Multi-Issue Bargaining with Multiple Encounters: An Evolutionary Social Simulation

E.H. Gerding[1*] J.A. La Poutré[1,2]

[1] *CWI, Centre for Mathematics and Computer Science*
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

[2] *School of Technology Management; Eindhoven University of Technology*
De Lismortel 2, 5600 MB Eindhoven, The Netherlands

Abstract

Negotiations have been extensively studied theoretically throughout the years. A well-known approach is the game-theoretic ultimatum game, where two agents negotiate on how to split a pie or a "dollar": the proposer (here called the seller) makes an offer and responder (the buyer) can choose to accept or reject. Usually the game is studied without taking into account backup opportunities if negotiations fail. In the real world, however, often several candidate sellers and buyers exist. In this paper a natural extension of the ultimatum game is presented, where both the sellers and the buyer can negotiate several times before their payoff is established. This way the basics of a competitive market are modelled where for instance a buyer can try several sellers before making a purchase. The game is investigated using an evolutionary simulation. We consider the outcomes when the agents know each other's number of remaining bargaining opportunities (the perfect information case), and the game where this information is not available (the imperfect information case). We find that the seller has the advantage in the first case, but that the buyer obtains more bargaining power in the latter game setting. For the perfect information case we also provide a game-theoretic analysis and compare the outcome with evolutionary results. Furthermore, the effects of search costs and allowing multiple issues to be negotiated simultaneously are investigated.

1 Introduction

In the advent of ubiquitous application of agent technology, bargaining agents are expected to play an essential role in electronic market places. Automated negotiations are therefore becoming an important field of research [3, 2, 10, 11]. The agents in the market are self-interested and often

*Corresponding author. E-mail: e.gerding@cwi.nl.

equipped with the ability to autonomously search for products and services and negotiate an agreement. Bargaining agents should also be capable of interacting with their environment and adapting to the behaviour of other agents.

In this paper we focus attention on the negotiation aspect and study the strategic behaviour of adaptive agents within a market-like setting. Within our model, the agents exchange a single good in a bilateral fashion. The agents, however, have several bargaining opportunities before reaching a final deal. Bilateral bargaining has been extensively researched, for instance in game theory [5, 15, 14, 22, 17, 16] . In contrast to for instance many applied auctions, bilateral bargaining allows for negotiations which involve not just the price of the good, but a much wider range of attributes.

We use the one-shot ultimatum game as the basic bargaining procedure for our model, a well-known approach within the field of game theory. In this game, the proposer makes an offer and the responder can only choose to accept or reject this offer. The ultimatum game has been extensively researched, both theoretically and by experiments using human subjects [15, 17, 9, 18]. Most traditional studies of this game concern an isolated pair of players, without taking into account the possibility of other (i.e., outside) opportunities. In the real world, however, players often have backup opportunities if negotiations fail.

This paper introduces an extension of the basic ultimatum game where both the proposing and the responding agents have several bargaining opportunities. This way a market place is modelled where for instance a number of sellers and buyers are available. Several issues can be negotiated simultaneously; not only the price, but also other important attributes such as delivery time, package deals, warranty, and other product-related aspects can be taken into account.

The market is computationally simulated by evolutionary algorithms (EAs). EAs are increasingly used to study the dynamic process of locally interacting, adaptive agents, as in the field of agent-based computational economics (ACE) [4, 19]. In contrast to for instance game theory, no explicit rationality assumptions are made in the EA; the agents are naive optimizers acting on limited information. In our simulation, the agents are also myopic: they do not have any forward-looking ability or memory. Nonetheless, surprisingly rational behaviour often emerges from such "low-rational" agents [21].

A key factor in the game is whether the agents know the number of bargaining opportunities of their current opponent; although the initial number of opportunities is set equal for buyers and sellers, two encountering agents can have a different so-called bargaining state. We study the effect of knowledge about an opponent's situation, and show some very interesting results. In particularly, when the number of bargaining opportunities of the opponent is hidden, results differ significantly from the ultimatum game. If the maximum number of encounters is sufficiently high, the responder's can obtain the largest share, depending on the amount of samples to determine the average payoff. This in contrast to the ultimatum game, where the proposer claims the maximum amount and the responder obtains nothing.

A general game-theoretical approach to this game is very difficult when the number of remaining bargaining opportunities of the opponent is not known. In this paper, however, we provide a game-theoretic analysis for the market game for a simpler case where the agents are perfectly informed. The concept of sub-game perfect equilibrium (SPE) is used. We compare the outcomes of the simulation with the SPE predictions using different EA settings. A very good match can be found for smaller games such as the one-shot game. Results can even be improved for longer games by reducing the amount of noise in the EA.

Furthermore, the effect of competitiveness of the game is studied. By allowing several issues to be negotiated at once, trade-offs can be made to obtain win-win situations and to reduce the competitive nature of the game. The competitiveness indicates to what extent trade-offs are beneficial for both. We study both distributive negotiations, where only a single issue is negotiated and various levels of integrative negotiations, where trade-offs between issues are possible.

Finally, the paper considers the effect of search costs in the market game. Search costs represent the amount of effort or money which is spent on finding a new negotiation partner. This is modelled by a fixed reduction of the payoffs each time an offer is rejected (and a new partner has to be found). Search costs seem to have little impact on the outcomes if the agents know the number of bargaining opportunities of their opponent. In the imperfect information case, however, even small search costs discourage the agents to explore further. Almost all deals are reached immediately, and agreements are similar to the one-shot game.

The remainder of the paper is organised as follows. In Section 2 the market game is described. Section 3 provides a game-theoretic analysis of the game in case of common knowledge of the agents' bargaining opportunities. Section 4 outlines the evolutionary simulation and Section 5 discusses the obtained results from the simulation. Lastly, Section 6 concludes.

2 Description of the Market Game

The market consists of buyers and sellers who exchange a single good through bilateral negotiations. A buyer and a seller negotiate the terms of the agreement in an ultimatum game-like setting, where the seller proposes an offer and the buyer can reject or accept the seller's offer. Whereas the ultimatum game ends after the buyer's decision, in the market game both agents have a number of additional bargaining opportunities if an agreement is not reached. Each such negotiation between a pair of agents is called an encounter. This game models a situation where for instance consumers can go to various sellers until a satisfactory deal is obtained (e.g. when buying a house, car, etc.).

