Multi-Issue Negotiation Processes by Evolutionary Simulation: Validation and Social Extensions*

E.H. Gerding[†] D.D.B. van Bragt J.A. La Poutré
CWI, Centre for Mathematics and Computer Science
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

Abstract

We describe a system for automated bilateral negotiations in which artificial agents are evolved by an evolutionary algorithm. The negotiations are governed by a finite-horizon version of the alternating-offers protocol. Several issues are negotiated simultaneously and negotiations can be broken off with a pre-defined probability.

In our experiments the bargaining agents have different preferences regarding the importance of the issues, which enables mutually beneficial outcomes. These optimal solutions are indeed discovered by the evolving agents. To further validate our system, the computational results are also compared to game-theoretic (subgame perfect equilibrium) predictions. The influence of important model settings, like the probability of breakdown or the length of the game, is investigated in detail in this validation part.

We also present an extension of the evolutionary system in which the agents use a "fairness" norm in the negotiations. This concept plays an important role in real-life negotiations. Fairness is implemented by re-evaluating a reached agreement and rejecting unfair agreements with a certain probability. In our model re-evaluation can take place in each round or only if the deadline of the negotiations is reached. When fairness is applied in each round, the agents reach equal utility levels. When the fairness of deals is only evaluated in the final round, both symmetric and asymmetric outcomes can occur, depending on the fairness model that is used.

1 Introduction

Lately, automated negotiations receive more and more attention, especially from the field of electronic trading [4, 7, 8, 11]. In the near future, an increasing use of bargaining agents in electronic market places is expected. Ideally, these agents should not only bargain over the price of a product, but also take into account aspects like the delivery time, quality, payment methods, return policies, or specific product properties. In such multi-issue negotiations, the agents should be able to negotiate outcomes that are mutually beneficial for both parties. The complexity of the bargaining problem increases rapidly, however, if the number of issues becomes larger than one. This explains the need for "intelligent" agents, which should be capable of negotiating successfully over multiple issues at the same time.

*Paper submitted to the Workshop on *Complex Behavior in Economics* at Aix-en-Provence, France, May 4-6, 2000. This research was part of the project "Autonomous Systems of Trade Agents in E-Commerce", funded by the Telematics Institute in the Netherlands.

[†]Corresponding author. E-mail: e.gerding@cwi.nl.

Negotiations are governed by a finite-stage version of Rubinstein's alternating-offers game [14] in this paper. The learning process of the bargaining agents is modelled with an evolutionary algorithm (EA). The strategies of the agents are updated in successive generations by this EA. EAs are powerful search algorithms (based on Darwin's evolution theory) which can be used to model social learning in societies of boundedly-rational agents [5, 12]. In an evolutionary setting, the adaptive agents typically learn in three different ways: (1) learning by imitation, (2) communication (exchange of strategic information), and (3) experimentation. It is important to note that EAs make no explicit assumptions of rationality. Basically, the fitness of the individual agents is used to determine whether a strategy will be used in future situations.

A small, but growing, body of literature already exists in this field. Oliver [11] was the first to demonstrate that, using an EA, artificial agents can learn effective negotiation strategies. In Oliver's model, the agents use quite elementary bargaining strategies. A more elaborate strategy representation is proposed and evaluated in [8]. Offers and counter offers are generated in this model by a linear combination of simple bargaining tactics (time-dependent, resource-dependent, or behaviour-dependent tactics). A recent, more fundamental, study is [15]. In [15], a systematic comparison between game-theoretic and evolutionary bargaining models is made.

In this paper we assess to what extent the behaviour of the adapting agents matches with game-theoretic models if multiple issues are involved. This work is therefore in line with the work reported in [15], where bargaining only concerned a single issue. We study models in which time plays no role, and models in which there is a pressure to reach agreements early (because a risk of breakdown in negotiations exists after each round).

When no time pressure is present, an extreme partitioning of the bargaining surplus occurs in the computer experiments (in agreement with game-theoretic predictions). Such extreme outcomes are not observed in real-life situations, where social norms such as fairness play an important role [3, 6, 16, 13]. We therefore incorporate a fairness measure in our evolutionary system. Our fairness models can be tuned from "weak fairness" (i.e., accept almost all agreements) to "strong fairness" (i.e., reject unfair deals). Also, the agents can apply fairness in each round or only if the deadline is reached. Results are sensitive to both aspects, but in general fair deals evolve more frequently if the agents evaluate the fairness of reached agreements in each round.

This evolutionary model is a first attempt to study complex bargaining situations which are more likely to occur in practical settings. A rigorous game-theoretic analysis is typically much more involved or even intractable under these conditions.

The remainder of this paper is organised as follows. Section 2 gives an outline of the setup of the computer experiments. A comparison of the computational results with game-theoretic predictions is presented in Section 3. Fairness is the topic of Section 4. Section 5 summarises the main results and concludes.

2 Experimental Setup

This section gives an overview of the setup of the computational experiments. The alternating-offers negotiation protocol is described in Section 2.1. Section 2.1 also discusses how the agents evaluate the outcome of the bargaining process using a multi-attribute utility function. Section 2.2 then describes the (genetic) representation of the agents' strategies and the EA which updates these strategies in successive generations.

