# A Qualitative Vickrey Auction* 

B. Paul Harrenstein, Mathijs M. de Weerdt, and Vincent Conitzer

May 4, 2009


#### Abstract

Restricting the preferences of the agents by assuming that their utility functions linearly depend on a payment allows for the positive results of the Vickrey auction and the Vickrey-Clarke-Groves mechanism. These results, however, are limited to settings where there is some commonly desired commodity or numeraire - money, shells, beads, etcetera-which is commensurable with utility. We propose a generalization of the Vickrey auction that does not assume that the agents' preferences are quasilinear, but nevertheless retains some of the Vickrey auction's desirable properties. In this auction, a bid can be any alternative, rather than just a monetary offer. As a consequence, the auction is also applicable to situations where there is a fixed budget, or no numeraire is available at all (or it is undesirable to use payments for other reasons) - such as, for example, in the allocation of the task of contributing a module to an open-source project. We show that in two general settings, this qualitative Vickrey auction has a dominant-strategy equilibrium, invariably yields a weakly Pareto efficient outcome in this equilibrium, and is individually rational. In the first setting, the center has a linear preference order over a finite set of alternatives, and in the second setting, the bidders' preferences can be represented by continuous utility functions over a closed metric space of alternatives and the center's utility is equipeaked. The traditional Vickrey auction turns out to be a special case of the qualitative Vickrey auction in this second setting.


## 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

[^0]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 this paper, we consider a general setting which covers all of the examples above. In this setting, we distinguish between a center who accepts bids - 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 center could ask for offers and engage in a bargaining process with each of the bidders separately. Another option would be to start an auction (or reverse auction). In this paper, we show that even without payments, it is possible to obtain a reasonable outcome, by using an auction in which bidders compete on other aspects of their offers. 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 center. We also show that if all bidders adhere to this dominant strategy, a weakly Pareto optimal outcome results.

To run such an auction without payments, the preferences of the center over the alternatives are made public. (In a typical auction with monetary bids, it can be assumed to be common knowledge that buyers prefer lower prices to higher ones, and sellers higher to lower ones; in general, the center's preferences may not be immediately obvious. However, in many cases it is reasonable to assume that the center's preferences should be common knowledge, for example if the center is a government that is transparently run. Otherwise, we assume, as is generally done in mechanism design, that the center can commit to the mechanism.) Our "qualitative" protocol closely follows the protocol of a Vickrey, or sealed-bid second-price, auction [16]. First, each bidder submits an offer (an alternative). The winner is the bidder who has submitted the offer that ranks highest in the center'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 center's preference order. This alternative is then the outcome of the auction.

For example, suppose an open-source project requires a particular module, and several companies want to contribute it for free, as being the contributor makes it easier to interface with one's existing code. Each party proposes various combinations of functionality, and the best proposal wins; the winner is only obligated to deliver something as good as the next best proposal.

In the next section some general notations and definitions from mechanism design are introduced, and in Section 3 we define the qualitative auction sketched above for the setting in which the bidders are indifferent among all outcomes where they do not win the auction (a no-externalities assumption). This restriction allows us to sidestep the negative conclusions of the impossibility result by Gibbard and Satterthwaite [4, 13]. In Section 4 we prove that a dominantstrategy equilibrium exists in the qualitative Vickrey auction when there are finitely many alternatives and the preference order of the center is a linear order, and that this yields a weakly Pareto efficient outcome. The remainder of
that section discusses several other properties, including a monotonicity property. Thereafter, in Section 5, we show that similar results hold when the bidders' preferences can be represented by continuous utility functions over a closed metric space of alternatives, and the center's utility is equipeaked. 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 terminology from mechanism design and fix some notations. For more extensive expositions, we refer the reader to [8], [9], and [15].

Let $N=\{1, \ldots, n\}$ be a finite set of agents with $n \geq 2$ and $\Omega$ a set of outcomes (possibly infinite). A preference relation $\succsim_{i}$ of agent $i$ is a transitive and total binary relation (that is, a weak order or a total preorder) on $\Omega$, 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 $\Omega$, for instance, quasilinear preferences or single-peaked preferences on $\Omega$. Let $\Theta_{i}$ be such a class for each $i \in N$; we let $\Theta$ denote $\Theta_{1} \times \cdots \times \Theta_{n}$. A preference profile $\succsim i n ~($ over $\Omega$ and $N$ ) is a vector $\left(\succsim_{1}, \ldots, \succsim_{n}\right)$ in $\Theta_{1} \times \cdots \times \Theta_{n}$, associating each agent with a preference relation over $\Omega$. We will assume that the preferences $\succsim_{i}$ of each player $i$ can be represented by a utility function $u_{i}: \Omega \rightarrow \mathbb{R}$. This will prove particularly convenient if the set $\Omega$ of outcomes is infinite.

Given a preference profile in $\Theta$ on $\Omega$, an outcome $\omega$ in $\Omega$ is said to be weakly Pareto efficient whenever there is no outcome $\omega^{\prime}$ in $\Omega$ such that all agents strictly prefer $\omega^{\prime}$ to $\omega$. Outcome $\omega$ is said to be Pareto efficient if there is no outcome $\omega^{\prime}$ in $\Omega$ such that that $\omega^{\prime}$ is weakly preferred to $\omega$ by all agents and strictly preferred by some.

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

A mechanism (or game form) $M$ on a set $\Omega$ of outcomes is a tuple $\left(N, R_{1}, \ldots, R_{n}, g\right.$ ), where $N$ is a set of $n$ agents; for each agent $i$ in $N, R_{i}$ is a set of actions available to $i$; and $g: R_{1} \times \cdots \times R_{n} \rightarrow \Omega$ is a function mapping each action profile in $R_{1} \times \cdots \times R_{n}$ to an outcome in $\Omega$. We will from now on refer to functions $s_{i}: \Theta_{i} \rightarrow R_{i}$ as strategies and vectors $s=\left(s_{1}, \ldots, s_{n}\right)$ of such functions, one for each agent, as strategy profiles.

In this paper we are primarily concerned with implementation in dominantstrategy equilibrium, which has been studied extensively in the context of mechanism design $[2,5]$. We say that a strategy $s_{i}^{*}$ is a dominant strategy for agent $i$ if, for any $\succsim_{i}$, no matter which actions $r_{-i}$ the other agents choose, $i$ is not worse off playing $s_{i}^{*}\left(\succsim_{i}\right)$ than any of her other actions, that is, for all $\succsim_{i} \in \Theta_{i}$,
$r_{-i} \in R_{-i}, r_{i} \in R_{i}$, we have

$$
g\left(r_{1}, \ldots, s_{i}^{*}\left(\succsim_{i}\right), \ldots, r_{n}\right) \succsim_{i} g\left(r_{1}, \ldots, r_{i}, \ldots, r_{n}\right)
$$

A strategy profile $s^{*}=\left(s_{1}^{*}, \ldots, s_{n}^{*}\right)$ is a dominant-strategy equilibrium if $s_{i}^{*}$ is a dominant strategy for all agents $i$ in $N$. In this case, we say that the mechanism implements in a dominant-strategy equilibrium the social choice function defined by $f(\succsim)=g\left(s^{*}(\succsim)\right)$.

