

Multi-Attribute Bilateral Bargaining in a One-to-Many Setting

E.H. Gerding¹, D.J.A. Somefun¹, and J.A. Han La Poutré^{1,2}

¹ Center for Mathematics and Computer Science (CWI)
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

² Eindhoven University of Technology, School of Technology Management,
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
{egerding, koye, hlp}@cwi.nl

Abstract. Negotiations are an important way of reaching agreements between selfish autonomous agents. In this paper we focus on one-to-many bargaining within the context of agent-mediated electronic commerce. We consider an approach where a seller negotiates over multiple interdependent attributes with many buyers individually. Bargaining is conducted in a bilateral fashion, using an alternating-offers protocol. In such a one-to-many setting, “fairness,” which corresponds to the notion of envy-freeness in auctions, may be an important business constraint. For the case of virtually unlimited supply (such as information goods), we present a number of one-to-many bargaining strategies for the seller, which take into account the fairness constraint, and consider multiple attributes simultaneously. We compare the performance of the bargaining strategies using an evolutionary simulation, especially for the case of impatient buyers and small premature bargaining break off probability. Several of the developed strategies are able to extract almost all the surplus; they utilize the fact that the setting is one-to-many, even though bargaining occurs in a bilateral fashion.

1 Introduction

It is common to characterize negotiations by their setting: bilateral, one-to-many, or many-to-many. In this paper we focus on the one-to-many bargaining setting, where a seller agent negotiates, on behalf of a seller, with many buyers individually in a *bilateral* fashion. We develop various strategies for this setting which enable a seller agent to bargain over *multiple interdependent* issues simultaneously and effectively.

In many cases, auctions can be used to effectively organize one-to-many bargaining. Depending on the setting, auctions can provide buyers with the incentive to reveal their preferences truthfully, and to allocate the goods efficiently. For various situations, however, auctions may not be the preferred protocol to bargainers. In situations of, for example, virtually unlimited supply, multiple issues, continuous sale, and/or repeated sales to the same buyers the appropriate auction protocol becomes, at best, much more complex. Consequently, businesses may opt for the intuitive and flexible bilateral bargaining protocol, where the seller agent negotiates bilaterally with one or more buyers simultaneously by exchanging offers and counter offers. In many electronic commerce domains supply is flexible and new goods can be reproduced quickly, at relatively low

costs; especially in these domains businesses may opt for a bilateral bargaining protocol. The sales of information goods, with its virtually unlimited supply, provide a particularly good example of an application domain where businesses may opt for bilateral bargaining.

Potentially, bargaining can lead to unsatisfied customers if buyers perceive the outcomes of the negotiations as unfair. This can occur when, for instance, two customers obtain similar goods at the same time but end up paying very different amounts. Fairness of negotiation outcomes is important for customer satisfaction, which in turn may be important for a business' long term profitability. The seller agent can prevent unfair outcomes by incorporating a fairness norm, comparable to the notion of envy-freeness in auctions [1], whereby customers are treated in a similar fashion. This fairness aspect emphasizes that bargaining is really one-to-many.

The challenge is to develop bargaining strategies for the seller agent that maximize overall revenue by utilizing differences in buyers' willingness to pay without violating the fairness constraint. In this paper we focus on strategies that can utilize differences in customers' time pressure. For the case of virtually unlimited supply, as for information goods, we present a number of one-to-many bargaining strategies for the seller, which take into account the fairness constraint and bargain over multiple attributes. In order to benefit from time pressure, these strategies specify—in addition to an actual (counter) proposal or an acceptance proposal—when to respond to an opponent's proposal. We compare the performance of the bargaining strategies using an evolutionary simulation, especially for the case of impatient buyers and small (exogenous) probability (per negotiation round) of a customer breaking off the negotiations. One set of strategies, the so-called “responsive threshold” strategies, are able to extract almost all the surplus, given sufficient time pressure. These strategies benefit from the fact that the setting is one-to-many, even though bargaining occurs in a bilateral fashion. In addition, the strategies are able to find win-win agreements (i.e., very little Pareto improvement is possible).

