Bargaining with Posterior Opportunities: An Evolutionary Social Simulation*

E.H. Gerding[1] and J.A. La Poutré[12]

[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

Summary. Negotiations have been extensively studied theoretically throughout the years. A well-known bilateral approach is the ultimatum game, where two agents negotiate on how to split a pie or a "dollar": the proposer makes an offer and responder can choose to accept or reject. In this paper a natural extension of the ultimatum game is presented, in which both agents can negotiate with other opponents in case of a disagreement. This way the basics of a competitive market are modelled where for instance a buyer can try several sellers before making a purchase decision. The game is investigated using an evolutionary simulation. The outcomes appear to depend largely on the information available to the agents. We find that if the agents' number of future bargaining opportunities is commonly known, the proposer has the advantage. If this information is held private, however, the responder can obtain a larger share of the pie. For the first 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 [4,2,10,11]. The agents in a competitive market are self-interested and can be equipped with the ability to autonomously search for products and services and negotiate the terms of an agreement. In this paper we focus on strategic aspects of bilateral bargaining within a market-like setting.

Bilateral bargaining has been extensively researched, for instance in game theory [6,13,14]. Negotiations are often stylised using the ultimatum game, a two-stage game in which an offer is proposed by player one (the proposer) in the first stage, and the second player (the responder) can only choose to accept or reject the offer. The ultimatum game has been extensively researched, both theoretically and by experiments using human subjects [14,15].

* This research was part of the project "Autonomous Systems of Trade Agents in E-Commerce", funded by the Telematics Institute in the Netherlands.

The ultimatum game models a negotiation between an isolated pair of players. In a market setting, however, an agent's behaviour can change if future opportunities are taken into account. This paper introduces a natural extension of the basic ultimatum game in which fallback opportunities are explicitly modelled. Both the proposing and the responding agents have several bargaining opportunities with different opponents before their payoff is determined. In this way a market place is modelled where several sellers and buyers are available.

The game is further extended to allow several issues to 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. This can reduce the competitive nature of the game since trade-offs can be made to obtain win-win solutions. Furthermore, the paper considers the effect of search costs if an offer is refused and a new opponent needs to be found.

An important aspect within this setting is the information available to the agents about their opponents. We distinguish between the complete information case, where an agent's current number of future bargaining opportunities is common knowledge, and the incomplete information case, where this information is known to the protagonist but hidden from the opponent.

The complete information case can be easily approached theoretically using game theoretic sub-game perfect equilibrium given reasonable assumptions. The incomplete information case, on the other hand, seems much more difficult to analyse. We therefore apply an evolutionary simulation to investigate this setting. We also compare the evolutionary and the theoretical approach in the complete information case.

In the field of computational economics, evolutionary simulations are increasingly applied to study the dynamic process of locally interacting, adaptive agents, particularly in the area of agent-based computational economics (ACE) [5,16]. In contrast to for instance game theory, no explicit rationality assumptions are made; the agents are naive optimisers acting on limited information. In these simulations, 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.

In the paper we present the evolutionary results for the settings described above and a game-theoretic analysis for the extended game with complete information. Sub-game perfect results predict an extreme split of the surplus similar to the ultimatum game: the proposer claims the entire surplus, and the responder accepts this deal. The evolutionary outcomes show a good match with these game-theoretic results. Moreover, the simulation shows that results differ significantly if information about the opponent's future bargaining opportunities is not available: if the number of bargaining opportunities is sufficiently high, the responder now obtains the largest share.

The outcomes in the incomplete information case, however, also depend on the existence of positive search costs. Search costs stimulate agents to reach agreements early and discourage both players to exploit the additional opportunities. In the evolutionary simulation, the agreements are then similar to the one-shot ultimatum game.

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

2 Description of the Bargaining Game

The modelled market consists of buyers and sellers who exchange a single good through bilateral negotiations. Each bargaining opportunity, an ultimatum-like game is played, where a seller proposes an offer and a buyer can reject or accept the seller's offer[1]. In our model an offer consists of one or more issues. If an agreement is reached, both agents obtain a payoff equal to their utility of the offer. In case of multiple issues, the utility is calculated as the weighted sum of the share obtained for each issue. The weights determine the *competitiveness* and are explained further in Section 5.3. The utilities of the agents are normalised between 0 and 1.

