A simulation study of two combinatorial auctions

David Nordström

Department of Economics
Lund University

Supervisor: Tommy Andersson
Co-supervisor: Albin Erlanson

May 24, 2012

Abstract

Combinatorial auctions allow buyers to express preferences over bundles of items. The Primal-Dual (PD) auction developed by de Vries et al. [5] is an efficient ascending combinatorial auction which, given certain conditions on buyers valuations, achieves Vickrey-Clarke-Groves (VCG) payments. The Universal Competitive Equilibrium (UCE) auction by Mishra and Parkes [6] is a generalization of the PD auction and achieves VCG payments under more general valuations. This study compares the PD and the UCE auction with respect to seller revenue and the number of iterations required to reach equilibrium. Simulations are performed for a fixed number of items over different levels of competition. The results indicate that for some numbers of buyers, the UCE auction yields slightly less revenue. There does not seem to be any difference in the number of iterations before termination.

Keywords: PD auction, UCE auction, Combinatorial auction, Ascending auction, VCG payments, Strategy-proof, Primal-dual algorithm

Contents

List of Figures

List of Tables

1 Introduction

Economic agents that compete over a set of items might value certain combinations differently than how they would have valued the items in solitude. For example, an airline company could have intricate preferences over different airport slots which can only be rightfully represented if they are allowed to express valuations over different combinations of slots. Or, the cost for a trucker to handle shippings in one lane could be contingent on the amount of loads in other lanes. Similarly, a construction company that participates in a procurement process for different projects might only realize economies of scale if they are assigned geographically adjacent projects. A socially optimal solution to these types of allocation problems can sometimes be solved by implementing a *combinatorial auction*. In a combinatorial auction buyers compete for many heterogeneous indivisible objects, and are allowed to express valuations over bundles of items rather than only stating preferences over single items. Combinatorial auctions have been applied in a wide variety of areas such as truckload transportation, facility sanitation, airport arrival slots, harbor planning, food programs, and allocation of spectrum licenses [4, 8].

Bikhchandani et al. [3] were among the first to formulate an assignment problem over bundles of items as a linear program. de Vries et al. [5] implemented this approach to formulate the assignment problem as a combinatorial auction. The linear program that implements the combinatorial auction is a maximization problem that obtains an *efficient allocation* with minimal seller revenue. An efficient allocation means that the buyer that values a bundle the most will win this bundle. In terms of a linear program, the primal function maximizes buyers social surplus given a set of conditions, and the dual minimizes seller revenue. Thus, the dual can be interpreted as the prices that buyers face over all possible combinations of bundles. When current prices do not fulfill an efficient allocation, prices (the dual) are adjusted upwards and towards such an allocation. This is effectively an interpretation of the combinatorial auction as an *ascending auction*, sometimes also called clock-auction.[1] An ascending auction is formally conducted by the seller or a third party auctioneer. At every round, the conductor reports current prices to the participating buyers who decides wether they are still willing to buy some bundle at the current prices. The conductor then raise prices for some set of adjusted buyers and again asks which buyers are prepared to buy bundles at the current prices. The auction finishes when prices have reached sufficiently high so that every buyers demand can be fulfilled.

Solving a combinatorial assignment problem as a linear program requires sufficient knowledge of duality theory. However, the combinatorial auction can also be solved with an algorithmic approach that implements the same solution. This is described in de Vries et al. [5] as the Primal-Dual (PD) auction. The PD auction is an iterative procedure that, via appropriate price adjustments, achieves an optimal allocation. The PD auction is efficient and will, under some

[1] see Demange et al. [7] for an early example of an ascending auction.

assumptions regarding buyers valuations, achieve truthful bidding which means that buyers can never do better than reporting their true preferences.

Mishra and Parkes [6] further develop the PD auction by examining if there is a way to achieve truthful bidding under a more general profile of buyers valuations. This resulted in the Universal Competitive Equilibrium (UCE) auction which is essentially a generalization of the PD auction. The UCE auction does indeed achieve truthful bidding under a more general setting, and is from that perspective an improvement of the PD auction. However, as will be made clear below, this property comes at the cost of stricter equilibrium requirements that might induce lower seller revenue, and a somewhat more complicated formulation that might require a higher number of iterations before an equilibrium is reached.

