Modeling one Human Decision Maker with a Multi-Agent System: the CODAGE Approach

Jean-Daniel Kant
Université Paris 6
Laboratoire d'Informatique (LIP6)
8 rue du Capitaine Scott
75015 Paris, France
Jean-Daniel.Kant@lip6.fr

Samuel Thiriot
Université Paris 6 and France Télécom R&D
Laboratoire d'Informatique (LIP6)
8 rue du Capitaine Scott
75015 Paris, France
Samuel.Thiriot@francetelecom.com

ABSTRACT

In this paper, we propose an agent-based model of human decision making: CODAGE (Cognitive Decision AGEnt). Classical Decision Theories have been widely used in multi-agent systems, but imply a too rational behavior when faced with real-world human data. Moreover, classical model usually exceeds human capabilities. Therefore, we derived our decision model from several cognitive psychological theories (e.g. Simon's decision theory, Montgomery's search of dominance structure, etc.) to take human bounded rationality into account. While most of existing cognitive agents use the BDI framework, we propose a new kind of architecture. In the CODAGE model, the decision maker is modeled by an entire multi-agent system, where each agent is in charge a particular sub-process of the whole decision. The architecture is intended to be as generic as possible. It could be viewed as an agent-based decision framework, in which different decision heuristics and biases could be implemented. We illustrate this approach with a simulation of a small experimental financial market, for which our model was able to replicate some human decision behaviors.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Human Factors

Keywords

Cognitive Modeling, Decision Making, Bounded Rationality

1. INTRODUCTION

Multi-agent systems (MAS) have been already used to model human decision making processes (e.g [19, 17]), and many successful applications have been derived from these works in various fields (e.g. economics [14], electronic commerce [7]). MAS seem appealing for human decision making because they capture the two levels of decision processes: the *individual level* (i.e. agent level), where each decision process could be modeled, from the reactive to the more cognitive one ; and the *collective level* (i.e. system level), where one has to model interactions among agents, communication between them, coordination, etc. For instance, in the field of Computational Economics, most works used rather reactive (zero-intelligence or zero-intelligence plus) agents and large-scale interactions in order to exhibit interesting market properties [14]. Other works would use more cognitive agents, mostly based on the BDI (Belief Desire Intentions) architecture [17].

When reactive (or zero-intelligence) agents are used, the emphasis is made on *emergence*: from a simple individual computation process, a global collective behavior emerges. This is consistent with the Artificial Life principles, stating that most of the complex phenomenon we observe in our world should be built from the bottom-up. This approach offers a variety of behaviors and phenomena at the macroscopic level. For instance, in the case of Financial Market Simulations, with few equations, reactive models could lead to equilibrium, oscillations, non equilibrium states, bubbles, that is to many phenomena observed in the real world at the macroscopic level. However, when one needs to explain these phenomena at the agent-miscroscopic level, the reactive approach suffers from its limitations to model the information processes the human agent may have used in the real world. In order to understand and explain complex social human functions (real markets, organization dynamics, coalition formation), more complex (i.e. more cognitive) models of human mind and actions are needed [4].

In the case of decision making, where can this more complex (more cognitive) theory be found ? In the last years, Classical Decision Theory and Game Theory have been used to design multi-agent systems. This includes classical (e.g. utility-based) theory models, extension of utilities, and Markovian Decision Processes (see [19] for a review), and yields several interesting applications. However, most of them are not compatible with the limitations of human capabilities, as stated by Simon with his concept of bounded rationality. Decision and Game Theories usually imply op-

timization processes, while agents are to explore a quasi-infinite (or exponential-growth) search space. And even if one relaxes this request of optimality, cognitive psychology has shown that humans violate most of the principles underlying Decision Theory Models. These works in psychology have been initiated by the seminal theories of Kahneman and Tversky, where the focus is made on the *heuristics* a decision maker used and the *biases* he/she had during a decision process [10]. Many recent models of decision making are based on these psychological facts, and focus on information processing [13].

This is the path we want to take. To do so, we follow a methodology based on what we called earlier *psychomimetism* [11]: (i) we need a cognitive model of decision based on psychological facts that have shown to be consistent with experimental observations, and (ii) we need to implement this cognitive model within an artificial system as faithfully as possible, so that the behavior of the artificial system gives some insights on the real-human cognitive process.

The remainder of this paper is organized as follows. In the next section, we present the cognitive model used in CODAGE, and then we give its implementation as a MAS architecture in section 3. To illustrate how CODAGE works, we present some results from an experimental financial market simulation in section 4. Finally, we conclude this paper with a discussion on our results, limitations and future work.

