Centrum voor Wiskunde en Informatica
REPORT RAPPORT

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

E.H. Gerding, D.D.B. van Bragt, J.A. La Poutré

Software Engineering (SEN)

SEN-R0024 October 31, 2000

Report SEN-R0024
ISSN 1386-369X

CWI
P.O. Box 94079
1090 GB Amsterdam
The Netherlands

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

E.H. Gerding D.D.B. van Bragt J.A. La Poutré
CWI
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.

We also present an extended model of the evolving agents in which the agents use a "fairness" norm in the negotiations. This concept plays an important role in real-life negotiations and experimental economics. In the implementation with fairness, agents also evaluate a potential agreement on its fairness and reject unfair proposals 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. In both cases, fair outcomes can be obtained. When fairness is applied in each round, the results become much more robust and rather insensitive to the actual fairness function.

To validate our system, the computational results are compared to game-theoretic (subgame perfect equilibrium) results. The influence of important model settings, like the probability of breakdown in negotiations, the length of the game, or the influence of fairness, as well as the proper settings of the EA parameters and their sensitivity are substantially investigated in this validation part.

2000 Mathematics Subject Classification: 68T05, 91A20, 91B26
1998 ACM Computing Classification System: I.2.6, I.2.11, J.4
Keywords and Phrases: multi-issue bargaining, evolutionary algorithms, fairness, game theory
Note: Work carried out under theme SEN4 "Evolutionary Systems and Applied Algorithmics". This paper has been presented at 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.

1. Introduction

Lately, automated negotiations have received more and more attention, especially from the field of electronic trading [5, 12, 13, 15]. 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.

In this paper, we consider negotiations that are governed by a finite-stage version of Rubinstein's alternating-offers game [21]. We investigate the computation of strategies of the agents by evolutionary algorithms (EAs). 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 [6, 18]. In an evolutionary setting, the adaptive agents typically learn in different ways: learning by imitation, communication (exchange of strategic information), and experimentation. It is important to note that EAs make no explicit assumptions or use of rationality. Basically, the fitness (i.e., quality) 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 [15] was the first to demonstrate that, using an EA, artificial agents can learn effective negotiation strategies. In Oliver's model, the agents use rather elementary bargaining strategies. A more elaborate strategy representation is proposed and evaluated in [13]. 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 [22]. In [22], a systematic comparison between game-theoretic and evolutionary bargaining models is made.

In this paper we first assess to what extent the evolutionary computation of agent strategies matches with game-theoretic results if multiple issues are involved. This work is therefore in line with the work reported in [22], 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 results). Such extreme outcomes are not observed in real-life situations, where social norms such as fairness play an important role [4, 11, 19, 23]. We therefore introduce the incorporation of a fairness norm in the agents' behaviour. The various fairness models that we incorporate can be tuned from "weak fairness" (i.e., accept almost all agreements) to "strong fairness" (i.e., even reject most of the "quasi-fair" deals). Also, the fairness mechanism can be active in each round or only if the deadline is reached. In this way, we achieve fair deals in our system. Results are depending on the actual fairness settings, but fair deals evolve for most fairness settings if the agents evaluate the fairness of potential agreements in each round; thus, results are much more robust when the latter model is used by the agents.

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 results 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 system and experiments. The alternating-offers negotiation protocol is described in Section 2.1. Section 2.1 also discusses the agents' negotiation strategies and 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 and Agent Model

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 in a disagreement, and both players receive nothing.

Agent Model An agent's strategy specifies the offers and counter offers proposed during the process of negotiation. In a game-theoretic context, a strategy is a plan which specifies an action for each history [3]. 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 [13, 15]. 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 $\vec{w_2} \cdot (\vec{1} - \vec{o})$. Here, $\vec{w_j}$ 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 $w_j^i \geq 0$. 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 focuses on real-coded problems [2].[2]

The different stages within an iteration of the evolutionary algorithm are depicted in Fig. 1. The

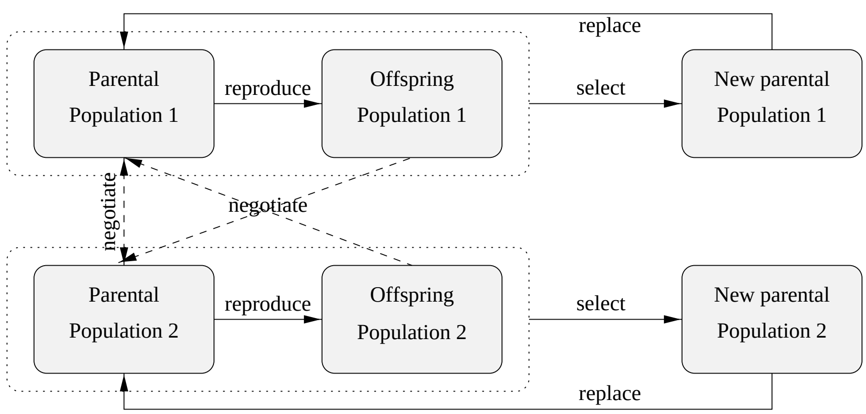

Figure 1: Iteration loop of the evolutionary algorithm. Two populations of agents, containing agents with different weight vectors, are evolved separately. Agents in population 1 always start the bargaining process. In the fitness evaluation, both the offspring and the parental agents compete against 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.

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 populations. Each agent

[1] The same approach was used in [15, 22].

[2] The widely-used genetic algorithms (GAs) are more tailored toward binary-coded search spaces [9, 14, 7].

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.

Selection model. Selection is performed using the $(\mu + \lambda)$-ES selection scheme [2]. 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 [15, 22]. Oliver [15], however, used binary-coded chromosomes. The agents' strategies are initialised at the beginning of each EA run by drawing a random number in the unit interval for each gene (from a uniform distribution).

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 [2, 22]. 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 [2, pp. 71-73][22]. This model allows the evolution of both the genes and the corresponding standard deviations

[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.

at the same time. More formally, an agent consists of object variables $[x_0, ..., x_{l-1}]$ and ES-parameters $[\sigma_0, ..., \sigma_{l-1}]$ in this model.

The mutation operator first updates an agent's ES-parameters σ_i into σ_i'-values in the following way:

$$\sigma_i' := \sigma_i \exp[\tau' N(0,1) + \tau N_i(0,1)], \tag{2.1}$$

where τ' and τ are the so-called "global" and "individual" learning rates, and $N(0,1)$ denotes a normally distributed random variable having expectation zero and standard deviation one. The indices i in N_i indicates that the variable is sampled anew for each value of i. We use commonly recommended settings for these parameters.[4] After the strategy parameters have been modified, the object variables are mutated:

$$x_i' := x_i + \sigma_i' N_i(0,1). \tag{2.2}$$

The initial standard deviations $\sigma_i(0)$ are set to a value of 0.1 (see Table 1). The particular value chosen for $\sigma_i(0)$ is not expected to be crucial, because the self-adaptation process (consisting of the parameter updating and the chromosome selection process) rapidly scales the step sizes into the proper range. To prevent complete convergence of the population, we force all standard deviations to remain larger than a small value $\varepsilon_\sigma = 0.025$ [2, pp. 72–73] (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 [22] showed little effect on the results when traditional recombination operators from ES (like discrete or intermediate recombination [2]) were applied. We therefore focus on mutation-based models in this paper. Note that communication between agents by duplication of (successful) strategies is already modelled by the selection mechanism.