The advantage of a 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, so they can be calculated on the basis of $i$ 's preferences alone. Moreover, there seems to be no reason why an agent would play a strategy that fails to be dominant if a dominant one is available. ${ }^{1}$ The downside is that it is not always possible to implement desirable social choice functions in dominant strategies. For example, the Gibbard-Satterthwaite theorem states that implementation in dominant-strategy equilibrium for three or more achievable alternatives allows only for social choice functions in which one of the players is a dictator, or in which at least one of the alternatives is never chosen, unless one imposes restrictions on the agents' preference relations $[4,13]$.

## 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 have a number of alternative implementations, denoted by $A$. The commission is assigned to one of the bidders, 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 let $\Omega_{i}$ denote $A \times\{i\}$, that is, the set of offers $i$ can make, and let $\bar{\Omega}_{i}$ be short for $\Omega \backslash \Omega_{i}$. Each offer is also an outcome, and vice versa, so we have $\Omega=\bigcup_{i \in N} \Omega_{i}$. We make the no-externalities assumption that each bidder is indifferent among outcomes in which the commission is assigned to another bidder. Formally, a bidder $i$ is an indifferent loser if $\omega \sim_{i} \omega^{\prime}$ for all outcomes $\omega, \omega^{\prime} \in \bar{\Omega}_{i}$, and we have no externalities if all bidders are indifferent losers. Without further loss of generality, we assume that $u_{i}(\omega)=0$ for all bidders $i$ and all outcomes $\omega \in \bar{\Omega}_{i}$. In what follows, we have $\Theta_{i}$ denote the set of $i$ 's preferences over $\Omega$ that comply with this restriction.

An outcome $\omega \in \Omega$ is said to be acceptable to $i$ if $\omega \succsim_{i} \omega^{\prime}$ for some $\omega^{\prime} \in \bar{\Omega}_{i}-$ that is, $u_{i}(\omega) \geq 0$-and unacceptable otherwise. That is, an outcome $\omega$ is acceptable to a bidder if she values it at least as much as any outcome in which she does not win the auction. We observe that according to this definition every outcome in $\bar{\Omega}_{i}$ is acceptable to $i$. Preferences $\succsim_{i}$ are said to be satisfiable if the set $\Omega_{i}$ contains at least one acceptable outcome $\omega$ (with $u_{i}(\omega) \geq 0$ ). Satisfiable

[^1]preferences can be argued for in contexts where a bidder is assumed not to partake in the auction if winning is sure to make her worse off. ${ }^{2}$

Let $\geq$ be a total preorder, that is, a reflexive, transitive, and total relation, on $\Omega$. We say that outcome $\omega$ is ranked at least as high as outcome $\omega^{\prime}$ in $\geq$ if $\omega \geq \omega^{\prime}$. The qualitative Vickrey auction on $\geq$ is then defined by the following protocol. First, the order $\geq$ is publicly announced. Then, each bidder $i$ submits a sealed offer $(a, i) \in \Omega_{i}$ to the center. The bidder $i^{*}$ who submitted the offer ranked highest in $\geq$ is declared the winner of the auction. Ties are broken by means of a tie-breaking rule (for the moment unspecified). Finally, the winner $i^{*}$ of the auction may choose from among her own offers in $\Omega_{i^{*}}$ any outcome that is ranked at least as high in $\geq$ as the offer that ranks second highest in $\geq$ among 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,

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

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 select from the outcomes $(a, 2)$ and $(b, 2)$, these being the only outcomes in $\Omega_{2}$ that rank at least as high as $(c, 1)$. In case bidder 2 prefers $(b, 2)$ to $(a, 2)$, she will be better off selecting $(b, 2)$, which would then also be the outcome of the auction.

Naturally, the qualitative Vickrey auction can yield different outcomes for different orders $\geq$ on the outcomes. So, we have actually defined a class of auctions. With a slight abuse of terminology, we will nonetheless speak of the qualitative Vickrey auction if the order $\geq$ can be taken as fixed. At first, we will consider $\geq$ an exogenous feature of the auction. Later, we will consider the case in which $\geq$ represents the preferences of the center.

The classic Vickrey or second-price auction [16] is strategy-proof-that is, bidding truthfully is a dominant strategy-because a 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. The situation is 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 specifies the offer $(a, i)$ in $\Omega_{i}$ to make, along with a contingency plan for which outcome to choose from among the outcomes in $\Omega_{i}$ that are ranked higher than the secondhighest offer submitted, in case $i$ happens to win the auction. Of course, bidder

[^2]$i$ 's choices can depend on her preferences $\succsim_{i}$ in $\Theta_{i}$. We call a strategy for $i$ straightforward if it satisfies the following properties:
( $i$ ) the offer $i$ submits is an outcome in $\Omega_{i}$ that is ranked highest in $\geq$ among those that are acceptable to $i$,
(ii) in case $\Omega_{i}$ contains no outcomes acceptable to her, $i$ submits an outcome in $\Omega_{i}$ that is ranked lowest in $\geq$,
(iii) in case $i$ wins the auction, she selects one of the outcomes in $\Omega_{i}$ she values most among those that are ranked at least as high as the second-highest offer submitted. If there are more than one such outcomes (equally valued by $i$ ), she selects the one that is highest ranked.

Making appropriate assumptions about $\Omega, \succsim_{i}$, and $\geq$, straightforward strategies can be guaranteed to exist. We say that the auction is strategy-proof if straightforward strategies exist, and all straightforward strategies are dominant strategies. We say that the auction is individually rational if for each bidder $i$ a straightforward strategy exist, and if $i$ plays any straightforward strategy the outcome will be acceptable to $i$.

Example 1 (continued) Let the preferences of the 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 all were to play a straightforward strategy, they would offer $(c, 1),(a, 2)$ and $(d, 3)$, respectively, because these are the highest-ranked acceptable offers for 1 and 2, and the lowest ranked offer for 3. Also, if the bidders adopt straightforward strategies, $(b, 2)$ is the outcome of the auction, as bidder 2 is the winner and may select any alternative ranked at least as high as $(c, 1)$.

In Sections 4 and 5 , we study straightforward strategies in two natural settings. The first is where the set of outcomes $\Omega$ is finite and the order $\geq$ is linear. Then, straightforward strategies exist and no ties can occur. Moreover, we find that all straightforward strategies are dominant without qualification, proving the strategy-proofness of the qualitative Vickrey auction in this setting. In the second setting, $\Omega$ is a possibly infinite but closed space, and the bidders' preferences as well as the center's order $\geq$ are representable by continuous utility functions. In this case, too, straightforward strategies prove to be dominant, if we assume that any local maximum in the center's order $\geq$ is also a global maximum, a condition we refer to as equipeakedness. Also, the tie-breaking rule has to comply with a rather natural restriction.

## 4 The Finite Linear Setting

In this section we consider the setting in which the set of alternatives is finite and the center's order $\geq$ is linear, that is, in addition to being total, transitive and reflexive, it also is anti-symmetric. Obviously, this type of order precludes ties. Intuitively, this can be understood as the center breaking possible ties in advance and including the results of this tie-breaking in $\geq$ when it is announced. If moreover the no-externalities assumption holds, we say the setting is finite linear. We now show that under these conditions, straightforward strategies are dominant, that is, the qualitative Vickrey auction is strategy-proof.

Theorem 1 If the set of outcomes is finite, $\geq$ is a linear order, and there are no externalities, then the qualitative Vickrey auction is strategy-proof.

