Heterogeneity of Memory Decay and Collective Learning in the El Farol Bar Problem

Shikhar Kumar[1] and Cleotilde Gonzalez[2]

[1] School of Information, The University of Arizona, Tucson, AZ 85721 USA
shikhark@email.arizona.edu
[2] Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213 USA
coty@cmu.edu

Abstract. In this paper, we present a multi-agent simulation of the El Farol bar problem (EFBP), where the agents are cognitive models relying on the mechanisms of Instance-based learning theory (IBLT) and ACT-R. EFBP is a well-known example that illustrates how complex systems and economies evolve from induction. We investigate the effects of the different memory abilities of the agents, as well as the effects of the population's heterogeneity on their cognitive abilities. Our results of the emergent dynamics of the bar attendance suggest that the multi-agent simulation based on the IBL model is able to capture the dynamical properties of the EFBP reported in earlier work (Arthur, 1994). But most importantly, we show that the dynamics of bar attendance are sensitively dependent on the agents' cognitive abilities as described by the decay parameter in the IBL model. Furthermore, our results show that heterogeneity in their cognitive abilities may lead to the control of emergent dynamics.

Keywords: Instance-based learning, ACT-R, agent-based modeling, El Farol bar problem, complex systems, complexity economics

1 Introduction

Traditional economic theories of decision making are based on the "rationality" assumption: that agents are able to form accurate expectations and make optimal decisions. In real life situations, however, we often make decisions that are bounded by our cognitive abilities [15]. Behavioral economists and psychologists know that humans are rationally bounded and can only behave according to the contingencies of their capacities and the demands of the environment [28]. This observation from individual decision makers also expands to large and complex economies where multiple agents act together [4]. But these two perspectives, the psychology of the individual and the economy of a society, have rarely been brought together [5].

Network science and complexity economics studies focus on the interactions between actors and decision makers and their emergent social and economic phenomena, but they over-simplify the cognitive aspects of the individuals involved. For example, to explain the complex dynamics seen in large economic systems like financial markets, researchers have often relied on agent-based models, but rarely on cognitive models. On the other hand, cognitive modelers often focus on explaining individual behavior, relying on detailed cognitive models/architectures that formalize invariant cognitive representations and mechanisms [1, 22], but they rarely model the behavior of a group of individuals (see [17, 27] for some exceptions). This paper seeks to bridge this gap by investigating the impact of the heterogeneity of memory decay on learning and reduction of stochasticity in a large economy through cognitive models.

We use a well-known example of how complex, multi-player systems may evolve from inductive reasoning [3]: The El Farol bar problem (EFBP). This example presents a binary choice situation. There is a bar with a fixed capacity and a large number of people that independently choose from week to week whether to patron the bar or stay at home. Those who go to the bar when it is overcrowded will not have fun and those that stay home when the bar is not overcrowded will also not have fun. Thus the goal is to maximize the enjoyable experience by going to the bar without overcrowding it.

Binary-choice is the core of every choice problem. This has been studied from the cognitive perspective using a theory of decisions from experience, Instance-based Learning Theory (IBLT) [18]. IBLT posits that a choice is made by maximizing the experienced expected values of the two options (e.g., going or not going to the bar) from experience. IBLT is in general agreement with Skinner's basic propositions (1985) and reinforcement learning models [12, 29]. The instance-based learning (IBL) model for binary choice proposed by Gonzalez and colleagues [16, 24] is a computational representation of some of the proposed mechanisms of IBLT. This model builds on the learning and memory mechanisms of the ACT-R architecture [2] and it has been demonstrated as a robust representation of human behavior in many variations of binary choice tasks [15]

In the EFBP, the agents' ability to successfully predict bar attendance is dependent on their ability to generate predictions based on their past experience (historical attendance). For a homogeneous population of agents, that ability can be the same across all agents; i.e., every agent has the same memory capacity to predict bar attendance.For a heterogeneous population on the other hand, this capacity would differ across agent population. Sensitivity towards recent events is controlled through a memory decay mechanism in the IBL model. Changing the values of the decay mechanism can greatly influence the model's sensitivity towards recent events.

In what follows, we first present the background concepts involving the EFBP and the generic IBL model of repeated binary choice. Then we explain how we created a multi-agent simulation where we manipulated the memory decay in the IBL model to investigate its effects in a homogeneous and in a heterogeneous population. Our approach presents exciting possibilities for bridging the gap between the macro-view of large economies and the richness of cognitive models for individual agents.

