

Strategy-Proof Contract Auctions and the Role of Ties[☆]

Mathijs M. de Weerd^a, Paul Harrenstein^b, Vincent Conitzer^c

^a*Delft University of Technology, Delft, The Netherlands*

^b*University of Oxford, Oxford, UK*

^c*Duke University, Durham, USA*

Abstract

A contract auction establishes a contract between a center and one of the bidders. As contracts may describe many terms, preferences over contracts typically display indifference. The Qualitative Vickrey Auction (QVA) selects the best contract for the winner that is at least as good for the center as any of the contracts offered by the non-winning players. When each bidder can always offer a contract with higher utility for the center at an arbitrarily small loss of her own utility, the QVA is the only mechanism that is individually rational, strategy-proof, selects stable outcomes, and is Pareto-efficient. For general continuous utility functions, a variant of the QVA involving fixed tie-breaking is strategy-proof and also selects stable outcomes. However, there is no mechanism in this setting that in addition also selects Pareto-efficient outcomes.

Key words: contract auction, indifference, social choice function, strategy-proofness, tiebreaking, weakly transferable utility

JEL: C70,D44,D61,D82

1. Introduction

In a standard auction, bidders compete only on the price they pay (or are paid, in the case of a reverse auction). However, in many situations, bidders also compete on other additional attributes, which could be laid down in a contract. In some settings, the bidders may not even compete on price at all, for instance, if the price is fixed in advance. In a *contract auction* one out of many possible contracts is selected between one player, which we call the *center*, and one winner out of a set of other players, called the bidders. A contract describes all terms of the arrangement between the center and the winner, such as the quality of service, deadlines, reputation, shipping, and payment method, but none of these are obligatory elements. For example, the center may be a company putting out a request for proposals, a governmental organization with a fixed budget acquiring a service from one of the available public transport companies, a hospital hiring one out

[☆]This paper gives a characterization of the mechanism presented at the 10th ACM Conference on Electronic Commerce in 2009 [10], and specifically analyzes the role of ties.

Email addresses: M.M.deWeerd^a@tudelft.nl (Mathijs M. de Weerd), paul.harrenstein@cs.ox.ac.uk (Paul Harrenstein), conitzer@cs.duke.edu (Vincent Conitzer)

Preprint submitted to Games and Economic Behavior

January 10, 2011

of many applying doctors, or a home owner requesting a construction bid for a house extension.

In such a contract auction setting we are interested in defining a mechanism that selects an outcome that is *stable*, that is, a contract such that the center cannot get a better deal with another bidder, nor a better deal with the winner without reducing the winner's utility. In particular, we would like to provide bidders with a strategy for this mechanism that is *dominant* (that is, the best the bidder can do irrespective of the other bidders) such that stable outcomes are guaranteed. When the contract consists just of a price for a certain item or service, in the Vickrey or sealed-bid second-price auction truthful bidding is a dominant strategy—and thus the Vickrey auction is *strategy-proof*—and selects the highest bidder as a winner [20]. In this paper we show how to generalize the Vickrey auction to the contract auction setting.

When the set of possible contracts is finite and the preferences of all players over possible contracts do not allow for indifferences, that is, are a linear order, we say the preferences are *strict*. Under this strict preference assumption, a contract auction can be seen as a special case of matching with contracts [13] where on one side there is only one player. This illustrates the powerful generality of the matching-with-contracts framework. However, these assumptions of finiteness and strictness of preferences are restrictive and they severely limit application of the framework. For example, the regular Vickrey auction is not a special case, because the possible payments constitute an infinite outcome space. Also, when there are two or more dimensions to a contract—such as both a payment and a delivery date—one would naturally expect indifferences in the center's preferences, if a lower value in the one dimension can be offset by a higher value in the other one. Rather, the center's preferences in such a situation are often lucidly modeled by indifference curves, where the center is indifferent across all the points that lie on the curve. Below we give an example of an application domain that illustrates such a possibly infinite domain of contracts, and preferences that are not strict.

Example 1. Deloxdu University has received a donation to build a new building to house a new center focusing on the intersection of computer science and economics. The donation specifies a fixed budget and requires that the university solicit bids from different firms for the construction of the building, according to an open and clearly specified process. Many attributes of the building, such as the numbers of offices and classrooms, are specified by the university. This leaves only two attributes for bidders to compete on: the energy usage $\eta(> 0)$ of the building and the amount of time $\mu(> 0)$ before the relevant people can move into the building. The firms cannot compete on price because the donation specifies a hard budget for the building that cannot be reallocated. After ample deliberation in committees, the university decides that its utility function is $u_0(\eta, \mu) = .9^\mu(1 - \eta)$, reflecting a discount factor of .9 and a long-run utility of occupying the building of $1 - \eta$. This utility function is communicated to the firms the university solicits bids from.

Two firms enter the competition. Firm 1 specializes in energy efficient building, but tends to be slow. Its utility function for winning a contract $\omega = (\eta, \mu)$ is $u_1(\eta, \mu) = 1 - 10\mu^{-2}\eta^{-1}$, which takes into account the resulting revenue for the firm. This reflects that taking either μ or η down to zero will take the cost to infinity and μ more quickly so. Firm 2 specializes in speedy construction, but tends to produce energy inefficient buildings. Its utility function for winning a contract is $u_2(\eta, \mu) = 1 - \mu^{-1}\eta^{-2}$. This reflects that, again, taking either μ or η down to zero will take the cost to infinity and,




Figure 1: The situation described in Example 1. The solid curves are indifference curves for the center, the dashed curves are indifference curves for Firm 1, and the dotted curve is an indifference curve for Firm 2.

in this case, η more quickly so.

Naturally, the university would welcome it if each firm were to put forward a proposal that maximizes the university's utility under the constraint that the firm does not have negative utility. Some calculation shows that for Firm 1, the corresponding bid would be $\mu = 6.33, \eta = .250$, resulting in $u_0 = .385$ and $u_1 = 0$. For Firm 2, the corresponding bid would be $\mu = 4.36, \eta = .479$, resulting in $u_0 = .329$ and $u_2 = 0$. The reader is referred to Figure 1 for a graphical representation of the situation. Firm 1 has the more desirable proposal: while it will take longer for the building to be ready to occupy, the gains in energy efficiency more than make up for it.

Clearly, however, asking each firm to put forward the best possible proposal that still gives it nonnegative utility and then requiring the winning proposal to be implemented, is not strategy-proof. In the above example, Firm 1 ends up with a utility of 0. Having won by a comfortable margin, it could submit a weaker proposal instead, increasing either μ or η or both, and still win, thereby obtaining positive utility. This is not surprising, because the above mechanism is a type of first-price auction. Can we create something closer to a Vickrey auction? A first idea might be to only require Firm 1 to implement the second-best proposal, the one submitted by Firm 2. This will not do: it turns out that implementing Firm 2's proposal would actually result in a utility of $-.0981$ for Firm 1. This should not be surprising, because Firm 2's proposal caters to Firm 2's strength (speed), not to Firm 1's (energy efficiency). The right approach is to allow the winning firm to choose *any contract that is at least as good for the university as the second-best proposal*. That is, Firm 1 can choose any contract that gives the university a utility

of at least .329. The contract that maximizes Firm 1's utility under this constraint is $\mu = 7.42, \eta = .281$, for a utility of $u_1 = .353$ (and $u_0 = .329$). Note that Firm 1 has slightly increased both μ and η relative to its original proposal. This is the outcome of the first mechanism that we describe in this paper. Figure 1 shows the bids, the final outcome, and the relevant indifference curves for both the university and the firms.

The main idea underlying the last auction in the example above—having the highest bidder win a contract she can select from a set of contracts determined by the highest bid among the other bidders—forms the basis of the *Qualitative Vickrey Auction (QVA)*, as introduced in previous work [10, 11]. A proposal by Goel et al. [9] relies on a similar conception. The possibility of indifferences in the center's preferences, however, introduces a number of issues, especially with respect to strategy-proofness and Pareto optimality. The auction in Example 1 worked properly because the utility functions had the property that a winning firm can always increase the university's utility by sacrificing an arbitrarily small amount of its own utility. This is because the university's utility has no local maximum when $\mu > 0, \eta > 0$, and firms' utility functions are continuous. If this assumption does not hold, it turns out that one has to be more careful in designing the auction, especially regarding how to break ties. This phenomenon can be illustrated if we slightly adapt Example 1.

Example 1 (continued). The university has discovered an additional option, namely, to renovate an existing building instead. As the firms discover, there are fundamental limitations on what can be achieved with this existing building: for this building, we must have $\mu \geq 4$ (because people are currently occupying it) and $\eta \geq .2$ (due to limitations of existing systems in the building). If indeed $\mu = 4$ and $\eta = .2$, this results in a utility of .525 to the university. Let us suppose that this renovation job is much easier for the firms than new construction and, hence, each of them would obtain a utility of .5 for doing it and achieving $\mu = 4$ and $\eta = .2$. The option of creating a new building still exists. However, in terms of u_0 , the renovation option (that is, where $\mu = 4$ and $\eta = .2$) dominates any new building that the two firms can propose at nonnegative utility to themselves. Therefore, the renovation option is the best proposal that each firm can put forward. Which firm will get the contract? It should be noted that each firm has a strictly positive utility for getting it and this is where we need to be more cautious in designing the mechanism. In particular, if the tie is broken against a firm with positive probability, then if we are not sufficiently careful, this firm may engage in the following manipulation: put forward a proposal for a *new* building with, say, $\mu = 3$ and $\eta = .1$. This proposal is more attractive for the university. While implementing it would give the firm negative utility, if the firm is only required to implement something as good as the next best proposal, then the firm could in fact just implement the renovation option. It would increase its utility by this manipulation because originally it did not (always) win the renovation contract.