Each agent initially has up to m bargaining opportunities to reach an agreement. In case of a disagreement the agents are newly matched with randomly selected opponents, until no more bargaining opportunities remain. The number of remaining bargaining opportunities we call an agent's *bargaining state*, denoted by $\gamma_s \in \{0, 1, \ldots, m\}$ for a seller and $\gamma_b \in \{0, 1, \ldots, m\}$ for a buyer. If an agent's bargaining state reaches zero, the agent obtains a disagreement payoff which is set to zero.

An example is shown in Fig. 1 from a buyer's perspective. The buyer, whose initial bargaining state is $\gamma_b = 2$, first encounters a seller, seller 1, with bargaining state $\gamma_s = 1$. The seller proposes $o = (0.5, 0.5)$ and the buyer refuses this offer. Because the seller has no more bargaining opportunities his bargaining game ends and he obtains the disagreement payoff. The buyer, on the other hand, can continue bargaining when matched with another opponent, seller 2. In the example this opponent with $\gamma_s = 2$ offers $(0.6, 0.6)$. The buyer now accepts and the bargaining game ends for both agents.

Note that although the agents initially have equal bargaining opportunities, the matched agents can have different bargaining states. Having agents

[1] Alternatively, a more complex bargaining game such as the alternating-offers game [14] involving multiple rounds can be used. Outcomes are equivalent to the ultimatum game, however, if no time pressure exists; agreements are delayed until in the final round a take-it-or-leave-it offer is made. This deadline effect was studied in [6] using an EA simulation.

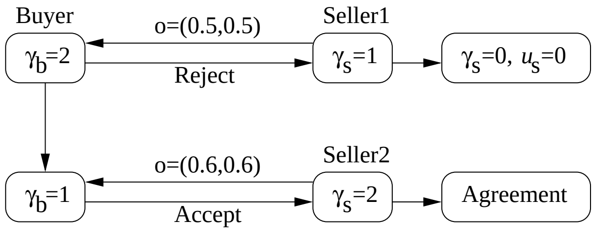

Fig. 1. A two-issue negotiation example in a market where each agent has two initial bargaining opportunities ($m = 2$).

with different states is an important aspect of the market game, particularly when agents are unaware of their opponent's remaining opportunities. Furthermore, once an offer is rejected, agents cannot go back on a previous offer[2]. Finally, for simplicity there are an equal number of buyers and sellers in the market. This in contrast to the work in e.g. [13], where markets are studied with unequal number of buyers and sellers.

3 Game-Theoretical Approach

This section considers the game-theoretic sub-game perfect equilibrium (SPE) of the above game where the agents' bargaining states are common knowledge. A game-theoretical analysis seems to be very difficult if the agents have incomplete information of their opponent's bargaining state. We will, however, drop the complete information assumption in the evolutionary approach (Section 4). In the following analysis we assume all agents of a specific type (i.e., buyer or seller) apply the same negotiation strategy. This assumption is reasonable since the preferences are homogeneous for a given type.

In case of a single opportunity, the bargaining game is reduced to the ultimatum game. The ultimatum game has a unique SPE where the seller (here the proposer) claims the total share for each issue, and the buyer (the responder) accepts this take-it-or-leave-it deal [14]. 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.

We argue that the SPE for multiple bargaining opportunities and complete information has the same outcome as the ultimatum game: the seller obtains the whole share, and the buyer receives the disagreement payoff, which is set to zero. Consider a buyer with $\gamma_b = 1$, i.e. with a final bargaining opportunity remaining. The buyer will then accept any positive amount offered by the seller. An anticipating seller will then claim the entire share,

[2] Agents are said to have no recall [18].

as in the ultimatum game, independent of γ_s. In SPE, the buyer's payoff for $\gamma_b = 1$ therefore equals zero.

If $\gamma_b = 2$, the buyer has two bargaining opportunities. Using the above, we can replace the payoff for refusing the seller's offer when $\gamma_b = 2$ by the disagreement payoff. The situation for $\gamma_b = 2$ is now equal to $\gamma_b = 1$: the buyer is indifferent between accepting and refusing a value of zero and in SPE the buyer accepts this deal, independent of γ_s. By backward induction the same holds for $\gamma_b = m$.