The PD and the UCE auction have attractive properties and as such they might be considered as candidate devices to allocate items in settings where buyers wishes to express combinatorial preferences. Thus it seems relevant to compare the performance of these auctions. The aim of this study is to compare differences between the PD and the UCE auction regarding seller revenue and the number of iterations required to reach equilibrium. This is done by constructing simulation models of the PD and the UCE auction that can be run over different number of buyers. To the best of my knowledge there has been no previous work on the comparative performance of two combinatorial auctions with respect to these issues.

The rest of this study is outlined as follows: In Section 2, I specify the general settings of a combinatorial auction. Section 3 explains truthful bidding and how this is linked to Vickrey-Clarke-Groves payments. Section 4 and 5 describe the PD and the UCE auction in more detail after which I, in Section 6, apply both auctions to an example. In Section 7, I state my research hypotheses and Section 8 describes the simulation models I use to test these hypotheses. Section 9 presents the results from the simulation. Section 10 provides a general discussion of the results and conclude by giving suggestions for future research.

For ease of exposition I will throughout this study denote the number of elements in a set X as $|X|$; and a set $X = \{1, 2, \dots, i-1, i+1, \dots\}$ as X_{-i}.

2 A general setting

Most features of the PD and the UCE auction are identical. In this section I will formulate the basic settings of a combinatorial auction that apply to both versions. The combinatorial auction described here can be seen as an ascending price auction. It is an ascending price auction in the sense that prices start at zero, then weakly increases in every round. For a technical summary, see Appendix A.

In an economy there is a finite set N of $n \geq 2$ buyers, who compete for a finite set of indivisible items, G. Every buyer $i \in N$, has a non-negative, integer valuation over a bundle $B \subseteq G$, equal to $v_i(B) \in \mathbb{R}_+$. Valuations are private information to every buyer. The set of all bundles is denoted as $\Omega = \{B \subseteq$

$G\}$. The seller values all items to zero, i.e., there are no reservation prices. Preferences are quasi-linear, so a buyer i who receives bundle B, and makes a payment of $p \in \mathbb{R}_+$ gets a net-utility of $v_i(B) - p$. Buyers have monotonic preferences so if $S \subseteq T$, then $v(S) \leq v(T)$. Other than monotonicity, no externalities are imposed on valuations.[2] Valuations over bundles do not need to be additive, i.e., it is allowed that $v_i(S) + v_i(T) \leq v_i(S + T)$. There is a null item, $\emptyset$, where $v_i(\emptyset) \equiv 0$ for every $i \in N$. The purpose of a null item is that when all prices equal valuations, a buyer can drop out of the auction by demanding the null item. In other words, receiving the null item is the same as not getting anything at all. Throughout the auction, $p(\emptyset) \equiv 0$.[3]

The combinatorial auction consists of a finite set of rounds $\mathbb{T} = \{0, 1, \ldots, T\}$. Any round $t \in \mathbb{T}$ is associated with a price vector $p^t \in \mathbb{R}_+^{|N| \times |\Omega|}$ (I will sometimes denote p^t simply as p). The price vector in the combinatorial auction is *non-anonymous*, meaning that it sets different prices for every buyer. This procedure is different from e.g. Demange et al. [7], where every buyer faces the same price for a bundle. One can think of p as a superset of every set $p_i^{|\Omega|}$ for $i \in N$ (call every such set p_i henceforth). At $t = 0$, $p_i^0 = 0$ for all $i \in N$.

2.1 The buyers

In every round, buyers report their demand set which consists of all utility-maximizing bundles, given the current price vector p. Since prices are non-anonymous, buyers only consider "their subset" $p_i \subset p$ when determining utility-maximizing demand. If a buyer demands the null item, such a buyer must receive zero net-utility from any other bundle as well. This implies that $p_i(B) = v_i(B) \ \forall \ B \in \Omega$. When a buyer demands the null item, call this buyer *inactive*.

