

Efficient Methods for Automated Multi-Issue Negotiation: Negotiating over a Two-Part Tariff

D.J.A. Somefun,^{1,*} E.H. Gerding,^{1,†} J.A. La Poutré^{1,2,‡}

¹*Centrum voor Wiskunde en Informatica, Centre for Mathematics and Computer Science, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands*

²*School of Technology Management, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands*

In this article, we consider the novel approach of a seller and customer negotiating bilaterally about a two-part tariff, using autonomous software agents. An advantage of this approach is that win-win opportunities can be generated while keeping the problem of preference elicitation as simple as possible. We develop bargaining strategies that software agents can use to conduct the actual bilateral negotiation on behalf of their owners. We present a decomposition of bargaining strategies into concession strategies and Pareto-efficient-search methods: Concession and Pareto-search strategies focus on the conceding and win-win aspect of bargaining, respectively. An important technical contribution of this article lies in the development of two Pareto-search methods. Computer experiments show, for various concession strategies, that the respective use of these two Pareto-search methods by the two negotiators results in very efficient bargaining outcomes while negotiators concede the amount specified by their concession strategy. © 2006 Wiley Periodicals, Inc.

1. INTRODUCTION

Automated bilateral negotiation forms an important type of interaction in agent based systems for electronic commerce.¹ It allows seller and customer to determine the terms and content of the trade iteratively and bilaterally. Consequently, deals may be highly customized (especially for complex goods or services) and highly adaptable to changing circumstances. Moreover, by automating the negotiation, the potentially time-consuming process is delegated to autonomous software agents who conduct the actual negotiation on behalf of their owners.

*Author to whom all correspondence should be addressed: e-mail: koye@cw.nl.

†e-mail: egerding@cw.nl.

‡e-mail: hlp@cw.nl.

INTERNATIONAL JOURNAL OF INTELLIGENT SYSTEMS, VOL. 21, 99–119 (2006)
© 2006 Wiley Periodicals, Inc. Published online in Wiley InterScience
(www.interscience.wiley.com). • DOI 10.1002/int.20122



In particular, information goods mark an application area where negotiating through the use of autonomous software agents may be beneficial. (For information goods, the cost of producing one additional unit is virtually zero.) Software, news articles, stock quotes, research reports, music, and video clips are all examples of information goods. Information goods are often sold through on-line services where customers pay a *nonnegotiable* subscription fee. Reuters News Services, Apple's iTunes Music Store, and Microsoft's DotNet (MSDN) are examples of this practice. Generally, customers have only an interest in a limited subset of, for example, all the news articles, music, or software available through such on-line services. Especially, with a high heterogeneity among customers and a lot of dynamics in customer preferences, offering nonnegotiable subscriptions may not be the best solution. Bargaining has, on the contrary, the advantage of providing more flexibility and adaptability toward the preferences of individual customers and daily dynamics. In this article, we therefore consider the problem of a seller and customer agent negotiating bilaterally about the subscription fee of such an on-line service.

Most commonly, businesses use an instance of an n -part tariff as the *nonnegotiable* subscription fee: An n -part tariff specifies a fixed price (or fee), $n - 1$ variable prices, and the corresponding $n - 1$ "intervals" in which the variable prices are applicable; the quantity or more generally the quality of the goods purchased is used to define these intervals. Given the charged tariff, a customer will choose the desired quantity or more generally the desired quality of the good or service on sale. A good example of a nonnegotiable two-part tariff is the practice of network operators for mobile phones; they charge a fixed subscription fee and a variable price per minute called. Charging a nonnegotiable two-part tariff is especially a natural choice in case of uncertainty in the quantity or quality customers will choose to purchase (e.g., minutes called in a month).

In this article, we consider the *novel* approach of software agents negotiating about a two-part tariff. For most practical applications, it suffices to focus on negotiating about a two-part tariff for the following three reasons (see Section 3 for a more in-depth discussion). First, negotiating about such tariffs has the advantage, over one-part tariffs, of generating win-win opportunities. Second, negotiating about more complex tariffs will often generate very few additional win-win opportunities, while the issue of preference elicitation may become problematic. Third, negotiating about a two-part tariff has the same discriminative power as nonnegotiable $m + 1$ -part tariffs ($m > 2$ being the number of customers), whereas individual customers do not have to understand the more complex $m + 1$ -part tariff. Therefore, price discrimination will not be a motive for negotiating about more complex tariffs.

We address the issue of negotiating about a two-part tariff by decomposing bargaining strategies into two parts: Pareto-search methods and concession strategies. Pareto-search methods aim at utilizing the win-win aspect of multi-issue bargaining; they determine the relative magnitude of the multiple issues, during the bargaining process. *At the same time* (but independently) a concession strategy determines the desired utility level of a deal during the bargaining process.

An *important* contribution of this article lies in the actual development of Pareto-search methods that result in efficient solutions while, at the same time, bargainers make concessions using a variety of concession strategies. To that end, we introduce the orthogonal and orthogonal-DF method: two Pareto-search methods. The orthogonal-DF combines the orthogonal search method with an amplifying mechanism, the Derivative Follower (with adaptive step size). We show through computer experiments that the respective use of these two Pareto-search methods by the two bargainers—combined with various concession strategies—results in very efficient bargaining outcomes when bargaining about two-part tariffs; that is, these outcomes closely approximate Pareto-efficient bargaining solutions.^a Note that we obtain these results without assuming any prior knowledge of the opponent's preferences: A simultaneous use of the orthogonal method by one and orthogonal-DF by the other bargainer results in bargainers learning to close efficient deals through an alternating exchange of offers and counteroffers.

Clearly, bargainers are generally free to choose strategies other than a Pareto-efficient-search method combined with a concession strategy. Fully rational bargainers who view the bargaining process as a one-shot game will not necessarily opt for such an approach. Bargaining between customer and seller is, however, generally conducted repeatedly. Obtaining win-win (Pareto-efficient) deals can therefore be a goal for both seller and buyer, to avoid unnecessary loss of payoff for the opponent. A seller may care, for instance, more about long-term profitability. Exploiting win-win opportunities can contribute to such a long-term goal, because it results in a higher customer satisfaction, which can lead to returning customers (without necessarily reducing the short-term profit). To give rise to the slightly cooperative aspect of repeated bargaining, negotiators can use the developed Pareto-search methods with a self-selected concession strategy. It goes beyond the scope of the article to discuss exhaustively under which exact conditions bargainers have an incentive to choose the approach of decomposing bargaining strategies. The point we make in this article is that the developed decomposition of the bargaining strategy into Pareto-search methods and concession strategies can be an important option for bargainers; it results in efficient deals while still leaving room for noncooperative behavior through bargainers' concession strategy.