Proof: Under the conditions stated in the theorem, for each bidder straightforward strategies are guaranteed to exist. All that remains to show is that every straightforward strategy is also dominant.

Let $i$ be an arbitrary bidder and let $s_{i}: \Theta_{i} \rightarrow S_{i}$ be an arbitrary straightforward strategy for $i$. First, we consider the case where there are no outcomes in $\Omega_{i}$ that are acceptable to $i$, so that $i$ adheres to $s_{i}$ by submitting the lowestranked offer in $\Omega_{i}$, denoted by $\left(a_{0}^{i}, i\right)$. If $i$ loses the auction, some other bidder $i^{*}$ ends up winning the auction and chooses some offer ( $a^{*}, i^{*}$ ) in $\Omega_{i^{*}}$ as the eventual outcome. We observe that $\left(a^{*}, i^{*}\right)$ is acceptable to $i$ and among her most preferred outcomes. If $i$ wins the auction, she may choose among all outcomes in $\Omega_{i}$ and, following $s_{i}$, she will select one that she likes best. Any other offer she could make would still make her win the auction and leave her the same range of outcomes to choose from. So, in both cases, $i$ cannot make herself better off by changing strategies.

For the remainder of the proof we may assume that there is at least one outcome in $\Omega_{i}$ which is acceptable to $i$. Let $\left(a^{i}, i\right)$ denote the highest-ranked offer in $\Omega_{i}$ that is still acceptable to $i$, that is, the offer $i$ would make if she followed the straightforward strategy $s_{i}$. First, we consider the case where submitting $\left(a^{i}, i\right)$ would make $i$ lose the auction, that is, where some other bidder $i^{*}$ would win the auction by offering $\left(a, i^{*}\right)$ and choose $\left(a^{*}, i^{*}\right)$ as the eventual outcome. Now, consider any other offer $\left(a^{\prime}, i\right)$ in $\Omega_{i}$ which $i$ could submit. Obviously, if $\left(a^{\prime}, i\right)$ were also a losing offer, $i^{*}$ would still win the auction and $i$ would be indifferent between the outcome $i^{*}$ would then choose and $\left(a^{*}, i^{*}\right)$. On the other hand, if $\left(a^{\prime}, i\right)$ would make $i$ win the auction, then we have $\left(a^{\prime}, i\right) \geq\left(a, i^{*}\right)$, rendering $\left(a, i^{*}\right)$ the second-highest offer. Then, $i$ has to choose from among the outcomes in $\Omega_{i}$ ranked higher than $\left(a, i^{*}\right)$. All of these outcomes, however, are unacceptable to $i$, that is, $\left(a^{*}, i^{*}\right) \succ_{i} \omega$ for all $\omega \in \Omega_{i}$ with $\omega \geq\left(a, i^{*}\right)$. Thus, also in this case, $i$ cannot make herself better off by changing strategies.

The final case to consider is where $i$ wins the auction by offering $\left(a^{i}, i\right)$ and $(b, j)$ is the second-highest offer. Let $\left(a^{*}, i\right)$ be the outcome $i$ chooses as her most preferred outcome among the outcomes in $\Omega_{i}$ that are ranked higher
than $(b, j)$. Then, $\left(a^{i}, i\right) \geq\left(a^{*}, i\right)>(b, j)$, because any outcome in $\Omega_{i}$ ranked higher than $\left(a^{i}, i\right)$ is unacceptable to $i$. We have $\left(a^{*}, i\right) \succsim_{i} \omega$ for any outcome $\omega \notin$ $\Omega_{i}$. If $i$ were to submit another offer that would still make her win, then the second-highest offer would remain the same, and so would the set of outcomes from which $i$ may choose. Thus, $i$ would do no better than by offering $\left(a^{i}, i\right)$ as prescribed by $s_{i}$. On the other hand, if $i$ were to submit a losing offer instead, some outcome $\omega \notin \Omega_{i}$ would result. Since $\left(a^{*}, i\right) \succsim_{i} \omega$, again $i$ would have done no worse by offering $\left(a^{i}, i\right)$. We can now conclude that $s_{i}$ is a dominant strategy for $i$.

The no-externalities requirement is meant to exclude examples like the following. Let there be two bidders, 1 and 2, and two alternatives, $a$ and $b$. Suppose the center's order is given by

$$
(a, 1)>(a, 2)>(b, 2)>(b, 1)
$$

Suppose further that bidder 1's preferences are such that

$$
(b, 1) \succ_{1}(b, 2) \succ_{1}(a, 1) \succ_{1}(a, 2) .
$$

Obviously, bidder 1 is no indifferent loser, that is, the no-externalities assumption is violated here. Also, she does not have a dominant strategy in this qualitative Vickrey auction. If bidder 2 submits $(b, 2)$, with the intention of also choosing $(b, 2)$ if she wins, bidder 1 is better off submitting $(b, 1)$ and losing the auction, than bidding $(a, 1)$, winning the auction and being forced to choose $(a, 1)$. If, on the other hand, bidder 2 were to submit $(a, 2)$, with the intention of also choosing $(a, 2)$ if she wins, bidder 1 prefers to avert disaster by submitting $(a, 1)$ and winning the auction. (By submitting $(b, 1)$, bidder 1 loses the auction and the outcome will be ( $a, 2$ ), an outcome less favorable to bidder 1 than ( $a, 1$ ).)

Without the requirement that the center's order is linear, one runs into all kinds of trouble concerning tie-breaking, at least in the finite case. Later we will see that in the continuous case, such problems can be side-stepped and no such restriction is necessary.

It is quite possible that, given a preference profile $\succsim$, if all bidders play a straightforward (and hence dominant) strategy, the outcome ( $a^{*}, i^{*}$ ) of the qualitative Vickrey auction is unacceptable to $i^{*}$, even though some submitted offers $(a, i)$ were acceptable to the respective bidder $i$. To see why, 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 be considered a serious weakness. This problem, however, can easily be side-stepped, by assuming all preferences to be satisfiable, that is, if for each bidder $i$ the set $\Omega_{i}$ contains at least one acceptable outcome.

Proposition 1 In the finite linear setting the qualitative Vickrey auction is individually rational if all bidders' preferences are satisfiable.

Proof: If a bidder $i$ plays a straightforward strategy, he submits an acceptable bid, which is guaranteed to exist since all preferences are satisfiable. If $i$ loses, the outcome is acceptable. If $i$ wins, he can at least choose the offer he submitted.

Besides individual rationality, another reason for requiring satisfiable preferences is that without this, the qualitative Vickrey auction fails to be (strongly) Pareto efficient among the bidders. In other words, for some preference profiles there could be an outcome $\left(a^{* *}, j\right)$ that is weakly preferred by all bidders over the straightforward outcome $\left(a^{*}, i^{*}\right)$, and strictly preferred by some.

Proposition 2 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 $\Omega_{j}$ are unacceptable to $j$ and such that $(a, i)$ is the only outcome in $\Omega_{i}$ 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 $\left(a^{*}, i^{*}\right)$ with $i^{*} \neq i$ : bidder $i^{*}$ strictly prefers $(a, i)$ to $\left(a^{*}, i^{*}\right)$ whereas all other bidders are at least indifferent.

In contrast to strong Pareto efficiency, weak Pareto efficiency among the bidders is satisfied almost trivially when there are at least three bidders. The mechanism is weakly Pareto efficient if there are no preference profiles and orders $\geq$ such that some outcome is strictly preferred over the straightforward 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 $(a, i) \sim_{k}(b, j)$.

