Artificial Intelligence Review (2005) 23:185–203
DOI 10.1007/s10462-004-6391-1

A Comparative Study of Game Theoretic and Evolutionary Models of Bargaining for Software Agents

SHAHEEN S. FATIMA[1], MICHAEL WOOLDRIDGE[1] & NICHOLAS R. JENNINGS[2]
[1]*Department of Computer Science, University of Liverpool, Liverpool L69 3BX, UK (E-mail: {S.S.Fatima, M.J.Wooldridge}@csc.liv.ac.uk);* [2]*School of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, UK (E-mail: nrj@ecs.soton.ac.uk)*

Abstract. Most of the existing work in the study of bargaining behavior uses techniques from game theory. Game theoretic models for bargaining assume that players are *perfectly rational* and that this rationality is *common knowledge*. However, the perfect rationality assumption does not hold for real-life bargaining scenarios with humans as players, since results from experimental economics show that humans find their way to the best strategy through trial and error, and not typically by means of rational deliberation. Such players are said to be *boundedly rational*. In playing a game against an opponent with bounded rationality, the most effective strategy of a player is not the equilibrium strategy but the one that is the best reply to the opponent's strategy. The evolutionary model provides a means for studying the bargaining behaviour of boundedly rational players. This paper provides a comprehensive comparison of the game theoretic and evolutionary approaches to bargaining by examining their assumptions, goals, and limitations. We then study the implications of these differences from the perspective of the software agent developer.

Keywords: bargaining, e-commerce, evolutionary algorithms, game theory, software agents

1. Introduction

Software agents are being used in a wide range of applications that include data allocation in large databases, monitoring electricity transformation networks, and vehicle routing among geographically dispersed dispatch centres (Sandholm 1983; Kraus 2001). More recently, software agents are being applied to electronic commerce for business-to-business, business-to-consumer, and consumer-to-consumer transactions (Maes et al. 1999; Sandholm 2000). The Internet and the world wide web provide the medium through which these transactions are carried out, and this has resulted in a large increase in the number of people buying and selling on the web. However, many of the steps of the buying process still remain unautomated. A human buyer is still responsible for collecting and interpreting information on merchants and products, making decisions about them, and, finally, entering purchase and payment information and instructions.

Software agents can be used to automate several of the most time consuming stages of the buying process (Ma 1999; Maes et al. 1999; Sandholm 2000). For example, a company that needs to order a product could have a buying agent that goes through all the stages of the buying process. The agent can be designed to automatically collect information on vendors and products that may fit the needs of the company, evaluate the various offers, make a decision on which merchants and products to investigate, negotiate the terms of transactions with these merchants, and finally place orders and make automated payments.

Of these, negotiation forms the key stage in all electronic transactions during which agents communicate and compromise to reach mutually beneficial agreements. The negotiators have a common interest in cooperating, but have conflicting interests over exactly *how* to cooperate. Put differently, agents can mutually benefit from reaching agreement on an outcome from a set of possible outcomes, but have conflicting interests over the outcome that they prefer. The main problem that confronts agents in such a situation is to decide how to cooperate before they actually enact the cooperation, and obtain the associated benefits. On the one hand, each agent would like to reach some agreement rather than disagree and not reach any agreement. But, on the other hand, each agent would like to reach an agreement that is as favourable to itself as possible. The agents therefore make a series of offers and counter-offers before an agreement is actually reached. This makes negotiation a time consuming process. Software agents not only save the labour time of human negotiators, but also find solutions that are as beneficial as possible to all the parties (Sandholm 2000).

On the basis of the extent to which agents cooperate with others, they can be divided into two main types: cooperative agents and competitive agents. Agents are said to be cooperative if they work towards a common goal. On the other hand, competitive or self interested agents do not share a common goal, but work towards maximizing their own benefits, as is the situation in all e-commerce applications. This paper therefore focuses on approaches for designing automated negotiation systems that are comprised of self-interested agents. In the context of self-interested agents, the term *bargaining* is used to refer to bilateral negotiations.

The problem of negotiation has long been studied by social scientists, resulting in the development of three main approaches that describe the strategy a negotiator should choose in different situations. The first approach is the formal *game theoretic* approach (Nash 1950; Harsanyi 1956; Roth 1979; Osborne and Rubinstein 1990). Game theory offers tools designed to help us understand how decision-makers behave when they interact. The basic assumptions that underlie the theory are that decision-makers pursue well defined objectives in accordance with their self-interest (i.e., they are rational) and that they take into account their knowledge or expectations of other decision-makers' behaviour (i.e., they reason strategically). Under these assumptions, it provides strategy recommendations and solutions to games. As it is based on a formal theory, this approach can in principle be applied to software agents (Sandholm 2000; Jennings et al. 2001; Kraus 2001; Lomuscio et al. 2003).

