A Qualitative Vickrey Auction

Paul Harrenstein, Tamás Máhr, and Mathijs de Weerdt

Abstract

The negative conclusions of the Gibbard-Satterthwaite theorem—that only dictatorial social choice functions are non-manipulable—can be overcome by restricting the class of admissible preference profiles. A common approach is to assume that the preferences of the agents can be represented by *quasilinear utility functions*. This restriction allows for the positive results of the Vickrey auction and the Vickrey-Clarke-Groves mechanism. Quasilinear preferences, however, involve the controversial assumption that there is some commonly desired commodity or numeraire—money, shells, beads, etcetera—the utility of which is commensurable with the utility of the other alternatives in question. We propose a generalization of the Vickrey auction, which does not assume the agents' preferences being quasilinear but still has some of its desirable properties. In this auction a bid can be any alternative, rather than just a monetary bid. Such an auction is also applicable to situations where no numeraire is available, when there is a fixed budget, or when money is no issue. In the presence of quasilinear preferences, however, the traditional Vickrey turns out to be a special case. In order to sidestep the Gibbard-Satterthwaite theorem, we restrict the preferences of the bidders. We show that this qualitative Vickrey auctions always has a dominant strategy equilibrium, which moreover invariably yields a weakly Pareto efficient outcome.

The work in this paper is an improved presentation of the idea introduced by Máhr and de Weerdt (2007) on auctions with arbitrary deals.

1 Introduction

Although it may often seem otherwise, even nowadays money is not always the primary issue in a negotiation. Consider, for instance, a buyer with a fixed budget, such as a government issuing a request for proposals for a specific (public) project, a scientist selecting a new computer using the limited monetary means of a recently acquired project, or an employee organizing a grand day out for her colleagues. In such settings, the buyer has preferences over all possible offers that can be made to him. A similar situation, in which the roles of buyers and sellers are reversed, occurs when a freelancer offers his services at a fixed hourly fee. If he is lucky, several clients may wish to engage him to do different assignments, only one of which he can carry out. Needless to say, the freelancer might like some assignments better than others. In the sequel we consider the general setting which covers all of the examples above and in which we distinguish between an issuer (of a commission)—the government, the scientist, the employee or the freelancer in the examples above—and a number of bidders.

In order to get the best deal, the issuer could ask for offers and engage in a bargaining process with each of the bidders separately. Another option would be to start an (inverted) auction. In this paper, we show that even without money, it is possible to obtain a reasonable outcome in this manner. We propose an auction protocol in which the dominant strategy for each bidder is to make the offer that, among the ones that are acceptable to her, is most liked by the issuer. We also show that if all bidders adhere to this dominant strategy a weakly Pareto optimal outcome results.

To run such an auction without money the preferences of the issuer are made public. Observe that if a single good is sold in an auction with monetary bids, the preferences of both bidders and sellers are commonly known: bidders prefer low prices to higher ones,

and for the sellers it is the other way round. Our protocol closely follows the protocol of a Vickrey, or closed-bid second-price, auction (Vickrey, 1961). First each bidder submits an offer. The winner is the bidder who has submitted the offer that ranks highest in the issuer's preference order. Subsequently the winner has the opportunity to select any other alternative as long as it is ranked at least as high as the second-highest offer in the issuer's preference order. This alternative is then the outcome of the auction.

In the next section some general notations and definitions from implementation theory are introduced, and in Section 3 we formally define the qualitative auction sketched above for the setting in which the bidders are indifferent to all outcomes where they do not win the auction. This makes that we can sidestep the negative conclusions of the impossibility result by Gibbard (1973); Satterthwaite (1975). We prove that a dominant strategy equilibrium exists in the qualitative Vickrey auction, which moreover yields a weakly Pareto efficient outcome for all preference profiles. In the rest of that section we describe several other properties of the qualitative Vickrey auction like monotonicity and (not) incentive compatibility regarding the issuer. We conclude the paper by relating our work to other general auction types such as multi-attribute auctions.

2   Definitions

In this section we review some of the usual terminology of mechanism design and fix some notations. For more extensive expositions the reader be referred to Moore (1992); Mas-Colell et al. (1995); Shoham and Leyton-Brown (forthcoming).

Let N be a finite set of agents and Ω a set of alternatives or outcomes. The players are commonly denoted by natural numbers. By a *preference relation* $\succsim_i$ *of player* i we understand a transitive and total binary relation (that is, a weak order or a total preorder) on Ω, with $\succ_i$ and $\sim_i$ denoting its strict and indifferent part, respectively. We use infix notation and write $a \succsim_i b$ to indicate that player i values alternative a at least as much as alternative b. It is not uncommon to restrict one's attention to particular subsets of preference relations on Ω, *e.g.*, the sets of quasilinear or single-peaked preferences on Ω. Be Θ_i such a class for each $i \in N$, we have Θ denote $\Theta_1 \times \cdots \times \Theta_n$. A *preference profile* $\succsim$ *in* Θ *(over* Ω *and* N*)* is a sequence $(\succsim_1, \ldots, \succsim_n)$ in $\Theta_1 \times \cdots \times \Theta_n$ associating each agent with a preference relation over Ω.

