La négociation automatisée sur les marchés électroniques avec agents intelligents : discussion

B. Leloup[1]
Benoit.Leloup@iaai.fr

[1]Institut des Applications Avancées de l'Internet,
12, avenue du Général Leclerc, 13003, Marseille, France
et Ecole Nationale Supérieure des Telecommunications de Bretagne, France

Résumé :
Les agents intelligents constituent une technologie puissante pour aborder les problèmes dynamiques complexes comme la tarification dynamique ou la négociation automatisée. La majorité de la littérature informatique sur les systèmes multi agents pour le commerce électronique considère des agents qui disposent de règles de choix très simples inspirées par les comportements humains. Dans le futur, il semble que les places de marché électronique verront cohabiter des agents humains et des agents artificiels programmés avec des heuristiques simples. Cet article se propose d'aborder certaines des questions posées par la conception de telles places de marché électronique.
Mots-clés : Agent intelligent, négociation automatisée, conception de places de marché pour le commerce électronique

Abstract:
Intelligent agent technology provides a powerful mechanism to address complex dynamical problems such as dynamic pricing or automated negotiation. Following an heuristic-based decision making approach, most of the computer science literature on *E*-commerce multi-agent systems focuses on behavioral agents endowed with simple rules of thumb inspired by observed human behaviors. It seems that in the future, electronic marketplaces will be composed by real agents as well as by artificial ones endowed with simple heuristics. Our paper aims to discuss some questions raised by the design of such electronic marketplaces.
Keywords: Intelligent agent, automated negotiation, electronic commerce market design

1 Introduction

Intelligent agents are software abstractions that are able to communicate and make decisions in an (semi) autonomous way. Depending on the problem addressed, they have various capacities and names such as *E*-commerce agents, shopbots, pricebots, intellibots and intelligent software robots. As *E*-commerce involves dynamics and heterogeneity, software agents will surely play crucial role in future *E*-commerce transactions. Unsurprisingly, the literature on *E*-commerce shows a growing interest for electronic agents that could assist firms and customers in their electronic commercial transactions (e.g., Faratin *et al.* [1998], Jennings and Wooldridge [1998], Leloup and Deveaux [2001], Jennings *et al.* [2001]). For customers, the use of a shopping agent can facilitate the search for products, the comparisons of prices and the task of processing product information. In the future, these shopping agents should evolve in order to directly contract and buy on the behalf of their owner. Such intelligent agents can be called "buying agents." Selling agents are interesting too, as they improve the productivity of online business interactions and allow the sellers to monitor and process several parallel commercial transactions with customers (human and / or artificial agents), thus increasing their profits. However, at the time being, the design of these intelligent agents remains a research issue. In particular, various transactions mechanisms (e.g., auctions, dynamic posted prices, bargaining) as well as their economic efficiency have to be explored. Let us note that currently, the Internet mainly involves simple negotiation protocols as posted prices (i.e., take it or leave it offers) or auctions. A growing number of researchers in the field of intelligent agent systems recognizes that there are certain problems and issues surrounding intelligent agent terminology and technology that must be resolved for the agent technology to develop. Representation, support and automation of negotiations on electronic markets are some of these important issues.

This paper aims to present and discuss some desirable properties of Internet agent-based negotiation models that increase their applicability to real *E*-business interactions: polymorph population of customers, multi-negotiation objects, taking into account reputations and/or learning effects and social interactions structures (i.e., situations in which customers can communicate their experiences to their neighborhood), introducing the possibility for the agents to restructure the negotiation problem if the current negotiation process fails on a particular dimension of the negotiation (i.e., to allow the software agents to 'open' the negotiation space to new negotiation objects so as to achieve an agreement), taking into account human / software agents interactions, etc.

The paper proceeds as follows. Section 2 aims to present formal definitions of the notions of heuristics and rules of thumb that are adapted to the study of Internet agent-based markets. We discuss about connections between heuristic-based decision making approaches and the standard dynamic programming procedure. Section 3 focuses attention to the horizon of time and to the timing of negotiation processes. Section 4 presents some important features of Internet agent-based market models. In particular, we discuss about the way sellers may perceive customers' behaviors on *E*-commerce marketplaces, the reasonable assumptions one can make about agents information sets and their rationality (game theory vs. interactive decision theory), etc. In section 5, we present various network interaction structures and we focus our attention to situations in which customers are able to communicate their experiences to their neighborhood. We design some stylized examples of *E*-commerce marketplaces and we discuss some important features of real negotiation problems such as restructuring a negotiation problem and human / software (intelligent) agents interactions.