The second approach, the *negotiation guides* approach, comprises informal theories that identify possible beneficial strategies for a negotiator (Fisher and Ury 1981; Raiffa 1982; Johnson 1993). Although this approach does not make such strong assumptions as the game theoretic approach, applying these methods to software agents is difficult since they are not based on a formal theory.

The third and more recent approach is the *evolutionary approach* (Andreoni and Miller 1995; Ellingsen 1997; Binmore et al. 1998; Anthony and Jennings 2002). This approach begins by dropping the assumption that players are rational. Instead of assuming that players always calculate the best strategy from a theoretical analysis of the game, players in the evolutionary approach learn how to play games through trial and error. They experiment with strategies, observe their payoffs, try other strategies and find their way to a strategy that works well. Such players are referred to as being *boundedly rational*, because they are not the perfect reasoners commonly assumed in classical game theory (Samuelson 1996).

In this paper we provide a comprehensive comparison of the game-theoretic and evolutionary approaches to bargaining by examining their assumptions, goals and limitations. We then study the implications of the differences between them from the perspective of the agent developer.

The remainder of the paper is structured as follows. sections 2 and 3 describe the game theoretic and evolutionary models respectively. In section 4 we compare these two approaches. Finally, section 5 ends with some conclusions.

2. The Game Theoretic Bargaining Model

We begin by giving a description of *games* and *solutions*. A game is a description of strategic interaction that includes the constraints on the actions that the players can take and the players' interests, but does not specify the actions that the players take (Osborne and Rubinstein 1994). A solution is a systematic description of the outcomes that may result in a game under the assumption that the players are *perfectly rational* and that this rationality is *common knowledge*. Generally speaking, game theory suggests reasonable solutions for games and examines their properties.

A game can be described in one of the two forms: normal form or extensive form. The normal form describes games of simultaneous offers. On the other hand, the extensive form describes games in which the players make a series of offers and counter-offers and corresponds more closely to most real life bargaining situations than games of simultaneous offers. The model of an extensive game specifies the orders of events; each player can consider his plan of action not only at the beginning of the game but also whenever he has to make a decision. The following elements constitute a bargaining model (Osborne and Rubinstein 1994):

1. The *bargaining protocol*.
2. The *bargaining strategies*.
3. The *information state* of agents.
4. The *bargaining equilibrium*.

2.1. *The bargaining protocol*

The *protocol* specifies the rules of encounter between the negotiation participants. That is, it defines the circumstances under which the interaction between agents takes place: what deals can be made and what sequences of offers are allowed. In general, agents must reach an agreement on the bargaining protocol to use before bargaining proper begins.The design of protocols for governing agent interactions is referred to as *mechanism design* (Rosenschein and Zlotkin 1994). The following is a list of the desirable attributes of a well designed mechanism.

1. *Stability*. No agent should have an incentive to deviate from the agreed upon strategies. In game theory, this is also known as the notion of strategies in *equilibrium* (Osborne and Rubinstein, 1994). The strategy that agents adopt forms a part of the negotiation mechanism. Once these strategies are determined, agents should not deviate from them. Otherwise agent developers would have the incentive to build agents with different manipulative strategies.
2. *Simplicity*. A well designed mechanism should make low computational demands on the agents, and require little communication overhead. This issue is related to both efficiency and stability. If the negotiation mechanism is simple, it increases the efficiency of the system, with fewer resources being used to carry out the negotiation itself. Similarly, with stable mechanisms, fewer resources need to be spent on outguessing the opponent, or trying to discover his optimal choices. Stability reveals the behaviour publicly, and agents have nothing better to do than just carry it out.
3. *Distribution*. The negotiation mechanism should not require a central decision-maker, since this causes a performance bottleneck, and can result in failure of the entire system due to the failure of a single node.

All these attributes together affect the agent designer's choice of a negotiation protocol.

2.2. *The bargaining strategies*

An agent's *strategy* is a specification of the sequence of actions (usually offers or responses) the agent plans to make during negotiation. At any instant, an agent's action depends on what the other players (including himself) have done at earlier points, which constitutes the history of negotiation. A strategy is therefore defined as a function from the history of negotiation to the agent's current action. There will usually be many strategies that are compatible with a particular protocol, each of which may produce a different outcome. For example, an agent could concede at the first round, or bargain very hard throughout negotiation until its timeout is reached (Pruitt 1981; Raiffa 1982; Faratin et al. 1998). It follows that the negotiation strategy an agent employs is crucial with respect to the outcome of negotiation. It should also be clear that strategies which perform well with certain protocols will not necessarily do so with others. The choice of a strategy to use is thus a function not just of the specifics of the negotiation scenario, but also the protocol in use. The strategy of a rational decision-maker always gives the optimal action (i.e., the one that maximizes its expected utility).

