evolutionary Computation*, 3, 199-230.

[9] Moss, S.J. and Edmonds, B. forthcoming. Modelling Economic Learning as Modelling. *Systems and Cybernetics*. (Also available at http://www.cpm.mmu.ac.uk/cpmrep03.html)

[10] Moss, S. J., Gaylard, H., Wallis, S. and Edmonds, B. forthcoming. SDML: A Multi-Agent Language for Organizational Modelling. *Computational and Mathematical Organization Theory*. (Also available at http://www.cpm.mmu.ac.uk/cpmrep16.html)

[11] Moss, S. J. and Kuznetsova, O. 1996. Modelling the Process of Market Emergence. In Owsinski, J. W. and Nahorski, Z. (eds.), *Modelling and Analysing Economies in Transition*, MODEST, Warsaw. 125-138. (Also available at http://www.cpm.mmu.ac.uk/cpmrep12.html)

[12] Zambrano, E. 1997. The Revelation Principle of Bounded Rationality. Sante Fe working paper 97-06-060. (Title page at http://www.santafe.edu/sfi/publications/Abstracts/97-06-060abs.html)