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 on three or more alternatives 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 offer. 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 auction turns out to be a special case. In order to sidestep the Gibbard-Satterthwaite theorem, we restrict the preferences of the agents. We show that this qualitative Vickrey auction always has a dominant strategy equilibrium, which moreover invariably yields a weakly Pareto efficient outcome, provided there are more than two agents.

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 a fixed budget earmarked for this purpose, 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 a (reverse) 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, provided there are three or more bidders.

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 it can be assumed to be common knowledge that bidders prefer low prices to higher ones, and sellers higher to lower ones. 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 between 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) and 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 with three or more bidders. The rest of that section concerns several other properties like weak monotonicity and incentive compatibility. 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), and Shoham and Leyton-Brown (forthcoming).

Let N be a finite set of agents and Ω a set of alternatives or outcomes. The agents are commonly denoted by natural numbers. By a *preference relation* $\succsim_i$ *of agent* 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 agent 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, for instance, the sets of quasilinear preferences 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 there is no outcome ω' in Ω such that that ω' is weakly preferred to ω by all agents and strictly preferred 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 agent 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 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, one for each agent, 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 $g^*(\succsim) = g(s^*(\succsim))$ for some selected equilibrium $s^*(\succsim)$ in $C(M, \succsim)$. Intuitively, M^* mimicks 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 $\succsim$ and returning the outcome $g(s_1^*(\succsim), \ldots, s_n^*(\succsim))$. 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 an agent i* in a game $(M, \succsim)$, whenever no matter which strategies the other agents adopt, i is not worse off playing s_i^* than any other of 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 agents 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 seems to be no reason why agents would play a strategy that fails to be dominant if a dominant one is available. On the downside is the Gibbard-Satterthwaite theorem, which says that implementation in dominant strategy equilibrium allows only for social choice functions in which one of the players is a dictator if one does not impose restrictions on the agents' preference relations.

3 A Qualitative Vickrey Auction

In the setting we consider, a commission is issued and auctioned among a set N of n agents, henceforth called *bidders*. The commission can get a number of alternative implementations denoted by A, which for presentational purposes we assume to be finite.[1] 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 between 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'

[1] The definitions and results of this paper can be extended so as to hold for infinite sets of outcomes as well, provided appropriate restrictions on the bidders' preferences are imposed. We believe, however, that doing so would technically complicate things while contributing only little to the conceptual content of this paper.

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 this restriction.

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. Finally, a preference profile $\succsim$ is said to be *positive* if each for each bidder i the set Ω_i contains at least one outcome (a,i) which i strictly prefers to losing the auction, that is, to any outcome not in Ω_i. Positive preference profiles could be argued for in contexts where a bidder is assumed not to partake in the auction if she is at best indifferent between winning and losing.[2]

Let $\geq$ be a linear (that is, a transitive, total and anti-symmetric) order over the *outcomes* Ω. 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 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.

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 $\geq$ represents the natural order over monetary bids—ranking higher bids higher than lower ones—together with a deterministic tie-breaking rule.[3] Since from each offer the bidder's entire preference relation can be derived, the traditional Vickrey auction could be considered a direct mechanism. Moreover, being a special case of the VCG mechanism, it is incentive compatible in dominant strategies.

[2]In a similar vein, one could introduce a *zero outcome* 0, which represents the possibility of no transaction taking place. A bidder i could also offer 0, which would intuitively mean that i refrains from participating in the auction. Such a zero outcome, however, brings along a number of intricacies, which lie beyond the scope of this paper.

[3]Not all tie-breaking rules $\tau : 2^{\Omega} \to \Omega$, however, can be represented by $\geq$. *E.g.*, if τ is such that $\tau(\omega_1,\omega_2) = \omega_1$, $\tau(\omega_2,\omega_3) = \omega_2$ and $\tau(\omega_1,\omega_3) = \omega_3$, it cannot be represented by an order $\geq$. Moreover, we also assume the number of possible offers in the Vickrey auction to be arbitrarily large but finite.