2 Heuristics, rules of thumb and dynamic programming

Intelligent agent technology provides a powerful mechanism to address complex dynamical problems such as dynamic pricing or negotiation. Because it is at the intersection of various academic domains as decision support systems, economics, psychology, decision theory, game theory and machine learning, such a research area should ideally embrace a multidisciplinary perspective. Some empirical evidences in psychology (e.g., Pruitt [1981]) and in economics (Simon [1992], Lettau and Uhlig [1999]) as well as some insights from evolutionary game theory (Weibull [1995]) tend to show that people base their strategies on simple heuristics. Most of the computer science literature on E-commerce multi-agent systems (e.g., Faratin et al. [1998]) focuses on behavioral agents endowed with simple rules of thumb inspired by observed human behaviors. On the contrary, the statistical (Bayesian) decision theory approach (based on the dynamic programming (DP) principle and optimal stochastic control methods) and the game theory approach which are more common in the economic literature assume that agents have large (huge) computational (and information processing) capabilities. These approaches (simple heuristics, DP) thus seem to be competing ones. Heuristics (also called "procedural knowledge") may be defined as "(…) criteria, methods, or principles for deciding which among several alternative courses of action promises to be the most effective in order to achieve some goal" (Pearl [1984]). This definition captures the essence of many other definitions, such as rule of thumb or a method not guaranteed to find the optimal solution but in practice often finds 'good' quality solutions. Unfortunately there is no correct formal definition of what a heuristic is and heuristics encompass a wide range of methods. The modern meaning is a procedure that may give an optimal (best) solution to a problem but offers no guarantee of doing so. If it can be proven that an exact solution exists, then this becomes an algorithm. Our analysis takes the sellers' viewpoint. Customers' viewpoint is similar (although there are important differences) and thus will not be discussed here. Let $\left\{ h_i \right\}_{i=1}^{k}$, be the set of (negotiation) heuristics a seller endowed his agent with. The selling agent may select these heuristics in any order but only one at a time. Following Leloup (2001a), one can define a negotiation heuristic h_i as,

i) A starting offer which is, by definition, the largest offer a selling agent can make when it follows heuristic h_i.
ii) A predetermined rule of thumb.

A rule of thumb is a mapping from a subset of the states of the negotiation process into the selling agent's action space (e.g., concede, do not concede, agree, do not agree). For example, a rule of thumb might say "when you are in state s_1 (e.g., when you have conceded once whereas your opponent has conceded twice), use action a_1 (e.g., concede); when it is in state s_2 (e.g., when you have conceded twice whereas your opponent has conceded twice), use action a_3 (e.g., wait), etc." If we impose no further restrictions, like limited complexity or bounded rationality (Simon [1992]), then implementing the DP solution may be viewed as a particular rule of thumb (Fogelman-Soulie *et al.* [1983]). Hence, if there are no computational limitations, then one can endow a selling agent with DP as one of its competing rules of thumb. As it has been shown in the literature (Lettau and Uhlig [1999]), an agent can learn to use a rule of thumb encoding suboptimal behavior even when that rule is competing against another rule which implements the DP solution. That is he can fail to learn rational behavior. Thus, if the set of predetermined heuristics a selling agent is endowed with is the result of a learning process (e.g., *via* genetic algorithms or reinforcement learning algorithms), then there is no guarantee for a seller to learn the optimal rule of thumb, that is there is no guarantee for a seller to endow his agent solely with a DP rule. As a consequence, heuristics-based decision making procedures and DP solutions should not be seen as competing approaches but rather as complementary ones. Thus, it is highly desirable to extend current Internet agent-based market models to allow the selling agents to follow, at any time, one of their several rules of thumb, they have been initially endowed with.