The Ultimatum Game. We first describe the ultimatum game played at each encounter in more detail. The ultimatum game is frequently used in game theory as a basic negotiation model[1]. The game consists of two consecutive steps: first the seller proposes an offer, and the buyer can then either accept or reject this offer. Often, only the price is negotiated and the negotiable surplus consists of a single issue. Instead, in our model bargaining can take place over multiple issues simultaneously: not only the price, but also other value-added services such as time, quality, return policy, or other product-related aspects can be taken into account. We use n to denote the total number of issues in the offer.

An offer is expressed as a vector $\vec{o}$, where the i-th component o^i specifies the share that the seller receives for issue i if the offer is accepted. Without loss of generality, the offer for each issue is normalized between zero and one; the value zero is most beneficial for the buyer, but still acceptable for the seller, whereas a value of one is the seller's favourite, and still acceptable for the buyer. The buyer's share is therefore $1 - o^i$ for each issue i if he accepts offer $\vec{o}$.

An offer $\vec{o}$ is evaluated using an additive multi-attribute utility function. The utility is calculated by the weighed sum of the share obtained for each issue. Formally, the seller's utility function is $\vec{w_s} \cdot \vec{o} = \sum_{i=1}^{n} w_s^i \cdot o^i$, where $\vec{w_s}$ denotes the weight preferences of the seller. Similarly, the buyer's utility function is $\vec{w_b} \cdot (\vec{1} - \vec{o}) = \sum_{i=1}^{n} w_b^i \cdot (1 - o^i)$, where $\vec{w_b}$ is the buyer's weight settings. The weights are normalized and larger then or equal to zero, i.e., $\sum_{i=1}^{n} w_s^i = \sum_{i=1}^{n} w_b^i = 1$ and $w_s^i, w_b^i \geq 0$. In our simulations we assume that all sellers have the same preferences, and also that the weights of the buyers are the same. This allows us to study the effect of multiple encounters in isolation without additional complications.

Multiple Encounters. The market game extends the simple bilateral negotiations by allowing multiple encounters. If two negotiating agents reach an agreement, both agents obtain a payoff equal to their utility of the offer. In case of a disagreement, however, the agents can be matched again with newly selected opponents. The matching of two agents occurs randomly, and is explained below. Each agent initially has up to m bargaining

[1] An alternative is a finite version of the multi-round alternating-offers game [15, 14]. This game is very similar to the ultimatum game, however, if no time pressure exists; the game-theoretic outcomes are equivalent, and also previous work using an EA simulation showed little influence of the additional rounds [5].

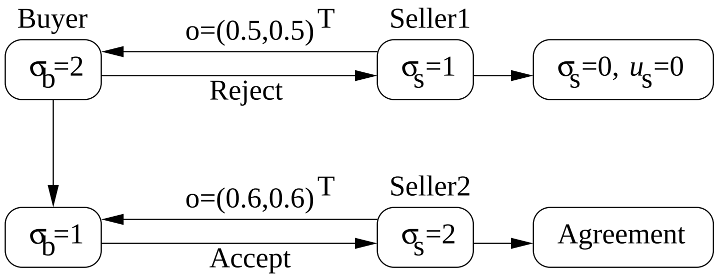

Figure 1: A two-issue negotiation game in a market with up to two bargaining opportunities ($n = 2, m = 2$).

opportunities to reach an agreement. We refer to an agent's number of remaining bargaining opportunities as the agent's *bargaining state*, denoted by $\sigma_s \in \{0, 1, \ldots, m\}$ if the agent is a seller and $\sigma_b \in \{0, 1, \ldots, m\}$ if the agent is a buyer. If an agent's bargaining state reaches zero, he obtains a disagreement payoff which is normalized to zero.

Matching procedure. There are two groups of agents: sellers and buyers. In each period, a pair of agents is randomly selected to negotiate. The agents are then first returned before a new pair is selected in the next period. We model market situation where the number of participants remains constant over time, also called a steady-state market in [14]. Therefore, if the pair reaches an agreement, their bargaining states are reset to the initial value and the agents can participate anew in the market. The same occurs if an agent's bargaining state reaches zero bargaining opportunities.

Each time an agent participates in a multiple-encounter game, he obtains a payoff value (either through agreement or by a disagreement in the final bargaining state). The payoff obtained in different multiple-encounter game can fluctuate due to variations in the bargaining states and the strategies of the randomly matched opponents. To obtain a more stable measure and a better indication of the "expected payoff", each agent plays at least r multiple-encounter games, and the final payoff is the average obtained in these r games. To maintain a steady state market, agents continue to participate until all agents have played r multiple-encounter games, even if the their final payoff has already been determined.

Example. An example of a game with $n = 2, m = 2$ is visualised in Fig. 1 from a buyer agent's perspective. The buyer, whose initial bargaining state is $\sigma_b = 2$, first encounters a seller, seller 1, with bargaining state $\sigma_s = 1$. The seller proposes $(0.5, 0.5)^T$ and the buyer refuses this offer. The seller has no more bargaining opportunities and obtains a utility u_s of zero. The buyer, on the other hand, can continue bargaining in a new encounter when matched with another opponent, seller 2. In the example this opponent has $\sigma_s = 2$ and offers $(0.6, 0.6)^T$. The buyer now accepts. Note that the matched agents can have different bargaining states, which is an important aspect of the market game. Furthermore, in the example the buyer finally accepts an offer which is worse than the first offer. Once an offer is rejected, agents cannot go back to a previous offer[2].

Note. It is important to note that in this paper both buyers and sellers *initially* have an equal number of bargaining opportunities. This way we can investigate the effect of the additional bargaining opportunities and abstract away from any other influences. This in contrast to the work in e.g. [14], where markets are studied with unequal number of buyers and sellers.