2.3. *The information state of agents*

An agent's *information state* describes the information it has about negotiation parameters such as the players' utility functions, their reserve values, their discounting factors, and their deadlines. The strategy selected by an agent depends on its information state. von Neumann and Morgenstern (1944) introduced the fundamental classification of games into those of *complete information* and those of *incomplete information*. The former category is basic. In these games the players are assumed to know all relevant information about the rules of the game and players' preferences that are represented by utility functions (Nash 1950).

The complete information assumption is limiting because uncertainty is endemic in most realistic applications. In the latter category, information may be lacking about a variety of factors in the bargaining problem. Each player may have some private information about his own situation that is unavailable to the other players, while having only probabilistic information about the private information about other players. Following Harsanyi (Harsanyi and Selten 1972), models of games of incomplete information proceed from the assumption that all players start with the same probability distribution on this private information and that these priors are *common knowledge*.[1] This is modelled by having the game begin with a probability distribution, known to all players. Thus players not only have priors over other players' private information, they also know what priors the other players have over their own private information. Strategic models of incomplete information thus include an extra level of detail, since they specify not only the actions and information available to the other players in the course of the game, but also their probability distributions and information prior to the start of the game.

Another important model of strategic bargaining is Rubenstein's infinite horizon alternating offers game (Rubinstein 1982). This takes the time preferences of bargainers into consideration, in the form of their discounting factors, but again assumes complete information. It was later extended in Rubinstein (1985) for bargaining with incomplete information about time preferences.

Other models of incomplete information include (Fudenberg and Tirole 1983; Fudenberg et al. 1985; Sandholm and Vulkan 1999;

Fatima et al. 2001). Fudenberg and Tirole (1983) analyse an infinite horizon bargaining game by taking the players' valuations and a probability distribution over them as common knowledge. Fudenberg et al. (1985) analyse buyer–seller infinite horizon bargaining games in which reserve prices are uncertain, but time preferences are known. Sandholm and Vulkan (1999) consider uncertainty over agent deadlines. A common feature of all these models is that they treat the information state of agents as common knowledge. Fatima et al. (200b, 2004) address uncertainty over two parameters, deadlines and reserve prices, by treating each agent's information as its private knowledge. Each of these models is formulated for a different environment and the strategic behavior of rational agents is studied. The equilibrium of a game depends on the players' information states, and changing the information state of a player results in a change in the equilibrium of the game (Fatima et al. 2002a).

2.4. *The bargaining equilibrium*

Perhaps the most crucial element of a bargaining model is its equilibrium. As mentioned earlier in this section, this is what makes a negotiation mechanism stable. The earliest concept of equilibrium was the Nash equilibrium. This was defined for games of simultaneous offers Nash 1950. Two strategies are in Nash equilibrium if each agent's strategy is the best response to its opponent's strategy. This is a necessary condition for system stability where both agents act strategically. For sequential offer protocols, the Nash equilibrium concept was strengthened in multiple ways by requiring that the strategies stay in equilibrium at every step of the game (van Damme 1983). In summary, rationality, as understood in game theory, requires that each agent will select an equilibrium strategy when choosing independently.

Finally, game theory not only provides the concept of equilibrium, but also studies the properties of the equilibrium solution. The following are desirable solution properties (Osborne and Rubinstein 1994; Rosenschein and Zlotkin 1994).

1. *Individual Rationality*. For both the players, the agreement must represent a situation at least as favourable as the conflict situation (i.e., no agreement). Otherwise the players would have no incentive to negotiate.
2. *Efficiency*. An agreement is efficient if there is no wasted utility, i.e., it is not possible to increase the utility of one of the agents

without decreasing the utility of other agents. Such an outcome is said to be Pareto-optimal.
3. *Uniqueness*. Uniqueness is a desirable property because it allows the solution to be identified unequivocally.
4. *Symmetry*. A bargaining mechanism is said to generate a symmetric solution if it does not treat the players differently on the basis of inappropriate criteria. Exactly what constitutes inappropriate criteria depends on the specific domain. For instance, a bargaining mechanism is said to possess the property of symmetry if the outcome does not depend on the identity of the first player, i.e., which player starts the process of negotiation.