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 results 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) results for 2-issue negotiations. Finally, Section 3.3 reports on additional experiments involving a more complex bargaining scenario and also concludes.

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%), and that 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 offering agent in turn demands almost the entire surplus (for each issue). Nevertheless, his opponent accepts this extreme take-it-or-leave-it deal. This division of the surplus agrees with game-theoretic results (see Appendix 1.1). Note, however, that it is rational for agents to accept this division in other rounds as well: We show in Appendix 1.1 with game theory that rational agents are indifferent between the round in which the agreement is reached when $p = 1$. The deadline-approaching behaviour, that is observed in our computational experiments when $p = 1$, corresponds better to "real-world" behaviour [20], however.

[4] Namely, $\tau' = (\sqrt{2l})^{-1}$ and $\tau = (\sqrt{2\sqrt{l}})^{-1}$ [2, p. 72], where l is the length of the chromosome.

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 results (see Appendix 1.2).

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 utility for

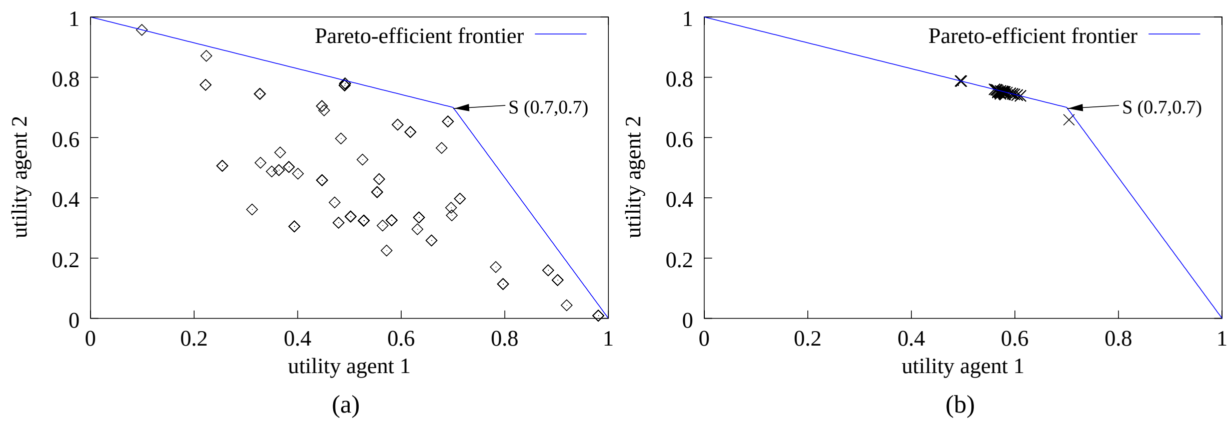

Figure 2: Agreements reached by the evolving agents at the start of a typical EA run (generation 0) are shown in Fig. 2a (for $p = 0.7$ and $n = 10$). Agreements reached after 100 generations are shown in Fig. 2b. Each agreement is indicated by a point in these two-dimensional planes. The Pareto-efficient 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.

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. 2b) 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 thus show that the adaptive agents learn to reach an agreement in an efficient way, viz. on the Pareto-efficient frontier.

3.2 Comparison with Game-Theoretic Results

The computational results are compared in more detail with game-theoretic results in this section. The key equilibrium concept used by game theorists to analyse extensive-form games[5] is the subgame perfect equilibrium (SPE) [16, pp. 97-101]. 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

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

bargaining games [21]. 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-theoretic analysis in some respects, as we show in Appendix 1.

It is important to note that we assume in the game-theoretic analysis of Appendix 1 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 adaptive agents in our computational experiments. The adaptive agents are boundedly rational because they only experience the profit of their interactions with other agents. 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, at subgame-perfect equilibrium the last agent in turn receives the entire bargaining surplus (for each issue) if $p = 1$ (see Appendix 1.1). 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

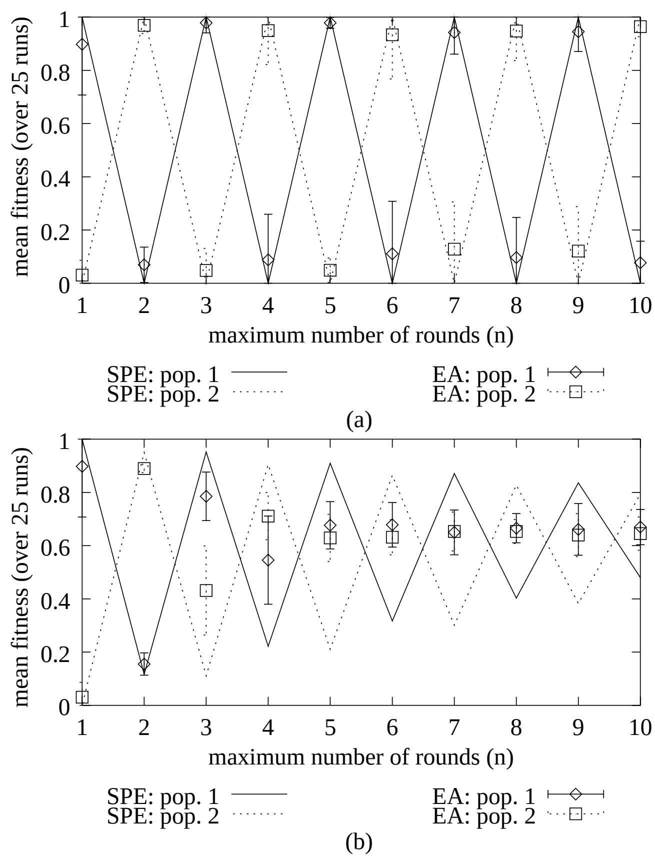

Figure 3: Comparison of the evolutionary results with SPE results for (a) $p = 1$ (time indifference) and (b) $p = 0.95$. The SPE predictions for successive values of n are connected to guide the eye. Figure 3a shows that the finite length of the game has a strong impact on the long-term behaviour if $p = 1$. If $p = 0.95$, 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 died out. The error bars indicate the standard deviations across 25 runs.

finite length of the game diminishes for longer games if $p = 0.95$. Notice for instance in Fig. 3b that the SPE partitioning is quite asymmetric (i.e., agents 1 and 2 obtain different payoffs) for small n, but more symmetric when n is large. 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 Appendix 1.2), 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; in fact, the game length is strongly overestimated (also cf. [22]).