Thus far, we have assumed that the center's order $\geq$ has been given externally. The order $\geq$ could of course also be construed as the preference relation of an additional player with an interest in the outcome of the auction, in particular, the center 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 satisfiable and linear over $\Omega_{i}$.

Proposition 3 In the finite linear setting the qualitative Vickrey auction is strongly Pareto efficient among the bidders and the center, provided preferences are satisfiable.

Proof: Let $\left(a^{*}, i^{*}\right)$ be an outcome of the qualitative Vickrey auction resulting from straightforward strategies. Having assumed the preferences to be satisfiable, $\left(a^{*}, i^{*}\right)$ is acceptable to $i^{*}$. We now show that $\left(a^{*}, i^{*}\right)$ is not Pareto dominated by any other outcome. For a contradiction, assume that ( $a, i$ ) weakly Pareto dominates $\left(a^{*}, i^{*}\right)$. Then, $(a, i)$ is distinct from $\left(a^{*}, i^{*}\right)$ and by linearity
of $\leq,(a, i)>\left(a^{*}, i^{*}\right)$. If $i=i^{*}$, we have $\left(a^{*}, i^{*}\right) \succ_{i^{*}}(a, i)$, for, otherwise, $\left(a^{*}, i^{*}\right)$ would not have been a choice of $i^{*}$ that is compatible with her playing a straightforward strategy. On the other hand, if $i \neq i^{*},(a, i)$ is ranked higher in $\geq$ than the bid $i$ submitted, that is, than the highest-ranked outcome that is still acceptable to $i$. Hence, $(a, i)$ is unacceptable to $i$, whereas $\left(a^{*}, i^{*}\right)$ is acceptable to $i$. Hence, $\left(a^{*}, i^{*}\right) \succ_{i}(a, i)$. In either case, $(a, i)$ does not Pareto dominate $\left(a^{*}, i^{*}\right)$.

Another interesting property of social choice functions implemented by the qualitative Vickrey auction is that of mononicity. A social choice function $f$ on $\Omega$ is said to be (weakly) monotonic on $\Theta$ if we have $f(\succsim)=f\left(\succsim^{\prime}\right)$ for any two preference profiles $\succsim$ and $\succsim^{\prime}$ in $\Omega$ that satisfy the following property: for all $i$, the orders $\succsim_{i}$ and $\succsim_{i}^{\prime}$ are identical, except for the position of $f(\succsim)$, which is ranked higher or the same in the rankings $\succsim_{i}^{\prime}$. In other words, if for all bidders $i$ in $N$, we have (1) for all outcomes $\omega$ and $\omega^{\prime}$ distinct from $f(\succsim), \omega \succsim_{i} \omega^{\prime}$ if and only if $\omega \succsim_{i}^{\prime} \omega^{\prime}$; and (2) for every outcome $\omega \neq f(\succsim), f(\succsim) \succsim_{i} \omega$ implies $f(\succsim) \succsim_{i}^{\prime} \omega$; then $f(\succsim)=f\left(\succsim^{\prime}\right)$. Intuitively, weak monotonicity captures the desirable property that if the current social choice $\omega^{*}$ becomes more preferred by some agents while the agents' preferences over the other outcomes stay the same, $\omega^{*}$ 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 no-externalities restriction on the individual preferences that a bidder is indifferent among all outcomes in which she does not win. As long as there are two or more alternatives or more than two bidders, it is impossible for a loser $i$ of the auction to move the outcome $\left(a^{*}, i^{*}\right)$ up in her preference order while keeping all her other preferences intact, without violating no-externalities (because the other outcomes in which she loses cannot also move up). Hence, for weak monotonicity on $\Theta$ we only have to consider preference profiles that only differ in that the outcome $\left(a^{*}, i^{*}\right)$ moves up in the preferences of the winner. We then find that the qualitative Vickrey auction is indeed weakly monotonic.

Proposition 4 In the finite linear setting, the qualitative Vickrey auction is weakly monotonic.

Proof: If there is only one alternative and no more than two bidders, then the proof is trivial. For any other case, let us consider two preference profiles $\succsim$ and $\succsim^{\prime}$ in $\Theta$, and let $\left(a^{*}, i^{*}\right)$ 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 $\succsim_{i}^{\prime}$ are identical for all bidders $i$ distinct from $i^{*}$, and that $\succsim_{i^{*}}$ and $\succsim_{i^{*}}^{\prime}$ only differ in that $\left(a^{*}, i^{*}\right)$ is moved up in $\succsim_{i^{*}}^{\prime}$. We now show that $\left(a^{*}, i^{*}\right)$ is also the outcome of the auction if the bidders' preferences are given by $\succsim^{\prime}$. For all bidders distinct from $i^{*}$, the sets of acceptable outcomes given $\succsim_{i}$ and $\succsim_{i}^{\prime}$ are the same. Hence, the highest-ranked offer $(a, i)$ submitted by any bidder distinct from $i^{*}$ will be identical given either $\succsim$ or $\succsim^{\prime}$. Now, either ( $a^{*}, i^{*}$ ) is acceptable in $\succsim$ if and only if $\left(a^{*}, i^{*}\right)$ is acceptable in $\succsim^{\prime}$, or $\left(a^{*}, i^{*}\right)$ is unacceptable in $\succsim$ but acceptable in $\succsim^{\prime}$. In the former case, the offer by $i^{*}$ given $\succsim^{\prime}$ will be
identical to her offer given $\succsim$. In the latter case, $i^{*}$ will offer ( $a^{*}, i^{*}$ ) when the preferences are given by $\succsim^{\prime}$. In either case, $i^{*}$ also wins the auction for $\succsim^{\prime}$. Because under $\succsim,\left(a^{*}, i^{*}\right)$ is one of $i$ 's most-preferred outcomes among those ranked higher in $\geq$ than $(a, i)$, it must be the case that under $\succsim^{\prime},\left(a^{*}, i^{*}\right)$ is uniquely $i$ 's most-preferred outcome among those ranked higher in $\geq$ than $(a, i)$. So, $\left(a^{*}, i^{*}\right)$ will be the outcome of the auction if the preferences are given by $\succsim^{\prime}$.

A social choice function $f$ is said to be strongly monotonic on $\Theta$ if $f(\succsim)=$ $f\left(\succsim^{\prime}\right)$ for all preference profiles $\succsim$ and $\succsim^{\prime}$ in $\Theta$ such that $f(\succsim) \succsim_{i} \omega$ implies $f(\succsim) \succsim_{i}^{\prime} \omega$, for all bidders $i$ and all outcomes $\omega$. 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 the following example involving two bidders and three outcomes shows.

