

# Approximating the Qualitative Vickrey Auction by a Negotiation Protocol

Koen V. Hindriks, Dmytro Tykhonov, and Mathijs de Weerd

Delft University of Technology, Delft 2628CD, The Netherlands  
{k.v.hindriks, d.tykhonov, m.m.deweerd}@tudelft.nl

**Abstract.** A result of Bulow and Klemperer has suggested that auctions may be a better tool to obtain an efficient outcome than negotiation. For example, some auction mechanisms can be shown to be efficient and strategy-proof. However, they generally also require that the preferences of at least one side of the auction are publicly known. However, sometimes it is very costly, impossible, or undesirable to publicly announce such preferences. It thus is interesting to find methods that do not impose this constraint but still approximate the outcome of the auction. In this paper we show that a multi-round multi-party negotiation protocol may be used to this end if the negotiating agents are capable of learning opponent preferences. The latter condition can be met by current state of the art negotiation technology. We show that this protocol approximates the theoretical outcome predicted by a so-called Qualitative Vickrey auction mechanism (even) on a complex multi-issue domain.

## 1 Introduction

A result of Bulow and Klemperer has suggested that auctions may be a better tool to obtain an efficient outcome than negotiation [1]. There are many different types of auctions, however, and although it is possible to show that these auctions have nice theoretical properties, some of the mechanisms proposed also require constraints to be in place which are not easy to meet in practice. This is particularly true for auctions which may result in complex, multi-issue outcomes. It thus becomes interesting to look for alternative methods that may be used that guarantee outcomes that approximate the auction mechanism. In this paper, we study a particular auction mechanism called a Qualitative Vickrey Auction (QVA) [2]. This is a generalization of the well-known Vickrey auction to a general complex multi-issue setting where payments are not essential. We show that the theoretical outcome predicted by this mechanism can be approximated by a specific negotiation protocol. The only assumption that we need to make to obtain this result is that the negotiating agents are able to learn the preferences of their opponents during a single negotiation session. Techniques to do so are available [3], making our proposal one that can be implemented given the current state of the art in negotiation.

The QVA studied here is particularly useful in a context where a single buyer tries to obtain a *complex deal* with one out of many sellers that are interested

in making such a deal. 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 may be fixed by a budget and therefore less interesting, supercomputers have many features (processing speed, memory, etc.) and requirements (power supply, cooling, etc.) that need to be settled to obtain a deal. Such a deal thus is complex as many issues have to be agreed upon. The QVA proposed in [2] provides an auction mechanism to obtain such a complex deal as it provides a method to obtain a multi-attribute outcome. Given some additional assumptions, this mechanism obtains a Pareto-efficient outcome that involves the seller that can make the best deal still acceptable to him. One assumption imposed by this mechanism, however, is hard to realize in practice. The mechanism requires that the buyer publicly announces her preferences. This is unrealistic for various reasons: (i) the buyer may not know the complete domain of possible outcomes (sometimes sellers come up with new options during the process), (ii) it may be hard to specify all preferences over a complex (possibly infinite) set of outcomes, or (iii) the buyer may not want to publicly reveal all her preferences in complete detail.

In previous work [4], a first step towards a negotiation protocol was made that can be used to replace the auction mechanism without this requirement. When the sellers and the buyer follow a particular strategy (staying close to their reservation value), this protocol leads to similar results as the auction. However, without this restriction, the results did not yield good approximations anymore. The latter result was to be expected as such a generic negotiation setting does not take into account the negotiation power the buyer has given that the buyer may choose from a range of available sellers.

In this paper, we propose two negotiation protocols that do take the negotiation power of the buyer into account but do not require parties to reveal their reservation value. The idea introduced here is that a protocol that consists of multiple negotiation rounds in which sellers are provided an opportunity to outbid the winner of the previous round may be used to approximate the QVA. 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).