3 Time and negotiation

3.1 Horizon of time

For the purpose of designing electronic marketplaces one can consider various kinds of time horizon: rate of time preference, continuous or discrete time approaches, finite horizon time, infinite horizon time with geometric discount factor, infinite horizon time with hyperbolic discount factor, stochastic discount factor, etc. Most of the time, agent-based markets have to be viewed from a prescriptive viewpoint and not from a descriptive one: intelligent agents are customized software (i.e., they take into account, as much as possible, customers or sellers' preferences about risk, prices, goods, delay, etc.) who do their best given their information about their environment and their (limited) computational capabilities. As is the case for mathematical finance (which also follows a prescriptive approach), inter-temporal consistency of choices is a highly desirable feature. In order to guarantee such an inter-temporal consistency of the agents' choices, when time is discrete (continue), it is well known that we have (necessary condition) to consider geometric (exponential) discounting. From an operational viewpoint and according to current computer technology, continuous time models are often not suitable and should, at best, be viewed as an approximation for discrete time models (with decision time intervals which tend towards [computer] zero) when the latter are too complex to deal with. Another important question concerns the horizon time itself: do we need to consider infinite or finite horizon time? With finite horizon time, computations and interpretations are often easier and thus most of models assume finite horizon time. However as is well known in (non cooperative) game theory, the use of finite horizon time may result in end game effects (Selten and Stoecker [1986]) and paradox of rationality may occur (e.g., the chain store paradox). One of the main problems with infinite horizon time is interpretation. Let a ($0 < a < 1$) denotes the discount factor. When $(1 - a)$, represents the probability for the market to collapse at each time, then the horizon time of the repeated negotiation problem may be viewed as stochastic. With this interpretation, one can easily show (Leloup [2001a]) that the expected number of periods is $1/(1 - a)$. In this setting, it is the expected number of periods (and thus the features of the market under study) which defines the value of the discount factor and not the converse. If a seller expects to play about 20 periods (e.g., to meet about 20 customers), then he should set his discount factor to 0.95. With this interpretation the discount factor a should not be interpreted as a rate of time preference but as a continuation probability. It is interesting to note that with this interpretation (stochastic discounting) it is suitable to assume the sellers and the customers to have a same discount factor (because it is based on the features of the market under study); an assumption which is usually not suitable with a rate of time preference interpretation. Hence, even if infinite horizon time models are often difficult to handle, one should (when it is possible) assume geometric discounting and infinite horizon time. Moreover, as underlined by Friedman (1991, p.118), both a rate of time preference and a stopping probability can fall within a same model: let us suppose the seller has a rate of time preference equals to $r > 0$ and a continuation probability equals to b. Then $a = b/(1+ r)$.

3.2 Timing of the negotiation process

Single instances of the repeated negotiation problem (with various customers) are called "negotiation processes." When time is discrete, it is convenient to focus attention to repeated negotiation processes in which customers arrive sequentially. In this setting, the maximal duration of a negotiation process can be normalized to one unit of time, each negotiation process yielding a reward (the reward is 0 if there is no agreement). With this convention, there are two types of time in the repeated negotiation problem. The time of the decision process: it is the arrivals' rate of customers and it is normalized to one. The second type of time is the time of negotiation processes. Of course, depending on the particular heuristic the selling agent has chosen and the behavior of its customer, some negotiation processes may take more time than others. As the sequence of negotiation processes is driven by customers arrivals, and as the maximal duration of a negotiation process is normalized to one, we don't have to care about random decision times (e.g., semi-markov decision processes) and/or parallel negotiation processes (with a same customer).

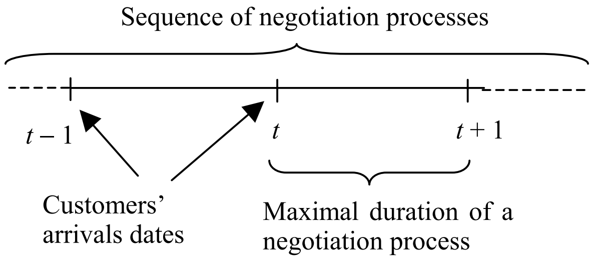

FIG. 1 – Timing of the decision process

In this setting, we allow the seller to carry out as much repeated negotiation problems in parallel as he needs [e.g., by relying on several agents] but we do not allow him to engage parallel negotiation processes within a given repeated negotiation problem [i.e., the selling agent cannot simultaneously carry out several negotiation processes with a same customer]. We believe this to be a reasonable and realistic assumption for e-markets composed with artificial selling agents who negotiate with real and/or artificial buying agents.