A number of related papers study bargaining using an evolutionary approach, e.g. [2–5]. Our approach extends previous research to multiple (types of) buyers and bilateral negotiation strategies for one-to-many multi-issue bargaining which can benefit from time pressure. In addition a growing body of literature exists on multi-issue negotiation, which focus on developing techniques a seller and/or a buyer can use to determine the relative magnitude of the various issues and consequently search for approximately win-win (or Pareto-efficient) deals [6–9]. Although we also consider the problem of how to determine the values for the various attributes of an offer, the focus of the paper lies on the development of “threshold” strategies for one-to-many negotiation. These strategies determine the desired utility level of a deal and can be used in conjunction with the techniques already developed in the literature.

The remainder of the paper is organized as follows. In Section 2 we describe the bargaining setting and introduce strategies for one-to-many bargaining. In Section 3 we discuss the simulation environment used for testing the performance of the strategies. We present the simulation results of the conducted computer experiments in Section 4. Conclusions follow in Section 5.

2 One-to-Many Bargaining

2.1 Fairness

An agent representing a business can be endowed with various bargaining strategies. The bargaining outcome should, however, be fair. Because fairness must be ensured by the seller and because buyer preferences are private, we define fairness as follows. Suppose at time t_d a buyer reaches a deal. We say that this deal is fair, relative to a fixed interval $\Delta > 0$, whenever there exist a start time t_s , with $t_d \in [t_s, t_s + \Delta]$, such that the seller is indifferent between any other deal reached within the interval $[t_s, t_s + \Delta]$. Whenever price is the only issue, a buyer does not strictly prefer any deal for which the seller is indifferent. In this case, we can give the following equivalent definition: a deal is fair, relative to a fixed interval $\Delta > 0$, whenever there exist a start time t_s , with $t_d \in [t_s, t_s + \Delta]$, such that the buyer does not strictly prefer any other deal which is reached within the interval $[t_s, t_s + \Delta]$. Note, that this definition of fairness is closely related to the notion of envy-free auctions in [1]; it adapts the notion of envy-freeness to the more continuous setting of bilateral bargaining.

2.2 Bargaining Protocol

The seller agent negotiates with many buyer agents simultaneously in a bilateral fashion by alternating offers and counter offers. An offer specifies a value for each attribute of the negotiation, such as the price, quality, quantity, and other relevant aspects. The protocol allows for multiple offers to be submitted simultaneously. Exchanging multiple offers can improve the (Pareto) efficiency of agreements made when several attributes are concerned. An offer constitutes a Pareto improvement over another offer whenever it makes one bargainer better off without making the other worse off. A bargainer proposing multiple offers can be indifferent between those offers whereas his opponent may prefer a particular offer and can improve efficiency by selecting this offer.

We call the set of offers combined with the preconditions a *proposal*. A bargainer can accept one of the submitted offers or reject all offers and place a counter proposal. Negotiations between a buyer and seller agent proceed to the next round whenever a proposal is submitted and terminates when one of the submitted offers is accepted or after a predefined period of time has elapsed. Note that a bargainer can introduce a delay before submitting a counter proposal. The duration of a round varies depending on the delay. Figure 1 depicts the alternating offer bargaining protocol.

2.3 Time Pressure

An important assumption is that buyers are impatient and prefer an early agreement. Time pressure or time impatience is a common assumption in bargaining, e.g. [10]. The seller agent is simultaneously and continuously negotiating with many buyers and is therefore less concerned with immediately reaching an agreement for a particular bargaining outcome, i.e., he is relatively patient. We model this relative time patience by assuming that the seller, unlike the buyers, has no *direct* time pressure: i.e., the seller is indifferent between selling now and later. In the experiments we do, however,




Fig. 1. The agents' bargaining protocol

consider a small exogenous probability (per negotiation round) of a customer breaking off the negotiation process; therefore the seller has an indirect incentive to speed up the negotiations (all else equal).

At least in theory, the seller can benefit from buyers' time-pressure by introducing a delay before submitting a counter proposal. An important question is then which bargaining strategies can most effectively utilize these potential benefits. Experimental results discussed in Section 4 show that responsive threshold strategies, which we will discuss in the next Section, are very effective: depending on the time pressure, they are capable of extracting very large shares of the seller surplus. (Reasoned from the seller's perspective the surplus is just the maximum utility he can realize by selling the goods or services.)

