

Bilateral Bargaining With Multiple Opportunities: Knowing Your Opponent's Bargaining Position

Enrico H. Gerding and Han La Poutré

Abstract—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 surplus 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 modeled, 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 remaining 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 surplus. For the first case we also provide a game-theoretic analysis and compare the outcome with evolutionary results. Furthermore, the effects of search costs, uncertainty about future opportunities, and allowing multiple issues to be negotiated simultaneously are investigated.

Index Terms—Automated negotiation, evolutionary algorithms (EAs), incomplete information, multiple opportunities, ultimatum game.

I. 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 [1]–[17]. 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 paper¹ we focus on strategic aspects of bilateral bargaining within a market-like setting, where the agents learn effective bargaining strategies by a process of trial and error.

Bilateral bargaining has been extensively researched, for instance, in game theory [19]–[21]. Negotiations are often stylized 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 experimentally, using human subjects [21]–[23].

The ultimatum game models a negotiation between an isolated pair of players. In a market setting, however, an agent's behavior 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 modeled. 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 modeled 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 tradeoffs 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. We also consider the case where uncertainty exists about future opportunities and a new opponent cannot always be found.

An important aspect within this setting is the information available to the agents regarding their opponent's bargaining position. We distinguish between two information conditions in this paper: the complete information case, where an agent's current number of remaining 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 condition can be approached theoretically using game theoretic subgame perfect equilibrium (SPE) given reasonable assumptions. The incomplete information condition, on the other hand, seems much more difficult to analyze. We therefore apply computational simulations using evolutionary algorithms (EAs) 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 [24] and [25]. In contrast to for instance game theory, the agents are not assumed to be completely rational. Rather, the agents are naive optimizers acting on limited information. In these simulations, the agents are also myopic, i.e., they do not have any forward-looking ability or memory. Nonetheless, the emerging behavior of such adaptive, “low-rational” agents often corresponds surprisingly well with

Manuscript received February 1, 2004; revised July 16, 2004, October 24, 2004. This research was part of the project “Autonomous Systems of Trade Agents in E-Commerce” funded by the Telematics Institute, The Netherlands.

E. H. Gerding is with CWI, Center for Mathematics and Computer Science, 1090 GB Amsterdam, The Netherlands (e-mail: eg@ecs.soton.ac.uk).

H. La Poutré is with CWI, Center for Mathematics and Computer Science, 1090 GB Amsterdam, The Netherlands, and also with the School of Technology Management, Eindhoven University of Technology, 5600 MB Eindhoven, The Netherlands (e-mail: hlp@cw.nl).

Digital Object Identifier 10.1109/TSMCC.2006.860574

¹A shorter version appeared in the conference proceedings of the Workshop on Economics with Heterogeneous Interacting Agents (WEHIA) [18].

game-theoretic outcomes for completely informed and perfectly rational agents.

In this paper we present the evolutionary results for the settings described above and a game-theoretic analysis for the extended game with complete information. Subgame 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 if settings are equal for the proposer and the responder. If however the responder has many more bargaining opportunities than the proposer, she/he obtains a larger share than predicted game theoretically.

The evolutionary simulation shows that results change significantly if information about the opponent's bargaining position is not available. If the number of bargaining opportunities is sufficiently high, the responder now obtains the largest share. This result is achieved even when the number of initial bargaining opportunities is equal for both the players. 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 the players to explore additional opportunities. In the evolutionary simulation, the agreements are then similar to the one-shot ultimatum game. A similar effect is observed if bargaining is terminated with a small probability because no new opponent can be found.

This paper is organized as follows. In Section II the bargaining game with multiple bargaining opportunities is described. Section III provides a game-theoretic analysis of the game in case of complete information. Section IV outlines the evolutionary simulation and Section V discusses the obtained results from the simulation. Section VI reflects on the ultimatum game as a model for bargaining and also contains a short overview of related work. Lastly, a conclusion is drawn in Section VII.

II. DESCRIPTION OF THE BARGAINING GAME

The modeled market consists of buyers and sellers who exchange a single good through bilateral negotiations. At 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.² 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. More formally, the seller's utility u_s for an offer \vec{o} can be written as $\vec{w}_s \cdot \vec{o} = \sum_{i=1}^n w_s^i \cdot o^i$, where \vec{w}_s is a vector containing the seller's weights for each issue and n is the number of issues. Similarly, the buyer's utility function $u_b = \vec{w}_b \cdot [\vec{1} - \vec{o}]$, where \vec{w}_b represents the buyer's weights. The utilities of the agents are normalized between 0 and 1. The differences in weights of the two players determine the degree of *competitiveness* of the negotiations (i.e., to what extent tradeoffs can be beneficial). We

²Alternatively, a more complex bargaining game such as the alternating-offers game [21] involving multiple rounds can be used. Outcomes are equivalent to the ultimatum game, 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 [19] using an EA simulation.




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

formalize the notion of competitiveness and address this issue further in Section V-C.

Each seller agent initially has up to m_s bargaining opportunities to reach an agreement, whereas a buyer agent has m_b bargaining opportunities to start with. In case of a disagreement the agents are newly matched with randomly selected opponents, until no more bargaining opportunities remain. An agent's bargaining position is characterized by the number of remaining bargaining opportunities, which is known as an agent's *bargaining state*. A seller's and buyer's bargaining state is denoted by $\gamma_s \in \{0, 1, \dots, m_s\}$, and $\gamma_b \in \{0, 1, \dots, m_b\}$, respectively. If an agent's bargaining state reaches zero, the agent obtains a disagreement payoff that is set to zero.