The paper is organised as follows. In Section 2 we define the general setting of a buyer and multiple sellers or service providers 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 a deal with any one out of a set of available service providers. Section 3 introduces the QVA mechanism that may be used to reach such a deal. In Section 4 we then propose two multilateral negotiation protocols that approximate the outcome predicted by the auction mechanism and thus may be used to replace it. In Section 5 we then experimentally validate that the outcomes of the negotiation protocols approximate the outcome predicted by the auction mechanism. Section 6 discusses some related work and Section 7 concludes with a discussion of the results obtained and outlines some directions for future research.

## 2 Problem Definition

The setting we are interested in here consists of a buyer that wants to obtain a service from one out of a potentially large number of sellers or service providers. An agreement that fixes the parameters of the service to be provided is a complex outcome  $x = \langle x_1, \dots, x_m \rangle \in X$  over  $m$  issues in a domain  $X = X_1 \times \dots \times X_m$ . These issues define all aspects of the agreement, such as price, quality, start time, duration, guarantees, penalty, etc. Buyer and sellers are assumed to associate a utility value with each outcome and have reservation values that determine when outcomes do not improve the status quo.

We introduce the following notation. The buyer is denoted by 0 and sellers are denoted by  $i \in \{1, \dots, n\}$ . Each party  $i$  has a *reservation value*  $v_i$  that represents the minimal utility value that an agreement should have to be an acceptable outcome for that party. Outcomes with a lower utility are called *unacceptable*. Each party  $i$  also has a utility function  $u_i : X \rightarrow \mathbb{R}$  which determines the utility that party associates with an outcome. The goal is to find an agreement between the buyer and one of the sellers that is not only acceptable to both, but that is also Pareto-optimal, 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 the next section a mechanism is introduced that has a dominant strategy equilibrium that yields a Pareto efficient outcome, provided there are two or more sellers.

## 3 Qualitative Vickrey Auctions

Earlier work has shown how to reach an agreement between a buyer and one of a set of sellers using a closed-bid auction-like mechanism [2]. This mechanism is called the Qualitative Vickrey Auction (QVA) for its similarity to the well-known Vickrey auction [6]. 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.<sup>1</sup> (If money is involved, the relative value of alternatives for other issues may be assigned a price, i.e. be translated into money.)

This auction mechanism can be thought of as consisting of two rounds. In the first round (1a-c), the buyer first publicly announces her preferences. Then potential service providers (sellers) submit offers in response, and a winner is selected by the buyer. In a second round (2a-b), the buyer first determines the second-best offer (from her perspective again) she received from another seller, 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

---

<sup>1</sup> If none of the issues is about money, a reverse auction is not different from a standard auction.

can be determined by the winner since the preferences of the buyer are publicly announced). It is assumed that the bids offered in the first round all go through a trusted third party, such as a solicitor, who can check whether the buyer follows the protocol. Summarizing, the steps of the procedure are:

- 1a. The buyer announces her preferences.
- 1b. Every seller submits an offer.
- 1c. The buyer selects the winner according to her preferences.
- 2a. The buyer announces the second-best offer she received.
- 2b. 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 by:

$$i^* = \arg \max_{i \in \{1, \dots, n\}} \max \{u_0(x) \mid x \in X, u_i(x) \geq v_i\},^2$$

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  $\hat{x}$  then is given by

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

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.,

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

This outcome of the auction-like mechanism 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 main problem with a realistic implementation of the QVA is that the buyer needs to communicate 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 quite a 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

---

<sup>2</sup> We assume ties are broken by the buyer using a given ordering over the sellers.

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 auction, or the qualitative Vickrey auction discussed above cannot be used, because in such auctions 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 mechanism and where there is no need to publicly announce the preferences of the buyer.