Given a preference profile Θ on Ω, an outcome ω in Ω is said to be *weakly Pareto efficient* whenever there is no outcome ω' in Ω such that all agents i strictly prefer ω' to ω. Outcome ω said to be *Pareto efficient* if for all outcomes ω' in Ω, ω is weakly preferred to ω' by all players and strictly preferred to ω' by some.

A *social choice function (on* Θ*)* is a map $f \colon \Theta \to \Omega$ associating each preference profile with an outcome in Ω. A social choice function on Θ is said to be (weakly) Pareto efficient whenever $f(\succsim)$ is (weakly) Pareto efficient for all preference profiles $\succsim$ in Θ.

A *mechanism (or game form)* M on a set Ω of outcomes is a tuple $(N, S_1, \ldots, S_n, g)$, where N is a set of n agents, for each player i in N, S_i is a set of strategies available to i, and $g \colon S_1 \times \cdots \times S_n \to \Omega$ is a function mapping each strategy profile s in $S_1 \times \cdots \times S_n$ on an outcome in Ω. A mechanism $(N, S_1, \ldots, S_n, g)$ is said to be *direct (on* Θ*)* if each agent's strategies are given by her possible preferences, that is, if $S_i = \Theta_i$ for each agent i in N. For Ω a set of outcomes, a pair $(M, \succsim)$ consisting of a mechanism M on Ω on a and a preference profile $\succsim$ on Ω we refer to as a *game (on* Ω*)*. With a slight abuse of terminology, we will also refer to functions $s_i \colon \Theta \to S_i$ as *strategies* and sequences $s = (s_1, \ldots, s_n)$ of such functions, for each agent one, as *strategy profiles*.

An *equilibrium concept* (or *solution concept*) associates each game with a subset of its strategy profiles; the set of strategy profiles thus associated may depend on the preference

profile. A mechanism M is said to *implement a social choice function f on Θ in an equilibrium concept C* whenever for all preference profiles $\succsim$ in Θ there is some $s^*(\succsim) \in C(M, \succsim)$ with $f(\succsim) = g(s^*(\succsim))$.

A direct mechanism $M = (N, \Theta_1, \ldots, \Theta_n, g)$ is said to be *truthful (or incentive compatible) in an equilibrium concept C* whenever for each preference profile $\succsim$ each agent i revealing her true preferences $\succsim_i$ is an equilibrium in $C(M, \succsim)$, that is, if $\succsim$ itself is in $C(M, \succsim)$.

If for a mechanism $M = (N, S_1, \ldots, S_n, g)$ and an equilibrium concept C, $C(M, \succsim)$ is nonempty for all preference profiles $\succsim$, we can associate with M a direct mechanism $M^* = (N, \Theta_1, \ldots, \Theta_n, g^*)$ where for each $\succsim$ in Θ we have $g^*(\succsim) = g(s^*)$ for some selected equilibrium $s^*(\succsim)$ in $C(M, \succsim)$. Intuitively, M^* mimics M by asking the agents to reveal their preferences, be it truthfully or untruthfully, calculating equilibrium strategies s_i^* in M for them given the revealed preferences and returning the outcome $g(s_1^*, \ldots, s_n^*)$. Thus we find that a social choice function f being implementable in C implies it being truthfully implementable in C, a fact better known as the *revelation principle*.

In this paper we will be primarily concerned with *dominant strategy equilibrium*, which is extensively studied in the context of mechanism design (Dasgupta et al., 1979; Green and Laffont, 1979) and in terms of which also the infamous Gibbard-Satterthwaite theorem is formulated. For the purposes of this paper we say that s_i^* is a *dominant strategy for a player i* in a game $(M, \succsim)$, whenever no matter which strategies the other players adopt i is not worse off playing s_i^* than any other if her strategies, that is, if for all strategy profiles $s \in S$ and all $t_i \in S_i$ we have

$$g(s_1, \ldots, s_{i-1}, s_i^*, s_{i+1}, \ldots, s_n) \succsim_i g(s_1, \ldots, s_{i-1}, t_i, s_{i+1}, \ldots, s_n).$$

A strategy profile $s^* = (s_1^*, \ldots, s_n^*)$ is then said to be a *dominant strategy equilibrium* if s_i^* is a dominant strategy for all players i in N. The advantage of dominant strategy equilibrium is that it is very robust. The dominant strategies of an agent i do not depend on the preferences of the other agents, they can be calculated on the basis of i's preferences alone. Moreover, there seem to be no reason why agents would not play a strategy that is not dominant if a dominant one is available. On the downside is the Gibbard-Satterthwaite theorem, which shows that implementation in dominant strategy equilibrium allows for extremely little flexibility if one does not impose restrictions on the players' preference relations.