A buyer is only allowed to change her demand set from one round to the next in a certain way. Either the demand set in round $t + 1$ must be identical to that in t, or her demand set has expanded so that it includes more bundles. Formally, if D_i^{t+1} is the demand set of buyer i in round $t + 1$, then $D_i^t \subseteq D_i^{t+1}$. This is a constraint imposed on the buyer and is called an *activity rule*.

2.2 The seller

Given buyers reported demand sets and p, the seller has to find every revenue maximizing allocation. Call one such allocation X, and the set of all such allocations $L(p)$[4]. A revenue maximizing allocation is one which maximizes the total sum of prices at p. Any allocation, be it revenue maximizing or not, can only assign an item to *one* buyer (the assignment must be feasible). At any price vector there may exist several revenue maximizing allocations. If the auction has not reached its final stage where it terminates, these allocations are only temporary. The allocation only applies for this round and the buyers will still be

[2] Items do not fulfill other characteristics such as being substitutes.

[3] Using a null item is a standard technique to achieve consistency in ascending auctions (see e.g. [7]).

[4] Call it $L(p)^*$ for the UCE auction. See Appendix A.

forced to compete for the items in the next round. Due to monotonicity (more is better), every buyer will demand G at round zero. For this reason, every revenue maximizing allocation at $t \in T$ will dispose all items to the buyers. Every buyer that does not receive a bundle under some allocation, receives the null item which is infinitely divisible. A revenue maximizing allocation must always give at least one buyer a bundle from her demand set. A buyer that receives a bundle in her demand set is satisfied. If a buyer does not receive a bundle from her demand set at some revenue maximizing allocation, this buyer is said to be *unsatisfied*.

Whenever it is not possible to satisfy all buyers simultaneously there might exist one or more sets of *minimally undersupplied buyers* (MUB). Denote the revenue maximizing allocation at p that minimizes the number of unsatisfied buyers X^*, and the number of unsatisfied buyers at this allocation as $|U|_{X^*}$. If $|U|_{X^*} > 0$, there exists at least one set of MUB. A set of buyers K are MUB if no revenue maximizing allocation can satisfy all buyers simultaneously, but there are revenue maximizing allocations that can satisfy every $K' \subset K$. For example, if $K = \{1, 2, 3\}$ is a set of MUB, there must necessarily exist at least three allocations that satisfy $\{1, 2\}$, $\{1, 3\}$, and $\{2, 3\}$; but no allocation that satisfy K. As with revenue maximizing allocations, at any p there might exist several sets of MUB. In such a case, the seller is free to choose any such set. A buyer that is inactive can never belong to K. Since such a buyer demands the null item, it will always be possible to satisfy this buyer.

2.3 Price adjustments and equilibrium

The set of MUB determine how prices will be adjusted till the next round. The seller is free to chose any such set and all buyers that are contained in that set will see a price increase. These buyers will see a price increase of 1 on every bundle that belong to their demand set at the current price vector. Bundles that did not belong to their demand set will not change in price; neither will any price change for buyers that do not belong to the selected set of MUB. Buyers that incur a price increase will either expand their demand set as new bundles become utility maximizing, or their demand set will be the same in the following round (the activity rule). Once prices have been adjusted by selecting a set of MUB, the auction moves to the next round by starting over with buyers reporting their updated demand sets. Since some prices have now been adjusted, the revenue maximizing allocations might change so that new allocations exist, and previous allocations disappear.

The procedure above is repeated until it is not possible to find any set of MUB. This directly implies that all buyers can be satisfied simultaneously in some revenue maximizing allocation. When there does not exist a set of MUB, the auction has reached a *competitive equilibrium* (CE) price vector. When a CE is reached the auction terminates. Upon termination, two things needs to be decided. First, *one* revenue maximizing allocation that satisfy all buyers must be chosen as a final allocation. At this stage all buyers and the seller will be indifferent between which such allocation is chosen. Buyers are assigned one element in their demand set, and the seller only choses allocations that maximize revenue.[5] Second, once an allocation has been decided, every buyer i makes a payment according to p_i and her assignment. If a buyer receives the null item, her payment is zero. Seller revenue is $\sum_{i \in N} p(X_i)$ where X_i is the assignment to buyer i ($X_i \in D_i$) in the final allocation. Here, payments correspond to the price vector p. This is exactly how final payments in the PD auction are determined. In the UCE auction, payments are not necessarily equivalent to p (see Section 3).