2. COGNITIVE MODEL

2.1 Simon's Descriptive Approach

There are three basic types of models: descriptive model, prescriptive and normative. Descriptive models tell us what people do or have done, prescriptive models suggest what people should and can do, and normative models tell us what people should ideally do (in theory). A typical example of normative theory is the Classical Decision Theory we mentioned above. In this paper we focus on *descriptive* theory, that is we try to model real-world human decision making. Broadly speaking, decision making could be defined as the cognitive process of selecting a course of action from among multiple alternatives. However, it should not be reduced to a simple selection of prototyped behavior or a case-based reasoning process, that is to a "déjà vu" kind of strategy. Decision Making is a *construction process*: the decision is built based on complex information processing mechanisms, and deals with various forms of knowledge representations (e.g. features, criteria, rules, etc.). Moreover, this process implies a commitment to the selected action.

To clarify what a decision making process is, we adopt the framework introduced by Simon in 1960 [22], as depicted in Figure 1. The decision process includes three steps. *Intelligence* means here the understanding of the decision situation, that it to identify what the decision is about (the need for a decision), the problem we have to solve and the context. Once this environment has been searched, the *design* phase begins, where a deep analysis of the problem domain and context leads to the identification of all possible options and alternatives. The final *choice* phase is the selection of the most appropriate course of action from the set of possible alternatives produced by the design phase.

The Choice phase may include an optimization process (choosing the best alternative) or not. According to the Satisficing model proposed by Simon[21], the decision maker


Figure 1: Decision Making as a three-step process (from Simon, 1960)

will chose the alternative that exceeds a particular criterion. Not surprisingly, the Satisficing model is fully consistent with Simon's idea of human Bounded Rationality[20]: humans have limited cognitive capacities, they usually are not able to optimize in most situations. A descriptive model of decision-making should take these limitations into account.

Finally, one should note that the decision process depicted in Figure 1 is not linear, as it includes feedback loops at each phase. For instance, the alternatives construction at the design phase might need further intelligence. Moreover, each phase incorporate a decision-making process itself, making Simon's decision model a quite complex system. From this framework, many models have been proposed (e.g. see [6, 13] for some reviews). Let us now present our decision model we implemented in CODAGE, which brings a synthesis of several cognitive theories.

2.2 Our decision model

The decision model implemented in CODAGE starts from the three phases of Simon's descriptive theory we described above:

Intelligence: due to limited perception abilities and thanks to its experience[1], the decision maker has the ability to quickly select the important informations. For instance, he/she might be more attentive to an information concerning himself or inducing some emotional reaction. The perception of the decision environment is based on a highly *selective attentional process*. The importance of attentional process in our model partly derives from the framing effects, as first observed by Kahneman and Tversky, where the formulation of a problem drastically influenced decision strategy [23]. This is because subjects focus primarily on a limited number of aspects (e.g. gain or loss in Kahneman and Tversky's experiments) and give the highest importance at these aspects to design alternatives and choose one of them. The concept of *salience* of information implements this principle (Cf. section 3.1.2). Moreover, we incorporate an *anchoring* mechanism, where some information is not continuously encoded : there is a set of relevant discrete value around which the decision maker anchors his/her judgment. This is consistent with the anchoring and adjustment heuristic observed in psychology of decision making [10].

Design: due to its bounded rationality, the decision maker cannot represent the whole world, he/she cannot produce all consequences and all possible actions. However, he/she needs an ability to anticipate short-term events, that is not

---

[1]As in many modeling works on decision making, we assume our decision maker to be what psychologist called an *experienced* subject. This subject is familiar with the decision task and has developed particular skills to solve similar problems. An experienced subject is not necessarily an expert, but never a novice.

only design a set of alternatives, but also produce some consequences of these alternatives. This will be implemented in CODAGE with a partial *tree of alternatives* where only some branches will be computed and explored (Cf. section 3.1.3).

Choice: in order to chose a satisficing alternative without requiring an optimization process, we adopt a heuristic proposed by Montgomery, the Search for Dominance Structure (SDS) [16]. Broadly speaking, the dominance structure is made of selected attributes that represent the good arguments by which one alternative dominates the others. The decision maker processes the information (the set of attributes) in order to find a dominant alternative. If one is found, it is chosen. Otherwise, he/she will look at a different set of attributes and alternatives. We present in section 3.2 our agent-based implementation of SDS.

Finally we slightly depart from the information flow in Simon's phase model. Instead of a sequential process combined with feedback loops between each phase, we view the three phases as a fully *concurrent and parallel* process. As soon as some information is perceived, the building of the alternatives tree starts and will be updated continuously as new information are coming. In parallel, dominance structures are built and as soon as one satisfying decision is found, the whole process is stopped to make the final decision. Hence this concurrency mechanism enables to take bounded capacities into account, by preventing the decision maker from exploring useless alternatives.

To sum up, our decision model incorporates the Simon's phase model augmented by five mechanisms:
  - selective attention (salience) ;
  - anchoring ;
  - partial tree of alternatives ;
  - search for dominance ;
  - concurrency.

