Revisiting the El Farol Problem: A Cognitive Modeling Approach

Davi D'Andréa Baccan and Luis Macedo

University of Coimbra,
Centre for Informatics and Systems of the University of Coimbra,
Department of Informatics Engineering, Polo II, 3030, Coimbra, Portugal
{baccan,macedo}@dei.uc.pt

Abstract. Decentralized market economies are complex systems that involve large numbers of heterogeneous participants. A good abstraction of this scenario is illustrated by the El Farol problem. In this problem, there is a bar with a fixed capacity and a given number of participants need to choose between either stay at home or go to the bar. However, if the attendance is above or equals the capacity of the bar, it becomes too crowded and the participants who attended did not have fun. In this paper we provide insight into the behaviour of the participants in those decentralized market economies scenarios by using a cognitive modelling approach in the El Farol problem. In three computer experiments we investigate, compare, and discuss a number of features of our agent-based model namely the relationship between beliefs and strategies, emotions of cognitive agents, as well as other aspects of market dynamics.

1 Introduction

Traditional economic theories tend to assume that agents are rational in the sense that they form expectations rationally and are able to make optimal decisions [10]. In other words, agents are considered to be able to correctly form probabilistic assessments, calculating which of the alternative courses of action maximize their expected utility (e.g. [26, 6]). On the other hand, observations regarding the behaviour of agents in real life scenarios, together with behavioral economics [14] findings constitute evidence that agents are not fully rational (e.g. [17, 11]). Agents do not always have enough time or the cognitive ability to process all the related information with accuracy, that is to say that they have bounded rationality.

A good example of a scenario in which agents have bounded rationality and need to make decisions essentially based on inductive reasoning and, therefore, cognitive agents might be used is illustrated by the El Farol problem [1]. In this problem, there is a bar with a fixed capacity and a certain number of people need to periodically and independently choose between two actions, namely go the bar or stay at home. However, if the number of people who go the bar is above or equals its capacity the bar becomes too crowded and those who attended did not have fun. In this problem agents generally make use of a strategy that provides

F. Giardini and F. Amblard (Eds.): MABS 2012, LNAI 7838, pp. 56–68, 2013.

them with a forecast for the next attendance that ultimately indicates whether the best action is to stay at home or go the bar. The only information available is the historical attendance and there is no communication between agents. In the context of those scenarios, unlike traditional economic theories, agents are somewhat forced to be heterogeneous in the sense they have to employ, for instance, different strategies or mechanisms for creating predictions about the next attendance. Nevertheless, as a strategic environment, the result heavily depends on the choice made by other agents.

The El Farol problem offers a rich set of possibilities for investigation as well as an interesting dynamics with respect to the behaviour of agents both in terms of micro and macro perspectives. It is important to stress that the interest in the El Farol problem is not new. On the contrary, a variety of different approaches have been proposed (e.g. [5, 19, 20]). For instance, Cross et al. [9] tried to incorporate minimal psychological factors to the El Farol, and observe whether they are able to reproduce some statistical regularities that are often found in real market data across different markets and periods of time, known as stylized facts [7]. Interestingly, despite its simplicity, their model was able to simulate a small number of fundamental phenomena.

In this paper we employ a cognitive modeling approach to observe the behaviour of agents in the context of the El Farol problem. It means that artificial agents will be empowered with mechanisms similar to or inspired in those used by humans. Therefore, the behaviour of artificial agents tends to be closer to the behaviour of humans in a similar scenario. In three computer experiments we investigate, compare, and discuss a number of features of our agent-based model namely the relationship between beliefs and strategies, emotions of cognitive agents, as well as other aspects of market dynamics.

The paper is organized as follows. In Section 2 we briefly present the cognitive emotion theories concepts related to our work. In Section 3 we present our agent-based model. In Section 4 we detail the experimental setup of our computer experiments and show our results, while in Section 5 we discuss our results. Finally, in Section 6 we conclude the paper and point out some future directions.

2 Cognitive Emotion Theories

Cognitive emotion theories (e.g. [15, 27]) rely on the assumption that emotions are mental states elicited as a result of the evaluation or appraisal of stimuli of all kinds (e.g. actions, events) and can be computed in terms of cognitions (beliefs) and motives (desires). Beliefs are mental states in which one holds a particular proposition to be true, whereas desires represent the motives or future states that one wants to accomplish.