2 El Farol Bar Problem

Every week, a set of N (N=100) people decide whether or not to go to a bar that offers live entertainment on Thursday nights. The bar has limited sitting capacity and may get too crowded if too many people (> 60) go at the same time. The evening is enjoyable if the bar is not over-crowded (< 60). Hence, the payoff of going to the bar is higher than staying at home. If the bar is too crowded though, the evening is less enjoyable and the payoff of staying at home is higher than going to the bar. There is no definite way to tell how many people will patron the bar in advance. Hence, a person will go to the bar if he expects to have a better time going to the bar than staying at home and will stay at home if he expects it to be better off than going to the bar.

One of the most important results of the EFBP is that agents are able to synchronize their actions such that the average attendance at the bar converges to the bar capacity, C, even though there is no communication between them [3]. This result has been reproduced in several other studies [6, 7, 13]. Although this is an interesting result, the stochasticity and oscillatory behavior demonstrated in past studies is costly and inefficient.

A less trivial question regarding the collective behavior of the agents is whether or not they can reduce stochastic fluctuations (whether or not they can learn). Researchers have investigated this question to determine what agent characteristics may lead to changes in the emergent dynamics [5, 9]. An example is the minority game [9], which involves heterogeneous inductive agents. Each of N players chooses between two options in every turn and those who are on the minority side win. The minority game has been extensively studied, with numerous variants analyzed [8, 10, 19–21, 23]. All results of the minority game can be directly applied to the EFBP [7].

Toy models such as the EFBP and the minority game have been used to explain market volatility based on the trading behavior of the agents [26]. Marsili and Challet showed that if the agents' learning rate is above a critical value, market dynamics become turbulent [26]. The reactivity of agents can be defined as their ability to react to the outcome of a game. Agents with high learning rates are more reactive than the ones with lower learning rates, leading to enhanced volatility in toy models. Similarly, with the use of reinforcement learning models in the EFBP [11, 30]; it was shown that agents can be divided into those who always attend the bar and those who always stay at home.

Here, we address two questions regarding the collective behavior of cognitive agents: How do homogeneous populations with different memory capacities reduce stochastic fluctuations in the EFBP? And how does the heterogeneity of memory capacities influence such learning? We address these questions by using the IBL model for repeated binary choice [15] in the context of the EFBP.

3 A generic IBL model of binary choice in the context of EFBP

The IBL model for repeated binary choice has accounted for human behavior in a large variety of tasks [15, 16, 25] . These model predictions are robust across different flavors of binary choice tasks, including probabilistic learning, dynamic probabilities, learning in multiple binary choice paradigms, and market entry games.

We used this generic model in the context of the EFBP. Each agent was defined with an IBL model consisting of four possible instances: [Go, +1], [Go, -1], [No-Go, +1], [No-Go, -1], representing the decision (Go, No-Go) of each agent and the possible outcome after each decision. For Go, -1 indicates the bar was overcrowded and +1 the bar was undercrowded; and for No-Go, -1 indicates the bar was undercrowded and +1 the bar was overcrowded.

The mechanisms of the IBL model for binary choice applied in the same way as they have been used in many other binary choice tasks [16, 17, 24]. These are summarized below for completeness.

At any time t, the model selects the option with the highest blended value (utility), which is computed as a weighted average of experienced outcomes in instances belonging to a given option. The blended value V of option j at time t is defined as:

$$V_j = \sum_{i=1}^{n} p_{ij} x_{ij} \tag{1}$$

Where x_{ij} is the outcome stored in an instance i for option j, and p_{ij} is the probability of retrieving the instance i from the memory. The n is the number of instances containing experienced outcomes on option j up to the last trial.

The probability of retrieving an instance i from memory is a function of its activation (A_i) relative to the activation of all other instances that correspond to option j. In each trial t, the retrieval probability is defined as:

$$P_i = \frac{e^{\frac{A_i}{\tau}}}{\sum_i \frac{A_i}{\tau}} \tag{2}$$

The activation of each instance in memory depends upon the Activation mechanism originally proposed in the ACT-R architecture [2]. In the IBL model, a simplified version that relies on recency and frequency of use of instances is used. In each trial t, activation of instance i is given as:

$$A_i = ln\left(\sum_{t_i \in \{1,..,t-1\}} (t - t_i)^{-d}\right) + \sigma ln\left(\frac{1 - Y_i}{Y_i}\right) \tag{3}$$

Where τ is random noise defined as $\sqrt{2}\sigma$, and σ is a free noise parameter accounting for the imprecision of recalling instances from memory for blending (adapted from ACT-R; [2]). Thus, a high *sigma* value implies more variability in retrieving instances from memory.