3. THE CODAGE MACRO-AGENT

In this section, we propose an agent architecture to implement our cognitive decision model. Broadly speaking, the CODAGE agent is a *macro-agent* managed by a cognitive multi-agent architecture. Let us first give some intuitions that support this idea. In section 2.2, we described our decision model with five mechanisms we incorporate into Simon's model. These mechanisms could be viewed as modules of decision subprocesses. As suggested by the concurrency mechanism, we do not believe in a sequential organization of theses modules. On the contrary, we claim that most of decision subprocesses are autonomous entities that interact at different phases of the decision process. Let us take the example of anchoring. In trading, round values (e.g. 50) are easier to be memorized and might be favored as decision thresholds. This number anchoring effect will not only bias the perception process, but also the alternative choice and probably the dominance search (round numbers will be favored to be the attributes of the dominance structure). Given our concurrency mechanism, we adopt the "Minsky's Society of Mind" point of view [15] and propose to model all these subprocess modules as autonomous agents in interactions.

The MAS comprises a set of specialized agents, we call *micro-agents* in order to distinguish them from the CODAGE macro-agent they belong to, and a tree of alternatives implemented as a blackboard system to facilitate information sharing, as depicted in Figure 3 below.

3.1 Knowledge representation

3.1.1 Facts

In CODAGE, we represent facts with a set of attributes, values, and predicates. For instance in a trading game, `$capital[capital_euros]=2501.2` means that the attribute "capital" has a value of 2501.2 and this numerical value is typed as "capital_euros"; `buy_proposition(alice,3,14.5)` encodes the fact that Alice proposed to buy 3 stocks at 14.5 euros each. We add two important mechanisms to encode the information processing prescribed by our cognitive model : salience and tree of alternatives.

3.1.2 Salience

The *salience*[2] of a fact represents its importance within the selective attention process. Each micro-agent ma of global agent pool $\mathcal{P}$ can vote to set the salience of a knowledge K within an alternative C (context, i.e. a possible state of the world). We denotes $v_{ma,K,C} \in [0,1]$ the resulting value of such a vote. If the value is strictly positive, K is added to C with the corresponding salience value $v_{ma,K,C}$ if K is new to C ; if K is already instantiated in C, then its value is simply updated in the equation 1 that gives the final value $S_{K,C}$ of the salience of a given knowledge K within the context of an alternative C as the mean of the micro-agents ' votes:

$$S_{K,C} = \frac{\sum_{ma \in \mathcal{P}} v_{ma,K,C}}{card(\mathcal{P})} \qquad (1)$$

Neurobiology supposes that a salient fact is processed more quickly than an non-salient one [8]. In our model, knowledge-source agents will focus their attention on salient facts. This is implemented with two kinds of delays: an *event propagation delay* $d_{K,C}$, which causes agents to be warned later for non-salient facts, and a *reaction delay* $d_{R,C}$ for each rule R activable in a knowledge-source agent.

The propagation delay is 0 if the salience is 1, and rises to a maximum level (γ) if the salience is 0. We use the following function:

$$d_{K,C} = -\gamma \cdot \left(\frac{S_{K,C} - 1}{S_{K,C} + 1} \right) \qquad (2)$$

Figure 2 shows how this delay evolves in function with salience.

Let $A_{R,C}$ be the activation of a rule R in alternative context C. It equals the mean of premise 's saliences:

$$A_{R,C} = \frac{\sum_{\pi \in Premises(R)} s_{\pi,C}}{card(Premises(R))} \qquad (3)$$

The agent reaction delay is inspired by the ACT-R theory [1, p.1043]:

$$d_{R,C} = I + Fe^{-A_{R,C}} \qquad (4)$$

where $I = 597ms$ and $F = 890ms$.

Finally, the total information propagation delay of knowledge K in context C is given by:

$$td_{K,C} = d_{K,C} + d_{R,C} = -\gamma \cdot \left(\frac{S_{K,C} - 1}{S_{K,C} + 1} \right) + I + Fe^{-A_{R,C}} \quad (5)$$

---

[2]Several psychological studies support the concept of salience. Due to lack of space, we suggest this review of salience effects [9]

Figure 2: Propagation delay as a function of salience ($\gamma = 10$)

Our mechanism of propagation delay has two major benefits:

- the system is more robust to information permutation: even if a low-salient information is added prior to a high-salient one, the latter will be considered first.
- it enables the partial exploration of the tree of alternatives, since the micro-agents will act based upon the most salient facts. Alternatives based on low-salient facts will not lead to further consequences exploration.

3.1.3 Tree of Alternatives (TA)

In CODAGE, the knowledge that the decision-maker has about the world is encoded into a decision tree[3], as the one depicted in Figure 4.