4 Evolutionary Approach

Evolutionary algorithms (EAs) are powerful search algorithms based on Darwin's theory of natural selection. 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. A number of related papers have demonstrated that, using an EA, artificial agents can learn effective negotiation strategies in similar negotiation games [6,17]. 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.

4.1 Evolutionary Algorithm

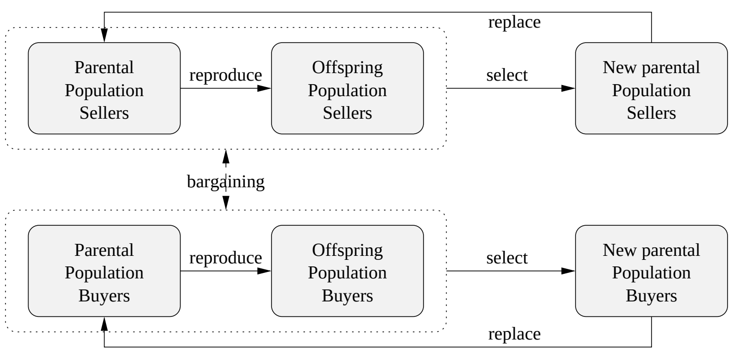

Fig. 2. Iteration loop of the evolutionary algorithm.

The evolutionary simulation works as follows. Sellers and buyers are grouped into separate populations. This way the two types of agents co-evolve. 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, depicted in Fig. 2, consists of three consecutive stages: reproduction, fitness evaluation, and selection.

In the reproduction stage, offspring agents are 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.

In the second stage, the outcomes of a series of bargaining games assess the quality or "fitness" of the agents. The parental and offspring populations are combined to form the group of seller and the group of buyer agents. Each bargaining opportunity, two agents are randomly selected (with replacement) and play the one-shot game. An agent obtains a payoff in case an agreement is reached or if no more opportunities are available. Because an outcome depends on many random factors, each strategy is evaluated a number of times and the fitness is the average of r payoff values. The parameter r is called the evaluation frequency. This way the fitness becomes a more accurate measure of the expected payoff.

Because both buyers and sellers start with the same bargaining state, in the first periods 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, a strategy's fitness is only measured after the first payoff is determined. A strategy is thus evaluated at least $r+1$ times. Furthermore, we model a market situation where the number of agents remains constant over time, also called a steady-state market in [13]. Therefore, once the fitness of a strategy has been established, the strategy can still be selected to play again but its fitness is no longer affected by the outcome. The bargaining games are continued until the fitness for each strategy has been established.

In the third and final stage (see Fig. 2), the fittest strategies (i.e., with the highest average payoff) 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 $\mu = \lambda$.

4.2 Strategy Encoding

An agent's strategy is encoded on a so-called chromosome. The implementation of the EA is based on "evolution strategies"(ES) [1], using real-encoding of the chromosome[3]. The chromosome specifies either an offer or a threshold for each bargaining state, depending on the type of the agent (i.e., seller or buyer). The threshold determines whether an offer of the opponent is accepted or rejected: if the utility falls below the threshold the offer is refused; otherwise an agreement is reached. A similar approach was used in [6,17].

[3] The widely-used genetic algorithms (GAs) are more tailored toward binary-coded search spaces [9,12,7].

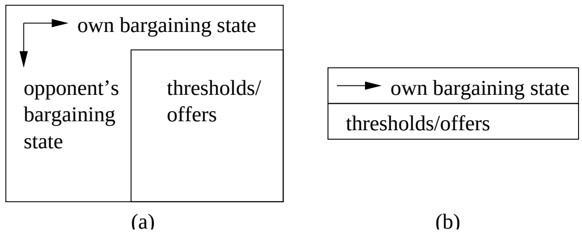

Fig. 3. Strategy representation for the complete information case (a) and incomplete information case (b).

We distinguish between the complete information setting and the incomplete information setting (see Section 2). The strategy representation for each setting is schematically depicted in Fig. 3. In the incomplete information case (Fig. 3b), different offers or thresholds are specified for different bargaining states of the agent. In case of complete information (Fig. 3a), the offers and thresholds also depend on the opponent's bargaining state.

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 [1]. Formally, $x_i' := x_i + \sigma N_i(0,1)$. All resulting values larger than unity (or smaller than zero) are set to unity (respectively zero). The standard deviation is initially set to 0.1 and gradually decays such that every t generations their value is reduced to half the size. We call t the half-life parameter. Although other mutation models such as self-adaptive control of the standard deviations [1,6] were tried as well, this model showed a closest match with game-theoretic results. Due to lack of space, we exclude the results of other models in this paper.