3 Game-Theoretical approach

Although the focus of this paper lies on evolutionary dynamics using computational simulations, it is also useful to derive theoretical results for perfectly rational agents and to compare these outcomes with the boundedly rational approach. Game theory has been extensively used to analyse bilateral negotiations and market situations. Game-theoretic outcomes are determined by the notion of equilibrium. A well-known equilibrium concept which is often applied for multi-stage games is the sub-game perfect equilibrium (SPE). SPE requires that a Nash equilibrium is reached in each sub game. We will apply this concept to the market game, being also a multi-stage game.

A general theoretic analysis of the above market game seems extremely difficult, and beyond the scope of this paper. We will therefore analyse a slightly simpler version by making some additional assumptions. First of all, we assume that an agent uses the same strategy in each multiple-encounter game. Note, however, that in each encounter the strategy may differ. Furthermore, we assume all agents of a specific type (i.e., buyer or seller) apply the same strategy. This assumption seems realistic for rational agents since all agents of the same type also have homogeneous preferences. Finally, and most importantly, the agents' bargaining states are common knowledge. A game-theoretical analysis seems to be very difficult if the agents have imperfect information on their opponent's bargaining state. This is even the case if the distribution of bargaining states is known and the maximum number of encounters is as small as two. We will, however, drop the second assumption in the evolutionary simulation (Section 4) and consider market games both with and without common knowledge of the bargaining states.

The Ultimatum Game. First we analyse the basic component of the market game: the one-shot ultimatum game for multiple issues. The ulti-

[2] Agents are said to have no recall [22].

matum game has been extensively studied in the past, see for instance [15, 17, 9]. It has a unique SPE where the seller claims the total share for each issue, and the buyer accepts this take-it-or-leave-it deal. This result can be obtained by applying backward induction. Intuitively, a rational buyer will accept any positive amount, which is always better than obtaining the zero payoff in case of a disagreement. The SPE is precisely the point where the buyer is indifferent between accepting and refusing.

Multiple encounters. We argue that the SPE for the game with multiple encounters and perfect information has the same outcome as for a single encounter (i.e., the ultimatum game): the seller obtains the whole share, and the buyer receives nothing. Rather than providing a detailed proof, we will show this intuitively by a simple example.

We take the simplest case where the maximum number of bargaining opportunities equals two. Consider an encounter where the bargaining states of the matched agents are $\sigma_s = 1$ and $\sigma_b = 1$ for the seller and the buyer respectively, i.e., both agents only have a single bargaining opportunity remaining. This situation is identical to the ultimatum game, where the SPE assigns the total surplus for each issue to the seller and zero to the buyer.

If agents are matched with $\sigma_s = 2$ and $\sigma_b = 1$ (i.e., the seller now has two bargaining opportunities), the situation is similar to above and an SPE exists where the buyer accepts a zero share for each issue. Note that the SPE always assigns a zero payoff to the buyer if $\sigma_b = 1$. Using backward induction, we can therefore replace the disagreement payoff for $\sigma_b = 2$ by zero. The situation for $\sigma_b = 2$ is now identical to $\sigma_b = 1$: the buyer is indifferent between accepting and refusing a value of zero and in SPE the buyer accepts this deal, independent of σ_s.

This way we have shown that the SPE in case of two bargaining opportunities equals the outcome for the ultimatum game. This analysis can be extended for longer games.

4 Evolutionary approach

Evolutionary algorithms (EAs) are powerful search algorithms based on Darwin's theory of natural selection. The evolutionary approach has taken up an enormous interest in different fields and applications. EAs have for instance been extensively used in the field of optimisation for hard problems. In recent years, more and more the evolutionary approach has been applied within the field of computational economics as a model for both social and individual decision making.

We apply an EA to the above market game and evolve the negotiation strategies of the agents. A number of related papers have demonstrated that, using an EA, artificial agents can learn effective negotiation strategies

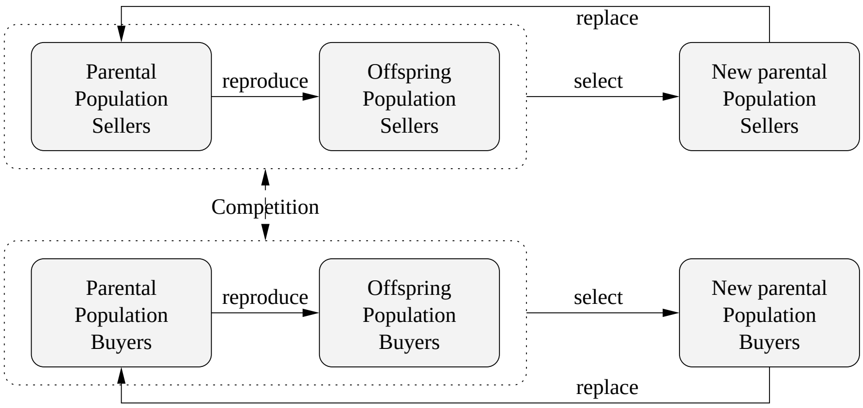

Figure 2: Iteration loop of the evolutionary algorithm. e next

in similar negotiation games [5, 20]. In [20], a comparison is also made with game-theoretical solutions to validate the EA. An important advantage of EAs is that they do not make any 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.

The implementation of the evolutionary system is outlined below. The evolutionary algorithm is first explained in Subsection 4.1. Subsection 4.2 describes how the negotiation strategies are encoded within the agents and Subsection 4.3 explains in detail the mutation operator used within the EA.

4.1 Evolutionary Algorithm

The sellers and buyers are grouped into separate populations. The system starts with randomly initialised "parental" populations of bargaining agents, having random bidding strategies. The EA is subsequently executed for a number of iterations or "generations". An iteration of the EA is depicted in Fig. 2 and consists of three stages: reproduction, fitness evaluation, and selection.

- **Reproduction.** At the beginning of each iteration, "offspring" agents are created by reproduction, see Fig. 2. An offspring agent is generated by first (randomly, with replacement) selecting an agent in the parental population, and then mutate its strategy to create a new offspring. The mutation operator is explained in more detail below.

- **Fitness evaluation.** In the second stage, the market game is executed to assess the quality or "fitness" of the agents. The parental and offspring populations are combined to form the group of sellers and the group of buyers. The fitness of an agent is his average payoff obtained in r multiple-encounter games. The market game continues until the fitness of all agents has been established (see also Section 2).