Example 2 Let $\geq$ be given by $(a, 1)>(a, 2)>(b, 1)>(b, 2)>(c, 1)>(c, 2)$ and let the preference profiles $\left(\succsim_{1}, \succsim_{2}\right)$ and $\left(\succsim_{1}^{\prime}, \succsim_{2}\right)$ be as follows.

| 1 | $1^{\prime}$ | 2 |
| :--- | :--- | :--- |
| $(c, 1)$ | $(c, 1)$ | $(b, 2)$ |
| $(b, 1)$ | $(x, i) \notin \Omega_{1}$ | $(a, 2)$ |
| $(x, i) \notin \Omega_{1}$ | $(b, 1)$ | $(c, 2)$ |
| $(a, 1)$ | $(a, 1)$ | $(x, i) \notin \Omega_{2}$ |

In the first profile, bidder 1 and bidder 2 offer $(b, 1)$ and $(a, 2)$, respectively, so that bidder 2 wins the auction and the outcome is $(a, 2)$. However, if we move $(a, 2)$ up in bidder 1's preference order, together with $(b, 2)$ and $(c, 2)$ so as to comply with no-externalities, and leave bidder 2 's preferences intact, then we obtain the profile $\left(\succsim_{1}^{\prime}, \succsim_{2}\right)$. Now, however, bidder 1 submits the losing offer $(c, 1)$, leaving bidder 2 in a position to choose her most preferred outcome (b, 2).

So far, we have assumed that the preference order of the center is publicly known. In some settings this is reasonable - for example, in a standard auction where it is common knowledge that the center prefers larger payments to smaller ones, or in the case where the center is a transparently run government. 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 center. Unfortunately, we can show that this is not the case, leaving an open problem for future work to investigate how much the center can profit by lying. ${ }^{3}$

[^3]Consider the following case where the mechanism is not incentive compatible for the center. As always, the winner can select any alternative that is ranked at least as highly as the second-highest offer in the center's ordering. Suppose that there is an alternative in this set that she strictly prefers to her own offer. This alternative is less preferred by the center than the agent's original offer. Had the center manipulated its order by moving the second-highest offer up and positioning it right under the winner's offer, then the winner would not have had any choice but to accept her original offer.

To make the example concrete, let us take the preferences and the offers from Example 1. Suppose the center moves the alternative $(c, 1)$ up in its order to the spot between $(a, 2)$ and $(a, 3)$. In that case, the (dominant) straightforward 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 she is only allowed to choose among the offers higher than or equal to $(c, 1)$, which now leaves $(a, 2)$ as the only allowed alternative. This outcome is better for the center than $(b, 2)$, which was the outcome resulting from its true preference order.

## 5 A Continuum of Alternatives

In the previous section, we have been concerned with the setting in which the alternative set is finite and the center's order is linear, and found that the qualitative Vickrey auction is strategy-proof in this case. If the number of alternatives is infinite, however, this no longer holds without certain restrictions being fulfilled. For instance, if $A=\mathbb{R}$ and the center's and bidders' preferences are all given by the natural order $\geq$ over $\mathbb{R}$, then no bidder's bid can be high enough, and even if a bidder were to win the auction, there would be no optimal way for her to choose an alternative. This particular example can easily be obviated by making some assumptions on the set of outcomes and the bidders' preferences. But even if we do so, an additional assumption needs to be made to render the qualitative Vickrey auction on continuous domains strategy-proof.

For the remainder of this section, we assume that each set $\Omega_{i}$ of outcomes constitutes a closed, but not necessary bounded, metric space with a metric $d_{i}: \Omega \times \Omega \rightarrow \mathbb{R}$. We also assume each bidder $i$ to be an indifferent loser (no externalities) and her preferences over $\Omega_{i}$ to be representable by a continuous utility function $u_{i}: \Omega_{i} \rightarrow \mathbb{R}$. We require $u_{i}(\omega)=0$ for all $\omega \in \bar{\Omega}_{i}$. Furthermore, for each bidder $i$, we assume the set $\Omega_{i}^{*}$ of outcomes in $\Omega_{i}$ that are acceptable to $i$ to be non-empty, bounded and closed. Whenever these restrictions are fulfilled, and the center's order $\geq$ can moreover be represented by a continuous utility function $u_{c}$, we say the setting is continuous. We call the center's ordering competitive if for every bidder $i$ and every $\omega \in \Omega_{i}$, there is, for every bidder $j$ distinct from $i$, some $\omega^{\prime} \in \Omega_{j}$ such that $u_{c}(\omega)=u_{c}\left(\omega^{\prime}\right)$.

The continuous setting does not exclude ties and thus sometimes tie-breaking is in order. Tie-breaking rules come in all sorts and kinds, some more reasonable than others. A mild condition is that of neutrality (with respect to the alternatives) which is fulfilled by a probabilistic tie-breaking rule $\tau: \times_{i \in N} \Omega_{i} \rightarrow \Delta(\Omega)$
if all tying bids of a bidder $i$ in $\Omega_{i}$ have an equal chance of winning. Formally, let $\tau_{i}\left(\omega_{1}, \ldots, \omega_{n}\right)$ denote the probability that $\omega_{i}$ is selected by $\tau\left(\omega_{1}, \ldots, \omega_{n}\right)$. Then, $\tau$ is neutral if for all bidders $i$, all $\left(\omega_{1}, \ldots, \omega_{n}\right) \in \times_{i \in N} \Omega_{i}$ and all $\omega_{i}^{\prime} \in \Omega_{i}$, $u_{c}\left(\omega_{i}\right)=u_{c}\left(\omega_{i}^{\prime}\right)$ implies

$$
\tau_{i}\left(\omega_{1}, \ldots, \omega_{n}\right)=\tau_{i}\left(\omega_{1}, \ldots, \omega_{i-1}, \omega_{i}^{\prime}, \omega_{i+1}, \ldots, \omega_{n}\right)
$$

Given a neutral tie-breaking mechanism a bidder cannot improve the probability of her winning a tie-breaking event by bidding $\omega^{\prime}$ instead of $\omega$ if $u_{c}(\omega)=u_{c}\left(\omega^{\prime}\right)$. In this section we will assume tie-breaking to be neutral with respect to the alternatives.

We say that $x \in A$ is a local maximum or a peak of a continuous function $f: A \rightarrow \mathbb{R}$ on a closed metric space with metric $d$ if there is an $\epsilon>0$ such that $f(y) \leq f(x)$ for all $y \in B_{\epsilon}(x)$, where $B_{\epsilon}(x)=\{y \in A: d(x, y) \leq \epsilon\}$ is the $\epsilon$-ball around $x$. We then say that a function $f: A \rightarrow \mathbb{R}$ is equipeaked if all local maxima of $f$ are also global maxima of $f$, that is, for all local maxima $x \in A$ of $f$ there is no $y \in A$ with $f(y)>f(x)$. For bounded and closed domains a continuous function $f$ is equipeaked if and only if $f(x)=f(y)$ for all local maxima $x, y \in A$ of $f$, hence the terminology. If an order $\leq$ over $A$ is represented by an equipeaked utility function $u: A \rightarrow \mathbb{R}$, we also say that $\leq$ is equipeaked. Note that equipeakedness is not too restrictive, since for example settings with no local maxima, or with a unique local maximum that is also the global maximum are special cases.