Similar results are obtained for other values of p. However, as p becomes smaller, the influence of the game length on the SPE outcome also becomes smaller (also see [22]). Therefore, if p is small (e.g., $p < 0.8$), the computational results automatically show a much better match with SPE outcomes than if p is large. To be precise, the achieved agreements match the SPE perfectly, while still a small number of disagreements occur.

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 corresponds much better to an infinite-horizon model than the finite-horizon model for $n \geq 8$. The same behaviour was observed for other EA settings (e.g. larger population size) and other negotiation situations (e.g. other weight settings).

3.3 Further Experiments and Conclusions

In this section we study the performance of the EA for more complex bargaining problems (with a larger number of issues). As a test case, we increased m to 8 (with randomly generated weight vectors $\vec{w}_i$ for each population of agents $i = 1, 2$). When all other parameters are kept the same as in Table 1, we observe that for $p = 1$ the long-term outcomes of the EA do not converge to the extreme partitionings shown in Fig. 3a (even for n as small as 5). In particular, the EA appears to be unstable. When we increase the population size from 25 to 100 agents,[6] the extreme partitionings in Fig. 3a reappear. Thus, for more complicated bargaining problems, the EA parameters must be adjusted. E.g., when m is larger, sufficiently large population sizes must be used to increase stability.

Other experiments are performed with $m = 8$ and $p < 1$, using the adjusted population size as described above. Similar observations are found as reported in Section 3.2 (like Fig. 3) for these experiments.

To conclude, we compared the computational results with results from game theory. Game-theoretic (SPE) results appear to be a very useful benchmark to investigate the results of the evolutionary simulations. We investigated in detail the influence of the finite length of the game. In computational simulations without a risk of breakdown, agreements are predominantly reached in the final round. This deadline effect is consistent with human-like behaviour. Furthermore, the last agent in turn successfully exploits his advantage and claims a take-it-or-leave-it deal (as occurs in SPE). In case of a small risk of breakdown ($p > 0.8$), on the other hand, this last-mover advantage is smaller than predicted by game theory. In fact, if the finite game becomes long enough, the deadline is no longer perceived by the evolving agents. In that case, the evolutionary system actually computes infinite-horizon results ($n = \infty$). This deviation from SPE becomes negligible, however, if the risk of breakdown is large ($p < 0.8$). Furthermore, it is important to work with sufficiently large population sizes in the case of large(r) number of issues, to increase stability and the quality of the actual outcomes.

[6] To avoid a quadratic increase in the number of fitness evaluations, each agent negotiates with 25 opponents from the other population in this case. These opponents are selected at random (without replacement).

4. Social Extensions: Fairness

We extend the agent model within our evolutionary system in this section to study the influence of "fairness", an important aspect of real-life bargaining situations. The motivation and description of this fairness model is given in Section 4.1. In the fairness model studied in Section 4.2 the evolving agents only take the fairness of a proposed deal into account when the deadline is reached. Section 4.3 presents results obtained when agents perform a "fairness check" in each round. Section 4.4 further analyses the model in Section 4.3. A comparison is made with game-theoretic (SPE) calculations for a simple case.

4.1 Motivation and description: the fairness model

Game-theoretic models for rational agents often yield very asymmetric outcomes for the two parties. We showed in Section 3.2 (see Fig. 3a) that such "unfair" behaviour can also emerge in a system of evolving agents. Game-theoretic (SPE) results have also been compared to human behaviour in laboratory experiments in the past (see [19] for an extensive overview). Large discrepancies between human behaviour and SPE outcomes were found in these experimental studies, both for ultimatum and multi-stage games [4, 11, 23]. These studies show for instance that a human proposer typically demands less in the ultimatum game than predicted by SPE (in SPE a proposer demands the *whole* surplus). A human responder also often rejects an unequal division of the surplus (even if the game is played only once against a particular opponent), whereas in SPE the responder accepts *all* proposals.

A possible explanation for the occurrence of these discrepancies between theory and practice is the strong influence of social or cultural norms on the individual decision-making process. Binmore et al. [4] argue, for example, that the *proposer* in the ultimatum game does not demand the whole surplus (as subgame perfection predicts), because he values the "fairness" of the partitioning. Another hypothesis is that the *responders* tends to reject unfair or "insultingly low" proposals [19, p. 264]. Therefore, an anticipating proposer should lower his demand in order to avoid a disagreement, this way taking into account the expectations about his opponent's behaviour. The importance of this phenomenon is also studied in [8].

Lin and Sunder [11] recently proposed a model in line with the second hypothesis stated above. In their model, 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 [11]. We will introduce a fairness model in our evolutionary system that has correspondence with Lin and Sunder's model.

The model described in [11] is used to evaluate ultimatum game experiments only (i.e., $n = 1$). We extend this approach to multiple-stage games in two different ways. On the one hand, we investigate the system's behaviour if the responding agent has a single fairness check in the last round. On the other hand, we investigate the system's behaviour if the responding agent has a fairness check independent of the round number. The first case is motivated by the deadline-effect observed in the experiments without a risk of breakdown (see Section 3.2): most agreements are reached in the last round and the final proposer then receives almost the whole bargaining surplus. The second case, however, is more likely to be an appropriate model for human behaviour.

We now introduce fairness by extending the agent model (see Section 2.1) as follows. If the value of an offer exceeds the responder's threshold, he has the opportunity to re-evaluate his decision. The probability that he finally accepts the agreement is a function of the acquired utility. This function is the so-called fairness function. This function is piece-wise linear, with up to three segments.[7] The instances that we use are shown in Fig. 4.[8] We now further distinguish between two different extended agent models. In the first model, the fairness function is effective at the deadline only (this is studied in Section 4.2). In the second model, the fairness function is effective at any moment (this is studied in Section 4.3).

[7] Piece-wise linear functions nicely fit the experimental data reported in [11].

[8] We want to remark here that, although the fairness function is the same for all agents, the actual fairness function can depend on cultural norms in the real world [11].

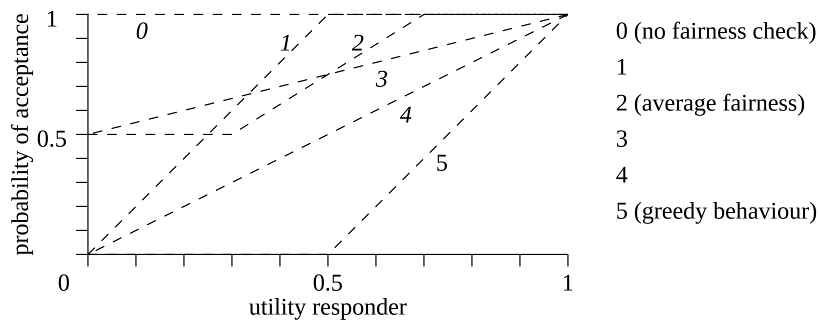

Figure 4: Different fairness functions used by the agents in the evolutionary experiments to determine the degree of fairness of agreements. Each function specifies the probability that the responding agent accepts a potential agreement as a function of the acquired utility. The functions range from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

Note that the fairness model is an extension on the responder's behaviour, and does not replace the responder's thresholds. In Appendix 1.3 we show that, for a simple bargaining game, the thresholds are indeed not superfluous and still play an important role in the outcome of the SPE when fairness is applied.