2.4 One-to-Many Bargaining Strategies

The challenge is to develop bargaining strategies for the seller agent that maximize overall revenue by utilizing differences in buyers' willingness to pay without violating the fairness constraint. Instead, these strategies utilize differences indirectly through buyers' time pressure. In order to benefit from time pressure, a strategy specifies, in addition to an actual (counter) proposal or an acceptance proposal, when to respond to an opponent's proposal.

The seller strategies as developed determine the offers of a proposal in two steps. First, they specify a threshold level which sets the utility level of the offers. Second, they generate the values for the individual attributes, given the threshold. Advanced techniques for multi-issue negotiation, such as discussed in [6–9], can be applied to the latter. The focus of this paper is on effective strategies for determining the threshold; we therefore only consider the relatively simple technique of randomly determining the attribute-values given a threshold. The probability with which a value of an attribute is determined may however depend on a buyer's corresponding offer (see below).

Besides specifying the utility of a (counter) proposal, the threshold is also used to determine when to respond to an outstanding proposal. More precisely, a seller strategy

responds with a fixed delay to all outstanding proposals which lie below the (current) threshold value; a proposal lies below the threshold value whenever the seller's utility for all offers in the proposal lies below the threshold value. While applying a delay to all proposals below the threshold value, the seller agent continues to negotiate with the remaining buyers by immediately responding with a counter proposal.

This negotiation without delay with the select group of buyers can, in principle, continue for several rounds. During these rounds the threshold is not adjusted. The goal at this time is to improve the Pareto-efficiency of the final agreement by finding mutually beneficial trade offs between the various attributes. In the simulation the seller agent only makes a single proposal to improve the efficiency of the deal. For a particular proposal of a buyer the seller strategy randomly generates offers within the neighborhood of the buyer's best offers, i.e., with the highest utility for the seller. This already suffices for very efficient outcomes. If a buyer does not accept one of the seller agent's offers, the seller agent will again respond with delay.

Another aspect that needs to be considered by the seller agent is the fairness of the agreements. Fairness prescribes that the seller should be indifferent between the deals made within the defined time interval. Whenever the seller agent almost simultaneously accepts two different offers a bargaining outcome may be unfair. The seller strategy ensures fairness by always making a (interesting) counter proposal, instead of accepting an offer directly. The seller agent can be equipped with a number of strategies for determining the threshold, which we introduce below.

Fixed and Time-Dependent Threshold Strategies For purpose of comparison we introduce a fixed threshold strategy. Clearly, the fixed threshold strategy is not capable of utilizing buyers' time pressure. The purpose of the strategy is to provide some insights in the minimal extractable profit, given strategic behavior of the buyers.

The second strategy we consider is a time-dependent threshold strategy: the current utility or threshold depends on time. The threshold only changes from one period to the next. Unlike the fixed-threshold strategy the time-dependent strategy is capable of utilizing buyers' time pressure. Its success, however, depends on how much it knows about buyers' preferences, or how easily more about buyers' preferences can be learned, in relation to time-based pricing strategies.

Responsive Threshold Strategies The fixed and time-dependent threshold strategies do not adjust the threshold based on the buyers' offers. Inspired by the first-price auction, we introduce another type of bargaining strategy with a responsive threshold. With this strategy, all offers submitted by the buyers within a certain fixed time interval are collected after the previous offers made by the seller agent. Then it determines the *current highest utility*, which is equal to the utility of the best offer from the collection of offers. The threshold is set to the current highest utility.

The success of the responsive threshold strategy does not depend on some (a priori) knowledge of buyer preferences, unlike the fixed and time-dependent strategies. Intuitively, buyers who— due to time pressure— suffer more from delay are inclined to bargain less “hard-headed” than other buyers. Consequently, these buyers may reach

a deal sooner and pay a higher price. Thus, at least potentially, the strategy is capable of utilizing buyers' time pressure without requiring (a priori) knowledge of buyer preferences. Unlike auctions, actual bargaining occurs in an alternating exchange of offers and counter offers, typically initiated by a buyer. Parties bargain over the price and other relevant aspects of the provided good or service. Even though the seller agent's strategy can be auction-inspired, buyers will be unaware of this fact. They do not know the opponent's bargaining strategy on forehand; they perceive the bargaining process as bilateral. Buyers may of course suspect some relationship with other ongoing negotiations. The point is that unlike a true auction the relationship with other simultaneously submitted offers is not specified up front, through a set of rules.