The Belief-Desire Theory of Emotions (BDTE) is a cognitive emotion theory consisted of propositions, beliefs, desires, new beliefs, and two hard-wired comparator mechanisms, namely the Belief-Belief Comparator (BBC) and the Belief-Desire Comparator (BDC) [27]. The conceptual framework of the BDTE is the same as the belief-desire theory of action which inspired the BDI (belief-desire-intention) approach to artificial agents [3].

A proposition p is represented as a tuple $\langle S, B, D\rangle$ where S is the mental language expressing the proposition p, B and D are quantities representing, respectively, the agent's degree of belief and desire regarding proposition p. The strength of a belief in a proposition p at time t, defined as $b(p,t)$, is a value $\in [0.0, 1.0]$, where 1.0 denotes certainty that p, 0.5 maximal uncertainty, and 0.0 certainty that not p. Similarly, the strength of a desire about a proposition p at time t, defined as $d(p,t)$, might be a value, for instance, $\in [-100, +100]$. Positive values denote desire in favor of p, negative values denote desire against p, and 0 denotes indifference. A new belief is the belief or fact in a proposition that agents receive basically through its sensors (e.g. vision and hearing in the case of a human agent). It is defined as a tuple $\langle S, B, *\rangle$, where * denotes that the desire is irrelevant for new beliefs.

The Belief-Belief Comparator (BBC) compares each newly acquired belief to all pre-existing beliefs, looking for either a match or a mismatch. A match means that a pre-existing belief was confirmed by the newly acquired belief, whereas a mismatch means that a pre-existing belief was disconfirmed. As a result, BBC yields either a belief-confirmation signal or a belief-disconfirmation signal. Similarly, the Belief-Desire Comparator (BDC) compares each newly acquired belief to all pre-existing desires, looking for either a match or a mismatch. A match means that a desire was "fulfilled", whereas a mismatch means that a desire was "frustrated". As a result, BDC yields either a desire-fulfillment signal or a desire-frustration signal.

BDTE defines emotions as products or signals produced by the BBC and BDC. Additionally, defining whether agents experience non neutral emotions (e.g. happiness, unhappiness) depends on the desire of agents regarding p. An agent would be happy about p at time t, if she/he is currently certain that p happens, and has a desire in favor of p, i.e. $d(p, t) > 0$. On the other hand, surprise is elicited only based on beliefs. Formally, surprise can be defined as a peculiar state of mind, usually of brief duration, caused by unexpected events, or proximally the detection of a contradiction or conflict between newly acquired and pre-existing beliefs (e.g. [24, 22]). Therefore, an agent would experience surprise regarding p, if at time t_{-1} she/he had some belief that p will happen, but receives the new belief that actually non p happens. In Table 1 we summarize how the emotions we use in this work are computed from a qualitative perspective.

However, the BDTE does not clearly define how to compute surprise. Therefore, in the context of artificial surprise for artificial agents two models can be stressed namely the model proposed by Macedo and Cardoso [23, 21] and the model proposed by Lorini and Castelfranchi (e.g. [18]). Both models were mainly inspired by a cognitive-psychoevolutionary model of surprise proposed by Meyer et al. A detailed description of the similarities and differences of the models, written by Macedo, Cardoso, Reisenzein, Lorini, and Castelfranchi, can be found at [22]. The model proposed by Macedo and Cardoso offers a straightforward and easy way of computing artificial surprise that we consider to be the most appropriate for this work. Macedo and Cardoso claim that the relation between the subjective probability and the intensity of surprise about an

Table 1. Belief-desire theory of emotions, qualitative formulation (adapted from [27]). The notation is as follows: $Bel(p,t)$ stands for "believes p at time t", $Certain(p,t)$ stands for "firmly believes p at t", $Des(p,t)$ stands for "desires p at t", and $Des(\neg p,t)$ stands for "desires not-p at t, $\neg$ is aversive against p at t".

Emotion	**if**	**Belief at t**	**Desire at t**	**Belief at t_{-1}**
$happy(p,t)$		$Certain(p,t)$	$Des(p,t)$	
$unhappy(p,t)$		$Certain(p,t)$	$Des(\neg p,t)$	
$surprised(p,t)$		$Certain(p,t)$	(irrelevant)	$Bel(\neg p,t_{-1})$

event E_g from a set of mutually exclusive events $\{E_1,...,E_n\}$ can be described by $Surprise(E_g) = log_2(1 + P(E_h) - P(E_g))$ where E_h is the event with the highest subjective probability in the set.