Each node is an alternative that represents a possible state of the world (past, current or future). TA is a decision tree, as in decision theory, but it will be only partially built and explored to be consistent with bounded rationality. TA works at a symbolic level: each alternative represents an instantiation context in which each micro-agent may add a fact and/or an action into the tree: this is a way to share information between micro-agents. Each fact in the tree has a salience that measure its degree of importance.

Arcs between alternatives nodes represent *transitions* in time, that what produce the transition from one alternative (parent) to another one (child). We implemented two types of transitions that triggers the change to a new state of world:

- *action* transition: a possible action, performed by the macro-agent (myself)
- *fact* transition: the probability that some attribute will have a certain value (e.g. the final stock value will be 56.2 Euros at the closing of the market) or that an other agent perform some action (e.g. bob has sold 5 stocks to alice at 14.6 Euros)

TA could be viewed as a *blackboard* system. As one knows, the opportunistic control of knowledge sources (the micro-agents in our case), that is running the right agent on the right data at the right time, is a tricky issue in blackboards [5]. In our model, there is no fixed agenda to select one agent at a time: each agent is autonomous, and is able to modify data on the decision tree whenever it needs to. From a computer implementation point of view, it is a full multi-threading process. To preserve data coherence and integrity inside the tree, we implemented a mechanism to *solve eventual con-*

---

[3]We do not assume that a human decision-maker actually has such a decision tree inside his/her head. This is just a convenient modeling tool to tackle alternatives management in our model.

tradictions. Any agent can signal a contradiction inside a given context C. In this case, the blackboard removes the two contradictory facts from C, creates two children of C (C_1 and C_2) and instantiates the two incompatible facts in two separated contexts. This method preserves the existence of the two solutions while avoiding the contradictions.

3.2 Agents

Each agent encodes a subprocess of the decision system, like an heuristic, an inference mechanism, perception, etc.

The *perception agent* (abbreviated as **PER** in the remaining part of this paper) imports informations from environment: e.g. buy and sell orders, accepted transactions and so. This knowledge is introduced at the root of the TA as symbols, predicates and variables. Initial salience values are set, depending on decision maker's habits and experience (what he/she is used to consider as important information)[4].

The *egocentric agent* (**EGO**) helps the macro-agent to selectively enhance the salience on every facts and actions he/she is involved in (e.g. the orders he gave, the proposals he made).

The *world rules agent* (**WRU**) contains the knowledge about the world rules. It encodes the main rules and constraints within the environment like the possible actions (e.g. in our simulated game, a trader can emit buy or sell order, or cancel a previous order), the forbidden actions (e.g. to buy with a null capital), and some anticipated consequences of actions (e.g. if an order is accepted, capital and bids count are updated according to a particular formula).

The *expertise agent* (**EXP**) contains a set of domain-specific heuristics and strategies the decision maker may use to perform his/her actions. In our example of a trading game, these strategies will increase the salience of critical attributes like total capital, gain and loss. They will give the relevant hypothesis to explore, like buying or selling a share. They also value the different facts (e.g. in term of expected outcomes).

The *anchoring agent* (**ANC**) gives the set of anchoring values, that will be used as reference points. In a predicate where some attribute value is unknown, the anchoring agent enumerates all possible values, and will propose to anchor to an already perceived value or to a given reference-point value, e.g. a value linked to the personal situation of the decision maker, or a constant specific to the problem domain (a national interest rate for instance).

The *uncertainty agent* (**UNC**) encodes the uncertainty of informations in the TA. It (i) sets probability p_K for a fact K to occur, and (ii) sets the probability $Pr(C)$ of alternative context C to occur in the real world.

The *decision agent* (**DEC**) monitors the decision tree and implements the search for dominance. When an alternative is added into the tree, it evaluates it. If this is a *satisficing solution*, the tree building process is stopped, and the action that created this branch is selected. If the alternative is too low (the aggregated utility of this alternative is lower than an elimination threshold), it is ignored. In other cases, the alternative is considered to be studied later, and added to

---

[4]In real-world applications, we could ask some experienced subjects to give their rankings importance for a set of domain facts, and derive the initial salience from this. However, when we will design a learning mechanism for the salience, the importance of these initial values will be much lowered (Cf. section 5.2 in the discussion)

an internal list. When this list is full, the alternative having the highest aggregated utility is selected. We compute the utility of an alternative A as follows :

$$AU(A) = f(\sum_{C \in Child(A)} Pr(C).u(C)) \qquad (6)$$

where $Child(A)$ is a the set of immediate children of A in the tree. $Pr(C)$ the probability given by UNC agent (see above) and f is an utility normalization function, a numerical function valued in $[0, 1]$. For instance, we could adopt CARA (Constant Absolute Risk Aversion) function for risk-averse subjects $f(x) = -(1/\rho).e^{-\rho.x}$, where $\rho \in [0, 1]$ is a risking factor.