Reservation Value A drawback of the responsive threshold strategy is that it becomes vulnerable whenever groups of buyers experience very little time pressure. Without time pressure buyers have no incentive to buy soon. They may all independently decide to initially submit very low offers; consequently utility will be very low for the seller. To circumvent this we also consider responsive threshold strategies with a reservation value. A seller agent is never willing to sell below the reservation value. This means we alter the earlier definition of the current highest utility. It now becomes the maximum of the reservation value and the best offer from the offers collected within a certain time interval. An interesting advantage of introducing a reservation value occurs when some but not all buyers experience very little time-pressure. The responsive threshold strategy can then still utilize the time-pressure of the other buyers.

We consider two approaches for determining the reservation value. Either the reservation value is fixed, like the fixed-threshold strategy, or it is time dependent, like the time-dependent threshold strategy. Thus the responsive threshold strategy with a reservation value is actually a combination of the responsive threshold strategy (without reservation value) and either the fixed or time-dependent strategies.

3 Bargaining Simulation Environment

We apply a simulation environment in order to evaluate the performance and robustness of the above negotiation strategies against many learning buyers. The agents in the simulation are assumed to be boundedly rational: they can learn and adapt their strategies by a process of trial and error, and they do not know the seller's strategy. The bargaining process is repeated many times, enabling buyers and the seller to learn from past interactions. An evolutionary algorithm is used to model the learning aspect of the agents. This is a common approach within the field of agent-based computational economics (ACE) [11]. A number of related papers study bargaining using an evolutionary approach, e.g. [2–5]. Our approach extends previous research to multiple (types of) buyers and bilateral negotiation strategies for one-to-many multi-issue bargaining which can benefit from time pressure.

3.1 The Bargaining Game

The seller agent negotiates with many buyer agents simultaneously by alternating offers and counter offers as described in Section 2.2, where the buyer agents initiate the ne-

negotiations. For our simulations we set a maximum number of n discrete periods, where n is set sufficiently large such that it has no significant impact on the results. For the analysis we assume that offers consist of two interdependent attributes, e.g. the price and the quality. We note that buyer agents in the simulation may leave the negotiations prematurely (due to a bargaining break off probability) but do not enter later. We also assume that, since buyers are impatient, buyer agents in the simulation will respond to the seller agent's counter offers without delay. This is modeled by having the buyer's counter proposal or acceptance proposal occur in the same period as the seller's proposal.

3.2 Buyers and their Agents

Buyers are interested in buying at most one good in each bargaining game. They can have different preferences regarding the time pressure and attribute value combinations, which together constitute the buyer *type*. For the analysis we assume a finite number of k types. Although k is fixed, the number of participating buyer agents of each type varies randomly for each negotiation game and is determined independently by a Poisson distribution with average λ .

To illustrate the feasibility of our approach for interdependent attributes, we use the well-known Cobb-Douglas utility function to represent a player's preferences for the two attributes. More specifically, the utility u_i for buyer type i in case of a disagreement equals zero and in case of an agreement u_i is defined as

$$u_i = (v_{1,i} - o_1)^{\alpha_i} (v_{2,i} - o_2)^{\beta_i} \delta_i^t,$$

where α_i and β_i are parameters that indicate the relative importance of the attributes; o_1 and o_2 are the negotiated values the seller receives for the attributes (and the buyer has to give in); and $v_{1,i}$ and $v_{2,i}$ represent the maximum buyer i 's is willing to give in on the individual attributes. For example, let attribute 1 and 2 refer to price and quality. Then o_1 represents the price and o_2 the difference between the maximum quality and the actual quality of the good received; $v_{1,i}$ then represents the maximum price buyer i is willing to pay and $v_{2,i}$ the maximum buyer i is willing to give in on the quality. Furthermore, δ_i is the discount factor used to model the time pressure, and t is the negotiation time. In the simulation depreciation occurs at discrete time intervals. Therefore, δ is the discrete representation of time pressure and t indicates the period in which an agreement is reached. Note that discount factors are commonly used for modeling time pressure, e.g. in the Rubinstein-Ståhl alternating-offers model [10].