Parental population size (μ)	30
Offspring population size (λ)	30
Initial standard deviations (σ)	0.1
Standard deviation half-life (t)	400
Number of generations	4000
Number of runs per experiment	30
Strategy evaluation frequency (r)	20

Table 1. Default settings of the evolutionary simulation.

5 Evolutionary Simulation Results

The results are organised as follows. First, the game with complete information is studied in Subsection 5.1 and the results are compared to the game-theoretic (SPE) predictions. Subsection 5.2 studies the incomplete information case. Subsection 5.3 introduces a measure of competitiveness for multi-issue negotiations and compares results for different levels of integrative negotiations. Finally, in Subsection 5.4 considers the effects of fixed search costs in the market game.

5.1 Game-Theoretic Validation

This section considers a competitive (i.e., single-issue) scenario with complete information of the agents' bargaining opportunities and compares the evolutionary algorithm (EA) outcomes to SPE predictions. Default parameter settings for the EA are shown in Table 1. Note that because of random fluctuations, the EA results are averaged over 30 runs using the same settings.

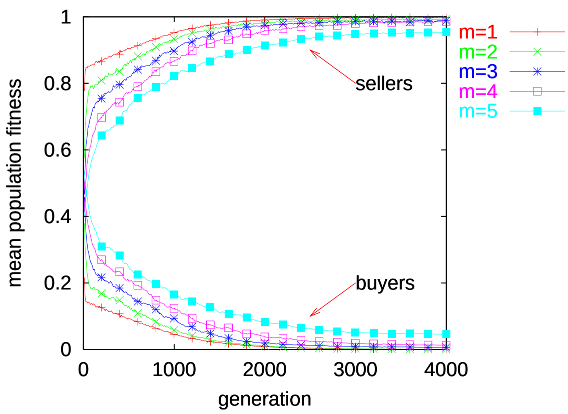

Fig. 4. Development of the mean fitness (averaged over 30 runs) for complete information setting with varying initial number of bargaining opportunities (m).

In SPE the share of the buyers is zero and the sellers obtain the whole surplus in case the initial number of bargaining opportunities of the players is equal and finite, and the bargaining state of the opponent is common knowledge (see also Section 3). Figure 4 shows the EA outcomes for different values of m (initial bargaining opportunities). The results indicate an almost perfect match between evolutionary outcomes after 4000 generations and game-theoretic outcomes, particularly when m is small.

For larger values of m we find that, using the same EA parameter settings, the evolutionary outcomes become somewhat less extreme. This is because as m becomes larger, the complexity of the problem increases due to a larger search space, making learning by an EA more difficult. However, a better match for larger values of m also appears by adjusting EA parameters, such as the evaluation frequency and the population size, to handle the increased complexity. We will not give the details here, due to space limitations. Henceforth, we present only experiments using uniform settings in this paper. We refer the interested reader to previous research [17,6], in which different EA settings are systematically studied for an alternating-offers bargaining game.

5.2 Incomplete Information

We now examine the results when the agents do not know their opponent's bargaining states; the agents only know their own bargaining states. Although no explicit information is available, the agents implicitly learn the distribution of the bargaining states in the opponent's population. This distribution is endogenously determined by the strategies of the agents. The strategies, in turn, adapt to the distribution of the bargaining states. This complex interaction is one reason why theoretical analysis is difficult. An EA, on the other hand, is well suited to find outcomes that emerge from such local interactions.

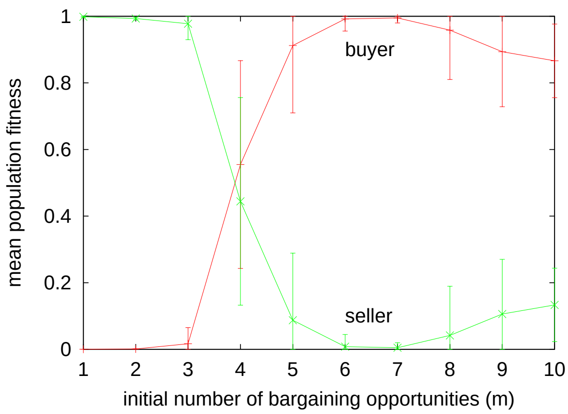

Fig. 5. Results after 4000 generations (averaged over 30 runs) for incomplete information settings with various m.