Our contributions concern two related contract auction mechanisms. First, we define the concept of weakly transferable utility, describing that for any bidder it is always possible to bid slightly higher at an arbitrarily small loss of utility. We show that under this condition the QVA is strategy-proof for the bidders and selects stable and Pareto efficient outcomes. Moreover, we find that all mechanisms that are strategy-proof for the bidders and select stable outcomes are equivalent to the QVA. We also show that the condition of weakly transferable utility is necessary for strategy-proofness. Second, for

general situations with continuous utility functions but without a condition on weakly transferable utility, we define the QVA^+ , where tie-breaking is fixed and the winner may select from the closure of all outcomes that would have made her a winner. We show that the QVA^+ is strategy-proof for the bidders and moreover selects stable outcomes. Furthermore, we find that there is no mechanism with these properties that is also Pareto efficient.

The setup of this paper is as follows. First we introduce notation and definitions in Section 2. Then, in Section 3 we introduce a condition called *weakly transferable utility* and characterize the QVA in settings in which this condition holds. Relaxing this condition, we introduce and analyze the QVA^+ in Section 4, followed by a discussion of related work in Section 5, and conclusions and future work in Section 6.

2. Preliminaries

Players and outcomes. In a contract auction we consider a finite set of players denoted by $\{0, 1, \dots, n\}$ with typical elements i, j . Player 0 is called the *center* and the players in the set $N = \{1, 2, \dots, n\}$ the *bidders*. The outcome of a contract auction is a contract between the center and one of the bidders, whom we call the *winner*. Let C_i be a set of contracts bidder i can offer. We have $\Omega_i = C_i \times \{i\}$ denote the set of (possible) outcomes ω_i in which i wins the auction. Thus, $\Omega_i \cap \Omega_j = \emptyset$, whenever $i \neq j$. The total set of outcomes is given by $\Omega = \Omega_1 \cup \dots \cup \Omega_n$. We assume each Ω_i to constitute a topological space. Moreover, for every $\Omega'_i \subseteq \Omega_i$, denote by $\text{cl}(\Omega'_i)$ the *closure* of Ω'_i , that is, the smallest closed superset of Ω'_i .

Preferences. We assume that the preferences of the center 0 over Ω are given by a *utility function* $u_0: \Omega \rightarrow \mathbb{R}$ such that the restriction of u_0 to each Ω_i is a continuous function. Similarly, the preferences of each player i over Ω are represented by a *utility function* $u_i: \Omega \rightarrow \mathbb{R}$ such that the restriction of u_i to Ω_i is a continuous function and $u_i(\omega) = 0$ for every $\omega \in \Omega \setminus \Omega_i$.¹ Thus, bidder i is indifferent between any two outcomes in which she is not the winner (a no-externalities assumption). A *utility profile* is a vector $u = (u_0, u_1, \dots, u_n)$ of utility functions, one for each player.

For player i and a set of outcomes $\Omega' \subseteq \Omega$ we write $\text{opt}_i(\Omega')$ to denote the *set of optimal outcomes for i in Ω'* , that is,

$$\begin{aligned} \text{opt}_i(\Omega') &= \underset{\omega \in \Omega'}{\text{argmax}} u_i(\omega) \\ &= \{\omega \in \Omega' : u_i(\omega) \geq u_i(\omega') \text{ for all } \omega' \in \Omega'\}. \end{aligned}$$

Given a utility profile, we say an outcome ω is *acceptable* for a bidder i if ω is not worse for i than losing the auction, that is, if $u_i(\omega) \geq 0$, and *unacceptable* otherwise. Formally, given a utility profile u , the acceptable outcomes for a bidder i in Ω_i are defined by

$$\Omega_i^+(u) = \{\omega \in \Omega_i : u_i(\omega) \geq 0\}.$$

¹The assumption of continuous utility functions is void on a discrete outcome space, since every function on a discrete space is continuous.

Throughout the paper, we assume the bidders' preferences to be *satisfiable* in the sense that for each bidder the set of acceptable outcomes is non-empty. Moreover, to guarantee the existence of maximal outcomes in certain subsets of outcomes, we make a *weak compactness* assumption and require that every sequence of acceptable outcomes in Ω_i for a bidder i with increasing utility for the center or with increasing utility for i must have a limit point within the acceptable outcomes.

Given a utility profile $u = (u_0, u_1, \dots, u_n)$, there is a number of sets of outcomes that play an important role in the mechanisms considered in this paper. A key concept is that of the set of outcomes for a bidder that maximize the center's utility under the constraint of being acceptable to that bidder. Thus, we define the set of *highest acceptable bids* as follows:

$$\overline{\Omega}_i(u) = \text{opt}_0(\Omega_i^+(u)).$$

We observe that for all bidders i we have $u_0(\omega_i) = u_0(\omega'_i)$ for all $\omega_i, \omega'_i \in \overline{\Omega}_i(u)$. Because bidders' preferences are satisfiable by assumption, by the weak compactness condition $\overline{\Omega}_i(u)$ is non-empty for all utility profiles u . As an auxiliary notion we define the set $N^*(u)$ of *potential winners* given a utility profile u . A bidder i is a potential winner if she has an acceptable bid that is ranked at least as high by the center as any other acceptable bid by any other bidder. Accordingly, we define

$$N^*(u) = \{i \in N : u_0(\omega_i) \geq u_0(\omega_j) \text{ for all bidders } j \neq i, \omega_i \in \overline{\Omega}_i(u), \text{ and } \omega_j \in \overline{\Omega}_j(u)\}.$$

We observe that the set $N^*(u)$ of potential winners is non-empty for every utility profile u , again by virtue of the weak compactness condition. To illustrate the above concepts, we include the following example.

Example 2. Consider the setting in which, apart from the center 0, there are two bidders, i and j . Let $\Omega_i = \mathbb{R} \times \{i\}$ and $\Omega_j = \mathbb{R} \times \{j\}$. Moreover, let the utility profile $u = (u_0, u_i, u_j)$ be such that for all $x \in \mathbb{R}$ and $k \in \{i, j\}$,

$$\begin{aligned} u_0(x, k) &= \frac{1}{2}x - 1 \\ u_i(x, k) &= \begin{cases} 3 - \frac{1}{3}x & \text{if } k = i, \\ 0 & \text{otherwise,} \end{cases} \\ u_j(x, k) &= \begin{cases} 4 - x & \text{if } k = j, \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

The situation is also depicted in Figure 2. Then,

$$\Omega_i^+(u) = \{(x, i) : x \leq 9\} \quad \text{and} \quad \Omega_j^+(u) = \{(x, j) : x \leq 4\}.$$

Moreover, $\overline{\Omega}_i(u) = \{(9, i)\}$ and $\overline{\Omega}_j(u) = \{(4, j)\}$. As $9 > 4$, also $u_0((9, i)) > u_0((4, j))$. Hence, $N^*(u) = \{i\}$.

Ties and tie-breaking. We will generally use the word *ties* to refer to indifferences in the center's preferences, that is, outcomes ω and ω' are said to be tied if $u_0(\omega) = u_0(\omega')$. The reasons for this are twofold. Firstly, we wish to avoid confusion with indifferences




Figure 2: Example illustrating the notations. Here $\bar{\omega}_i = (9, i)$ and $\bar{\omega}_j = (4, j)$ and, thus, $\bar{\Omega}_i(u) = \{(9, i)\}$ and $\bar{\Omega}_j(u) = \{(4, j)\}$. The outcome in the darker gray area are unacceptable for bidder i . The outcomes that lie either in the lighter or the darker gray area are unacceptable to bidder j .

in the bidders' preferences. Secondly, the center is not a strategic player in this context; its preferences serve only to define the mechanism. Specifically, its preferences define an objective function that the mechanism seeks to maximize, and if multiple outcomes achieve the maximum objective value, the mechanism must specify how to break this tie. Therefore, the center will not make any decisions once the mechanism is defined, so that at that point it is no longer strictly necessary to think of u_0 as reflecting preferences (and, *a fortiori*, indifferences).

In its most general form, by a *winner determination (including tie-breaking) rule* we understand a function t that associates each utility profile u with a bidder $t(u)$. Numerous restrictions on tie-breaking rules are of course possible. In this paper, we assume that tie-breaking rules break ties among the highest bidders, that is, for all utility profiles u we have $t(u) \in N^*(u)$.

Mechanisms. By a *social choice function* or a *direct mechanism* we understand a function f that maps each utility profile $u = (u_0, u_1, \dots, u_n)$ to an outcome $f(u)$ in Ω . A social choice function f is *individually rational* if for all utility profiles u and all bidders i , $f(u)$ is acceptable, that is, $u_i(f(u)) \geq 0$. Moreover, f is *strategy-proof* if for all bidders $i > 0$ (note that the center is not treated as a strategic agent) and all utility profiles u and u' ,

$$u_i(f(u)) \geq u_i(f(u_0, u_1, \dots, u_{i-1}, u'_i, u_{i+1}, \dots, u_n)).$$

We then also say that u_i is a *dominant strategy* for each bidder i . By an *indirect (revelation) mechanism* we understand a game form $G = (\{1, \dots, n\}, A_1, \dots, A_n, h)$, where A_i is a set of actions for bidder i and $h: A_1 \times \dots \times A_n \rightarrow \Omega$ is an *outcome function*.

(Again, the center is not treated as a strategic bidder.) Given a utility profile $u = (u_1, \dots, u_n)$, an action a_i for bidder i in an indirect mechanism G is said to be *dominant* if for all actions b_1, \dots, b_n it is the case that

$$u_i(h(b_1, \dots, b_{i-1}, a_i, b_{i+1}, \dots, b_n)) \geq u_i(h(b_1, \dots, b_{i-1}, b_i, b_{i+1}, \dots, b_n)).$$

In a slightly informal fashion, we say that an indirect mechanism is *strategy-proof* if an appropriate concept of *straightforward strategies* can be defined such that straightforward strategies are both dominant and guaranteed to exist.