An example for a two-issue negotiation 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 an offer $\vec{o} = (0.5, 0.5)$ and the buyer refuses this offer. Because the seller has no more bargaining opportunities, her/his bargaining game ends and she/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, even if the agents initially have equal bargaining opportunities ($m_s = m_b$), the matched agents can have different bargaining states. Having agents 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.³

III. GAME-THEORETICAL APPROACH

This section considers the game-theoretic SPE of the above game where the agents' bargaining states (i.e., bargaining positions) 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 IV). 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

³Agents are said to have no recall [26].

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 [21]. 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 game with multiple bargaining opportunities and complete information has an SPE with the same outcome as the ultimatum game: the seller obtains the entire share and the buyer receives the disagreement payoff, which is set to zero.⁴ 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, as in the ultimatum game, independent of γ_s . In SPE, the buyer's payoff for $\gamma_b = 1$ therefore equals zero. Note that this holds only if the seller is informed about the buyer's bargaining state.

If $\gamma_b = 2$, the buyer has two bargaining opportunities. Following 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 similar 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 $\gamma_b = m_b$.

We note that, because the agents are indifferent to the bargaining state in which the agreement is reached, actually several subgame perfect equilibria exist. In all cases, however, the divisions are the same. Note also that the above argument holds only if the seller is informed about the buyer's number of remaining bargaining opportunities. If this information is not available, a game-theoretic analysis seems much more difficult. An evolutionary simulation, however, is very apt to analyze the case of incomplete information. We analyze both the completely informed and the uninformed case in Section V. First, the evolutionary system is described in detail.

IV. EVOLUTIONARY APPROACH

Evolutionary algorithms (EAs) are powerful search algorithms inspired by Darwin's theory of natural selection [27]. In recent years, the evolutionary approach has been applied more and more 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 [19], [28]. An important advantage of EAs is that they do not assume complete rationality of the agents. Basically, the fitness (i.e., quality) of the individual agents is used to determine whether a strategy will be used in future situations.

A. Evolutionary Algorithm

The implementation used is based on a branch within EAs called *evolution strategies* (ES) [29], originally developed by Rechenberg [30] and Schwefel [31]. The ES were developed

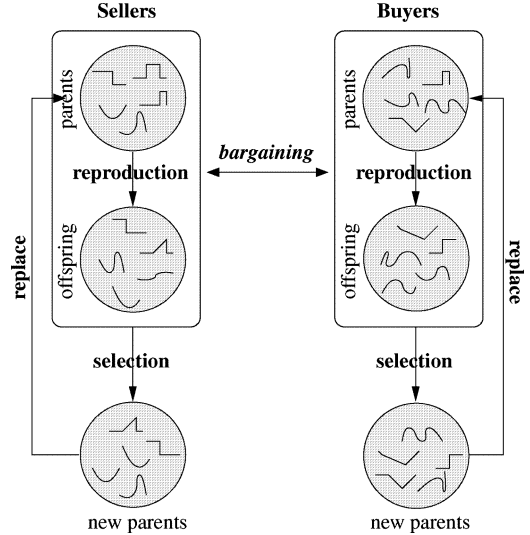


Fig. 2. Iteration loop of the evolutionary algorithm.

independently from the well-known genetic algorithms (GAs) [27], [32], introduced by Holland [33]. While GAs are more tailored toward binary-coded search spaces, ES are originally designed for real-encoded representations, which is a more natural encoding for the type of bargaining strategies we employ in the simulations.

The evolutionary simulation is depicted in Fig. 2 and works in the following way. Sellers and buyers are grouped into separate populations. This way the two types of agents coevolve. The system starts with randomly initialized “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 later.

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. For 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 measured only after the first payoff is

⁴This holds for continuous divisions of the surplus.

<i>Seller</i>	$\bar{o}(1 1)$	$\bar{o}(2 1)$...	$\bar{o}(m_s 1)$
<i>Strategy</i>	$\bar{o}(1 2)$	$\bar{o}(2 2)$...	$\bar{o}(m_s 2)$

	$\bar{o}(1 m_b)$	$\bar{o}(2 m_b)$...	$\bar{o}(m_s m_b)$

<i>Buyer</i>	$t(1 1)$	$t(2 1)$...	$t(m_b 1)$
<i>Strategy</i>	$t(1 2)$	$t(2 2)$...	$t(m_b 2)$

	$t(1 m_s)$	$t(2 m_s)$...	$t(m_b m_s)$

Fig. 3. Strategies of a seller and a buyer for the market game with complete information about the opponent's bargaining state. The offers $\bar{o}(\gamma_s|\gamma_b)$ and thresholds $t(\gamma_b|\gamma_s)$ are conditional on the bargaining state of the opponent, where $\gamma_s \in \{1, \dots, m_s\}$ and $\gamma_b \in \{1, \dots, m_b\}$.

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 [20]. 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) [29], where μ is the number of parents and λ is the number of generated offspring. In our simulation, we take $\mu = \lambda$.

B. Strategy Encoding

An agent's strategy is encoded on a so-called chromosome. The implementation of the EA is based on ES [29], using real-encoding of the chromosome.⁵ 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 representation was used in [19], [28] for an alternating-offers game.