The utility $u(C)$ of an alternative C is given using a classical multi-attribute utility model, where we use the salience to weight each fact :

$$u(C) = \sum_{K \in C} p_K.S_{K,C}.v(K) \qquad (7)$$

where K is a knowledge fact in C, p_K his probability, and $v(K)$ its associated value (e.g. expected outcome) as given by EXP agent. It is worth noticing this influence of salience (computed by the other agents) into the decision process : the most salient knowledge will have the highest weight in the utility function.

3.3 Decision process overview

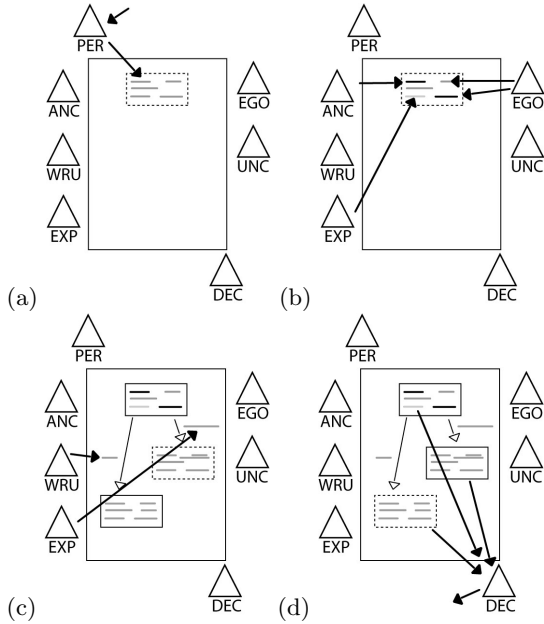We summarize the decision process in CODAGE with the flow charts depicted in Figure 3.

Figure 3: CODAGE decision process overview

Intelligence (a) The perception agent represents the current world in the root of the tree TA. (b) As soon as information appears, the EGO agent look for personal concerns and increases the corresponding saliences. Expertise agent may also update salience based on new information and its heuristics, while ANC agent increase the saliences of anchored values.

Design (c) Based on the most salient facts, agents use the TA to simulate actions, and to anticipate events and other decision maker's actions in a short or medium term. New alternatives are added to the TA, from EXP, WRU and DEC among others.

Decision (d) In parallel with (b) and (c), the decision agent assesses alternatives (utility computation), apply dominance search that leads either to the choice of an action or a selection of alternatives to be further explored.

4. SIMULATION RESULTS

4.1 Simulating an experimental market

This architecture has been instantiated in the economy field. We did not reuse the classical benchmarks used by the *Agent-based Computationnal Economics* community, like the well-known Santa-Fe Artificial Stock Market (SF-ASM), since we want to focus on cognitive aspects of decision-making within a simulated market, while SF-ASM focus on conditions of equilibrium and market behaviors using reactive agents.

We have selected an experimental financial market conducted by Biais, Hilton, Mazurier and Pouget [2]. This experimental market is aimed to study the effects of cognitive biases on the decision of traders placed on a market under asymmetric information. On this market, traders can publish at any time buy or sell orders (fixing the count and the limit price), accept an offer or cancel a previous order. There is a single risky asset, which pays a liquidating dividend at the end of the game which can be A, B or C with equal probability (in the experiment, 50, 240 and 490). Before trading starts the players receive heterogeneous private signals. For instance, if the final dividend is B, half the participants are privately informed it is not A, while the others know it is not C. There exists no communication between participants.