An outcome ω is said to be *Pareto efficient* if there is no other outcome ω' such that $u_i(\omega') \geq u_i(\omega)$ for all players $i \in \{0, 1, \dots, n\}$ and $u_j(\omega') > u_j(\omega)$ for at least one player $j \in \{0, 1, \dots, n\}$. A mechanism f is said to be Pareto efficient if it generally yields Pareto efficient outcomes, that is, if $f(u)$ is Pareto efficient for all utility profiles u .

Finally, we say an outcome ω is (strongly) *stable* if it is acceptable to the winner and there is no bidder i and no outcome $\omega' \in \Omega_i$ such that both

$$(i) \ u_0(\omega') > u_0(\omega), \text{ and}$$

$$(ii) \ u_i(\omega') \geq u_i(\omega),$$

that is, if the outcome is acceptable to the winner, and the center cannot get a better deal with another bidder, nor a better deal with the winner without reducing the winner's utility. From the acceptability of stable outcomes, it follows directly that any mechanism that always selects stable outcomes is individually rational for the bidders.

The above definition of stability in contract auctions relates to the notion of a (strongly) stable matching M , which is a matching in which there is no couple (i, j) such that j strictly prefers i to his/her partner in M , and i either strictly prefers j to her/his partner in M or is indifferent between them [15, 17]. In our definition we fix the matching to a single contract and one of the partners to the center, and we extend the definition with the possibility to select such a contract between the center and a bidder from a set of possibilities over which both have a preference order. As in matching theory, it is possible to define a weaker notion of stability where the inequalities are all strict. Naturally, the positive results in this paper all hold for this weaker notion as well.

The reader is referred to Table 1 for a summary overview of the notations used in this paper.

3. Contract Auctions with Weakly Transferable Utility

We first consider a specific generalization of the Vickrey auction, called the Qualitative Vickrey Auction. We show when it is strategy-proof and selects stable outcomes. Then we give a characterization of the complete set of social choice functions with these properties.

3.1. A Strategy-Proof Contract Auction

The *Qualitative Vickrey Auction (QVA)* is an indirect mechanism in which each bidder i with utility function u_i submits a sealed bid $\hat{\omega}_i \in \Omega_i$ to the center and was proposed by Harrenstein et al. [10]. For each bidder i and each profile $\hat{\omega} = (\hat{\omega}_1, \dots, \hat{\omega}_n)$

symbol	meaning	definition
0	center	page 5
$1, \dots, n$	bidders	page 5
ω_i	outcome where i wins	page 5
Ω_i	set of outcomes with i as winner	page 5
Ω	set of outcomes	$\Omega_1 \cup \dots \cup \Omega_n$
$u_i : \Omega \rightarrow \mathbb{R}$	player i 's utility function	page 5
$\text{opt}_i(\Omega')$	outcomes in Ω' optimal for i	$\text{argmax}_{\omega \in \Omega'} u_i(\omega)$
$\Omega_i^+(u)$	outcomes acceptable to i	$\{\omega \in \Omega_i : u_i(\omega) \geq 0\}$
$\bar{\Omega}_i(u)$	highest acceptable bids of i	$\text{opt}_0(\Omega_i^+(u))$
$N^*(u)$	potential winners	page 6
$t(u)$	winner-determination (incl. tie-breaking) rule	page 7
$t(\hat{\omega})$	winner-determination (incl. tie-breaking) order	page 18
f	social choice function	page 7
$\Omega_i^*(\hat{\omega})$	potentially winning outcomes for i (QVA-i)	page 9
$\Omega_i^*(u)$	potentially winning outcomes for i (QVA-d)	page 13
$\Omega_i^*(\hat{\omega}, t)$	potentially winning outcomes for i (indirect QVA ⁺)	page 18
$\Omega_i^*(u, t)$	potentially winning outcomes for i (direct QVA ⁺)	page 17

Table 1: Overview of notations

of bids define $\Omega_i^*(\hat{\omega})$ as the set of outcomes in Ω_i that yield the center at least the utility of any other bid $\hat{\omega}_j$ submitted by some bidder j distinct from i , that is,

$$\Omega_i^*(\hat{\omega}) = \{\omega_i \in \Omega_i : u_0(\omega_i) \geq u_0(\hat{\omega}_j) \text{ for all } j \neq i\}.$$

These outcomes we refer to as the *potentially winning outcomes for i* in the QVA. The QVA then consists of the following steps.

- (1) First, the utility function u_0 of the center is announced to all bidders.
- (2) Then, each bidder i submits a sealed bid $\hat{\omega}_i \in \Omega_i$ to the center.
- (3) The bidder i^* who submitted the bid with the highest utility for the center is declared the winner of the auction, that is, i^* is selected from $\arg \max_{i \in N} u_0(\hat{\omega}_i)$. Ties are broken arbitrarily.
- (4) Finally, the winner i^* may choose from among her outcomes in Ω_{i^*} any outcome that yields the center a utility that is at least as high as the center's utility for the next-highest bid, that is, i^* may choose from $\Omega_{i^*}^*(\hat{\omega}_1, \dots, \hat{\omega}_n)$. The outcome she chooses is the outcome of the auction.

The QVA is guaranteed to yield an outcome. There being only a finite number of bidders, bids with a highest utility for the center always exist and a winner i^* can be selected. Moreover, bidder i^* 's initial offer $\hat{\omega}_{i^*}$ is witness to the fact that $\Omega_{i^*}^*(\hat{\omega}_1, \dots, \hat{\omega}_n)$

is non-empty. Furthermore, if the winner's bid $\hat{\omega}_{i^*}$ is acceptable to her, then it is also witness to the fact that an acceptable outcome of the QVA exists.

The intuition why the classical Vickrey or second-price auction [20] is strategy-proof—that is, bidding truthfully is a dominant strategy—is that a bidder's monetary bid only determines whether she turns out to be the winner, but not what price she has to pay if she does. The situation is similar in the qualitative Vickrey auction. Again, a bidder's bid determines whether she emerges as the winner, but the set of alternatives from among which she may choose is decided by the second-highest bid.

A strategy for a bidder i in the QVA specifies the bid $\omega_i \in \Omega_i$ to make, along with a contingency plan prescribing which outcome to choose from among the outcomes in Ω_i that have at least as high a utility for the center as the second-highest offer submitted, in case i happens to win the auction. Of course, bidder i 's strategy may depend on her utility function u_i .

We call a strategy for i *straightforward* if it satisfies the following properties:

- (i) the bid i submits is an outcome $\hat{\omega}_i$ with the highest utility for the center among those that are acceptable to i , that is, $\hat{\omega}_i \in \bar{\Omega}_i(u)$,
- (ii) in case i wins the auction, she selects one of the outcomes ω_i^* in Ω_i she values most among those that have a center's utility at least as high as the second-highest bid submitted, that is,

$$\omega_i^* \in \text{opt}_i(\Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n)),$$

where, for bidders j distinct from i , $\hat{\omega}_j$ denotes the actual bid submitted. If there are multiple such outcomes in this set, equally valued by i^* , she selects one with the highest utility for the center.

Example 2 (continued). Consider the situation depicted in Figure 2. If both bidders play straightforward strategies, bidder i bids outcome $\bar{\omega}_i = (9, i)$ and bidder j outcome $\bar{\omega}_j = (4, j)$. Moreover, bidder i then emerges as the winner and she will subsequently choose an outcome from $\text{opt}_i(\{(x, i) : x \geq 4\})$, that is, outcome $(4, i)$. This yields a utility of $1\frac{2}{3}$ to i and a utility of 1 to the center.

We say that the auction is *strategy-proof* if such straightforward strategies exist and all straightforward strategies are dominant strategies.

The QVA may not be strategy-proof when a bidder has a strictly positive utility for her highest acceptable bid. In case she is one of the potential winners, but the tie is broken against her, she may submit a bid that is more preferred by the center, but unacceptable to her. In this way she will win the auction, but still be able to choose her original bid. However, under a mild condition we call weakly transferable utility we can generalize the Vickrey auction in a particularly natural way.

We say that a domain with a utility profile $u = (u_0, u_1, \dots, u_n)$ has *weakly transferable utility* if for every bidder i , every outcome $\omega \in \Omega_i$, and every $\epsilon > 0$, there is an $\omega' \in \Omega_i$ such that both

- (i) $u_0(\omega') > u_0(\omega)$, and
- (ii) $u_i(\omega') \geq u_i(\omega) - \epsilon$.

Intuitively, this condition says that, given an outcome $\omega_i \in \Omega_i$, bidder i can always please the center slightly more with another outcome in Ω_i at an arbitrarily small loss of utility to herself. This condition generalizes the common concept of transferable utility as well as that of the more general *quasi*-transferable utility, which occurs in domains where there is “a good which is divisible, always desirable and a substitute for any other good” [3, page 237]. With the condition of weakly transferable utility we include other circumstances, in which one player can accommodate another, such as finishing a project a bit earlier. Note that with weakly transferable utility we introduce an even weaker variant, where this transfer only has to be possible from the bidders to the center.

For utility *not* to be weakly transferable, there must be an outcome which, intuitively, is locally optimal for both the center as well as the bidder. Then, there is no way the bidder can concede so as to increase the center’s utility. An example of this phenomenon can be found in Figure 4 below.

The following lemma, which is crucial for the proofs of Theorem 1 and Theorem 2, says that under weakly transferable utility, the highest bid which is acceptable to a bidder has 0 utility for her.