4 Sophisticated behaviors and interactive decision theory

4.1 Interactive decision theory vs. game theory

Economics has a growing interest for dynamic complex systems. As a social science it has to deal with heterogeneity in human behaviors, (individual and collective) learning, path dependency and uncertainty. Within such a complex and uncertain environment, economists are primarily concerned with qualitative results such as limit-theoretic behaviors (e.g., convergence of beliefs and/or of actions). Economic agents very rarely have complete knowledge of all variables affecting their strategies. However, as it is underlined by Aghion et al. (1991), "the theory of imperfect competition generally assumes that individual firms know all relevant aspects of demand function." As firms cannot observe all the variables affecting consumer choices, the best they can do is to represent consumer decision rules by constructing choice probabilities (i.e., stochastic individual demand functions). The set of probabilities associated with the various choices the firms can make corresponds to the (usually parametric) response function of the consumers. This response function may be known or unknown to the firms. When this function is unknown, firms can learn its parameters by experimentation.

Individual decision theory (utility theory) deals primarily with behaviors in which the outcome depends on the decision maker's own behavior. As this is underlined by Harsanyi (1977), individual decision theory "can also handle situations where the outcome does depend on other individual's beliefs – as long as it is assumed that their behavior is governed by well defined deterministic or probabilistic laws rather than by their own rational strategy choices and their own rational expectations about other participants' behaviors." For these problems, one can formulate the interactive decision problem as a game against nature.

The interactive decision theory approach is less exigent about the information the agents have about each other than game theory does. By its very nature, the study of agent-based electronic marketplaces has to deal with polymorph population of agents (that is with agents who have distinct forms of rationality). Indeed, autonomous agents are (or will be) designed by competing providers who endow their agents with different kinds of technologies. Moreover, at the time being, it is far from obvious for a seller to distinguish, by his sole observations, if he negotiates with a human agent or with an artificial agent. As human agents exhibit different behaviors than artificial ones (the latter being usually more consistent than the former), common knowledge of rationality (a standard assumption in game theory) is not a suitable assumption (i.e., incomplete information games are not appropriate) to study agent-based electronic markets.

4.2 Sellers' perception of customers behaviors

Sellers usually neither know the negotiation skills of their customers nor their reservation prices. As already underlined, so as to take into account the lack of information a seller has about his customers reservation prices, one can follow an interactive decision theory approach, that is one can assume that the best the seller can do is to represent customers' decision rules by constructing choice probabilities (Anderson *et al.* [1992]). There are various assumptions one can make to define such choice probabilities. For an economic viewpoint, the interested reader may refer to Anderson *et al.* (1992) or Leloup (2002). In this setting, sellers affect, for each one of their negotiation heuristics (DP included), agreement probabilities to customers. Depending on the customers' negotiating behaviors, each heuristic has its own strengths and weaknesses. Thus, the probability for a given heuristic to reach a particular agreement point may not be the same for all heuristics. Let $\pi^i_j(n)$, be the probability for an agreement point j to be reached when following heuristic h_i for the n^{th} time. We can distinguish "simple situations," in which the probabilities of agreements $\pi^i_j(n)$ are constant with n, from "sophisticated situations" in which they are not. In risky situations, the $\pi^i_j(n)$ are all known to the sellers. It is more interesting and realistic to focus attention to sophisticated situations in which the probabilities of agreement that are associated to each heuristic are not necessarily constant with the number of times a selling agent used them (e.g., situations in which customers have learning and/or punishment and/or communication skills). Since customers behaviors are sophisticated, the seller then cannot make a stationary assumption on the probabilities of purchase he affects to customers. Because they involve non stationary agreement probabilities, sophisticated situations are complex to deal analytically with and one must often rely on multi-agent system simulations to understand the dynamics of the market (Leloup [2002]).