Results produced after 4000 generations of the EA for the incomplete information case are shown in Fig. 5, for different values of m (the initial number of bargaining opportunities). These results are averaged over 30 runs. The error bars indicate the standard deviation. Whereas in the complete information case the seller obtains almost the entire surplus, the responder (i.e., buyer) has the best bargaining position in the incomplete information case (see Fig. 5). This holds as long as the initial number of bargaining opportunities are sufficiently large (i.e., ≥ 5). Note that these results are obtained even though the buyers' and sellers' initial settings are equal.

The results can be explained as follows. If the buyer is in her final state, she will accept any deal (as in the ultimatum game). In other states, however, the buyer can try to find a better deal elsewhere. Consider a seller in his last bargaining state. Because he does not know the buyer's bargaining state, he can no longer anticipate the buyer's behaviour. In order to prevent a disagreement, the sellers will then concede in the last bargaining state. The expected payoff in case of a disagreement and the offers in earlier bargaining states will then also decrease. After many generations, the simulation converges to an outcome where the seller concedes almost his entire surplus in each bargaining state. We also observe that the seller concedes slightly less if he has more bargaining opportunities remaining, resulting in less extreme deals if m becomes large, as shown in Fig. 5.

In the incomplete information setting the first-mover (here the seller), has no information about his opponent. The responder, on the other hand, can make a relatively more informative decision based on the seller's offer. Whereas in the ultimatum game the proposer seems to dominate the outcome, a more competitive setting allows the responder to obtain a considerable advantage. This result, however, holds only if the number of bargaining opportunities is finite. Furthermore, the players incur no costs for refusing a deal. As we will show in Section 5.4, even slight costs completely change these results.

In the above experiments, each strategy is evaluated $r = 20$ times. We find that if r is much lower the fitness measure becomes exceedingly stochastic and as a result outcomes are less extreme. If a high value of r is used, the average utility becomes a more accurate measure of the expected payoff (see Section 4.1). In the experiments we find that more extreme outcomes are found by increasing r, particularly when m is large and the final outcomes depend on a large number of interactions. Note that, in game theory, the payoff is based on the expected outcome, i.e., the average outcome of a strategy evaluated infinitely often.

When the number of initial bargaining opportunities is set higher than three, a sudden transition in the long-term outcomes can be observed in Fig. 5: up to $m = 3$, the seller obtains almost all, whereas the buyer obtains the largest share if $m > 3$. By increasing m, the number of possible states increases, making the buyer's behaviour less predictable for the seller. The value for which the transition occurs depends on game parameters such as r and the competitiveness of the negotiation. The latter will be discussed further in the next section.

5.3 Integrative Negotiations

An advantage of bilateral negotiation is the ability to negotiate complex contracts with several issues. When mutually beneficial solutions are available, negotiations are called *integrative [8]*. We consider integrative two-issue negotiations in this section and introduce the notion of competitiveness. We show that the information in the integrative case has a very similar impact as in the competitive case. Due to increased complexity, however, the evolutionary results are less extreme when the number of bargaining opportunities is large.

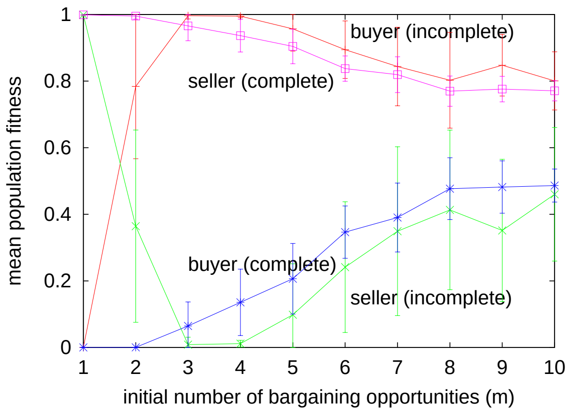

Fig. 6. Mean long-term outcomes for two-issue negotiations and $\alpha = 0.2$.