Buyer Agent's Strategy The buyer agents in the simulation apply time-dependent strategies similar to the seller's time-dependent threshold strategy described in Section 2.4. The buyer agent also uses an analogous (random) strategy for determining the values of the attributes given the threshold. The time-dependent strategy consists of a piece-wise linear function to determine the threshold. The parameters that determine the function are adaptive: using an evolutionary algorithm they evolve such that the performance of the strategy increases.

We also applied an extended strategy in our experiments by using two separate piece-wise linear functions: one produces the threshold for determining the utility level

of the offers and the other function determines the threshold for accepting or rejecting the seller's offers. The separation of the two functions enhances the bargaining capabilities of the buyer agent. Results using the two representations are very similar. The outcomes presented in this paper are based on the extended strategy.

3.3 Seller Agent

The seller agent bargains with a number of buyers simultaneously, without knowing the type of these buyers. The seller agent's utility in case of an agreement equals $u_s = o_1^{\alpha_s} o_2^{\beta_s}$, and is zero in case of a disagreement (recall from Section 2.3 that we can assume the seller has no time pressure). The total utility equals the sum of utilities obtained over all buyers. Production costs are set to zero.

We consider five strategies for the seller agent: fixed threshold, time-based threshold, responsive threshold and two combined strategies (see also Section 2.4). The time-based threshold strategy is similar to the strategy used by the buyer. The first two strategies and the combined strategies have parameters which determine respectively the threshold value and the reservation value during a bargaining game. These parameters are *adaptive*: optimal values are learned using an evolutionary algorithm, explained below. The responsive threshold strategy does not have any parameters that need to be learned.




Fig. 2. The EA cycle for negotiations with two buyer types and an adaptive seller

3.4 The Evolutionary Algorithm

Evolutionary algorithms (EAs) are a class of search algorithms inspired by Darwin's theory on variation and natural selection, and are becoming increasingly popular for modeling economic behavior, particularly within the field of agent-based computational economics (ACE), see e.g. [11]. We use an implementation based on "evolution strategies" [12], which is typical for real-valued encoding of the strategies (whereas the more popular branch of "genetic algorithms" is originally based on binary encoding).

The EAs are used to produce effective bargaining strategies for the buyer agents. Strategies for the agents of different buyer *types* are produced by separate EAs, which operate in parallel. This allows for heterogeneous strategies to emerge. Furthermore, in case of an adaptive seller agent, a separate EA is also used to produce strategies for the seller agent. A graphical representation of the evolutionary simulation with two buyer types and an adaptive seller agent is given in Fig. 2.

Each EA starts with a population of *parent* strategies, which are randomly generated. The EA then performs the following cycle to improve the quality or *fitness* of the strategies. First, the reproduction operator generates a population of *offspring* strategies by randomly selecting strategies from the parent population and slightly mutating the strategy to obtain variation.

In the next step, the fitness of the strategies is determined by the average utility obtained in a number of bargaining games. At the start of each bargaining game, the number of participating buyers of each type is determined randomly using a Poisson distribution as described above. Buyer agents are then generated for each buyer and are assigned a randomly selected strategy from either the parent or offspring population of the corresponding type. Similarly, a strategy is selected randomly for the seller agent (in case of an adaptive seller). The bargaining game is played for a fixed number of times, determining the number of buyers and assigning new strategies at the start of each game.

In the final stage of the cycle, a deterministic selection scheme called $(\mu + \lambda)$ -selection chooses the strategies with the highest fitness from both the parents and the offspring populations as the new parents for the next generation [12]. The cycle is repeated for a fixed number of generations.

Strategy encoding As mentioned in Section 3.2, the buyer agent's strategy consists of two piece-wise linear functions: an offer and a threshold function. The functions are encoded using real values, where each bending point of a function is encoded by two real values (i.e., the period and the corresponding threshold value). Additionally, two end points mark the values for the first and last period. For example, 8 real values are needed to encode a pair of functions with two line pieces each.

The same representation is used for the seller agent if he uses a time-based threshold strategy. If a fixed threshold is used, only a single real value is needed to encode this. Note that the seller agent uses the same function for both the threshold and for producing offers.