Because all agents of both populations start with their first encounter, in the first periods of the market game the opponent's bargaining states do not represent an ongoing bargaining society. To prevent so-called initiatory effects and to model an on-going bargaining society, the agent's fitness is only measured after his first multiple-encounter game (each agent thus plays at least $r+1$ games).

- **Selection.** In the third and final stage (see Fig. 2), the fittest agents from each group are selected as the new parents for the next iteration. This selection scheme is also known as $(\mu+\lambda)$-selection for evolutionary strategies (ES) [1], where μ is the number of parents and λ is the number of generated offspring. In our simulation, we take $\mu=\lambda$. The selection step completes one iteration (or "generation") of the EA.

4.2 Strategy Encoding

An agent's strategy is encoded on the so-called chromosome of the agent. The implementation of the EA is based on "evolution strategies" (ES) [1], using real-encoding of the chromosome[3]. In a game-theoretic context, a strategy is a plan which specifies an action for each history. In our model, an agent's strategy specifies either an offer $\vec{o}$ or a threshold t for each bargaining state, depending on his type (i.e., seller or buyer). The threshold determines whether an offer of the opponent is accepted or rejected: if the utility of the offer falls below the threshold the offer is refused; otherwise an agreement is reached. A similar approach has been used in [20, 17, 13].

In our simulation we distinguish between two different settings: (1) the market game with perfect information of the bargaining states, and (2) the market game where the opponent's bargaining state is unknown. In the first case, the offers and thresholds of the agents are conditional on the opponent's bargaining state. The strategy representation for this game is depicted in Fig. 3. In the latter case, the offers and thresholds are determined by the agent's own bargaining state only.

Seller	$\vec{o}(1\|1)$	$\vec{o}(1\|2)$	...	$\vec{o}(1\|m)$	$\vec{o}(2\|1)$	...	$\vec{o}(m\|m)$
Buyer	$t(1\|1)$	$t(1\|2)$	...	$t(1\|m)$	$t(2\|1)$	...	$t(m\|m)$

Figure 3: The strategies of a seller and a buyer for the market game with perfect information about the opponent's bargaining state. The offers $\vec{o}(\sigma_s|\sigma_b)$ and thresholds $t(\sigma_b|\sigma_s)$ are conditional on the bargaining state of the opponent, where $\sigma_s, \sigma_b \in \{1, \ldots, m\}$.

[3]The widely-used genetic algorithms (GAs) are more tailored toward binary-coded search spaces [8, 12, 6].

4.3 Mutation Operator

The mutation operator produces random changes in a chromosome in the following way. Each real value x_i on the chromosome position i is mutated by adding a zero-mean Gaussian variable with a standard deviation σ_i [1]. Formally, $x'_i := x_i + \sigma'_i N_i(0,1)$. All resulting values larger than unity (or smaller than zero) are set to unity (respectively zero).

We experiment with several mutation models, including (1) mutation with fixed standard deviations, (2) mutation with exponential decay of the standard deviations and (3) self-adaptive control of the standard deviations. In (1) the standard deviations are initially set to σ and remain constant over time. In (2) the standard deviations gradually diminish such that every t generations their value is reduced to half the size. We call t the half-life parameter. The third mutation model has self-adaptive control of the standard deviations; both the strategy and the corresponding standard deviations evolve at the same time. A detailed description of the latter mutation model can be found in [1, pp.71-73],[20, 5].

We mainly report results using (2) in this paper, showing a closest match with game-theoretic predictions. We briefly mention some of the results using the first and the third model.

EA settings	Parental population size (μ)	30
	Offspring population size (λ)	30
	Mutation model	exponential decay
	Initial standard deviations (σ)	0.1
	Standard deviation half-life (t)	400
	Number of generations	4000
	Number of runs per experiment	30
Game settings	Maximum number of encounters (m)	5
	Number multiple-encounter games for final payoff (r)	5
	Number of issues (n)	1
	Compititiveness for $n = 2$ (α)	0.2

Table 1: Default settings of the evolutionary simulation.

5 Evolutionary Simulation Results

This section discusses the results from the evolutionary simulation. Since the outcomes often depend on random factors, the results are averaged over 30 runs using the same settings. All default settings which are used in the experiments are listed in Table 1. Deviations from these values are mentioned in the text and in the captions.

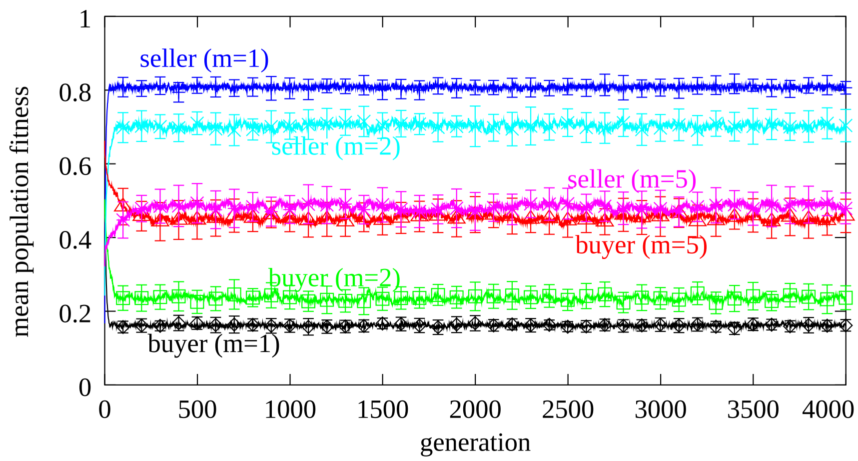

Figure 4: Development of the mean fitness of the seller and buyer populations for $m = 1$,$m = 2$ and $m = 5$ and using a fixed mutation model with $\sigma = 0.1$. Other settings are shown in Table 1. Results are averaged over 15 runs and the error bars show the standard deviations.