For simple situations (e.g., situations in which customers use non adaptive or non reactive rules of thumb), the selling agent can learn on line these probabilities by experimentation (Leloup and Deveaux [2001]) or the seller can elicit them before the repeated negotiation problem starts with off line learning techniques (e.g., by carrying out market studies). When off line learning is possible one should preferably elicit the probability of agreement *via* laboratory experiments and/or multi-agent systems simulations. For repeated negotiation problems that lie within an interactive decision making approach and for which the behaviors of the customers do not change as time goes on (i.e., for repeated negotiation problems with non sophisticated behaviors), the interested reader may refer to Fogelman-Soulie *et al.* (1983). Let us note that when a seller negotiates with non sophisticated customers, then he has to (or can only) make a (quasi) stationary assumption on the probabilities of agreement he affects to customers so that his optimal negotiation strategy can be obtained as a solution of a stochastic control problem. For the case of a single selling agent who is randomly matched with several (sophisticated) buying agents (risky environment and off line learning allowed), Leloup (2001a) addressed the issue of the design of a monitoring algorithm that is able to sequentially and optimally choose between several negotiation heuristics in a repeated negotiation problem.

4.3 The selling agent's monitoring policy

For the single selling agent's monitoring decision problem, that is for the problem of finding a meta (negotiation)-heuristic that sequentially chooses the negotiation heuristics the selling agent follows during its repeated interactions with buying agents, it can be shown that the selling agent's optimal strategy (i.e., the optimal sequence of heuristics) exists and that it can theoretically be determined by standard DP (see Leloup [2001a]). However, the decision tree of the repeated negotiation problem grows rapidly and thus becomes quickly intractable when the set of heuristics and/or the number of interactions is large. Thus, for operational purposes, an approach based on Bellman's DP equations is usually not appropriate (intractability of the optimal solution). To answer operational requirements, we thus have to focus on monitoring policies that both are tractable and computable in reasonable time by the (intelligent) agents.

5 Interactions: illustrations

5.1 Neighborhood and sophisticated behaviors

On traditional markets, customers usually communicate their purchase experience to their neighborhood. Such information about the conditions and commercial advantages the customer bought (or not) the good (price, timings, quality, etc.) often has a significant influence on his neighborhood (future) purchasing behaviors. In particular, it is likely that his neighbors will refuse to buy the good (with the same commercial advantages and conditions) at a higher price than the one he was charged, even if their initial reservation price was higher than or equal to it (i.e., even if, without this information, they would have bought the good at this price). Thus, if we assume that customers communicate their purchase experience to their neighborhood, it is likely for the prices the seller can charge to be sticky upwards. Of course, prices are more or less sticky upwards depending on the customers' vicinity and on the diffusion rate of the information within the population of customers. Properties of local and global network interactions models are of particular interest when designing electronic commerce marketplaces. Various definitions of neighborhoods are possible. Let ϑ_j, be customer j's neighborhood. ϑ_j is the number of people customer j is connected with in the social network involved (population of customers). We also need to define interaction structures. Concerning the two dimensional lattice the following definitions are common in the automata literature (e.g., Wolfram [1994]),

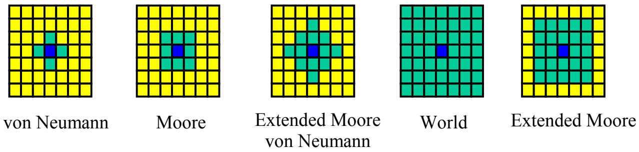

FIG. 2 – Interaction structures on two dimensional lattices

5.2 Sophisticated customers vs. a single selling agent

Let us consider an agent-based market in which a single selling agent is randomly matched with buying agents. Buying agents are all endowed with the same negotiation skills (i.e., they have the same set of heuristics and the same monitoring algorithm). As they are customized to fit their owner preferences they however do not behave the same. These agents may come from a same provider (e.g., IBM) and thus are endowed with the same technology. At the end of its commercial transaction (e.g., posted price negotiation protocol or bargaining), each buying agent communicates to its Moore neighborhood (8 neighbors) the price displayed (or the agreement point achieved with) by the selling agent. For simplicity, one can assume that buying agents have a reservation price rule and an *ad hoc* policy which consists in rejecting all prices that are strictly higher than a price that has been charged (in the past) by the selling agent to a member of their Moore neighborhood, even if these prices are lower than or equal to their initial reservation price.