We distinguish between the complete information setting and the incomplete information setting (see Section II). The strategy representation for each setting is schematically depicted in Figs. 3 and 4, respectively. In the incomplete information case (Fig. 4), an offer or threshold is specified for each bargaining states of the agent. In case of complete information (Fig. 3), an offer or threshold is also conditional on the opponent's bargaining state.

C. 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 [29]. Formally, $x'_i = x_i + \sigma N_i(0, 1)$. All resulting values larger than unity (or

⁵The widely-used GAs are more tailored toward binary-coded search spaces [27], [32], [33].

<i>Seller Strategy</i>	$\bar{o}(1)$	$\bar{o}(2)$...	$\bar{o}(m_s)$
<i>Buyer Strategy</i>	$t(1)$	$t(2)$...	$t(m_b)$

Fig. 4. Strategies of a seller and a buyer for the market game, where the players are uninformed about the opponent's bargaining state. An offer $\bar{o}(\gamma_s)$ or threshold $t(\gamma_b)$ is determined only by an agent's own bargaining state, since more information is not available.

TABLE I
DEFAULT SETTINGS OF THE EVOLUTIONARY SIMULATION

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

smaller than zero) are set to unity (respectively zero). The standard deviation is initially set to 0.1 and decays exponentially such that at 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 [19], [29] were tried as well, this model showed a closest match with game-theoretic results. In this paper we show only the results obtained using the exponential decay model.

V. EVOLUTIONARY SIMULATION RESULTS

The results are organized as follows. First, the game with complete information is studied in Section V-A and the results are compared to the game-theoretic SPE predictions. Section V-B studies the incomplete information case. Section V-C introduces a measure of competitiveness for multi-issue negotiations and compares results for different levels of integrative negotiations. Finally, Section V-D considers the effects of fixed search costs in the market game and uncertainty about future opportunities.

A. Game-Theoretic Validation: The Complete Information Condition

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

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 finite and the bargaining state of the opponent is common knowledge (see also Section III). Note that this result holds for both the symmetric setup and if agents have an unequal initial number of bargaining opportunities (i.e., for $m_b \neq m_s$).

Fig. 5 shows the EA outcomes for a symmetric setup, where the initial number of bargaining opportunities is the same for both sellers and buyers. In the following, we use m to denote the initial bargaining opportunities in the symmetric case. The results indicate an almost perfect match between evolutionary

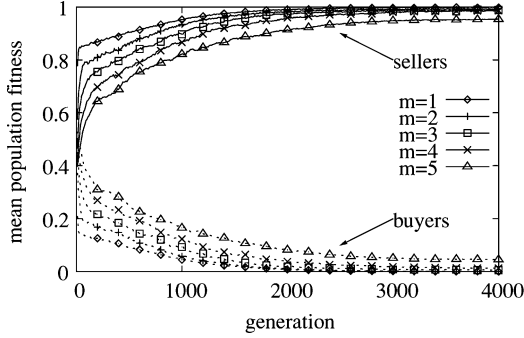


Fig. 5. Development of the mean fitness (averaged over 30 runs) for complete information setting with varying initial number of bargaining opportunities ($m = m_s = m_b$).

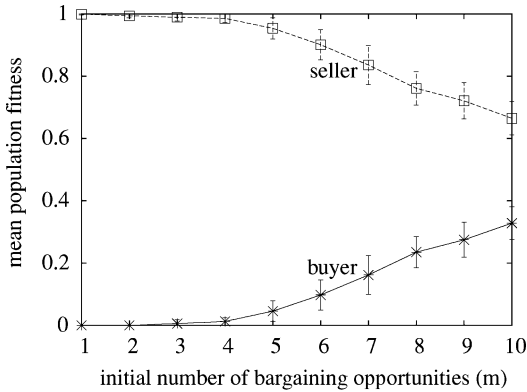


Fig. 6. Results after 4000 generations (averaged over 30 runs) in case of complete information and symmetric settings ($m = m_s = m_b$).

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. Fig. 6 shows the long-term EA outcomes (after 4000 generations) for m up to 10. 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. Details on tuning the EA are beyond the scope of this paper. Instead, we refer the interested reader to previous research [19], [28], in which different EA settings are systematically studied for an alternating-offers bargaining game. Henceforth, we present only experiments using uniform EA settings in this paper.

For the asymmetric case, SPE also predicts an extreme division, where the seller obtains the entire share (see Section III). This holds even for instance when $m_s = 1$ and $m_b = 5$.⁶

Results using the evolutionary simulation are shown in Fig. 7. Interestingly, the simulation results deviate from game-theoretic

⁶To see this, recall that (using backward induction) the disagreement payoff for the buyer is always zero, independent of her/his bargaining state; the buyer is therefore indifferent between accepting and refusing. To ensure an agreement immediately, a seller could bid an infinitely small amount to the buyer, which the buyer will then accept. In the limit, this amount converges to zero.