For calculating non neutral emotions we rely on the BDTE [27], whereas for calculating surprise we rely on the artificial surprise model proposed by Macedo and Cardoso [23]. Similar to Table 1, in Table 2 we summarize how the emotions we use in this work are computed from a qualitative perspective.

Table 2. Belief-desire theory of emotions, quantitative formulation (adapted from [27]). The $b(p,t)$ represents the strength of belief in p at time t, with 1 denoting certainty that p, 0.5 maximal uncertainty, and 0 certainty that not-p. $d(p,t)$ represents the direction and strength of the desire for p at time t, with values > 0 denoting positive desire, 0 indifference, and values < 0 aversion against p. $Happiness(p,t)$, $Unhappiness(p,t)$ are the emotion intensities, ranging from 0 (absence of the emotion) to some maximum number, in this work 100. Regarding the calculation of surprise, $P(E_h)$ is the highest subjective probability attributed to an event from a set of mutually exclusive events, and $P(E_g)$ is the subjective probability of the event that actually happened [23].

Emotion Intensity in function of d and b	**for domain subset (else emotion intensity = 0)**
$Happiness(p,t)$	$b(p,t) = 1$ & $d(p,t) > 0$
$Unhappiness(p,t)$	$b(p,t) = 1$ & $d(p,t) < 0$
$Surprise(p,t)$	$log_2(1 + P(E_h) - P(E_g))$

3 Agent-Based Model

We distinguish two main aspects related with our agent model we used: (i) empowering agents with the BDTE; and, (ii) providing the agents with ways for dealing with information, for representing preferences, and for learning and evolving. These aspects are described in more detail as follows.

Firstly, our model was developed in the JABM (Java Agent Based Modeling) [25] that is a powerful Java API for developing agent-based simulation models using a discrete-event simulation framework. Therefore, we empowered JABM artificial agents with the Belief-Desire Theory of Emotions (BDTE). It included the implementation of the underlying mechanisms for dealing with propositions, beliefs, desires, new beliefs, with the BBC and BDC comparators, and also with the model proposed by Macedo and Cardoso. Additionally, inspired in the highly sophisticated, complex and dynamic human memory mechanism (see [8, 2] for an extensive review), we empowered agents with two different memory systems namely short-term memory and long-term memory as well as with the processes of encoding, storing (including forgetting), and retrieving memories. Therefore, agents are able to deal with previous knowledge with respect to whether the current strategy succeeded or failed and use such knowledge to calculate its current belief in the strategy.

Secondly, in designing the artificial agents we addressed the following three main design questions (e.g. [16].

The first question is how artificial agents deal with information. Consistent with the classical El Farol problem definition [1], agents receive only endogenous information that is the only information available is the historical attendance.

The second question includes all issues related to how to represent the preferences of artificial agents. The first issue is the definition of not only which strategies will be available for agents to forecast the next attendance but also the specification of all parameters related to those strategies. Agents have available six strategies commonly used in the context of the El Farol problem, namely noise trader strategy (NT), simple moving average strategy (SMA), exponential moving average strategy (EMA), opposite strategy (OPS), same strategy (SAS), and lagged strategy (LAS). The NT generates a uniformly distributed forecast between 0 and 100. The SMA generates a forecast using a simple moving average with a given window size, uniformly chosen between 2 and 100. The EMA generates a forecast using a moving average with a given window size in which recent values referring to the attendance gain more weight as opposed to old values. The OPS generates a forecast that is the opposite of the last week, whereas the SAS generates a forecast that is the same of the last week. Last but not least, the LAS generates a forecast that is exactly the same as a given past week, uniformly chosen between 1 and 5. The second issue is the definition of whether agents will use a fixed strategy or if they will be allowed to change from the current strategy to a new strategy based on some criteria. Agents have available two different scenarios. In the first, agents use a fixed strategy (henceforth referred to as FS) that means that once the strategy is defined, before the beginning of the experiment, the artificial agent will use this strategy until the end. In the second, agents can change from the current strategy to a new strategy (henceforth referred to as CS) based on their belief regarding whether they believe the strategy works or not, as we will explain in details later in this section.

