

Qualitative One-to-Many Multi-Issue Negotiation: Approximating the QVA

Koen V. Hindriks · Dmytro Tykhonov ·
Mathijs M. de Weerd

Published online: 9 January 2010
© Springer Science+Business Media B.V. 2010

Abstract When there is one buyer interested in obtaining a service from one of a set of sellers, multi-attribute or multi-issue auctions can ensure an allocation that is efficient. Even when there is no transferable utility (e.g., money), a recent qualitative version of the Vickrey auction may be used, the QVA, to obtain a Pareto-efficient outcome where the best seller wins. However, auctions generally require that the preferences of at least one party participating in the auction are publicly known, while often making this information public is costly, undesirable, or even impossible. It would therefore be useful to have a method that does not impose such a requirement, but is still able to approximate the outcome of such an auction. The main question addressed here is whether the Pareto-efficient best-seller outcome in multi-issue settings without transferable utility (such as determined by the QVA) can be reasonably approximated by multi-bilateral closed negotiation between a buyer and multiple sellers. In these closed negotiations parties do not reveal their preferences explicitly, but make alternating offers. The main idea is to have multiple rounds of such negotiations. We study three different variants of such a protocol: one that restricts the set of allowed offers for both the buyer and the seller, one where the winning offer is announced after every round, and one where the sellers are only told whether they have won or not after every round. It is shown experimentally that this protocol enables agents that can learn preferences to obtain agreements that approximate the Pareto-efficient best-seller outcome

K. V. Hindriks · D. Tykhonov · M. M. de Weerd (✉)
Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology,
P.O. Box 5031, 2600 GA Delft, The Netherlands
e-mail: m.m.deweerd@tudelft.nl

K. V. Hindriks
e-mail: k.v.hindriks@tudelft.nl

D. Tykhonov
e-mail: d.tykhonov@tudelft.nl

as defined by the auction mechanism. We also show that the strategy that exploits such a learning capability in negotiation is robust against and dominates a Zero Intelligence strategy. It thus follows that the requirement to publicly announce preferences can be removed when negotiating parties are equipped with the proper learning capabilities and negotiate using the proposed multi-round multi-bilateral negotiation protocol.

Keywords Qualitative auction · Multi-bilateral negotiation · Bayesian learning · Approximation · Procurement · Multi-attribute auction · Simulations

1 Introduction

In a procurement setting in which a buyer faces several sellers an auction may provide an effective mechanism to reach an agreement. Auctions may also be used when the outcome that needs to be reached is complex and consists of multiple issues that need to be settled, such as when a Request For Quote (RFQ) is issued by a corporation or government organization (Teich et al. 2004, 2006). A bid in such a reverse auction may involve, for example, the desired quality of service, the quantity demanded, the terms and time of delivery, and so forth. Such a setting with one buyer and multiple sellers (i.e., a reverse auction) is used throughout this paper, but all results directly transfer to a forward auction with one seller and multiple buyers as well.

The various types of auctions have desirable theoretical properties such as yielding an efficient outcome and being strategy-proof. However, some of these mechanisms impose requirements which are not easy to meet in practice. One of these requirements generally associated with (reverse) auctions is that the preferences of the buyer have to be known by all bidders. This requirement is often not realistic in practice. First of all, the explicit elicitation of a buyer's value function may be difficult (Bichler et al. 2001). Even modeling such preferences is a very complex problem, and very relevant in the context of auctions (Teich et al. 2004). The buyer may not know the complete domain of possible outcomes as sellers may come up with new options during the process, and it usually is very hard to specify preferences completely over a complex and possibly infinite set of outcomes. This is particularly true for auctions that are used to settle multiple issues, e.g., related to an RFQ. Finally, the buyer may not want to publicly reveal his preferences to the extent required by multi-issue auctions. It may be disadvantageous to do so given that it is not unlikely that future encounters with similar parties will take place.

In negotiations, on the other hand, the preferences of one party are only partially revealed to the other in the course of the process. However, in settings where one buyer may make a choice among a set of sellers, the competitive aspect is not explicitly taken into account in negotiation protocols. In recent work on multiple (parallel) bilateral negotiations (Nguyen and Jennings 2003; Rahwan et al. 2002; Li et al. 2004), this is partly taken care of by informing all other sellers when a provisional agreement is reached in one of the negotiation sessions. However, when other sellers improve upon this outcome, the seller who has reached the first agreement does not get any opportunity anymore to improve upon this. Such negotiation protocols consequently may often lead to inefficient agreements.

It thus becomes interesting to look for alternative methods that may be used that guarantee outcomes that approximate the efficient outcome of an auction mechanism. The problem we study in this paper is whether alternative mechanisms based on multiple bilateral (also called multi-bilateral) negotiations can be used to reduce the preference information that needs to be made public but that also retain some of the desired theoretical properties such as efficiency of agreements as guaranteed by the auction mechanism. Studying mechanisms based on multilateral negotiations is interesting in its own right (Thomas and Wilson 2005), but also because their relationship to various auction formats has implications for institutional design. Studying the factors that relate and differentiate auctions from negotiation mechanisms may lead to a more informed selection of a transaction mechanism.

In this paper, we consider a particularly interesting instance of this more general one-to-many multi-issue allocation problem where there is not necessarily the possibility to transfer utility (such as when the government announces a fixed budget in a RFQ).¹ Without transferable utility, most of the known auction mechanisms cannot be used. However, one particular auction mechanism, called the *Qualitative Vickrey Auction* (QVA) (Harrenstein et al. 2009), can obtain a Pareto-efficient outcome where the best seller wins. However, the QVA requires the buyer to publicly announce its preferences.

We study various multi-bilateral negotiation mechanisms. The main idea is that the (efficient) outcome of the QVA may be approximated by a negotiation protocol that consists of multiple negotiation rounds in which sellers are provided an opportunity to outbid the winner of the previous round. We show experimentally that each of these mechanisms is able to approximate the efficient outcome as defined by the QVA. The main assumption that we need to make to obtain this result is that the negotiating agents are able to (privately) learn part of the preferences of their opponents during a negotiation session. Techniques to do so are available (Hindriks and Tykhonov 2008a), making our proposal one that can be implemented given the current state of the art in negotiation. Additionally, experiments are performed that show that a negotiating agent that exploits learning outperforms a Zero Intelligence strategy (Gode and Sunder 1993).

The paper is organised as follows. In Sect. 2 we define the general setting of a buyer and multiple sellers that aim to reach an agreement settling multiple issues. This setting is generic in the sense that it covers arbitrary situations where one buyer wants to obtain an efficient multi-issue agreement with any one out of a set of available sellers. Section 3 introduces the QVA auction that may be used to reach such an agreement. In Sect. 4 we then propose three variants of a multi-bilateral negotiation protocol as alternative mechanisms to the QVA. Each of these protocols is related to the QVA in the sense that it approximates the outcome defined by the QVA. The different protocols introduced moreover progressively require less information to be revealed publicly by the buyer. Section 5 presents experimental results to evaluate how well these protocols approximate the outcome defined by the QVA mechanism. The results validate our claim that the QVA may be replaced by a multi-bilateral negotiation protocol while still obtaining agreements that are similar. The tradeoff that has to be made concerns

¹ This also means that 'pricing out' is not an option to elicitate preferences (Teich et al. 2004).

the amount of effort and time that needs to be invested in reaching an agreement. Finally, Sect. 6 discusses related work and Sect. 7 concludes with a discussion of the results obtained and outlines directions for future research.

2 Definitions

The setting we consider in our work consists of a buyer that wants to procure a service or product from one out of a potentially large number of sellers. An agreement in this setting is an outcome that fixes the parameters of the service to be provided. Formally, the space of all possible outcomes is defined as all tuples $x = \langle x_1, \dots, x_m \rangle \in X$ over m issues in a domain $X = X_1 \times \dots \times X_m$. These issues define all aspects of the agreement, such as quality, start time, duration, guarantees, penalty, etc. Buyer and sellers are assumed to associate a *utility value* with each outcome and to have a *reservation value* that determines when an outcome does not improve the status quo for that party, i.e. the buyer or one of the sellers. In this paper we concentrate on a setting where we do not assume that one of these issues is a price. In other words, we relax the condition that there is a transferable utility (and we thus also relax the usual assumption that utility functions are quasi-linear in price).

We introduce the following notation. The buyer is denoted by 0 and sellers are denoted by $i \in \{1, \dots, n\}$. The *reservation value* of each party i is denoted by v_i and represents the minimal utility value that an agreement should have to be an acceptable outcome for that party. Outcomes with a utility below the reservation value are called *unacceptable*. Each party i also has a utility function $u_i : X \rightarrow \mathbb{R}$ which represents the utility that party associates with an outcome.

The goal is to find an agreement between the buyer and the sellers that is not only acceptable to both, but that is also Pareto efficient, i.e., there should not be another agreement with the same or higher utility for both players, and strictly higher for at least one of them. In addition to Pareto-efficiency of the final agreement between the buyer and the winning seller, we are interested in an agreement with the seller that can make the best offer to the buyer (that is still acceptable to the seller), often called allocative efficiency in the context of auctions. In this paper we call an outcome that meets both these efficiency conditions simply *efficient*.

An example is a buyer that is interested in buying a supercomputer. A range of potential suppliers is available that may provide a supercomputer. Apart from price (which is often limited by a given budget), supercomputers have many features (processing speed, memory, etc.) and requirements (regarding power supply, cooling, etc.) that need to be settled to obtain an agreement. Such an agreement thus is complex as many issues have to be agreed upon and finding an efficient outcome can be a complex process. In the next section an auction mechanism is summarized that has a dominant strategy equilibrium that yields a such an efficient outcome.