3   A Qualitative Vickrey Auction

In the setting we consider, a commission is issued and auctioned. The commission can get a number of alternative implementations denoted by A. The commission is then assigned to a bidder who commits herself to implement it in a particular way. Thus the outcomes of the auction are given by pairs (a, i) of alternatives $a \in A$ and bidders i in N, that is, $\Omega = A \times N$. Intuitively, (a, i) is the outcome in which i wins the auction and implements alternative a. For each bidder i in N we have Ω_i denote $A \times \{i\}$, the set of *offers i can make*. Obviously, each offer is also an outcome, rather, we have $\Omega = \bigcup_{i \in N} \Omega_i$. We assume each bidder to be indifferent to outcomes in which the commission is assigned to another bidder, that is, $\omega \sim_i \omega'$ for all bidders i in N and all outcomes ω and ω' in $\Omega \setminus \Omega_i$. In what follows we have Θ_i denote the set of i's preference profiles over Ω which comply with these restrictions. However idealistic, we do not believe that for our auction setting this is an overly unreasonable assumption to make. No other restrictions on the bidders' preferences are imposed, unless expressly stated otherwise. If $(a, i) \succsim_i (x, j)$ for some alternative x and some bidder j distinct from i, outcome (a, i) is said to be *acceptable to i*, and *unacceptable to i*, otherwise. That is, an outcome ω is acceptable to bidder i if i values at least as much

as any outcome in which she does not win the auction. Observe that if $i \neq j$, any outcome $(a, j) \in \Omega_j$ is acceptable to i. We assume there to be at least two bidders.

Let $\geq$ be a linear (that is, a transitive, total and anti-symmetric) order over the *outcomes* $\Omega = N \times A$. The *qualitative Vickrey auction on* $\geq$ is defined then by the following protocol. First, the order $\geq$ is publicly announced. For $\omega \geq \omega'$ we say that *outcome ω is ranked at least as high as outcome ω' in* $\geq$. Then, each bidder i submits a secret *offer* $(a, i) \in \Omega_i$ to the auctioneer. The bidder i^* who submitted the offer ranked highest in $\geq$ is declared the winner of the auction. Observe that ties are precluded because of the linearity of $\geq$. Finally, i^* may choose from among her own offers in Ω_{i^*} any outcome that are is ranked at least as high as the offer that ranks *second highest* in $\geq$ among all the ones submitted. The outcome she chooses is then the outcome of the auction. The winner's initial offer is witness to the fact that such an outcome always exists.[1]

Example 1 *Let $N = \{1, 2, 3\}$ and $A = \{a, b, c, d\}$. Let us further suppose that the order $\geq$ on the alternatives is lexicographic, that is,*

$$(a, 1) > (a, 2) > (a, 3) > (b, 1) > \cdots > (c, 3) > (d, 1) > (d, 2) > (d, 3).$$

Suppose the three bidders 1, 2, and 3 submit the offers $(c, 1)$, $(a, 2)$ and $(d, 3)$, respectively. Bidder 2 then emerges as the winner, as $(a, 2) > (c, 1) > (d, 3)$. Since $(c, 1)$ is the second-highest offer, bidder 2 may now choose from the outcomes $(a, 2)$ and $(b, 2)$, these being the only outcomes in Ω_2 that rank higher than $(c, 1)$. In case bidder 2 prefers $(b, 2)$ to $(a, 2)$ she would only do well selecting $(b, 2)$, which would then also be the outcome of the auction.

For different orders $\geq$ on the outcomes, the qualitative Vickrey auction can obviously yield different outcomes. So, actually, we have defined a class of auctions. With a slight abuse of terminology we will nevertheless speak of *the* qualitative Vickrey auction if the respective order $\geq$ can be taken as fixed. At first we will consider $\geq$ an extraneous feature of the auction. Later we will come to consider the case in which $\geq$ represents the preferences of the issuer of the commission.

The traditional second-price or Vickrey auction, in which a single item is allocated, is a special case of the above protocol, when the alternatives are taken to be monetary bids for a single good, the bidders have quasilinear preferences over the outcomes and $\geq$ representing the natural order over monetary bids—ranking higher bids higher than lower ones—and a tie-breaking rule.[2] Since from each bid the bidder's entire preference relation can be derived, the traditional Vickrey auction could be considered a direct mechanism. Moreover, it being a special case of the VCG mechanism, it is incentive compatible in dominant strategies.