The remainder of the article is organized as follows. In Section 2, we relate the article to the literature. In Section 3, we explain in greater detail that, when negotiating about a tariff, it suffices to consider only two-part tariffs for most practical applications. In Section 4, we discuss the customer and seller agent. In Section 5, we introduce the approach of disentangling multi-issue bargaining strategies into Pareto-search methods and concession strategies. Moreover, we develop the orthogonal and orthogonal-DF methods: two Pareto-search methods. To assess the relative performance of the developed Pareto-search method we consider other standard techniques as candidate Pareto-search methods in Section 6. In Section 7,

^aAn offer constitutes a Pareto improvement compared to another offer whenever it makes one bargainer better off without making the other worse off. An offer is said to be Pareto efficient whenever there does not exist a different offer that constitutes a Pareto improvement.

we study in greater detail the Pareto-search aspects of bargaining. Through computer experiments we investigate the efficiency of the introduced bargaining approach. Conclusions follow in Section 8.

2. RELATED WORK

With one-issue negotiation, bargainers only need to determine the (desired) utility level of a (counter) offer. Negotiating over more than one issue complicates things; bargainers also need to determine the magnitude of the various issues. By making clever trade-offs between various issues, win-win or Pareto improvements may occur. In Ref. 2 a heuristic is introduced that tries to find Pareto improvements by searching for the most similar offer, given a desired utility level. They use fuzzy similarity criteria because most of the considered issues take on very limited discrete values. Based on these similarity criteria, an iterative hill-climbing algorithm is used to find the most similar offer. This hill-climbing algorithm is limited, however, to linearly additive utility functions.

In this article, we develop the orthogonal and orthogonal-DF, two Pareto search methods. Like the heuristic developed in Ref. 2, the orthogonal search method tries to find Pareto improvements by searching for the most similar offer, given a desired utility level. We, however, consider negotiation over a continuous issue (or issue that can take on many values). For this problem domain, Euclidean distance is a more natural choice as the similarity criterion. With Euclidean distance, standard mathematical techniques (from fields such as convex analysis) are immediately at our disposal. Moreover, implementing the similarity criterion entails a straightforward application of standard techniques from numerical analysis. The orthogonal search method finds—from the collection of offers that have the desired utility level—the offer closest to the opponent's last offer, measured in Euclidean distance. Unlike the heuristics developed in Ref. 2, it is not restricted to linearly additive utility functions.

An important limitation of any search method using similarity criteria only—such as the orthogonal search method and the heuristics developed in Ref. 2—is that the rate of convergence depends solely on bargainers' preferences and the chosen similarity criteria. This convergence rate is often very slow. The orthogonal-DF combines the orthogonal search method with an amplifying mechanism, the Derivative Follower (with adaptive step size). Consequently, it can very quickly converge to a Pareto-efficient solution.

Slow rate of convergence is especially a problem whenever software agents are not a priori restricted. In that case, bargainers may simultaneously make a concession (i.e., adjust the desired utility level) and search for clever trade-offs. Due to these concessions, the Pareto-efficient solution, to which a search method converges, also changes. With a very slow convergence rate, very little improvement in the efficiency can be realized before a deal is closed. Thus the heuristic developed in Ref. 2 may need a lot of time to find a Pareto-efficient solution, time that may not be available due to bargainers making concessions simultaneously. The developed orthogonal-DF needs less time and consequently can also work very well in conjunction with concessions.

In the literature, the difficulties with bargainers simultaneously making concessions and searching for clever trade-offs is generally avoided by assuming an intermediary.³⁻⁹ The mediator is inspired by the idea of a *single negotiation text* (SNT). SNT is a mediation device suggested by Roger Fisher.¹⁰ During negotiation, the mediator first devises and proposes a deal (SNT-1) for the two bargainers to consider. The mediator is not trying to promote the first proposal; rather, it is meant to serve as an initial, single negotiation text—a version to be criticized by both parties and then modified in an iterative manner. Modifications to the SNT-1 will be made by the mediator based on the criticisms from the two sides. Thus, both parties need to reveal (aspects of) their preferences to the mediator; hence, trust becomes an important issue. Furthermore, additional costs are often involved with a mediator.

The orthogonal-DF method is somewhat related to the work of Ehtamo et al.^{3,4} They develop the *method of improving directions*, which is a mathematical formalization of the SNT method (with a mediator). In essence it is a multicriteria decision-making gradient search method. Given an SNT, bargainers give their most preferred direction of the next SNT, which is just the gradient. The mediator then uses some relatively straightforward procedure to determine the jointly improving direction that is then used to determine the next SNT. The orthogonal-DF also searches for such a jointly improving direction, but without the use of a mediator.

3. NEGOTIATING ABOUT AN n -PART TARIFF

A typical businesses practice is to use an instance of an n -part tariff as the *nonnegotiable* subscription fee. Based on the observations discussed in the following two subsections, we can conclude that, for most practical applications, it suffices to focus on a two-part tariff when negotiating about a tariff. In the remainder of this article, we therefore focus on software agents negotiating about a two-part tariff.

3.1. Win-Win Opportunities

An advantage of negotiating about a two-part tariff is that, unlike a one-part tariff, win-win opportunities are possible. This win-win or Pareto-efficiency aspect arises due to the fact that, unlike one-issue negotiation, trade-offs between issues can be made. Mutually beneficial trade-offs are possible, for example, when seller and buyer predict the quantity the buyer will purchase, differently. Consequently, buyer (seller) can be indifferent between certain fixed and variable price combinations whereas the opponent strictly prefers one of these combinations over the other.

These win-win opportunities can arise whenever buyer and seller negotiate about an n -part tariff (for any $n > 1$). Two-part tariffs are, however, commonly used as a subscription fee; customers will generally be familiar with the pricing scheme. Due to this familiarity and the limited complexity of the scheme, it will be relatively straightforward for a customer to instruct an agent to negotiate on his

behalf, that is, preference elicitation¹¹ is less of an issue than when negotiation is about more complex tariffs. Moreover, often more complex tariffs will generate very few additional win–win opportunities; the magnitude of these extra win–win opportunities depends on the extent to which both negotiators believe higher order tariff changes (such as changes in the variable price) influence the (expected) quantity or quality purchased by the buyer.