3 The Qualitative Vickrey Auction

The Qualitative Vickrey Auction (QVA) (Harrenstein et al. 2008, 2009) is particularly useful in a context where a single buyer tries to obtain a *complex agreement* with one

out of many sellers that are interested in making such an agreement. This generalization of a Vickrey auction (Vickrey 1961) is strategy-proof, and under most realistic settings (when money is involved, or the set of outcomes is discrete and linearly ordered by the buyer, or the preferences of the suppliers are equipeaked), it is efficient, i.e., it obtains a Pareto-efficient outcome that involves the seller that can make the best agreement (for the buyer) still acceptable to him.

Intuitively, this mechanism captures the negotiation power of the buyer. If there are many sellers, the buyer will end up with some very good offers, but if there is only one seller that has a sufficiently good offer, the agreement is not that good for the buyer. This interpretation can be given to most auction mechanisms. This mechanism, summarized below, has the special feature that it also works if none of the issues is about money.²

This auction mechanism can be thought of as consisting of two rounds. In the first round (1.1–1.3 below), the buyer publicly announces her preferences. Then potential sellers submit offers in response, and a winner is selected by the buyer. In the second round (2.1–2.2 below), the buyer determines the second-best offer (from her perspective again) she received from another seller, and announces this publicly. Finally, the winner is allowed to select any agreement that has at least the same utility to the buyer as the second-best offer (which can be determined by the winner since the preferences of the buyer are publicly announced). If the bids offered in the first round all are made public afterwards, anyone can check whether the buyer follows the protocol. The steps of this procedure can be found in Algorithm 1.

Algorithm 1 The qualitative Vickrey auction

1. Round 1: Winner selection
 - 1.1 The buyer announces her preferences.
 - 1.2 Every seller submits an offer.
 - 1.3 The buyer selects the winner according to her preferences.
 2. Round 2: Agreement selection
 - 2.1 The buyer announces the second-best offer she received.
 - 2.2 The winner may select any agreement that has at least the same utility for the buyer as the second-best offer.
-

The properties that make this mechanism interesting are not only Pareto efficiency, and that the seller wins that can make the best offer, but also that it is a dominant strategy for a seller to bid an offer that is just acceptable to itself and ranks highest in the buyer's preferences. In the problem domain defined in the previous section, this dominant strategy comes down to proposing an offer with exactly the same utility as its reservation value. Formally, the winner in a given problem domain X then can be defined by³:

$$i^* = \arg \max_{i \in \{1, \dots, n\}} \max \{u_0(x) \mid x \in X, u_i(x) \geq v_i\},$$

² If none of the issues is about money, a reverse auction is not different from a standard auction.

³ We assume ties are broken by the buyer using a given ordering over the sellers.

where v_i denotes the reservation value of seller i .

To determine the outcome, we also need the second-best offer. Assuming all sellers follow the dominant strategy, the second-best offer \hat{x} is given by

$$\hat{x} = \arg \max_{x \in \{x | u_i(x) \geq v_i, i \in \{1, \dots, n\} \setminus \{i^*\}\}} u_0(x).$$

The outcome then is the best possible for the winner i^* , given that it is at least as good for the buyer as the second-best offer \hat{x} , i.e.,

$$\omega = \arg \max_{x \in \{x | u_0(x) \geq u_0(\hat{x})\}} u_{i^*}(x).$$

Intuitively, this auction-like mechanism selects the best seller, because the sellers have a dominant strategy to submit their best offer, and the best of these is chosen by the mechanism in the first step. Moreover, the outcome selected is Pareto-efficient, because in the last step the winner maximizes its utility given a constraint on the utility for the buyer (and full knowledge of both preferences). The exact conditions and the proofs for efficiency and the dominant strategy equilibrium can be found in [Harrenstein et al. \(2009\)](#).

The main problem with a realistic implementation of the QVA is that the buyer needs to communicate all her preferences to all sellers. This is impractical for various reasons. Firstly, in many settings it is undesirable for the buyer to communicate all her preferences to all sellers, because the buyer may not want to disclose all details for strategic reasons. Secondly, this preference function can be a quite complex function over a large domain, which is difficult to communicate efficiently. Finally, a buyer may not even know the complete domain of agreements on forehand, even though she is able to rank any given subset of agreements. The latter holds for example when a government sends out a request for proposals to construct a bridge over a river within a given budget. It is impossible to list all possible types of bridges designers may come up with. But also in domains such as the super-computer domain, sellers usually come up with new options and alternatives in a negotiation process. If only the limited domain known by a buyer is used, the resulting outcomes will generally not be efficient. Therefore, in the complex multi-issue domains we consider in this paper, a standard ascending/descending auction, or the qualitative Vickrey auction discussed above cannot be used, because in such an auction the sellers require complete knowledge of the preferences of the buyer. In the next section we describe an approach based on negotiation that may be used to approximate the efficient outcome of such an auction and where there is no need to publicly announce the preferences of the buyer.

4 Negotiation Protocols

In a QVA, all sellers propose an offer to the buyer. The buyer then determines which of the sellers has the winning offer. That seller is then allowed to change his offer to improve his utility value while taking into account that the buyer's utility value may not be lower than the second-best offer. In this setting, the dominant strategy for sellers in the first step is to propose an outcome that has a utility value equal to

their reservation value with a maximal utility for the buyer. We could see this as an indication of the negotiation power of the buyer in a QVA. It means that sellers need to be aware that they are one out of potentially many other sellers that the buyer may reach an agreement with.

This negotiation power of the buyer explains why a QVA cannot simply be replaced by multiple bilateral negotiations based on e.g. an alternating offers protocol between the buyer and each of the sellers as this would not take into account that multiple sellers are contending for an agreement with the buyer. In [Hindriks et al. \(2008\)](#) it was shown that a negotiation using the alternating offers protocol without any additional assumptions except for the fact that agents were able to learn opponent preferences does not result in a good approximation of the efficient outcome of the QVA.

In order to relax the constraint of the QVA that a buyer has to publicly announce its preferences, we propose three different negotiation protocols that take the negotiation power of the buyer into account. The first protocol we study tries to stay as close as possible to the QVA and imposes quite strict constraints on the moves of the negotiating parties. In fact, this protocol may be viewed as a variant of the QVA that does not require disclosing the preferences of the buyer. That is, this protocol consists of two negotiation rounds where in the first round sellers are constrained and required to propose offers that have a utility equal to their reservation value to the buyer and in the second negotiation round the buyer is constrained and required to propose offers that have a utility that is equal to that of the second-best outcome of the first round. Although this protocol is an improvement over the QVA in the sense that it does not require the public announcement of complete preferences, it still requires the negotiating parties to reveal their reservation value.

In order to remove the requirement to reveal reservation values, a second and a third protocol are studied that involve multiple negotiation rounds instead of just two rounds. The main idea is that in these future rounds sellers are provided an opportunity to outbid the winner of the previous round. The negotiation power of the buyer is represented in this protocol by the fact that negotiation continues over multiple rounds until no seller is willing to outbid the best outcome of the previous round (from the buyer's perspective). Both protocols are variants of this idea, where the second protocol requires the buyer to announce the winning bid at the end of each round and the third protocol only tells each player whether he or she is the winner at the end of each round.

4.1 A Protocol Based on two Negotiation Rounds

The first negotiation protocol consists of two rounds which closely match the QVA mechanism (see Algorithm 2 below for details). The buyer, however, is not required to announce his preferences. In the first round (step 1 of the algorithm) bilateral negotiation sessions are performed between the buyer and every potential seller. The idea is that in the first round negotiating parties try to learn a model $\tilde{u}(x)$ of each others' preferences, both in order to win the first round as well as to be able to perform well in the second round. Throughout the paper we use an alternating offers protocol ([Osborne and Rubinstein 1994](#)) in the bilateral negotiation sessions. Furthermore, we assume

that information about a negotiation session is not used in negotiation sessions with other sellers. At the end of the first round a winner (one of the sellers) is determined by the buyer. Then a second negotiation round (step 2) between the buyer and the winner is performed. Before starting this second round, however, the agreement between the *second-best* offer from one of the sellers and the buyer (from the perspective of the buyer) is revealed to all sellers. This is in particular useful for the winner who continues negotiation with the buyer. In the second round a final agreement between the winner and the buyer is established.

Algorithm 2 A protocol based on two negotiation rounds

1. Round 1: The buyer bilaterally negotiates with every seller.
 - 1.1 The buyer starts negotiation with each seller i .
 - 1.2 The seller proposes bids that have a utility equal to the seller's reservation value.
 2. Round 2: The buyer determines the winner and negotiates a final agreement with the winner.
 - 2.1 The buyer determines the seller with the highest agreement, as well as the second-best outcome.
 - 2.2 The winning seller and the buyer negotiate the final agreement. In this round the buyer is required to make offers with a utility equal to that of the second-best outcome of the previous round.
-

For this protocol to obtain efficient agreements, there are additional constraints on the negotiation strategies or procedure that the buyer and seller are required to use. In fact, the parties are required to propose offers that implement the steps of the QVA mechanism quite closely. In the first round all offers proposed by sellers are required to have a *utility equal to their reservation value* (see step 1.2.). This constraint is derived from the fact that the dominant strategy in the QVA for sellers is to propose such offers. In the second round all offers proposed by the buyer in the alternating-offers negotiation session are required to have a *utility equal to that of the second-best outcome of the first round* (see step 2.2.). This constraint is derived from the rule to compute the final outcome in the QVA. The intuition is that the buyer knows that an agreement with another seller of a certain quality can be reached. This should induce the winner to reduce the negotiation space it considers. An alternative way of putting this is that the winner of the first round is required to adjust its reservation value and increase it to the utility it associates with the second-best outcome as revealed by the buyer (if that outcome has a higher utility than its initial reservation value; otherwise, the seller would not change its reservation value).