Lemma 1. *If utility is weakly transferable, then $u_i(\omega_i) = 0$ for every highest acceptable bid $\omega_i \in \bar{\Omega}_i(u)$.*

Proof. Assume $u_i(\omega_i) \neq 0$. As ω_i is acceptable, then $u_i(\omega_i) > 0$. Then, let $\epsilon = u_i(\omega_i)$. Since utility is weakly transferable, we may conclude that there exists an ω'_i with $u_i(\omega'_i) \geq u_i(\omega_i) - \epsilon = 0$ and $u_0(\omega'_i) > u_0(\omega_i)$. Then ω'_i is acceptable for i , contradicting that ω_i is a highest acceptable bid, that is, that $\omega_i \in \bar{\Omega}_i(u)$. \square

We are now in a position to prove Theorem 1, our first main result.

Theorem 1. *If utility is weakly transferable, then the Qualitative Vickrey Auction is strategy-proof and selects stable and Pareto efficient outcomes.*

Proof. It suffices to show that for every bidder i a straightforward strategy as defined above exists and that these straightforward strategies never result in negative utility for i , are dominant, and result in stable and Pareto efficient outcomes. To this end, let $u = (u_0, u_1, \dots, u_n)$ be an arbitrary utility profile.

To show the existence of a straightforward strategy for each bidder, we first note that $\bar{\Omega}_i(u)$ is nonempty by the weak compactness condition, so the bidder is in fact able to bid according to a straightforward strategy. Next, let $\hat{\omega}_j$ denote the bid submitted by bidder j in the second stage of the auction. We argue that $\text{opt}_i(\{\omega_i \in \Omega_i : u_0(\omega_i) \geq u_0(\hat{\omega}_j) \text{ for all } j \neq i\})$ is non-empty if i emerges as the winner in the third stage of the auction. So assume that $u_0(\hat{\omega}_i) \geq u_0(\hat{\omega}_j)$ for all bidders j distinct from i (and that i wins the tie if there is one). It follows that there exists an infinite sequence (possibly with repetitions) of outcomes $\omega_i \in \Omega_i$ such that $u_0(\omega_i) \geq u_0(\hat{\omega}_j)$ for all $j \neq i$, and in which bidder i ’s utility is nondecreasing and converges either to infinity or to the least upper bound on utilities that bidder i can obtain in this set. By the weak compactness assumption, this sequence must have a limit point ω_i^∞ . By continuity of the center’s utility function on Ω_i , we have $u_0(\omega_i^\infty) \geq u_0(\hat{\omega}_j)$ for all $j \neq i$. By the continuity of the bidder’s utility function on Ω_i , outcome ω_i^∞ must achieve the least upper bound on utilities that bidder i can obtain in $\Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n)$, that is, $\omega_i^\infty \in \text{opt}_i(\Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n))$. This establishes the existence of straightforward strategies.

For strategy-proofness, we show that every straightforward strategy is also dominant. We first observe that by changing her bid, a bidder i can only affect whether she wins or loses, but cannot improve her utility when she wins, because if she wins, a straightforward strategy already selects an optimal outcome for her from a set she cannot change. Suppose that a bidder i would win with her straightforward bid; by individual rationality, she gets utility at least 0, so she has no incentive to change her bid to lose, which would give her utility 0.

Now suppose that a bidder i would lose with her straightforward bid. We distinguish two cases.

If bidder i lost because her highest acceptable bid has a strictly lower utility to the center than the bid of another bidder, then, if she changed her bid to win, she would have to choose an outcome at least as desirable to the center as that other bid, which must be unacceptable to her and, hence, result in negative utility. So, she would prefer to lose.

If bidder i lost a tie with her highest acceptable bid ω_i , there is another bidder $j \neq i$ with $u_0(\omega_j) = u_0(\omega_i)$. If she changes her bid to win, then, after winning, she has to choose an outcome at least as desirable to the center as ω_j , or, equivalently, at least as desirable to the center as ω_i . Hence, she must choose one of her highest acceptable bids (or an unacceptable bid, but that would be worse). By Lemma 1, however, this will yield her utility 0 and she will not be better off. We may conclude that the mechanism is strategy-proof.

To prove that the mechanism chooses stable outcomes, we first show that no straightforward strategy ever results in negative utility to the winner. When using a straightforward strategy, a bidder i will only bid acceptable outcomes in $\Omega_i^+(u)$, that is, $\hat{\omega}_i \in \Omega_i^+(u)$. If the bidder wins, then she can at the very least choose the outcome that she bid, that is, $\hat{\omega}_i$, which must give her nonnegative utility. If bidder i loses the auction her utility will be 0 and thus be nonnegative as well. Then observe that it is not possible to make the center better off with the same bidder without making that bidder worse off, because among the outcomes that give the bidder maximal utility, ties are broken in favor of the center (in step (ii) of the straightforward strategy). Moreover, the outcomes from which this selection is made are all at least as good for the center as any acceptable outcome for any other bidder. Hence, no acceptable outcome for any other bidder would make the center better off.

To further conclude that any outcome ω chosen by the mechanism is also Pareto efficient, we show, in addition to stability, that no outcome ω' exists that gives one bidder i a strictly higher utility without reducing the center's utility. In case there is just one potential winner i , outcome ω is already selected to be optimal for i , and since i is a strict winner, all acceptable bids by all other bidders are strictly worse for the center than ω . In case of a tie among multiple bidders, for any one of these bidders we would have to choose one of her highest acceptable outcomes to keep the center from being worse off. By Lemma 1, however, such an outcome would yield the bidder utility 0. \square

3.2. Characterizing the Qualitative Vickrey Auction

We now set out to prove that the Qualitative Vickrey Auction is, in an important sense, the unique mechanism satisfying the desirable properties it has according to Theorem 1 in this context. To do so, we first appeal to the revelation principle and convert

the QVA to a direct-revelation mechanism. When it is helpful to distinguish, we refer to the indirect version described above as QVA-i, and the direct version that we are about to describe as QVA-d. In fact, there are really multiple versions of QVA-d, depending on the tie-breaking rule and the concept of straightforward strategy chosen for the QVA-i. tie-breaking rule chosen and, in the second place, because the version of QVA-d that results from applying the revelation principle depends on the precise dominant-strategies solution (that is, which straightforward strategies are used) of the QVA-i mechanism from which we start. We therefore define a class of QVA-d mechanisms.

To this end, we have to adapt the concept of potentially winning outcomes given utility profiles rather than bidding profiles. Given a utility profile $u = (u_0, u_1, \dots, u_n)$, let the set of *potentially winning outcomes* $\Omega_i^*(u)$ for a bidder contain all outcomes in Ω_i with a utility for the center that is as least as high as the highest acceptable outcome of any other bidder, that is,

$$\Omega_i^*(u) = \{\omega_i \in \Omega_i : u_0(\omega_i) \geq u_0(\omega'_j) \text{ for all bidders } j \neq i \text{ and all } \omega'_j \in \bar{\Omega}_j(u)\}.$$

In the setting of Example 2, we would thus have $\Omega_i^*(u) = \{(x, i) : x \geq 4\}$ and $\Omega_j^*(u) = \{(x, j) : x \geq 9\}$ (also see Figure 2).

We say that a social choice function f is a *QVA-d mechanism* if for all utility profiles $u = (u_0, u_1, \dots, u_n)$ the outcome $f(u)$ is always for a potential winner $i \in N^*(u)$ and among the best optimal potentially winning outcomes for i from the ones that are optimal for the center, that is, if

$$f(u) \in \bigcup_{i \in N^*(u)} \text{opt}_0(\text{opt}_i(\Omega_i^*(u))).$$

Via the revelation principle, QVA-d mechanisms correspond exactly to combinations of a QVA-i mechanism together with a profile of straightforward strategies for it. From this, we obtain the following corollary.

Corollary 1. *If utility is weakly transferable, then every QVA-d mechanism is strategy-proof, and selects stable and Pareto efficient outcomes.*

We now set out to prove that every direct-revelation mechanism with the desired properties—even without requiring Pareto efficiency—must be a QVA-d mechanism. First, we show that given a utility profile u , for any social choice function that chooses stable outcomes, these outcomes must be for a potential winner $i \in N^*(u)$, and acceptable (thus in $\Omega_i^+(u)$) as well as potentially winning (in $\Omega_i^*(u)$, that is, with a utility for the center at least as high as the highest acceptable outcome of any other bidder).

Lemma 2. *If a social choice function f chooses stable outcomes, then for every utility profile u , $f(u) \in \Omega_i^+(u) \cap \Omega_i^*(u)$ for some bidder $i \in N^*(u)$.*

Proof. Consider an arbitrary utility profile $u = (u_0, u_1, \dots, u_n)$ along with an arbitrary highest acceptable bid $\omega_i \in \bar{\Omega}_i(u)$ for the winner i . As $f(u)$ is acceptable for i —and thus immediately $f(u) \in \Omega_i^+(u)$ by virtue of stability (which implies individual rationality)—, we have $u_0(\omega_i) \geq u_0(f(u))$. Also, consider an arbitrary bidder j distinct from i and an arbitrary highest acceptable bid $\omega_j \in \bar{\Omega}_j(u)$. Then, by satisfiability of preferences, $u_j(\omega_j) \geq 0 = u_j(f(u))$. As f chooses stable outcomes, it follows that $u_0(f(u)) \geq u_0(\omega_j)$. Hence, $u_0(\omega_i) \geq u_0(f(u)) \geq u_0(\omega_j)$ and we may conclude that both $f(u) \in \Omega_i^*(u)$ and $i \in N^*(u)$. \square

We now present our characterization result.