The results are organised as follows. First, the game with common knowledge of the agent's bargaining state is studied in Subsection 5.1 and the results are compared to the game-theoretic predictions. In Subsection 5.2 introduces a measure of competitiveness for the negotiations and compares results for different levels of integrative negotiations. In Subsection 5.3 the assumption of common knowledge is dropped. Subsection 5.4 considers the effects of fixed search costs in the market game. Results with perfect and imperfect knowledge are compared. A summary of the results is given in Subsection 5.5.

5.1 Perfectly Informed Agents and Game-Theoretic Validation

We first consider a distributive (i.e., single-issue) scenario with perfect information. Figure 4 shows the simulation results for different values of m (maximum bargaining opportunities) and using a fixed mutation scheme with $\sigma = 0.1$. For $m = 1$, the results are close to SPE outcomes: the seller obtains almost the entire surplus. However, as m increases, the deviation from SPE outcomes becomes more pronounced.

The explanations for the disagreement with game-theoretic results are two-fold. First of all, longer games (i.e. with more encounters) require a longer chromosome, and the optimization problem therefore becomes more complex. I.e., the space of possible solutions increases. Another factor which depreciates the results is the noise introduced by the mutation operator. The noise affects longer games more than shorter ones. In longer games, an

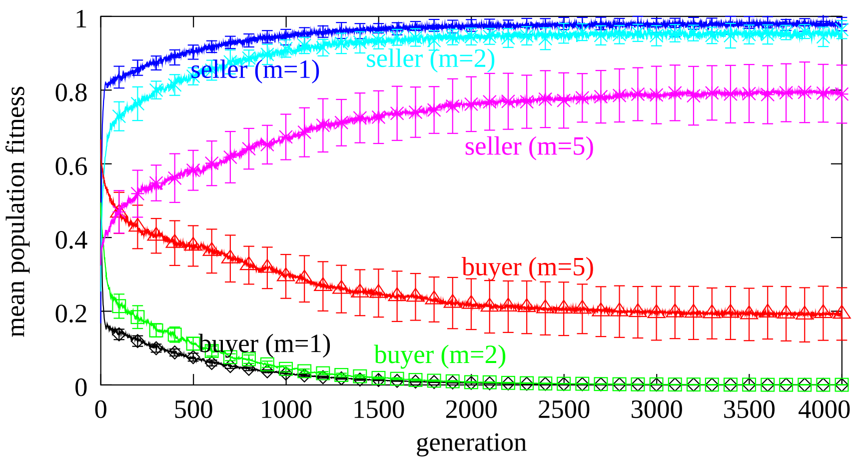

Figure 5: Development of the mean fitness of the seller and buyer populations for $m = 1$,$m = 2$ and $m = 5$ and using exponential decay of the mutation standard deviations. Results are averaged over 15 runs of the same experiment.

agent has a higher probability of encountering an opponent with suboptimal behaviour. Buyers, for instance, will learn not to accept extreme deals, but to try their luck in the next encounter. In the long-term, this also affects the offers made by the sellers, and the effect carries on to the buyer's thresholds in earlier encounters.

To test whether the noise indeed influences the long-term fitness of the agents, we apply the exponential decay mutation model, where the mutation rate is slowly diminished. The results in Fig. 5 show a much better match with game-theoretic SPE expectations. Furthermore, lowering the number of random perturbations stabilises the long-term results. We also ran simulations with lower fixed standard deviation σ, and we used the model with self-adaptive control of the standard deviations. The model with exponential decay, however, showed a much better match with SPE outcomes, also when varying other settings (e.g. the number of issues). In what follows, we therefore use this mutation model as the default setting.

Although a better match is found by optimizing EA parameters such as the mutation model, increasing m still results in less extreme deals. We also tried different population sizes μ and λ and various values for r. Only minor improvements were found, however. The outcomes of the EA are therefore quite robust. The evolutionary approach has lately been proposed as an explanation for the discrepancy between game-theoretic outcomes and laboratory experiments with human subjects, for instance in [17]. The evolutionary results could be considered more "fair" and closer to real-world outcomes, where for instance an extreme split of the surplus is usually not found [17, 9].

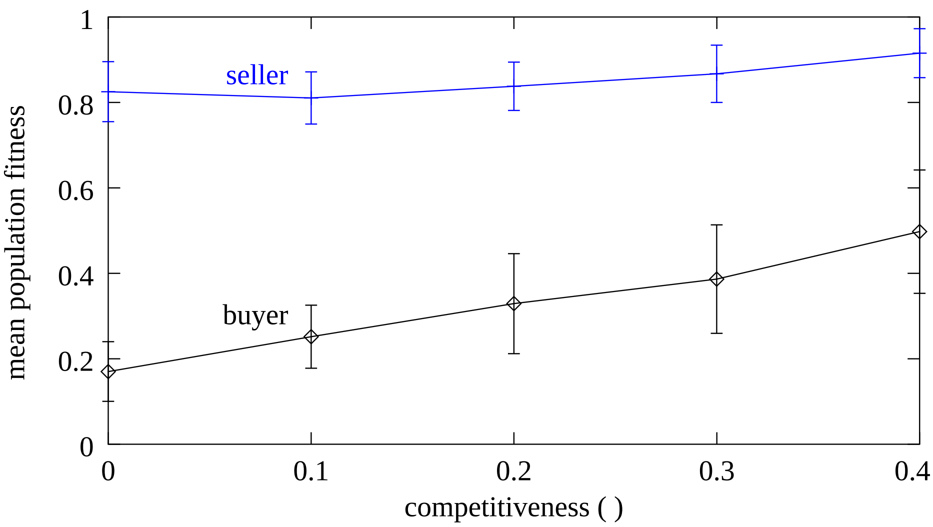

Figure 6: Fitness outcomes after 4000 generations for the seller and the buyer populations for $n = 2$ and different levels of competitiveness α.

Interestingly, in [17] a less extreme outcome is also found for the one-shot ultimatum game using an evolutionary simulation, even though the setup is very similar. In [17] "proposers do not play a sub-game perfect strategy but instead 'learn' to make offers of about 20 to 25 % of the total amount to their opponents". The main distinction is the selection scheme used; in [17] selection and reproduction of a strategy depends on a probability equal to the relative fitness. In our case, reproduction occurs randomly and selection is deterministic (see Section 4.1). Previous experiments [20] also showed inferior matching results when fitness proportional selection was used.