Fig. 7. Results after 4000 generations (averaged over 30 runs) in case of complete information and varying the seller's initial number of bargaining opportunities m_s . The buyer's initial number of bargaining opportunities $m_b = 5$.

outcomes if the buyer has many more opportunities to bargain than the seller. In the extreme case, where the seller has only a single bargaining opportunity, the buyer obtains a larger share on average (note that the standard deviation is very high, around 0.20). The outcomes are robust to other EA settings. To understand why this occurs, recall that the game-theoretic derivation is based on backward induction, starting at the final state where $\gamma_b = 1$. When $m_s = 1$, however, the bargaining state $\gamma_b = 1$ is never reached in practice, since the seller always prefers an agreement immediately. This is confirmed by the experiments that show that in the long run no agreements are reached in $\gamma_b = 1$ and $\gamma_b = 2$.

The adaptive agents do not reason forward and never reach the final state using trial and error. The asymmetric outcomes are therefore less predictable, which also explains the high standard deviations. The results seem to be more realistic, however. 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 [34]. The evolutionary results could be considered closer to real-world outcomes, where for instance an extreme split of the surplus is usually not found [23], [34].

B. Incomplete Information Condition

We now examine the results when the agents do not know their opponent's bargaining states; the agents know only 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.

Results produced after 4000 generations of the EA for the incomplete information case are shown in Fig. 8, 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. In the complete information case,




Fig. 8. Results after 4000 generations (averaged over 30 runs) for incomplete information settings with various $m = m_b = m_s$. The error bars indicate the standard deviation of the averaged results.

the seller obtains almost the entire surplus, whereas in the incomplete information case, the responder (i.e., buyer) obtains a larger share (see Fig. 8). 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/his final state, she/he 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 her/his last bargaining state. Because the seller does not know the buyer's bargaining state, she/he can no longer anticipate the buyer's behavior. In order to prevent a disagreement, the seller will then concede in the last bargaining state. The expected payoff in case of a disagreement and the offers in earlier bargaining states will also then decrease. After many generations, the simulation converges to an outcome where the seller concedes almost her/his entire surplus in each bargaining state. We also observe that the seller concedes slightly less if she/he has more bargaining opportunities remaining, resulting in less extreme deals if m becomes large, as shown in Fig. 8.

In the incomplete information condition, the first-mover (here the seller) has no information about her/his opponent. The responder, on the other hand, can make a relatively more informative decision based on the seller's offer. In the ultimatum game the proposer seems to dominate the outcome, and 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 and equal for both players. Furthermore, the players incur no costs for refusing a deal. As we will show in Section V-D, even slight costs completely change these results.

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. 8. 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 behavior less predictable for the seller. The value for which the transition occurs depends on game parameters such

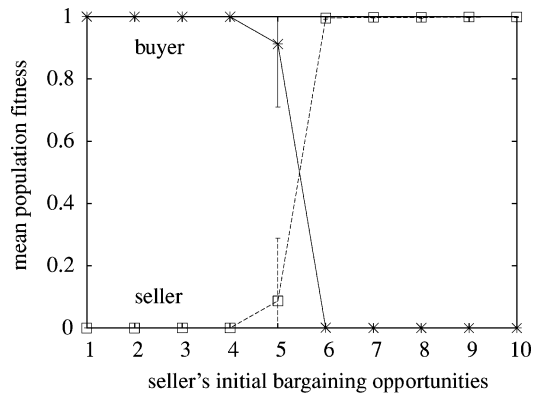


Fig. 9. Results after 4000 generations (averaged over 30 runs) for incomplete information settings with various m_s and $m_b = 5$. The error bars indicate the standard deviation of the averaged results.

as r and the competitiveness of the negotiation. The latter will be discussed further in the next section.

We also consider the asymmetric case, see Fig. 9. The evolutionary outcomes appear to be in favor of the buyer when $m_b > m_s$, and of the seller when $m_b < m_s$. Note that, the outcomes are "all or nothing" for the buyer or the seller. An interesting case is when $m_b = m_s$, where a sudden transition seems to occur. Although the asymmetric case is very stable, results for $m_b = m_s$ are more sensitive to EA parameters such as the evaluation frequency. The EA also takes longer to converge.

C. 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* [35]. 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. Because of the increased complexity, however, the evolutionary results are less extreme when the number of bargaining opportunities is large.

The utility of an offer is an additive, weighted function of the share obtained for each issue (see also Section II). The weights for sellers and buyers for the two issues are $\vec{w}_s = (0.5 - \alpha, 0.5 + \alpha)$ and $\vec{w}_b = (0.5 + \alpha, 0.5 - \alpha)$, 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 $2(0.5 + \alpha)$, where each agent obtains $(0.5 + \alpha)$.

Results for $\alpha = 0.2$ are shown in Fig. 10. 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. 10). Only two bargaining opportunities are needed to obtain an advantage for the responder, as opposed to four in the single-issue game (Fig. 8).

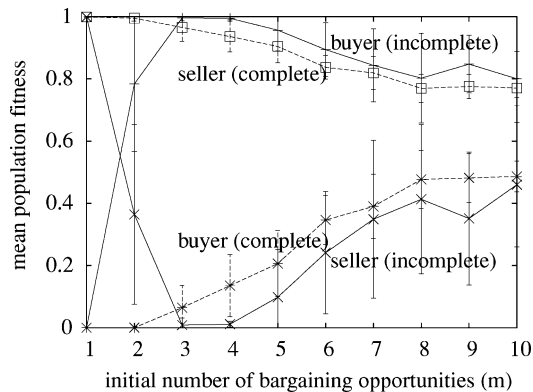


Fig. 10. Mean long-term outcomes for two-issue negotiations, $\alpha = 0.2$, and $m = m_s = m_b$.