The decay d is a free parameter, and t_i (i.e., a timestamp) refers to the previous trial in which the outcome i was observed. Hence, the activation of an instance containing an outcome is dependent on the frequency of observing that outcome. The d parameter accounts for the rate of forgetting information: a higher value leads to a faster decay of an instance's activation in memory and signifies reliance on recently observed outcomes. The Y_i term is a random draw from a uniform distribution $U(0,1)$, and the $\sigma.ln((1\text{-}Y_i)/Y_i)$ in term represents the Gaussian noise for capturing the participant-to-participant variability in activation.

4 The multi-agent simulation of EFBP

An agent-based model was developed with a total of N (N=100) agents, where each agent was implemented with an IBL model as described above. The model was implemented in the multi-agent simulation environment, Netlogo [31]. Details about the Netlogo simulation environment for EFBP can be found elsewhere [14]. Each agent had one free parameter (d), which controlled its sensitivity towards recent outcome. In order to generate less noisy results, the noise parameter in the IBL model was fixed to a very low value of 0.001.

For each time step, every agent decides whether or not to go to the bar, according to the IBL's blending mechanism and what every agent has learned from its past experience.

We generated two types of simulations involving a (1) homogeneous population, and (2) a heterogeneous population. For each simulation in the homogenous case, all N agents had a fixed d parameter, producing a population of identical agents in terms of decay capabilities. In order to study the effect of d parameter on the emergent dynamics of the bar's attendance, we ran different simulations with different values of d (0.1, 1.0, 2.5). For low values of d when all agents relied on a longer-lasting memory, agents were expected to be able to account for more past experiences and make more informed decisions. We expected rapid learning and reduced oscillations around the optimal capacity level. With high values of d, agents would rely only on more recent experiences, producing more oscillations, less coordination, and less learning about the bar's capacity level.

Our investigation of the heterogeneous populations (involving a mix of agents with high and low memory values) was more exploratory. We generated different mixtures of agents with different values of d. Each simulation was run for a period of 200 cycles.

5 Dynamics of bar attendance with a homogeneous population

Our results show that bar attendance fluctuates around the capacity (C=60). But the amplitude and frequency of the fluctuation decreases with time, regardless of the value of d, and replicates learning patterns seen before in the EFBP.

More interestingly, a novel contribution is this simulation's sensitivity to the value of d. The stochasticity of bar attendance increased with an increase in the value of the d parameter. Figure 1 shows the average attendance level for each of the 200 weeks. Each panel shows the results of one value of d, from 0.1 (first panel) to 2.5 (last panel).

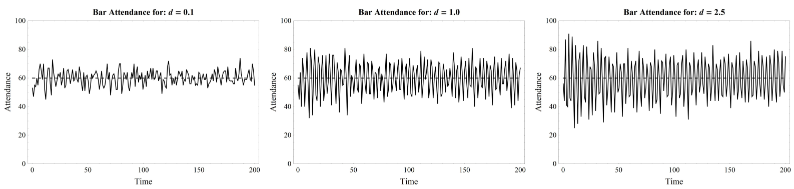

Fig. 1. Typical patterns of the dynamics of bar attendance with homogeneous population for different values of d.

To analyze the degree of stochasticity for different values of d, we computed the mean deviation of the bar attendance from the bar capacity (C = 60). The mean deviation of the bar attendance from the capacity for N trials was computed as follows:

$$\frac{1}{N}\sum |Attendance - C| \tag{4}$$

We ran each simulation 20 times and compared the mean estimates of deviation from bar capacity across different values of d. Table 1 shows an increase in mean deviation as d increased from 0.1 to 5.0. Our results suggest that d influence the dynamics of bar attendance and collective learning as shown in Figure 1.

Table 1. Mean and standard deviation of the mean deviation estimates across 20 simulation runs for different values of d

d	Mean deviation from C
0.10	M=4.56, SD=0.23
1.00	M=10.50, SD=0.29
2.50	M=12.97,SD=0.31
5.00	M=11.14, SD=0.40

6 Dynamics of bar attendance with heterogeneous population

Expectations of a homogeneous population in any social scenario are unrealistic since individuals vary greatly in their cognitive abilities. In order to start exploring the effect of heterogeneous population on the dynamics of bar attendance, different levels of diversity in the population were generated by creating two groups in the population varying on the value of d for each subgroup.

To produce a heterogeneous mix of population, we chose two of d values (0.1 and 5.0) and generated a population mix with different fractions of agents having a $d = 0.1$ and a $d = 5.0$. We can name the two kinds of agents in the population as A and B, with A being the agents with $d = 0.1$ and B being the agents with $d = 5.0$. We varied the number of A and B agents to produce a balanced ($A = 50$, $B = 50$) or imbalanced ($A = 25/75$, $B = 75/25$) mix of agents.