Theorem 2. *Given a utility profile u and a social choice function f , if utility is weakly transferable, and f is strategy-proof for the bidders and chooses stable outcomes, then f is a QVA-d mechanism.*

Proof. By Lemma 2, we know that $f(u) \in \Omega_i^+(u) \cap \Omega_i^*(u)$ for some bidder $i \in N^*(u)$. We first show that $f(u) \in \text{opt}_i(\Omega_i^*(u))$ for the winner $i \in N^*(u)$ both with and without a tie, and then conclude by showing that $f(u) \in \text{opt}_0(\text{opt}_i(\Omega_i^*(u)))$ and, thus, that f is a QVA-d mechanism.

In the case where there is a tie, that is, $|N^*(u)| > 1$, we show that $f(u) \in \text{opt}_i(\Omega_i^*(u))$ by proving that $\text{opt}_i(\Omega_i^*(u)) = \Omega_i^+(u) \cap \Omega_i^*(u)$ and then applying Lemma 2 again. First, observe that for the winner $i \in N^*(u)$ it holds that its potentially winning outcomes are acceptable, and thus $\text{opt}_i(\Omega_i^*(u)) \subseteq \Omega_i^+(u) \cap \Omega_i^*(u)$. Moreover, $\Omega_i^+(u) \cap \Omega_i^*(u) \subseteq \text{opt}_i(\Omega_i^*(u))$, because the only outcomes in $\Omega_i^+(u) \cap \Omega_i^*(u)$ are the highest acceptable bids for i , that is, the bids in $\bar{\Omega}_i(u)$, for which, by Lemma 1, we know i has utility 0. Hence, $\text{opt}_i(\Omega_i^*(u)) = \Omega_i^+(u) \cap \Omega_i^*(u)$, and thus with Lemma 2 we have $f(u) \in \text{opt}_i(\Omega_i^*(u))$.

In the case where there is no tie, we show by contradiction that the outcome should be optimal for i among the potentially winning outcomes, that is, $f(u) \in \text{opt}_i(\Omega_i^*(u))$, because otherwise i can report an alternative utility function \hat{u}_i such that $u_i(f(\hat{u})) > u_i(f(u))$ (also see Figure 3). Suppose that $f(u) \notin \text{opt}_i(\Omega_i^*(u))$. Let ω_i^* be any outcome in $\text{opt}_i(\Omega_i^*(u))$. Informally, we define a (shifted) utility function \hat{u}_i below that renders $f(u)$ unacceptable for i . This only leaves acceptable outcomes for i with a higher utility for i than $f(u)$. We then use weakly transferable utility to show that there exists an acceptable outcome ω' ‘in between’ ω_i^* and $f(u)$ that is strictly better for the center than ω_i^* . This leaves i as the only potential winner.

Formally, let δ be the difference between i ’s utility for ω_i^* and $f(u)$, that is, $\delta = u_i(\omega_i^*) - u_i(f(u)) > 0$. We now define a utility function \hat{u}_i that is similar to u_i , but shifted downward by $u_i(f(u)) + \epsilon$ for some $0 < \epsilon < \delta$ such that $f(u)$ becomes unacceptable, that is, for all $\omega_i \in \Omega_i$,

$$\hat{u}_i(\omega_i) = u_i(\omega_i) - u_i(f(u)) - \epsilon.$$

Let $\hat{u} = (u_0, u_1, \dots, \hat{u}_i, \dots, u_n)$. Then, because utility is weakly transferable,² there is an outcome $\omega'_i \in \Omega_i$ such that $u_0(\omega'_i) > u_0(\omega_i^*)$ and $u_i(\omega'_i) \geq u_i(\omega_i^*) - \epsilon'$, where $\epsilon' < \delta - \epsilon$. This outcome ω'_i is still acceptable for bidder i under utility function \hat{u}_i , because

$$\begin{aligned} \hat{u}_i(\omega'_i) &= u_i(\omega'_i) - u_i(f(u)) - \epsilon \\ &\geq u_i(\omega_i^*) - \epsilon' - u_i(f(u)) - \epsilon \\ &= \delta - \epsilon' - \epsilon \\ &> 0. \end{aligned}$$

Moreover, since $\omega_i^* \in \Omega_i^*(u)$, we know that for all bidders $j \neq i$ and all $\omega_j \in \bar{\Omega}_j(u)$ it holds that $u_0(\omega_i^*) \geq u_0(\omega_j)$. Since $\bar{\Omega}_j(\hat{u}) = \bar{\Omega}_j(u)$ and $u_0(\omega'_i) > u_0(\omega_i^*)$, we then conclude that for all bidders $j \neq i$ and all $\omega_j \in \bar{\Omega}_j(\hat{u})$ it holds that $u_0(\omega'_i) > u_0(\omega_j)$ and thus $N^*(\hat{u}) = \{i\}$, that is, i is also the only potential winner with \hat{u} . So $f(\hat{u}) \in \Omega_i^*(\hat{u})$

²Continuity of utility functions is insufficient, because u_0 could have a (local) maximum at ω_i^* .




Figure 3: Illustrating the argument in the proof of Theorem 2, showing that, if we assume $f(u) \neq \omega_i^*$, we arrive at a contradiction. The gray zones on the right represent the unacceptable outcomes for bidder i under u_i and \hat{u}_i . The outcomes in the gray area on the left bidder i cannot choose, because they yield a lower payoff to the center than the maximal highest acceptable bids of the other bidders. If bidder i would receive $u_i(f(u))$, she could increase her utility by misrepresenting her true preferences u_i by \hat{u}_i . Then, outcome $f(u)$ would no longer be acceptable (light gray zone). Moreover, there would still be an outcome ω_i' that is acceptable under \hat{u}_i and for which $u_0(\omega_i') > u_0(\bar{\omega}_j)$. Thus, by bidding ω_i' , bidder i would still win. Bidder i would then receive an outcome that is acceptable for her under \hat{u}_i (in the white area) and which yields i a strictly higher utility under u_i than $f(u)$.

and since $f(\hat{u})$ is stable with respect to \hat{u} (and thus individually rational), $\hat{u}_i(f(\hat{u})) \geq 0$ and thus

$$\begin{aligned} u_i(f(\hat{u})) &= \hat{u}_i(f(\hat{u})) + u_i(f(u)) + \epsilon \\ &\geq u_i(f(u)) + \epsilon. \end{aligned}$$

This contradicts f being strategy-proof and it follows that $f(u) \in \text{opt}_i(\Omega_i^*(u))$.

Finally, we show that $f(u) \in \text{opt}_0(\text{opt}_i(\Omega_i^*(u)))$. For a contradiction assume that there is an outcome $\omega_i \in \text{opt}_i(\Omega_i^*(u))$ such that $u_0(\omega_i) > u_0(f(u))$. Since both $f(u)$ and ω_i are taken from $\text{opt}_i(\Omega_i^*(u))$ they have equal utility for i . This, however, contradicts $f(u)$ being stable. \square

3.3. Necessity of weakly transferable utility

In this section, we show why weakly transferable utility is necessary for the strategy-proofness of the QVA. The intuition is as follows. Without weakly transferable utility, Lemma 1 no longer holds and a bidder can have positive utility for her highest acceptable bid. Consider a bidder who has positive utility for her highest acceptable bid, is one of the potential winners, but has the tie is broken against her. This bidder has now an incentive to bid another, unacceptable, outcome that is preferred by the center to her highest acceptable bid: the bidder will then win and still be able to choose her original bid after all. The next proposition makes this precise. (The proof uses the language of the indirect mechanism, but this is immaterial.)




Figure 4: The setting used in the proof of Proposition 1, illustrating that the QVA is not strategy-proof if utility is not weakly transferable. Moreover, in the proof of Proposition 2 this example is used to show that there is no mechanism that is strategy-proof, and also chooses outcomes that are stable and Pareto efficient. Again, the gray zone indicates the unacceptable outcomes to the bidders.

Proposition 1. *In the absence of weakly transferable utility, the QVA is not strategy-proof in general.*

Proof. Consider a setting with two bidders, 1 and 2. Let $\Omega_1 = \mathbb{R} \times \{1\}$ and $\Omega_2 = \mathbb{R} \times \{2\}$. Suppose the center's utility function is such that for both bidders i and all $(x, i) \in \Omega_i$,

$$u_0((x, i)) = 1 - x^2 + x^3.$$

Moreover, let for both bidders i the utility function u_i be such that for all $(x, i) \in \Omega_i$,

$$u_i((x, i)) = 1 - 2x^2.$$

See Figure 4 for an illustration of these utility functions. The highest acceptable bid for both bidders i is $(0, i)$, which, if realized, would give the center utility 1 and bidder i utility 1. (The only bids that the center would prefer to these are bids of at least $(1, i)$, but, if realized, such bids result in negative utility to bidder i .) Outcome $(0, i)$ bears witness to the fact that in this situation utility is not weakly transferable, since it is impossible to improve the center's utility by reducing the bidder's utility only marginally, that is, there exists an $\epsilon > 0$ (in this case any ϵ with $0 < \epsilon \leq 2$) for which there is not an outcome ω_i with $u_0(\omega_i) > u_0((0, i))$ and $u_i(\omega_i) \geq u_i((0, i)) - \epsilon$. Suppose, without loss of generality, that if both bidders bid straightforwardly, the tie is broken in favor of bidder 1. Then, bidder 2 would be better off bidding an outcome of (say) $(2, 2)$: this will cause bidder 2 to win, and then bidder 2 can choose outcome $(0, 2)$ after all, because this is at least as good for the center as bidder 1's bid. Consequently, the QVA is not strategy-proof in this context. \square

It turns out that we can, in fact, obtain a version of the QVA that is strategy-proof

even without requiring weakly transferable utility, but we need to be considerably more careful about tie-breaking.

4. Ties and Strategy-proofness in General Settings

The Qualitative Vickrey Auction in Section 3 generalizes the classical Vickrey auction under weakly transferable utility. Example 1, however, illustrated the complexity of the issues that result when dealing with indifferences if this condition does not hold.