4.2 Fairness check at the deadline

In this section, fairness is applied by the responding agents on potential agreements in the last round. We use $p = 1$ and $n = 3$ in these experiments (other settings are as before, see Table 1). Therefore, the fairness function is only used by the agents in population 2 in the third round. Figure 5 shows that if the agents in population 2 use fairness function 1 (i.e., a weak fairness model), the partitioning is much less extreme than in case of no fairness check (function 0). However, the agents in population 1 still reach a relatively high fitness (utility) level. Fair agreements evolve, on the other hand, when the agents in population 2 use function 2 (a case 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

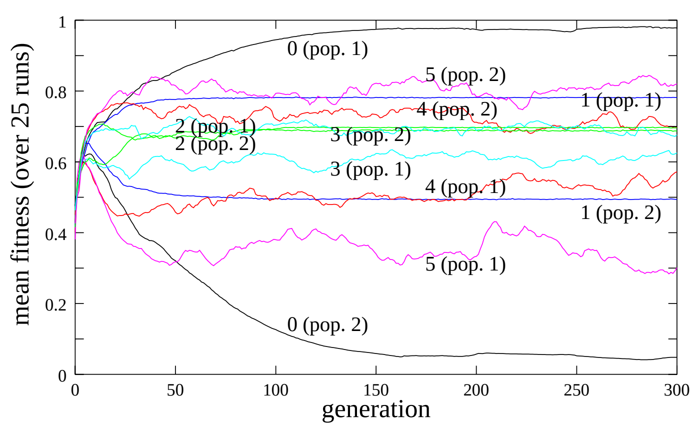

Figure 5: Influence of different fairness functions on the mean fitness of the evolving agents (for $p = 1$ and $n = 3$). Agreements are only re-evaluated if the deadline of the game is actually reached. Equal payoffs for both agents is achieved when fairness function 2 is used. In case of a "regular" fairness function, such as functions 3 to 5, the roles are switched and agent 2 receives a *larger* share. The different fairness functions are specified in Fig. 4.

(functions 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. 5). Note that, in these extreme fairness functions, many last-round agreements are rejected. As a result, the deadline seems to be effectively reached one round earlier since the fairness is active in the last round only. This indeed became apparent in our experiments: The agents in population 2 then demand and get a large share of the surplus in the round before last (note that if the agents in population 1 do not accept this proposal, chances are few that their counter proposal in the last round is accepted due to the strong fairness function).

Concluding, our results show that fair outcomes can evolve in an evolutionary system with a fairness model in the last round. However, there is a rather large sensitivity to the actual fairness function that is used by the adaptive agents; an "average" fairness function[9] yields the fair results, whereas more extreme ones give larger deviations.

4.3 Fairness check in each round

This section studies the second fairness model, in which the responding agent re-evaluate potential agreements each round (i.e., not only at the deadline). The EA settings are the same as in the previous section. The results are shown in Fig. 6. For fairness functions 1, similar results are observed as in Fig. 5. However, when the adaptive agents use fairness functions 2 through 5, the agents in both populations reach almost identical fitness levels (close to point S in Fig. 2). Also, we observe that agreements are usually reached in each round. Thus, Fig. 6 shows that the agents' long-term behaviour is much less sensitive to the shape of the fairness function: the various "stronger" fairness functions all yield similar results. Also note that the adaptive agents have no explicit knowledge about the location of the symmetric point S, and that this knowledge is also not incorporated within the fairness functions.

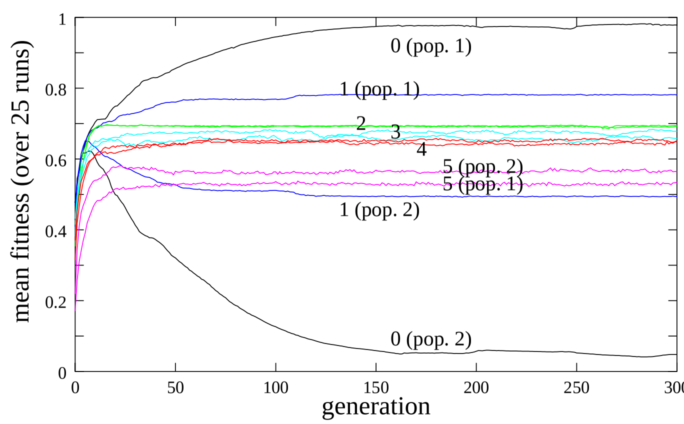

Figure 6: Influence of different fairness functions on the mean fitness of the evolving agents (for $p = 1$ and $n = 3$). All potential agreements are re-evaluated before the offer is finally accepted. Notice that for functions 2 to 5 the fitnesses of the evolving populations are approximately equal to each other (in the long run). Standard deviations are omitted for clarity. The different fairness functions are specified in Fig. 4.

Figure 6 also shows that when the agents use fairness function 5, the mean fitness of the agents decreases. This decrease in performance is caused by an increasing number of disagreements. In the experiments we observed that for fairness function 5, the agents do propose a symmetric partitioning, close to point S, in each round. A proposal is then accepted in a certain round with a probability of 0.4 (see Fig. 4). Thus, approximately 22% of all bargaining games ends up in a disagreement with this fairness function and the maximum possible fitness is approximately equal to 0.55. Figure 6 shows that the long-term fitness of the agents is indeed close to this value.

[9]This function corresponds most to the shape of the various data reported in [11].

From our experiments it thus follows that fair agreements can evolve if fairness is evaluated each round, even when the acceptance probability becomes relatively low: the fairness of the deals is much more stable w.r.t. the actual choice of the fairness function when applied each round. Of course, the number of actual agreements drops if a strong fairness function is used, resulting in a (mutual) lower fitness.