3.2. Price Discrimination

Charging a nonnegotiable two-part tariff is especially a natural choice in case of uncertainty in the quantity or quality customers will choose to purchase (e.g., minutes called in a month, in the case of a network operator for a mobile phone). Typically, businesses use more complex nonnegotiable tariff schemes to apply price discrimination, that is, charging different prices to different customers. They offer customers a menu of options where a nonnegotiable tariff assigns a price to each option in the menu. Customers can then self-select the best purchase (cf. Refs. 12–15). An n -part tariff can, for example, be used to introduce quantity discounts, where the variable price declines in successive intervals (see Figure 1a for an example with a four-part tariff).

In this article, we introduce the possibility of negotiating about the tariff. Price discrimination is now no longer a motive for introducing more complex tariffs, because negotiating about a two-part tariff can have the same discriminative power as a *nonnegotiable* $m + 1$ -part tariff (where m is equal to the total number of customers). To see this, note that negotiating with m customers about a two-part tariff can result in m different two-part tariffs. Moreover for any $n + 1$ -part tariff there exist n two-part tariffs (with $n > 1$) that, given rational behavior, will result in an equivalent revenue (cf. Ref. 12). Thus, negotiating with m customers about their two-part tariff can mimic the situation of a nonnegotiable $m + 1$ -part tariff.

Figure 1 illustrates this revenue equivalence between a $m + 1$ -part tariff and m two-part tariffs: Figure 1a depicts a four-part tariff with one fixed price and, depending on the quantity purchased, three different variable prices; Figure 1b depicts the corresponding 3 two-part tariffs with one fixed and one variable price. A rational customer will be indifferent between either the 3 two-part tariffs or the four-part tariff. Faced with the 3 two-part tariffs, in Figure 1b, a rational customer (who spends the smallest amount possible to buy the desired goods) will only consider the solid parts of the 3 two-part tariffs, which is the four-part tariff in Figure 1a. Consequently, both the four-part and the 3 two-part tariffs will result in the same revenue.

4. AGENTS AND BARGAINING PROTOCOL

4.1. Bargaining Protocol

Following Ref. 16, bargaining occurs in an alternating exchange of offers and counteroffers, typically initiated by the customer. An offer specifies the fixed price

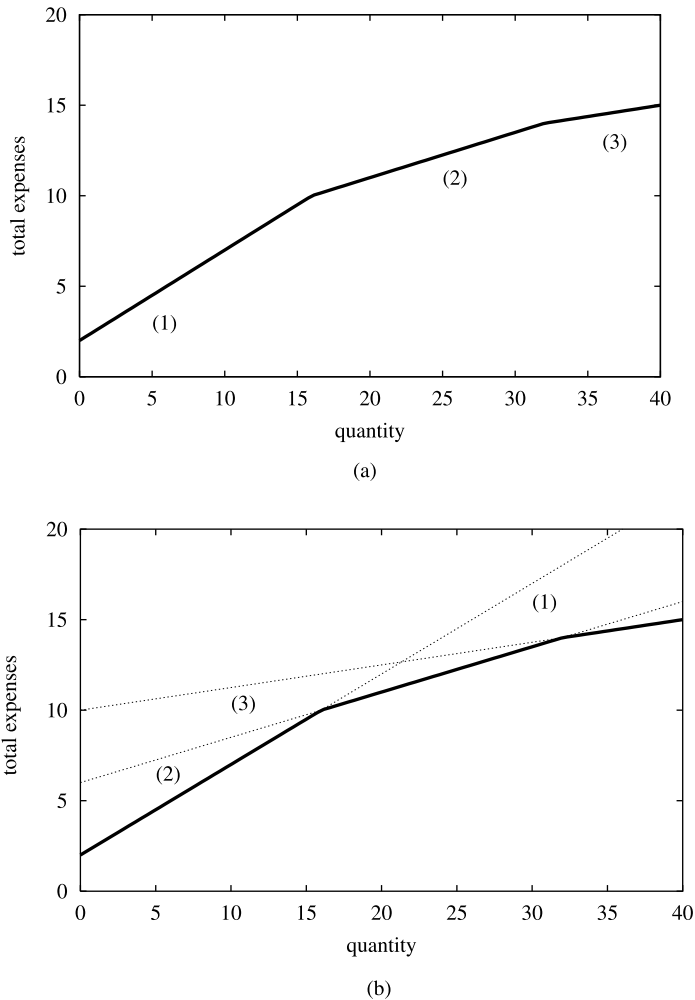


Figure 1. (a) A 4-part tariff. (b) The corresponding 3 two-part tariffs (dotted lines); labels (1), (2), and (3) identify the two-part tariffs that correspond to the line segments of the four-part tariff. When faced with the 3 two-part tariffs in (b), a rational customer will only consider the solid parts of the 3 two-part tariffs, which is the four-part tariff in (a).

and the variable price. The fixed price represents the subscription fee and the variable price gives the price for every unit purchased. Attached to an offer are preconditions that specify until when the offer is valid. We call the offer combined with the preconditions a *proposal*. The bargaining process continues until an agreement is reached or one of the bargainers terminates the process. Based on this bargaining process, Figure 2 depicts the alternating-offers bargaining protocol a customer agent and seller agent use to do the actual bargaining.

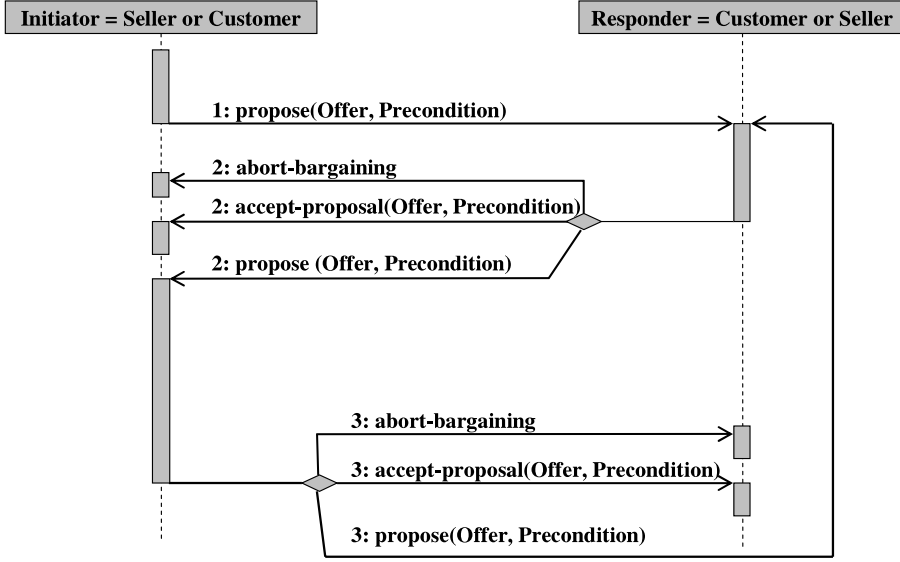


Figure 2. The bilateral bargaining protocol of alternating offers: numbers 1, 2, and 3 mark the first, second, and third offers submitted in a sequence of three offers: the second and third offers are the counteroffers to the first and second offers. After the third offer, the process repeats itself.