In this section, we present a class of mechanisms related to the QVA, which are likewise strategy-proof and choose stable outcomes even in settings where the condition of weakly transferable utility is violated. We find, however, that these properties can be achieved only if the tie-breaking rule is made known in advance. Moreover, we show that it is impossible to guarantee Pareto-efficiency in all situations. Intuitively, in the QVA-i, the winner may choose from among all outcomes that are at least as good for the center as all other bids. In the mechanisms in this section, the winner can instead choose any outcome that, if she had bid that outcome, would have won her the auction, *taking into account the tie-breaking rule at hand*. Accordingly, we (re)define for a utility profile $u = (u_0, u_n, \dots, u_n)$, bidder i , and winner determination (including tie-breaking) rule t , the set of such *potentially winning outcomes* $\Omega_i^*(u, t)$ as

$$\Omega_i^*(u, t) = \{\omega \in \bar{\Omega}_i(u') : u' = (u_0, u_1, \dots, u_{i-1}, u'_i, u_{i+1}, \dots, u_n) \text{ for some } u'_i \text{ and } t(u') = i\}.$$

In the QVA, the outcome is selected from the set of potentially winning outcomes. In particular, to ensure strategy-proofness, the outcome should be among the potentially winning outcomes with the highest utility for the winner. There is, however, a small technical point to observe in the definition of the class of mechanisms in this section. The set of potentially winning outcomes $\Omega_i^*(u, t)$ is not guaranteed to include its limit points and, therefore, may not include an outcome that maximizes the utility of the winner. This is illustrated by the following example.

Example 3. Consider a setting with two bidders, 1 and 2, wherein $\Omega_1 = [0, 1] \times \{1\}$ and $\Omega_2 = [0, 1] \times \{2\}$. Assume further that for all $(x, i), (y, j) \in \Omega$, $u_0(x, i) \leq u_0(y, j)$ if and only if $x \leq y$. The utilities of bidders i are such that for all $(x, i) \in \Omega$,

$$u_1(x, i) = \begin{cases} 1 - x & \text{if } i = 1, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad u_2(x, i) = \begin{cases} 1 - 2x & \text{if } i = 2, \\ 0 & \text{otherwise.} \end{cases}$$

Furthermore, let t break all ties in favor of bidder 2. Assume that both bidders adhere to their straightforward strategies and bid $(1, 1)$ and $(\frac{1}{2}, 2)$, respectively. Then, bidder 1 wins the auction. Still, $\Omega_1^*(u, t) = (\frac{1}{2}, 1] \times \{1\}$ is an open set without an optimal outcome for bidder 1. In particular, she cannot choose $(\frac{1}{2}, 1)$ because she would have lost the tie against bidder 2 with this bid.

It turns out that this problem can be overcome by simply allowing the winner to choose the limit points of $\Omega_i^*(u, t)$ as well. We are now in a position to introduce a new class of direct mechanisms closely related but in essential details different from the QVA-d mechanisms. We define for each utility profile $u = (u_0, u_1, \dots, u_n)$ and each winner determination (including tie-breaking) rule t ,

$$\Omega_{\text{QVA}^+}(u, t) = \text{opt}_0(\text{opt}_{t(u)}(\text{cl}(\Omega_{t(u)}^*(u, t)))).$$

We say that a social choice function f is a *direct QVA⁺ mechanism* if

- (i) $f(u) \in \Omega_{\text{QVA}^+}(u, t)$ for all utility profiles u , and
- (ii) $t(u)$ breaks ties according to a fixed order over the outcomes in Ω such that for all bidders i , $t(u) = i$ implies $i \in N^*(u)$.

The variant we consider in the remainder of this section, which we refer to as *QVA⁺*, is an indirect version of a direct QVA⁺ mechanism. By a *winner determination (including tie-breaking) order* we understand a function t that associates with each bidding profile $\hat{\omega} = (\hat{\omega}_1, \dots, \hat{\omega}_n)$ a bidder i whose bid $\hat{\omega}_i$ is ranked highest among $\hat{\omega}_1, \dots, \hat{\omega}_n$ in a fixed strict (i.e., linear) order over all outcomes in Ω . We generally assume that $t(\hat{\omega}) \in \arg \max_{i \in N} u_0(\hat{\omega}_i)$. To avoid cluttered terminology, in the remainder we will also refer to a winner determination (including tie-breaking) order simply as a *tie-breaking order*.

First, we (re)define the set $\Omega_i^*(\hat{\omega}, t)$ of *potentially winning outcomes* for bidder i given center's utility function u_0 , bid profile $\hat{\omega} = (\hat{\omega}_1, \dots, \hat{\omega}_n)$, and tie-breaking order t , as

$$\Omega_i^*(\hat{\omega}, t) = \{\omega_i \in \Omega_i : t(\hat{\omega}_1, \dots, \hat{\omega}_{i-1}, \omega_i, \hat{\omega}_{i+1}, \dots, \hat{\omega}_n) = i.\}$$

Observe that generally $\omega_i \in \Omega_i^*(\hat{\omega}, t)$ implies that $u_0(\omega_i) \geq u_0(\hat{\omega}_j)$ for all bidders j distinct from i . The indirect mechanism is then formulated as follows. The parts where the mechanism differs from the QVA are italicized.

- (1) First, the utility function u_0 of the center is announced to all bidders *together with a fixed tie-breaking order t over the outcomes*.
- (2) Then, each bidder i submits a sealed bid $\hat{\omega}_i \in \Omega_i$ to the center.
- (3) A bidder i^* who submitted the bid with the highest utility for the center is declared the winner of the auction, that is, i^* is selected from $\arg \max_{i \in N} u_0(\hat{\omega}_i)$. *Ties are broken according to the tie-breaking order t .*
- (4) Finally, the winner i^* may choose any outcome from *the closure of the set of outcomes in Ω_{i^*} that, given the bids of the other bidders, would have made her the winner of the auction, taking into account the tie-breaking order t , that is, from $\text{cl}(\Omega_{i^*}^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t))$* . The outcome she chooses is the outcome of the auction.

The winner's initial offer is witness to the fact that such an outcome always exists. Moreover, if her bid was acceptable to her, then it is witness to the fact that an acceptable such outcome exists.

In this auction a strategy for a bidder i specifies a bid $\hat{\omega}_i \in \Omega_i$ along with a contingency plan prescribing which outcome to choose from $\text{cl}(\Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t))$ for all possible bids $\hat{\omega}_1, \dots, \hat{\omega}_n$ of the other bidders in case i happens to win the auction. Of course, bidder i 's choices can depend on her utility function u_i . We call a strategy for i *straightforward* if it satisfies the following properties:

- (i) the bid i submits is an outcome $\hat{\omega}_i$ with the highest center's utility among those that are acceptable to i , that is, $\hat{\omega}_i \in \bar{\Omega}_i(u)$; *if this leaves multiple outcomes in $\bar{\Omega}_i(u)$ to choose from, i chooses one that is highest in the center's tie-breaking order,*




Figure 5: The QVA-i is not strategy-proof in case there is a tie at $\bar{\omega}_i$ where i loses the tie-breaking, because $\hat{\omega}_i$ is a manipulation where i may still select $\bar{\omega}_i$ as a final outcome. The QVA⁺, however, is strategy-proof, because a winner may only select among outcomes that are potentially winning (or that are limit points of sequences of potentially winning outcomes).

- (ii) in case i wins the auction, she selects one of the outcomes in $\text{cl}(\Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t))$ she values most. If there are multiple such outcomes (equally valued by i), she selects one with the highest utility for the center.

The following two examples illustrate differences between the QVA and QVA⁺ in settings where weakly transferable utility does not hold.

Example 4. Consider the situation illustrated in Figure 5. There, the condition of weakly transferable utility is violated, as there is no outcome in Ω_i with a higher utility than $u_0(\bar{\omega}_i)$ for the center for which i 's utility is less than ϵ less than i 's utility for $\bar{\omega}_i$. First consider the QVA-i in this setting. Outcome $\bar{\omega}_i$ would be bidder i 's straightforward bid. Suppose that this bid ties with another bidder's bid $\hat{\omega}_j$ and that i loses the tie. Consequently she would receive utility 0. However, i could manipulate by bidding the non-straightforward bid $\hat{\omega}_i$ instead and win the auction. Moreover, as the center still has the same utility for $\bar{\omega}_i$ as the second-highest bid $\hat{\omega}_j$, bidder i can now choose $\bar{\omega}_i$ and gain positive utility. Thus, in this setting the QVA is not strategy-proof.

Now consider the QVA⁺ and suppose that i submits the straightforward but tying and losing bid $\bar{\omega}_i$. Any effort to manipulate, however, now fails. For instance, by non-straightforwardly bidding $\hat{\omega}_i$ instead, bidder i would emerge as winner. Yet, the set $\text{cl}(\Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t))$ only contains outcomes that give her negative utility. In particular, it does not contain $\bar{\omega}_i$. Observe that $\bar{\omega}_i$ is not in $\Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t)$, because bidding $\bar{\omega}_i$ makes her lose the auction, and neither is it in the latter's closure, even if ω_i , which has the same utility for the center, is.

Example 5. Another illustrative example involves a discrete domain of outcomes with two bidders, 1 and 2, with $\Omega_1 = \{H, L\} \times \{1\}$ and $\Omega_2 = \{H, L\} \times \{2\}$. Note that any function on this domain is continuous. Assume that, for all $i, j \in \{1, 2\}$, the center strictly prefers (H, i) to (L, j) and is indifferent otherwise, whereas for either bidder i , outcome (L, i) has strictly positive utility and (H, i) is unacceptable. In the QVA, (L, i) is the only straightforward bid for either bidder i . Suppose, without loss of generality,

that the resulting tie is broken in favor of bidder 1. Then, bidder 2 has an incentive to non-straightforwardly bid $(H, 2)$, win the auction, and choose $(L, 2)$, which means an increase in her utility. On the other hand, in the QVA^+ the only straightforward bid for either bidder i is (L, i) . Again suppose bidder 1 to win the tie-break and to choose (L, i) , which is the best i can possibly expect. Now, bidder 2 loses the auction and consequently receives utility 0. However, bidding $(H, 2)$ makes her win the auction but also forces her to choose $(H, 2)$, which yields her a negative utility.