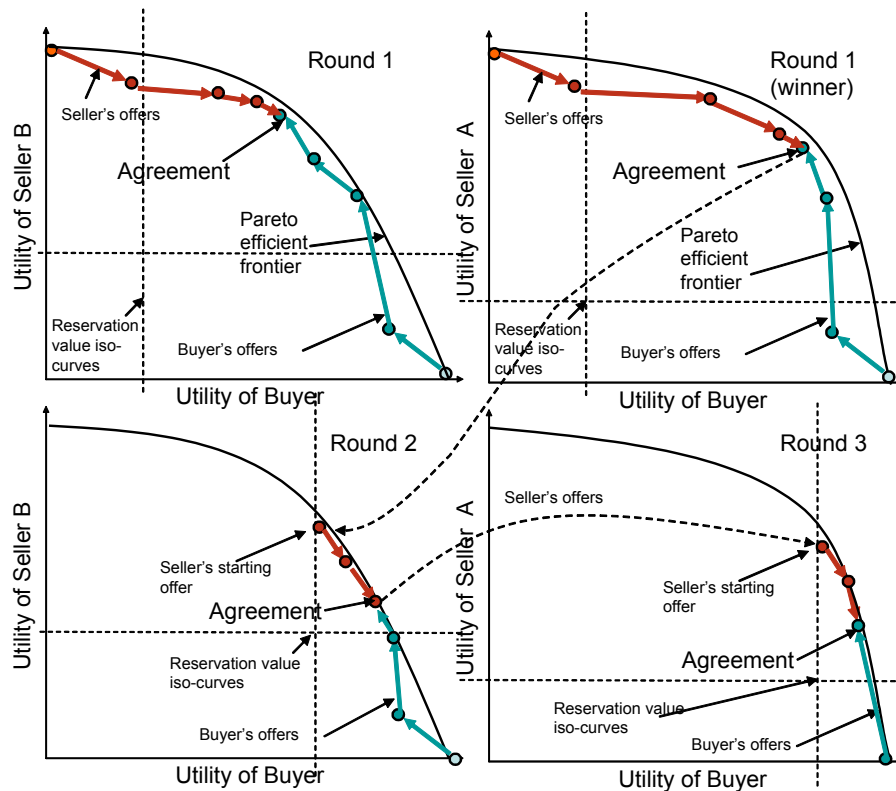
## 4 Negotiation Protocols

The QVA requires sellers to first propose an offer to the buyer, who then determines a winner that is allowed to improve utility value for itself as long as the buyer's utility is not lowered. In this setting, the dominant strategy for sellers is to propose an outcome that has a utility value equal to their reservation value in the first round. This represents the negotiation power of the buyer. Sellers need to be aware that they are one out of many other sellers the buyer may choose from.

### 4.1 Multiple negotiation rounds with multiple sellers

Using bilateral negotiation based on a simple alternating offers protocol between the buyer and each of the sellers does not by itself take the fact into account that multiple sellers are contending for a deal with the seller. To incorporate this fact, we introduce a protocol that consists of multiple rounds of (parallel) bilateral negotiations between the buyer and the sellers. After each round  $r$ , the buyer communicates the *winning agreement* 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 the last round 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, and that he thinks has a higher utility to the 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 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 and the setup we used in our experiments are illustrated in Figure 1.

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 can be used to assess if an offer



**Fig. 1.** 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.

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 1 also illustrates that the size of the negotiation space is decreased in every next round. This is explained by the fact that the buyer will only accept offers that improve the winning agreement reached in previous round. 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:

**Hypothesis 1** *The agreement reached using the proposed negotiation protocol converges to that of the theoretical 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 service providers. The process of reaching such an agreement is more complicated than that of the Vickrey Auction but does not require the complicating step to publicly announce 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. Instead of making the buyer's preferences public, in this case 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.

#### 4.2 A variant without making intermediate agreements public

The same 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, we do require 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 they can check whether they believe they could have improved this final agreement to obtain a deal.

Consequently, in this second variant the sellers have less information on how to outbid 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 the variant.

**Hypothesis 2** *The agreement reached without revealing the winning agreement in each round converges to that of the theoretical 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 second 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 second variant compared to that needed in the first:

**Hypothesis 3** *On average more rounds will be needed to reach a final winning agreement using the second setup than the first.*

Both setups should work with any kind of negotiation strategy. However, because the sellers have to reach an agreement after all it is most likely that they would use some kind of a concession-based negotiating tactic. This makes it also possible to learn the preferences of the opponent during the negotiation. In the experiments the buyer therefore maintains a separate opponent model for every seller. The model is updated when the buyer negotiates with the corresponding seller. The seller has to maintain only a model of the buyer’s preferences.