Mutation with Exponential Decay The mutation operator changes the strategy of an agent as follows. Each real value x_i is mutated by adding a zero-mean Gaussian variable with a standard deviation σ [12]: $x'_i := x_i + \sigma N_i(0, 1)$. All resulting values larger than unity (or smaller than zero) are set to unity (respectively zero). In our simulations, we use a model of exponentially decaying standard deviations. This approach ensures convergence and is analogous to simulated annealing, where a temperature parameter determines the variation of the solution. A half-life parameter determines the number of generations that the mutation standard changes to half the value.

4 Experimental Results

4.1 Settings

In this paper we report on two series of experiments with either a 0% or a 1% exogenous customer break off probability: a break off probability of 1% means that at every new round there is a 1% probability of an active customer breaking off (or leaving) the current bargaining game. The following settings are used for both series of experiments. (We note that also experiments are carried out using other settings, e.g. with a different number of participating buyers and buyer valuations, resulting in very similar outcomes, but are omitted due to space limitations.) Buyers are grouped into three types ($k = 3$), each type having adaptive bargaining strategies evolving in separate populations. The time pressure (discount factor) for each type is set as a control parameter. The values $v_{1,i}$ and $v_{2,i}$, and the parameters α_i and β_i are randomly generated from a uniform distribution at the beginning of each experiment, such that $v_{1,i}, v_{2,i} \in [100, 300]$ and $\alpha_i, \beta_i \in [0.7, 0.9]$. A buyer furthermore has a *minimum threshold value*, which is a minimum acceptable utility and is fixed at 10% of the most favourable utility (i.e., the utility u_i when $o_1 = o_2 = 0$, see Section 3.2).

The piece-wise linear functions of the buyer agents, and of the seller agent in case of time-based threshold strategy, consist of two line pieces. The number of buyers of each type participating in a bargaining game is determined randomly by a Poisson distribution with the average $\lambda = 10$. Buyers and sellers produce 3 offers in each round, which are randomly selected given a threshold value. However, when the seller produces counter offers without delay to improve Pareto efficiency (see Section 2.4), the seller generates 5 offers in the vicinity of the buyer's best offers. The length of a bargaining game is set to 40 periods.

The EA settings are chosen such that results are robust and the EAs are able to find good solutions. All buyer types use equal settings, with 20 strategies in the parent populations and 20 offspring strategies. The mutation standard deviation (see Section 3.4) is initially set to 0.2, and decays with a half-life value of 50 generations. The EA settings for the seller are the same, except that each seller population only contains 10 strategies. Buyers have larger populations because more buyers than sellers participate each game, and because in case of the extended buyer strategy (with two functions) the search space for the buyer is larger (a higher population size is often recommended for larger search spaces). The fitness of the strategies for a single generation is determined by 100 bargaining games. For these settings the EAs are able to find almost optimal solutions for simple test cases.

4.2 Results

The results reported in this Section are obtained after a process of learning, when the strategies have converged. It is important to note that, during learning, the preferences of the buyers remain unchanged, although the number and composition (i.e., number of each type) of buyers can differ in each bargaining game. Experiments are run for 40000 bargaining games (400 generations). Results are averaged over the last 1000 bargaining games of an experiment, and over 30 experiments, accounting for random settings such as the number of participating buyers and the buyer's preferences.




Fig. 3. Seller's obtained fraction of total surplus, with (a) 0% (on the left) and (b) 1% (on the right) break off probability, using 5 threshold strategies: (1) fixed threshold, (2) time-dependent threshold, (3) responsive threshold, (4) combined (3) and (1), and (5) combined (3) and (2).

Figure 3 compares—for a break off probability of (a) 0% (on the left) and (b) 1% (on the right)—the obtained fraction of the total *seller surplus* for different seller threshold strategies and buyer discount factors (buyers have equal discount factors). We define the seller surplus of a bilateral negotiation as the seller's maximum feasible utility, i.e., when the buyer offers her minimum threshold value and the offer is Pareto-efficient. As shown in Fig. 3, the fixed threshold strategy (1) is able to extract around 75% of the seller surplus. The outcomes are relatively independent of the break off probability; this is because almost all deals will be closed in the first or second round (the average round of agreement lies between 0.14 and 0.24 in case of time pressure). Note that these outcomes are independent of the discount factor. Clearly, the fixed threshold strategy is unable to benefit from the buyers' time pressure.