3. The Evolutionary Model

While a perfectly rational player determines the equilibrium strategy from a theoretical analysis of the game, a boundedly rational player finds the most effective strategy by learning through trial and error. Evolutionary or genetic algorithms (GAs) provide a metaphor that can be used for economic learning (Reichmann 1999). GA learning is now being widely studied for applications that involve bargaining and markets (Andreoni and Miller 1995; Dawid 1996; Bullard and Duffy 1998; Anthony and Jennings 2002).

GAs are an abstraction of biological evolution, and provide a framework for learning and adaptation by moving from one population of genes to a new population by using *natural selection* together with the genetics-inspired operators of *crossover* and *mutation*. Before describing how learning is accomplished through these operators, we list the elements of an evolutionary model and then describe each element in detail. The following are the key elements of an evolutionary model (Goldberg 1989; Mitchell 2001).

1. A *population* of individuals
2. The *player interactions*
3. The set of *evolutionary operators*
4. The *stable state*

3.1. *Population of individuals*

An evolutionary model imagines a game being played not by a single set of players, but by large populations of players. These players are repeatedly, randomly matched to play the game. The players are analogous to genes, and strategies are the behaviours or characteristics with which these genes endow their hosts. The payoff of a gene is its ability to produce offspring. Those genes that lead to higher rates of reproduction will flourish in a population at the expense of those that lead to lower reproduction rates. Each agent is characterized by a strategy that it plays when it is matched against another player. As a game proceeds, each agent observes the payoff of its strategy. It also observes the payoffs and strategies of others, and has access to information concerning how others have played. In the light of these observations, it adjusts its strategies. These adjustments involve experiments with strategies that it has not tried, but are overall designed to switch away from strategies that give low payoffs to strategies that give high payoffs. The task for the evolutionary model is to study this process of dynamic strategy adjustment.

Each genetic individual within the genetic population is assigned a *fitness value*. The fitness of an individual gives information about his performance according to the problem to be solved. The fitness usually equals the value of a function, optimized by means of the genetic algorithm. For instance, in a bargaining scenario, the utility associated with a strategy represents the fitness value of the agent playing that strategy.

There are two types of population, viz., *symmetric* and *asymmetric*. The symmetric population is the simpler of the two and involves all the players having the same fitness function (Young 1993; Ellingsen 1997; Binmore et al. 1998). In the asymmetric population, players have different fitness functions. For example, in bargaining between a buyer and a seller, the utility functions of agents differ. Such scenarios are handled by having two separate subpopulations, one representing the buyer and the other representing the seller.

3.2. *Player interactions*

The kind of player interactions in an evolutionary model are different from the player interactions in a game theoretic model. In the game theoretic model, a single set of players interact with each other. In contrast, in every generation, players in the evolutionary model are repeatedly, randomly matched to play the game. The matched players then negotiate using some predefined protocol. The utility that results from these negotiations forms an individual's fitness value. Individuals for the next generation are then selected on the basis of their fitness values. This is done by means of the evolutionary operators.

3.3. *Evolutionary operators*

In the evolutionary model, the players have access to the strategies played by other individuals and the associated payoffs. In the light of this information, they learn to play the most effective strategy. Learning thus forms an essential element of the evolutionary model. GA learning is a way of *social* rather than *individual* learning. Social learning means learning from others, and is conceptually different from the conventional AI based learning, which involves learning by single, isolated agents, as typified by Lucas (1986) and Heinemann (2000). Learning in the evolutionary model is a compound of three basic types of learning, and is realised by means of three basic operators – *selection*, *crossover*, and *mutation*. The three different learning techniques that correspond to these three operators are: *learning by imitation* (selection/reproduction), *learning by communication* (crossover) and *learning by experimentation* (mutation) (Reichmann 1999). Each of these is explained below.

Reproduction is a means of deriving a new population from an old one. Reproduction is done by selecting particular individuals out of the pool of the old population. The assignment of fitness to each of the individuals is the crucial part of the learning process. It is the fitness of a strategy that determines whether it is reproduced. Selection is the process of passing on agents that have a high fitness value, relative to the other agents, to the next generation. This is analogous to learning by imitation, since the good strategies are passed on without any change to the next generation. Learning by pure imitation will lead to stability, but optimality will occur only by pure chance.