2.1 Negotiation Protocol

During the negotiation process, the agents exchange offers and counter offers in an alternating fashion. In the following, the agent starting the negotiations is called "agent 1", whereas his opponent is called "agent 2".

Bargaining takes place over multiple issues simultaneously. An offer can then be denoted as a vector $\vec{o}$. The i-th component of this vector, denoted as o^i, specifies the share of issue i that agent 1 receives if the offer is accepted. We assume (without loss of generality) that the total bargaining surplus available per issue is equal to unity. Agent 2 then receives $1 - o^i$ for issue i in case of an agreement. The index i ranges from 1 to m (the total number of issues).

As stated above, agent 1 makes the initial offer. If agent 2 accepts this offer, an agreement is reached and the negotiations stop. Otherwise, play continues with a certain continuation probability p ($0 \leq p \leq 1$). When a negotiation is broken off prematurely both agents receive nothing. Such a breakdown in negotiations may occur in reality when agents get dissatisfied as negotiations take too long, and therefore walk away from the negotiation table, or when intervention of a third party results in a vanishing bargaining surplus.

If negotiations proceed to the next round, agent 2 needs to propose a counter offer, which agent 1 can then either accept or refuse. This process of alternating bidding continues for a limited number of n rounds. When this deadline is reached the negotiations end up in a disagreement and both players get nothing.

The offers and counter offers are specified by the agents' strategies. In a game-theoretic context, a strategy is a plan which specifies an action for each history [2]. In our model, a strategy specifies the offers and thresholds for each round in the negotiation process. A threshold specifies whether an offer should be accepted or rejected: If the value of the offer falls below the threshold the offer is refused; otherwise an agreement is reached[1].

The agents evaluate the offers of their opponents using an additive multi-attribute utility function [8, 11]. We assume that all agents are risk neutral. Agent 1's utility function u_1 is then equal to $\vec{w}_1 \cdot \vec{o} = \sum_{i=1}^{m} w_1^i \cdot o^i$. Agent 2's utility function u_2 is equal to $\vec{w}_2 \cdot (\vec{1} - \vec{o})$. $\vec{w}_j$ is a vector containing agent j's weights w_j^i for each issue i. The weights are normalised and larger than zero, i.e., $\sum_{i=1}^{m} w_j^i = 1$ and $w_j^i > 0$. Because we assume that $0 \leq o_i \leq 1$ for all i, $0 \leq u_j(\vec{o}) \leq 1$.

2.2 The Evolutionary System

We use an EA to evolve the negotiation strategies of the agents. This section discusses how the EA has been implemented, and how the system can be interpreted as a model for social or economic learning processes. The implementation is based on "evolution strategies"(ES), a branch of evolutionary computation that traditionally focusses on real-coded problems [1][2].

The different stages within an iteration of the evolutionary algorithm are depicted in Fig. 1. The system initially starts with two separate and randomly initialised "parental" populations. Agents in population 1 start the bargaining process (i.e., they are of the "agent 1" type). The fitness of the parental agents is determined by competition between the agents in the two

[1]The same approach was used in [11, 15].

[2]The widely-used genetic algorithms (GAs) are more tailored toward binary-coded search spaces [9].

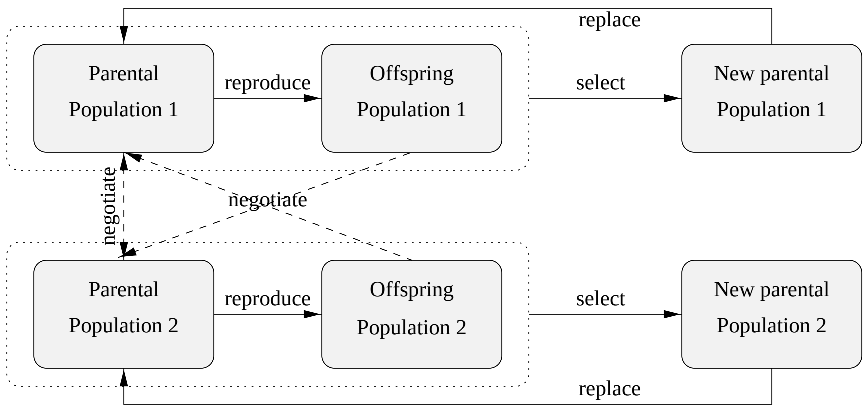

Figure 1: Iteration loop of the evolutionary algorithm. Two populations of agents, containing agents with different weight vectors, are evolved separately. Agents in population 1 always start the bargaining process. In the fitness evaluation, both the offspring and the parental agents compete against agents in the two parental populations. The best candidates of the union of parents and offspring are then selected to be the parents in the next iteration.

populations. Each agent competes against all agents in the other population. The average utility obtained in these bilateral negotiations is then used as the agent's fitness value.

In the next stage (see Fig. 1), "offspring" agents are created. An offspring agent is generated in two steps. First, an agent in the parental population is (randomly, with replacement) selected. This agent's strategy is then mutated to create a new offspring agent (the mutation model is specified below). The fitness of the new offspring is evaluated by interaction with the parental agents[3]. A social or economic interpretation of this parent-offspring interaction is that new agents can only be evaluated by competing against existing or "proven" strategies.