4.2. Agents

4.2.1. Seller Agent

Bargaining with a customer is done based on the seller agent's desired aspiration level expressed in expected utils. We define the expected utility U_s of the seller agent as the expected revenue from selling a tariff (p_f, p_v) , where p_f denotes the fixed and p_v the variable price, that is,

$$U_s = p_f + p_v \cdot \rho_s(p_v) \quad (1)$$

where $\rho_s(p_v)$ denotes the expected number of additional units consumed for a variable price of p_v .

4.2.2. Customer Agent

The customer agent acts on behalf of the customer. The customer can indicate her preferences by specifying the number of units she expects to purchase ρ_c and a maximum budget b_{max} . The budget provides the agent with a mandate for the negotiation; the total expected costs should not exceed b_{max} . As an approximation of a customer's utility, the agent uses the following function whenever a deal is closed:

$$U_c = b_{max} - (p_f + p_v \cdot \rho_c) \quad (2)$$

Moreover, without loss of generality, we assume a utility level of 0 in case of a disagreement.

5. DECOMPOSING BARGAINING STRATEGY

We decompose bargaining strategies into concession strategies and “Pareto-search” methods. Concession strategies determine what the desired utility level of an offer will be, given a particular sequence of offers and counteroffers. Heuristics that implement Pareto-search methods determine, for a particular utility level and a particular history of offers and counteroffers, what the multi-issue offer will be, that is, the fixed price p_f and the variable price p_v of the two part tariff. In terms of a multivariable utility function, a counteroffer entails both a movement off the iso-utility line and a movement along the iso-utility line. (Given a specified utility level, an iso-utility line specifies all the p_f and p_v points that generate the same utility.) Concession strategies determine the movement of an iso-utility line; Pareto-search strategies determine the movement along an iso-utility line.

Pareto-search methods aim at reaching agreement as soon as the respective “concession” strategies permit this. Therefore, it may be good for both parties to use it. In more economic terms a negotiated tariff is called Pareto efficient if it is impossible to change the tariff without making one of the bargainers worse off, that is, one of the bargainers will always attach a lower (or equal) utility to the adjusted tariff. From a system design perspective, Pareto efficiency of the negotiated tariffs is clearly desirable.

In Sections 5.1 and 5.2 we introduce the orthogonal and orthogonal-DF search methods, two Pareto-search methods. The experiments in Section 7 show that if the customer agent uses one Pareto-search method and the seller agent uses the other one—while both simultaneously make concessions—then the bargaining outcome will closely approximate a Pareto-efficient solution for a wide variety of concession strategies.

5.1. Orthogonal Search

Both customer agent and seller agent may use—what we call—an orthogonal strategy as the Pareto-search algorithm. This strategy is probably best explained through an example. Suppose the customer (with whom the seller bargains over the tariff) placed the t th offer of $(p_f(t), p_v(t))$. Moreover, the seller’s concession strategy dictates a *utility level* of $U_s(t + 1)$, that is, in utils the (counter) offer should be worth $U_s(t + 1)$. Based on this information the orthogonal strategy determines $(p_f(t + 1), p_v(t + 1))$, the counteroffer of the seller, by choosing a (p_f, p_v) combination that generates $U_s(t + 1)$ utils and lies closest (measured in Euclidean distance) to the point $(p_f(t), p_v(t))$. Figure 3 gives the graphical representation of the orthogonal strategy.

The use of the orthogonal strategy by both parties results in a mapping f from a bargainer’s desired utility level at t to the desired utility level at $t + 2$. Given convex preferences (cf. Ref. 17) and fixed desired utility levels, the mapping f can be shown to satisfy the Lipschitz condition $\|f(x) - f(y)\| \leq \|x - y\|$ for all x and y




Figure 3. Example of orthogonal strategy, which determines $(p_f(t+1), p_v(t+1))$, the counteroffer of the seller, by choosing a (p_f, p_v) combination that generates $U_s(t+1)$ utils and lies closest (measured in Euclidean distance) to the point $(p_f(t), p_v(t))$.

in the domain of f .^b Thus, given fixed aspiration levels and convex preferences, the orthogonal strategy does imply that consecutive offers do not diverge. Figure 4 illustrates the use of the orthogonal strategy by both parties for the case of tangent iso-utility lines. It depicts a sequence of two offers and counteroffers with convex preferences and a fixed aspiration level. (U_s and U_c denote the iso-utility lines of the seller and customer, respectively.) The figure illustrates, for instance, how the customer's offer at time 1 (with aspiration level $U_c(1) = U_c$) is transformed into an offer at time 3 (with aspiration level $U_c(3) = U_c$).

5.2. Orthogonal-DF Search

The use of just the orthogonal strategy by both parties may lead to very slow convergence to Pareto-efficient bargaining outcomes. To speed up the convergence process we can add an amplifying mechanism to the orthogonal method. As the amplifying mechanism, we use the derivative follower with adaptive step size (ADF). (Henceforth we will call this the orthogonal-DF.)

The derivative follower (DF) is a local search algorithm (cf. Ref. 19). The DF modifies the variable price p_v found by the orthogonal strategy by either increasing or decreasing p_v with a step-size $\delta > 0$ in the direction $s_{dr} \in \{-1, +1\}$.

^bThe proof is a straightforward application of convex analysis (cf. Ref. 18) given that, without loss of generality, we can assume that the preferences are bounded. That is, negative and extremely high $(p_f(t), p_v(t))$ combinations can be discarded, without loss of generality.




Figure 4. Example of the use of the orthogonal strategy by both parties for a sequence of two offers and counteroffers, with fixed aspiration levels and convex preferences. The combination (p_f^*, p_v^*) denotes a Pareto-efficient tariff.

Before the DF modifies p_v the search direction s_{dr} is updated, however. Suppose last time the DF increased the price p_v (i.e., $s_{dr} = +1$). The DF then continues to increase the price whenever the opponent's last offer lies "closer" to the current iso-utility level than the opponent's before last offer; otherwise the DF changes the search direction by setting s_{dr} equal to -1 . (The current iso-utility curve specifies all (p_f, p_v) combinations that resulted in the utility currently desired.) We say that a DF "turns" whenever it adjust the search direction. Algorithm 1 specifies the orthogonal-DF in greater detail and Figure 5 illustrates the use of the orthogonal-DF by the seller (the customer uses the orthogonal strategy only).