Theorem 3. *The QVA^+ is strategy-proof for the bidders, and selects stable outcomes.*

Proof. Consider an arbitrary utility profile $u = (u_0, u_1, \dots, u_n)$ and let t be a winner determination (including tie-breaking) order.

For strategy-proofness, consider arbitrary bids $\hat{\omega}_1, \dots, \hat{\omega}_n$, one for each bidder. Observe that, as in the case of the QVA, for the QVA^+ by changing her bid, a bidder i can only influence whether she wins or loses, but not the set of outcomes from which she may choose if she happens to win. By individual rationality, the winning bidder receives at least utility zero. Hence, by submitting a losing bid, she cannot be better off. Moreover, the straightforward strategy chooses an outcome optimally for the winner, so there is also no incentive to change her strategy in a way that would still make her win.

Now consider the case in which a bidder i plays a straightforward strategy, bids $\hat{\omega}_i$ accordingly, but loses the auction. Obviously, submitting another losing bid will not increase her utility. It suffices to show that $u_i(\omega_i) \leq 0$ for all $\omega_i \in \text{cl}(\Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t))$.

First, let us assume that there is some bidder j whose bid $\hat{\omega}_j$ involves an outcome with a strictly higher utility for the center than $\hat{\omega}_i$. In that case, all bids ω'_i that would have won bidder i the auction would have negative utility for her, that is, for all $\omega_i \in \Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t)$ we have $u_i(\omega_i) < 0$. Hence, by continuity, for all $\omega_i \in \text{cl}(\Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t))$, $u_i(\omega_i) \leq 0$ and i is not better off submitting a winning bid.

Second, assume that i loses the auction while submitting a straightforward but tying bid $\hat{\omega}_i$. Because a straightforward strategy requires $\hat{\omega}_i$ to be chosen for maximal tie-breaking success, submitting any other tying bid would not render i the winner of the auction. Rather, to win i has to submit a bid ω'_i with strictly higher utility for the center than her original bid, that is, such that $u_0(\omega'_i) > u_0(\hat{\omega}_i)$. Now again we may conclude that $u_i(\omega_i) < 0$ for all $\omega_i \in \Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t)$ and $u_i(\omega_i) \leq 0$ for all $\omega_i \in \text{cl}(\Omega_i^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t))$ by continuity. We conclude that the mechanism is strategy-proof.

Finally, we prove that the QVA^+ chooses stable outcomes. First, to prove individual rationality, observe that if a bidder loses the auction, her utility is 0, whereas if she wins the auction, the outcome she bid is among the ones she can also choose from. When using a straightforward strategy, this outcome would already yield her a non-negative utility, and the strategy will always select an outcome for which her utility is at least as high. Let now i^* be the winner of the auction. We observe that all outcomes $\omega_{i^*} \in \Omega_{i^*}^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t)$ result in at least as high a utility for the center as any acceptable bid of any other bidder. By continuity of the center's utility function, this also holds for every outcome $\omega_{i^*} \in \text{cl}(\Omega_{i^*}^*(\hat{\omega}_1, \dots, \hat{\omega}_n, t))$. Hence, no acceptable outcome for any other bidder would make the center better off. It remains to be shown that it is impossible to make the center better off with the winner i^* without making i^* worse off. Suppose $\omega_{i^*}^0$ is an outcome that makes the center better off. By the above, the center prefers $\omega_{i^*}^0$ to any acceptable bid of any other bidder. Hence, $\omega_{i^*}^0$ was available to i^* to select after winning. Straightforward strategies, however, prescribe bidders to

choose from among the outcomes that yield i^* maximal utility one that also yields the highest utility to the center. Since i^* in fact chose an outcome that is worse for the center than ω_i^0 , it follows that choosing ω_i^0 would have made i^* worse off. \square

However, in some situations where there is a tie among bidders who have a positive utility for winning the tie, the QVA⁺ is not Pareto efficient. This follows immediately from our next result, which is that in some such settings there is *no* mechanism that is strategy-proof and also chooses outcomes that are stable and Pareto efficient.

Proposition 2. *There exists a setting (where utility is not weakly transferable) with two bidders, 1 and 2, with $\Omega_1 = \mathbb{R} \times \{1\}$, $\Omega_2 = \mathbb{R} \times \{2\}$ and a particular center utility function for which there is no mechanism that strategy-proof, and also chooses outcomes that are both stable and Pareto efficient.*

Proof. We prove this is impossible in a restricted setting with a specific center utility function and two bidders. The bidders' outcome sets and the center's utility function are just as in Proposition 1 and Figure 4, that is, $\Omega_i = \mathbb{R} \times \{i\}$ for both bidders i and for all $(x, i) \in \Omega_i$,

$$u_0((x, i)) = 1 - x^2 + x^3.$$

For the sake of contradiction, assume that the mechanism satisfies all three properties: strategy-proofness, choosing stable outcomes, and Pareto efficiency. Consider again the utility profile where for both bidders i the utility function is such that for all $(x, i) \in \Omega_i$,

$$u_i((x, i)) = 1 - 2x^2.$$

By stability, the mechanism cannot choose an outcome (x, i) with $x > \sqrt{\frac{1}{2}}$, since this would otherwise not be individually rational. Hence, the best possible outcomes for the center are $(0, 1)$ and $(0, 2)$, both with utility 1. Furthermore, the mechanism must select one of these outcomes. Otherwise, it would be possible to make the center happier with the other bidder (whom this would also make happier). Without loss of generality, suppose that bidder 1 wins in this situation and the mechanism chooses $(0, 1)$. We note that bidder 1 obtains positive utility, whereas bidder 2 has utility 0.

Define for each $\delta \geq 0$ a modified utility function u_2^δ in which bidder 2's utility has been shifted up by δ , that is, let for all $\delta \geq 0$ and all $(x, 2) \in \Omega_2$,

$$u_2^\delta((x, 2)) = 1 - 2x^2 + \delta.$$

We derive a contradiction by showing that the mechanism is not strategy-proof, as bidder 2 can increase her utility by reporting u_2^δ for some $\delta > 1$ instead of truthfully reporting u_2 .

Under u_2^δ , outcomes $(x, 2)$ with $x > 1$ are acceptable for bidder 2. These outcomes are preferable to the center to outcome $(0, 1)$. Therefore, by the stable outcomes property, for the profile (u_0, u_1, u_2^δ) , bidder 2 must win. Moreover, again by the stable outcomes property, the mechanism must provide an outcome that gives the center utility at least 1.

In fact, by strategy-proofness, we see that bidder 2 must win with an outcome that gives the center utility exactly 1. To arrive at a contradiction, let ω_2 be the outcome that results when bidder 2 reports u_2^δ . This yields the center a utility greater than 1. Then, bidder 2 could report $u_2^{\delta'}$ for a δ' with $1 < \delta' < \delta$, so that ω_2 becomes unacceptable.

Since $\delta' > 1$, however, bidder 2 must still win, so bidder 2 would get an outcome giving her higher utility than ω_2 . This contradicts strategy-proofness, and thus bidder 2 must win with an outcome that gives the center utility exactly 1.

There are only two such outcomes, namely, $(0, 2)$, and $(1, 2)$. Outcome $(1, 2)$ yields the center utility 1 and, under u_2^δ , bidder 2 utility $\delta - 1$. As such, however, outcome $(1, 2)$ is not Pareto efficient: outcome $(0, 2)$ also gives the center utility 1 but bidder 2 utility $1 + \delta$ (under u_2^δ). Therefore, for profile (u_0, u_1, u_2^δ) , the mechanism chooses $(0, 2)$. This, however, gives bidder 2 utility 1 under u_2 , rather than utility 0, which she would have received, had she truthfully reported u_2 . It follows that the mechanism is not strategy-proof, a contradiction. \square

5. Related work

In this section, we relate the results in this article to work by others, as well as to the earlier versions of our own work. We initially described the idea of applying the principle of the Vickrey auction to a more general setting without explicit payments in [11, 16]. In [11, 10], we described the QVA in the contract auction setting with finitely many outcomes when the center has no ties among outcomes. In [10], we also considered a continuous setting for the QVA with ties for the center. There we needed the center's utility function to be *equipeaked*, a weaker condition than weakly transferable utility which required all local optima, or peaks, in the center's utility function to yield the same utility to the center. Both of these settings are special cases of the setting with general utility functions (Section 4). This contrasts with the results in the current paper, where we show that (i) under the marginally stronger condition of weakly transferable utility we can give a full characterization of the QVA with arbitrary tie-breaking, and (ii) with general continuous utility functions (that is, without weakly transferable utility), the QVA is strategy-proof when the winner chooses from the closure of strictly better (incl. tie-breaking) bids, instead of from outcomes that are at least as good for the center as the other bids.

This also illustrates the main differences from a paper using the QVA for sponsored search [9]. There, under the condition of consistent valuations, which is implied by our weakly transferable utility condition, it is also shown that the QVA is strategy-proof. That paper additionally shows how this mechanism can be applied to sponsored search. Besides the consistency condition, there are no restrictions on what a contract constitutes, and this generality is used to model contracts with a cost per impression in addition to a cost per click. The utility functions of the publisher (center) and the advertisers then can be expressed given their estimates of the click-through rate. Working in this model, the authors derive the Impression-Plus-Click Auction and show that it is strategy-proof.