The utility of an offer is an additive, weighted function of the share obtained for each issue (see also Section 2). The weights for sellers and buyers for the two issues are $w_s = (0.5 - \alpha, 0.5 + \alpha)$ and $w_b = (0.5 + \alpha, 0.5 - \alpha)^T$ respectively, where $\alpha \in [0.0, 0.5]$ is the so-called degree of *competitiveness*. When the parameter α is set equal to 0, negotiations are purely competitive; if $\alpha = 0.5$ there is no competition at all. 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 $\alpha = 0.2$ are visualised in Fig. 6. The results show that, as in the competitive case, a transition occurs to a buyer-dominated outcome for sufficiently large m and incomplete information. We find, however, that this transition already occurs when $m = 2$ (see Fig. 6). Only two bargaining opportunities are needed to obtain an advantage for the responder, as supposed to four in the single-issue game (Fig. 5).

Figure 6 also shows a less extreme split compared to competitive negotiations, particularly for large m. This occurs firstly since the strategy search space is increased (a value for each issue needs to be learned), making learning more difficult. Moreover, the win-win possibilities are fully exploited: if one of the agents slightly concedes, the other agent can obtain a relatively large gain by negotiating a Pareto-efficient deal. This effect becomes stronger as α increases. In the extreme case, where $\alpha = 0.5$, both agents can obtain the full surplus without any concession.

Note that the EA parameters are fixed for the various game settings. As we mentioned in Section 5.1 we can adjust the parameters to handle more complex bargaining settings as a result of a larger m and an increased number of issues. By increasing the population size and adjusting other parameters of the EA, we obtain results which are closer to game-theoretic predictions.

5.4 Search Costs

We further extend the bargaining game in this section and introduce search or negotiation costs each time an offer is refused and agents engage in a new negotiation. These costs can represent the amount of money, time, or effort that an agent may incur for finding a new opponent. It is shown theoretically that if buyers have search costs, the sellers charge monopolistic prices in equilibrium [3, Ch.7]. We consider the impact of search costs on the bargaining game where both buyers and sellers have equal search costs β. The final utility is reduced by fixed search costs β for each new bargaining opportunity. Only the first bargaining opportunity has no costs.

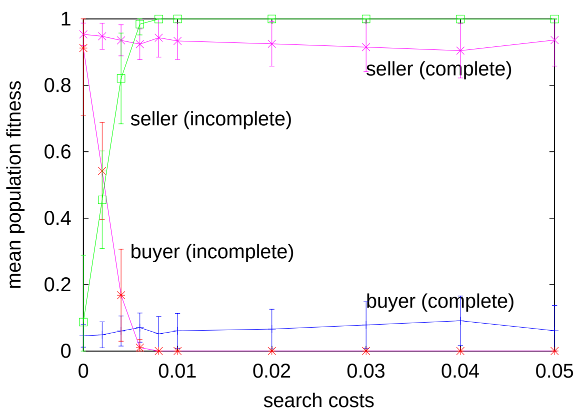

Fig. 7. Mean long-term results as a function of the search costs (β) for $m = 5$.

Evolutionary outcomes for the complete and incomplete information settings with different search costs are depicted in Fig 7. Negotiations are competitive and buyers and sellers each have 5 initial bargaining opportunities. Search costs seem to have little impact on the fitness in the complete information case; variations are not statistically relevant. Although the fitness does not change, the actual behaviour of the agents does: most agreements are now reached immediately. Without search costs, agreements reached are distributed over the various bargaining states.

In the incomplete information case, on the other hand, even small search costs have a drastic impact on the fitness of the agents, see Fig. 7. The sellers claim almost the entire 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. These outcomes are consistent with economic theory, which states that prices become monopolistic even if search costs are infinitely small.

As in the complete information case, both buyers and sellers are stimulated to reach agreements early in case of search costs. The final opportunity of the seller is therefore almost never reached, removing the advantage for the buyer. The game changes from a game with incomplete information, to a game where almost all players complete a deal in their first bargaining opportunity. Now the seller can again claim the entire surplus as in the one-shot game.

6 Conclusion

We study the evolutionary dynamics of a market-like game in this paper, 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. We furthermore investigate multi-issue integrative negotiations and the effects of search costs if a disagreement occurs.