The third design question refers to how artificial agents learn from mistakes and evolve. For each round, one of the strategies mentioned earlier is used by agents to predict whether the bar will be crowded or not and therefore indicates to them if the best action is either to stay at home or go the bar. Therefore, a strategy succeeds when it indicates the correct action or, in other words, if the strategy predicted that the bar will be crowded (or not) and it turned out to be crowded (or not) the action indicated by the strategy was the correct (wrong) one. Therefore, when a strategy succeeds (fails) an agent increases (decreases) its belief in the correctness of the strategy, based on a Bayesian process. According to the BDTE, the previous scenario can be modelled as follows. S that is the mental language expressing the proposition p is defined as "My strategy works", and the strength of a belief in the proposition p at time t, defined as *b(p,t)*, is a value $\in [0.0, 1.0]$, where 1.0 denotes certainty that the strategy really works, 0.5 denotes maximal uncertainty that is the agent does not know whether the strategy works or not, and 0.0 denotes certainty that the strategy does not work. The *b(p,t)* is calculated considering the experience of the agent in using the current strategy in a given number of last rounds, that is its memory size (henceforth referred to as *MS*). For example, suppose the unlikely scenario in which an agent is using a strategy that worked in the last 100 rounds. In this case, the *b(p,t)* would be close to 1.0, meaning that the agent "firmly believes" its strategy works.

Practically, on the one hand, a $b(p,t) > 0.5$ means that the agent has some degree of belief in the fact that its current strategy works and so it makes sense to a "rational" agent to maintain using the current strategy. On the other hand, a *b(p,t)* < 0.5 means that the agent has some degree of belief in the fact that its current strategy does not work and so it makes sense changing to a new strategy. Finally, in the context of our experiments, if *b(p,t)* $==$ 0.5 the agent maintain using the current strategy. In this context, for the *CS* scenario, we defined a belief threshold (henceforth referred to as BT) of 0.5 by which the agent must change its current strategy. Therefore, an agent only changes its current strategy if and only if *b(p,t)* < 0.5. Additionally, when an agent starts using a strategy its initial *b(p,t)* $= 0.5$.

It is also important to present some underlying concepts we employed, namely the concept of global belief in the strategy, global surprise, and cumulative happiness. Global belief in the strategy (henceforth referred to as *GBS*) is the sum of all *b(p,t)* and that is the "global belief that the strategy works". Global surprise (henceforth referred to as *GSu*) is the sum of all surprise "felt" by agents that is the "global surprise felt by agents with respect to whether their strategy works or not". Cumulative happiness (henceforth referred to as *CuH*) is the cumulative sum of all happiness "felt" by agents. An agent "feel" happiness when its strategy works or, in other words, when it indicates the correct action. We assume all agents "firmly desire" the strategy to work that is to say, according to the BDTE, that each agent has a *d(p,t) = +100.* In Table 3 we summarize the acronyms used throughout the paper, describe its meanings, and explain how we compute each of them.

To illustrate how *CuH*, *GBS*, and *GSu* work, consider the following example. Assume that there are two groups of agents namely G1 and G2. G1 consists of 59 agents using a fixed strategy that indicates the action go to the bar, whereas G2 consists of 41 agents using a fixed strategy that indicates the action stay at home. In this context, the attendance would be 59 and therefore the right action to take would be go to the bar. Therefore, for each round all agents of G1 would "feel" happiness, while all agents of G2 agents would "feel" unhappiness. Practically, in the first round, *CuH = 590*, in the second round, *CuH = 1180 (590+590)* and so forth. It is worth noting that this is the scenario that provides optimal results in terms of *CuH* and that such optimal values are used by us as references for calculating and plotting the results of *CuH* throughout the paper. Additionally, the $b(p, t)$ of all agents of G1 would be close to 1.0, while the $b(p, t)$ of all agents of G2 agents would be close to 0.0. *GBS* in this example tends to its maximum possible value that is 59. Regarding *GSu*, for all agents the $surprise(p, t)$ would be 0.0 and consequently *GSu* would be also 0.0.