Given that utilities are either fixed for the buyer or the seller, it is rational for the parties to try to propose the best agreement possible for the other party. In general this is the case since negotiators need to take into account that an offer needs to be reasonable for the other party in order to reach an agreement at all. In particular, this is the case for sellers in the first round, because they need to win in this round to go through to the second.

Given this setup, our hypothesis about the feasibility to approximate the mechanism outcome by means of negotiation is the following.

Hypothesis 1 The outcome determined by the mechanism can be approximated by a negotiation setup in which: (i) the buyer does not reveal her preferences, (ii) the



Fig. 1 Negotiation moves in the first protocol: on the *left-hand side* the first round is illustrated where a seller stays on its reservation value iso-curve, and on the *right-hand side* it is shown how the buyer stays on the iso-utility curve of the utility of the second-best agreement of the first round

negotiating agents can learn an opponent's preference profile, and (iii) these agents use the negotiation procedure discussed above (Algorithm 2).

Note that in the first round, the buyer is free to choose the offers he proposes. This makes it possible for the seller to learn the preferences the buyer has during this negotiation. Also note that as the seller is supposed to propose offers with a fixed utility value (equal to its reservation value) it is difficult if not impossible for the buyer to learn the preferences of sellers in this round. Figure 1 illustrates the first protocol and the constraints on the negotiation moves of the sellers in the first round and the buyer in the second round.

The major drawback of this first protocol is that there is no way to dictate or control the restrictions for bidding behavior of the sellers and the buyer. Either the buyer and the sellers have to trust each other that they comply with the negotiation protocol or a third party trusted by all agents has to be invited to control the bidding. There is, however, some incentive for the sellers, which can be derived from the similarity to the QVA where proposing an offer at the reservation value is a dominant strategy. In the next section the need for such a trusted third party is removed.

4.2 Multiple Negotiation Rounds with Multiple Sellers

To remove the restrictions on the negotiation imposed in the first protocol, we introduce a protocol (see Algorithm 3 for details) that consists of multiple rounds of (parallel) bilateral negotiations between the buyer and the sellers. After each round r (step 2.2), the buyer communicates the *winning agreement* ω_i^r (where i is index of the winning seller) of round r to the sellers that did not win (i.e. they did not reach an agreement that was best from the buyer's perspective). All of the sellers then are provided with the opportunity to improve the agreement they reached with the buyer in a next round of negotiation sessions. A seller will do so if he can make an offer that has a utility value above his reservation value v_i , which he supposes has a higher utility to the

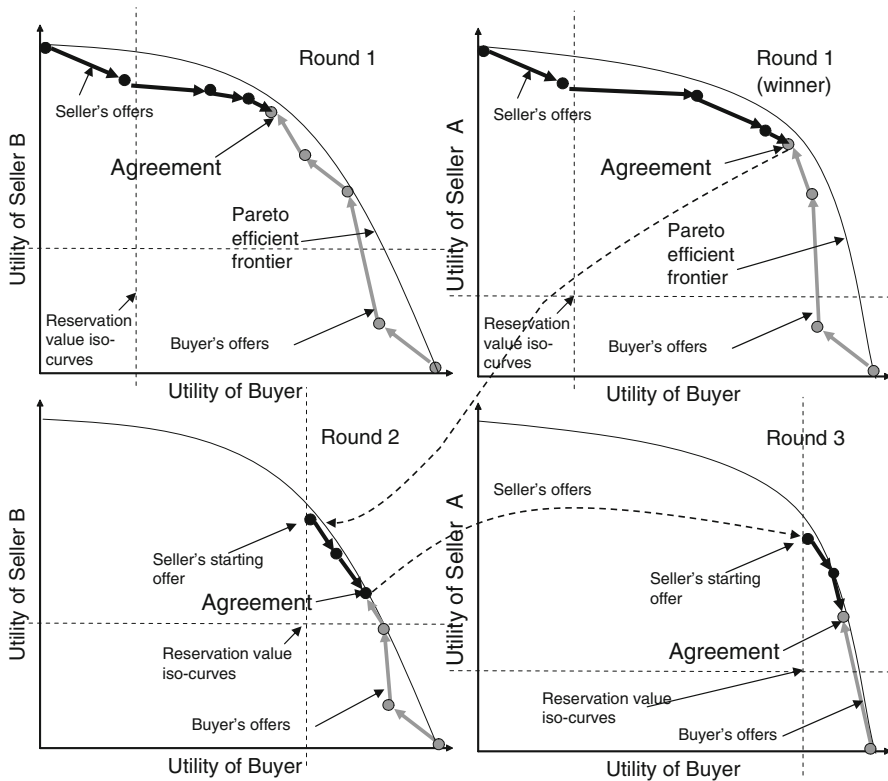


Fig. 2 In round 2 Seller B aims to improve the agreement reached between the buyer and Seller A in round 1, and then in round 3 Seller A tries to improve upon this agreement

buyer than the winning agreement of the last round. Negotiation is therefore assumed to resume for the seller in a next round starting with the agreement reached in the last round. This process continues until no seller (except for the winner) is prepared to negotiate in a next round to improve their last offer. The winning agreement of the last round then is the final agreement of the negotiation process. The details of this process are given below and are illustrated in Fig. 2.

It is advantageous for a seller to understand the buyer's preferences in this process, because this can be used to reach an agreement that satisfies the buyer as best as possible while at the same time maximizing the utility for the seller itself. In particular, such an opponent model \tilde{u} can be used to assess if an offer can be made that has the same utility value as the winning agreement from the point of view of the seller but that has a higher utility for the buyer. Only if such an offer cannot be made, an additional concession has to be made. Without the ability to learn an opponent model such an assessment cannot be made, and the seller will drop out of the negotiation process.

Figure 2 illustrates that the size of the negotiation space is decreased in every next round. This is explained by the fact that the buyer only accepts offers that improve the winning agreement reached in the previous round (see step 2.1). This process

forces the final agreement closer to that of the reservation value of the sellers, in line with the dominant strategy sellers have in the QVA. We thus formulate the following hypothesis.

Algorithm 3 A protocol with multiple negotiation rounds

1. Set the utility of ω^0 to be 0. Set the round number r to be 1.
 2. While no final agreement ω has been found, do the following.
 - 2.1 The buyer bilaterally negotiates an agreement with every seller i , accepting only offers with a utility at least as high as the utility of ω^{r-1} .
 - 2.2 If only one seller has reached an agreement, then the agreement of the previous round is the final agreement $\omega = \omega^{r-1}$.
 - 2.3 Otherwise, publicly announce the winning agreement ω^r .
 - 2.4 Start a new round, setting $r = r + 1$.
-

Hypothesis 2 The agreement reached using the proposed negotiation protocol converges to that of the efficient outcome of the QVA, assuming the negotiating parties are able to learn the preferences of their opponent.

The proposed negotiation protocol does not require the buyer to publicly announce his preferences. The protocol thus provides a realistic alternative for the QVA, that, given the hypothesis formulated above, can be used in settings where a buyer aims to reach an agreement with one out of multiple sellers. The process of reaching such an agreement is more complicated than that of the Vickrey auction but does not require publicly announcing the preferences of the buyer. Somehow the situation is reversed, however, as the protocol outlined above requires the public announcement of the winning agreement in every negotiation round (step 2.3). Instead of making the buyer's preferences public, in this case some information about the sellers' preferences is made public. We believe that this is not a prohibitive feature of the protocol as this only provides limited information to the sellers, but it still is interesting to investigate if this step in the protocol can be replaced by one that reveals even less information.

4.3 A Variant without Making Intermediate Agreements Public

A similar protocol can also be applied without informing sellers about intermediate agreements. In this case, the buyer only indicates to a seller that it did not win in the last round. The winning agreement of the previous round thus can no longer be used as a reference point that needs to be improved upon from the buyer's point of view, and a seller instead continues negotiation in the next round with the agreement it reached itself in the previous round. Moreover, in the previous protocol where a winning agreement is made public, a seller can estimate—given the opponent model it learns during a negotiation session—how much it has to concede to improve that winning agreement. This is no longer possible in this second setup. However, it is required that when the negotiation protocol terminates and a final agreement is reached that this agreement is made public in order to allow sellers to verify that the buyer has

not manipulated the process. Making only the final agreement public is sufficient for sellers that have a reasonable opponent model to assess whether the process has been fair, as there should be at least one seller that can make an offer with approximately the same utility to the buyer at his own reservation value.

Consequently, in this second variant the sellers have less information on how to out-bid the winning seller of the previous round. Still, the buyer does have this information as it knows the winning agreement of the previous round and, therefore, would only accept offers of a seller that improve the winning agreement of the previous round. Given this, we formulate the following hypothesis concerning this variant.

Hypothesis 3 The agreement reached without revealing the winning agreement in each round converges to that of the efficient outcome of the QVA, assuming the negotiating parties are able to learn the preferences of their opponent.

As the sellers have less information in this third setup, they will have more difficulty in proposing offers that improve the winning agreement of previous rounds and more rounds may be needed to explore options to find such offers. We therefore formulate the following hypothesis about the number of rounds needed to reach a final agreement in the third variant compared to that needed in the second.

Hypothesis 4 On average more rounds will be needed to reach a final winning agreement using the third setup than the second.

We have argued that it is important that parties are able to learn opponent preferences. One question that remains is whether the sellers have an incentive to learn, or that they can achieve the same or even higher utility without learning.

Hypothesis 5 An agent will be better off by learning the preferences of the opponent than without learning.