The qualitative Vickrey auction, however, is not in a general a direct mechanism, as from an offer the full preference relation of a player cannot be derived in general. As such incentive compatibility is not a concept that directly applies to it. Instead we prove the existence of a dominant strategy equilibrium $s^*(\succsim)$ for each preference profile $\succsim$ in Θ. Thus, the qualitative Vickrey auction implements a social choice function f^*, which is defined such

---

[1] A closely related but slightly simpler auction could be defined as follows. First a linear order $\geq$ on the *alternatives* A is announced and each bidder submits a closed bid consisting of an *alternative* a in A. The bidder who submitted the alternative a^* that is ranked highest in $\geq$ wins the auction and may choose from all alternatives that are ordered *strictly* higher than the second-highest bid in $\geq$. In case of a tie, that is, if there is no unique bidder who submitted the alternative that is ranked highest in $\geq$, the tie is broken via some tie-breaking rule and the winner of the tie may choose an outcome from among all outcomes ranked *at least as high as his own bid*. In this auction an dominant strategy equilibrium also always exists and is based on the same underlying principles and intuitions. We are indebted to Kevin Leyton-Brown for this observation.

[2] Not all tie-breaking rules τ, however, can thus be represented. *E.g.*, if τ is such that $\tau(1, 2) = 1$, $\tau(2, 3) = 2$ and $\tau(1, 3) = 3$, it cannot be represented by an order $\geq$.

that for all preference profiles $\succsim$ in Θ, $f^*(\succsim)$ is the outcome of the equilibrium $s^*(\succsim)$. We will then study the formal properties of this social choice function.

Intuitively, the classic Vickrey auction is truthful because an agent's bid only determines whether she turns out the winner but not the price she has to pay if she does so. Things are much similar in the qualitative Vickrey auction. Again, the bidder's offer determines whether she emerges as the winner, but the range of alternatives from among which she may choose is decided by the second-highest offer.

A strategy for a bidder i in the qualitative Vickrey auction consists of an offer (a, i) in Ω_i along with a contingency plan which outcome to choose from among the outcomes in Ω_i that are ranked higher than the second-highest offer submitted in case i happens to win the auction. Any such strategy may depend on a preference profile $\succsim$ in Θ. We call a strategy for i *adequate* if it fulfills the following properties:

  (*i*) the offer i submits is the outcome in Ω_i that is ranked highest in $\geq$, and that is still acceptable to i,

 (*ii*) in case Ω_i contains no outcomes acceptable to her, i submits the outcome in Ω_i that is ranked lowest in $\geq$,

(*iii*) in case i wins the auction, she chooses one of the outcomes in Ω_i she values most among those that are ranked higher than the second-highest offer submitted.

Given a preference profile $\succsim$ items (*i*) and (*ii*) completely determine the offer i is to submit, but (*iii*) may allow for a little flexibility, if i's preferences over Ω_i are not linear. Also observe that whether an offer is acceptable to a bidder i or not can be read off immediately from i's preference relation and does not depend on the preferences of the other players or other extraneous features.

Example 1 (continued) *Let the preferences of the three bidders 1, 2 and 3 be given by the following table, where higher placed outcomes are more preferred.*

1	2	3
$(c, 1)$	$(d, 2)$	$(x, i) \notin \Omega_3$
$(d, 1)$	$(b, 2)$	$(a, 3)$
$(x, i) \notin \Omega_1$	$(a, 2)$	$(d, 3)$
$(b, 1)$	$(x, i) \notin \Omega_2$	$(c, 3)$
$(a, 1)$	$(c, 2)$	$(b, 3)$

If the bidders 1, 2 and 3 were all to play an adequate strategy, they would offer $(c, 1)$, $(a, 2)$ and $(d, 3)$, respectively, since these are for 1 and 2 their highest-ranking acceptable bid and for 3 the lowest-ranking bid overall. In this case $(b, 2)$ would be the outcome of the auction, because bidder 2 is the winner and may select any alternative ranked above $(c, 1)$. It might be worth observing that it can happen that, if all of her offers are unacceptable to her, a player adhering to the strategy bids her least preferred outcome. Bidder 3, e.g., would do so if the outcomes $(b, 3)$ and $(d, 3)$ had been interchanged in her preference order.

We are now in a position prove that a bidders' adequate strategies are dominant in the qualitative Vickrey auction.