The second operator (i.e., crossover) is a recombination operator. Crossover randomly chooses two genetic individuals, called parents, from the population and creates an offspring by combining parts of the bit strings of the two parents. Crossover can be interpreted as a form of learning by communication. Two agents meet, talk to each other about their strategies, and adapt parts of each other's behaviour. A complete strategy that is produced by crossover may not be the same as any of the strategies in the initial population, but parts of the strategy are those that are already there in the initial population.

The third operator is mutation. Mutation randomly alters single bits of the bit string by which a genetic individual is coded. It can be viewed as learning by experimentation. While selection (imitation) and crossover (communication) can reproduce strategies already in use (at least partially) by other individuals, mutation (experimentation) is able to find strategies that have never been used before. Mutation is the only operation that introduces completely new strategies into the system (i.e., those that were not present in the initial population). The initial population evolves as a result of selection, crossover, and mutation, and reaches a stable state in which a large majority of individuals in the population play the most effective strategy.

3.4. *Stable state*

Stability is a situation in which – after some evolutionary changes – a state is established in which the population, as a whole, shows a uniform behaviour. The population is allowed to evolve until it reaches a stable state (i.e., almost all the individuals exhibit the same behaviour and play the same strategy). In this state, although new strategies may get generated through crossover and mutation, they cannot invade the population on account of their being inferior to the stable strategy played by a significant proportion of the population. The outcome generated by the stable strategy is the stable outcome.

Recall that the population in the evolutionary model can be either symmetric or asymmetric. If the population is symmetric, then no matter what strategies individuals play in the initial state, the three evolutionary operators always cause the population to evolve to a stable state in which the players play the most effective strategy. For instance, a study of the evolutionary stability in a symmetric population has been carried out in Young (1993), Ellingsen (1997) and Binmore et al. (1998).

On the other hand, if the population is asymmetric, and there are two subpopulations, then the evolution of strategies in each subpopulation depends on the evolution of strategies in the other subpopulation (i.e., the strategies coevolve). As noted in Binmore et al. (1998), the stability results that hold for the symmetric information scenario may not hold for an asymmetric population. Moreover, in such a scenario, the stable strategy of each subpopulation can depend on how the two subpopulations are initialised. In this context, Fatima et al.'s (2003) study the coevolution of negotiation strategies for bilateral negotiation. There are two subpopulations, one representing the buyer and the other representing the seller. The results of this study demonstrate that in some negotiation scenarios the stable state changes if the initial population changes, while in others the stable state is independent of how the two populations are initialized.

4. A Comparison of Game Theoretic and Evolutionary Models

The similarity between the two models is that agents in both game theoretic and evolutionary models behave as expected utility maximizers. While game theoretic agents show this behaviour by virtue of their self-interest and rationality, agents in the evolutionary model (although boundedly rational) show a similar behaviour on account of learning, since only the fittest (given that fitness is measured in terms of utility) agents survive and reach the stable state. The difference lies in how the agents maximize their utility. This is caused by the differences that underlie these approaches, and their main characteristics. These are discussed in detail below. Table 1 provides a summary of the comparison.

The first difference lies in the number of agents. In the game theoretic model a game is viewed as being played by a single set of players, while the evolutionary model views the game as being played by a large[2] population of players, so that the effect of an individual's actions on the entire population can be ignored.

Second, the two models differ in terms of their basic assumptions. The main assumption in the game theoretic model is that the players are perfectly rational, and that this rationality is common knowledge. On the basis of this assumption, agents find the equilibrium strategy from a theoretical analysis of the game, and always play the equilibrium strategy. Game theory need not apply to human behaviour, since in practice humans do not always behave as the theory suggests (Roth 1995). An agent's optimal actions may be quite different depending upon whether it is playing against a perfectly rational agent or an ordinary person. The outcome prescribed by game theory for the former may not be valid when playing against the latter. For instance, in a game of poker one plays differently against an inexperienced player than against an experienced one. If an experienced player plays against an inexperienced player, he can win the game by using a strategy that is the best response to the opponent's strategy and not by using the equilibrium strategy. Game theory thus cannot always be used as a guide to behaviour (Samuelson 1996).