To test this hypothesis we present some evidence where we compare the results of using a negotiation strategy that uses (Bayesian) learning to another strategy that does not.

5 Experimental Evaluation

In this section, we first discuss the design of the experimental setup and then present the obtained results. We present experimental results to evaluate how well each of the three multi-bilateral negotiation mechanisms approximate the QVA, although they do not make the buyer's preferences public. We also investigate the number of rounds required in the second and third protocol, and we investigate whether learning dominates not learning.

5.1 Experimental Design

The first experimental design choice concerns the number of sellers that participate in the negotiations. While the mechanism nor the protocol limit the number of sellers, in the experiments we use only two sellers with distinct preference profiles. This

already simulates the competition between sellers since only the best and second-best can influence the outcome. Admittedly, when more sellers participate, the buyer may make a mistake in some of the rounds in finding the best seller, because it usually cannot perfectly learn the preferences of all sellers. However, in the multi-round protocols, such mistakes are easily repaired in future rounds. In addition, increasing the number of sellers reduces the expected difference between the best and second-best seller. This may give better results (regarding the deviation from the outcome of the QVA), because when accidentally the second-best is chosen as the best, the outcome will only marginally differ.

The second choice concerns the domain of negotiation. We have deliberately chosen a very generic domain and even relaxed natural constraints on this domain to further ensure generality. In the experiments we have used the so-called service-oriented negotiation domain taken from (Faratin et al. 2003). This domain consists of four issues that need to be settled, which represent the various attributes considered relevant with respect to the service offered, and include for example delivery time, quality, duration, and a penalty. Although we did use the generic four-issue structure of this domain we did not impose specific restrictions on the preferences such as that a lower penalty is always preferred by a seller as would be natural in this domain. As a result we have more variation in the preference profiles than one would typically expect in this domain. This variation in preference profiles ensures the relevance of our results for other domains as well.

For the experiments we have created a set of 12 preference profiles per role each, 12 for the buyer role and 12 for the seller role. Preference profiles were represented as piece-wise linear additive utility functions and each party in addition was assigned a reservation value. The remaining parameters such as the relative importance of an issue (weights), the utility associated with the alternatives for each issue (called an evaluation function), and the reservation values are set as follows:

1. To model the relative importance of the value of the issues, two different sets of weights are used. One representing equal importance of all issues, using 0.25 as weight for each of the four issues, and a set of weights representing dominance of two issues over the other two, using the weights 0.30, 0.50, 0.05, and 0.15.
2. The utility associated with each of the alternatives associated with an issue were modeled by either a linear “uphill” function, a linear “downhill” function, or a combination of the two (resulting in a triangular shape). Two of the three types of evaluation functions are illustrated in Fig. 3.
3. The reservation value for the buyer and sellers was set to either 0.3 or 0.6.

In Fig. 3 an example of a preference profile for a buyer can be found. The relative scaling of the evaluation functions of the individual issue in the figure indicates its corresponding weight. The utility of a complete bid can be calculated by the summation of the utilities of individual issues.

Tables 1 and 2 show the predefined profiles that were created using variations of the three preference profile parameters defined above. The reservation value was varied with the preference profiles and set to either 0.3 and 0.6, and, as explained above, two weights vectors were associated with issues ($\langle 0.30, 0.50, 0.05, 0.15 \rangle$ and $\langle 0.25, 0.25, 0.25, 0.25 \rangle$). In a typical negotiation scenario it is normal to assume at




Fig. 3 Example of a preference profile of a buyer with weights 0.30, 0.50, 0.05, and 0.15. Issues 1, 3, and 4 have “uphill” utility function, issue 2 has a “triangular” shape utility function

Table 1 Predefined buyer profiles

| Profile | w_1 | w_2 | w_3 | w_4 | <i>eval.fn</i> | v_i |
|---------|-------|-------|-------|-------|----------------|-------|
| Buyer1 | 0.25 | 0.25 | 0.25 | 0.25 | uphill | 0.3 |
| Buyer2 | 0.30 | 0.50 | 0.05 | 0.15 | uphill | 0.3 |
| Buyer3 | 0.25 | 0.25 | 0.25 | 0.25 | downhill | 0.3 |
| Buyer4 | 0.30 | 0.50 | 0.05 | 0.15 | downhill | 0.3 |
| Buyer5 | 0.25 | 0.25 | 0.25 | 0.25 | triangle | 0.3 |
| Buyer6 | 0.30 | 0.50 | 0.05 | 0.15 | triangle | 0.3 |
| Buyer7 | 0.25 | 0.25 | 0.25 | 0.25 | uphill | 0.6 |
| Buyer8 | 0.30 | 0.50 | 0.05 | 0.15 | uphill | 0.6 |
| Buyer9 | 0.25 | 0.25 | 0.25 | 0.25 | downhill | 0.6 |
| Buyer10 | 0.30 | 0.50 | 0.05 | 0.15 | downhill | 0.6 |
| Buyer11 | 0.25 | 0.25 | 0.25 | 0.25 | triangle | 0.6 |
| Buyer12 | 0.30 | 0.50 | 0.05 | 0.15 | triangle | 0.6 |

least some level of opposition between the buyer’s and the seller’s preferences. To ensure this, evaluation functions for the issues 1, 3, and 4 of the buyer’s profiles are set to the “uphill” type and the seller’s evaluation functions for the issues 2, 3, and 4 are fixed to the “downhill” type. To vary the level of opposition between the buyer’s and the seller’s profiles the type of the evaluation function of the remaining issue is set to one of the three possible types “uphill”, “downhill”, and “triangle”. These variations result in a total of $2 * 2 * 3 = 12$ possible profiles per role.

A sample of 50 different negotiation setups is created by means of a random selection out of the twelve profiles from Tables 1 and 2 for each of the three roles (one buyer, two sellers). Moreover, as a seller with a lower reservation value in such a setup has a higher chance of winning the first round (due to convexity of the Pareto efficient frontier), the sample is balanced such that in 80% of the cases the sellers have equal reservation values. To generate 20% of the negotiation setups where sellers have unequal reservation values, a complete set of all possible seller pairs with unequal reservation values is build. This set is used for the random selection of the negotiation setups. The rest of the sample (80%) of the seller profiles with equal reservation values was generated in a similar way.

Finally, a choice has to be made concerning the type of negotiating agent and the strategy that agent uses. As we have argued above, learning a preference profile is

Table 2 Predefined seller profiles

| Profile | w_1 | w_2 | w_3 | w_4 | $eval.fn$ | v_i |
|----------|-------|-------|-------|-------|-----------|-------|
| Seller1 | 0.50 | 0.30 | 0.15 | 0.05 | uphill | 0.3 |
| Seller2 | 0.25 | 0.25 | 0.25 | 0.25 | uphill | 0.3 |
| Seller3 | 0.50 | 0.30 | 0.15 | 0.05 | downhill | 0.3 |
| Seller4 | 0.25 | 0.25 | 0.25 | 0.25 | downhill | 0.3 |
| Seller5 | 0.50 | 0.30 | 0.15 | 0.05 | triangle | 0.3 |
| Seller6 | 0.25 | 0.25 | 0.25 | 0.25 | triangle | 0.3 |
| Seller7 | 0.50 | 0.30 | 0.15 | 0.05 | uphill | 0.6 |
| Seller8 | 0.25 | 0.25 | 0.25 | 0.25 | uphill | 0.6 |
| Seller9 | 0.50 | 0.30 | 0.15 | 0.05 | downhill | 0.6 |
| Seller10 | 0.25 | 0.25 | 0.25 | 0.25 | downhill | 0.6 |
| Seller11 | 0.50 | 0.30 | 0.15 | 0.05 | triangle | 0.6 |
| Seller12 | 0.25 | 0.25 | 0.25 | 0.25 | triangle | 0.6 |

an important capability required when the preferences of the buyer are not publicly known. For this reason, we use an agent capable of learning a preference profile in a single negotiation session using Bayesian learning introduced in [Hindriks and Tykhonov \(2008a\)](#). In the experiments, this negotiating agent builds a model of opponent preferences by learning a probability distribution over a set of hypotheses about the utility function of its opponent. In our case the agent has to learn the weights of issues and the corresponding evaluation functions. These structural assumptions make the learning task feasible.

We briefly explain the learning mechanism itself, for details please see [\(Hindriks and Tykhonov 2008a\)](#). During a negotiation session every time a new bid is received from the opponent the probability of each hypothesis about the opponent's utility function is updated using Bayes' rule. To be able to use Bayes' rule the conditional probability that the bid might have been proposed given a hypothesis is used. The utility of the bid according to the current hypotheses is computed and compared with a predicted utility based on the assumption that the opponent uses a concession-based tactic. This assumption is rational as in general any negotiator has to concede to reach an agreement. The details of the strategies in combination with the protocols as used in our simulation experiments can be found in [Appendix A](#).

5.2 First Negotiation Protocol

This first set of experiments should test our hypothesis that negotiating agents that use the first protocol and can learn a preference profile on the fly are able to approximate the outcome determined by the QVA mechanism quite well. For this, we study two results. Firstly, we compare the number of times the same winner is selected by the first multi-bilateral negotiation mechanism as by the QVA. Secondly, we study the differences in utility for both the winner and the buyer in case the same winner is selected.




Fig. 4 The distribution of the difference in utility for the buyer (*left*) and the seller (*right*) between the outcome of the first negotiation protocol and the outcome selected by the QVA

First of all, in our experiments the winner defined by the QVA mechanism and the winner of the multi-lateral negotiation protocol in the experiments completely coincide. This means that in the first round (of this first negotiation protocol) the same seller is selected as a winner as in the QVA.