Hilton et al. suggest than the participants try to analyze the actions of the others to find the final asset price. Traders are reasoning in a high-uncertainty context, making them more influenced by **cognitive biases**. The authors study two biases : *overconfidence* and *self-monitoring*. Overconfidence makes the decision-maker to overestimate the representativity of his current informations. Traders suffering of self-monitoring are more attentive to the image they present to others, making them more manipulative.
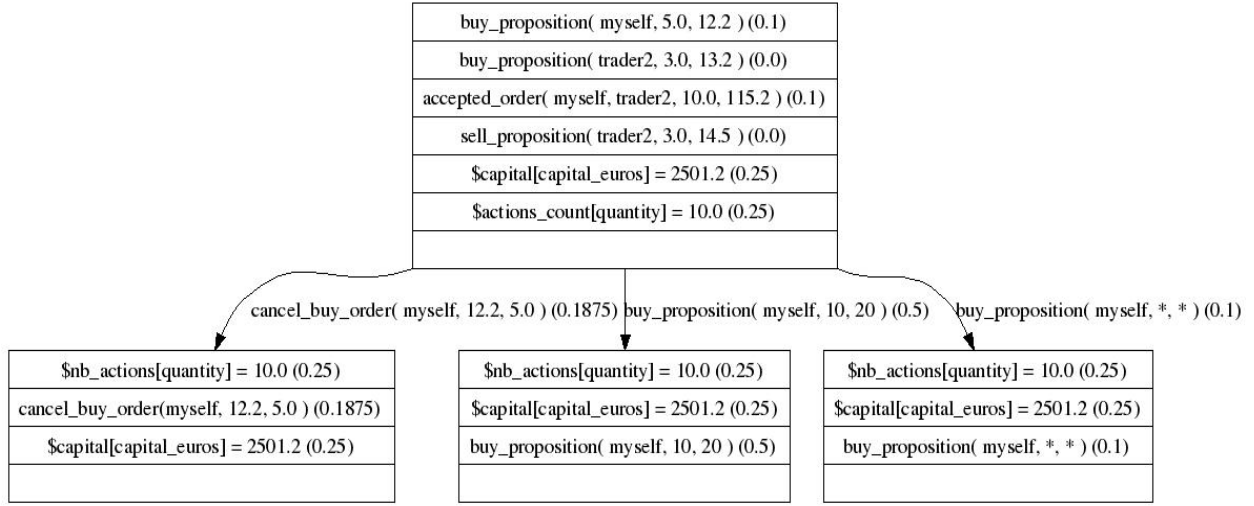
Figure 4: A Tree of Alternatives

We implement the **overconfidence bias** in CODAGE. To do so, we decrease the importance of initial probabilities (to favor current informations). Giving p_i the initial probability, $nb\_observations_K$ the number of times K is observed by the macro-agent, and $total_nb\_observations$ the total number of observations, the probability of a fact K is given by :

$$p_K = \frac{\beta.p_i + nb\_observations_K}{\beta + total_nb\_observations} \qquad (8)$$

The modification of the β parameter of UNC agent modifies the sensitivity of personal experience. The self-monitoring bias seemed to be too general to be implemented yet.

The other experimental settings are as follows. We use the decision equations (7)-(8) described in section 3.2, with f set to a simple mean function, and $v(K)$ set to fixed randomly chosen values (no prior knowledge). Finally, we use here two instantiations of ANC agent : one for quantity values anchoring, and the second for prices anchoring. Each ANC agent only favor a set of discrete value (e.g. price value or quantity value), according to a salience anchoring curves like the one depicted in Figure 5. In this Figure, the quantity values are discretized using a step equals to 5, other values will a null salience value :


Figure 5: salience anchoring curve for quantity

4.2   Example of simulation

To see how CODAGE implements the experience described above, let us examine a tree generated by our program de-picted in Figure 4 above. The process that generated this tree is the following:

- the *perception agent* added facts in the root alterna-tive. At this time, no fact is salient, it is the raw perception.

- when a fact is added, an event "NewNonSalientFactEvent" is sent to all agents. Agents dealing with salience vote for facts : EGO agent votes for all facts concerning this trader, the ANC agents vote for salience according to their own salience curve, and the EXP agent highlights facts useful for trading (salience is displayed between brackets on the figure).

- each time a salience is modified, an event "NewSalient-FactEvent" is sent with a latency, computed using Equation 5 . Each agent can react. Here, the WRU agent has proposed to cancel a previous offer or to emit a new buy order.

- when an agent proposes a new action, the TA copies salient facts (over an given recopy threshold) in the new alternative.

- WRU has added an incomplete predicate $buy\_proposition(myself, *, *)$, which contains two un-defined variables : count and price. The ANC agents propose first the most salient values, here 10 unities and 20.0.

- at each alternative modification, the DEC agent eval-uates - if possible - the alternative, and selects it if it is a satisficing one.

4.3   Overview of market simulation

An overview of the market simulation is displayed in Fig-ure 6, which shows offers (plain) and demand (dotted) curves, and the trades (squares). During a primarily period, the agents put orders that are too riskless for being accepted (low offers, high demands). Then the EXP agent modifies salience of facts leading to a compromise (we supposed it was one of the trader's general heuristics). Traders will trade on this basis. Since traders use values generated by ANC agents, only anchored values will be used in the market.
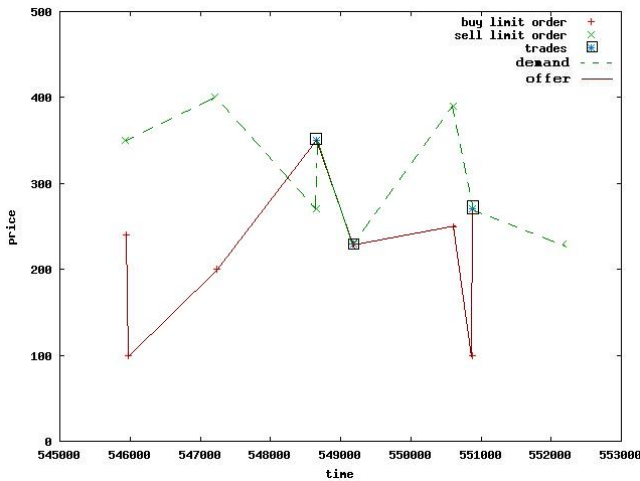
Figure 6: Biais et al. 's market simulation (extract)