Proposition 1 *In the qualitative Vickrey auction and given a preference profile $\succsim$ in Θ, all adequate strategies for a bidder i are dominant.*

Proof: Let i be an arbitrary bidder and $s(\succsim)$ an arbitrary adequate strategy for i. First assume that there are no outcomes in Ω_i that are acceptable to i and that i adheres to $s_i(\succsim)$ submitting the lowest ranked offer in Ω_i, denoted by (a_0^i, i). If i loses the auction, some other bidder i^* ends up winning the auction and chooses some offer (a^*, i^*) in Ω_{i^*} as the eventual outcome. Observe that (a^*, i^*) is acceptable to i and among her most preferred outcomes. If i wins the auction, she may choose among *all* outcomes in Ω_i and, following $s_i(\succsim)$ she will select one that she likes best. Any other offer she could make would still make her win the auction and leaving her the same range of outcomes to choose from. So, obviously, in both cases, $s_i(\succsim)$ is a dominant strategy.

For the remainder of the proof we may assume that there are outcomes in Ω_i which are acceptable to i. Let (a^i, i) denote the highest ranked offer in Ω_i that is still acceptable to i, that is, the offer i would make if she follows the adequate strategy $s_i(\succsim)$. First assume that submitting (a^i, i) would make i lose the auction, that is, that some other bidder i^* would win the auction by offering (a, i^*) and choose (a^*, i^*) as the eventual outcome. Now consider any other offer (a', i^*) in Ω_i which i could submit. Obviously, if (a', i^*) were also a losing offer, i^* would still win the auction and i would be indifferent between the outcome i^* would then choose and (a^*, i^*). On the other hand, if (a', i) would make i win the auction, we have $(a', i) \geq (a, i^*)$, rendering (a, i^*) the second-highest offer. Then, i has to choose from among the outcomes in Ω_i ranked higher than (a, i^*). All of these outcomes, however, are unacceptable to i, that is, $(a^*, i^*) \succ_i \omega$ for all $\omega \in \Omega_i$ with $\omega \geq (a, i^*)$. Thus, also in this case we may conclude that $s_i(\succsim)$ is a dominant strategy for i.

Finally, assume that i wins the auction by offering (a^i, i) and that (b, j) is the second-highest offer. Let (a^*, i) be the outcome she chooses as her most preferred outcome among the outcomes in Ω_i that are ranked higher than (b, j). Then, $(a^i, i) \geq (a^*, i) > (b, j)$, because any outcome in Ω_i ranked higher than (a^i, i) is unacceptable to i. Obviously, $(a^*, i) \succsim_i \omega$ for any outcome $\omega \notin \Omega_i$. For any other winning offer, the second-highest offer would remain the same and so does the set of outcomes from which i may choose. Thus, i would do no better than by offering (a^i, i) as prescribed by $s_i(\succsim)$. On the other hand, if i were to submit a losing offer, some outcome $\omega \notin \Omega_i$ would result. Then $(a^*, i) \succsim_i \omega$, again i would have done better by offering (a^i, i). Hence, $s_i(\succsim)$ is a dominant strategy for i. $\qquad\square$

Among the adequate strategies of a bidder i one stands out, namely, the one in which she selects the highest ranked outcome among those which she prefers most and, of course, are ranked higher than the second-highest offer submitted. For each preference profile $\succsim$ in Θ we denote this strategy by $s_i^*(\succsim)$. In virtue of Proposition 1, the strategy profile $s^*(\succsim)$ in which each agent i adheres to $s_i^*(\succsim)$ for each $\succsim$ in Θ, is a dominant strategy equilibrium. Accordingly, the qualitative Vickrey auction implements the social choice function f^*, which is such that for all preference profiles $\succsim$ in Θ, $f^*(\succsim)$ equals the outcome the strategy profile $s^*(\succsim)$ gives rise to.

Now we are also in a position to define a *direct* mechanism $M^* = (N, \Theta_1, \ldots, \Theta_n, g^*)$ such that N are the bidders participating in the qualitative Vickrey auction we are considering, Θ_i the possible preference relations over Ω (restricted as in the beginning of this section), and g^* such that for all $\succsim$ in $\Theta_1 \times \cdots \times \Theta_n$ we have $g^*(\succsim) = f^*(\succsim)$.