Algorithm 1: The orthogonal-DF algorithm. The following is given: (a) the opponent's last and before last offer $O_1 = (p_f(t), p_v(t))$ and $O_2 = (p_f(t-2), p_v(t-2))$, respectively; (b) bargainer's utility function $u(p_f, p_v)$ and next aspiration level $U(t+1)$; (c) the step-size δ ; and (d) the search direction $s_{dr} \in \{-1, +1\}$. Based on this information the orthogonal-DF computes the next counteroffer $O = (p_f(t+1), p_v(t+1))$.

- (1) Use the orthogonal strategy to compute $O'_1 = (p'_f(t), p'_v(t))$ and $O'_2 = (p'_f(t-2), p'_v(t-2))$, that is, the points in the (p_f, p_v) plane that generate $U(t+1)$ utils and lie closest to O_1 and O_2 , respectively.
- (2) Compute d_1 and d_2 , the distance of the opponent's last two offers, that is, $d_1 = \|O_1 - O'_1\|$ and $d_2 = \|O_2 - O'_2\|$ ($\|\cdot\|$ denotes Euclidian distance).
- (3) Update s_{dr} : Whenever $d_1 > d_2$, the orthogonal-DF "turns," that is, $s_{dr} = -1 \cdot s_{dr}$; otherwise $s_{dr} = s_{dr}$.




Figure 5. Sequence of two offers and counteroffers with fixed aspiration levels where the seller uses the orthogonal-DF and the customer only uses the orthogonal strategy.

- (4) Update δ : Decrease δ whenever the orthogonal-DF turns. For a number of periods directly after a turn δ is not increased, and otherwise δ is increased (cf. Ref. 20 for the details).
- (5) Compute the counteroffer $O = (p_f(t+1), p_v(t+1))$: Set $p_v(t+1) = p'_v(t) + \delta \cdot s_{dr}$. Next choose $p_f(t+1)$ such that, given $p_v(t+1)$, the counteroffer generates $U(t+1)$ utils.

The orthogonal-DF uses the ADF. The difference between ADF and DF is that the step-size δ becomes adaptive.^{20,21} We use the ADF proposed in Ref. 20. The ADF starts by (exponentially) increasing the step-size δ as long as it continues to search in the same direction. After a turn occurs (i.e., it changes the search direction), the ADF reduces its step size (exponentially). Because a turn typically signals that an optimum has been passed, the ADF does not increase the step size for a number of steps directly after a turn (to avoid a repeating “overshoot” of the optimum). This pause in increasing the step size is set long enough such that, within the vicinity of an optimum, the ADF does not increase its step size. Consequently, the ADF converges very quickly to local optima whenever the optimization problem is static (cf. Ref. 20 for the proofs). Due to concessions, finding a Pareto-efficient deal is not necessarily a static problem. Using the ADF as an amplifier for the orthogonal method can still work very well, however. Opponent’s concessions can cause the orthogonal-DF to overshoot an optimum a bit more than it otherwise would. Because it increases the step size exponentially, it will, however, turn almost as quickly. Consequently, concession may cause very little loss in performance.

6. ALTERNATIVE PARETO-SEARCH METHODS

To fully appreciate the Pareto-search methods developed in Section 5, it is insightful to consider other techniques as candidate Pareto-search methods. In spirit, the work of Faratin et al.² is probably most closely related to the methods developed in Section 5. Their algorithm is, however, specifically designed to handle (and exploit) discrete issues and moreover is limited to a linear additive utility function (see also Section 2). Our methods are designed to handle (and exploit) continuous issues and nonlinear/general utility functions. Therefore, a direct comparison is not possible. Other techniques developed in the multi-issue negotiation literature assume a mediator (see also Section 2); therefore they also do not qualify as candidate Pareto-search methods.^c Due to a lack of comparable techniques in the multi-issue negotiation literature, we will focus instead on more standard numerical techniques in this section.

To understand which type of standard techniques are most appropriate, we explain in Section 6.1 the strong links between the Pareto-search problem and finding a minimum of a one-dimensional function (without using derivatives). Subsequently, in Section 6.2 we suggest the downhill simplex method as a potentially good alternative to especially the orthogonal-DF Pareto-search method. In the experiments reported in Section 7.3, we actually use the downhill simplex method as a benchmark. In Section 6.3, we briefly explain why we do not consider candidate Pareto-search methods that explicitly take into account the effect concessions have on the location of Pareto-efficient offers.

6.1. Minimizing a Convex Function

In computer experiments reported in Section 7.3, both seller and customer have convex preferences. In the absence of concessions, finding Pareto-efficient solutions can, consequently, be interpreted as finding an (absolute) minimum of a convex one-dimensional function f . Moreover, the effect of one negotiator—say the customer—using the orthogonal strategy is that the seller gets consistent feedback. This feedback can be interpreted as evaluations of the function f at various points in its domain. Figures 6 and 7 illustrate this correspondence between finding a Pareto-efficient solution and minimizing a convex function f : Figure 6 gives the bargaining setup and Figure 7 the corresponding minimization problem.

In Figure 6, U_s and U_c denote the seller's and customer's iso-utility lines. The numbers 1 and 2 indicate the seller's offer and the customer's corresponding counteroffer. (The arrow highlights the sequence of events.) Figure 7 depicts the function f . We assume (without loss of generality) that the seller determines p_v , the variable price first; the corresponding fixed price follows directly from the iso-utility

^cWe may be able to translate the idea in Ref. 9 of using annealers to our non-single-negotiation text setting and subsequently make an experimental comparison. However, their approach is specifically designed for contracts with many issues (like, e.g., 100 or more issues) as opposed to our approach, where we aim for reaching high-quality outcomes in only a few rounds, for a limited number of issues.




Figure 6. Depiction of the Pareto-search problem: The customer uses the orthogonal strategy. U_s and U_c denote the seller's and customer's aspiration levels. The numbers 1 and 2 indicate the seller's offer and the customer's corresponding counteroffer; dis' denotes the difference between the two offers. Moreover, (p_f^*, p_v^*) denotes the Pareto-efficient shop offer and dis^* the distance to the corresponding Pareto-efficient customer offer.