Once a CE price vector has been reached, the combinatorial auction has found an allocation that is efficient [5, 6]. It is efficient in the sense that every bundle is rewarded to the buyer that had the highest valuation for that bundle. Though a CE price vector has been reached, it does not need to be unique. Depending on how sets of MUB are chosen, the auction can take different paths towards a CE. For example, if selecting the same set of MUB over and over prices within this group of buyers might be driven "too high" [5]. There might be other buyers outside of this set that have lower valuations, and thus might become inactive at an early stage. It is "good" to obtain inactive buyers at an early stage in the sense that these buyers can never be in a set of MUB. Focusing on making buyers inactive within fewer iterations might speed up convergence to equilibrium. Though this is an important aspect of any auction there is no such "choice rule" of MUB in neither the PD nor the UCE auction, hence I will not discuss it further here.

3 Vickrey-Clarke-Groves payments and truthful bidding

In the previous section I outlined the general settings of a combinatorial auction. The auction terminated when a CE price vector had been reached and payments were then made from the buyers to the seller according to assignments in the final allocation. In any auction, final payments will affect how buyers behave in terms of bidding patterns. If an auction can achieve Vickrey-Clarke-Groves (VCG) payments then such an auction might induce truthful bidding. Buyers bid truthfully if they behave according to their actual valuations. In this section I will explain the idea of VCG payments and how these relate to the PD and the UCE auction.

The notion of VCG payments can best be understood by revisiting the seminal Vickrey (or second-price) auction (Vickrey [12]). In the Vickrey auction buyers compete for a single indivisible good. The buyer submitting the highest bid wins and pay the price of the second highest bid. In this auction it is a dominant strategy to bid truthfully. A buyers payment is independent of her bid, and she can not do anything better than to bid her valuation. For a set of buyers

[5]Note that the mechanism actually minimizes possible seller revenue given the condition that every buyer must be satisfied. This is due to formulation of the combinatorial auction as a linear program (see [6, 5, 3]).

N, let $i \in N$ be the winning buyer in a Vickrey auction. Buyer i's net-utility is $v_i - p$ where p is the second highest bid. Since no other buyer receives the item, their utility is zero. Denote the total utility of all N buyers when buyer i wins as $V(N) = \sum_{k \in N} v_k$. Now, if I remove buyer i from this auction, buyer j who submitted the second highest bid will win and total utility is $V(N_{-i})$. A *VCG payoff* for buyer i is defined as $V(N) - V(N_{-i})$ and is in this example equal to $v_i - p$, where p is buyer i's *VCG payment*. Examining the net-utility for the winning buyer i, it is clear that buyer i's payment is determined by the externality she imposes on the other buyers, $V(N_{-i})$. Since $V(N_{-i})$ is given and can in no way be influenced by buyer i, there is no possibility for buyer i to affect her payment. This is exactly what makes VCG payments desirable in an auction. Under certain settings, an allocation mechanism with VCG payments induces truthful bidding by buyers (the mechanism is said to be *strategy-proof*).

The PD auction can only achieve VCG payments under rather strict assumptions whereas the UCE auction achieves VCG payments under a more general setting. For the PD auction to terminate with VCG payments buyers valuations must fulfill a *submodularity condition*[6]. This assumption requires that a buyer contributes more to the social surplus in a small coalition than in a larger coalition. The UCE auction achieves VCG payments under any profile of valuations. It does so by departing from the payment schedule discussed in the previous section. Instead of paying the final price, buyer i who receives a final assignment X_i makes a payment of $p_i = v_i(X_i) - [V(N) - V(N_{-i})]$. In a combinatorial auction this payment is at par with the VCG payment in a Vickrey-auction. Since there are many items in a combinatorial auction, the utility of all other buyers do not need to be zero either when buyer i is present or not.