The game-theoretic outcome using sub-game perfect equilibrium (SPE) for the one-shot ultimatum game predicts an extreme split of the surplus: the seller obtains the whole surplus whereas the buyer obtains her disagreement payoff. We extend the analysis for multiple bargaining opportunities with complete information of the opponent's bargaining state and find an equivalent outcome. A theoretical analysis seems to be very difficult, however, if the bargaining states of the agents are not common knowledge. An evolutionary simulation, on the other hand, is very well suited to investigate such games with incomplete information.

We first compare the evolutionary results with the game-theoretical outcomes for the game with complete information to validate the EA approach. If the number of bargaining opportunities is small, a very good match is found. In larger games or when the negotiations become less competitive, the EA shows somewhat deviating outcomes due to larger search space and the limited computational capacity of the EA. We note that we mainly report experiments using uniform EA settings in this paper. However, adjusting EA settings appear to improve results even further for more complex games.

The evolutionary simulation shows a large impact of the additional bargaining opportunities if the agents have no information on their opponent's number of future opportunities. Whereas in the complete information game the seller dominates the market, the buyer is better off in the incomplete information setting, as long as the number of bargaining opportunities is sufficiently high. By increasing the initial number of bargaining opportunities a sudden transition is observed where the buyer obtains the largest share instead of the seller. This occurs because the seller can then no longer anticipate the buyer's response and gives in to avoid a disagreement.

A similar large impact is found for two-issue integrative negotiations. At the same time, integrative negotiations produce less extreme evolutionary outcomes, both in the game with complete and incomplete information, particularly if the number of initial bargaining opportunities is large. This mainly occurs since the space of possible deals increases. Moreover, the agents find win-win situations which benefit one agent without affecting the payoff obtained by the opponent.

An integrative setting also already affects small games with incomplete information: we find that for certain settings, a transition from a seller to a buyer dominated payoff occurs even in case both agents merely have two initial bargaining opportunities, whereas in the competitive case more bargaining opportunities are needed to achieve the same result.

We also study the effect of search or negotiation costs in case a negotiation fails and the agent needs to find a new opponent. Search costs induce players to reach an agreement in the very first bargaining opportunity. This changes an incomplete information game into an ultimatum-like game with only a single bargaining opportunity. Even very small search costs result in an extreme split where the seller obtains almost the entire share, similar to the ultimatum game outcome. This is consistent with economic theory which states that even infinitely small search costs produce monopolistic prices.

In this paper we have shown that evolutionary simulations are extremely useful to investigate negotiations with incomplete information, which are unwieldy to analyse theoretically. Using evolutionary algorithms, we can simulate complex interactions involving a large number of agents, as is the case in bargaining with posterior opportunities. It is interesting to further refine the model to specific real-world settings, where for instance agents have incomplete information about their own future number of bargaining opportunities. Another interesting extension is allowing agents to return to previously encountered opponents.

Acknowledgements

We are grateful to all the participants in the EXYSTENCE thematic institute 2002 at Trieste, Italy for their feedback, in particularly Fernando Vega-Redondo. We also would like to thank Koye Somefun at the CWI for the useful insights and valuable discussions.

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. S.Y. Choi, D.O. Stahl, and A.B. Whinston. *The Economics of Electronic Commerce.* Macmillan Technical Publishing, 1997.
4. F. Dignum and U. Cortés, editors. *Agent-Mediated Electronic Commerce III, Current Issues in Agent-Based Eletronic Commerce Systems.* Springer-Verlag, 2001.
5. J.M. Epstein and R. Axtell. *Growing Arificial Societies: Social Science from the Bottom Up.* The MIT Press, 1996.
6. E.H. Gerding, D.D.B. van Bragt, and J.A. La Poutré. Multi-issue negotiation processes by evolutionary simulation: Validation and social extensions. *Computational Economics*, forthcoming. A preliminary version appeared as Technical Report SEN-R0024, CWI, Amsterdam, 2000.
7. D.E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Wesley, Reading, 1989.
8. 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.
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. 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. M.J. Osborne and A. Rubinstein. *Bargaining and Markets.* Academic Press, San Diego, CA, 1990.
14. M.J. Osborne and A. Rubinstein. *A Course in Game Theory.* The MIT Press, Cambridge, MA, 1994.
15. 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.
16. 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*, pages 40–43, 2002.
17. 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).
18. R. Zwick and C.C. Lee. Bargaining and search: An experimental study. *Group Decision and Negotiation*, 8(6):463–487, 1999.