Next, consider the differences in the utility of the outcome for both the buyer and the winning seller, represented by the histograms in Fig. 4. These results are obtained in the second round of the protocol and show that in general the outcomes obtained via the negotiation protocol approximate those of the QVA mechanism. In 78% of the experiments the difference is less than 5%. The average difference in utility for the buyer between the efficient outcome of the QVA and the experimental results is only -0.09% and the standard deviation is 4%. Moreover, in 94% of the experiment the difference was not more than 10%. For the (winning) seller the differences are even smaller (0.01% on average with a standard deviation of 5%), indicating that overall the outcomes were good approximations. Moreover, some of the bigger deviations could be traced back to difficulties with learning an opponent's preference profile.

To summarize, these two observations indicate that there is no reason to conclude that the QVA and the first negotiation protocol are significantly different, supporting Hypothesis 1. The (small) difference in utility from the outcome selected by the QVA can be explained as follows. In this protocol all agents try to maximize the opponent's utility while staying above their reservation value. For this, the ability of an agent to learn the preferences of an opponent is a key factor in a successful approximation of the auction mechanism. First, the selection of the winning (as well as the second-best) offer mainly depends on the ability of a seller to learn the preference profile of the buyer, because otherwise acceptable offers that maximize the buyer utility cannot be found. Second, the utility of the winning seller in the final agreement is determined by the buyer's ability to learn the seller's preference profile, because otherwise the outcome will not be near the Pareto front of the winning seller and the buyer. The difference from the utility of the QVA outcome can thus be explained by approximation errors in the used learning method.

This first protocol requires sellers to make only offers at their reservation value in the first round, and in the second round it requires that the buyer makes offers at the value of the second-best offer in the first round (see an example of a negotiation trace in Fig. 5). In many cases, imposing such requirements is unrealistic. In previous work (Hindriks et al. 2008) we have seen that simply relaxing these requirements does not result in outcomes that are still approximating the outcome of the QVA. In this paper we therefore propose a multi-round protocol, which is evaluated hereafter.

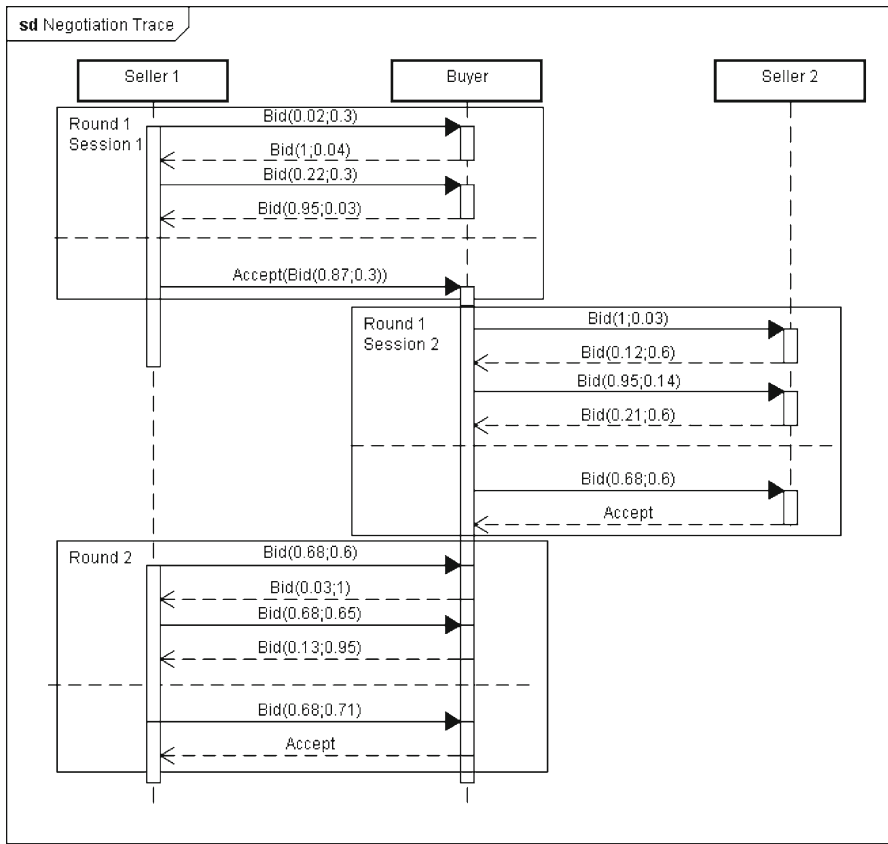


Fig. 5 Example of a negotiation trace for the first negotiation protocol (each bid is denoted by the utility of the buyer and the seller). Seller 1 wins with an outcome with utility 0.68 for the buyer and 0.71 for the seller

5.3 Second Negotiation Protocol

This second set of experiments tests the hypothesis that the second (multi-round) protocol approximates the outcome of the QVA. Figure 6 shows an example of a negotiation trace for the second negotiation protocol. As above, the winner defined by the QVA and the winner in the negotiation experiments coincide in all of the runs. Again the outcomes obtained by using the negotiation protocol are quite close to those determined by the mechanism. Figure 7 shows the histograms of the differences of the utility of the outcomes. The average difference with the buyer's utility for the QVA outcome is 0.01% (with a standard deviation of 1.5%). The utility of the sellers differs from their utility of the QVA outcome by -0.37% (with a standard deviation of 1.6%). According to the t -test the difference between the means of the utilities of the QVA outcome and the experimental results are not only very small, but also insignificant (for the buyer: $t = 0.054$, $P(T < t) = 0.957$, for the seller: $t = 1.648$, $P(T < t) = 0.106$). This supports Hypothesis 2.

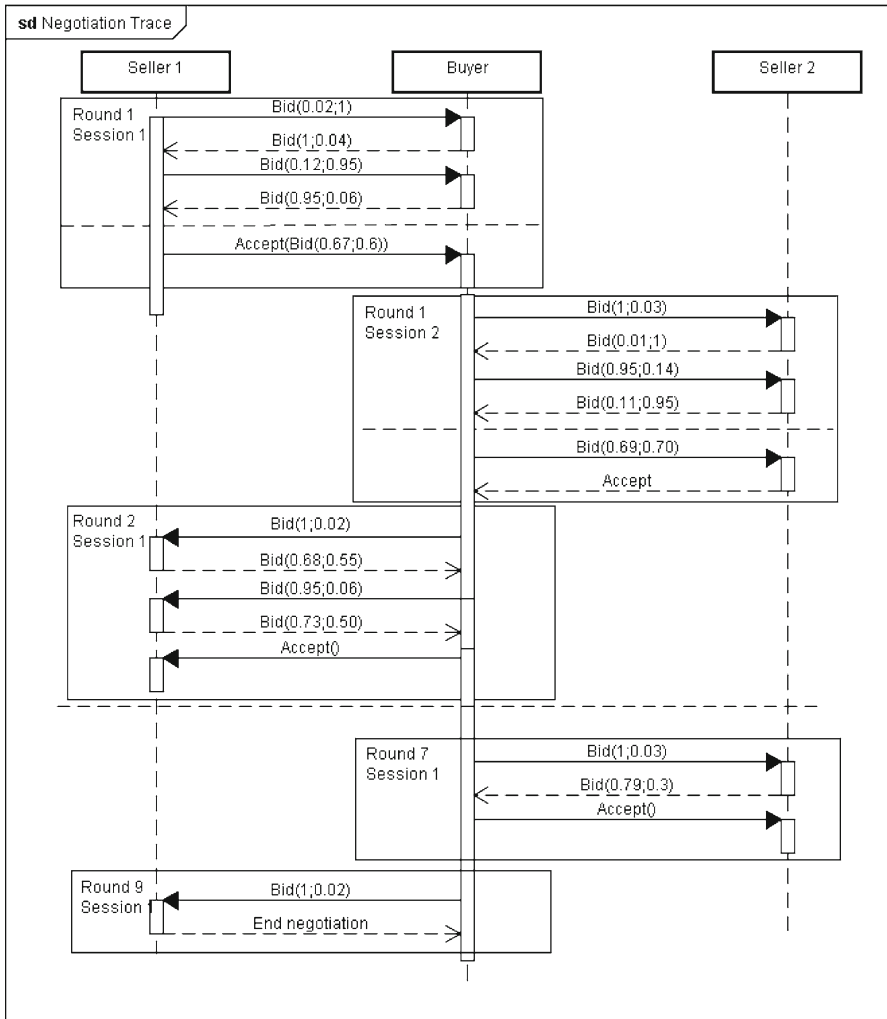


Fig. 6 Example of a negotiation trace for the second negotiation protocol. Here Seller 2 wins negotiation with an outcome of (0.79, 0.3)




Fig. 7 The distribution of the difference in utility for the buyer (left) and the seller (right) between the outcome of the second negotiation protocol and the outcome selected by the QVA




Fig. 8 The distribution of the difference in utility for the buyer (*left*) and the seller (*right*) between the outcome of the third negotiation protocol and the outcome selected by the QVA

Moreover, as before, the difference of the experimental results and the utility of the outcome selected by the QVA can be explained by approximation errors in the learning method used.

5.4 Third Negotiation Protocol

Even the experimental results using the third negotiation protocol show only a small deviation from the outcome selected by the QVA. Figure 8 shows the histograms of the differences in utility for these outcomes. The average difference from the utility of the buyer's outcomes in the QVA is 1.39% (with a standard deviation of 2.2%). The utility of the sellers differ from the utility of the outcome in the QVA by -1.28% (with a standard deviation of 2.3%).