In the final stage of the iteration (see Fig. 1), the fittest agents are selected as the new "parents" for the next iteration (the selection procedure is explained in detail below). This final step completes one iteration (or "generation") of the EA. All relevant settings of the evolutionary system are listed in Table 1.

EA Parameters	Parental population size (μ)	25
	Offspring population size (λ)	25
	Selection scheme	$(\mu + \lambda)$-ES
	Mutation model	self-adaptive
	Initial standard deviations ($\sigma_i(0)$)	0.1
	Minimum standard deviation (ϵ_σ)	0.025
Negotiation parameters	Number of issues (m)	2
	Weights of agents in population 1 ($\vec{w_1}$)	$(0.7, 0.3)^T$
	Weights of agents in population 2 ($\vec{w_2}$)	$(0.3, 0.7)^T$

Table 1: Default settings of the evolutionary system.

[3]In an alternative model, not only the parental agents are used as opponents, but also the newly-formed offspring. This leads to a much more diverse collection of opponents. The fitness of the agents therefore becomes more subject to noise. Nevertheless, similar dynamics have been observed in this alternative model.

Selection model. Selection is performed using the $(\mu + \lambda)$-ES selection scheme [1]. In conventional notation, μ is the number of parents and λ is the number of generated offspring (μ=λ=25, see Table 1). The μ survivors with the highest fitness are selected (deterministically) from the union of parental and offspring agents. The $(\mu + \lambda)$-ES selection scheme is an example of an "overlapping generations" model, in which successful agents can survive for multiple generations. In an economic context, selection can be interpreted as imitation of behaviour which seems promising. In general, EAs use two additional operators: mutation and recombination. These operators are explained in detail below.

Mutation and recombination model. Mutation operates directly on the "chromosome" of an agent. The chromosome specifies the strategy an agent uses in the bargaining game. An agent's chromosome consists of a collection of "genes". These genes contain the values for the offers and thresholds (per round). In multi-issue negotiations, a sequence of m genes specifies an offer. Threshold values are represented by a single gene. Each gene is real-valued with a range between 0 and 1. A similar strategy representation was used in [11, 15]. Oliver [11], however, used binary-coded chromosomes.

The offspring's genes x_i are created by adding a zero-mean Gaussian variable with a standard deviation σ_i to each corresponding gene of the parent [1, 15]. All offspring genes with a value larger than unity (or smaller than zero) are set equal to unity (respectively zero). In our simulations, we use an elegant mutation model with self-adaptive control of the standard deviations σ_i [1, pp. 71-73][15]. At the beginning of each EA run, the standard deviations σ_i are set equal to $\sigma_i(0) = 0.1$ (see Table 1). To prevent complete convergence of the population, we force all standard deviations to remain larger than a small value $\epsilon_\sigma = 0.025$ (see Table 1).

Mutation can be interpreted as undirected exploration of new strategies, or as mistakes made during imitation. Communication between the agents is often modelled by a recombination operator, which typically exchanges parts of the parental chromosomes to produce new offspring. Earlier experiments [15] showed little effect on the results when traditional recombination operators from ES (like discrete or intermediate recombination [1]) were applied. We therefore focus on mutation-based models in this paper.

3 Validation and Interpretation of the Evolutionary Experiments

Experimental results obtained with the evolutionary system are presented in this section. A comparison with game-theoretic predictions is made to validate the evolutionary approach. Section 3.1 discusses the ability of the evolutionary system to (a) avoid the occurrence of disagreements and (b) to discover agreements in the neighbourhood of the Pareto-efficient frontier. Section 3.2 compares, for different settings, the (mean) long-term behaviour of the evolving agents with game-theoretic (subgame perfect equilibrium) predictions.

3.1 The Evolution of Pareto-Efficient Agreements

First, we investigate whether the adaptive agents learn to avoid the occurrence of disagreements. If we set the continuation probability p equal to 1, and the number of rounds n equal to 10, disagreements can only occur when the deadline is reached (i.e., after 10 rounds). The computer experiments show that the percentage of disagreements is very small (around 0.1%). This is somewhat surprising, because in the long run almost all agreements are reached in the very last round (after 1000 generations, about 80% of all agreements are reached just before the deadline). Furthermore, the last agent in turn demands almost the entire surplus (for each issue). Nevertheless, his opponent accepts this extreme take-it-or-leave-it deal. This behaviour agrees with game-theoretic predictions (see Section A.1).

Next, we study a model with a risk of breakdown in the negotiations ($p = 0.7$). Initially, the percentage of disagreements is approximately equal to 23%. This percentage rapidly decreases to a value between 1% and 10%. The number of disagreements decreases because in the long run most agreements are reached in the first round (after 1000 generations, approximately 75% of the agreements are reached immediately). Again, this behaviour is consistent with game-theoretic predictions (see Section A.2).