The behaviour of agents that have bounded rationality can only be studied by dropping the assumption of perfect rationality. Instead of being perfectly rational agents, which always play the equilibrium strategy from a theoretical analysis of the game, the players in real-life scenarios must learn how to play games through trial and error. The players must be able to learn the effective strategy by experimenting

Table 1. A comparison of game theoretic and evolutionary models

	Criterion	Game theoretic model	Evolutionary model
1	Number of players	Single set	Large populations
2	Basic assumption about rationality	Perfect rationality	No assumption
3	Strategy set	Players have strategy sets from which they choose particular strategies	Society has the strategy set and individuals inherit or choose from it
4	Strategy	Defined in terms of the history of the play	Depends only on the genes and not the previous history
5	Outcome	Determined by equilibrium strategies	Determined by stable state
6	Player interactions	One shot games	Repeated random pairing of agents
7	Learning	No learning; equilibrium outcome is determined from a theoretical analysis of the game;	Players learn the best strategy through repeated interaction with other players
8	Sharing of knowledge	No shared knowledge	Players have access to strategies played by others and the associated payoffs which they use for learning

with strategies, observing their payoffs, and trying other strategies in the search space. The evolutionary algorithms make no assumption about the rationality of players. In this model, agents learn the best strategy by playing against a large population of agents and observing their payoffs as well as those of others. The central assumption in the evolutionary model is that the population is very large, so that the effect of the behaviour of a single individual on the entire population can be ignored.

The third difference lies in the way that agents select strategies. In the game theoretic model, each agent has a set of strategies from which it chooses a particular strategy. On the other hand, in the evolutionary model the entire population has a strategy set and individuals inherit or choose from it.

The fourth difference lies in the way that strategies are defined. While the game theoretic model defines a strategy in terms of the history of the play (i.e., the sequence of offers and counter offers), the evolutionary model defines a strategy in terms of the genes that constitute the strategy, and not in terms of previous history.

The fifth difference lies in the way in which an outcome is determined. In the game theoretic model, the outcome is determined by the equilibrium strategies, while in the evolutionary model it is determined by the stable state of the population.

The sixth difference lies in way in which players interact. In the game theoretic model, a single set of players engage in one shot or repeated games, but in the evolutionary model there is repeated random pairing of agents.

The seventh difference lies in whether the agents can learn strategies through trial and error. In the game theoretic model, agents do not learn, but instead determine the outcome of a game from a theoretical analysis, which is possible because of the perfect rationality assumption. On the other hand, agents in the evolutionary model learn the best strategy through repeated interaction with other players, and through experimentation using selection, crossover, and mutation as the operators.

The final difference lies in whether agents have a shared knowledge of the strategies played by the other agents and the associated payoffs. In the game theoretic model, agents have no such shared knowledge, but in the evolutionary model agents learn to play better strategies on the basis of this shared knowledge. We now study the implications of these differences to the agent developer. From the differences listed above, it can be seen that 8developing software agents using the game theoretic approach is easier than developing agents using the evolutionary approach. This is mainly due to the following reasons. First, the game theoretic model involves a single set of players. So if there are n players in a game, then the number of agents is also n. On the other hand, the evolutionary model involves large populations of players. If the population is symmetric, then there is a single large population of players. But if the population is asymmetric, then each player type is represented by a large population. Consequently, a small increase in the number of player types leads to a considerable increase in the population size.

Second, in the game theoretic model, a game is played just once. In contrast, players in the evolutionary model need to repeatedly interact with each other until the stable state is reached. As the population size increases, the time it takes for the population to stabilize also increases.

The third issue is associated with learning, and in particular, whether or not it can be done offline. If learning can be done offline, the best strategy can be determined statically, and agents can be coded with this strategy. The actual bargaining process becomes simpler, since it does not involve online learning. Whether learning in the evolutionary model can be done offline depends on the population. If the population is symmetric (i.e., all individuals have the same utility function), then the stable outcome does not depend on the initial state of the population. It is therefore common to initialise the population randomly (Goldberg 1989; Koza et al. 1991). This allows learning to be done offline and the agents can be encoded with the stable strategy.

On the other hand, when the population is asymmetric, it is composed of subpopulations with individuals in different subpopulations having different utility functions. In such situations, the subpopulations coevolve; the evolution of strategies in each subpopulation depends on the evolution of strategies in the others. Moreover, the stable outcome can depend on how the subpopulations are initialised (Fatima et al.'s 2003). Consequently, it may not be possible to determine the stable outcome offline. For instance, if we consider a simple bilateral negotiation scenario, we have two subpopulations – one representing the buyer and the other representing the seller. The buyer (seller) cannot determine the stable outcome without knowing how the seller (buyer) subpopulation is initialised. In such situations, learning can only be done online. In other words there need not be a single stable outcome (as in the case of symmetric population); changing the initialization of individuals in any population can result in a change in the stable outcome. Since software agents may be constructed by separate designers and/or may represent different real world parties, at the time of development, the software agent developer may not know the actual initialization of different populations. Consequently, the evolutionary algorithm needs to be run each time there is a negotiation. This results in excessive computation, communication, and time overheads, which can become prohibitive, particularly in applications that require negotiation to end by a deadline. For one-to-many and many-to-many negotiations, these overheads become higher as the number of player types increases. Thus in terms of the properties of the negotiation mechanism we see that the game theoretic model may be easier to implement than the evolutionary model.