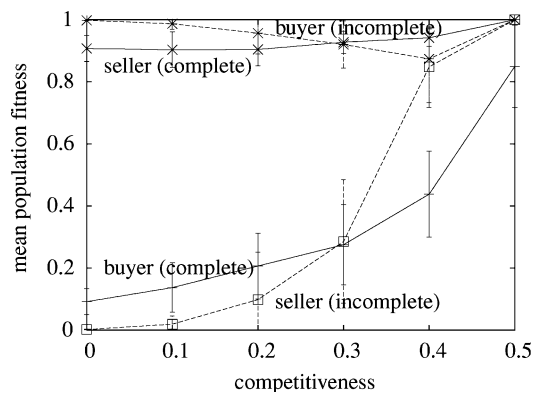


Fig. 11. Mean long-term outcomes for $m = m_s = m_b = 5$ and different values for the competitiveness (α).

Fig. 10 also shows a less extreme split compared to competitive negotiations, particularly for large m . This occurs 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. As shown in Fig. 11, 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 V-A 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 that are closer to game-theoretic predictions.

D. Search Costs and Premature Termination

We further extend the bargaining game in two ways. First, we introduce search or negotiation costs each time an offer is refused and agents engage in a new negotiation. Subsequently, we consider the case where there exists uncertainty about whether a new bargaining opponent can be found. While we have assumed




Fig. 12. Mean long-term results as a function of the search costs (β) for $m = m_s = m_b = 5$.

until now that the number of bargaining opportunities remains fixed, there can be external factors that influence the number of opportunities (e.g., if a seller has in the meanwhile sold the good to another buyer). This is modeled as a probability that negotiations terminate prematurely, i.e., before the final number of bargaining opportunities is completely exhausted.

Search 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 [36, 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.

Evolutionary outcomes for the complete and incomplete information settings with different search costs are depicted in Fig. 12. Negotiations are competitive and buyers and sellers each have five 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 behavior 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. 12). 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.




Fig. 13. Long-term fitness values for $m = m_b = m_s = 5$ and incomplete information, when negotiations are discontinued with a certain probability after each disagreement.

Similar outcomes are observed when bargaining for a buyer and/or a seller is discontinued with a certain probability after each disagreement.⁷ Fig. 13 shows the long-term outcomes for different probabilities of premature termination after each bargaining opportunity. The probability is set equal for buyers and sellers, and for each bargaining opportunity, but drawn independently. As with search costs, the seller obtains the largest share if the probability is sufficiently high.

Note that, the effect of premature termination is less extreme, however, this is because search costs also affect the utility if an agreement is *not* reached, providing an additional incentive to reach agreements (otherwise, a negative utility is obtained). In case of premature termination, on the other hand, an agent is indifferent between termination after the first bargaining opportunity and a disagreement in the last bargaining opportunity.

VI. DISCUSSION AND RELATED RESEARCH

The ultimatum game is frequently applied in economic and social literature and is an elegant model of negotiation. Despite its apparent simplicity, it captures the essence of negotiation—the division of a surplus between two players that can be created through cooperation. A more extensive model and well-known model is the Rubinstein–Ståhl alternating-offers game consisting of several rounds where offers and counter offers are exchanged [37]. In absence of any time pressure, however, the division of surplus in subgame-perfect equilibrium is equivalent to the ultimatum game—the player that has the opportunity to make the offer in the final round obtains the entire surplus. We therefore opted for the more simple ultimatum game as the basic model for negotiation, which allows us to focus on the aspect of multiple opportunities and the information condition of the participants.

In this paper, we have implicitly assumed that the size of the surplus is known to both the players and is equal for each negotiation game. These are typical assumptions made both in the ultimatum and alternating-offer game (e.g., in [20] this as-

sumption is also made within a market-like setting, see also below). In case of price negotiations, e.g., this implies that the reservation prices of both players are common knowledge (the surplus is then the difference between the reservation price of the buyer and the seller). Although an estimate of the surplus can usually be made, in practice there is often uncertainty about the precise value. In recent years, an increasing amount of research is performed on bargaining with incomplete information about the size of the surplus, see e.g. [38]. Incomplete information provides a likely explanation for the occurrence of inefficiencies such as costly delays and disagreements. While in this paper the focus is on multiple opportunities and uncertainty regarding the opponent's bargaining position, an interesting direction for future research is to include also uncertainty about the size of the surplus and to use evolutionary simulations to investigate situations where the surplus differs at various bargaining opportunities.

Closely related to our research is the work in [20, Part 2], and the references therein, where various market models are presented and a market equilibrium is calculated using the so-called sequential equilibrium solution concept; especially the model described in [20, Sec. 8.2] corresponds to our model in a number of ways. More specifically, as in our case, buyers and sellers are randomly matched, after which the ultimatum game is played to divide a known surplus. In their model, however, the proposer is selected randomly, whereas in our model the seller is always the proposer. Furthermore, an important aspect of their model is that buyers and sellers are matched simultaneously and that there are fewer sellers than buyers (thus not all buyers are matched). The matching period is equal for all buyers and sellers. As a result, in equilibrium all agreements are reached immediately and the seller obtains the entire surplus.⁸ By contrast, in our model, buyers and sellers are matched sequentially and each agent can be in a different period or *bargaining state*. This makes an analysis much more difficult. Moreover, we also consider the symmetric case where the *initial* number of bargaining opportunities is equal for all agents, and we investigate the effects of search costs and uncertainty about future opportunities. Furthermore, we find that an important factor determining the outcome is whether a player is informed about the opponent's bargaining state (i.e., the number of remaining bargaining opportunities).