Figure 2 maps the agreements reached in the evolutionary system (for the same settings, i.e., $p = 0.7$ and $n = 10$) onto a two-dimensional plane. Each point in this plane shows the

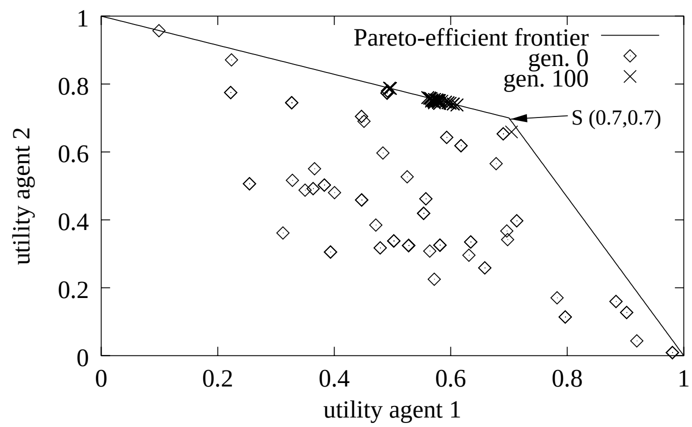

Figure 2: Agreements reached by the evolving agents at the start of a typical EA run (generation 0), and after 100 generations (for $p = 0.7$ and $n = 10$). Each agreement is indicated by a point in this two-dimensional plane. The Pareto-efficient frontier is indicated with a solid line. In point S [at $(0.7, 0.7)$] both agents obtain the maximum share for their most important issue, and receive nothing for the other issue. Note that, after 100 generations, almost all agreements are located in the vicinity of the Pareto-efficient frontier.

utility for both agents which results from an agreement. Agreements can never be located above the so-called "Pareto-efficient frontier" in this plane (indicated by the solid line). An agreement is located on the Pareto-efficient frontier when an increase of utility for one agent necessarily results in a decrease of utility for the other agent. A special point on the Pareto-efficient frontier is S. In this symmetric point [at $(0.7, 0.7)$] both agents obtain the maximum share of the issue they value the most, and receive nothing of the less important issue. Initially, at generation 0, many agreements are located far from the Pareto-efficient frontier. After 100 generations, however, the agents have learned to coordinate their behaviour and most agreements are Pareto-efficient. It is important to note that, even in the long run, the agents keep exploring the search space. This results in continuing movements of the "cloud" of agreements (visible in Fig. 2) along the Pareto-efficient frontier. The (mean) long-term behaviour of the evolving agents is studied in more detail in the next section.

Results in this section already show that the adaptive agents learn to avoid the occurrence of disagreements and, furthermore, are able to coordinate their behaviour in an efficient way.

3.2 Comparison of Long-Term Behaviour with Game-Theoretic Predictions

The computational results are compared in more detail with game-theoretic predictions in this section. The key equilibrium concept used by game theorists to analyze extensive-form games[4] is the subgame perfect equilibrium (SPE) [14, 2]. Two strategies are in SPE if they constitute a Nash equilibrium in any subgame which remains after an arbitrary sequence of offers and replies made from the beginning of the game. Rubinstein successfully applied this notion of subgame-perfection to bargaining games [14]. His main theorem states that the infinite-horizon alternating-offers game has a unique SPE in which the agents agree immediately on a deal (if time is valuable). Our experimental setup differs in two respects from Rubinstein's model (see Section 2.1). First, we use a finite-length instead of an infinite-length bargaining game. Second, the agents bargain over multiple issues instead of a single issue. This changes the game-theoretical analysis in some respects, as we show in the Appendix.

It is important to note that we assume in the game-theoretic analysis of the Appendix that the bargaining agents behave fully rational and have complete information (for instance about the importance of the different issues for their opponent). Both assumptions are obviously not valid for the evolving agents in our computational experiments (who learn by trial-and-error instead of abstract reasoning). The SPE behaviour of fully rational agents will nevertheless serve as a useful theoretical benchmark to interpret the behaviour of the boundedly-rational agents in our experiments.

As we mentioned before in Section 3.1, it is optimal (i.e., subgame-perfect) to propose a take-it-or-leave-it deal in the last round if $p = 1$. The logic of subgame-perfection then requires that the responder in the final round accepts this extreme deal (see Section A.1). Hence, we expect the fitness of agents in population 1 to converge to unity if n is odd, and to converge to zero if n is even (the opposite holds for the agents in population 2). This tendency is indeed clearly visible in Fig. 3a, even for games as long as 10 rounds. Figure 3b shows that game theory predicts that the influence of the finite length of the game diminishes for longer games if $p < 1$. Notice for instance in Fig. 3b that the SPE partitioning is quite asymmetric for small n, but more symmetric for larger n. This effect is actually much stronger in the evolutionary system (see Fig. 3b). The evolving agents do not reason backwards from the deadline (as is done in game theory, see Section A.2), but focus on the first few rounds, where expected utility is relatively high. This means that only few agreements are reached in later rounds. As a result, the deadline is not perceived accurately by the evolving agents.

It is interesting to note that, in the limit of $n \to \infty$, game theory predicts that the agents in population 1 have a fitness of ≈ 0.71, whereas the agents in population 2 have a fitness of ≈ 0.68. This corresponds to a point in the vicinity of the symmetric point S indicated in Fig. 2. The computational experiments reported in Fig. 3b show that the behaviour of the agents is in fact much better predicted by an infinite-horizon model than the finite-horizon model for $n \geq 8$. This supports our previous claim that the boundedly-rational agents do not accurately perceive and exploit the finite deadline of the game. In fact, they strongly overestimate the game length (see [15]).