The qualitative Vickrey auction, however, is not a direct mechanism, as from an offer the full preference relation of a bidder 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 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 bidder's monetary offer only determines whether she turns out to be the winner, but not what price she has to pay if she does. 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 satisfies 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 selects 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) leaves some room for flexibility when i's preferences over Ω_i contain indifferences. Also observe that whether an offer is acceptable to a bidder i can be read off immediately from i's preference relation and does not depend on the preferences of the other bidders 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 offer and for 3 the lowest-ranking offer 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 bidder adhering to the strategy offers her* least *preferred outcome. Bidder* 3, *for instance, would do so if the outcomes* $(b, 3)$ *and* $(d, 3)$ *had been interchanged in her preference order.*

We are now in a position to prove that the 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. Since $(a^*, i) \succsim_i \omega$, again i would have done better by offering (a^i, i). Hence, $s_i(\succsim)$ is a dominant strategy for i. □

Among the adequate strategies of a bidder i one stands out, namely, the one in which she selects from her most preferred outcomes that ranked higher than the second highest, the one that is ranked highest. For each preference profile $\succsim$ in Θ we denote this strategy by $s_i^*(\succsim)$. Let further s^* be the strategy profile such that $s^*(\succsim) = (s_1^*(\succsim), \ldots, s^*(\succsim))$ for each preference profile $\succsim$. Then, in virtue of Proposition 1, $s^*(\succsim)$ is a dominant strategy equilibrium for each $\succsim$ in Θ. Accordingly, the qualitative Vickrey auction on $\geq$ implements the social choice function $f_{\geq}^*$, which is such that for all preference profiles $\succsim$ in Θ, $f_{\geq}^*(\succsim)$ equals the outcome the strategy profile $s^*(\succsim)$ gives rise to. If $\geq$ is clear from the context we omit the subscript $\geq$ in $f_{\geq}^*$.

We are now 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. □

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 selecting the winner from the bidders i with acceptable outcomes among their set Ω_i of possible offers, if such bidders exist. The problem can obviously also be sidestepped by assuming all preferences to be *positive*, that is, if for each bidder i the set Ω_i contains at least one acceptable outcome which i strictly prefers to losing the auction.

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.

Proposition 3 *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 $\geq$ be any order on the outcomes and let (a, i) be the *lowest* ranked outcome therein. Now define the preference profile $\succsim$ such that for all bidders 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 i strictly prefers to losing the auction. 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^*) with $i^* \neq i$: bidder 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. A mechanism is weakly Pareto efficient if there are no preference profiles and orders $\geq$ such that some outcome is *strictly* preferred over the dominant equilibrium outcome by all bidders. If there are three or more bidders, for any two outcomes (a, i) and (b, j) there is some bidder k distinct from both i and j and thus $(a, i) \sim_k (b, j)$. In words, bidder k will never strictly prefer any outcome where she is not a winner. In the case with only two (distinct) bidders, say i and j, 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. However, (b, j) can turn out the dominant strategy equilibrium outcome only if j has no acceptable offers at all, which is a rather uninteresting borderline case.

Thus far, we have assumed the order $\geq$ to have been given externally. The order $\geq$ could of course also be construed as the preference relation of an additional bidder with a interest in the outcome of the auction, for instance, the issuer of the commission. Extending the concepts of Pareto efficiency so as to include the preferences of this new party, we find that the qualitative Vickrey auction is both weakly and strongly Pareto efficient provided that the preferences of each bidder i are positive and linear over Ω_i. Linearity can be dropped if we consider the direct mechanism M^*.

Proposition 4 *The qualitative Vickrey auction is strongly Pareto efficient among the bidders and $\geq$, if the preferences of each bidder i are positive and linear over Ω_i.*

Proof: Let (a^*, i^*) be a dominant strategy equilibrium outcome of the qualitative Vickrey auction. Having assumed the preferences to be positive, (a^*, i^*) is acceptable to i^*. We now show that (a^*, i^*) is not dominated by any other outcome. Consider an arbitrary outcome (a, i) in Ω distinct from (a^*, i^*). Without loss of generality we may assume that $(a, i) > (a^*, i^*)$. If $i = i^*$, then $(a^*, i^*) \succ_{i^*} (a, i)$ by linearity and the observation that otherwise, i^* would have not have selected (a^*, i^*). On the other hand, if $i \neq i^*$, the outcome (a, i) is ranked higher in $\geq$ than the second-highest offer. As such (a, i) is not acceptable to i, whereas (a^*, i^*) is. Hence, $(a^*, i^*) \succ_i (a, i)$. In either case, (a, i) does not Pareto dominate (a^*, i^*) strongly. □