Proposition 2 *The direct mechanism M^* truthfully implements the social choice function f^*.*

Proof: That M^* truthfully implements f^* is an almost immediate consequence of Proposition 1 by an argument much similar to that for the revelation principle. $\qquad\square$

It is quite possible that, given a preference profile $\succsim$, if all bidders play an adequate (and hence dominant) strategy, the outcome (a^*, i^*) of the qualitative Vickrey auction is

unacceptable to i^* although some submitted offers (a, i) were acceptable to the respective bidder i. To appreciate this consider once more Example 1 but now suppose that the bidders' preferences are such that all offers are unacceptable to them, apart from $(d, 2)$, which is acceptable to bidder 2. Then, bidder 1 would win the auction and be forced to select some outcome $(x, 1)$ that is unacceptable to her. This could, and probably should, be considered a serious weakness. Fortunately, this defect can easily be remedied in the direct mechanism M^* by making the following adaptation. In principle, the bidder i^* with the highest ranked offer (a, i^*) in Ω_{i^*} that is still acceptable is declared the winner of the auction, but in case there is no such bidder, that is, if for all bidders i all offers in Ω_i are unacceptable, the bidder i whose lowest ranked offer in Ω_i is still higher ranked than the lowest ranked offers of any of the other bidders is declared the winner. The outcome of the auction is then i's most preferred outcome in Ω_i.

3.1 Pareto efficiency

The generalized Vickrey auction fails to be *(strongly) Pareto efficient among the bidders*, in the sense that for some preference profiles there could be an outcome (a^{**}, j) that is weakly preferred by all bidders over the dominant equilibrium outcome (a^*, i^*), and strictly preferred by some.

Fact 1 *For any order $\geq$ on the outcomes, there is a preference profile for which the outcome of the qualitative Vickrey auction on $\geq$ is not Pareto efficient among the bidders.*

Proof: Let $\geq$ any order on the outcomes and let (a, i) the *lowest* ranked outcome therein. Now define the preference profile $\succsim$ such that for all players j distinct from i all outcomes in Ω_j are unacceptable to j and that (a, i) is the only outcome in Ω_i that is acceptable to i. Obviously, there is no way in which (a, i) can be the outcome of the auction. Still, (a, i) Pareto dominates any other outcome (a^*, i^*): i^* strictly prefers (a, i) to (a^*, i^*) whereas all other bidders are at least indifferent.

In contrast to strong Pareto efficiency, *weak Pareto efficiency among the bidders* is satisfied almost trivially, that is, there are no preference profiles and orders $\geq$ such that some outcome is *strictly* preferred over the dominant equilibrium outcome by all players. If there are three or more bidders, for any (a, i) and (b, j) there is some bidder k distinct from both i and j and thus $(a, i) \sim_k (b, j)$. With two distinct bidders, i and j say, we have $(a, i) \sim_j (b, i)$ and $(a, j) \sim_i (b, j)$ for all $a, b \in A$. The only way, moreover, in which it can happen that both $(a, i) \succ_i (b, j)$ and $(a, i) \succ_j (b, j)$ is that (a, i) is acceptable to i and (b, j) unacceptable to j. In that case, however, (b, j) will not be the dominant strategy equilibrium, as we remarked previously. The case of one bidder is of course trivial.

3.2 Monotonicity

Another property of the social choice function implemented by the qualitative Vickrey auction is that of mononicity. A social choice function f is said to be *monotonic* if the outcome ω stays the same for all preference profiles where the order relative to ω is kept, but everything else may change. In other words, for every player the order above ω may be completely overthrown, as well as the order below ω in its preference order, but no alternative may change from below to above or vice-versa. Formally, we define this property as follows: for any $\omega \in \Omega$ and any preference profile $\succsim$ with $f(\succsim) = \omega$, for any other preference profile $\succsim'$ with the property that $\forall i \in N$, $\forall \omega' \in \Omega$, if it holds that $\omega \succsim_i' \omega'$ if $\omega \succsim_i \omega'$, then it must be that $f(\succsim') = \omega$.

Fact 2 *The social choice function implemented by the qualitative Vickrey auction is monotonic.*

Proof: Let a preference profile $\succsim$ be given, and let $\omega = (a^*, i^*)$ be the outcome of the social choice function implemented by the qualitative Vickrey auction. Let now another preference profile $\succsim'$ be given, such that $\forall i \in N$, $\forall \omega' \in \Omega$ it holds that $\omega \succsim'_i \omega'$ if $\omega \succsim_i \omega'$. Then it holds for any bidder $j \neq i^*$ that set of acceptable bids in $\succsim'_j$ is exactly the same as the set of acceptable bids in $\succsim'_j$, namely, all bids (b, j) for which $(b, j) \succsim_j (a^*, i^*)$. For player i^* it holds that when (a^*, i^*) was acceptable in $\succsim_{i^*}$, then it is also acceptable in $\succsim'_{i^*}$. Consequently, the acceptable offer from i^* that is highest in $\geq$ is at least as high in $\geq$ as (a^*, i^*), and thus player i^* is still the winner.