The time-based threshold strategy (2), on the other hand, shows that higher profits can be obtained if the threshold changes over time, see Fig. 3a. Buyers with a high valuation will purchase relatively early, since waiting for a better deal does not compensate the loss due to time discounting. Buyers with a low valuation, on the other hand, have the incentive to reach an agreement in a later stage if they can get a better price for it. This way the seller can indirectly discriminate between buyers with different valuations and time pressures. The performance of the time-based threshold strategy (2) is, however, vulnerable to an increase in the bargaining break off probability (see Fig. 3b).

Note that with no time discounting (i.e., when $\delta = 1$) the fixed threshold strategy performs better. This is due to the difference in strategy complexity: only a single value needs to be optimized in case of a fixed threshold, whereas an entire function (encoded by 4 values) needs to be learned in case of the time-based threshold. This is clearly more difficult, especially within a dynamic environment with learning buyers.




Fig. 4. Average inefficiency of the agreements using different strategies with (a) 0% (on the left) and (b) 1% (on the right) break off probability. The numbers correspond to the seller threshold strategies in Fig. 3.

Outcomes using the responsive threshold bargaining strategies (see Fig. 3 (3),(4), and (5)) show an impressive increase in the fraction of surplus when buyers are impatient. If the time pressure becomes sufficiently high, the seller obtains almost the entire surplus. Even for lower time pressure, results are much better for the seller compared to the fixed and time-based threshold strategies. The consequence of increasing the bargaining break off probability from 0% to 1% is that buyers' time pressure needs to be a bit higher before the threshold strategies will dominate the other strategies and the maximal attainable performance drops a bit. This drop in performance is, however, mainly caused by lower sales due to the premature departure of customers; thus the seller's bargaining position does not change fundamentally.

For the case of no or very low time pressure, the results also show that simple auction-like mechanisms such as the responsive threshold strategy are not sufficient in case of unlimited supply. Without competition between buyers, the market price goes to cost level, resulting in a zero surplus for the seller. This problem can be resolved by combining the responsive threshold strategy with an adaptive reserve value. As shown in Fig. 3, this results in very good outcomes, even if buyers are very patient. This makes the combined strategy very versatile. We note that these outcomes also generalize to settings where buyer types have different time preferences, assuming that buyers with higher valuation have a higher time pressure. The outcomes are not shown here due to space limitations.

The mean (in)efficiency of the obtained bargaining results is depicted in Fig. 4. The inefficiency is measured as the seller's maximum Pareto-improvement of a given outcome divided by the seller's actual utility plus the improvement. The outcomes show

low inefficiencies for all strategies (the highest inefficiency is around 2.7% of the total utility). However, the responsive threshold strategies result in relatively the most Pareto-efficient deals. Unlike the other strategies, the responsive strategies set the threshold exactly to the best offer. Since this is the buyer's best offer, it is already quite efficient. In case of the other strategies, however, the utility of the offers usually exceed the seller's threshold, and the seller first needs to map the buyer's offers to the right utility level when making counter offers. This mapping results in additional inefficiencies of the outcomes. Even with a reasonably simple strategy for determining the relative magnitude of the attribute values, already good results are found. The Pareto-efficiency is expected to improve even further by incorporating more advanced strategies as described in e.g. [7, 6, 8, 9]. This is however left for future work.

4.3 Bargaining Revisited

A possible strategy of the buyer agent is to bid very low, and then accept the counter offer of the seller. Such a strategy could be beneficial in case the seller's counter offer is influenced by the buyers' offers, as with the responsive threshold strategies. This could then result in low profits for the seller. To see if indeed buyers benefit from such a strategy, the strategy representation for buyers was extended by using two separate functions: one produces the threshold for determining the utility level of the offers and the other function determines the threshold for accepting or rejecting the seller's offers (see Section 3.2). Even with separated function, however, the responsive threshold strategy performs very much in favor of the seller (as shown by the results). This occurs because the counter offer is delayed by the seller whenever offers fall below the (seller's) threshold, hence providing the buyers with an incentive to try and get an agreement without delay.

5 Concluding Remarks

In this paper, we consider strategies for a seller agent who negotiates with many buyers simultaneously in a bilateral fashion over multiple interdependent attributes. These strategies respect a notion of fairness such that buyers are treated similarly. An important aspect of the developed strategies is their ability to benefit from impatient buyers that prefer early agreements. Buyers can have different valuations and time preferences. A buyer's actual valuation and time preference is only known to himself (i.e., a buyer's type constitutes private information).