Another application of the QVA can be seen in a mechanism for selecting one out of several providers in e-commerce [4], taking into account privacy aspects of a transaction, such as birth date, gender, address, credit card information, email addresses, etc. The main difference is that in the model used, preferences are partial orders over subsets of the privacy aspects. This domain is a good example of a discrete space, so if such preferences can be translated to a utility function, that function will be continuous and the results on QVA⁺ apply immediately. However, it would be interesting to see whether QVA⁺ can be changed to deal with any partially ordered preference function.

Also related is the work on many-to-one matching with contracts [13, 12]. Specifically, if we restrict that setting to a single hospital (the center) and an auction mechanism resulting in exactly one contract (between a doctor and a hospital), we obtain a setting that corresponds closely to ours, albeit one with only finitely many possible contracts and strict preferences (no ties) for all agents (including the center)—precisely the types of restrictions that we try to avoid here.

There is some work on indifferences in matching *without* contracts [1, 8]. Erdil and Ergin [8] show that when ties are allowed (on the side of the hospitals) in one-to-many matching (without contracts), there does not exist a truth-telling dominant strategy, even though a truth-telling Bayesian-Nash equilibrium does.

The paper by Alkan and Gale [2] considers the setting of one-to-one matching with continuous utility functions. In their work they restrict the outcome set to outcomes with nonnegative utility and that are not Pareto-dominated for the two parties involved, that is, the outcomes for all players can be represented by the graph of a continuous, nonnegative decreasing function in the first quadrant. Note that this implies the condition that utility is weakly transferable. For matching under these conditions, they show that the (strong) core is non-empty and has a strong connectedness property (that is, for any two points in the core there is a monotone path between them where every coordinate along the path either increases, decreases, or remains constant). Dominant strategies for the players are not discussed by Alkan and Gale [2].

Moving back from matching to auctions, in extant work on auctions, payments are almost invariably seen as a special attribute for which the preferences of the center and the bidders are related: a lower price for the bidder means a worse outcome for the center. In our framework, a payment can be part of the specification of an outcome (although we do not require the existence of payments at all). Weakly transferable utility is an immediate generalization of the above property of payments, and as such, our characterization includes not only the traditional Vickrey auction (in which payments are in fact the only parameter of the outcomes), but also multi-attribute auctions. For example, Che analyzed situations where a bid consists of a price and a quality attribute, and proposed first-price and second-price sealed-bid auction mechanisms [5]. His work was extended by David et al. for situations where the good is described by two attributes and a price [6]. They analyzed the first-price sealed-bid and English auction, and derived strategies for bids in a Bayesian-Nash equilibrium. In addition, they studied a setting where the center can also strategize, and they showed when and how much the center can profit from lying about its valuations of the different attributes.

Parkes and Kalagnanam [18] concentrated on iterative multi-attribute reverse English auctions [18]. In their work, prices of attribute-value combinations (a full specification of the good) are initially set high, and bidders submit bids on some attribute-value combinations to lower the prices. The auction finishes when there are no more bids. Such auctions allow the bidders to have any (non-linear) cost structure, and the authors claim that myopic best-response bidding results in an ex-post Nash equilibrium for bidders, and that the auction then yields an efficient outcome. All of the above multi-attribute auctions try to capture the value of non-price-related attributes in auction mechanisms. While they share motivation with our work, those models require the availability of payments to transfer utility. Our work, on the other hand, allows for, but does not require, such payments.

6. Conclusions and future work

In this paper, we considered contract auctions in two different settings. In the first setting, we made an assumption that utility is weakly transferable. This assumption, which still allows the Vickrey auction as a special case, allowed us to dispense with conditions on tie-breaking. Moreover, we were able to prove a characterization showing that the family of QVA mechanisms consists exactly of the mechanisms satisfying strategy-proofness and stability (and they also satisfy Pareto efficiency). In the second setting, we no longer made any assumption about transferable utility at all. Here, we showed how the desirable properties (other than Pareto efficiency) can still be obtained with careful tie-breaking as long as utility functions are continuous, which includes utility functions on discrete spaces.

Still, a number of interesting questions have been left unanswered. First, it would be desirable to have a characterization result in the second setting as well, perhaps showing that in this general setting the QVA^+ is the only mechanism satisfying certain desirable properties. We expect that the tie-breaking rule used would play an important role in such a characterization.

Additionally, it would be desirable to generalize the results for contract auctions presented here to matching with contracts and continuous utility functions, perhaps combining our results with the work by Alkan and Gale [2].

It would also be desirable to study the relation to known impossibility results in matching theory with respect to individual rationality, Pareto-efficiency, and the core [19].

Another avenue for further work is to generalize to combinatorial auctions (package auctions), and to see how this relates to the work by Day and Milgrom [7], which is quite pessimistic about the Vickrey auction in that context. Another possible direction is to analyze the QVA when preferences are partially ordered but not necessarily complete.

Finally, it would be desirable to further study potential computational and communication problems, which will undoubtedly arise when using this type of auction in a variety of realistic applications. For example, as often in public procurement, for the QVA-i the utility function of the center should be known by all bidders, giving rise to the question of how to communicate these preferences efficiently, but also whether it is possible to create a similar mechanism in case the center is not able to communicate, or even know, its own preferences beforehand in every single detail. Relaxing this requirement may give rise to different kinds of indirect implementations of the QVA, for example, ones more similar to multi-bilateral negotiation, such as in [14].

We believe that techniques that successfully address these issues could be effectively used in many real-world settings, ranging from the assignment of programming and development tasks in open-source projects, sponsored search [9], incentivizing privacy in e-commerce [4], to allocating an infrastructural project with a fixed budget to one of a number of competing construction companies.

Acknowledgements

We would like to thank Tamás Máhr and the three anonymous referees for helpful and constructive comments. Mathijs de Weerdts was partly supported by the Next Generation Infrastructures Foundation (<http://www.nextgenerationinfrastructures.eu/>). Paul Harrenstein was supported by the Deutsche Forschungsgemeinschaft under grants

BR 2312/3-2, BR 2312/3-3, and BR 2312/9-1 as well as the ERC under Advanced Grant 291528 ("RACE"). Part of the work by Vincent Conitzer was performed while visiting CWI Amsterdam. Moreover, Vincent Conitzer is grateful for support through NSF IIS-0812113, IIS-0953756, and CCF-1101659, ARO W911NF-12-1-0550 and W911NF-11-1-0332, and an Alfred P. Sloan Research Fellowship.

References

- [1] Abdulkadiroglu, A., Pathak, P. A., Roth, A. E., 2009. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *The American Economic Review* 99 (5), 1954–1978.
- [2] Alkan, A., Gale, D., 1990. The core of the matching game. *Games and Economic Behavior* 2 (3), 203–212.
- [3] Aumann, R. J., Wooders, M. H., 1999. *Topics in Mathematical Economics and Game Theory: Essays in Honor of Robert J. Aumann*. American Mathematical Soc.
- [4] Bonatti, P. A., Faella, M., Galdi, C., Sauro, L., 2011. Towards a mechanism for incentivating privacy. In: *Computer Security—ESORICS 2011*. Springer, pp. 472–488.
- [5] Che, Y. K., 1993. Design competition through multidimensional auctions. *RAND Journal of Economics* 24 (4), 668–680.
- [6] 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, NY, USA, pp. 77–85.
- [7] Day, R. W., Milgrom, P., 2008. Core-selecting package auctions. *International Journal of Game Theory* 36, 393–407.
- [8] Erdil, A., Ergin, H., 2008. What's the matter with tie-breaking? Improving efficiency in school choice. *American Economic Review* 98 (3), 669–689.
- [9] Goel, S., Lahaie, S., Vassilvitskii, S., 2009. Contract auctions for sponsored search. In: Leonardi, S. (Ed.), *Proceedings of the 5th International Workshop on Internet and Network Economics*. Vol. 5929 of *Lecture Notes in Computer Science*. Springer, pp. 196–207.
- [10] Harrenstein, P., De Weerd, M. M., Conitzer, V., 2009. A qualitative Vickrey auction. In: *Proceedings of the ACM Conference on Electronic Commerce*. ACM, pp. 197–206.
- [11] Harrenstein, P., Máhr, T., de Weerd, M. M., 2008. A qualitative Vickrey auction. In: Endriss, U., W, G. P. (Eds.), *Proceedings of the 2nd International Workshop on Computational Social Choice*. University of Liverpool, pp. 289–301.
- [12] Hatfield, J. W., Kominers, S. D., 2012. *Contract Design and Stability in Many-to-Many Matching*. Mimeo.
- [13] Hatfield, J. W., Milgrom, P. R., 2005. Matching with contracts. *American Economic Review* 95 (4), 913–935.
- [14] Hindriks, K. V., Tykhonov, D., de Weerd, M. M., 2012. Qualitative One-to-Many Multi-Issue Negotiation: Approximating the QVA. *Group Decision and Negotiation* 21 (1), 49–77.
- [15] Irving, R. W., 1994. Stable Marriage and Indifference. *Discrete Applied Mathematics* 48 3, 261–272.
- [16] Máhr, T., de Weerd, M. M., 2007. Auctions with arbitrary deals. In: Marek, V., Vyatkin, V., Colombo, A. W. (Eds.), *HoloMAS 2007*. Vol. 4659 of *LNAI*. Springer, pp. 37–46.
- [17] Manlove, D., 2013. *Algorithmics of Matching Under Preferences*. World Scientific Publishing.
- [18] Parkes, D. C., Kalagnanam, J., 2005. Models for iterative multiattribute procurement auctions. *Management Science* 51 (3), 435–451.
- [19] Sönmez, T., 1999. Strategy-proofness and essentially single-valued cores. *Econometrica* 67 (3), 677–689.
- [20] Vickrey, W. S., 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16 (1), 8–37.