We compared the computational results with predictions made by game theory in this section. Game-theoretic (SPE) predictions appear to be very useful to interpret the behaviour of the adaptive agents in the evolutionary simulations. We investigated the influence of the finite length of the game. In games without a risk of breakdown, the agents successfully exploit

[4] That is, games with a tree structure.

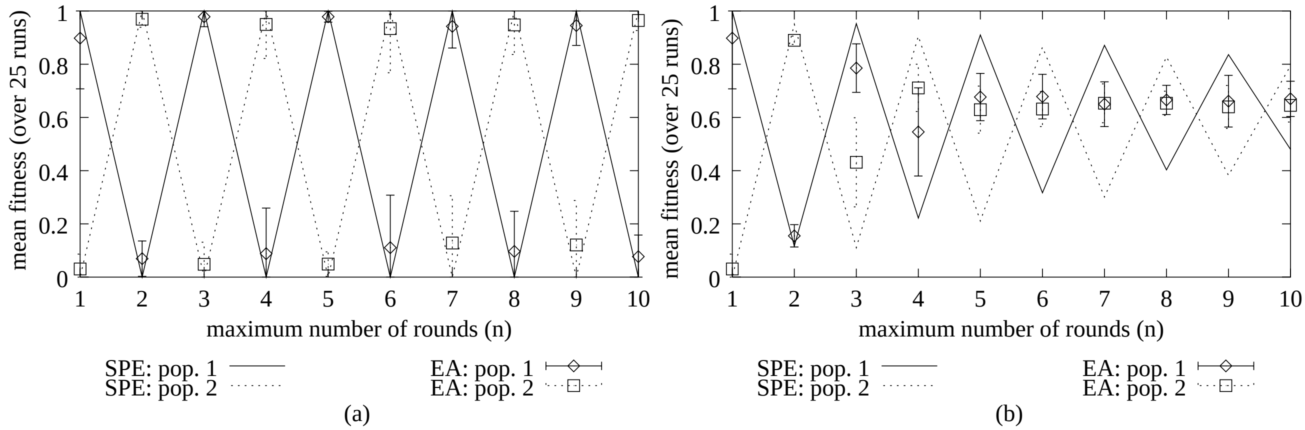

Figure 3: Comparison of the evolutionary results with SPE predictions for (a) $p = 1$ (time indifference) and (b) $p = 0.95$. The SPE predictions for successive values of n are connected to guide the eye. Figure 3a shows that the finite length of the game has a strong impact on the long-term behaviour if $p = 1$. If $p < 1$, see Fig. 3b, the finite character of the game is not fully exploited by the boundedly-rational agents in our computational experiments. We measured the mean fitness of agents in populations 1 and 2 after the initial transients have died out. The error bars indicate the standard deviations across 25 runs.

their last-mover advantage. In games with a risk of breakdown, on the other hand, this last-mover advantage is smaller than predicted by game theory. Moreover, if the game becomes long enough, the finite deadline is no longer perceived by the evolving agents. In that case, their behaviour actually agrees better with game-theoretic predictions for the infinite-horizon game.

4 Social Extensions: Fairness

Game-theoretical models for rational agents often predict very asymmetric outcomes for the two parties. We showed in Section 3.2, especially in Fig. 3a, that such "unfair" behaviour can also emerge in a system of adaptive agents. Game-theoretic (SPE) predictions have also been compared to human behaviour in laboratory experiments in the past (see [13] for an extensive overview). Large discrepancies between human behaviour and SPE predictions were found in these experimental studies, both for ultimatum and multi-stage games [3, 6, 16].

A possible explanation for the occurrence of these discrepancies is the strong influence of social or cultural norms on the individual decision making process. Binmore et al. [3] argued, for example, that the agents in the strongest bargaining position demand less than the SPE amount because they value the "fairness" of the outcomes. In a recent paper, Lin and Sunder [6] propose a model in which the probability of acceptance of a proposal increases with the amount offered to the responder. Such a model, making more realistic assumptions about the agents' behaviour, appears to organise the data from experiments with humans better than the SPE model [6].

We implemented Lin and Sunder's model in our evolutionary system. In this extended model, the negotiation protocol is as described in Section 2.1 with one exception: If an agreement is reached, the responder has the opportunity to re-evaluate his decision. The probability that he finally accepts the agreement then depends on the utility he acquires. If the agreement is rejected, the game continues (unless the deadline has been reached). Several "fairness" functions, determining the probability of acceptance, are shown in Fig. 4. We use

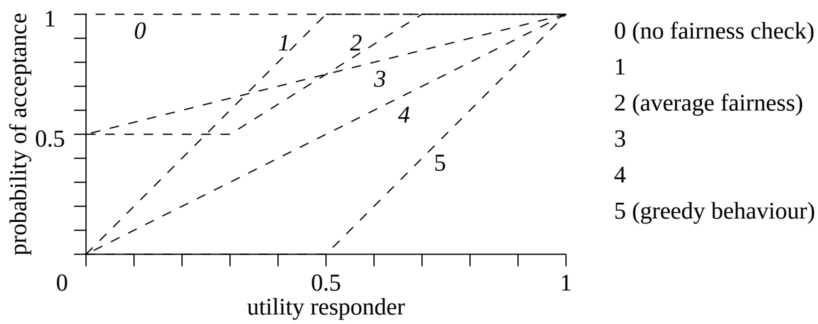

Figure 4: Different fairness functions used by the agents in the evolutionary experiments to determine the degree of fairness of agreements. Each model specifies the probability that the responding agent accepts an agreement as a function of the acquired utility. The models range from a model without a fairness check (model no. 0) to a model (no. 5) in which even fair agreements are rejected with a high probability.

piece-wise linear functions with up to three segments[5].