## 5 Experimental Evaluation


In this section, we first discuss the design of the experiment, following the experimental setup of [4]. We then present the experimental results of the first and second setup and discuss our findings.

### 5.1 Experimental Design

While the mechanism nor the protocol limit the number of sellers, in the experiments we use two sellers with distinct preference profiles. We expect that a higher number of sellers will only obscure errors (e.g., due to learning or noise), because in the end all that matters is the difference between the best and the second-best offer.

In the experiments, we use the so-called Service-Oriented Negotiation (SON) domain [7]. This domain consists of four issues. These issues are supposed to represent various attributes relevant to a service being offered, such as price, quality, time, and penalty. We did use the structure of this domain but did not restrict preferences such that a higher price is always preferred by a seller as would typically be expected. This means that we have more variation in the profiles we used in the experiments than one would typically expect in this domain. As a consequence, our results are more easily generalized to other settings. For the experiments we created a set of 12 preference profiles per role. The following values for each of the parameters (reservation values, weights for each issue, and an evaluation function for each issue) were used:

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 reservation value of each agent is set to either 0.3 or 0.6.



**Fig. 2.** 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.

Buyer							Seller						
Profile	$w_1$	$w_2$	$w_3$	$w_4$	$eval.fn_2$	$v_0$	Profile	$w_1$	$w_2$	$w_3$	$w_4$	$eval.fn_1$	$v_i$
Buyer1	0.25	0.25	0.25	0.25	uphill	0.3	Seller1	0.50	0.30	0.15	0.05	uphill	0.3
Buyer2	0.30	0.50	0.05	0.15	uphill	0.3	Seller2	0.25	0.25	0.25	0.25	uphill	0.3
Buyer3	0.25	0.25	0.25	0.25	downhill	0.3	Seller3	0.50	0.30	0.15	0.05	downhill	0.3
Buyer4	0.30	0.50	0.05	0.15	downhill	0.3	Seller4	0.25	0.25	0.25	0.25	downhill	0.3
Buyer5	0.25	0.25	0.25	0.25	triangle	0.3	Seller5	0.50	0.30	0.15	0.05	triangle	0.3
Buyer6	0.30	0.50	0.05	0.15	triangle	0.3	Seller6	0.25	0.25	0.25	0.25	triangle	0.3
Buyer7	0.25	0.25	0.25	0.25	uphill	0.6	Seller7	0.50	0.30	0.15	0.05	uphill	0.6
Buyer8	0.30	0.50	0.05	0.15	uphill	0.6	Seller8	0.25	0.25	0.25	0.25	uphill	0.6
Buyer9	0.25	0.25	0.25	0.25	downhill	0.6	Seller9	0.50	0.30	0.15	0.05	downhill	0.6
Buyer10	0.30	0.50	0.05	0.15	downhill	0.6	Seller10	0.25	0.25	0.25	0.25	downhill	0.6
Buyer11	0.25	0.25	0.25	0.25	triangle	0.6	Seller11	0.50	0.30	0.15	0.05	triangle	0.6
Buyer12	0.30	0.50	0.05	0.15	triangle	0.6	Seller12	0.25	0.25	0.25	0.25	triangle	0.6

**Table 1.** Predefined buyer and seller profiles.

- The relative value associated with each of alternatives for an issue is represented by a linear “uphill” function, a linear “downhill” function, or a combination of the two (resulting in a triangular shape). That is, three types of evaluation functions were used to assign relative value to issue alternatives.

Figure 2 shows an example of a preference profile for a buyer. 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 summation of the utilities of individual issues.