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 on Ω is said to be *(weakly) monotonic on Θ* if $f(\succsim) = f(\succsim')$ for any preference profiles $\succsim$ and $\succsim'$ in Ω that only differ in that the social choice $f(\succsim)$ under $\succsim$ is possibly moved up in the individual preference orders $\succsim'_i$. In other words, is for all bidders i in N and all outcomes ω and ω' distinct from $f(\succsim)$, $\omega \succsim_i \omega'$ if and only if $\omega \succsim'_i \omega'$ and $f(\succsim) \succsim_i \omega$ implies $f(\succsim) \succsim'_i \omega$, then $f(\succsim) = f(\succsim')$. Intuitively, weak monotonicity captures the desirable property that if the social choice ω^* becomes more preferred by some or more bidders while the bidders' preferences over the other outcomes stay the same, ω^* remains the social choice. A mechanism is said to be weakly monotonic if the social choice functions it implements are weakly monotonic.

For the qualitative Vickrey auction we have imposed the restriction on the individual preferences that a bidder is indifferent between any outcome in which she does not win. In case there are two or more alternatives or more than two bidders, this makes that a loser i of the auction cannot move the outcome (a^*, i^*) up in his preference order, keeping all her other preferences intact, without violating this restriction. Hence, for weak monotonicity on Θ we only have to consider preference profiles that only differ in that the outcome (a^*, i^*) moves up in the preferences of the winner. We then find that the qualitative Vickrey auction is indeed weakly monotonic.

Proposition 5 *The qualitative Vickrey auction is weakly monotonic.*

Proof: If there is only one alternative and no more than two players the proof is trivial. For any other case consider two preference profiles $\succsim$ and $\succsim'$ in Θ and let (a^*, i^*) be the outcome of the auction if the bidders' preferences are given by $\succsim$. Without loss of generality we may assume that $\succsim_i$ and $\succsim'_i$ are identical for all bidders i distinct from i^*. Also assume that $\succsim_{i^*}$ and $\succsim'_{i^*}$ only differ in that (a^*, i^*) is moved up in $\succsim'_{i^*}$. We now show that (a^*, i^*) is also the outcome of the auction if the bidders' preferences are given by $\succsim'$. Observe that for all bidders distinct from i^* the sets of acceptable outcomes given $\succsim_i$ and $\succsim'_i$ remain the same. Hence, the highest-ranked offer (a, i) submitted by any bidder distinct from i^* will be identical given either $\succsim$ or $\succsim'$. Now either (a^*, i^*) is acceptable in $\succsim$ if and only if (a^*, i^*) is in $\succsim'$, or (a^*, i^*) is unacceptable in $\succsim$ but acceptable in $\succsim'$. In the former case, the offer by i^* given $\succsim'$ will be identical to her offer given $\succsim$. In the latter case i^* will offer (a^*, i^*) when the preferences are given by $\succsim'$. In either case i^* also wins the auction for $\succsim'$. Morever, (a^*, i^*) is one of the outcomes among those ranked higher in $\geq$ than (a, i) that i^* prefers most. By moving (a^*, i^*) up in i^*'s preference order, this remains the case and (a^*, i^*) will also be the outcome of the auction if the preferences are given by $\succsim'$. □

A social choice function f is said to be *strongly monotonic on Θ* if $f(\succsim) = f(\succsim')$ for all preference profiles $\succsim$ and $\succsim'$ in Θ such that $f(\succsim) \succsim_i \omega$ implies $f(\succsim) \succsim'_i \omega$ for all bidders i and all outcomes ω. This is a very strong property that is satisfied by hardly any reasonable social choice function. It is therefore not very surprising that the qualitative Vickrey auction fails to be strongly monotonic as well, as witness the following example involving two bidders and three outcomes.

Example 2 *Let* $\geq$ *be given by* $(a,1) > (a,2) > (b,1) > (b,2) > (c,1) > (c,2)$ *and the preference profiles* $(\succsim_1, \succsim_2)$ *and* $(\succsim_1', \succsim_2)$ *as follows.*

1	1′	2
$(c,1)$	$(c,1)$	$(b,2)$
$(b,1)$	$(x,i) \notin \Omega_1$	$(a,2)$
$(x,i) \notin \Omega_1$	$(b,1)$	$(c,2)$
$(a,1)$	$(a,1)$	$(x,i) \notin \Omega_2$