We observed in the former experiments without a risk of breakdown ($p = 1$) (see Section 3.2) that, in the long run, almost all reached agreements are highly asymmetric (the last agent in turn demands and receives almost the whole bargaining surplus). In these experiments almost all agreements are delayed until the very last round. It is therefore interesting to investigate the system's behaviour if the agent in the weakest bargaining position rejects this take-it-or-leave-it deal with a certain probability. The effect of such a fairness model, with only a single fairness check in the last round, is studied in Fig. 5a. Figure 5a shows (for

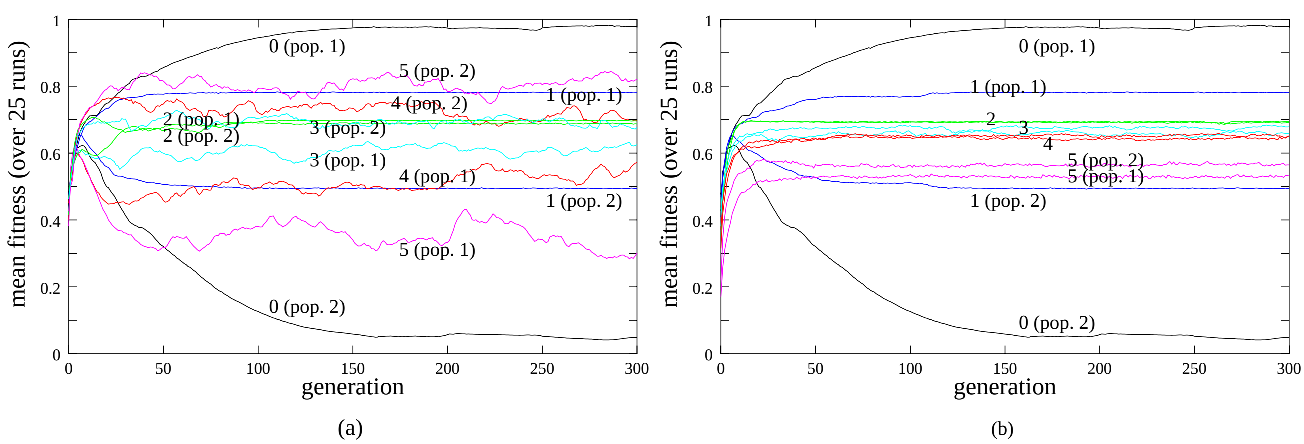

Figure 5: Influence of different fairness models on the mean fitness of the evolving agents (for $p = 1$ and $n = 3$). In Fig. 5a agreements are only re-evaluated if the deadline of the game is actually reached. In Fig. 5b *all* agreements are re-evaluated. Notice in Fig. 5b that for models 2, 3 and 4 the fitnesses of the evolving populations are approximately equal to each other (in the long run). Standard deviations are omitted for clarity. The different fairness models are specified in Fig. 4.

[5]Piece-wise linear functions nicely fit the experimental data reported in [6]. Note, however, that [6] evaluates ultimatum game experiments, as opposed to the multiple-stage bargaining experiments reported here.

$p = 1$ and $n = 3$) that if the agents in population 2 (the last responders) use fairness model 0 (i.e., no fairness check) or 1 (i.e., a weak fairness model), the agents in population 1 reach a relatively high fitness level. Fair agreements evolve, on the other hand, when the agents in population 2 use model 2 (a model with average fairness). In this case the mean long-term fitness is approximately equal to 0.7 for all agents (most agreements are located close to the symmetric point S in Fig. 2). However, when a stronger fairness function is used by the agents (models 3 through 5), the roles reverse and the agents in population 2 reach a *higher* fitness level than their opponents in population 1 (see Fig. 5a). In these extreme models, many last-round agreements are rejected. As a result, the deadline is effectively reached one round earlier. The agents in population 2 can then demand a large share of the surplus in the round before last. If the agents in population 1 do not accept this proposal, chances are few that their counter proposal in the last round is accepted, even if this counter proposal can be considered as fair. Such behaviour by agents in population 2 is therefore better characterised as "greedy" rather than fair. These results show that fair outcomes can in principle evolve in an evolutionary system. However, we also observe a rather large sensitivity to the actual fairness function that is used by the adaptive agents.

An alternative fairness model, in which *each* agreement is re-evaluated, is studied in Fig. 5b. For fairness model 1, similar results are observed as in Fig. 5a. However, when the adaptive agents use fairness models 2 through 4, the agents in both populations reach almost identical fitness levels (close to point S). Agreements are also predominantly reached in the first round, indicating that delaying the negotiations is no longer advantageous. Figure 5b clearly shows that the agents' long-term behaviour is much less sensitive to the shape of the fairness function. Note that the adaptive agents have no explicit knowledge about the location of the symmetric point S. This knowledge is also not incorporated within the fairness models.