We find that within the continuous setting, equipeakedness of the center's order over each $\Omega_{i}$ and neutral tie-breaking are sufficient conditions for the qualitative Vickrey auction to be strategy-proof. If moreover the center's order is assumed to be competitive, then equipeakedness is also necessary. The intuition behind the proof of sufficiency is captured by the following informal argument that no strategies can yield a higher payoff to a bidder $i$ than the straightforward ones. Let $\hat{\omega}_{1}, \ldots, \hat{\omega}_{n}$ be the bids submitted by bidders $1, \ldots, n$, respectively. Suppose that $i$ followed a straightforward strategy when bidding $\hat{\omega}_{i}$. If $\hat{\omega}_{i}$ does not tie the highest other $\operatorname{bid}(\mathrm{s})$, the argument is basically as in the finite linear setting. So we can assume that $\hat{\omega}_{i}$ is among the highest-ranked bids submitted but that it is tied with at least one other bid. The tie-breaking rule then determines the winning bid. We distinguish two cases. Either $\hat{\omega}_{i}$ is a local maximum in $u_{c}$, or it is not. In the former case, it is impossible for $i$ to submit a higher-ranked bid because of the assumption that $u_{c}$ is equipeaked over $\Omega_{i}$. Neither would it help $i$ to offer a bid that is ranked just as high as $\hat{\omega}_{i}$ by the center so as to manipulate the tie-breaking, as tie-breaking is assumed to be neutral with respect to the alternatives. If, on the other hand, $\hat{\omega}_{i}$ is not a local maximum in $u_{c}$, then it can be shown that bidder $i$ 's utility is zero no matter which bid she submits, that is, $i$ is indifferent between winning and losing the auction.

Theorem 2 In the continuous setting, the qualitative Vickrey auction is strategyproof if the center's order $\geq$ is equipeaked over each $\Omega_{i}$ and tie-breaking is neutral with respect to the alternatives. Moreover, if the center's order is competitive,
and tie-breaking gives every tied bidder a positive chance of winning, then equipeakedness of $\geq$ is also a necessary condition.

Proof: For the first claim, we need to prove that straightforward strategies are dominant, that is, no other strategy ever gives a higher utility. So let $\hat{\omega}_{1}, \ldots, \hat{\omega}_{n}$ be the bids submitted by bidders $1, \ldots, n$, respectively. Let $i$ be an arbitrary bidder who plays a straightforward strategy in submitting $\hat{\omega}_{i}$. Having assumed $\Omega_{i}^{*}$, that is the outcomes in $\Omega_{i}$ that are acceptable to $i$, to be non-empty, $\hat{\omega}_{i}$ is acceptable to $i$ and maximizes the center's utility function $u_{c}$ in $\Omega_{i}^{*}$. Let $v \in \bar{\Omega}_{i}$ be one of the highest-ranked offers among the other bids, that is, $v \in \operatorname{argmax}_{\omega \in\left\{\hat{\omega}_{j}: j \neq i\right\}} u_{c}(\omega)$. Now consider the set $\Upsilon_{i}=\left\{\omega \in \Omega_{i}^{*}: u_{c}(\omega) \geq\right.$ $\left.u_{c}(v)\right\}$ of acceptable outcomes from among which $i$ can choose if she submits a winning bid. Since $\Omega_{i}^{*}$ is bounded and closed, so is $\Upsilon_{i}$. Hence, if $i$ wins the auction-be it by submitting the highest bid or by tying and subsequently winning the tie-break-an outcome maximizing $i$ 's utility function $u_{i}$ in $\Upsilon_{i}$ exists. Let $\omega_{i}^{*} \in \operatorname{argmax}_{\omega \in \Upsilon_{i}} u_{i}(\omega)$. Without loss of generality we may assume that $i$ submits $\hat{\omega}_{i}$ and chooses outcome $\omega_{i}^{*}$ if she wins.

We distinguish three cases: $u_{c}\left(\hat{\omega}_{i}\right)>u_{c}(v), u_{c}\left(\hat{\omega}_{i}\right)<u_{c}(v)$ and $u_{c}\left(\hat{\omega}_{i}\right)=$ $u_{c}(v)$. In the first case, $i$ wins the auction, chooses $\omega_{i}^{*}$, and $i$ 's utility is $u_{i}\left(\omega_{i}^{*}\right)$. By submitting any other winning bid, her utility will likewise be $u_{i}\left(\omega^{*}\right)$. If she were to submit a losing bid, her utility would drop to zero, whereas a bid which ties $v$ would yield her a utility of at most $u_{i}\left(\omega_{i}^{*}\right)$.

Next, we consider the case in which $u_{c}\left(\hat{\omega}_{i}\right)<u_{c}(v)$. Then, $i$ loses the auction and her utility is zero. Moreover, $\Upsilon_{i}=\emptyset$, that is, there are no outcomes acceptable to $i$ that $i$ can choose from if she were win the auction. Accordingly, by submitting any other bid than $\hat{\omega}_{i}, i$ 's utility would be zero or less.

Finally, let us consider the case where $u_{c}\left(\hat{\omega}_{i}\right)=u_{c}(v)$. Then, $\omega_{i}^{*}$ exists and $i$ obtains utility $u_{i}\left(\omega_{i}^{*}\right) \geq 0$ with some probability $p$. Now, either $\hat{\omega}_{i}$ is a local maximum of the center's utility function $u_{c}$ on $\Omega_{i}$, or $\hat{\omega}_{i}$ is not. In the former case, consider an arbitrary $\omega_{i} \in \Omega_{i}$. Then $u_{c}\left(\omega_{i}\right) \leq u_{c}\left(\hat{\omega}_{i}\right)$, due to $u_{c}$ being equipeaked on $\Omega_{i}$. If $u_{c}\left(\omega_{i}\right)<u_{c}(v), i$ would lose the auction bidding $\omega_{i}$, and her utility would be zero. If, on the other hand, $u_{c}\left(\omega_{i}\right)=u_{c}\left(\hat{\omega}_{i}\right)$, then by bidding $\omega_{i}, i$ still obtains utility $u_{i}\left(\omega^{*}\right)$ with probability $p$ by virtue of tiebreaking being neutral.

So, for the remainder of the proof of the first claim, we may assume that $\Omega_{i}^{*}$ contains no local maxima in $u_{c}$. We first prove for all $\tilde{\omega}_{i} \in \operatorname{argmax}_{\omega \in \Omega_{i}^{*}} u_{c}(\omega)$ that $u_{i}(\tilde{\omega})=0$. For a contradiction, assume that there is some $\tilde{\omega}_{i} \in \operatorname{argmax}_{\omega \in \Omega_{i}^{*}} u_{c}(\omega)$ with $u_{i}\left(\tilde{\omega}_{i}\right) \neq 0$. Because $\tilde{\omega}_{i}$ is acceptable to $i$, clearly, $u_{i}\left(\tilde{\omega}_{i}\right)>0$. By continuity of $u_{i}$ on $\Omega_{i}$, there is an $\epsilon>0$ such that $u_{i}(\omega)>0$ for each outcome $\omega \in \Omega_{i}$ with $d_{i}\left(\tilde{\omega}_{i}, \omega\right) \leq \epsilon$, that is, there is some some ball $B_{\epsilon}\left(\tilde{\omega}_{i}\right)$ of outcomes around $\tilde{\omega}_{i}$ that are all acceptable to $i$. Moreover, since $\tilde{\omega}_{i}$ is not a local maximum of $u_{c}$ in $\Omega_{i}$, there is some $\omega^{\prime} \in B_{\epsilon}\left(\tilde{\omega}_{i}\right)$ with $u_{c}\left(\omega^{\prime}\right)>u_{c}\left(\tilde{\omega}_{i}\right)$. This, however, is at variance with our assumption that $\tilde{\omega}_{i}$ maximizes $u_{c}$ in $\Omega_{i}^{*}$, that is, that $\tilde{\omega}_{i} \in \operatorname{argmax}_{\omega \in \Omega_{i}^{*}} u_{c}(\omega)$.