Furthermore, we show that (a^*, i^*) is the outcome that is most preferred by i^* in $\succsim'_i$ from the outcomes ranked above the second-highest offer in $\geq$. Firstly, we note that this second-highest offer is the same with $\succsim'$ as with $\succsim$, since all other bidders have exactly the same set of acceptable bids. Secondly, (a^*, i^*) is still ranked above this second-highest offer in $\geq$. Finally, suppose for the sake of contradiction that there is a offer (b, i^*) that is ranked higher in $\succsim'_{i^*}$ than (a^*, i^*) and also higher than the second-highest offer in $\geq$. Then this would have been the winning offer for $\succsim$ as well. Consequently, $(b, i^*) = (a^*, i^*)$.

In the border case where no offer is acceptable to any bidder, the outcome for $\succsim$ is the highest offer of the bidder i^* whose highest offer is ranked highest in $\geq$. Since, $\geq$ does not change, the winner is also i^* for $\succsim'$, and by definition of $\succsim'$, we can immediately conclude that (a^*, i^*) is still his highest offer. So for the profile $\succsim'$, the social choice function implemented by the qualitative Vickrey auction selects $\omega = (a^*, i^*)$ as the outcome. □

Since the qualitative Vickrey auction is both weakly Pareto efficient and monotonic, one might think that the Muller and Satterthwaite (1977) theorem applies, saying that it must be that this social choice function is dictatorial when there are three or more alternatives. This, however, is not the case here, because this theorem talks about social choice functions on general preferences. The restriction to the distinct subsets of outcomes we introduced in this paper appears sufficient to thwart the application of this theorem. Note however that when the order $\geq$ ranks all bids of one bidder above all others, this mechanism is in fact a dictatorship.

3.3 Incentive compatibility for the issuer

So far we have assumed that the preference order of the issuer is publicly known, like the fact the a seller likes to get a higher price. In some settings however, this order $\geq$ may not be common knowledge. Therefore we should also study whether the proposed mechanism is incentive compatible for the issuer as well. Unfortunately, we can show that this is not the case, leaving an open problem for future work.

Fact 3 *The qualitative Vickrey auction is not incentive compatible for the issuer.*

Proof: Consider the following case as a counter example. As always, the winner can select an alternative that is equally or more preferred than the second-highest bid in the publicly known ordering. Suppose that there is an alternative in this set he prefers over his own bid. By definition, this alternative is less preferred by the seller than the highest bid. Had the seller manipulated its order, the winner wouldn't have had any other choice than to accept his original bid.

For example, take the preferences and the bids from Example 1. Suppose the buyer moves the alternative $(c, 1)$ up in its order to the spot between $(a, 2)$ and $(a, 3)$. In that case the dominant strategies for the bidders would still lead to the same bid, and the winner

would still be player 2 with his bid $(a, 2)$, but now he is only allowed to choose among the bids higher than or equal to $(c, 1)$, which leaves $(a, 2)$ as the only acceptable alternative. This outcome is better for the buyer than $(b, 2)$, which was the outcome based on his true preference order. $\qquad \square$

4   Extensions and variants

In this section we consider a number of extensions and variants of the idea put forward in this paper.

4.1   Other auction types

To start with, similar results on incentive compatibility and Pareto-efficiency can be obtained for the English auction in a straightforward manner. In this setting the auctioneer accepts only bids in increasing order of the global ordering until no bidder is interested anymore. The dominant strategy for a bidder i is then to bid the highest acceptable alternative in his preference order that is higher in $\geq$ than the current accepted bid. The effect of this strategy is equivalent to the dominant strategy described earlier for the qualitative Vickrey auction: the winner is the bidder that has an acceptable bid that is highest in the public order and the winning alternative is not dominated by any acceptable bid by any other player.

The qualitative auction protocol can also be rephrased for Dutch auctions, or first-price sealed bid auctions, but those are not incentive compatible. But then, neither are traditional variants of these auction, when preferences are assumed to be quasilinear.

4.2   Multi-attribute auctions

The qualitative Vickrey auction we proposed does not assume that preferences of players can be expressed as quasilinear utility functions. This can be a feature for applications where preferences cannot easily be expressed in terms of money. Similar considerations play an important role the related field of multi-attribute auctions. In a multi-attribute auction the good is defined by a set of attributes which can take different values. A bid consist of a value for each attribute and a price. Che (1993) analyzed situations where a bid consists of a price and a quality attribute, and proposed first-price and second-price sealed-bid auction mechanisms. His work was extended by David et al. (2002) for situations where the good is described by two attributes and a price. They analyzed first-price sealed-bid, and English auction formats, and derived strategies for bids in a Bayesian-Nash equilibrium. In addition, they studied a setting where the issuer can also strategize, and they showed when and how much the issuer can profit from lying about his valuations of the different attributes.