5. Conclusions

A number of game-theoretic models have been explored that solve the bargaining problem. More recently, evolutionary approaches have been applied to the same problem. This paper describes these two approaches and provides a comparative study in terms of their underlying assumptions and their main characteristics, such as the requirement for common knowledge, the definition of strategy, the solution concepts, the kind of player interactions, the presence of learning, and whether it can be done offline. We then study the implications of these differences from the perspective of the software agent developer. More specifically, the discussion suggests the following main conclusions.

Although the game-theoretic model assumes perfect rationality, designing software agents is easier since the agents only need to play the equilibrium strategy. On the other hand, the evolutionary approach does not assume perfect rationality, but may require the software agents to learn the stable strategy every time there is a negotiation. If the stable outcome depends on the initial population, then learning cannot be done offline. On line learning results in excessive computation, communication, and time overheads that can become prohibitive, particularly in applications that require negotiation to end by a deadline. For one-to-many and many-to-many negotiations, these overheads become higher as the number of player types increases. Thus in terms of the properties of the negotiation mechanism we se that the game theoretic model may be easier to implement than the evolutionary model.

Notes

1. An event is said to be mutual knowledge in some state if in that state each player knows the event. An event is said to be common knowledge if not only is it mutual knowledge but also each player knows that all other players know it, each player knows that all other players know that all the players know it, and so on.
2. Yao and Darwen (2000) studies how the stability of an evolutionary model changes with a change in population size.

References

Andreoni, J. & Miller, J. H. (1995). Auctions with Artificial Adaptive Agents. *Games and Economic Behavior* **10**: 39–64.

Anthony, P. Jennings, N. R. (2002). Evolving Bidding Strategies for Multiple Auctions. In *Proceedings of the 15th European Conferenceon Artificial Intelligence (ECAI-2002)* 178–182.

Binmore, K., Piccione, M. & Samuelson, L. (1998). Evolutionary Stability in Alternating-Offers bargaining Games. *Journal of Economic Theory* **80**: 257–291.

Bullard, J. Duffy, J. (1998). A Model of Learning and Emulation with Artificial Adaptive Agents. *Journal of Economic Dynamics and Control* **22**: 179–207.

Dawid, H. (1996). *Adaptive Learning by Genetic Algorithms*: Springer. Berlin: Heidelberg, New York.

Ellingsen, T. (1997). The Evolution of Bargaining Behavior. *The Quarterly Journal of Economics* 581–602.

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

Fatima, S. S., Wooldridge, M. & Jennings, N. R. (2001). Optimal Negotiation Strategies for Agents with Incomplete Information. In Meyer J. J. & Tambe, M. (eds.) *Intelligent Agent VIII* Springer Lecture Notes in AI, 377–392.

Fatima, S. S., Wooldridge, M. & Jennings, N. R. (2002a). The Influence of Information on Negotiation Equilibrium. In Padget, J., & Shehory, O. Parkes, D., Sadeh, N. & Walsh, W. E. (eds.) *Agent Mediated Electronic Commerce IV - Designing Mechanisms and Systems*, 180–193, Springer: Berlin.

Fatima, S. S., Wooldridge, M. & Jennings, N. R. (2002b). Multi-Issue Negotiation Under Time Constraints. In Castelfranchi, C. & Johnson, W. L. (eds.) *International Conference on AutonomousAgents and Multi-Agent Systems*. Bologna, Italy, 143–150.

Fatima, S. S., Wooldridge, M. & Jennings, N. R. (2003). Comparing Equilibria for Game Theoretic and Evolutionary Bargaining Models. In proceedings of the *Fifth International Workshop on Agent Mediated Electronic Commerce (AMEC-03)*. Melbourne, Australia, 70–77.

Fatima, S. S., Wooldridge, M. & Jennings, N. R. (2004). An Agenda Based Framework for Multi-Issue Negotiation. *Artificial Intelligence Journal* **152**(1): 1–45.

Fisher, R. & Ury, W. (1981). *Getting to Yes: Negotiating Agreement Without Giving in*. Houghton Mifflin, Boston.

Fudenberg, D., Levine, D. & Tirole, J. (1985). Infinite Horizon Models of Bargaining with One Sided Incomplete Information. In Roth A. (ed.) *Game Theoretic Models of Bargaining*. University of Cambridge Press: Cambridge.