Table 3. Summary of the main acronyms, meanings, and its respective forms of calculation

Acronym	Meaning	Calculation
CuH	Cumulative happiness	Cumulative sum of all happiness (i.e. its strategy worked)
GBS	Global belief in the strategy	Sum of the individual $b(p, t)$ of all agents
GSu	Global surprise	Sum of the individual $surprise(p, t)$ "felt" by all agents

4 Experiments and Results

We conducted three computer experiments to explore how the cognitive agents we modelled behave in the context of the El Farol problem. In Table 4 we summarize the features of the experiments. The experiments are defined in terms of the Strategies, and Fixed Strategy (FS) or Changing Strategy (CS) scenario. For all experiments memory size (MS) is 100, $BT = 0.5$, initial *b(p,t) = 0.5*, *d(p, t) = +100*, number of rounds is 2000 and, consistent with the seminal paper on the El Farol [1], the number of agents is 100 and the capacity is of 60. Due to the nature of the experiments, we run E1 five times so that we have five different configurations concerning the distribution of the strategies. Conversely, we run E2 and E3 just one time.

In the context of the CS scenario, the basic algoritm for changing a strategy is as follows. At the start of the simulation, each agent needs to select a strategy. During the simulation, if and only if $b(p, t) < BT$ an agent needs to change its current strategy by selecting one of the remaining strategies. When an agent has tested all the strategies, he/she changes to the strategy selected at the start of the simulation, creating a cycle. All strategies have the same probability of being selected.

Table 4. Experiments are defined as follows: **Strategy(ies)**: noise trader strategy (NT), simple moving average strategy (SMA), exponential moving average strategy (EMA), opposite strategy (OPS), same strategy (SAS), and lagged strategy (LAS); and **FS/CS**: fixed strategy (FS) or changing strategy (CS) scenario. For all experiments, number of rounds is 2000, belief threshold (BT) is 0.5, initial belief in strategy ($b(p,t)$) = 0.5, $d(p,\ t) = +100$, capacity of the bar is 60, the number of agents is 100, and memory size (MS) is 100. We run E1 five times, E2 and E3 one time.

Exp.	**Strategy(ies)**	**FS/CS**
1	NT, SMA, EMA, OPS, SAS, LAS	FS
2	NT, SMA, EMA, OPS, SAS, LAS	CS
3	NT	FS

We show and compare the results of our experiments both over time and in general. All outliers were removed and in some situations we smoothed and magnified the data in order to make the presentation clearer and the understanding easier. Such modifications are clearly indicated in graphics, otherwise the data is in its original scale.

We first show in Figure 1 the results regarding the attendance. In Figure 2 we show the results regarding CuH (cumulative happiness). GBS (global belief in strategy) and GSu (global surprise) "felt" by agents are shown in Figures 3 and 4, respectively.

We can see in Figure 1 that the attendances of E1 are quite similar, while the attendance of E2 (red) resembles the attendance of E3 (green). Additionally, E1 exhibits more volatile attendances than E2 and E3. Regarding CuH (cumulative happiness), we can see in Figure 2 that E1 has values that are considerably lower