5.2 Perfect Information and Integrative Negotiations

Negotiations are called integrative when mutually beneficial solutions are available [7]. In the simulation we use a simple model with two issues where the buyers and sellers have opposing weight settings. More formally, the weight settings can be described as $\vec{w}_s = (0.5 - \alpha, 0.5 + \alpha)^T$ and $w_b = (0.5 + \alpha, 0.5 - \alpha)^T$ for the seller and buyer respectively, where $\alpha \in [0.0, 0.5]$ is the degree of competitiveness. When the parameter α is set equal to 0, negotiations are purely distributive (and equivalent to a single issue); if $\alpha = 0.5$ there is no competition at all. We use $\alpha = 0.2$ as the default value. Note that the maximum social welfare, i.e. the maximum total utility that a seller and a buyer can achieve together equals $2 \cdot (0.5 + \alpha)$, where each agent obtains $(0.5 + \alpha)$.

Results for $m = 5$ and different values of α are shown in Figure 6. It is interesting to see that, in case of mutually beneficial opportunities, the

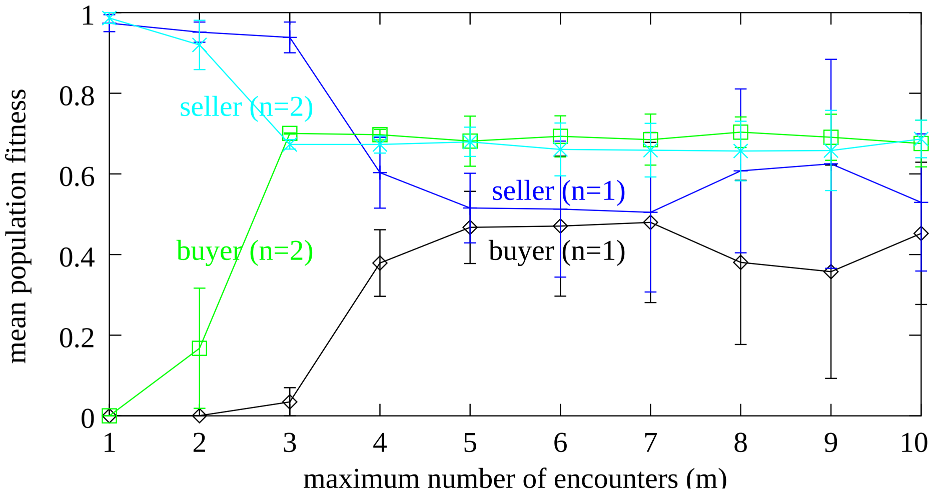

Figure 7: Average long-term fitness outcomes after 4000 generations for different values of m (the maximum number of encounters). Both distributive ($n = 1$) and integrative ($n = 2$,$\alpha = 0.2$) negotiation results are shown.

boundedly rational agents find win-win situations. The buyer particularly benefits from this integrative setting.

5.3 Imperfect information

In this section we examine the results when the agents do not share the knowledge of their opponents' bargaining states; the agents only know their own bargaining state. The agents do implicitly learn the distribution of bargaining states in the opponent's population. However, this distribution is endogenously determined by the strategies of the agents. The strategies, in turn, adapt to the to the distribution of the bargaining states. This cyclic behaviour is one reason why theoretical analysis is difficult.

Long-term results for distributive and integrative experiments with varying m and $r = 5$ are shown in Fig. 7. Notice the large dissimilarity with the imperfect information case (Fig. 5): outcomes are now symmetric if $m = 5$. When negotiations are integrative, symmetric (win-win) outcomes are reached even for $m = 3$. Furthermore, results are much more stable (i.e., the standard deviations are smaller) in case of integrative negotiations. We note that these results are obtained even though the chromosome in the imperfect information case is much shorter and the search space is reduced.

We also experiment with other values of r, i.e., the number of outcomes or "samples" used to determine the final payoff. If r is small, the seller obtains the largest payoff, see Fig. 8. If r becomes larger, however, the buyer has the advantage. By increasing r the stochasticity of fitness value is reduced and the fitness becomes a more accurate measure of the expected

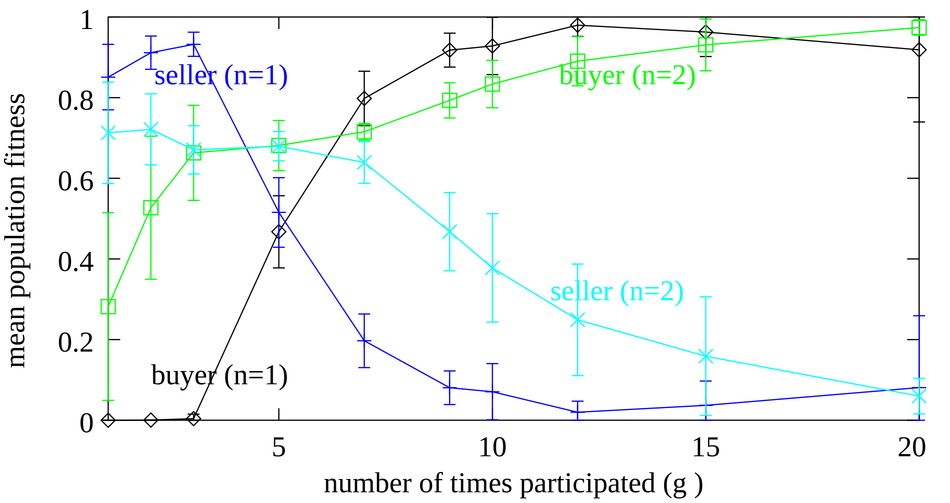

Figure 8: Experimental results after 4000 generations for different values of r (the number of times an agent participates in the market game). Results for a distributive setting are shown ($n = 1$) as well as an integrative setting ($n = 2$,$\alpha = 0.2$). Notice that the average buyer fitness becomes higher as r increases.

payoff (see Section 2). Figure 8 shows the results for various r, both for a distributive and integrative negotiation setting where $m = 5$. As in the game with perfect information, an extreme result is again observed if r becomes large. Now the *buyer* obtains the largest share. This occurs both for distributive and integrative negotiations, although r has less influence in the latter case[4].