Fudenberg, D. & Tirole, J. (1983). Sequential Bargaining with Incomplete Information. *Review of Economic Studies* **50**: 221–247.

Goldberg, D. E. (1989). *Genetic Algorithms.* Addison-Wesley, Newyork.

Harsanyi, J. & Selten, R. (1972). A Generalized Nash Solution for Two-Person Bargaining Games with Incomplete Information. *Management Science* **18**(5): 80–106.

Harsanyi, J. C. (1956). Approaches to the Bargaining Problem before and After the Theory of Games. *Econometrica* **24**: 144–157.

Heinemann, M. (2000). Adaptive Learning of Rational Expectations Using Neural Networks. *Journal of Economic Dynamics and Control* **24**, 1007–1026.

Jennings, N. R., Faratin, P., Lomuscio, A. R., Parsons, S. & Wooldridge, M. (2001). Automated Negotiation: Prospects Methods and challenges. *Group Decision and Negotiation* **10**(2): 199–215.

Johnson, R. (1993). *Negotiation Basics* Sage: Newbury Park.

Koza, J. R., Bennett, F. H., Andre, D. & Keane, M. A. (1999). *Genetic Programming III Darwinian Invention andProblem Solving*. Morgan Kaufmann Publishers.

Kraus, S. (2001). *Strategic negotiation in multi-agentenvironments* The MIT Press: Cambridge, MA.

Lomuscio, A., Wooldridge, M. & Jennings, N. R. (2003). A Classification Scheme for Negotiation in Electronic Commerce. *International Journal of Group Decision and Negotiation* **12**(1): 31–56.

Lucas, R. E. (1986). Adaptive Behavior and Economic Theory. *Journal of Business* **59**: S401–S426.

Ma, M. (1999). Agents in E-commerce. *Communications of the ACM* **42**: 79–80.

Maes, P., Guttman, R. & Moukas, A. (1999). Agents that Buy and Sell. *Communications of the ACM* **42**(3): 81–91.

Mitchell, M. (2001). *An Introduction to Genetic Algorithms*. The MIT Press, Cambridge, MA.

Nash, J. F. (1950). The Bargaining Problem. *Econometrica* **18**: 155–162.

Osborne, M. J & Rubinstein, A. (1990). *Bargaining and Markets*. Academic Press: San Diego California.

Osborne, M. J.: Rubinstein, A. (1994). *A Course in Game Theory*. The MIT Press: Cambridge.

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

Raiffa, H. (1982). *The Art and Science of Negotiation*. Harvard: University Press Cambridge, USA.

Reichmann, T. (1999). Learning and Behavioral Stability – An economic Interpretation of Genetic Algorithms. *Journal of Evolutionary Economics* **9**: 225–242.

Rosenschein, J. S. & Zlotkin, G. (1994). *Rules of Encounter*. MIT Press:Cambridge.

Roth, A. (1979). *Axiomatic Models of Bargaining*. Springer-Verlag Berlin.

Roth, A. E. (1995). Bargaining Experiments. In Kagel, J. & Roth, A. E. (eds.) *Handbook of Experimental Economics*. Princeton University Press: Princeton, NJ.

Rubinstein, A. (1982). Perfect Equilibrium in a Bargaining Model. *Econometrica* **50**(1): 97–109.

Rubinstein, A. (1985). A Bargaining Model with Incomplete Information About Time Preferences. *Econometrica* **53**: 1151–1172.

Samuelson, L. (1996). Bounded Rationality and Game Theory. *The Quarterly Review of Economics and Finance* **36**: 17–35.

Sandholm, T. (1983). An Implementation of the Contract Net Protocol Based on Marginal Cost Calculations. In *Proceedings of the AAAI-93* Washington, DC, 256–262.

Sandholm, T. (2000). Agents in Electronic Commerce: Component Technologies for Automated Negotiation and Coalition Formation. *Autonomous Agents and Multi-Agent Systems* **3**(1): 73–96.

Sandholm, T. & Vulkan, N. (1999). Bargaining with Deadlines. In *Procceding of the AAAI-99*. Orlando, FL, 44–51.

van Damme, E. (1983). *Refinements of the Nash equilibrium concept*. Springer-Verlag: Berlin.

von Neumann, J. & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press: Princeton NJ.

Yao, X. & Darwen, P. (2000). Genetic Algorithms and Evolutionary Games. In *Commerce, Complexity and Evolution*. Cambridge University Press: Cambridge, 313–333.

Young, H. P. (1993). An Evolutionary Model of Bargaining. *Journal of Economic Theory* **59**: 145–168.