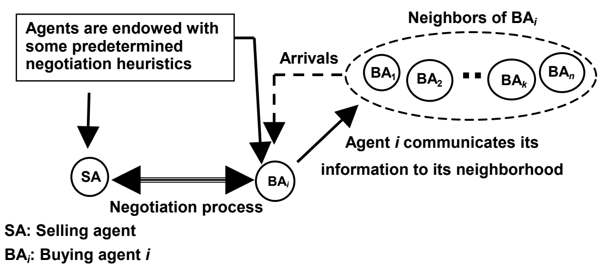

FIG. 3 – Artificial vs. Artificial

Due to the behavioral rules followed by the population of buying agents the response function of the population of customers (which, with posted price negotiation protocols, usually corresponds to the de-cumulative distribution function of the probabilities of purchase) is no longer stationary (i.e., the probabilities of agreement that are associated to each heuristic are not constant with the number of times a selling agent used them). Such sophisticated situations are complex to deal analytically with. In case of global interaction (world neighborhood) one can establish analytically the optimal policy of a selling agent who perceives the behaviors of the buying agents through an unknown (and non stationary) distribution function (Leloup [2001b]). However, to our knowledge, for other kinds of interaction structures there are no analytical results. Thus, to get insights about the dynamics of such marketplaces and to parameterize their selling agents, sellers have to carry out multi-agent system simulations of interactions between their selling agents and buying agents. For a presentation of multi-agent system simulations of agent-based electronic commerce marketplaces with Moore neighborhood (local interactions on a two dimensional lattice) and posted price negotiation protocol, the reader may refer to Leloup (2002) and the references therein. Of course, on traditional markets the impact of social interactions on customers behaviors is not as simple as the one postulated in the preceding example. Indeed, there are several (usually unknown to the observer) reasons that may lead a human agent to continue to accept, with positive probabilities, to buy the good at a higher price than the one that has been charged in the past by the selling agent to a member of his neighborhood. On traditional markets, interaction structures are also usually more complicated than the stylized notions of neighborhood usually considered in the social network interaction literature: there are strong and weak connections between customers, the information provided may be ambiguous, some customers have more influence than others (gurus phenomena), etc. There are many other kinds of sophisticated situations that may be pertinent for *E*-commerce market design. An interesting situation concerns the case in which a selling agent is matched with identified (e.g., by means of an identification number or using cookies) human customers. For this particular decision problem, the seller can test sequentially and preferably in laboratory experiment conditions, each one of his heuristics. The probabilities of agreement thus elicited are then the empirical frequencies observed in the population of experimental subjects. In this setting, the seller needs to assume that the customers he negotiates with behave like an average and representative customer of the whole population of experimental subjects, he tested before the repeated negotiation problem starts. Of course, for a selling agent to be fully operational in such environments, most of the hard work is still to be done: for real applications, one still needs to define the pertinent demand pattern, the timing of customers' arrivals, the matrix of the various probabilities of agreement, etc.

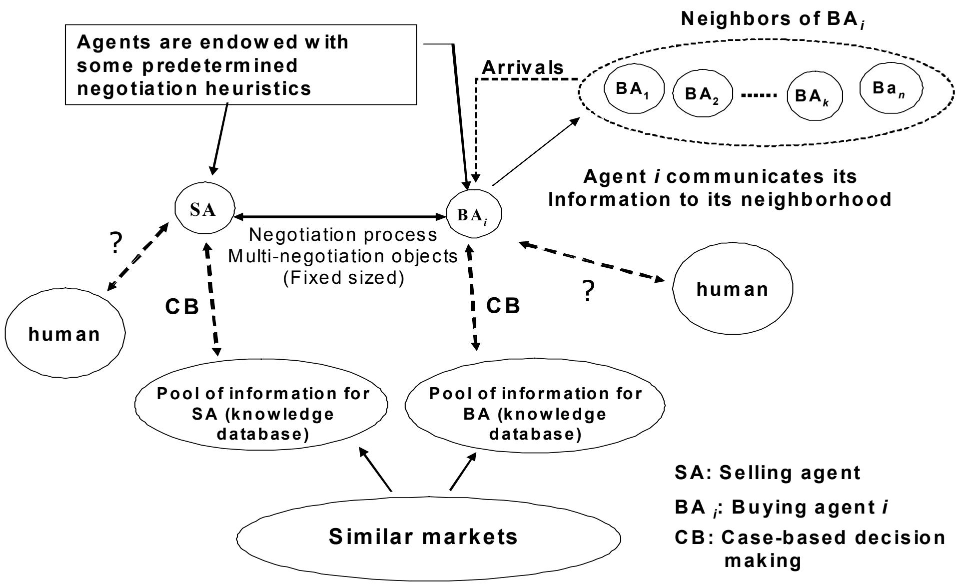

FIG. 4 – Multi-negotiation objects and restructuring in negotiation: human and artificial