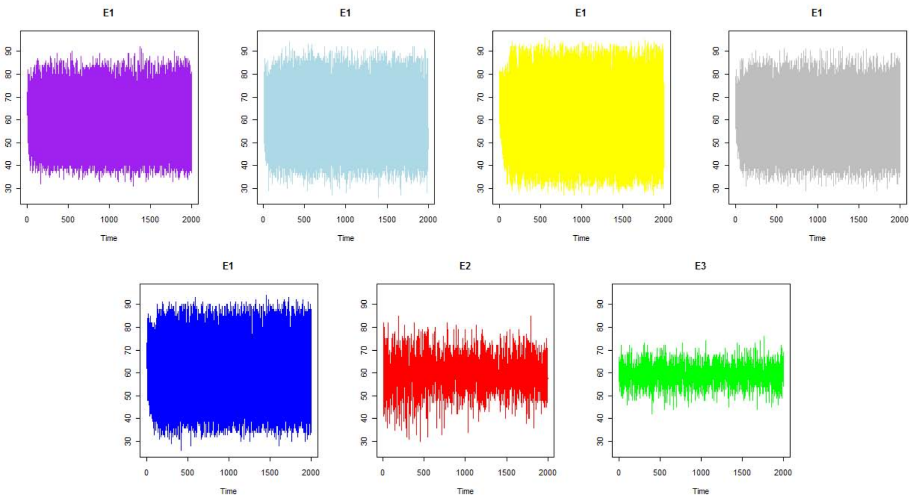

Fig. 1. Attendance of all experiments

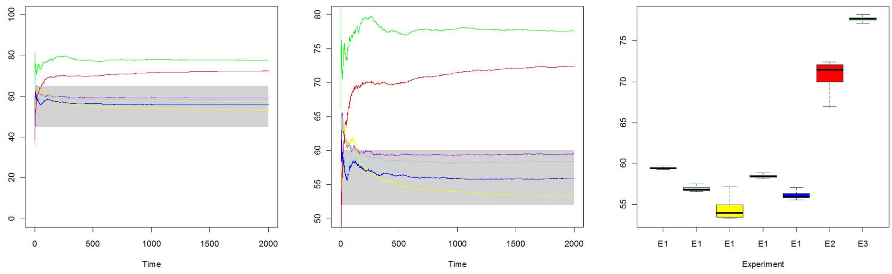

Fig. 2. *CuH* (cumulative happiness): original scale (left), zoomed in (center), boxplot (right). Gray area indicates runs of E1.

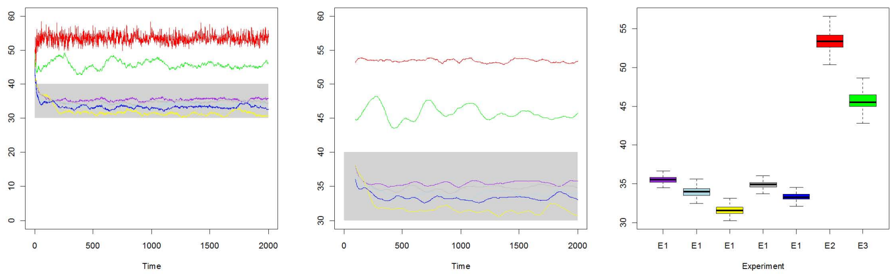

Fig. 3. *GBS* (global belief in strategy): zoomed in (left), SMA=100 and zoomed in (center), boxplot (right). Gray area indicates runs of E1.

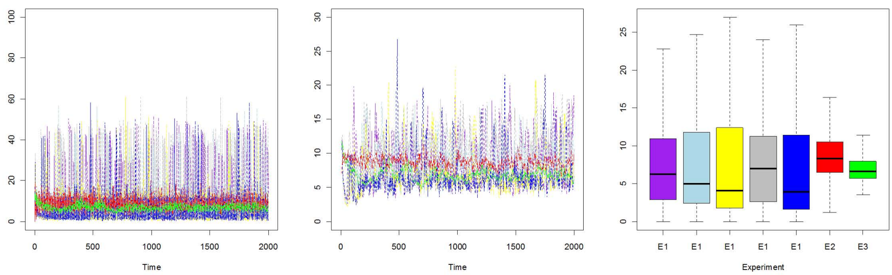

Fig. 4. *GSu* (global surprise): original scale (left), SMA=10 and zoomed in (center), boxplot (right)

when compared to E2 and E3. Similarly, we can see in Figure 3 that *GBS* values of E1 are lower than E2 values that are, in their turn, higher than E3 values. In terms of *GSu* (global surprise), we can see in Figure 4 that *GSu* oscillates in a narrow range between 0 and 30 for all experiments.

5 Discussion

In analyzing the results we were especially interested in observing the relationship between beliefs and strategies that is the belief in strategies (GBS), emotions of cognitive agents namely happiness (CuH) and surprise (GSu), as well as other aspects of market dynamics.

First of all, we need to bear in mind that there is an inherent relationship between memory size (MS), belief in strategy (*b(p,t)*) and consequently the global belief in strategy (GBS), and global surprise (GSu). The memory size (MS) is used to store the experience in using the strategy, specifically whether the action indicated by the strategy proved to be right or wrong. Practically, the MS refers to a given number of last rounds which the agent is able to "remember", as we mentioned earlier in Section 3. Such knowledge is latter retrieved so that the belief in strategy (*b(p,t)*) can be calculated and consequently its global sum that is the global belief in strategy (GBS). Additionally, according to the artificial surprise model proposed by Macedo and Cardoso, surprise varies from 0.0 to 1.0 and the closer the *b(p,t)* is to the point of maximal uncertainty, that means a belief in strategy (*b(p,t)*) equals 0.5, the lower is the individual surprise and consequently its global value (GSu).

From the results we can draw the following conclusions.