We also performed several additional experiments, for instance with different weight vectors and with $m > 2$. A general finding is that extreme outcomes of the evolutionary process do not occur if the agents apply a fairness check. This does not mean, however, that in the long term the agents in population 1 and population 2 always reach the exact same fitness level (as is suggested in Fig. 6, especially for fairness functions 2, 3 and 4). In particular, unequal outcomes can occur if the Pareto-efficient frontier is asymmetric.[10]

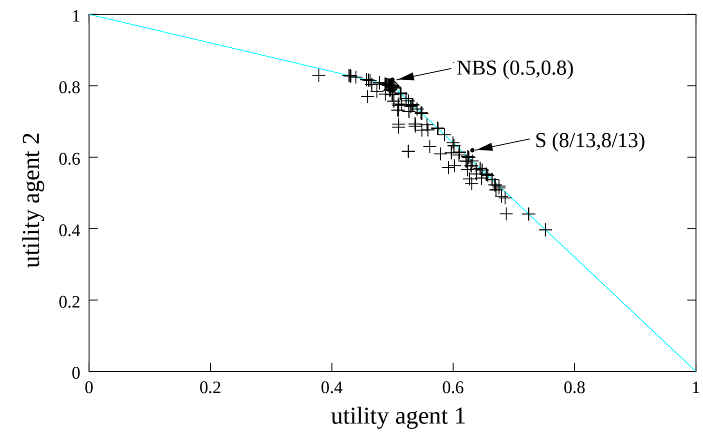

Figure 7: Agreements reached at generation 300 by the evolving agents when agent 1 has weight settings $(0.5, 0.5)$, and agent 2 uses $(0.2, 0.8)$. The responding agents evaluate fairness using function no. 4. Notice that with these weight settings the Pareto-efficient frontier is asymmetric. The point NBS is called the Nash bargaining solution, and maximises the product of the agents' utilities (see [3, Ch. 5] for further details). The point S denotes the symmetric point on the Pareto-efficient frontier, where the agents obtain equal payoffs.

An asymmetric Pareto-efficient frontier of a 2-issue negotiation problem is shown in Fig. 7. In this case, agent 1 values both issues equally important, whereas agent 2 has different valuations for each issue (his weights are 0.2 and 0.8 for issues 1 and 2 respectively). This creates an advantage for agent 2. If each agent obtains the whole surplus on his most important issue, agent 1 obtains 0.5, whereas agent 2 gets 0.8.[11] The symmetric point, on the other hand, is located at $(\frac{8}{13}, \frac{8}{13})$.[12]

Both solutions can be considered to be fair outcomes in different ways: the first solution maximises the product of the agents' utilities and also splits the surplus equally, whereas in the second case equal utility levels are obtained for both agents (see [17, Ch. 16] for a related discussion). In the computational results, we observe that, when fairness functions 2-5 are applied, the agreements are divided and are usually concentrated in two separate clusters (clouds). An exemplar snapshot is shown in Fig. 7, where the agents apply fairness function no. 4. It follows that, on average, the agents in population 2 obtain a significantly larger payoff in these experiments. The issue of the choice of and distribution over multiple "fair" agreement points seems an important issue for further research, in both a computational setting as well as in experimental economics.

[10] A symmetric Pareto-frontier is mirror-symmetric with respect to the diagonal in the two-dimensional utility plane (e.g. Fig. 2).

[11] This outcome corresponds to the Nash bargaining solution [3, Ch. 5], see Fig. 7.

[12] This outcome corresponds to the Kalai-Smorodinsky solution [10].

	Payoff agent 1	Payoff agent 2
SPE	0.419	0.391
EA	0.391 ($\pm$0.022)	0.412 ($\pm$0.014)

Table 2: Comparison of the adaptive agents' payoffs in the EA with SPE results. Results for the adaptive agents are obtained after 300 generations (averaged over 25 runs).

To conclude, in this section we showed that the outcome of the agreements are quite robust and independent of the fairness function when the agents apply fairness each round. In case of two-issue negotiations with a symmetric Pareto-efficient frontier, most agreements are reached at the symmetric point. In the asymmetric case, fair solutions can also be obtained. The solutions are then distributed over various possible outcomes, which can all be considered fair in different ways.

In the next section, we investigate the evolving strategies of the agents in more detail. The computational results for models with fairness are also complemented by a game-theoretic analysis for single-issue negotiations in this section.

4.4 Comparison with Game-Theoretic Results

Although incorporating fairness aspects makes the analysis more complicated, SPE strategies can again be derived for finite-length games using a backward induction mechanism. The equations for calculating the SPE strategies and outcomes are given in Appendix 1.3. A solution for $m = 1$, $n = 3$ and $p = 1$ is also given in Appendix 1.3. In this example, offers are evaluated with fairness function no. 4 in Fig. 4. This function, as well as $m = 1$, were chosen because of mathematical feasibility. We note that solving these above equations for multiple issues becomes increasingly complex, and is beyond the scope of this paper; this is left for further research. We henceforth focus on the manageable single-issue negotiations.

Section 4.4 compares the game-theoretic and EA results for single-issue negotiations and analyses the strategies in detail. Only small deviations between theoretical and computational results are found. Section 4.4 shows that even better results can be achieved by altering the mutation scheme.

Strategy analysis Table 2 shows both the SPE results and the payoffs obtained by the evolving agents (in the long run) in the game with $m = 1$, $n = 3$, and $p = 1$, and with (the rather strong) fairness function 4. Notice that the SPE payoffs are in good agreement with the outcome of the evolutionary experiments. However, in SPE agent 1's payoff is slightly larger than agent 2's payoff. Exactly the opposite effect is visible in the EA, although Table 2 shows that differences between theory and experiment are small (we will further address this in Section 4.4). Henceforth, we will analyse the evolving strategies in more detail. This analysis shows in what sense the negotiation strategies of the evolving agents are affected by fairness considerations.

Round	Offer (SPE)	Offer (EA)	Threshold (SPE)	Threshold (EA)
1	0.609	0.58 ± 0.06	0.391	0.23 ± 0.21
2	0.375	0.39 ± 0.07	0.250	0.14 ± 0.13
3	0.500	0.48 ± 0.09	0.000	0.13 ± 0.13

Table 3: Comparison of the evolved strategies with game-theoretic (SPE) results. The offer of the proposing agent is shown for each round, together with the acceptance threshold of the responder in the same round. Results for the adaptive agents are obtained after 300 generations (averaged over 25 runs).

Table 3 compares the offers of the evolving agents (for each round) with SPE results. Recall that offers are labelled in terms of the amount received by player 1 (see Section 2.1). Table 3 shows that the offers made by the evolving agents agree well with SPE results for all rounds. Moreover, because of the fairness function, agreements are reached in *all* rounds (both in SPE and in the EA agreements per round are in the range of $20-40\%$), with some emphasis on the first round. This can be calculated from Table 3 for SPE, while similar acceptance rates for each round were observed in the EA.[13]

Table 3 also shows the acceptance thresholds for each round as predicted by game theory and as generated by the EA. In SPE, the threshold of the (responding) agent is equal to the payoff of this agent when he rejects the proposer's offer and waits until the next round (see Appendix 1.3). Table 3 shows that in rounds 2 and 3 the thresholds in SPE are not really relevant. This explains the large variance in the thresholds of the evolutionary agents in rounds 2 and 3 (see Table 3), viz.: in the evolutionary experiments the value of the threshold does not affect the results (in rounds 2 and 3) as long as the threshold value is below the payoff received by the responder. In round 1, however, the threshold is important in SPE and influences the offer made. The computational experiments seem to indicate otherwise, and show a lower average threshold value compared to SPE results. Notice, however, that the variance is rather high for this variable. The latter (0.23 ± 0.21) forces the offers to a similar level as in SPE. We further discuss the influence of the threshold below.