5.3 Multi-negotiation objects

The negotiation object (i.e., the range of issues over which an agreement can be reached) we considered until now is only a single issue (the price) while "it may cover hundreds of issues (related to price, quality, timings, penalties, (…), etc.)" (Jennings et al. [2001]). An extension of the theory to multi-negotiation objects is a highly desirable feature that would increase its applicability to real *E*-business interactions. However, multi-negotiation objects raises new important questions such as: how a selling (buying) agent may restructure its negotiation problem if the negotiation fails on a particular dimension (i.e., how to design intelligent agents who are able to 'open' the negotiation to new negotiation objects so as to achieve an agreement). For the task of restructuring the negotiation problem one can reason by similarity with a case-based reasoning approach (Gilboa and Schmeidler [1996]) and/or use some restructuring heuristics as defined in Shakun (1996).

We next describe a plausible algorithm for restructuring the negotiation problem. If, there are no possible agreement points that can be achievable, then

- Check the knowledge database. If there are similar problems, then follow a case-based reasoning approach to restructure the negotiation problem,

- If there are no similar problems, then delay the negotiation and ask for a feedback from a human agent (e.g., via mobile phone communication) and apply his new proposed solution,
- Add this new solution to the knowledge database.

For implementing such an algorithm, we need to develop new technologies that connect mobile commerce (*M*-commerce) and *E*-commerce. Such a connection between *E* and *M*-commerce may, at the same time, create new valuable services for customers and make it easier the use of intelligent agents thus accelerating the adoption of this technology.

6 Concluding remarks

In this paper, we have discussed some important features for the design of agent-based *E*-commerce marketplaces. We argued that heuristics-based decision making procedures and DP solutions should not be seen as competing approaches but rather as complementary ones. We have discussed about the horizon of time of negotiation processes. We emphasized that in order to avoid end game effects and rationality paradox the horizon of time of the negotiation problem should be stochastic rather than finite. Moreover, as is the case in mathematical finance, in order to guarantee an inter-temporal consistency of the agents' choices, the discount parameter must be geometric or exponential. In this setting and with a stochastic discounting interpretation, it may then be acceptable to assume the sellers and the customers to have a same discount factor. An important distinction between game theory and interactive decision theory has been presented. As interactive decision theory is less exigent than game theory about the information the agents have about each other, and because common knowledge of rationality, which is a necessary condition for game theory to apply, is a too strong assumption when studying polymorph populations, it has been argued that an interactive decision theory approach is more appropriate than a game theory one to study agent-based *E*-commerce marketplaces. Sellers usually neither know the negotiation skills of their customers nor their reservation prices. When following an interactive decision theory approach one can take into account the lack of information a seller has about his customers reservation prices, by assuming that the best he can do is to represent customers' decision rules by constructing choice probabilities. With this interpretation, one still needs to explicit the way the sellers elicit the probabilities of agreement that are associated to each one of their heuristics. Ideally, a model of learning should both allow for on line learning and for off line learning. When off line learning is possible one should preferably elicit the probability of agreement *via* laboratory experiments and/or with multi-agent system simulations.

We introduced a distinction between "simple situations," in which the probabilities of agreement that are associated with each one of the selling agent's heuristics are constant with the number of times it used them, and "sophisticated situations" in which they are not. As is the case on traditional markets, reputation effects and/or learning effects cannot be avoided when designing *E*-commerce marketplaces. The topology of network interactions between customers is of primary importance, both to understand and to predict (from the sellers viewpoint) the behaviors of the customers. It was not possible to present in details all aspects and concepts that have been developed in the social network interaction literature. We thus have focused our attention to some specific notions of neighborhood which are very common in the automata literature. However, we have underlined that real life interaction structures are usually more complicated than these stylized notions of neighborhood. A lot of work is still to be done on this research area. On the contrary to traditional markets, *E*-commerce marketplaces may involve global interaction structures such as world neighborhood. Indeed, the current pace of *E*-commerce, the development of virtual communities, and the emergence of the agent technology will allow faster communication and will facilitate exchanges of information between customers. In this setting, global interaction structures like world neighborhood may become more and more relevant to characterize customers' interactions on the Internet. However, even if a concept of world neighborhood may be pertinent for Internet agent-based markets, one must be cautious: the diffusion rate of the information within a population of customers depends on their vicinity, their connectivity and on the size of the population. Finally, we briefly addressed some important questions raised by the extension of current agent-based negotiation models to multi-negotiation objects. Many important aspects were not discussed in this paper. In particular, we did not attempt to answer questions about the economic efficiency of the aggregate agent-based market, the dynamics of the prices one can observe with sophisticated negotiation behaviors, and the effect of competition among sellers on their selling agents' optimal monitoring algorithms. Another interesting question concerns whether how and why the various heuristics the agents are endowed with might be changed in the course of time.