In our earlier work [19], [28], we investigated evolutionary simulations of the alternating-offers game, using a related approach. Although the focus in these papers is on isolated games between two agents, in [19] we also present initial results of the market-like setting with multiple opportunities as an example of the applicability of evolutionary simulations. Only results where players are not informed about their opponent's bargaining state are discussed, however. In this paper, we delve deep into the aspect of multiple opportunities itself and present a much more extensive analysis, such as a game-theoretic solution for the complete information case and a comparison between two

⁷This is analogous to discount factors or a probability of break down in case of multiround bargaining, as used in [19].

⁸There are actually many equilibria, since there is no time pressure and a seller always has the opportunity to bargain in the next matching period in case of a disagreement.

information conditions. Furthermore, in this paper we extend the results to the asymmetric case (i.e., where the initial number of bargaining opportunities varies for a seller and a buyer) and we investigate a model where uncertainty exists about future opportunities.

A number of related papers investigate situations of bargaining using an evolutionary computation approach. The ultimatum game is investigated in [34] using an EA. Oliver [39] was the first to demonstrate that a system of adaptive agents can learn effective negotiation strategies using EAs in the alternating-offers game. More elaborate strategy representations are proposed in [10]. Offers and counter offers are generated in this model by a linear combination of simple bargaining tactics (time-dependent, resource-dependent, or behavior-dependent tactics). The parameters of these different negotiation tactics and their relative importance weightings are calculated using a GA. The time-dependent tactics are further investigated in [40] using GAs, for the case that negotiating agents have different time preferences. GAs are also used in [8] to search for appropriate (counter) offers during negotiation. Whether an offer is appropriate depends on various social, environmental, and cognitive traits of the agents. The traits determine, e.g., the type of concession tactics, knowledge of opponent preferences, cost of delaying settlement, and the magnitude of initial offers. These papers study the basic negotiation game between two players, but do not consider market-like settings where buyers and sellers may have several opportunities to reach an agreement.

Bilateral bargaining within a market-like setting is studied in [13]–[15] by introducing so-called *market-driven agents*. Rather than using fixed strategies, these agents are able to flexibly adjust their strategies according to changing market conditions, such as the number of trading opportunities and the amount of competition. In a buyer's market, supply is greater than demand and a market-driven buyer will make relatively small concessions, whereas in a seller's market the opposite holds. Concessions also depend on the negotiators' desire to trade or *eagerness*: if an agent is eager to reach an agreement, he/she makes relatively large concessions to increase the probability of an agreement. An indifferent agent, on the other hand, is more willing to risk a conflict and tries to get a better deal instead. Other factors that impact the concessions made are time pressure and the gap or *spread* between the utility of the offers made by the negotiating agents. In [13], a theoretical analysis is provided of how these market-driven agents respond to different market conditions, and in [14] the market-driven strategy is also assessed using computational simulations. The market-driven strategy is augmented in [15] by using fuzzy rules to guide agents in relaxing their trade conditions (e.g., aspiration or reservation price) in the face of intense negotiation pressure (e.g., due to very stiff competition). The strategies described in [13]–[15] are based on heuristics for determining the concession and, although flexible, the agents do not learn from past experience. This differs considerably from our work, where the strategies are adapted using an EA. Also, in this paper we apply game-theoretic analysis to validate our simulation approach. Furthermore, we explicitly study the impact of future bargaining

opportunities, whereas in [13]–[15] the focus is on concurrent opportunities (and competition).

VII. CONCLUSION

We studied 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 modeled. We also investigated multi-issue integrative negotiations and the effects of search costs and premature termination if a disagreement occurs.

The game-theoretic outcome using SPE for the one-shot ultimatum game predicts an extreme split of surplus: The seller obtains the whole surplus, whereas the buyer obtains her/his disagreement payoff. We extended the analysis for multiple bargaining opportunities with complete information of the opponent's remaining number of bargaining opportunities or *bargaining state* and find an equivalent outcome. However, 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 compared the evolutionary results with the game-theoretical outcomes for the game with complete information to validate the evolutionary approach. If the initial number of bargaining opportunities is small and equal for both players, 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. This can be improved, however, by adjusting EA settings. An interesting discrepancy from SPE also occurs when buyers have many more bargaining opportunities than do sellers. Although SPE outcomes are similar to the symmetric case, the EA results are more in favor of the buyer. These more realistic outcomes occur because the final state, where the buyer has exhausted all bargaining options, is never reached in the experiments. The theoretical outcomes, however, are based on backward induction from this final state.

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. In the complete information game, the seller dominates the market, whereas the buyer is better off in the incomplete information setting, provided that the buyer has *at least as many* opportunities as the seller and 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. In case the seller has more bargaining opportunities, however, the seller agent obtains the entire surplus.

Similar outcomes were 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 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 studied 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 the economic theory, which states that even infinitely small search costs produce monopolistic prices. The outcomes are similar, but less extreme if search costs are replaced by a probability that bargaining is discontinued after a disagreement. This models the situation where uncertainty exists about future opportunities.