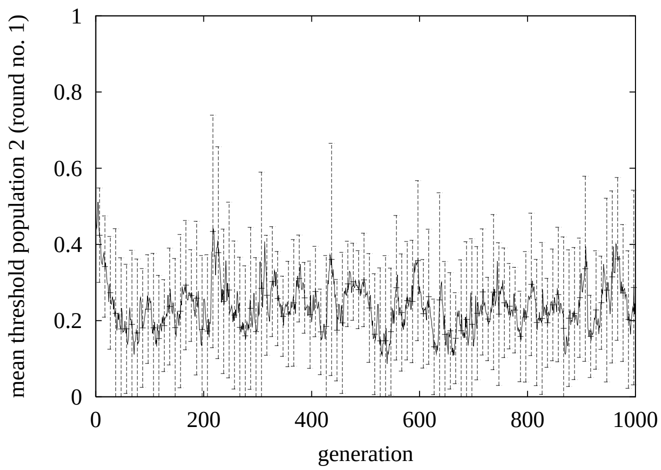

Figure 8: Evolving threshold values for the first round in the population of initial responders (population 2) for a single run of the EA. The dispersion from the mean value is indicated with error-bars. Notice the significant variance, both within the population and during the course of evolution.

Figure 8 shows the evolution of the threshold value for the first round in the population of initial responders for a single experiment. The error-bars in this figure indicate the dispersion from the mean value of the threshold. Notice the significant variance, both within the population and during the course of evolution. One the one hand, this is the reason that although the mean threshold is lower than in SPE, the effect on the offers of the evolving agents in population 1 is similar. On the other hand, the slight deviations (from SPE) of the offers made by the adaptive agents can be explained by these variations in the threshold values: because of the occurrence of frequent (stochastic) positive peaks in this threshold value, agent 1 has a slightly lower offer than that from SPE to avoid a too large number of rejected proposals.

We also checked the "consistency" of the strategies of the evolutionary agents in longer games. In SPE, the offers and thresholds in the final rounds are identical regardless of the game length. We performed an experiment with $n = 6$. The strategies followed by the evolutionary agents in the last 3 rounds of both games (i.e., $n = 3$ and $n = 6$) indeed turned out to be similar, with similar thresholds

[13]Acceptance rates are approximately 39%, 22%, 20% for the SPE rounds 1-3, and $36 \pm 4\%$, $25 \pm 3\%$, $20 \pm 2\%$ for the EA rounds 1-3.

and offers.

To conclude, this relatively simple bargaining situation shows a good match between theoretical (SPE) and experimental results. Furthermore, when fairness norms are applied, the outcome of the negotiation process comes to depend on the actual round in which an agreement is finally reached. We will consider improvements in the EA performance in Section 4.4.

Tuning the EA Parameters We now investigate the further optimization of EA parameters. We especially focus on increasing stability in the populations and improving the match between EA and SPE as discussed in Section 4.4.

Instead of self-adaptive mutation step-sizes σ_i (see Section 2.2), we use an annealing mechanism for the mutation step sizes in our EA.[14] Here, the mutation step-size gradually decreases as follows. At the beginning of each EA run, the mutation step-sizes are set equal to $\sigma_i(0) = 0.1$ (as before, see Table 1). The step sizes then exponentially decrease until $\sigma_i = 0.01$ after 1000 generations. This procedure indeed reduces the fluctuations in the threshold values and the offers in the long run. Results for experiments with this EA setting appear to be in excellent agreement with SPE results, see Table 4.

	Payoff agent 1	Payoff agent 2
SPE	0.419	0.391
EA with decreasing σ_i	0.416 ± 0.012	0.395 ± 0.009

Table 4: Comparison of the evolutionary agents' payoffs with SPE results when the mutation step-sizes σ_i of the adaptive agents exponentially decrease in the course of evolution. Results for the evolutionary agents are obtained after 1000 generations. Notice the excellent agreement between theory and experiment.

We also applied this method on the symmetric two-issue negotiation problem as described in Section 4.3. In this case ($m = 2$), we found no significant effect of the new mutation scheme on the average fitness of the populations. We suspect that this is due to the integrative nature of the negotiation problem: the results obtained in case of two issues are beneficial for both parties and are therefore already much more stable.

In this section we have shown that by choosing other EA parameters, the EA can be fine-tuned to a more stable situation if needed. This rendered an optimal match with the SPE results for $m = 1$.

5. Conclusions

We have investigated a system for negotiations, in which agents learn effective negotiation strategies using evolutionary algorithms (EAs). 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 applied 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 results appear to be very useful to interpret the behaviour of the (boundedly-rational) evolutionary agents. When no risk of breakdown exists, the evolutionary 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. This extreme division of the surplus is also predicted by game-theoretic SPE, whereas the deadline timing is not (SPE outcomes are indifferent on timing here). Notice that this timing in the evolving agents actually corresponds very well to real-life situations.

[14] A similar approach was applied in [1] for a genetic algorithm.

When a risk of breakdown exists, on the other hand, most agreements are reached in the very first round. If the the finite game becomes long enough, the deadline is therefore no longer perceived by the evolving agents. In that case, the evolutionary system deviates from SPE outcomes and actually computes infinite-horizon results. However, if the risk of breakdown becomes very small, the influence of the game length on SPE outcomes also becomes smaller. Henceforth, the deviation between EA and game-theoretic results then becomes negligible. Note that the deviation from game theory may correspond better to real-life situations.

In the second part of the paper, we reconsider bargaining with no risk of breakdown. An extreme division of the bargaining surplus, predicted by theoretical (SPE) models in this case, appears to be unrealistic in many real-life bargaining situations. We therefore introduced and modelled the important concept of "fairness" in our evolutionary system. In this extended model, a responding agent carries out a fairness check before an agreement is definitely accepted. This fairness check was modelled in two ways. In the first model, a responding agent only considers fairness at the deadline. In the second model, fairness is considered for any potential agreement, independent of the round in which it occurs.

In both cases, fair outcomes can be obtained. In the first case, fairness results are very sensitive to the actual fairness function (which can be cultural dependent [11]). In the second case, however, the outcomes are rather insensitive to the actual choice of fairness function. Thus, much more robust results are obtained; these become rather independent of reasonable variations in the actual fairness function. This second case also seems to correspond more to real-world situations. The second fairness model was also applied for an asymmetric bargaining situation (where the players have asymmetric preferences). Multiple outcomes then exist which can be considered "fair" in different ways. Agreements reached by the evolutionary agents are then usually divided in separate clusters around these fair points.

We also performed a theoretical analysis of the fairness models. As an example, we compared theoretical results for a simple bargaining game (concerning a single issue) with the outcome of evolutionary experiments. A good match between theoretical and experimental results was observed. Interestingly, both the theoretical and the experimental approach shows that several outcomes can co-exist when such a fairness model is used: The deal then depends on the round in which the agreement is reached. These results are encouraging and open new prospects to capture human behaviour (as observed in laboratory experiments) in computational models with artificial agents.