This can be explained as follows. Recall from Section 3 the game-theoretic analysis using backward induction. In the buyer's final encounter, a rational buyer will accept any offer. The seller can anticipate the buyer's behaviour in the buyer's last encounter if he knows his bargaining state; the seller will then ask the maximum amount. He cannot know what the buyer will do, however, if the bargaining state is unknown. Because the first step of the induction cannot be made, the outcome in earlier encounters cannot be derived either. In the simulation we observe in the beginning of a run that the buyer accepts any deal in his last encounter. However, this no longer affects the behaviour in earlier encounters.

With imperfect information, a buyer will no longer accept a bad deal because there is a good chance of getting a better deal in the next encounter. Consider a seller in his last encounter. The seller will not propose a bad offer, since a buyer is likely to refuse it. To prevent a disagreement, the seller will

[4]Note that the average distribution of bargaining states could also change due to an increased r, which could also cause the observed results. However, experiments where several outcomes are omitted before the final payoff is measured show no significant deviations.

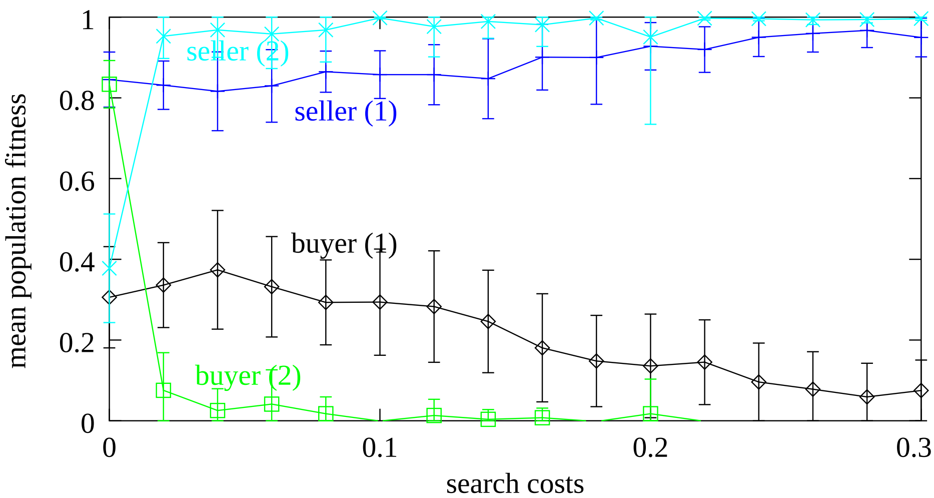

Figure 9: Average long-term population fitnesses for different search costs. Here, $n = 2$ and $r = 10$. In experiments (1) the agents have perfect information of their opponent's bargaining opportunities, whereas in (2) the bargaining states are not common knowledge.

propose a higher share to the buyer in his last encounter. A similar offer is then made in earlier encounters. In the process of evolution, these offers slowly increase until the buyer obtains the entire share.

It is interesting to see that if r is around 5 the agents agree on an even split on average. This "fair" division remains if other parameters of the EA (such as the mutation scheme) are varied. In game theory, the payoff is based on the expected outcome, i.e., the average outcome when the game is repeated infinitely. In reality, however, learning is based on a limited number of samples. As is shown here, the number of samples has a large impact on the outcome when strategies are based on imperfect information.

5.4 Search costs

In the experiments discussed so far, there is no pressure for the agents to reach an early agreement; as long as bargaining opportunities are abundant, a disagreement has no direct impact on an agent's payoff. In this section we introduce search or negotiation costs into the market game. For each time an offer is rejected, the final payoff of both agents is reduced with a fixed amount β. We consider the case where both agents have equal search costs.

Perfect information. If search costs are excessive, e.g. larger or equal to one, the game is effectively the single-encounter game; the seller claims the whole surplus by a take-it-or-leave-it deal in very the first encounter. A similar outcome, however, is already obtained for much smaller search costs.

Figure 9 shows the fitness results after 4000 generations for different search costs, where $m = 5$ and $n = 2$ with default competitiveness ($\alpha = 0.2$). The fitness of the two agents gradually grows apart as the search costs increase. Notice that search costs below 0.10 hardly have any impact on the division of the surplus in this case.

It is interesting to note that, even though most agreements are reached immediately, some agreements are still reached in later encounters. The agents keep exploring other opportunities. These later encounters play an important role in the division of payoffs if search costs are low. The buyer, for instance, still obtains a considerable share of the surplus in the first encounter if the maximum number of bargaining opportunities m is equal to five. This share increases if $m > 5$.

Imperfect information. Figure 9 shows a much larger impact of the search costs if the bargaining state of the opponent is unknown. Whereas with zero search costs the buyer has the advantage, the sellers claims the largest share even if search costs are very small (e.g. 0.01) and equal for both agents. Results are robust for different settings of the EA.

With zero search costs, a buyer has nothing to lose by waiting for a better payoff in the next encounter. In the simulation we observe that agreements are distributed over the various encounters. Both buyers and sellers, however, are stimulated to reach agreements early in case of search costs. The final encounter of the seller is therefore almost never reached, removing the advantage for the buyer. During the course of a run, agreements are steadily reached in earlier encounters and at the same time the seller's share increases. Whereas in the perfect information case the additional bargaining opportunities still play an important role, these no longer affect the payoff in case of imperfectly informed agents.

5.5 Summary of results

Results are sensitive to a number of parameter settings of the EA such as the mutation model and the selection scheme. Game-theoretic SPE outcomes are used as a benchmark and to tune the EA in order to obtain a best match. In particularly we find that deviations from SPE are mainly caused by noise due to exploration of new strategies. If we slowly reduce the noise by an exponential decay of the mutation standard deviation, in the long term a good match can be found for smaller games.