It now follows that $i$ cannot do better than by bidding $\hat{\omega}_{i}$, because even if $i$ wins the auction using another bid, she must select from outcomes in
$\operatorname{argmax}_{\omega \in \Omega_{i}^{*}} u_{c}(\omega)$ to provide an acceptable outcome at least as good as the highest other bid, but we have just shown this gives her utility zero. This completes the proof of the first claim.

For the second claim, we may assume that $\geq$ is competitive and tie-breaking gives every tied winner a positive chance of winning. Let us assume that $u_{c}$ is not equipeaked over some $\Omega_{i}$, that is, there is some local maximum $\hat{\omega}_{i} \in \Omega_{i}$ of $u_{c}$ and a non-empty set $\Omega_{i}^{\prime}=\left\{\omega \in \Omega_{i}: u_{c}(\omega)>u_{c}\left(\hat{\omega}_{i}\right)\right\}$. Then, there is some $\epsilon>0$ such that $u_{c}(\omega) \leq u_{c}\left(\hat{\omega}_{i}\right)$ for all $\omega \in \Omega_{i}$ with $d\left(\omega, \hat{\omega}_{i}\right) \leq \epsilon$. Now, define $i$ 's preference so that $u_{i}(\omega) \geq 0$ if and only if $d\left(\omega, \hat{\omega}_{i}\right) \leq \epsilon$, and so that $\hat{\omega}_{i}$ uniquely maximizes $i$ 's utility.

We show now that bidder $i$ has no dominant strategy in this situation. By competitiveness, for each bidder $j$ there is a bid $\hat{\omega}_{j} \in \Omega_{j}$ such that $u_{c}\left(\hat{\omega}_{j}\right)=$ $u_{c}\left(\hat{\omega}_{i}\right)$. First, consider the case in which all bidders $j$ distinct from $i$ offer $\hat{\omega}_{j}$. If $i$ also offers $\hat{\omega}_{i}$, a tie-breaking event results and the best she can hope for is a utility of $u_{i}\left(\hat{\omega}_{i}\right)$ with probability $p<1$. By bidding some $\omega \in \Omega_{i}^{\prime}$ instead, however, she will get $u_{i}\left(\hat{\omega}_{i}\right)$ for certain. It follows that any dominant strategy of $i$ will prescribe $i$ to submit some offer from $\Omega_{i}^{\prime}$. Now consider an arbitrary $\omega_{i}^{\prime} \in \Omega_{i}^{\prime}$. By competitiveness, for each bidder $j$ there is some $\omega_{j}^{\prime}$ with $u_{c}\left(\hat{\omega}_{i}\right)<$ $u_{c}\left(\omega_{j}^{\prime}\right)=u_{c}\left(\omega_{i}^{\prime}\right)$. If any bidder $j$ other than $i$ was to submit $\omega_{j}^{\prime}$, bidder $i$ had better lose the auction and be satisfied with a utility of zero, because each outcome $\omega \in \Omega_{i}$ with $u_{c}(\omega) \geq u_{c}\left(\omega_{j}^{\prime}\right)$ yields a negative outcome. It follows that bidding some $\omega_{i}^{\prime} \in \Omega_{i}^{\prime}$ is not part of any dominant strategy for $i$ and, hence, there is no dominant strategy for $i$ at all.

We observe that in the continuous case, the set of outcomes is not assumed to be bounded. Accordingly, it also includes scenarios in which the center always wants "more." One such setting is the standard setting of auctioning a single good for a real-valued amount. We find that there, the qualitative Vickrey auction and the traditional Vickrey or second-price auction [16] are equivalent.

Example 3 (Vickrey Auction) Consider the standard setting in which a single item is auctioned for a real-valued amount. The outcomes of such an auction are given by a positive real number and a bidder, the latter specifying the winner of the auction and the former the amount the winner has to pay. Formally, let $A=[0, \infty)$. We assume each bidder $i$ to entertain a private value $v_{i} \in[0, \infty)$ for the object, and her utility function $u_{i}$ to be such that for all $(x, j) \in \Omega$,

$$
u_{i}(x, j)=\begin{array}{ll}
v_{i}-x & \text { if } i=j \\
0 & \text { otherwise } .
\end{array}
$$

Moreover, the center's order is such that for all $x, y \in \mathbb{R}$ and all bidders $i$ and $j$ we have $u_{c}(x, i) \geq u_{c}(y, j)$ if and only if $x \geq y$.

In the Vickrey auction, each bidder $i$ submits an offer $\left(a_{i}, i\right)$ from $\Omega_{i}$, where we say that $a_{i}$ is $i$ 's bid. The winner of the Vickrey auction is then the bidder $i^{*}$ with the highest bid, where possible ties are broken by some tie-breaking mechanism. The winner then has to pay the second-highest bid $a_{j}$ submitted.

As is well-known, for each bidder $i$, bidding her private value $v_{i}$ is a dominant strategy, that is, the unique bid $x$ such that $u_{i}(x, i)=0$.

We observe that in this setting each bidder's set of acceptable outcomes is bounded and closed. Moreover, $u_{c}$ has no local maxima, so the center's order is equipeaked. Also, for no bidder $i$ are there distinct $\omega, \omega^{\prime} \in \Omega_{i}$ such that $u_{c}(\omega)=$ $u_{c}(\omega)$. Hence, tie-breaking is vacuously neutral. The conditions specified in Theorem 2 are thus fulfilled. As we are now in the case where there are no local maxima in $u_{c}$, it follows from the proof of Theorem 2 that every offer $\hat{\omega}_{i}$ submitted by any bidder $i$ who plays a straightforward strategy in the qualitative Vickrey auction gives i a payoff of zero. Observe that in the present setting there is only one such offer for each bidder. The winner $i^{*}$ of the qualitative Vickrey auction can now choose from all outcomes that are ranked at least as high by the center as the second-highest bid. Choosing $\left(x, i^{*}\right)$, where $(x, j)$ is the second-highest bid, is the most profitable choice $i^{*}$ can make. Hence, the qualitative Vickrey auction and the traditional Vickrey auction are equivalent.

## 6 Related Work

The idea of applying the principle of the Vickrey auction without payments was originally introduced by Máhr and de Weerdt in a paper on auctions with arbitrary deals [7] and presented in a more formal way in a later workshop [6]. The paper at hand not only improves the presentation of this idea, placing it in a solid theoretical context, it also shows that this approach works both when the domain is finite and the center's order is linear, and when the center's order is equipeaked in the continuous setting, and gives separate proofs for these two settings.

In our framework, a payment can be part of the specification of an alternative. The results given for the continuous case thus not only generalize the traditional Vickrey auction in which payments are in fact the only parameter of the alternatives, it also generalizes multi-attribute auctions. In a multi-attribute auction each alternative is defined by a set of values (the attributes). In extant work the payments, however, are always 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. 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 [1]. His work was extended by David et al. for situations where the good is described by two attributes and a price [3]. 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 concentrated on iterative multi-attribute reverse English auctions [11]. 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-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. 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, the most important difference is that those models require a payment to transfer utility.