5. DISCUSSION

In this paper, we propose the CODAGE approach to model human decision-making, where the decision-maker is modeled with a multi-agent system. Following psychomimetism, we designed a cognitive decision model based on psychological theories. We implemented this model with a specific multi-agent system architecture, made of a set of specialized cognitive agents that share information using a tree of alternatives TA. We gave some simulation results, and have shown how a cognitive bias could be implemented into the system. Now that we have fully described our approach, it is easier to compare to related works. Due to lack of space, we limit this comparison to the most used cognitive architecture in MAS : BDI.

5.1 CODAGE and BDI

Belief-Desire-Intention (BDI) agents have been very popular to implement cognitive processes, including decision making [18, 17]. Based on the philosophical concepts of intentions, plans and practical reason as developed by Bratman [3], it involves two important processes: deliberation where the agent commits to a particular plan, and means-end reasoning where the agent builds the possible options (plans). The agent *intends* to do what will achieve its *desires* (goals) given its *beliefs* about the world. An intention means an action the agent commits to. The question arises whether a BDI architecture could implement our cognitive decision model, and to what extent the CODAGE agent differs from the BDI framework. BDI and CODAGE do have things in common. First, they both try to account for bounded rationality: they incorporate partial exploration of states (alternatives in CODAGE, plans in BDI). Second, they both implement Simon's three stages decision model: in BDI, intelligence includes beliefs building and beliefs revision, design includes means-end reasoner to generate options and filtering, and choice includes deliberation.

However, they differ from several points. At first, it should be noted that BDI theory intends to capture rational behavior as a combination of deliberative planning and alternatives selection. It is strongly goal-oriented, and is aimed to produce plans, i.e. organizing a sequence of actions through time. In CODAGE, we adopt a broader - more abstract

view of decision-making, where actions could be selective (e.g. buying a particular stock at a given price) or more complex (e.g. executing a plan). We focus on how possible actions (alternatives) are built, explored, biased and interact with other subprocess like perception and choice. We propose a descriptive theory on how a decision-maker may process information to do so. Second, CODAGE includes several mechanisms derived from psychological theories of decision-making (selective information, anchoring, search for dominance, miscalibration bias) which is not the case in original BDI architecture[5]. Last but not least, as we mentioned above, BDI is a sequential process. In CODAGE, we believe that most of the subprocesses involved in decision-making are independent and concurrent, they influence each-other at several phases within the Simon's model and therefore should be implemented as autonomous agents in interaction within a MAS. At this early stage of our work, we are aware that this strong concurrency assumption needs further validation (we discuss the validation issue in 5.2 below).

To sum up, BDI and CODAGE do not tackle the same issues, and are not exactly concerned with the same kind of cognition process: CODAGE focuses on alternatives and decision-making while BDI focuses on plans and reasoning. Therefore, when the emphasis is made on planning, nothing prevent from using a CODAGE agent to select a plan within a BDI architecture. Similarly, one could add a BDI agent within the MAS architecture in CODAGE in order to implement some practical reasoning and planning process within a complex decision making. Again, this is yet a theoretical view, and further work is to be done in that direction to study the mutual benefits of BDI and CODAGE.

5.2 Limitations and future work

At first, we need a more complete experimental validation of our model. In section 4, we presented some experimental results of the Biais et al 's experimental market that are encouraging as the global results look like the one observed in the experiment. CODAGE is a "white box" which produces a decision process to explain human data in a human-readable form. However, we lack many informations and expertise about this experiment, since we did not have access to it. As many symbolic systems aimed to model mental behavior, we need many informations (heuristics, anchors, initial probabilities, game rules...) to be inserted into the agents. We are now looking for an experimental market to set up, with an associated methodology to define how the necessary knowledge will be obtained. Concerning the validation process, we could use external experts (e.g. economists, traders) to assess to model and/or we could put real decision makers "into the loop", that is we have real decision makers that play against our artificial macro-agents.

Second, concerning the current model itself, one shall note that we introduced saliences only for facts and not for alternatives. But, when the tree of alternatives is enough developed, too many branches might be explored simultaneously. This could affects the stability of the system, and is not "cognitive-like" as it violates bounded rationality. There-

---

[5]Recently, Norling proposed to incorporate some elements of Folk Psychology into BDI, by adding an ability to learn to recognize situations and select appropriate plan based upon this [17]. This is done by adding a metal-level plan with a Q-learning reinforcement algorithm to select plans. However, the rest of BDI architecture remains unchanged

fore, we need to implement the salience notion at alternative level, allowing only one alternative to be salient at the same time. We need to determine the rules to manage this salience of an alternative, that is to define what makes an alternative salient.