Furthermore, we improved the results for several cases by fine-tuning the EA parameters. In particular, we used other mutation schemes or population sizes to increase stability in the EA population. These optimisations in the EA parameters resulted in better (optimal) matches with the corresponding game-theoretical results.

Note that the results are obtained by the EA without any a priori knowledge of the opponent's preferences. These preferences are implicitly learned within the negotiation strategy in the course of evolution.

A simulation environment as described here facilitates the study of cases for which a rigorous mathematical approach is unwieldy or even intractable. An interesting line of research is to further explore the notion of fairness and to compare the computational outcomes with results from experimental studies with human subjects. Of particular interest is the study of asymmetric multi-issue bargaining situations, where more than one outcome can be considered "fair". This raises several new research questions for experimental economics as well as computational sciences.

APPENDIX

1. 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 [22], but extend the analysis to multi-issue negotiations and to models which take fairness considerations into account. In Appendix 1.1, we study a model without a risk of breakdown. The more general model (with a risk of breakdown) is then investigated in Appendix 1.2. This model is extended to the case in which the agents perform an additional fairness check in Appendix 1.3.

1.1 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 (see also the discussion in [3, 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. Due to the absence of time pressure multiple subgame perfect equilibria exist in this case. Although these equilibria differ in the timing of the agreements, the all result in the same outcome (i.e., the agent in turn at $t = n-1$ always receives the entire surplus for all issues). It is for instance subgame perfect for the last responder to concede the entire surplus (for all issues) to his opponent before the deadline is actually reached or, alternatively, to accept a take-it-or-leave-it deal from the opponent at any point in time.

1.2 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 agent 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 $\pi_2(t = n-1) = p^{n-1}$, where $\pi_i(t)$ denotes agent i's payoff in the bargaining game starting at time t. 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 $\pi_2(t = n-2)$ 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 [i.e., $\pi_2(t = n-2) = \pi_2(t = n-1)$], . This implies that $u_2[\vec{o}_1(t = n-2)]$ should be equal to p. Agent 1's payoff $\pi_1(t = n-2)$ 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.[15] At $t = n-3$, agent 2 can, in a similar fashion, propose an equivalent offer (in terms of payoff) and receive a payoff of $\pi_2(t = n-3) = p^{n-3} f_2[p f_1(p)]$. (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). As in the infinite-horizon game [21], the agents agree immediately on a deal. Table 3.2 shows the SPE partitionings for different game lengths.

[15]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 $u_2 \leq 0.7$, $f_1(u_2) = 1 - \frac{0.3}{0.7} u_2$.

n	Payoff agent 1 $[\pi_1(t=0)]$ (SPE)	Payoff agent 2 $[\pi_2(t=0)]$ (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 5: Payoffs for agent 1 and agent 2 for different lengths n of the alternating-offers game, assuming that both agents use SPE strategies.

1.3 Model with fairness

The fairness models evaluated in Section 4.2 (i.e., with a fairness check at the deadline only) and in Section 4.3 (i.e., with a fairness check in each round) are analysed in this appendix. As in Appendix 1.2, we apply backward induction to deduce the SPE partitioning. The fairness function is now formally denoted as $g_t(u)$. This (real-valued) function returns the probability of acceptance of a proposal in round t in case the responding agent's utility is equal to u. If a fairness check is performed only in the last round, $g_t(u) = 1$ for all $t < n-1$. In case the same fairness check is performed each round, $g_t(u)$ is independent of t. We assume that the fairness function is a monotonic non-decreasing function of u and that $g_t(u=1) = 1$. Let agent i be the agent proposing a deal at time t and agent $-i$ the responder. We then abbreviate $g_t[u_{-i}(\vec{o}_i(t))]$ (the probability of acceptance of offer $\vec{o}$ in round t) as $p_t^{acc}(\vec{o})$.

If n is even, agent 2 will be the proposer in the last round (at $t = n-1$). Agent 2 will then propose an offer $\vec{o}_2(t=n-1)$ which maximises his payoff (in SPE). Agent 2 receives a payoff of $p^{n-1}u_2[\vec{o}_2(t=n-1)]$ if his offer is accepted. The acceptance probability is equal to $p_{n-1}^{acc}[\vec{o}_2(t=n-1)]$. Agent 2's payoff in round $t = n-1$ is therefore:

$$\pi_2(t=n-1) = \max_{\vec{o}_2(t=n-1)\in\mathcal{P}} p^{n-1}u_2[\vec{o}_2(t=n-1)]p_{n-1}^{acc}[\vec{o}_2(t=n-1)], \tag{1.1}$$

where $\mathcal{P} \subset [0,1]^m$ is the set containing all Pareto-efficient offers. Analogously, the payoff for agent 1 in round $t = n-1$ is equal to:

$$\pi_1(t=n-1) = p^{n-1}u_1[\vec{o}_2(t=n-1)]p_{n-1}^{acc}[\vec{o}_2(t=n-1)]. \tag{1.2}$$

It is again straightforward to show that it is optimal to propose a Pareto-efficient deal. Assume for instance, that a Pareto-*in*efficient offer is made. The proposer of this offer can then improve his payoff by selecting an offer on the Pareto-frontier which yields his opponent the same payoff. Because the probability of acceptance only depends on the responder's utility of this offer, this will not affect the fairness evaluation.

We now analyse the previous round ($t = n-2$). In SPE, at $t = n-2$ agent 2 only accepts a deal which is at least equal to the payoff $\pi_2(t=n-1)$ that he receives in the next round (in SPE). Therefore, $\pi_2(t=n-2) \geq \pi_2(t=n-1)$ in SPE. Effectively, $\pi_2(t=n-1)$ acts as a threshold used by agent 2 to determine the minimal acceptable offer at $t = n-2$. Some elementary manipulations then show that in SPE agent 1 should make an offer $\vec{o}_1(t=n-2)$ such that

$$p^{n-2}u_2[\vec{o}_1(t=n-2)] \geq \pi_2(t=n-1), \tag{1.3}$$

otherwise, agent 2 rejects the proposal at $t = n-2$ to earn $\pi_2(t=n-1)$ in the last round. We now define $\mathcal{R} \subset [0,1]^m$ to be the set of offers for which Eq. 1.3 is not violated. In SPE, agent 1's payoff in round $t = n-2$ then equals

$$\begin{aligned}\pi_1(t=n-2) &= \max_{\vec{o}_1(t=n-2)\in\mathcal{P}\cap\mathcal{R}} p^{n-2}u_1[\vec{o}_1(t=n-2)]p^{acc}_{n-2}[\vec{o}_1(t=n-2)] \\ &\quad +\{1-p^{acc}_{n-2}[\vec{o}_1(t=n-2)]\}\pi_1(t=n-1). \end{aligned} \tag{1.4}$$

In a similar fashion, we can calculate agent 2's payoff at $t = n-2$ in SPE:

$$\begin{aligned}\pi_2(t=n-2) &= p^{n-2}u_2[\vec{o}_1(t=n-2)]p^{acc}_{n-2}[\vec{o}_1(t=n-2)] \\ &\quad +\{1-p^{acc}_{n-2}[\vec{o}_1(t=n-2)]\}\pi_2(t=n-1). \end{aligned} \tag{1.5}$$

For $t = n-3$ expressions very similar to Eqs. 1.4 and 1.5 can be derived (but the roles of the two agents switch). 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 the previous model without fairness (see Appendix 1.2) all agreements occur in the first round in SPE (for $p < 1$). When the agents apply a fairness check in each round, however, even in SPE a significant number of agreements occurs after the first round. In this case, the strategy followed in all rounds comes to play a role in determining the outcome of the game.