On average, the buyer gets a slightly better outcome in the proposed negotiation setup compared to the QVA outcome ($t = -3.8$, $P(T < t) = 0.00043$). This results in somewhat lower utilities for the sellers ($t = -3.9$, $P(T < t) = 0.00027$). These differences thus, although small, are significant. This can be explained by the fact that unlike in the auction mechanism, where the final agreement always corresponds to the reservation value of the second-best seller, in this last setup the sellers are not aware of each other's reservation value. Therefore, on the one hand, the deviation of the utility of the outcome is influenced by the size of the concessions made by the winning seller. As a result, the buyer can benefit from the seller's concessions. On the other hand, due to imperfection of the learned model of the opponent preferences, the second-best seller might drop out of the negotiation too early. The winning seller can benefit from this because no more concessions on her behalf are necessary. In such a case, the final agreement has a lower utility for the buyer. This relationship between the buyer's and the seller's utility of the final agreement can be observed in Fig. 9. There we can see that one negotiating party can benefit from the underperformance of the other.

Even though the utilities of the outcome of this third negotiation protocol and the QVA are significantly different, they are still quite close. In addition, the same seller is always selected as a winner. We therefore conclude that also this third protocol is a reasonable approximation of the QVA, supporting our third hypothesis.

On average the number of the negotiation rounds in the third protocol is significantly higher than in the second setup (3.5 against 11.3, respectively, $t = -9.39$, $P(T < t) = 1.9 \cdot 10^{-12}$). Moreover, per round, using the third protocol almost two times more offers were made than in the second setup. That is, on average 50 offers were made in the




Fig. 9 The relation between the difference in utility for the buyer (*vertical axis*) and the seller (*horizontal axis*) for the second (*left*) and the third (*right*) negotiation protocol

second setup against 23 offers in the first one ($t = -14.4$, $P(T < t) = 4.14 \cdot 10^{-19}$). This confirms our fourth hypothesis.

All these experiments are done on an Intel Pentium 4 3.0GHz processor. For each experiment we measured the total run time of all rounds, including the Bayesian learning and the alternating-offers negotiation with two sellers in each round. For this third protocol, the run time is about 15.6 min on average, with a standard deviation of 23%. For the second protocol this is about 5.4 min on average with a standard deviation of 13%.

5.5 Negotiation Strategy

In the experiments discussed above we have shown that when all agents use the Bayesian learning strategy it is possible to approximate the outcome of the QVA. We also argued that learning is an essential part of any strategy that is able to realize outcomes similar to the auction. The question of quality of learning of the used learning technique was studied in [Hindriks and Tykhonov \(2008b\)](#) and, hence, is not covered in this paper. An important question that remains is whether such a strategy *dominates* other (non-learning) strategies. A strategy is said to *dominate* another strategy if it outperforms that strategy.

In order to (partially) answer this question we perform some additional experiments. These experiments are similar to the experimental setup described above but we replace one of the sellers with a seller that uses the Zero-Intelligence strategy ([Gode and Sunder 1993](#)). The Zero Intelligence (ZI) strategy randomly proposes bids above its reservation value. On average, it is difficult for the ZI strategy to achieve a better agreement than its reservation value and any effective negotiation strategy is expected to outperform it. However, accidentally, the ZI strategy may make very smart moves and if a certain strategy always outperforms ZI, it may be concluded that this strategy also outperforms many other strategies. As before, 50 negotiation setups with two sellers and one buyer are used but this time for each of these setups two variants are run: (i) one where the first seller uses the ZI strategy and the second seller uses the

Table 3 By changing to a ZI strategy, sellers do not win more often and do not receive a higher utility

| Strategy of the winner in the experiments | Matches the QVA winner | Number of sessions | Deviation of the outcomes | |
|---|------------------------|--------------------|--|------------------------------------|
| | | | Winning seller | Buyer |
| Bayesian | Yes | 50 | 6% (always > 0%), st. dev. = 3.5% | −6% (always < 0%), st. dev. = 3.1% |
| | No | 15 | N/A (utilities of the sellers' should not be compared) | −8% st. dev. = 6% |
| Zero-intelligence | Yes | 35 | 1% st. dev. = 2.4% | −10% st. dev. = 6.5% |
| | No | 0 | | |

When their competitor uses a ZI strategy, the winner using a Bayesian strategy stays the same and has a higher utility

Bayesian strategy, and, vice versa, (ii) where the first seller uses the Bayesian strategy and the other one uses the ZI strategy.

A summary of the experimental results of these 100 sessions is presented in Table 3. Every outcome of a negotiation session is classified into one of four possible cases, depending on the strategy of the winner (the first column), and whether the winner matches the winner in the QVA outcome (the second column). The third column provides the number of the negotiation sessions in each case. The deviation in the utility of the outcome for the winning seller and the buyer are given in the fourth and fifth column, respectively.

Most importantly these results show that there is not even a single negotiation session where a seller that loses in the QVA wins by switching from the Bayesian strategy to the ZI strategy (see the fourth row in the table). Moreover, a seller using the ZI strategy loses in 15 out of 50 runs, while it could have won using the Bayesian strategy. In addition, in cases where a seller wins in spite of using ZI, the utility is lower on average (1% better than the QVA outcome versus 6% better). These differences in utility are shown in more detail in Fig. 10. Most of these differences can be explained by the fact that adding a ZI strategy complicates learning (the ZI strategy itself does not learn). It explains why on average the buyer's utility is significantly lower as it is impossible to learn anything from the offers proposed by a ZI strategy.

A second conclusion from these results is that sellers that follow the Bayesian strategy do not obtain worse outcomes when other sellers use the ZI strategy. This can be seen by looking closely at the results for the 50 cases where the seller following the Bayesian strategy wins both when using the negotiation protocol as well as according to the QVA. In those cases, the same seller is still a winner, and on average obtains a higher utility (6% difference from the utility of the QVA compared to 1.28% more utility on average in the previous section).

To conclude, the results show that outcome utilities of a seller that uses the Bayesian strategy do not get worse when the second-best seller switches to the ZI strategy, nor do they get better when the seller itself switches to the ZI




Fig. 10 The relation between the difference in utility for the buyer (vertical axis) and the seller (horizontal axis) for the third protocol for the Bayesian (left, 50 cases) and the ZI (right, 35 cases) strategy

strategy. Therefore, regardless of the choice of a strategy by the opponent, the most rational choice for the seller among these two is to stick with the Bayesian strategy. The large number of the conducted experiments with a seller that uses the ZI strategy provide a significant variation of the negotiation behavior of that seller. Given the fact that in all cases the ZI strategy did not perform better than the Bayesian strategy we can derive that the Bayesian strategy is a good choice for rational sellers regardless of the strategy of the other sellers.

6 Related Work

The protocols presented in this paper relate to both auctions (because they select one winner among a set of sellers and the final agreements are efficient) as well as to negotiations (because the set of possible agreements is not agreed upon on forehand, and the parties do not know each other's preferences). In this section we therefore discuss both earlier work on (multi-attribute) auctions, as well as on multi-lateral negotiation, and we briefly discuss existing work on learning this context. Regarding the relation to auctions, the relation to the Vickrey auction and its generalization to multiple issues (possibly without money), the QVA, has already been discussed in the first two sections of this paper. The final two protocols, however, show stronger similarity to an auction that is even more familiar, i.e., the English auction.

Like in an English auction, in each round of the multiple bilateral negotiations, new agreements will only be accepted by the buyer if they are better than the best agreement in the previous round. When no other seller can make a better proposal, the process stops and the winning agreement is this best agreement. This process is very similar to an English auction, where bids are increased until all but one bidder stop bidding. The only difference is that in the setting discussed in this paper, the utility of the buyer is not known, while in a (reverse) English auction the item is fixed and the utility of the seller (buyer) is assumed to be linear in the price. This makes it very easy to come up with a better bid in an English auction, but quite hard to do so in our situation, where the sellers really need to learn the preferences of the buyer. The English auction is ex-post efficient, meaning that under the assumption that other sellers are rational as

well, a seller can never do better than to follow a straightforward strategy of making an offer that is ranked slightly higher than the previous bid as long as this is above its own reservation value. We believe a similar result for our setting can be derived, supporting also theoretically that sellers can never do better than learning the buyer's preferences as well as possible, and then making concessions until either its offer is higher than the best offer of the previous round, or its own reservation value has been reached (i.e., the fifth hypothesis).

In its basic form, the English auction, just like the Vickrey auction, is on one item that is completely described, but several generalizations of auctions have been proposed where some of the attributes of the item are left open for "negotiation". However, in extant work the payments are always seen as a special attribute for which the preferences of the buyer and the sellers are related: a lower price for the seller means a worse outcome for the buyer. For example, (Che 1993) analyzed situations where a bid consists of a price and a quality attribute, and proposed both first-price and second-price sealed-bid (e.g., Vickrey) auction mechanisms. His work was extended by David et al. (2002) for situations where the good is described by two attributes and a price. They analyzed the first-price sealed-bid and the English auction, and derived strategies for bids in a Bayesian-Nash equilibrium. In addition, they studied a setting where the buyer can also strategize, and they showed when and how much the buyer can profit from lying about its valuations of the different attributes.

Later work on iterative multi-attribute auctions has focused on a finite (discrete) domain with quasi-linear utility (Parkes and Kalagnanam 2005). For this domain two related protocols are proposed. In Nonlinear&Discrete (NLD) a reverse English (or Japanese) auction is held simultaneously for every combination of attribute values (called a bundle). In such an auction the price is dropped for each bundle until just one seller remains. The winning bundle is the one that maximizes the difference between the valuation of the buyer and the price. Straightforward bidding for sellers in this auction is defined by bidding the ask price for a bundle if that has still a positive utility. Straightforward bidding is shown to be an ex-post Nash equilibrium for the sellers and to result in an efficient outcome, maximizing the gains from trade (equal to the one-side VCG mechanism). On the buyer's side, strategizing can bring a benefit of at most the marginal value to the economy contributed by the winner. In addition, for the setting where the utility of the buyer is additive over all attributes, NLD can be simplified. The mechanism Additive&Discrete (AD) does not hold an auction for every possible bundle, but just for every value (level) of each attribute separately.