Figure 7. Depiction of the convex one-dimensional function f that corresponds to the Pareto-search problem depicted in Figure 6. For the Pareto-efficient shop offer (p_f^*, p_v^*) , we have $f(p_v) = dis^*$; dis^* is the distance to the corresponding Pareto-efficient customer offer. Moreover, $f(p_v') = dis'$, which is the distance between offers 1 and 2 in Figure 6.

line. For any pair of offers on the seller's iso-utility line, f specifies which offer represents a Pareto improvement; that is, (p'_f, p'_v) represent a Pareto improvement over (p_f, p_v) if and only if $f(p'_v) \leq f(p_v)$.

6.2. Standard Minimization Techniques

Thus, in the absence of concessions, the seller can find Pareto-efficient deals by minimizing the corresponding convex function f . For this problem, we can use standard numerical techniques for minimization without derivatives (combined with the orthogonal method to transform the problem to minimizing f). These techniques require function evaluations only (as oppose to information about the derivative of f) because negotiators do not know f ; all they can get are point evaluations of f , by considering the distance between a seller's offer and customer's corresponding offer (see Section 6.1). Brent's method is such a standard technique. This technique is an obvious choice because of its very good convergence properties (cf. Refs. 22 and 23). Roughly, the idea of Brent's method is to combines golden section search with parabolic interpolation.

Conducted computer experiments confirm that Brent's method works very well, as long as bargainers do *not* make any concession (i.e., the iso-utility line remains fixed throughout the bargaining process). (However, the Pareto-search methods developed in Section 5 perform just as well.) In the experiments, the customer and seller use the orthogonal strategy and Brent's method, respectively. As soon as bargainers start making concessions, the performance of the orthogonal strategy combined with Brent's method deteriorates. An essential aspect of Brent's method is the bracketing of the optimum: Through successive steps it tries to decrease the interval in which the optimum is know to lie. This bracketing of the optimum is hard to maintain when bargainers make concessions. In that case, the curvature and optimum of f may change through time. Consequently, the optimum may no longer lie in the bracket interval.

Thus Brent's method is ill suited to serve as an alternative for the Pareto-search method developed in Section 5. The downhill simplex method (DSM) is another well known and widely used optimization technique. (Note that DSM is unrelated to the simplex method of linear programming.) It is a multidimensional optimization method. Like Brent's method, it does only require function evaluations, not derivatives. Unlike Brent's method, it does not require the bracketing of the minimum. This makes the DSM better suited to handle nonstationarity. Given an n -dimensional problem space, the DSM must be started with $n + 1$ points, defining an initial simplex. For these points, it gets the corresponding function values. By manipulating the shape of the simplex, the algorithm searches for the point that results in the lowest value of f . For further details on the DSM see Ref. 23. In Section 7.3 we report on the experiments conducted with the DSM. In these experiments one bargainer (typically the customer) uses orthogonal search. On the other hand, the seller uses the orthogonal-DSM, that is, orthogonal search to obtain function evaluations of f and the DSM to find a minimum of f (i.e., a Pareto-efficient offer).

6.3. Opponent Modeling

Note that all the Pareto-search methods developed in Section 5 do not explicitly take into account the effect concessions (especially made by the opponent) can have on the location of Pareto-efficient offers. To take into account these effects requires elaborate opponent modeling, and updating of these models through trial and error. Given uncertainty about the opponent's preferences, opponent modeling can only be fruitful if seller and buyer negotiate very frequently with one another. Second, negotiators should be able to identify each other (e.g., on the web a person can often conceal his identity). Third, updating a model of the opponent is difficult because a bargainer needs to disentangle the concession and Pareto-improvement effects.

Despite (or because of) these concerns, opponent modeling is still an interesting line of research. In this article, the focus lies on a Pareto-search method that does not require opponent modeling, however. Consequently, we limit our search for appropriate alternative Pareto-search methods to techniques that also do not require opponent modeling.

7. EXPERIMENTAL SETUP AND RESULTS

As discussed in Section 5, negotiation essentially consists of two strategic aspects: the concession of the agents and the Pareto-search method. In this section we focus on the latter aspect of the negotiations. By means of computer experiments, we investigate the effectiveness and robustness of the orthogonal and orthogonal-DF approaches to find Pareto-efficient solutions for a wide variety of settings. We evaluate the robustness of the search strategy by experimenting with various concession strategies on the customer side.

Section 4.2 provides a general specification of the customer agent and the seller agent. Sections 7.1 and 7.2 describe the agent settings that are specifically used within the experimental setup. In particular the agents' preferences and concession strategies are specified in detail in Sections 7.1 and 7.2, respectively. The experimental results are discussed in Section 7.3.

7.1. Agent Preference Settings

We simulate the negotiation with a variety of customer and seller preferences, expressed by the agents' utility functions. The *customer agent's* expected utility depends on ρ_c , the total number of articles the customer expects to purchase (see also Section 4.2.2). The value ρ_c is assumed to be a constant, set randomly between 1 and 20 at the beginning of the experiment. Note that this results in a linear iso-utility curve in the (p_f, p_v) plane (see, e.g., Figure 4). Furthermore, because the purpose is to demonstrate the efficiency of the final deals reached, we set the customer agent's mandate b_{max} for the bundle such that an agreement is always reached.

The expected utility (i.e., expected revenue) for the *seller agent* is based on ρ_s , the expected number of articles that the customers will purchase *on average*.

In contrast to the customer agent, the expectation is not a constant but a function of the variable price p_v . It is assumed that customers who are willing to pay a high variable price are expected to read less than customers with a low variable price (i.e., we assume the law of demand holds; see Ref. 17). In the experiments we use the linear function $\rho_s(p_v) = b - a \cdot p_v$ with $b = 20$ and a set randomly between 0.03 and 0.07 at the beginning of an experiment. Note that the seller's iso-utility curve is now convex (toward the origin).

7.2. Concession Strategies

The customers and the seller can each select their own concession strategies. Although a seller's concession in the main system can depend on the behavior of all customers (i.e., one to many), in the experiments the seller agent's strategy is simply to decrease the desired utility level or *aspiration level* with a fixed amount each round. The initial aspiration level is randomly varied. Note that the number of customers and their behavior does not affect the seller's concession when this strategy is used.

On the customer side, on the other hand, we implemented four classes of concession strategies to investigate the robustness of the Pareto-search method:

- (1) *Hardhead*. The customer agent does not concede when this strategy is used; the aspiration level remains the same during the negotiations.
- (2) *Fixed*. A fixed amount c in utils is conceded each round.
- (3) *Fraction*. The customer concedes the fraction γ of the difference between the current desired utility level and the utility of the opponent's last offer.
- (4) *Tit-for-tat*. This strategy mimics the concession behavior of the opponent, based on subjective utility improvement. If the utility level of the seller's offers increases, the same amount is conceded by the customer. Note that it is the increment in the utility level perceived by the customer. The seller's actual concession is shielded from the customer agent, as an improvement could also occur when the seller moves along his iso-utility curve. Furthermore, note that the perceived utility improvement could also be negative. To make the concession behavior less chaotic, no negative concessions are made by the customer.