First, in all experiments, agents need to create a belief with respect to whether the current strategy works or not. Therefore there are only two mutually exclusive outcomes about the proposition. For instance, assume the outcome "strategy works" referred to as *O1* and the outcome "strategy does not work" referred to as *O2*, and an agent that has a *b(p,t)* = 0.7 in *O1* and therefore a complementary belief of 0.3 in *O2*. On the one hand, if the strategy succeeded, the surprise "felt" by the agent would be 0 ($Surprise(E_g) = log_2(1 + 0.70 - 0.70)$), according to the artificial surprise model proposed by Macedo and Cardoso. On the other hand, if the strategy failed, the surprise "felt" by the agent would be approximately 0.48 ($Surprise(E_g) = log_2(1 + 0.70 - 0.30)$). As expected and in accordance with the nature of the El Farol problem, considering the fact that it is a strategic environment in which there is no dominant strategy, we did not find high GBS values and consequently GSu values are relatively low.

Second, not surprinsingly and in accordance with the literature, we can observe that E3 (green) has "better" results in terms of attendance, *CuE*, and "good" results in terms of *GSu*. Attendance oscillates in a narrow range around the capacity of the bar. Higher CuE values means that agents in E3 agents are "feeling" higher happiness than those of E1 and E2. Last but not least, a lower GSu can be considered good because agents are not "feeling" surprised about the actions they choose to take. It means that agents are happier if all use a noise trader strategy (NTS) rather than trying to forecast the next attendance by using a "technical" strategy such as a simple moving average (SMA). However, in real life, it is difficult to imagine a scenario in which a cognitive agent may use a noise strategy to generate the next attendance, just "ignoring" the historical attendance, specifically when we consider the fact that humans intuitively try to discover patterns and predict things, even in random sequences [29, 28].

Third, as we briefly pointed out earlier, regarding *GBS*, the higher values are found in E1, while the lowest values are found in E1. Interestingly, it means that in E1 some agents have a degree of belief that is lower than 0.5. In such situation, we can expect a cognitive agent to change from its current strategy, that he/she believes that is not working ($b(p, t) < 0.5$), to a new strategy, instead of maintain using it. Therefore, despite the results of E3, the scenario illustrated by E2 as well as its results are more realistic ones.

Nevertheless, it is important to bear in mind that our results were obtained in a particular given setting, with specific configurations, for example, in terms of memory size, belief threhold, process for increasing and decreasing beliefs, and set of strategies available. Perhaps with other configurations the results might be different than ours. This means that, although we do not know which are the preferences of agents and, as expected, we need to make assumptions, relying on some assumptions, such as using the same memory size for all agents, might be a drawback of our approach, especially with respect to yielding results as realistic as possible.

6 Conclusion and Future Work

In this paper, consistent with findings from behavioral economics research and with real life observations, as well as departing from traditional economic theories, we take into account the fact that agents are actually heterogeneous, have bounded rationality, and are not fully rational. In this context, we provided non-trivial insights into the behaviour of agents in such scenarios by using a cognitive modelling approach in the El Farol problem. In three computer experiments we investigated, compared, and discussed a number of features of our agent-based model, specifically the relationship between beliefs and strategies, the emotions of happiness and surprise of cognitive agents, as well as other aspects of market dynamics.

We consider that the current work opens up a novel set of possibilities. We envision at least three future works. First, we could enhance the current work by incorporating more realistic findings with respect to how humans use both memory and past experience in decision-making. For instance, Kahneman and Tversky [13] pointed out that in revising their beliefs, people tend to overweight recent information and underweight prior information, while Griffin and Tversky [12] reported that people update their beliefs based on the "strength" and the "weight" of new evidence, where strength refers to aspects such as the salience and extremity, and weight refers to statistical informativeness such as the sample size. Second, we are interested in testing several processes for increasing and decreasing beliefs, as well as introducing new forms of forgetting (e.g. decay functions), in order to investigate if the results are similar to those we found in the current work. Third, we are interested in applying the ideas and concepts presented in this work to minority games [4] and ultimately to artificial financial markets as well as to compare our results to other cognitive approaches in the same context.

Acknowledgements. This work is supported by a Ph.D. scholarship from Portuguese Foundation for Science and Technology from Ministry of Science, Technology and Higher Education (FCT-MCTES), Portugal, reference SFRH/BD/60700/2009.