Our results shown in Figure 2 indicate that as the fraction of A agents ($d = 0.1$) decreased, the dynamics of bar attendance became more stochastic. This result is interesting in the sense that it shows that as the memory capacity of the majority of agents decreased, the emergent dynamics became more stochastic. This suggests that the right amount of variation in the cognitive abilities of agents can lead to better control (reduction in randomness) in the EFBP. We further analyzed this effect by computing mean deviation from the bar capacity across trials. Similar to the homogeneous case, our results show that learning is more prominent when the majority of agents have high memory capacity (high d).

As can be seen in Table 2, mean deviation is highest for the case with $A = 10$ and $B = 90$, and lowest for $A = 90$ and $B = 10$. This result further solidifies the claim that the d parameter significantly affects the emergent dynamics and learning in this problem. This result has far reaching implications in the field of complexity economics.

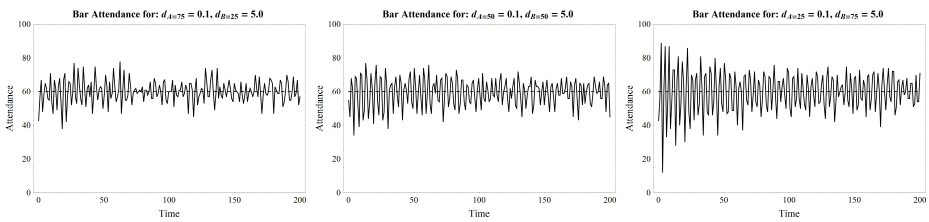

Fig. 2. Typical patterns of the dynamics of bar attendance with different mixes of agents with $d = 0.1$ and $d = 5.0$.

Table 2. Mean and standard deviation of mean deviation estimates across 20 simulation runs for different values of A and B

A, B	Mean deviation from C
90, 10	M=4.86, SD=0.26
75, 25	M=5.78, SD=0.26
50, 50	M=7.37,SD=0.28
25, 75	M=9.09, SD=0.39
10, 90	M=10.34, SD=0.31

7 Discussion

The major goal of the current paper was to simulate a well-known problem (EFBP) in complexity economics using cognitive agents employing the mechanisms of a well-known cognitive model (IBL), so as to bridge the gap between complexity economics and cognitive psychology.

We presented a simulation of how cognitive agents with heterogeneous or homogeneous memory abilities learn in EFBP. Cognitive agents were modeled with the IBL model of binary choice [15, 16, 24]. We demonstrate the dynamics' sensitivity to the agents' cognitive parameters (d). This result is important because it shows how the dynamics of the collective behavior of all agents is dependent on the cognitive abilities of individual agents. When all the agents are sensitive to only recent outcomes (high d parameter), this results in high stochasticity; but less stochasticity results when all agents can account for a larger set of experiences (low d parameter). Furthermore the heterogeneity of the population with more people with longer-lasting memory is also beneficial to reducing stochasticity.

These results have far reaching implications in the financial markets, where a diverse range of individuals participate, each with different levels of cognitive abilities. In the context of markets, our results imply that when the market is composed of impatient individuals (high d), a high stochasticity in the market dynamics will be seen because the individuals would predict the market based on their most recent experience, instead of accounting for all the previous outcomes. Such dynamics are often seen in financial markets and our results indicate that this may be due to individuals that predict the market based on the most recent information. On the other extreme, if the market is composed of experts or statistical agents (low d), the market would be less stochastic because the individuals would predict the market based on its complete past history. In the case of heterogeneous population, if the majority of individuals are reactive (high d), the dynamics lean towards stochasticity. On the other hand, if the majority of individuals are experts (low d), more stable dynamics are found. These results show how diversity in the agent population can affect market dynamics. We also observed that the right mix of heterogeneous agents can reduce the randomness in the emergent dynamics of the bar attendance as compared to the homogeneous population with the same d parameters, suggesting that diversity in cognitive abilities can lead to improved control.

In this paper, we only used two groups of agents for the heterogeneous according to the value of d. This case, however, can be extended to include different distributions of the d parameter across the agent population. We leave this investigation for future work.

We expect to extend this work to address various other issues that arise in social systems that can be answered by employing cognitive models/architectures in the complex systems framework, by using agent-based modeling.