Third, we need to incorporate *learning* into the system, so the CODAGE agent becomes more adaptive. In order to overcome the influence of prior knowledge, and to enhance the reliability of the decision system, we plan to introduce reinforcement learning, where the critical parameters will be affected by the results of actions. The first step would be to have salience values learned. To do so, we could use the reinforcement learning of the RALF neural network, a cognitive architecture designed to set the parameter values in order to get a positive reward as much a possible (RALF does not optimize reward, but find a satisficing solution) [12].

Finally, the CODAGE architecture is intended to be as much generic as possible. It could be viewed as an agent-based decision framework where different decision heuristics and biases could be implemented. We should move further into that direction, and try to incorporate new types of agent (e.g. emotional), new kinds of biases and heuristics.

Acknowledgments

6. REFERENCES

[1] J. R. Anderson, D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere, and Y. Qin. An integrated theory of the mind. *Psychological Review*, 111(4):1036–1060., 2004.

[2] B. Biais, D. Hilton, K. Mazurier, and S. Pouget. Judgmental overconfidence, self-monitoring and trading performance in an experimental financial market. *Review of Economic Studies*, 72(2), April 2005.

[3] M. Bratman. *Intention, Plans and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.

[4] C. Castelfranchi. The theory of social functions. challenges for multi-agent-based social simlation and multi-agent learning. *Cognitive Systems Research*, 2(1):5–38, 2001.

[5] D. D. Corkill. Collaborating Software: Blackboard and Multi-Agent Systems & the Future. In *Proceedings of the International Lisp Conference*, New York, New York, October 2003.

[6] S. Dillon. Descriptive decision making: Comparing theory with practice. In 33^{rd} *Annual Conference of the Operational Research Society of New Zealand*, Auckland, 1998.

[7] F. Guttman, A. G. Moukas, and P. Maes. Agent-mediated electronic commerce. *The Knowledge Engineering Review*, 13(2):147–159, 2002.

[8] N. H.-C. Salience from feature contrast : temporal properties of saliency mechanisms. *Vision Research*, 40(18):2421–2435, 2000.

[9] C. Haynes and S. Kachelmeier. The effects of accounting contexts on accounting decisions: A synthesis of cognitive and economic perspectives in accounting experimentation. *Journal of Accounting Literature*, 1998.

[10] K. Kahneman, P. Slovic, and A. Tversky, editors. *Judgement under uncertainty: Heuristics and Biases*. Cambridge University Press, New-York, 1982.

[11] J.-D. Kant. Modeling human cognition with artificial systems : some methodological considerations. In *Proceedings of HCP 99 (Human Centered Processes)*, pages 501–508, Brest, September 1999.

[12] J.-D. Kant and D. Levine. Ralf : A simplified neural network model of rule formation in the prefrontal cortex. In *Proceedings of JCIS98,*, volume II, pages 8–14, Research Triangle Park, USA, october 1998.

[13] D. Koehler and N. Harvey, editors. *Blackwell Handbook of Judgment and Decision Making*. Blackwell, 2004.

[14] B. LeBaron. Agent-based computational finance: suggested readings and early research. *Journal of Economic Dynamics & Control*, 24:679–702, 2002.

[15] M. Minsky. *The society of mind*. Simon & Schuster, Inc., New York, NY, USA, 1986.

[16] H. Montgomery. Decision rules and the search for a dominance structure: toward a process model of decision-making. In P. Humphrey, O. Svenson, and A. Vari, editors, *Analyzing and Aiding Decision Process*. North-Holland, Amsterdam, 1983.

[17] E. Norling. Folk Psychology for Human Modeling: Extending the BDI Paradigm. In *AAMAS 04*, New York, July 2004.

[18] S. Parsons and N. R. Jennings. Argumentation and multi-agent decision making. In *Proceedings of the AAAI Spring Symposium on Mixed-Initiative Decision Theoretic Systems*, Stanford, 1998.

[19] S. Parsons and M. Wooldridge. Game theory and decision theory. In S. Parsons, P. Gymtrasiewicz, and M. Wooldridge, editors, *Game Theory and Decision Theory in Agent-Based Systems*. Springer Verlag, Berlin, Heidelberg, New-York, 2002.

[20] H. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69:99–118, 1955.

[21] H. Simon. Rational choice and the structure of the environment. In H. Simon, editor, *Models of Man*. John Wiley, New-York, 1957.

[22] H. Simon. *The New Science of Management Decision*. Harper and Brothers, New York, 1960.

[23] A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211:453–458, 1981.