We also remark that, although a responder's fairness considerations determines for a large part the offers made by a proposer, this does not make the responder's thresholds superfluous in SPE. Recall that the role of the threshold is reflected in Eq. 1.3. This will also become apparent in the example shown below.

Simple example (see Section 4.4) We will now apply the general approach presented above to a simple single-issue (i.e., $m = 1$) bargaining problem. Because $m = 1$, the offer vector $\vec{o}(t)$ has only a single component. We denote the value of this component as $x(t)$ in the remainder of this appendix. It is obvious (because the agents are assumed to be risk neutral, see Section 2.1) that $u_1[x(t)] = x(t)$, and $u_2[x(t)] = 1-x(t)$ for $0 \leq t \leq n-1$. The agents evaluate the fairness of the offers (in each round) using fairness function 4 in Fig. 4 [i.e., $g_t(u) = u$]. Furthermore, we take $n = 3$ and $p = 1$. Notice that, because the number of rounds n is odd in this example, we need to switch the roles of agent 1 and agent 2 when we apply Eqs. 1.1-1.5 in the following.

Agent 1 makes an offer to agent 2 in the final round (at $t = 2$). In SPE, agent 1 applies Eq. 1.1 to maximise his payoff $\pi_1(t = 2)$. Substituting parameters for this problem, the product term on the RHS of Eq. 1.1 becomes $u_1[x(t=2)]g_2[u_2(x(t=2)]$, which can be simplified further to $x(t=2)[1-x(t=2)]$. This term is maximised for $x(t = 2) = 0.5$, which results in $\pi_1(t = 2) = 0.25$. Using Eq. 1.2, the payoff of agent 2, $\pi_2(t = 2)$, is then also equal to $[1-x(t=2)]x(t=2) = 0.25$.

Agent 2 makes a move at $t = 1$. We initially assume that the condition stated in Eq. 1.3 is not violated by agent 2's offer. Agent 2's payoff is then determined by applying Eq. 1.4. Substituting the parameters of this problem and simplifying, the term that should be maximised in Eq. 1.4 becomes equal to $[1-x(t)]x(t) + [1-x(t)]0.25$. This term is maximised for $x(t = 1) = 0.375$. The condition stated in Eq. 1.3 is not violated because $u_1(0.375) = 0.375 \geq 0.25$. Our initial assumption therefore turns out to be valid. We can now apply Eqs. 1.4 and 1.5 to derive that $\pi_1(t = 1) \approx 0.297$ and $\pi_2(t = 1) \approx 0.391$.

Agent 1 proposes an offer in the first round (at $t = 0$). Again, we initially ignore Eq. 1.3. Using Eq. 1.4, agent 1 then maximises his payoff $\pi_1(t = 0)$. This results in $x(t = 0) \approx 0.648$. However, this offer violates the condition in Eq. 1.3, since $u_2(0.648) = 0.352 < \pi_2(t = 1)$. Agent 1 should therefore propose a payoff-equivalent deal to agent 2 [i.e., $\pi_2(t = 0) = \pi_2(t = 1)$]. For $x \approx 0.609$ this condition is satisfied and agent 2 becomes indifferent between accepting or refusing this deal. Subgame perfection then predicts that agent 2 accepts this proposal, yielding agent 1 a payoff of ≈ 0.419. Tables 2 and 3 summarise these theoretical predictions. Notice that, in this example, Eq. 1.3 (i.e., the responder's threshold) indeed plays a role in round 1, whereas in the rounds 2 and 3 the equation does not influence the proposals made in SPE.

References

1. J. Andreoni and J.H. Miller. Auctions with artificial adaptive agents. *Games and Economic Behavior*, 10:39–64, 1995.
2. Th. Bäck. *Evolutionary Algorithms in Theory and Practice*. Oxford University Press, New York and Oxford, 1996.
3. K. Binmore. *Fun and Games*. D.C. Heath and Company, Lexington, MA, 1992.
4. K. Binmore, A. Shaked, and J. Sutton. Testing noncooperative bargaining theory: A preliminary study. *American Economic Review*, 75:1178–1180, 1985.
5. K. Binmore and N. Vulkan. Applying game theory to automated negotiation. *Netnomics*, 1(1):1–9, 1999.
6. 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.
7. D.E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, 1989.
8. G.W. Harrison and K.A. McCabe. Expectations and fairness in a simple bargaining experiment. *International Journal of Game Theory*, 25:303–327, 1996.
9. J.H. Holland. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. The University of Michigan Press/Longman Canada, Ann Arbor, 1975.
10. E. Kalai and M. Smorodinsky. Other solutions to the Nash bargaining problem. *Econometrica*, 43:513–518, 1975.
11. H. Lin and S. Sunder. Modeling bargaining behavior in ultimatum games. Working paper, 1999. Available at http://www.som.yale.edu/Faculty/sunder/research.html.
12. P. Maes, R.H. Guttman, and A.G. Moukas. Agents that buy and sell. *Communications of the ACM*, 42(3):81–91, 1999.
13. 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.

14. M. Mitchell. *An Introduction to Genetic Algorithms.* The MIT Press, Cambridge, MA, 1996.
15. 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.
16. M.J. Osborne and A. Rubinstein. *A Course in Game Theory.* The MIT Press, Cambridge, MA, 1994.
17. H. Raiffa. *The Art and Science of Negotiation.* Harvard University Press, Cambridge, MA, 1982.
18. T. Reichmann. Learning and behavioral stability: An economic interpretation of genetic algorithms. *Journal of Evolutionary Economics*, 9(2):225–242, 1999.
19. A.E. Roth. Bargaining experiments. In J. Kagel and A.E. Roth, editors, *Handbook of Experimental Economics*, pages 253–348, Princeton University Press, Princeton, NJ, 1995.
20. A.E. Roth, J.K. Murnighan, and F. Schoumaker. The deadline effect in bargaining: Some experimental evidence. *American Economic Review*, 78(4):806–823, 1988.
21. A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):155–162, 1982.
22. D.D.B. van Bragt, E.H. Gerding, and J.A. La Poutré. Equilibrium selection in alternating-offers bargaining models: The evolutionary computing approach. Technical Report SEN-R0013, CWI, Amsterdam, 2000. Presented at the 6th International Conference on Computing in Economics and Finance (CEF'2000).
23. 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.