Also, for single-item and multi-unit auctions, there is work studying the Vickrey auction in the case where the utility functions of the agents are not linear in the payments [12]. This allows for example for situations where the value of an item is less when the payment is high because the bidder has no money left to fully exploit the item (e.g. refurnishing a house, exploiting radio frequencies, or vehicle ownership licenses in Singapore). Our main contribution in light of this paper is again that the Vickrey auction generalizes even further, to settings where there is no need for any payment. Their conditions of continuity and finiteness seem to be the equivalents of our conditions in their restricted setting with payments.

A limited number of settings is known that allow for efficient and strategyproof mechanisms without requiring transferable utility. We discuss the best known of these below: situations where the preferences of the players are singlepeaked, the house allocation problem, and the stable matching problem [14].

Preferences are single-peaked when the outcomes can be mapped onto a one-dimensional domain, each player has one preferred outcome in this domain, and each player's preferences are strictly decreasing as one moves away from her preferred outcome. In such a setting a mechanism is strategy-proof, onto and anonymous if and only if it gives the median of all peaks (possibly after inserting a number of artificial peaks) [10].

As in our case, in the house allocation problem the preferences of each player only depend on their own allocation. The house allocation problem is to find an allocation of houses to players for which there is no blocking coalition, that is, , a set of players that can be better off by trading houses among each other (of which at least one should be strictly better off). It is shown that there is exactly one allocation for which no such blocking coalition exists. This setting is quite different from ours, since in our case the problem is to select one winner.

The standard application of the stable matching problem is the problem of matching men to women or medical students to residencies. As in the house allocation problem, a group of players can block a (proposed) matching. In this case a blocking pair is a man and a woman who prefer each other over their partners in the (proposed) matching. The goal here is to find a matching without such blocking pairs. Although this is another example of a mechanism without money, it appears to have no further relation to the qualitative Vickrey auction proposed in this paper.

## 7 Conclusions and Future Work

In this paper we showed that there is another way to deal with the impossibility theorem by Gibbard and Satterthwaite besides requiring quasilinear utility functions. For settings where there is only one winner, the most important requirement is that all bidders are indifferent between all outcomes where they are not the winner (the no-externalities requirement). We proposed a protocol for settings where the preference order of the center 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 settings where payments are not necessarily possible.

We defined a class of dominant strategies for this qualitative Vickrey auctionthe straightforward strategies - and saw that the resulting outcome is weakly Pareto efficient, provided that the center's order is linear and the domain of alternatives is finite. We also found that the qualitative Vickrey auction for this setting is weakly, but not strongly, monotonic. In the case of continuous domains, we showed that the qualitative Vickrey auction is strategy-proof, and individually rational, provided that the center's utility function is equipeaked. Still, there are a number of interesting questions left unanswered regarding the properties of qualitative mechanisms such as the one presented here.

Firstly, we expect to be able to generalize the English auction in a similar manner to a qualitative auction, obtaining similar results on incentive compatibility and Pareto-efficiency. In such a setting the center accepts only bids in increasing order of the global ordering until no bidder is interested anymore, and the outcome is the last alternative that is bid. A straightforward strategy for a bidder $i$ is then to offer the highest alternative (if it is acceptable) in her preference order that is higher in $\geq$ than the last submitted bid. Another extension of this work that we would like to pursue is to generalize to multi-unit auctions, where $n$ identical items need to be allocated to at least $n+1$ agents. We do not expect to be able to generalize combinatorial auctions in a similar way, but showing the exact reasons for this is part of our future work.

Furthermore, as in many mechanism design settings, we require the center to follow the protocol. This is important, because in principle (given enough knowledge about the bidders' preferences) the center can lie about its preference order by moving the second-best bid up in its order up to just below the best bid. This will force the winner to choose her original bid, which is generally strictly better than the second-best bid. For many settings the preference order of the center will be publicly known (such as the case of a government that is transparently run, or when the center is a seller that just wants to maximize the payment), but for some settings the center may indeed behave strategically. We would like to study how to modify the mechanism to incentivize the center to be truthful in such settings (if possible).

Finally, we would like to further study potential computational an communication problems that will undoubtedly arise when using this type of auction in a variety of realistic applications. For example, in this paper we require that the
total pre-order of the center is 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 itself may not even know its preferences beforehand explicitly. If we succeed in dealing with such issues, we believe the result could be effectively used in many real-world settings, ranging from assigning programming and development tasks in open-source projects, to allocating an infrastructural project with a fixed budget to one of a number of competing construction companies.

## Acknowledgements

We are indebted to Tamás Máhr, Felix Brandt, Kevin Leyton-Brown, and the anonymous COMSOC and EC referees for inspiring discussions and valuable comments. This work is partially supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/3-2. It is also partially supported by an Alfred P. Sloan Research Fellowship and by the National Science Foundation under IIS-0812113.

## References

[1] Y. Che. Design competition through multidimensional auctions. RAND Journal of Economics, 24(4):668-680, 1993.
[2] 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.
[3] 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.
[4] A. Gibbard. Manipulation of voting schemes: a general result. Econometrica, 41(4):587-602, July 1973.
[5] J. Green and J.-J. Laffont. Incentives in Public Decision Making. NorthHolland, Amsterdam, 1979.
[6] P. Harrenstein, T. Mahr, and M. M. de Weerdt. A qualitative vickrey auction. In U. Endriss and G. Paul W, editors, Proceedings of the 2nd International Workshop on Computational Social Choice, pages 289-301. University of Liverpool, 2008.
[7] T. Máhr and M. de Weerdt. Auctions with arbitrary deals. In V. Marek, V. Vyatkin, and A. Colombo, editors, HoloMAS 2007, volume 4659 of LNAI, pages 37-46. Springer, 2007.
[8] A. Mas-Colell, M. D. Whinston, and J. R. Green. Microeconomic Theory. Oxford University Press, Inc., 1995.
[9] J. Moore. Implementation, contracts, renegotiations in environments with complete information. In J. Laffont, editor, Advances in Economic Theory, chapter 5, pages 182-282. Cambridge University Press, 1992.
[10] H. Moulin. On strategy-proofness and single peakedness. Public Choice, 35(4):437-455, 1980.
[11] D. Parkes and J. Kalagnanam. Models for iterative multiattribute procurement auctions. Management Science, 51(3):435-451, 2005. Special Issue on Electronic Markets.
[12] H. Saitoh and S. Serizawa. Vickrey allocation rule with income effect. Economic Theory, 35(2):391-401, 2008.
[13] 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(2):187-217, 1975.
[14] J. Schummer and R. Vohra. Mechanism Design without Money. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, Algorithmic Game Theory, chapter 10, pages 243-266. Cambridge University Press, 2007.
[15] Y. Shoham and K. Leyton-Brown. Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations. Cambridge University Press, Cambridge, 2008.
[16] W. Vickrey. Counter speculation, auctions, and competitive sealed tenders. Journal of Finance, 16(1):8-37, 1961.


[^0]:    *(CACM, 2009. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in Proceedings of the ACM Conference on Electronic Commerce, 2009.

[^1]:    ${ }^{1}$ An exception is if the agents can make a binding agreement with each other to collude.

[^2]:    ${ }^{2}$ 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.

[^3]:    ${ }^{3}$ It should be noted that this does not affect anything from the perspective of the bidders, in the following sense. It is generally assumed that the center can commit to the mechanism. Hence, if the center commits to a qualitative Vickrey auction that uses an order $\geq$ that does not correspond to the center's true preferences, from the perspective of the bidders, this is no different from the case where $\geq$ does correspond to the center's true preferences.