Our iterative multi-bilateral negotiations with increasing utility for the buyer differ from this work in the following aspects. First, we consider piece-wise linear domains that may be continuous, which is a strict generalization of the finite domains. Second, we consider price just as one of the issues, which allows us to consider also utility functions that are not quasi-linear, but also makes it hard to define the gains from trade. We focus on Pareto-efficiency instead. Third, the mechanism differs in that we use bilateral negotiations between the buyer and each seller over all possible bundles. Such a negotiation usually involves only the exchange of a very limited subset of possible bundles, which is much more efficient compared to holding an auction for each possible bundle. Fourth, in these negotiations both the buyer and the sellers try to

learn each other's utility function in order to make better proposals. In general, there is not enough information to learn these perfectly, while in the finite setting at the end of the auction the utility function of one of the sides is completely known to the other party.

How much information is revealed exactly is an important topic for our future work. In [Parkes and Kalagnanam \(2005\)](#) this is measured using the normalized size of the set of possible weights for the issues, since in both NLD and AD a utility function is defined by the weights over the issues. For the sellers, this set is given by the constraints determined from the bids under the assumption that they follow the straightforward strategy. The size of the resulting convex set is approximated by a simple Monte Carlo algorithm.

We are aware of one other paper explicitly discussing the use of learning in the context of multilateral negotiation or multi-attribute auctions. In [Beil and Wein \(2003\)](#) the buyer learns the cost function of the sellers under the assumption that they all have a fixed form where only P parameters need to be determined. This can be exactly computed using inverse optimization after P auction rounds, where in each round, the buyer makes a different utility function public and assumes the sellers are bidding according to a straightforward strategy regarding the announced utility function. Only round $P + 1$ is for real. In that last round, the buyer constructs a utility function that maximizes its revenue. In the Bayesian learning approach used in this paper the buyer does not announce, nor change its utility function between rounds. If he did, the sellers would have more difficulty learning and in fact, we have seen that this usually reduces the buyer's utility. Moreover, as discussed above, in our work there is an incentive for the sellers to learn as well as possible and then use a straightforward concession strategy. In contrast, in the setting of [Beil and Wein \(2003\)](#), there is a clear incentive for the sellers to hide their utility in the first P rounds, to prevent the buyer to exploit them in the final round.

A multi-unit version of multi-attribute auctions is discussed in [Teich et al. \(2006\)](#). In their setting, two types of attributes are distinguished: those that relate to the seller (called bid attributes), and those that do not (called negotiable bid issues). Utility information regarding the negotiable bid issues is communicated to the sellers in the form of linear bonuses and penalties. Sellers indicate what the values of the issues are and how much they are willing to sell. The mechanism returns the current price, and the seller can then accept, reject, or change the submitted issues. The price is obtained by the mechanism using a straightforward strategy.

The extension to also deal with multiple units is straightforward, but significantly increases the applicability. We believe our mechanism can also easily be adapted to deal with multiple (identical) units. In [Teich et al. \(2006\)](#) the utility function of the buyer is assumed to be a weighted sum of the negotiable bid issues. There is no assumption on the utility of the sellers, except that it is quasi-linear (since all issues are related to the price). Again, quasi-linearity is a restriction compared to our model. The most relevant difference from our approach, however, is the fact that the utility function of the buyer is communicated, except for a fixed discount that may be different for each seller, while in our approach we explicitly do not allow this.

Related work on negotiation mechanisms that deal with multiple players is reported in [Li et al. \(2004\)](#), [Nguyen and Jennings \(2003\)](#), [Rahwan et al. \(2002\)](#). Our approach

differs in at least two regards. First, our aim has been to reach an agreement that is as close as possible to an efficient agreement as obtained by the QVA. Second, we propose a new negotiation protocol that is based on several rounds of multiple standard bilateral negotiation sessions where all participants that lost in an earlier round are allowed to make a proposal that is better than the winning proposal of the earlier round. Below we consider these existing negotiation approaches in a bit more detail.

[Rahwan et al. \(2002\)](#) and later [Nguyen and Jennings \(2003\)](#) have proposed a negotiation framework where the buyer negotiates with a number of sellers concurrently, and updates its reservation value in all other negotiation threads with the value of an agreement, whenever one is made. The latter work presents experimental results on the effect of a number of negotiation strategies in a setting where each utility function is a standard linear combination of the issues. It seems that in such a parallel setting the speed of the negotiation threads may influence the changes in reservation value of the buyer and thus the result. In our work this is resolved because there is always a next round until all sellers except one decide to end the negotiation.

Another line of work in this field includes an expectation about results obtained in other threads ([Li et al. 2004](#)). Like in the work discussed above, the reservation value for the buyer is set based on events in the other threads. The interesting extension here is that the reservation value can be set at the expected best offer in other threads, or even in future threads.

As a final topic to discuss here, we briefly review an empirical study of comparing an auction mechanism with a negotiation mechanism ([Kjerstad 2005](#)). The aim of this study is to estimate the impact of trading mechanisms on the price of an agreement. The paper considers data of 216 trades that result from either using an auction or a negotiation as the trade mechanism. It is concluded that trade mechanisms can have an influence on the number of suppliers due to, i.e., costs associated with a particular mechanism. The results show that the choice of the trading mechanism does not influence the price of an agreement, however. This seems to contradict our results for the third protocol, where we see a (albeit small) difference in utility. We believe this can be attributed to imperfect learning, but we leave a thorough investigation of the cause of this difference for future work.

7 Conclusion

In general, negotiations facilitate the expression of agreements in greater detail than auctions do, making it possible to arrive at better win–win solutions between the buyer and the seller. However, (reverse) auctions on the other hand can guarantee that the deal is with the best seller, and some auctions, such as the English auction or the Vickrey auction, remove the need for bidders to strategize, making it a lot easier to participate (following a straightforward strategy).

Sandholm, among others, acknowledges this, and proposes a combinatorial auction that allows for as much details in bids as the buyers and sellers would find useful, a method he called expressive commerce ([Sandholm 2007](#)). However, even in that approach, the preferences of the buyer are quasi linear and need to be given on

forehand to allow all participants to make successful bids.⁴ The main problem this paper deals with is the fact that the preferences of the buyer in a one-to-many multi-issue setting may not be given on forehand and may not even be quasi linear, which is required for most auction types. Even the QVA requires that the utility function of the buyer is made publicly known. The main contribution of this paper is the idea of combining a negotiation protocol with the ability of the parties to learn the preferences of the opponent, removing the need to make these preferences public. The protocol proposed introduces multiple negotiation rounds in which sellers that lost in the previous round are given an opportunity to improve their offers and possibly outbid the winner.

We have discussed three variants of multi-round negotiations to approximate an auction. We showed experimentally that both the outcomes of the two-round protocol as well as the second (multi-round) protocol are not significantly different from the Pareto-efficient best-seller outcome of the QVA. The results of the third set of experiments indicate that even if no information is made public until the end of the negotiation the protocol closely approximates the outcome of the QVA. The number of rounds needed to find the winning contract, however, is significantly higher in the third than that used in the second protocol. This can be explained by the fact that sellers have no information about the winning agreement of the previous negotiation round and have to make more offers to be able to explore the outcome space before they are able to outbid the winner. Our results thereby show that a trade-off needs to be made between revealing preference information and the average amount of time needed to complete the negotiation. Our final set of experiments show that Bayesian learning in combination with a concession strategy dominates a Zero Intelligence strategy with random offers. This supports our hypothesis that the Bayesian strategy is dominant. A full proof of this claim is left for our continued studies.

For other future work, we are interested in potential forms of manipulation that may be available to the buyer in the third protocol in case the process cannot be monitored by a trusted third party. If a buyer has complete knowledge about the winner to be, he could lie about an offer in an earlier round. This “second-highest offer” can then be chosen in such a way that the negotiation space of the final agreement will be very small, in favor of the buyer. Finally, we also want to study how to modify the ideas presented in this paper to make the protocols presented applicable to a broader range of real-world one-to-many multi-issue negotiations over complex domains where preferences cannot completely be made public in advance.

Acknowledgments Dmytro Tykhonov is supported by the Technology Foundation STW, applied science division of NWO, and the Ministry of Economic Affairs of the Netherlands.

⁴ Minor changes in the preferences are allowed afterwards (scenario navigation), but may influence the efficiency.