7.3. Results

In the experiments, the seller and customer negotiate in an alternating fashion until an agreement is reached. The efficiency of the agreement is then evaluated based on the distance of the final offer from a Pareto-efficient solution. We measure an offer's distance from a Pareto-efficient solution as the maximum possible utility improvement for the customer if a Pareto-efficient offer was made, all else remaining equal. To evaluate the quality of the results we compare the outcomes using various search strategies and concession strategies of the customer. Tables I and II provide an overview of the results. The row labeled "Random" in Table II contains the outcomes when both seller and customer use a random search strategy. This strategy selects a random point on the iso-utility curve.^d The distance of

^dOnly the downward sloping part of the seller's iso-utility curve is used.

Table I. Average distance from the Pareto-efficient solution for various customer concession strategies (rows) and customer/seller search strategies (columns). Results are averaged over 5000 experiments with random parameter settings. Standard deviations are indicated in parentheses. Best results (see column Orthogonal/DF) are obtained if the customer and seller use orthogonal search and the seller’s search is amplified with a derivative follower.

Concession strategy	Orthogonal/DF search	Random search
Hardhead	3.64 (± 8.75)	18.81 (± 26.86)
Fixed ($c = 20$)	14.15 (± 19.91)	27.61 (± 37.10)
Fixed ($c = 40$)	22.77 (± 28.71)	34.40 (± 45.38)
Fixed ($c = 80$)	35.22 (± 45.42)	44.53 (± 59.61)
Fraction ($\gamma = 0.025$)	7.27 (± 14.98)	25.57 (± 34.89)
Fraction ($\gamma = 0.05$)	8.33 (± 15.64)	29.99 (± 41.11)
Fraction ($\gamma = 0.1$)	17.80 (± 44.28)	36.00 (± 50.09)
Tit-for-tat	20.73 (± 26.59)	35.91 (± 47.59)

the final offer, when random search is used, lies between 3 and 6% of the maximum Pareto inefficiency (that is possible when keeping the seller’s iso-utility fixed).

Although the inefficiency with random search is only small compared to the maximum inefficiency, even better results are obtained when one bargainer (typically the customer) uses orthogonal search and the other (the seller) uses orthogonal-DF search (i.e., orthogonal search combined with a derivative follower). The results are shown in the column labeled “Orthogonal/DF” of Table I. The improvements are considerable. The performance is up to more than six times better than using the random cooperative strategy. The distance of the final offer as a percentage of maximum Pareto inefficiency lies now between 1 and 4%.

Table II shows the results if both customer and seller use orthogonal-DF search (column *DF/DF*). The distance of the final offer lies between 2 and 9% of the maximum Pareto inefficiency. These results are slightly worse than random. The derivative follower relies on a consistent response from the opponent to signal the right direction. If both use a derivative follower, this signal is distorted.

Table II. Average distance from the Pareto-efficient solution for various customer concession strategies (rows) and customer/seller search strategies (columns). Results are averaged over 5000 experiments with random parameter settings. Standard deviations are indicated in parentheses.

Concession strategy	DF/DF search	Orthogonal/DSM search
Hardhead	19.98 (± 23.92)	31.01 (± 45.11)
Fixed ($c = 20$)	48.04 (± 92.28)	41.33 (± 56.15)
Fixed ($c = 40$)	70.25 (± 144.92)	49.64 (± 66.94)
Fixed ($c = 80$)	72.30 (± 137.69)	57.79 (± 78.72)
Fraction ($\gamma = 0.025$)	20.91 (± 37.15)	38.94 (± 53.92)
Fraction ($\gamma = 0.05$)	21.67 (± 41.52)	44.87 (± 66.22)
Fraction ($\gamma = 0.1$)	36.90 (± 65.16)	50.37 (± 75.46)
Tit-for-tat	34.52 (± 50.04)	44.17 (± 55.48)

Additionally, the column labeled “Orthogonal/DSM” of Table II shows the results if one bargainer (typically the customer) uses orthogonal search and the other (the seller) uses the orthogonal-DSM (i.e., orthogonal search combined with the downhill simplex method; see Section 6). The distance of the final offer lies between 5 and 9% of the maximum Pareto inefficiency. These results are slightly worse than the DF/DF and much worse than the results of the Orthogonal/DF reported in Table I.

The orthogonal-DSM search performs much worse than the orthogonal-DF search (both are combined with the orthogonal search). The simplicity of the DF with adaptive step size, which underlies the orthogonal-DF search, is also its strength: It stores much less information (storage requirement are kept to a minimum) and updates this information more frequently than the DSM, which underlies the orthogonal-DSM search. Consequently, the DF (with adaptive step size) handles nonstationarity much better than the DSM.

The standard deviations reported in Tables I and II are relatively large (this is probably due to the fact that concessions are being made). However, this does not necessarily weaken our conclusions because the standard deviations are high for all the results reported. Indeed, using the *t*-test for population average confirms that, with a 0.9995 confidence, we can accept the hypothesis that the average distance for the orthogonal/DF search is smaller than the corresponding average distance for the random, DF/DF, and orthogonal/DSM searches.^c

Notice that the average distance depends on the concession strategy used by the customer. Although in individual cases Pareto-efficient agreements (with zero distance) are reached using the orthogonal/DF search, the average distance consistently shows some (usually slight) inefficiencies, even when the customer makes no concessions (i.e., the *hardhead* strategy). The reason for this is twofold. First, the DF accelerates finding the efficient solution by making, at times, large steps on the iso-utility curve. At a certain point the algorithm passes the Pareto-efficient point, and then turns. This way the offers keep oscillating around the optimal point. If the concessions are sufficiently large, an agreement can be reached at a point that is less than optimal.

Second, the direction and step size of the DF are based on changes in the Euclidean distance between the seller and customer offers through time. The distance can be influenced by both concessions and movements along the iso-utility curve. As the opponent’s iso-utility curve is unknown, the agents are unable to distinguish between the two. This can mislead the DF whenever concessions are very large. Two possible solutions are to make either small concessions or have intervals with no concessions, allowing the search algorithm to find the best deal.