Table 1 shows the predefined profiles that we created using variations of the three preference profile parameters defined above. First, we vary the reservation value of the preference profiles: 0.3 and 0.6. Second, two weights vectors are used:  $\langle 0.30, 0.50, 0.05, 0.15 \rangle$  and  $\langle 0.25, 0.25, 0.25, 0.25 \rangle$ . Finally, a typical negotiation scenario, such as the SON domain, would assume some level of opposition between the buyer’s and the seller’s preferences. To this end, evaluation functions for the issues 1, 3, and 4 of the buyer’s profiles are fixed 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 total of  $2 * 2 * 3 = 12$  possible profiles.

A sample of 50 different negotiation setups is created by means of random selection of one of the twelve profiles from Table 1 for each role. 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 is generated in a similar way.


In the experiments, the negotiating agents build a model of the opponent preferences, by learning a probability distribution over a set of hypotheses about the utility function of the opponent [3], in this case, these are the evaluation functions and the weights of issues. These structural assumptions about the evaluation functions and weights are made to decrease the number of parameters to be learned and simplify the learning task.

During a negotiation session, the probability of each hypothesis is updated using Bayes’ rule every time a new bid is received from the opponent. This requires a conditional probability that represents the probability that the bid might have been proposed given a hypothesis. Therefore the utility of the bid is calculated according to this hypothesis and compared with the predicted utility according to the assumption about rationality of the opponent’s negotiation tactics. To estimate the predicted utility value an assumption about the opponent concession tactics is used based on a linear function. The linear function is the most plausible model of the negotiation concession tactic used by the opponent.


## 5.2 First setup

The results for the first set of experiments validate our hypothesis that negotiating agents that can learn are able to approximate the outcome determined by the mechanism quite well. The winner predicted by the mechanism and the negotiation setup coincide 100%. Moreover, the outcomes obtained by negotiation are also quite close to those determined by the mechanism. Figure 3 shows the histograms of the differences in the outcomes. The average difference with the theoretical utility of the buyer’s outcomes is 0.01% (with a standard deviation of 1.5%). The utility of the sellers differs from the theoretical utility by -0.37% (with a standard deviation of 1.6%). According to the t-test the difference between the means of the theoretical outcomes utilities and the experimental results are insignificant (for the buyer:  $t = 0.054$ ,  $P(T < t) = 0.957$ , for the seller:  $t = 1.648$ ,  $P(T < t) = 0.106$ ). This confirms Hypothesis 1.

In this setup 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 prefer-



**Fig. 3.** First negotiation setup: histograms of the differences in the utilities of experimental and theoretical outcomes for the buyer (left) and the seller (right).




**Fig. 4.** Second negotiation setup: histograms of the differences in the utilities of experimental and theoretical outcomes for the buyer (left) and the seller (right).

ences 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 buyer to learn the preference profile of the seller, because otherwise acceptable offers that maximize the seller’s utility cannot be found. Second, the utility of the winning buyer in the final agreement is determined by the seller’s ability to learn the buyer’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 theoretical utility can thus be explained by approximation errors in the used learning method.

### 5.3 Second setup

Experimental results of the second negotiation setup show also only a small deviation from the theoretical results of the QVA. Figure 4 shows the histograms of the differences in the outcomes. The average difference with the theoretical utility of the buyer’s outcomes is 1.39% (with a standard deviation of 2.2%). The utility of the sellers differ from the theoretical outcome by -1.28% (with a standard deviation of 2.3%). This confirms our second hypothesis.

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 ( $t = -3.9$ ,  $P(T < t) = 0.00027$ ). This can be explained by the fact that unlike the auction mechanism where the contract always corresponds to the reservation value of the second-best seller,



**Fig. 5.** Relationship between buyer’s outcome deviation (vertical axis) and the seller’s outcome deviation (horizontal axis) for the first (left) and the second (right) negotiation setup.

in the second setup the sellers are not aware of the reservation values of each other. Therefore, deviation of the contract utility depends on 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 Figure 5. There we can see that one negotiating party can benefit from the underperformance of the other.