References

1. Anderson, J.R., Lebiere, C.: The newell test for a theory of cognition. Behavioral and brain Sciences 26(05), 587–601 (2003)
2. Anderson, J., Lebiere, C.: The atomic components of thought lawrence erlbaum. Mathway, NJ (1998)
3. Arthur, W.B.: Inductive reasoning and bounded rationality. The American economic review 84(2), 406–411 (1994)
4. Arthur, W.B.: Complexity and the economy. science 284(5411), 107–109 (1999)
5. Baccan, D.D., Macedo, L.: Revisiting the el farol problem: A cognitive modeling approach. In: Multi-Agent-Based Simulation XIII, pp. 56–68. Springer (2012)
6. Casti, J.L.: Seeing the light at el farol: a look at the most important problem in complex systems theory. Complexity 1(5), 7–10 (1996)
7. Challet, D., Marsili, M., Ottino, G.: Shedding light on el farol. Physica A: Statistical Mechanics and Its Applications 332, 469–482 (2004)
8. Challet, D., Marsili, M., Zecchina, R.: Statistical mechanics of systems with heterogeneous agents: Minority games. Physical Review Letters 84(8), 1824 (2000)
9. Challet, D., Zhang, Y.C.: Emergence of cooperation and organization in an evolutionary game. arXiv preprint adap-org/9708006 (1997)
10. Chow, F., Chau, H.: Multiple choice minority game. Physica A: Statistical Mechanics and its Applications 319, 601–615 (2003)
11. Duffy, J., Hopkins, E.: Learning, information, and sorting in market entry games: theory and evidence. Games and Economic behavior 51(1), 31–62 (2005)
12. Erev, I., Roth, A.E.: Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. American economic review pp. 848–881 (1998)
13. Franke, R.: Reinforcement learning in the el farol model. Journal of Economic Behavior & Organization 51(3), 367–388 (2003)
14. Garofalo, M.: Modeling the?el farol bar problem?in netlogo. PRELIMINARY DRAFT, Dexia Bank Belgium (2006)
15. Gonzalez, C.: The boundaries of instance-based learning theory for explaining decisions from experience. Decision making: Neural and behavioural approaches 202, 7398 (2013)
16. Gonzalez, C., Dutt, V.: Instance-based learning: Integrating sampling and repeated decisions from experience. Psychological review 118(4), 523 (2011)
17. Gonzalez, C., Dutt, V., Lejarraga, T.: A loser can be a winner: Comparison of two instance-based learning models in a market entry competition. Games 2(1), 136–162 (2011)
18. Gonzalez, C., Lerch, J.F., Lebiere, C.: Instance-based learning in dynamic decision making. Cognitive Science 27(4), 591–635 (2003)

19. Hart, M., Jefferies, P., Johnson, N., Hui, P.: Generalized strategies in the minority game. Physical Review E 63(1), 017102 (2000)
20. Heimel, J., Coolen, A.: Generating functional analysis of the dynamics of the batch minority game with random external information. Physical Review E 63(5), 056121 (2001)
21. Johnson, N., Hui, P., Zheng, D., Tai, C.: Minority game with arbitrary cutoffs. Physica A: Statistical Mechanics and its Applications 269(2), 493–502 (1999)
22. Laird, J.E., Newell, A., Rosenbloom, P.S.: Soar: An architecture for general intelligence. Artificial intelligence 33(1), 1–64 (1987)
23. Lee, K., Hui, P., Johnson, N.: The minority game with different payoff functions: crowd–anticrowd theory. Physica A: Statistical Mechanics and its Applications 321(1), 309–317 (2003)
24. Lejarraga, T., Dutt, V., Gonzalez, C.: Instance-based learning: A general model of repeated binary choice. Journal of Behavioral Decision Making 25(2), 143–153 (2012)
25. Lejarraga, T., Hertwig, R., Gonzalez, C.: How choice ecology influences search in decisions from experience. Cognition 124(3), 334–342 (2012)
26. Marsili, M., Challet, D.: Trading behavior and excess volatility in toy markets. Advances in Complex Systems 4(01), 3–17 (2001)
27. Reitter, D., Lebiere, C.: Social cognition: Memory decay and adaptive information filtering for robust information maintenance. In: AAAI (2012)
28. Simon, H.A.: A behavioral model of rational choice. The quarterly journal of economics pp. 99–118 (1955)
29. Sutton, R.S., Barto, A.G.: Introduction to reinforcement learning, vol. 135. MIT Press Cambridge (1998)
30. Whitehead, D., et al.: The el farol bar problem revisited: Reinforcement learning in a potential game. ESE discussion papers 186 (2008)
31. Wilensky, U.: {NetLogo} (1999)