Références

Aghion, P., Bolton, P., Harris, C., B. Jullien, (1991) "Optimal Learning by Experimentation" *Review of Economic Studies*, *58*, 621-654.

Anderson, S.P., de Palma, A., and Thisse, J.-F (1992), *Discrete Choice Theory of Product Differentiation*, MIT Press, Cambridge MA.

Bellman, R.E. (1957), *Dynamic Programming*, Princeton, Princeton University Press.

Faratin, P., Sierra, C. and Jennings, N.R. (1998), "Negotiation Decision Functions for Autonomous Agents." *Journal of Robotics and Autonomous Systems*, 24 (3-4), 159-182.

Fogelman-Soulie, F., Munier, B. and Shakun, M.F. (1983), "Bivariate Negotiations as a Problem of Stochastic Terminal Control." *Management Science*, 29 (7), 840-855.

Friedman, J.W. (1991), *Game Theory with Applications to Economics*, Oxford University Press.

Gilboa, I., Schmeidler, D., (1996), "Case-Based Decision Theory." *Quarterly Journal of Economics*, 110, 605-639.

Harsanyi, J.C. (1977), *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, Cambridge University Press.

Jennings, N.R., Wooldridge, M.J. (1998), *Agent Technology,* Springer, Berlin.

Jennings, N.R., Faratin, P., Lomuscio, A.R., Parsons, S., Wooldridge, M.J., and Sierra, C. (2001), "Automated Negotiation: Prospects, Methods and Challenges." *Group Decision and Negotiation*, 10, 199-215.

Leloup, B. (2001a), "Learning to Negotiate on an Internet Agent-Based Market." *Proceedings of Group Decision and Negotiation* 2001, June 4-7, La Rochelle, France, 189-194.

Leloup, B. (2001b), *Second Order Uncertainty in Economics: The Exploration vs. Exploitation Compromise*. Ph.D. thesis, Ecole Normale Supérieure (Cachan), France.

Leloup, B. (2002), "Pricing with Local Interactions on Agent-Based *E*-Marketplaces." *Journal of Electronic Commerce Research and Applications* (forthcoming).Leloup, B., Deveaux, L. (2001) "Dynamic Pricing on the Internet: Theory and Simulation." *Electronic Commerce Research*, 1 (3), 53-64.

Lettau, M., Uhlig, H. (1999), "Rules of Thumb versus Dynamic Programming," *The American Economic Review*, 89 (1), 148-174.

Pearl., J. (1984), *Heuristics - Intelligent Search Strategies for Computer Problem Solving*. Addison-Wesley.

Pruitt, D.G. (1981), *Negotiation Behavior*, Academic Press, New-York.

Selten, R., Stoecker, R. (1986), "End behavior in Sequences of Finite Prisoner's Dilemma Supergames." *Journal of Economic Behavior and Organization*, 7, 47-70.

Shakun, M.F. (1996), "Modeling and Supporting Task-Oriented Group Processes: Purposeful Complex Adaptive Systems and Evolutionary Systems Design," *Group Decision and Negotiation*, 5, 305-317.

Simon, H. (1992), *Economics, Bounded Rationality, and the Cognitive Revolution*, Cheltenham, Edward Elgar.

Weibull, J. (1995), *Evolutionary Game Theory*, Cambridge MIT Press.

Wolfram, S. (1994), *Cellular Automata and Complexity*. Addison-Wesley, Menlo Park.