In this paper, we have shown that evolutionary simulations are extremely useful to investigate negotiations with incomplete information, which are unwieldy to analyze theoretically. Using EAs, we can simulate complex interactions involving a large number of agents, as is the case in bargaining with multiple 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 or about the size of the negotiated surplus. Another interesting extension is allowing agents to return to previously encountered opponents.

ACKNOWLEDGMENT

The authors would like to thank the anonymous referees and the guest editor of this special issue for their valuable comments and suggestions.

REFERENCES

- [1] K. Binmore and N. Vulkan, "Applying game theory to automated negotiation," *Netnomics*, vol. 1, no. 1, pp. 1–9, 1999.
- [2] D. D. B. v. Bragt and J. A. La Poutré, "Generating efficient automata for negotiations—An exploration with evolutionary algorithms," in *Proc. Genetic and Evolutionary Computation Conf., GECCO-2001*, L. Spector, E. Goodman, A. Wu, W. Langdon, H.-M. Voigt, M. Gen, S. Sen, M. Dorigo, S. Pezeshk, M. Garzon, and E. Burke, Eds., San Francisco, CA: Morgan Kaufmann, 2001, p. 1093.
- [3] F. Dignum and U. Cortés, *Agent-Mediated Electronic Commerce III, Current Issues in Agent-Based Electronic Commerce Systems*. New York: Springer-Verlag, 2001.
- [4] P. Faratin, C. Sierra, and N. R. Jennings, "Using similarity criteria to make issue trade-offs," *Artif. Intell.*, vol. 142, no. 2, pp. 205–237, 2002.
- [5] R. H. Guttman, A. Moukas, and P. Maes, "Agent-mediated electronic commerce: A survey," *Knowl. Eng. Rev.*, vol. 13, no. 2, pp. 147–159, Jul. 1998.
- [6] N. R. Jennings, P. Faratin, A. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge, "Automated negotiation: Prospects, methods and challenges," *Group Decision Negotiation*, vol. 10, pp. 199–215, 2001.
- [7] G. E. Kersten and S. Noronha, "Rational agents, contract curves, and inefficient compromises," *IEEE Trans. Syst., Man, Cybern. A, Syst. Humans*, vol. 28, no. 3, pp. 326–338, May 1998.
- [8] R. Krovi, A. C. Graesser, and W. E. Pracht, "Agent behaviors in virtual negotiation environments," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 29, pp. 15–25, Feb. 1999.
- [9] P. Maes, R. H. Guttman, and A. G. Moukas, "Agents that buy and sell," *Commun. ACM*, vol. 42, no. 3, pp. 81–91, 1999.
- [10] N. Matos, C. Sierra, and N. R. Jennings, "Determining successful negotiation strategies: An evolutionary approach," in *Proc. 3rd Int. Conf. Multi-Agent Systems (ICMAS-98)*, Paris, France, 1998, pp. 182–189.
- [11] T. D. Nguyen and N. R. Jennings, "Coordinating multiple concurrent negotiations," in *Proc. 3rd Int. Conf. Autonomous Agents and Multi-Agent Systems (AAMAS2004)*, New York: ACM Press, 2004.
- [12] S. Paurobally, P. J. Turner, and N. R. Jennings, "Automated negotiation for m-services," *IEEE Trans. Syst., Man, Cybern. A, Syst. Humans*, vol. 33, pp. 709–724, Nov. 2003.
- [13] K. M. Sim, "A market-driven model for designing negotiation agents," *Computational Intelligence*, vol. 18, no. 4, pp. 618–637, Nov. 2002.
- [14] K. M. Sim and C. Y. Choi, "Agents that react to changing market situations," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 33, pp. 188–201, Apr. 2003.
- [15] K. M. Sim and S. Y. Wang, "Flexible negotiation agents with relaxed decision rules," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 34, pp. 1602–1608, Jun. 2004.
- [16] K. Somefun, E. H. Gerding, S. Bohte, and J. A. La Poutré, "Automated negotiation and bundling of information goods," *Agent-Mediated Electronic Commerce V: Designing Mechanisms and Systems*, P. Faratin, D. C. Parkes, J. A. Rodríguez-Aguilar, and W. E. Walsh, Eds., Berlin, Germany: Springer-Verlag, 2004.
- [17] D. Zeng and K. Sycara, "Bayesian learning in negotiation," *Int. J. Human Comput. Syst.*, vol. 48, pp. 125–141, 1998.
- [18] E. H. Gerding and J. A. La Poutré, "Bargaining with posterior opportunities: An evolutionary social simulation," in *The Complex Dynamics of Economic Interaction*, M. Gallegati, A. Kirman, and M. Marsili, Eds., Berlin, Germany: Springer-Verlag, 2004, vol. 531, pp. 241–256.
- [19] E. H. Gerding, D. D. B. v. Bragt, and J. A. La Poutré, "Multi-issue negotiation processes by evolutionary simulation: Validation and social extensions," *Comput. Econom.*, vol. 22, pp. 39–63, 2003.
- [20] M. J. Osborne and A. Rubinstein, *Bargaining and Markets*. San Diego, CA: Academic, 1990.
- [21] —, *A Course in Game Theory*. Cambridge, MA: MIT Press, 1994.
- [22] A. E. Roth, "Bargaining experiments," in *Handbook of Experimental Economics*, J. Kagel and A. E. Roth, Eds., Princeton, NJ: Princeton University Press, 1995, pp. 253–348.
- [23] H. Lin and S. Sunder, "Using experimental data to model bargaining behavior in ultimatum games," in *Experimental Business Research*, R. Zwick and A. Rapoport, Eds., Norwell, MA: Kluwer, 2002, pp. 373–397.
- [24] J. M. Epstein and R. Axtell, *Growing Artificial Societies: Social Science from the Bottom Up*. Cambridge, MA: MIT Press, 1996.
- [25] L. Tesfatsion, "Agent-based computational economics: Modelling economies as complex adaptive systems," in *Proc. 6th Joint Conf. Information Sciences*, North Carolina, Mar. 2002, pp. 40–43.
- [26] R. Zwick and C. C. Lee, "Bargaining and search: An experimental study," *Group Decision Negotiation*, vol. 8, no. 6, pp. 463–487, 1999.
- [27] M. Mitchell, *An Introduction to Genetic Algorithms*. Cambridge, MA: MIT Press, 1996.
- [28] D. D. B. v. Bragt, E. H. Gerding, and J. A. La Poutré, "Equilibrium selection in alternating-offers bargaining models: The evolutionary computing approach," *The Electronic Journal of Evolutionary Modeling and Economic Dynamics (e-JEMED)*, vol. 1, 2002.
- [29] T. Bäck, *Evolutionary Algorithms in Theory and Practice*. New York: Oxford Univ. Press, 1996.
- [30] I. Rechenberg, *Evolutionstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Stuttgart: Frommann-Holzboog, 1973.
- [31] H.-P. Schwefel, *Evolution and Optimum Seeking*. New York: Wiley, 1995. Sixth-Generation Computer Technology Series.