Figure 5b shows that when the agents use fairness model 5, the mean fitness of the agents decreases. This decrease in performance is caused by an increasing number of disagreements in the last round. In the experiments we observe that, for fairness model 5, the agents propose a symmetric partitioning, close to point S, in each round. In this case, a proposal is only accepted with a probability of 0.4 (see Fig. 4). Approximately 22% of all bargaining games ends up in a disagreement in this case. The maximum fitness which can be obtained is therefore approximately equal to 0.55. Figure 5b shows that the long-term fitness of the agents is indeed close to this value.

Fair agreements can therefore evolve when the agents check the fairness of agreements each round, although relatively strong fairness models can lead to an increasing number of disagreements. These results show that incorporating a fairness norm in an evolutionary system is feasible. This opens new prospects to capture human behaviour, observed in laboratory experiments, in computational models with artificial agents.

5 Conclusions

We investigate a system for automated negotiations, in which artificial agents learn effective negotiation strategies using evolutionary algorithms. Negotiations are governed by a finite-horizon version of the alternating-offers game. Both negotiations with and without a risk of breakdown are studied. Furthermore, several issues are negotiated simultaneously. In our experiments, the agents have different preferences regarding the importance of the issues. This implies that the agents can reach mutually beneficial agreements if they coordinate their behaviour in an optimal manner.

To validate the evolutionary approach, we apply techniques from the field of game theory. In this validation part, the long-run behaviour of the evolving bargaining agents is compared with subgame perfect equilibrium (SPE) predictions for fully rational and completely informed agents. These SPE predictions appear to be very useful to interpret the behaviour of the boundedly-rational agents. When no risk of breakdown exists, the agents delay almost all agreements until the deadline is reached. The last agent in turn then proposes a take-it-or-leave-it offer and demands most of the surplus for each issue. Similar behaviour is predicted by game theory. When a risk of breakdown exists, on the other hand, most agreements are reached in the very first round. Experiments show that the adaptive agents do not accurately perceive and exploit the finite deadline of the game in this case, especially in long bargaining games. In such games, their behaviour tends more towards theoretical outcomes for infinite-horizon bargaining games.

An asymmetric division of the bargaining surplus, predicted by theoretical (SPE) models, appears to be unrealistic in many real-life bargaining situations. We therefore modelled the important concept of "fairness" in our evolutionary system. In this extended model, the agents carry out a fairness check before an agreement is finally accepted. This fairness check is carried out either solely for agreements which are reached just before the deadline, or for all agreements. In both cases, fair outcomes (i.e. agreements with equal utilities for both players) can be obtained. In the second case, however, the outcomes are much less sensitive to the actual fairness model. This is an interesting result, considering that the agents can reach such fair agreements without any explicit information about each other's preferences. These results are encouraging and open new prospects to capture human behaviour (as observed in laboratory experiments) in computational models with artificial agents.

We intend to study more complex bargaining scenarios in the future. Many extensions of the basic alternating-offers model have already been proposed in the game-theoretic literature [10]. Our simulation environment enables us to study cases for which a rigorous mathematical approach is unwieldy or even intractable. We also plan to further investigate the various experimental fairness models that are available in economic literature.

A Game-Theoretic Analysis of Multi-Issue Negotiations

Subgame perfect equilibrium strategies for multiple-stage games with complete information can be derived using a backward induction approach. In this appendix we follow the same approach as in [15], but extend the analysis to multi-issue negotiations. In Section A.1, we study a model without a risk of breakdown. The more general model (with a risk of breakdown) is then investigated in Section A.2.

A.1 Model without a Risk of Breakdown ($p = 1$)

Because time plays no role in this model, the last agent in turn has the opportunity to reject all proposals from his opponent and demand the entire surplus (for each issue) in the last round. In subgame perfect equilibrium, the other agent accepts this proposal. It should be noted that there does not exist a SPE in which the responder rejects the proposer's take-it-or-leave-it deal (follow the discussion in [2, pp. 200-201]). If the maximum number of rounds n is odd, agent 1 will therefore receive the entire surplus, whereas agent 2 receives all in case n is even.

A.2 Model with a Risk of Breakdown ($p < 1$)

Assume that agent i makes a proposal $\vec{o}_i(t)$ to his opponent, agent j, in round t of the negotiations ($t < n$). Assume also that agent i knows that agent j's threshold is equal to $\tau_j(t)$. It is then a best response for agent i to propose a Pareto-efficient deal to agent j. Consider for example the two-issue bargaining problem depicted in Fig. 2. Suppose agent i proposes an equal partitioning for both issues to agent j. In case of an agreement, this would yield the utility pair (0.5, 0.5) in Fig. 2. However, agent j would be indifferent if agent i demanded the whole surplus for his most important issue and 6/21 for the other issue. This way, agent i's utility would increase from 0.5 to 11/14, whereas agent j's utility would remain the same. This latter agreement is located on the Pareto-efficient frontier. A similar argument holds if the roles of the agents are reversed and player j makes a proposal to agent i.