A distinct difference between the outcomes of the market game using perfectly informed and imperfectly informed agents is observed in the simulation. If the agents know one another's bargaining opportunities, the outcomes resemble the SPE predictions when m is small: the seller obtains almost all of the surplus. As m becomes larger, the space of possible deals becomes too large to be fully explored. The results are then much less extreme and the division of surplus can be considered more “fair”. These outcomes are quite robust to different EA setting. Less extreme deals are also made if negotiations are integrative. The EA results can be considered more realistic for boundedly rational agents with only limited computational capacity.

If the bargaining state is not known, on the other hand, the seller can no longer anticipate the buyer's behaviour and the buyer acquires the advantage; he can reject unpropitious offer of the seller and wait for a better offer in the next encounter. The seller gives in to avoid a disagreement in the final round.

In case of imperfect information, the outcomes are much more sensitive to the number of samples used to determine the final payoff. If an agent's final payoff is averaged over many outcomes, the buyer obtains almost the entire surplus (when m is at least 5). If the number of samples around 5, on the other hand, the agents reach an equal fitness on average. With even less samples, the seller obtains the largest share.

The effects of fixed search costs per encounter have also been investigated. The search costs stimulate both buyers and sellers to reach an early agreement. In case of perfectly informed agents, the results gradually move towards the outcome for the single-encounter game. If search costs are small, the agents keep exploring other opportunities. In case of imperfect knowledge even small search costs cause an extreme split of the surplus as before, where the seller obtains the entire partition. The remaining opportunities no longer affect the outcome.

6 Conclusion

We study the evolutionary dynamics of a market-like game, where a seller sells a single good and has several opportunities to do so. At the same time, a buyer wishes to buy an item by trying several sellers. The terms of an agreement are negotiated using an ultimatum-like game, where the seller proposes an offer and the buyer can choose to accept or reject the offer. The game is extended to allow for multiple opportunities for both the seller and the buyer if the deal is rejected. This way a competitive market is modelled. To investigate the effect of multiple encounters only, settings such as the maximum number of encounters and the population sizes are set equal for both buyers and sellers.

The game-theoretic outcome using sub-game perfect equilibrium (SPE) for the one-shot ultimatum game predicts an extreme splits of the surplus: the seller obtains the whole surplus whereas the buyer obtains his disagreement payoff. We extended the analysis for multiple encounters with perfect information of the opponent's number of bargaining opportunities and found the same outcome. A theoretical analysis seems to be very difficult, however, if the bargaining states of the agents are not common knowledge.

We first compared the evolutionary results with the game-theoretical outcomes for the perfect information game. If the maximum number of bargaining opportunities is small, both approaches yield very similar results. In larger games or when the negotiations become less competitive, the EA shows deviating outcomes due to the limited computational capacity of the agents; the payoffs of the buyer and the seller become less extreme and can be considered more fair and realistic than the game-theoretic outcomes.

A large impact of the additional bargaining opportunities is found if the agents have no information on their opponent's number of bargaining opportunities, even though both buyer and seller initially have equal opportunities available. In that case, the seller cannot predict the buyer's response and gives in to avoid a disagreement.

The outcomes are however sensitive to the number of samples available to determine an agent's final payoff. If the multiple-encounter game is played more often, an agent's average payoff more accurately represents the expected payoff. If that case, the buyer obtains almost the entire surplus for each issue. If the number of samples is very low, on the other hand, the seller obtains the largest share. A "fair division, where both agents obtain the same payoff, can also be obtained, in particularly when negotiation are integrative. These results are quite robust to different settings of the EA.

The influence of search costs also varies depending on the information available to the players. Small search costs have little influence on perfectly informed agents. In the incomplete information case, the effect is much more present: even small search costs result in extreme splits, similar to the one-shot ultimatum game.

We observe a large impact of the additional opportunities, particularly if the agents are uninformed about their opponent's number of future opportunities. In situations where multiple sellers and buyers are available, and the agents have no search costs, the outcomes of the ultimatum game therefore do not seem to hold. If agents endure even small search costs, however, results seem to indicate a similar outcome to the one-shot ultimatum game.

References

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

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

[3] F. Dignum and U. Cortés, editors. *Agent-Mediated Electronic Commerce III, Current Issues in Agent-Based Eletronic Commerce Systems.* Springer-Verlag, 2001.

[4] J.M. Epstein and R. Axtell. *Growing Arificial Societies: Social Science from the Bottom Up.* The MIT Press, 1996.

[5] E.H. Gerding, D.D.B. van Bragt, and J.A. La Poutré. Multi-issue negotiation processes by evolutionary simulation: Validation and social extensions. Technical Report SEN-R0024, CWI, Amsterdam, 2000. Workshop on Complex Behaviour in Economics, Aix-en-Provence, France, May 4-6, 2000.

[6] D.E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Wesley, Reading, 1989.

[7] R.H. Guttman and P. Maes. Agent-mediated integrative negotiation for retail electronic commerce. In *Proceedings of the First International Workshop on Agent Mediated Electronic Trading (AMET)*, Minneapolis, May 1998.

[8] 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.

[9] H. Lin and S. Sunder. Using experimental data to model bargaining behavior in ultimatum games. Working paper, 2001. Available at http://www.som.yale.edu/Faculty/sunder/research.html.

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

[11] 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.

[12] M. Mitchell. *An Introduction to Genetic Algorithms.* The MIT Press, Cambridge, MA, 1996.

[13] 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.

[14] M.J. Osborne and A. Rubinstein. *Bargaining and Markets.* Academic Press, San Diego, CA, 1990.

[15] M.J. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT Press, 1994.

[16] M.G. Raith. *The Structure of Fair-Division Problems and the Design of Fair-Negotiation Procedures*, chapter 14. Kluwer Academic Publishers, 1999.

[17] T. Riechmann. Evolutionary learning in the ultimatum game. Computing in Economics and Finance 2001.

[18] 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.

[19] L. Tesfatsion. Agent-based computational economics: Modelling economies as complex adaptive systems. In *Proceedings of the 6th Joint Conference on Information Sciences, March 8-13, North Carolina, USA, 2002*, pages 40–43.

[20] 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).

[21] H.P. Young. *Individual Strategy and Social Structure.* Princeton University Press, 1998.

[22] R. Zwick and C.C. Lee. Bargaining and search: An experimental study. *Group Decision and Negotiation*, 8(6):463–487, 1999.