- [32] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley, 1989.
- [33] J. H. Holland, *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. Ann Arbor, MI: The University of Michigan Press/Longman Canada, 1975.
- [34] T. Riechmann, *Learning in Economics: Analysis and Application of Genetic Algorithms*. Heidelberg, Germany: Physica-Verlag, 2001.
- [35] R. H. Guttman and P. Maes, "Agent-mediated integrative negotiation for retail electronic commerce," in *Proc. 1st Int. Workshop Agent Mediated Electronic Trading (AMET)*, Minneapolis, May 1998.
- [36] S. Y. Choi, D. O. Stahl, and A. B. Whinston *The Economics of Electronic Commerce*. Macmillan Technical Publishing, 1997.
- [37] A. Rubinstein, "Perfect equilibrium in a bargaining model," *Econometrica*, vol. 50, no. 1, pp. 155–162, 1982.
- [38] L. M. Ausubel, P. Cramton, and R. J. Deneckere, "Bargaining with incomplete information" in *Handbook of Game Theory, with Economic Applications*, R. J. Aumann and S. Hart, Eds., North-Holland: Elsevier Science Publisher, vol. 3, 2002.
- [39] J. R. Oliver, "A machine learning approach to automated negotiation and prospects for electronic commerce," *J. Manag. Inf. Syst.*, vol. 13, no. 3, pp. 83–112, 1996.
- [40] S. Fatima, M. Wooldridge, and N. R. Jennings, "Comparing equilibria for game theoretic and evolutionary bargaining models," in *Proc. 5th Int. Workshop Agent-Mediated E-Commerce (AMEC V)*, Melbourne, Australia, 2003, pp. 70–77.



Enrico H. Gerding received the M.Sc. degree in artificial intelligence from the Free University, Amsterdam, The Netherlands, in 1999, and the Ph.D. degree from the Eindhoven University of Technology, Eindhoven, The Netherlands, in 2004. The research that resulted in his Ph.D. thesis entitled "Autonomous Agents in Bargaining Games: An Evolutionary Investigation of Fundamentals, Strategies, and Business Applications" was carried out at the Centre for Mathematics and Computer Science (CWI), Amsterdam, The Netherlands.

Presently, he is a Research Fellow in the School of Electronics and Computer Science, University of Southampton, Southampton, U.K. He is currently working on a project to develop novel market-based solutions for the control of complex computational systems. His research interests include mechanism design, automated negotiation, and evolutionary algorithms.



Han La Poutré received the M.Sc. degree in mathematics (cum laude) from the Eindhoven University of Technology, the Netherlands, and the Ph.D. degree in computer science from the Utrecht University, the Netherlands.

He is currently a research group leader at the Center for Mathematics and Computer Science (CWI), Amsterdam, the Netherlands. At CWI, he is heading the research group "Computational Intelligence and Multiagent Games." He is also a Full Professor at the Eindhoven University of Technology, the Netherlands. He has also been a Research at the Princeton University. His research interests include multiagent games (negotiations, auctions, market games), computational intelligence, adaptive agents, and agent-based computational economics.

Dr. Poutré is a member of the editorial boards of *Netnomics*, *Computational Management Science*, and *e-Jemed* (Electronic Journal of Evolutionary Modeling and Economic Dynamics). He received the Best Ph.D. Paper Award at the 1990 ACM Symposium on Theory of Computing (STOC), a NATO Science Fellowship, and a Fellowship of the KNAW (the Royal Netherlands Academy of Sciences and Arts). Both in 1999 and 2005, his research group at CWI was rated excellent in six-yearly evaluation by NWO (the Netherlands Organisation for Scientific Research).