References

[1] Arthur, W.B.: Inductive reasoning and bounded rationality. The American Economic Review 84(2), 406–411 (1994)
[2] Baddeley, A., Eysenck, M., Anderson, M.C.: Memory, 1st edn. Psychology Press (February 2009)
[3] Bratman, M., Israel, D., Pollack, M.: Plans and resource-bounded practical reasoning. Computational Intelligence 4(3), 349–355 (1988)
[4] Challet, D., Zhang, Y.C.: Emergence of cooperation and organization in an evolutionary game. Physica A: Statistical and Theoretical Physics 246(3-4), 407–418 (1997)
[5] Challet, D., Marsili, M., Ottino, G.: Shedding light on el farol. Physica A: Statistical Mechanics and its Applications 332, 469–482 (2004)
[6] Chiodo, A., Guidolin, M., Owyang, M.T., Shimoji, M.: Subjective probabilities: psychological evidence and economic applications. Technical report, Federal Reserve Bank of St. Louis (2003)
[7] Cont, R.: Empirical properties of asset returns: stylized facts and statistical issues. Quantitative Finance 1(2), 223 (2001)
[8] Conway, M.A.: Episodic memories. Neuropsychologia 47(11), 2305–2313 (2009)
[9] Cross, R., Grinfeld, M., Lamba, H., Seaman, T.: A threshold model of investor psychology. Physica A: Statistical Mechanics and its Applications 354, 463–478 (2005)
[10] Fama, E.F.: Efficient capital markets: A review of theory and empirical work. The Journal of Finance 25(2), 383–417 (1970)
[11] Farmer, J.D., Foley, D.: The economy needs agent-based modelling. Nature 460(7256), 685–686 (2009)
[12] Griffin, D., Tversky, A.: The weighing of evidence and the determinants of confidence. Cognitive Psychology 24(3), 411–435 (1992)
[13] Kahneman, D., Tversky, A.: Intuitive prediction: Biases and corrective procedures. Management Science 12, 313–327 (1979)
[14] Kahneman, D., Tversky, A.: Prospect theory: An analysis of decision under risk. Econometrica 47(2), 263–291 (1979)
[15] Lazarus, R.S.: Emotion and Adaptation. Oxford University Press, New York (1991)
[16] LeBaron, B.: A builder's guide to agent-based financial markets. Quantitative Finance 1(2), 254 (2001)
[17] Lo, A.W.: The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. Journal of Portfolio Management, 15–29 (2004)
[18] Lorini, E., Castelfranchi, C.: The unexpected aspects of surprise. International Journal of Pattern Recognition and Artificial Intelligence, 817–835 (2006)
[19] Lus, H., AydIn, C.O., Keten, S., Ünsal, H.I., AtIlgan, A.R.: El farol revisited. Physica A: Statistical Mechanics and its Applications 346(3-4), 651–656 (2005)
[20] Lustosa, B.C., Cajueiro, D.O.: Constrained information minority game: How was the night at el farol? Physica A: Statistical Mechanics and its Applications 389(6), 1230–1238 (2010)

[21] Macedo, L.: The Exploration of Unknown Environments by Affective Agents. PhD thesis, University of Coimbra (2006)

[22] Macedo, L., Cardoso, A., Reisenzein, R., Lorini, E., Castelfranchi, C.: Artificial surprise. In: Handbook of Research on Synthetic Emotions and Sociable Robotics: New Applications in Affective Computing and Artificial Intelligence, pp. 267–291 (2009)

[23] Macedo, L., Reisenzein, R., Cardoso, A.: Modeling forms of surprise in artificial agents: empirical and theoretical study of surprise functions. In: 26th Annual Conference of the Cognitive Science Society, pp. 588–593 (2004)

[24] Meyer, W.-U., Reisenzein, R., Schutzwohl, A.: Toward a process analysis of emotions: The case of surprise. Motivation and Emotion 21, 251–274 (1997)

[25] Phelps, S.: Evolutionary Mechanism Design. PhD thesis, University of Liverpool (2007)

[26] Rabin, M.: A perspective on psychology and economics. European Economic Review 46, 657–685 (2002)

[27] Reisenzein, R.: Emotions as metarepresentational states of mind: Naturalizing the belief-desire theory of emotion. Cognitive Systems Research 10(1), 6–20 (2009)

[28] Taleb, N.N.: Fooled by Randomness: The Hidden Role of Chance in Life and in the Markets, 2nd updated edn. Random House (October 2008)

[29] Tversky, A., Kahneman, D.: Judgment under uncertainty: Heuristics and biases. Science 185(4157), 1124–1131 (1974)