The strategies introduced determine three aspects: a threshold, multi-attribute offers with a utility level corresponding to the threshold, and a scheme for determining when to respond. Five different threshold strategies for the seller agent are evaluated and compared: (1) fixed threshold, (2) time-dependent threshold strategies, (3) responsive, (4) responsive with fixed reservation value, and (5) responsive with time-dependent reservation value. The last two strategies are actually a combination of the responsive threshold strategy with the first two strategies.

We use an evolutionary simulation to analyze the performance of the different strategies when the buyers and seller bargain over two interdependent attributes. The buyers'

bargaining strategies adapt and learn through the use of an evolutionary algorithm (EA). The seller's strategies (1) and (2), and the combined strategies (4) and (5) also adapt and learn using an EA. The responsive threshold strategy (3), on the other hand, determines the threshold value based exclusively on the offers received by the buyers, and does not require any learning.

The outcomes show that bilaterally exchanging multiple offers combined with a random offer generation mechanism suffices for closely approximating Pareto-efficiency. Furthermore, given a small probability (per negotiation round) of a customer breaking of the negotiations the responsive threshold strategies appear to be very successful in utilizing time pressure and consequently extract a very high share of the surplus. For sufficiently high time pressure, the seller obtains almost all surplus, indicating that buyers submit and/or accept offers close to their reservation value. Thus buyers self-select to pay their valuation, while the bargaining outcomes respect our notion of fairness. The results also show superior performance of the combined strategies (4 and 5) compared to the auction-inspired strategy (3), in case some or all buyers have very little time pressure. In other words, the combined strategy is very versatile.

References

1. Goldberg, A.V., Hartline, J.D.: Envy-free auctions for digital goods. In: Proceedings of the 4th conference on Electronic commerce, ACM, ACM (2003) 29–35
2. Gerding, E., van Bragt, D., Poutré, J.L.: Multi-issue negotiation processes by evolutionary simulation: Validation and social extensions. *Computational Economics* **22** (2003) 39–63
3. Fatima, S., Wooldridge, M., Jennings, N.R.: Comparing equilibria for game theoretic and evolutionary bargaining models. In: Proc. 5th Int. Workshop on Agent-Mediated E-Commerce, Melbourne, Australia. (2003) 70–77
4. Phelps, S., McBurney, P., Parsons, S., Sklar, E.: Co-evolutionary auction mechanism design: A preliminary report. In Padget, J., Shehory, O., Parkes, D., Sadeh, N., Walsh, W., eds.: *Agent-Mediated Electronic Commerce IV: Designing Mechanisms and Systems*. Volume 2531 of Lecture Notes in Computer Science. Springer-Verlag, Berlin (2002) 123–142
5. Cliff, D.: Explorations in evolutionary design of online auction market mechanisms. *Electronic Commerce Research and Applications* **2** (2003) 162–175
6. Klein, M., Faratin, P., Sayama, H., Bar-Yam, Y.: Negotiating complex contracts. *Group Decision and Negotiation* **12** (2003) 111–125
7. Faratin, P., Sierra, C., Jennings, N.R.: Using similarity criteria to make issue trade-offs. *Journal of Artificial Intelligence* **142** (2003) 205–237
8. Ehtamo, H., Hämäläinen, R.P.: Interactive multiple-criteria methods for reaching pareto optimal agreements in negotiations. *Group Decision and Negotiation* **10** (2001) 475–491
9. Somefun, D., Gerding, E., Bohte, S., Poutré, J.L.: Automated negotiation and bundling of information goods. In Rodriguez-Aguilar, J.A., Faratin, P., Walsh, W.E., Parkes, D., eds.: *Workshop on Agent Mediated Electronic Commerce V: Designing Mechanisms and Systems*. Volume 3048 of Lecture Notes in AI., Heidelberg, Springer-Verlag (2004)
10. Rubinstein, A.: Perfect equilibrium in a bargaining model. *Econometrica* **50** (1982) 97–109
11. Tesfatsion, L.: (Guest Editor), Special issue on agent-based computational economics. *Journal of Economic Dynamics and Control* **25** (2001)
12. Bäck, T.: *Evolutionary Algorithms in Theory and Practice*. Oxford University Press, New York and Oxford (1996)