Bidder 1 *and bidder* 2 *then offer* $(b,1)$ *and* $(a,2)$*, respectively, so that bidder* 2 *wins the auction and the outcome is* $(a,2)$*. However, moving* $(a,2)$ *up in bidder* 1*'s preference order, together with* $(b,2)$ *and* $(c,2)$ *so as to comply with the restriction set on preference profiles, and leaving bidder* 2*'s preferences intact results in the profile* $(\succsim_1', \succsim_2)$*. Now, however, bidder* 1 *submits the losing offer* $(c,1)$*, leaving bidder* 2 *in a position to choose her most preferred outcome* $(b,2)$*.*

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 that a seller likes to get a higher price. In some settings however, this order $\geq$ may not be common knowledge. Therefore, we should also investigate 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 to investigate how much the issuer can profit by lying.

Consider the following case where the mechanism is not incentive compatible for the issuer. As always, the winner can select an alternative that is equally or more preferred than the second-highest offer in the publicly known ordering. Suppose that there is an alternative in this set she strictly prefers to her own offer. By definition, this alternative is less preferred by the issuer than the highest offer. Had the issuer manipulated its order by moving the second highest offer up and position it right under the winner's offer, the winner would not have had any other choice than to accept her original offer.

For example, take the preferences and the offers from Example 1. Suppose the issuer 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 offers, and the winner would still be bidder 2 with her offer $(a,2)$, but now she is only allowed to choose among the offers higher than or equal to $(c,1)$, which leaves $(a,2)$ as the only acceptable alternative. This outcome is better for the issuer than $(b,2)$, which was the outcome based on his true preference order.

4 Extensions and variants

In this section we consider a number of extensions and variants of the ideas underlying the qualitative Vickrey auction.

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 offer the highest acceptable alternative in her 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 offer that is highest in $\geq$ and the winning alternative is not dominated by any acceptable offer by any other bidder.

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 auctions, when preferences are assumed to be quasilinear.

4.2 Multi-attribute auctions

The qualitative Vickrey auction does not assume that preferences of bidders 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 in 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 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 issuer can also strategize, and they showed when and how much the issuer can profit from lying about his valuations of the different attributes. The main difference from our work is that in their approach the preferences of the auctioneer (issuer) and the bidders are related: better for the bidder means generally worse for the auctioneer.

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—that is, the strategy always to 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 theirs proposing an iterative protocol and using an ex-post Nash equilibrium as solution concept, is that they use quasilinear utility functions. To the best of our knowledge, such a restriction on the preferences of the issuer and the bidders being weakly inverse is essential to all of the existing work on (multi-attribute) auction mechanisms.

5 Discussion

In this paper we showed that there is another way of dealing with the impossibility theorem by Gibbard (1973) and 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 between 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 class of dominant strategies for this qualitative auction and saw that it is weakly Pareto efficient in the resulting equilibrium, provided there are more than three bidders. We also found that the social choice function implemented by the qualitative Vickrey auction is weakly, but not strongly, monotonic. Furthermore, we showed that the mechanism is not incentive compatible for the issuer. We also briefly discussed the relation of the qualitative Vickrey auction 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.

We would like to show how much worse off the bidders can be if the issuer turns out to be malicious. Another direction stems from the observation that we defined qualitative Vickrey auctions as a class of mechanisms, some of which are dictatorships, for instance when all outcomes with a particular winner are ranked above all outcomes where another bidder wins. It would be interesting to see precisely under which conditions on the issuer's order $\geq$ the qualitative Vickrey auction is not dictatorial. Also, we are interested in the properties of the qualitative Vickrey auction if additional restrictions on $\geq$ and the class of preference orders are imposed, for instance, if the preference orders of the bidders are assumed to be the weak inverse of $\geq$. Finally, we are interested in other qualitative generalizations of quasilinear mechanisms, for example of online auctions (Hajiaghayi et al., 2005).

Acknowledgements

We are indebted to Felix Brandt, Kevin Leyton-Brown and three anonymous referees for inspiring discussions and valuable comments. This work is partially supported by the Technology Foundation STW, applied science division of NWO, and the Ministry of Economic Affairs of the Netherlands as well as by the Deutsche Forschungsgemeinschaft under grants BR 2312/3-1 and BR 2312/3-2.

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.

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

T. Máhr and M.M. 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.

D.C. Parkes and J. 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-Universität München
80538 Munich, Germany
Email: paul.harrenstein@ifi.lmu.de

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

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