Appendix

A Instantiations of the Negotiation Protocols with a Negotiation Strategy

Below we give the complete description of the negotiation processes used in this paper. In this description, the following notation is used.

| | |
|----------------------|---|
| $u_0(x)$ | utility of the buyer |
| $u_i(x)$ | utility of the i -th seller |
| $\tilde{u}_0(x)$ | model of the buyer's preferences learned by a seller |
| $\tilde{u}^{s_i}(x)$ | model of the i -th seller's preferences learned by the buyer |
| ω_i | outcome of the first round in negotiation with i -th seller (Algorithm 2) |
| i^* | winner of the first round (Algorithm 2) |
| $\hat{\omega}$ | the second-best outcome of the first round (Algorithm 2) |
| ω_i^r | outcome of the r -th round in negotiation with i -th seller (Algorithm 3) |

Algorithm 2 A protocol based on two negotiation rounds instantiated with a negotiation strategy which uses Bayesian learning to model opponent's preferences

1. Round 1: The buyer bilaterally negotiates with every seller.
 - 1.1 The buyer starts negotiation with seller i with its best bid $x_0^b = \arg \max u_0(x)$ and $t = 1$.
 - 2.2 Seller i updates the buyer's preferences model $\tilde{u}_0(x)$ given the buyer's bid x_{t-1}^b .
 - 3.3 The seller searches for a bid that has utility of the seller's reservation value and maximizes the expected utility of the buyer: $x_t^{s_i} = \arg \max_{x \in \{x | u_i(x) = v_i\}} \tilde{u}_0(x)$
 - 4.4 If the found bid has not been proposed before then the seller sends it to the buyer, otherwise stop the bidding.
 - 5.5 If seller stopped bidding, use the latest seller's bid $x_t^{s_i}$ as an agreement ω^{s_i} and go to step 2.
 - 6.6 Otherwise, the buyer updates the seller's i preference model $\tilde{u}^{s_i}(x)$ given x^{s_i} .
 - 7.7 The buyer makes a concession in its own utility space $c(t)$ and sends a bid that maximizes the expected utility of the seller i given this concession: $x_t^b = \arg \max_{x \in \{x | u_0(x) = u_0(x_0^b) - c(t)\}} \tilde{u}^{s_i}(x)$
 - 8.8 Increase t and go to step 1.2.
 2. Round 2: The buyer determines the winner and negotiates a final agreement with the winner.
 - 2.1 The buyer determines the winner $i^* = \arg \max_{i \in \{1, \dots, n\}} u_0(\omega_i)$ and the second-best outcome $\hat{\omega} = \arg \max_{i=1, \dots, i^*-1, i^*+1, \dots, n} u_0(\omega_i)$.
 - 2.2 The winning seller i^* starts negotiation with its best bid $x_0^{s_{i^*}} = \arg \max u_{s_{i^*}}(x)$ and $t = 1$.
 - 2.3 The buyer updates the seller's preferences model $\tilde{u}_{i^*}(x)$ given the seller's bid $x_{t-1}^{s_{i^*}}$.
 - 2.4 The buyer searches for a bid that has utility of the second-best outcome and maximizes the expected utility of the seller: $x_t^b = \arg \max_{x \in \{x | u_0(x) = u_0(\hat{\omega})\}} \tilde{u}_{s_{i^*}}(x)$
 - 2.5 If the found bid has not been proposed before then the buyer sends it to the seller, otherwise stop the bidding and use the buyer's latest bid x_t^b as a final agreement ω . Exit.
 - 2.6 Otherwise, the seller updates the buyer's preference model $\tilde{u}^{s_1}(x)$ given x^{s_i} .
 - 2.7 The seller makes a concession in its own utility space $c(t)$ and send a bid that maximizes the expected utility of the buyer: $x_t^s = \arg \max_{x \in \{x | u_i(x) = u_i(x_0^{s_{i^*}}) - c(t)\}} \tilde{u}^b(x)$
 - 2.8 Increase t and go to step 2.3.
-

Algorithm 3 A protocol with multiple negotiation rounds instantiated with a negotiation strategy which uses Bayesian learning to model opponent's preferences

1. Set the utility of ω^0 to be 0. Set the round number r to be 1.
 2. While no final agreement ω has been found, do the following for every seller i that still participates.
 - 2.1 The buyer starts negotiation with seller i with its best bid $x_0^b = \arg \max u_0(x)$ and $t = 1$.
 - 2.2 While no agreement in round r has been reached:
 - 2.2.1 Seller i updates the buyer's preferences model $\tilde{u}_0(x)$ given the buyer's bid x_{t-1}^b .
 - 2.2.2 Seller i determines the utility of the next bid $u_t^{s_i}$:
 - A. If $t = 1$, then set $u_t^{s_i}$ to be the utility of ω^{r-1} or $\max u_{s_i^*}(x)$ in case $r = 1$.
 - B. Otherwise, make a concession $u_t^{s_i} = u_{t-1}^{s_i} - c(t)$.
 - 2.2.3 If utility of $u_t^{s_i}$ is less than the reservation value v_i , withdraw from negotiation.
 - 2.2.4 Determine a bid that has utility $u_t^{s_i}$ and maximizes the expected utility of the buyer: $x_t^{s_i} = \arg \max_{x \in \{x | u_i(x) = u_t^{s_i}\}} \tilde{u}_0(x)$ and send this to the buyer.
 - 2.2.5 The buyer updates the seller's preferences model $\tilde{u}_{s_i^*}(x)$ given the seller's bid $x_t^{s_i}$.
 - 2.2.6 If the buyer's utility $u_0(x_t^{s_i})$ is more than the utility of the best outcome of the previous round, accept the seller's offer, i.e., $\omega_i^r = x_t^{s_i}$.
 - 2.2.7 The buyer makes a concession $c(t)$ and determines a bid that maximizes the expected utility of the seller i : $x_t^b = \arg \max_{x \in \{x | u_0(x) = u_0(x_{t-1}^b) - c(t)\}} \tilde{u}^{s_i}(x)$.
 - 2.2.8 If the utility of u_t^b is less than the utility of $u_t^{s_i}$, accept the seller's bid anyway.
 - 2.2.9 Otherwise, send u_t^b .
 - 2.2.10 If the utility of the buyer's next bid $u_0(x_t^b)$ is less than its reservation value v_0 then send a bid with the reservation value: $x_t^b = \arg \max_{x \in \{x | u_0(x) = v_0\}} \tilde{u}^{s_i}(x)$.
 - 2.3 If only one seller has reached an agreement, then the agreement of the previous round is the final agreement $\omega = \omega^{r-1}$.
 - 2.4 Otherwise, publicly announce the winning agreement $\omega^r = \omega_{\arg \max_i \{u_0(\omega_i^r)\}}^r$.
 - 2.5 Start a new round, setting $r = r + 1$.
-

References

- Beil D, Wein L (2003) An inverse-optimization-based auction mechanism to support a multiattribute RFQ process. *Manag Sci* 49(11):1529–1545
- Bichler M, Lee J, Lee H, Chung J (2001) Absolute: an intelligent decision making framework for e-sourcing. In: *Proceedings of the 3rd international workshop on advanced issues of e-commerce and web-based information systems*, San Jose, CA, pp 195–201
- Che Y (1993) Design competition through multidimensional auctions. *RAND J Econ* 24(4):668–680
- David E, Azoulay-Schwartz R, Kraus S (2002) Protocols and strategies for automated multi-attribute auctions. In: *Proceedings of the 1st international joint conference on autonomous agents and multiagent systems*. ACM, New York, pp 77–85
- Faratin P, Sierra C, Jennings NR (2003) Using similarity criteria to make negotiation trade-offs. *J Artificial Intell* 142(2):205–237
- Gode DK, Sunder S (1993) Allocative efficiency in markets with zero intelligence traders: market as a partial substitute for individual rationality. *J Polit Econ* 101(1):119–137
- Harrenstein P, Mahr T, de Weerd MM (2008) A qualitative vickrey auction. In: Endriss U, Paul WG (eds) *Proceedings of the 2nd international workshop on computational social choice*. University of Liverpool, pp 289–301
- Harrenstein P, de Weerd MM, Conitzer V (2009) A Qualitative vickrey auction. In: *Proceedings of the ACM conference on electronic commerce*. ACM Press, pp 197–206
- Hindriks KV, Tykhonov D (2008a) Opponent modelling in automated multi-issue negotiation using bayesian learning. In: *Proceedings of the 7th international conference on autonomous agents and multiagent systems*, pp 331–338

- Hindriks KV, Tykhonov D (2008b) Towards a quality assessment method for learning preference profiles in negotiation. In: *Proceeding of workshop on agent-based complex automated negotiations*, pp 64–71
- Hindriks KV, Tykhonov D, de Weerd MM (2008) Approximating an auction mechanism by multi-issue negotiation. In: Hindriks KV, Brinkman WP (eds) *Proceedings of the 1st international working conference on human factors and computational models in negotiation*, pp 33–38
- Kjerstad E (2005) Auctions vs negotiations: a study of price differentials. *Econometrics Health Econ* 14(12):1239–1251
- Li C, Giampapa J, Sycara K (2004) Bilateral negotiation decisions with uncertain dynamic outside options. In: *Proceedings of 1st IEEE international workshop on electronic contracting*, pp 54–61
- Nguyen TD, Jennings N (2003) Concurrent bilateral negotiation in agent systems. In: *Proceedings of 14th international workshop on database and expert systems applications*, pp 844–849
- Osborne MJ, Rubinstein A (1994) *A course in game theory*. MIT Press, Cambridge
- Parkes D, Kalagnanam J (2005) Models for iterative multiattribute procurement auctions. *Manag Sci* 51(3):435–451
- Rahwan I, Kowalczyk R, Pham H (2002) Intelligent agents for automated one-to-many e-commerce negotiation. In: *Proceedings of 25th Australasian conference on computer science*, Australian computer society, Inc., Darlinghurst, Australia, pp 197–204
- Sandholm T (2007) Expressive commerce and its application to sourcing: how we conducted \$35 billion of generalized combinatorial auctions. *AI Mag* 28(3):45–58
- Teich JE, Wallenius H, Wallenius J, Koppius OR (2004) Emerging multiple issue e-auctions. *Eur J Oper Res* 159(1):1–16
- Teich JE, Wallenius H, Wallenius J, Zaitsev A (2006) A multi-attribute e-auction mechanism for procurement: theoretical foundations. *Eur J Oper Res* 175(1):90–100
- Thomas CJ, Wilson BJ (2005) Verifiable offers and the relationship between auctions and multilateral negotiations. *Econ J* 115(506):1016–1031
- Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J Finance* 16(1):8–37