To conclude, the orthogonal/DF strategy clearly outperforms other combinations of search strategies in the experiments. Inefficiencies still occur, especially if

^cLet m_a , s_a^2 , and n_a denote the sample mean, standard deviations, and number of elements of sample *a*; moreover, m_b , s_b^2 , and n_b are defined similarly for sample *b*. Then the test value *t* is defined as $t = (m_a - m_b) / ((s_a^2/n_a) + s_b^2/n_b)^{1/2}$. (Note that degrees of freedom are not really an issue with $n_a = n_b = 5000$.)

the concessions are large. A trade-off therefore exists between reaching an agreement fast (by making large concessions) and reaching an efficient agreement.

8. CONCLUSIONS AND FUTURE WORK

In particular, so-called information goods mark an application area where bargaining may be beneficial. Information goods are often sold through on-line services where customers pay a *nonnegotiable* subscription fee. Especially, with a high heterogeneity among customers and a lot of dynamics in customer preferences, offering standardized nonnegotiable subscriptions may not be the best solution. Negotiating has, conversely, the advantage of providing more flexibility and adaptability toward the preferences of individual customers and daily dynamics. In this article, we therefore consider the problem of a seller and customer agent negotiating about the subscription fee of such an on-line service.

We point out in the article that for most practical applications, it suffices to concentrate on negotiating about a two-part tariff. Consequently, the focus in the remainder of the article lies on negotiating about a two-part tariff. Autonomous software agents execute the negotiation on behalf of the bargainers. To perform the actual negotiation these agents make use of bargaining strategies. We present a decomposition of bargaining strategies into concession strategies and Pareto-efficient search methods. An important technical contribution of the article lies in the development of two Pareto-search methods: the orthogonal and orthogonal-DF methods. Computer experiments show, for various concession strategies, that the respective use of these two Pareto-search methods by the two bargainers results in very efficient bargaining outcomes, while negotiators concede the amount specified by their concession strategy.

An interesting extension of our work will be to develop the orthogonal-DF for multi-issue negotiation with more than two issues. Moreover, it will be interesting to investigate the performance of such a search method experimentally.

Acknowledgments

This article is an extended and significantly modified version of Ref. 24.

References

1. Lomuscio AR, Wooldridge M, Jennings NR. A classification scheme for negotiation in electronic commerce. *Group Decis Negot* 2003;12:31–56.
2. Faratin P, Sierra C, Jennings NR. Using similarity criteria to make issue trade-offs. *J Artif Intell* 2003;142:205–237.
3. Ehtamo H, Verkama M, Hämäläinen RP. On contracting under incomplete information using linear proposals. In: *Proc Fifth Int Symp on Dynamic Games and Application*, Grimentz, Switzerland; 1992. pp 128–133.
4. Ehtamo H, Verkama M, Hämäläinen RP. How to select fair improving directions in a negotiation model over continuous issues. *IEEE Trans Syst Man Cybern C* 1999;29:26–33.
5. Ehtamo H, Kettunen E, Hämäläinen RP. Searching for joint gains of in multi-party negotiations. *Eur J Oper Res* 2001;130:54–69.

6. Ehtamo H, Hämäläinen RP. Interactive multiple-criteria methods for reaching Pareto optimal agreements in negotiations. *Group Decis Negot* 2001;10:475–491.
7. Lin RJ, Chou ST. Mediating a bilateral multi-issue negotiation. In: *Proc IEEE Int Conf on E-Commerce (CEC'03)*, Newport Beach, CA; 2003. pp 76–84.
8. Lin RJ, Chou ST. Bilateral multi-issue negotiations in a dynamic environment. In: Faratin P, Parkes D, Rodriguez J, Walsh W, editors. *Proc Agent-Mediated Electronic Commerce V (AMEC'03)*, Melbourne, Australia; 2003. pp 117–123.
9. Klein M, Faratin P, Sayama H, Bar-Yam Y. Negotiating complex contracts. *Group Decis Negot* 2003;12:111–125.
10. Fisher R. *International mediation: A working guide*. New York: International Peace Academy; 1978.
11. Luo X, Jennings NR, Shadbolt N. Acquiring trade-off preferences for automated negotiations: A case study. In: *Proc Agent-Mediated Electronic Commerce V (AMEC'03)*, Melbourne, Australia; 2003. pp 37–55.
12. Wilson R. *Nonlinear pricing*. Oxford, UK: Oxford University Press; 1993.
13. Courcoubertis C, Weber R. *Pricing communication networks: Economic, technology and modeling*. West Sussex, England: Wiley; 2003.
14. Brooks CH, Fay S, Das R, MacKie-Mason JK, Kephart JO, Durfee EH. Automated strategy searches in an electronic goods market: Learning complex price schedules. In: *Proc ACM Conference on Electronic Commerce*. New York: ACM Press; 1999. pp 31–41.
15. Kephart JO, Brooks CH, Das R. Pricing information bundles in a dynamic environment. In: *Proc 3rd ACM Conf on Electronic Commerce*. New York: ACM Press; 2001. pp 180–190.
16. Rubinstein A. Perfect equilibrium in a bargaining model. *Econometrica* 1982;50:97–109.
17. Mas-Collé A, Whinston MD, Green JR. *Microeconomic theory*. Oxford, UK: Oxford University Press; 1995.
18. Webster R. *Convexity*. Oxford, UK: Oxford Science Publications; 1994.
19. Kephart JO, Hanson JE, Greenwald AR. Dynamic pricing by software agents. *Comput Netw* 2000;36:731–752.
20. van Bragt DDB, Somefun DJA, Kutschinski E, La Poutré JA. An algorithm for on-line price discrimination. Technical report SEN-R0213, Centrum voor Wiskunde en Informatica; 2002.
21. Dasgupta P, Das S. Dynamic pricing with limited competitor information in a multi-agent economy. In: Eztion O, Scheuermann P, editors. *Proc 7th Int Conf on Cooperative Information Systems*, Eilat, Israel. *Lecture Notes in Computer Science*, vol. 1906. Berlin: Springer; 2000. pp 291–310.
22. Brent RP. *Algorithms for minimization without derivatives*. Englewood Cliffs, NJ: Prentice-Hall; 1973.
23. Press WH. *Numerical recipes*. Cambridge, UK: Cambridge University Press; 1992.
24. Somefun DJA, Gerding EH, Bohte S, La Poutré JA. Automated negotiation and bundling of information goods. In: Faratin P, Parkes DC, Rodríguez-Aguilar JA, editors. *Agent-mediated electronic commerce V: Designing mechanisms and systems. AAMAS 2003 Workshop, AMEC 2003*, Melbourne, Australia, July 15, Revised Selected Papers. *Lecture Notes in Computer Science*, vol 3048. Heidelberg: Springer-Verlag; 2004. pp 1–17.