Parkes and Kalagnanam (2005) concentrated on iterative multi-attribute reverse English auctions. Here 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 (always bid a little bit below the current ask price) results in an ex-post Nash equilibrium for bidders, and that the auction then yields an efficient outcome. One of the main differences with our approach, besides their proposing an iterative protocol and using Bayes-Nash equilibria as solution concept, is that they assume quasilinear preferences in that the bids always have to include a price. Moreover, the monetary prices are essential to the mechanism working properly. In our model, however, price could be handled as one of the attributes, but could also be left out altogether.

5 Discussion

In this paper we showed that there is another way of dealing with the impossibility theorem by Gibbard (1973); Satterthwaite (1975) besides requiring quasilinear utility functions. For settings where there is only one winner, all that is required is that all bidders are indifferent to all outcomes where they are not the winner. We proposed a protocol for settings where the preference order of the issuer is publicly known (in a way similar to the public knowledge that sellers prefer high prices and buyers low prices). This protocol is called the qualitative Vickrey auction since it can be seen as a generalization of the Vickrey auction to a setting without quasilinear utility functions.

We defined a dominant strategy for this qualitative auction, and showed that it is weakly Pareto efficient in the resulting equilibrium, and that the implemented social choice function is monotonic. Furthermore we showed that the mechanism is not incentive compatible for the issuer, and we briefly discussed the relation to other auction types. Still there are a number of interesting questions left unanswered regarding the properties of qualitative mechanisms such as the one presented here.

To start with, we would like to show what the price is of having a lying issuer. Another direction stems from the fact that we defined qualitative Vickrey auctions as a class of mechanisms, and that this class includes dictatorships (i.e., when all outcomes where a player i wins are ranked above all others). It is interesting to study under which condition on the social order does the dictatorship property holds. Also, other restrictions may lead to specific properties, such as when the preference order of each bidders is the (weak) inverse of the public order $\geq$. Finally, we are interested in seeing whether we can introduce other qualitative generalizations of quasilinear mechanisms (for example of online auctions (Hajiaghayi et al., 2005)).

Acknowledgements

References

Y. K. Che. Design competition through multidimensional auctions. *RAND Journal of Economics*, 24:668–680, 1993.

P. Dasgupta, P. Hammond, and E. Maskin. The implementation of social choice rules: Some general results on incentive compatibility. *The Review of Economic Studies*, 46:185–216, 1979.

E. David, R. Azoulay-Schwartz, and S. Kraus. Protocols and strategies for automated multi-attribute auctions. In *Proceedings of the First Joint Conference on Autonomous Agents and Multiagent Systems*, pages 77–85, 2002.

A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41(4):587–602, July 1973.

J.R. Green and J.-J. Laffont. *Incentives in Public Decision Making*. North-Holland, Amsterdam, 1979.

Mohammad Hajiaghayi, Robert Kleinberg, Mohammad Mahdian, and David Parkes. Online auctions with re-usable goods. In *In Proc. 6th ACM Conf. on Electronic Commerce (EC-05)*, pages 165–174, 2005.

Tamás Máhr and Mathijs de Weerdt. Auctions with arbitrary deals. In V. Marek, V. Vyatkin, and A.W. Colombo, editors, *HoloMAS 2007*, volume 4659 of *LNAI*, pages 37–46. Springer-Verlag Berlin Heidelberg, 2007. ISBN 978-3-540-74478-8.

A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, Inc., 1995.

J. Moore. Implementation, contracts, renegotiations in environments with complete information. In J.J. Laffont, editor, *Advances in Economic Theory*, chapter 5, pages 182–282. Cambridge University Press, 1992.

Eitan Muller and Mark A. Satterthwaite. The equivalence of strong positive association and strategy-proofness. *Journal of Economic Theory*, 14(2):412–418, 1977.

David C. Parkes and Jayant Kalagnanam. Models for iterative multiattribute procurement auctions. *Management Science*, 51(3):435–451, 2005. Special Issue on Electronic Markets.

M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, Cambridge, forthcoming.

W. Vickrey. Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

Paul Harrenstein
Theoretische Informatik
Institut für Informatik
Ludwig-Maximilians-Univerität München
80538 Munich, Germany
Email: `paul.harrenstein@ifi.lmu.de`

Tamás Máhr
Algorithmics Group
Faculty of Engineering, Mathematics and Computer Science (EWI)
Delft University of Technology
2628 CD, Delft, The Netherlands
Email: `t.mahr@tudelft.nl`

Mathijs de Weerdt
Algorithmics Group
Faculty of Engineering, Mathematics and Computer Science (EWI)
Delft University of Technology
2628 CD, Delft, The Netherlands
Email: `M.M.deWeerdt@tudelft.nl`