On average the number of the rounds in the second setup is significantly higher than in the first setup (3.5 against 11.3 respectively,  $t = -9.39$ ,  $P(T < t) = 1.9 \cdot 10^{-12}$ ). Furthermore, on average the second setup required almost two times more offers than in the first setup, that is 50 offers in the second setup against 23 offers in the first one ( $t = -14.4$ ,  $P(T < t) = 4.14 \cdot 10^{-19}$ ). This confirms our third hypothesis. Although this requires significantly more time, the absolute numbers of the rounds and offers in the second setup can still be considered acceptable we believe.

## 6 Related Work

Related work on negotiation mechanisms that deal with multiple players is reported in [8–10]. Our approach differs in at least two regards. First, our aim has been to reach an agreement that is as close as possible to a theoretical agreement as defined 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.

Other work has dealt with this problem only from the buyer’s side. Rahwan et al. [10] and later Nguyen and Jennings [9] 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 [8]. 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.

## 7 Conclusion

In this paper we proposed two setups based on multi-round negotiations that are capable of approximating the Qualitative Vickrey Auction (QVA), obtaining a (near) Pareto-efficient outcome where the best seller wins (in the eyes of the buyer). In both setups the buyer as well as the sellers can use any reasonable negotiation strategy. 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 showed experimentally that both setups converge to the results of the QVA.

The results of the second experiment indicate that even if no information is made public until the end of the negotiation the protocol converges to the results of the QVA. The number of rounds needed to find the winning contract is, however, significantly higher than in the first setup. This can be explained by the fact that sellers have no information about the winning agreement of the previous negotiation round and would have to make several offers before they can outbid the winner.

An interesting direction for future research is to study robustness of the proposed setups for the negotiation strategies used by the sellers and the buyer. We are also interested in potential forms of manipulation of the buyer in the second setup 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. Of course, in such a case the first setup could be used, or, as an intermediate solution, the winners and winning bids of all rounds can be made public after the final agreement, and the buyer should pay all involved a significant penalty upon detection. However, we would like to study whether there are other modifications of the idea presented in this paper that make it

applicable to a broader range of real-world multi-player single-winner multi-issue negotiations over complex domains where preferences cannot completely be made public in advance.

## References

1. Bulow, J., Klemperer, P.: Auctions versus negotiations. *The American Economic Review* **86**(1) (1996) 180–194
2. Harrenstein, P., Mahr, T., de Weerd, M.M.: A qualitative vickrey auction. In Endriss, U., Paul W, G., eds.: *Proceedings of the 2nd International Workshop on Computational Social Choice*, University of Liverpool (2008) 289–301
3. Hindriks, K.V., Tykhonov, D.: Opponent modelling in automated multi-issue negotiation using bayesian learning. In: *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*. (2008) 331–338
4. Hindriks, K.V., Tykhonov, D., de Weerd, M.: Approximating an auction mechanism by multi-issue negotiation. In Hindriks, K.V., Brinkman, W.P., eds.: *Proceedings of the First International Working Conference on Human Factors and Computational Models in Negotiation (HuCom2008)*. (December 2008) 33–38
5. Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. The MIT Press (1994)
6. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* **16**(1) (1961) 8–37
7. Faratin, P., Sierra, C., Jennings, N.R.: Using similarity criteria to make negotiation trade-offs. *Journal of Artificial Intelligence* **142**(2) (2003) 205–237
8. Li, C., Giampapa, J., Sycara, K.: Bilateral negotiation decisions with uncertain dynamic outside options. In: *Proc. of 1st IEEE Int. Workshop on Electronic Contracting*. (2004) 54–61
9. Nguyen, T.D., Jennings, N.: Concurrent bilateral negotiation in agent systems. In: *Proc. of 14th Int. Workshop on Database and Expert Systems Applications*. (2003) 844–849
10. Rahwan, I., Kowalczyk, R., Pham, H.: Intelligent agents for automated one-to-many e-commerce negotiation. In: *Proc. of 25th Australasian Conf. on Computer Science*, Australian Computer Society, Inc., Darlinghurst, Australia (2002) 197–204
11. Gode, D.K., Sunder, S.: Allocative efficiency in markets with zero intelligence (zi) traders: Market as a partial substitute for individual rationality. *Journal of Political Economy* **101**(1) (1993) 119–137