The SPE partitioning can now be calculated as follows. If the maximum number of rounds n is even, agent 2 will be the proposer in the last round (i.e., at $t = n - 1$). Agent 2 will then demand the whole surplus for each issue and agent 1 will receive nothing. This division of the surplus would yield agent 2 a payoff (expected utility) of p^{n-1}. We now analyse the previous round ($t = n - 2$). Suppose agent 1's offer to agent 2 is $\vec{o}_1(t = n - 2)$. Agent 2's payoff would then be $p^{n-2}u_2(\vec{o}_1(t = n - 2))$. In equilibrium, at $t = n - 2$ agent 1 should propose agent 2 a payoff-equivalent deal. This implies that $u_2(\vec{o}_1(t = n - 2))$ should be equal to p. Agent 1's payoff is then $p^{n-2}f_1(p)$, where $f_1(u_2)$ describes the location of the Pareto-efficient frontier. This function returns the utility of agent 1 when agent 2's utility is equal to u_2 and the agreement is Pareto-efficient[6]. At $t = n - 3$, agent 2 can, in a similar fashion, propose an equivalent offer (in terms of payoff) and receive a payoff of $p^{n-3}f_2(pf_1(p))$. (The $f_2(u_1)$ function is the inverse of the f_1 function.)

This procedure is then repeated until the beginning of the game is reached (at $t = 0$). The same line of reasoning holds if the number of rounds is odd (simply switch the roles of agent 1 and agent 2). In equilibrium, agent 1's payoff at $t = 0$ is then equal to $x_1^*(n)$ and agent 2 receives $x_2^*(n)$. As in the infinite-horizon game [14], the agents agree immediately on a deal. Table 3.2 shows the SPE partitionings for different game lengths.

n	Payoff agent 1 (x_1^*) (SPE)	Payoff agent 2 (x_2^*) (SPE)
1	1	0
2	$f_1(p)$	p
3	$f_1(pf_2(p))$	$pf_2(p)$
4	$f_1(pf_2(pf_1(p)))$	$pf_2(pf_1(p))$
5	$f_1(pf_2(pf_1(pf_2(p))))$	$pf_2(pf_1(pf_2(p)))$
6	$f_1(pf_2(pf_1(pf_2(pf_1(p)))))$	$pf_2(pf_1(pf_2(pf_1(p))))$
...	...	...

Table 2: Payoffs for agent 1 and agent 2 for different lengths (n) of the alternating-offers game, assuming that both agents use SPE strategies.

[6] For the bargaining problem studied in this paper (depicted in Fig. 2), the Pareto-efficient frontier is described by the function $f_1(u_2) = \frac{0.7}{0.3}(1 - u_2)$ for $u_2 > 0.7$. For $u_2 \leq 0.7$, $f_1(u_2) = 1 - \frac{0.3}{0.7}u_2$.

References

[1] Th. Bäck. *Evolutionary Algorithms in Theory and Practice.* Oxford University Press, 1996.

[2] K. Binmore. *Fun and Games, A Text on Game Theory.* D.C. Heath and Company, 1992.

[3] K. Binmore, A. Shaked, and J. Sutton. Testing noncooperative bargaining theory: A preliminary study. *American Economic Review*, 75:1178–1180, 1985.

[4] K. Binmore and N. Vulkan. Applying game theory to automated negotiation. *Netnomics*, 1(1):1–9, 1999.

[5] H. Dawid. *Adaptive Learning by Genetic Algorithms: Analytical Results and Applications to Economic Models.* Lecture Notes in Economics and Mathematical Systems, No. 441. Springer-Verlag, Berlin, 1996.

[6] H. Lin and S. Sunder. Modeling bargaining behavior in ultimatum games. Working paper, avaiable at http://www.som.yale.edu/Faculty/sunder/research.html, 1999.

[7] P. Maes, R.H. Guttman, and A.G. Moukas. Agents that buy and sell. *Communications of the ACM*, 42(3):81–91, 1999.

[8] N. Matos, C. Sierra, and N.R. Jennings. Determining successful negotiation strategies: An evolutionary approach. In *Proceedings of the 3rd International Conference on Multi-Agent Systems (ICMAS-98)*, pages 182–189, Paris, France, 1998.

[9] M. Mitchell. *An Introduction to Genetic Algorithms.* MIT Press, 1996.

[10] A. Muthoo. *Bargaining and Applications.* Cambridge University Press, Cambridge, 1999.

[11] J.R. Oliver. A machine learning approach to automated negotiation and prospects for electronic commerce. *Journal of Management Information Systems*, 13(3):83–112, 1996.

[12] T. Reichmann. Learning and behavioral stability: An economic interpretation of genetic algorithms. *Journal of Evolutionary Economics*, 9(2), 1999.

[13] A.E. Roth. Bargaining experiments. In J. Kagel and A.E. Roth, editors, *Handbook of Experimental Economics*, pages 253–348, Princeton, NJ, 1995. Princeton University Press.

[14] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):155–162, 1982.

[15] D.D.B. van Bragt, E.H. Gerding, and J.A. La Poutré. Equilibrium selection in alternating-offers bargaining models: The evolutionary computing approach. In *Proceedings of the 6th International Conference on Computing in Economics and Finance (CEF'2000)*, Barcelona, Spain, july 2000. forthcoming.

[16] E. Weg, A. Rapoport, and D.S. Felsenthal. Two-person bargaining behavior in fixed discounting factors games with infinite horizon. *Games and Economic Behavior*, 2:76–95, 1990.