Though it is possible to attain VCG payments in both auctions, this does not necessarily imply that truthful bidding is a dominant strategy. If VCG payments are achieved and if a buyer would behave according to a false valuation profile she can never do better than if she reported truthfully. In this sense VCG payments induce truthful bidding. However, if a buyer would bid in a way that did not correspond to any valuation profile then truthful bidding can not be realized. Such a bidding pattern would occur if a buyer completely changed her demand set from one round to the next or at some round of the auction did not fulfill bundle monotonicity. For the PD and the UCE auction to achieve truthful bidding, rules that forbid such behavior must be imposed. One of these rules, the activity rule, was mentioned in the previous section. The other rule, the *bundle rule*, forbids behavior that is inconsistent with the monotonicity assumption: for any buyer i, if $B \subseteq T$ and $B \in D_i$, then $T \in D_i$. By imposing the activity rule, and the bundle rule, the PD and UCE auction achieves truthful bidding whenever VCG payments can be realized. There are however some issues over this notion. As valuations are private information buyers will not know if they satisfy the submodularity condition in the PD auction. If buyers believe that this condition is not fulfilled and that VCG payments will not be reached, they might

[6] See Appendix A.

choose to behave accordingly and report false valuations. The UCE auction is robust to such considerations as it achieves VCG payments under any valuation profile. Buyers know this and will hence behave thereafter.

4 The Primal-Dual auction

The PD auction is conducted in the same way as was explained in Section 2. Given the current price vector, buyers report their demand sets, the seller then examine revenue maximizing allocations, a set of MUB is selected, and prices are updated accordingly. Once a price vector has been reached where it is possible to satisfy every buyer, the auction terminates. Though I use the notion of a CE price vector as equilibrium condition, this concept is not used in de Vries et al. [5]. Rather, the focus is on the primal-dual algorithm and how this program is solved. No matter which interpretation that is used, the PD auction do terminate in a CE price vector (which is the optimal dual solution in the linear program). When the PD auction terminates, an efficient and feasible allocation is obtained. If buyers fulfill the submodularity condition then the PD auction terminates with VCG payments.

de Vries et al. [5] were the first to select adjusted buyers by using the notion of MUB. Adjusted buyers are those that incur a price increase from one round to the next. In the literature there are different ways of selecting such a set of adjusted buyers. MUB can be seen as a generalization of *minimal overdemanded sets of objects* developed by Demange et al. [7]. If $B \subseteq G$ then B is overdemanded if $|\{i \in N : S \subseteq B \ \forall \ S \in D_i\}| > |B|$. If overdemand holds, $\{i \in N : S \subseteq B \ \forall \ S \in D_i\}$ must necessarily contain at least one coalition of MUB. An appropriate subset can then be chosen so that it is possible to satisfy all but one buyer simultaneously.

A useful rule in the PD auction is how seller revenue updates from round to round. At any round t of the auction, where the set K^t is selected as MUB, the change in the sellers maximum revenue from t to $t+1$ is $(|K^t|-1)$. Since K^t is a set of MUB, every set K^t_{-i} can be satisfied. When prices are increased, K^t will either be MUB in $t+1$ or not. If K^t is MUB, then total price increase (change in revenue) is $|K^t|-1$ for every allocation containing $K^t_{-i} \ \forall \ i \in K^t$. If K^t is no longer MUB, then one buyer i must have increased her demand set, making it possible for all buyers in K^t to be satisfied. As the demand set of i increased, the bundle $B = X_i^{t+1}$ has the same price for buyer i in $t+1$ as in t. I will use this rule in Section 8.2 where I develop a simulation model of the PD auction.

5 The Universal Competitive Equilibrium auction

Unlike the PD auction, the UCE auction has a few important deviations from the procedure discussed in Section 2.

First, final payments do not need to correspond to the price vector (as explained in Section 3). The UCE auction maintains a single price path as in the PD auction; however, prices in the UCE auction only serves to elicit buyers preferences and does not need to correspond to final payments.

Second, every *marginal economy* where one buyer is excluded has to achieve a CE price vector. The UCE auction is a generalization of the PD auction as it considers a richer environment by including every marginal economy in the allocation problem. Let $\mathbf{N} = \{N, N_{-1}, \ldots, N_{-n}\}$. Call a group of buyers $M \subseteq N$ economy $E(M)$. If $M \in \mathbf{N} \backslash N$, $E(M)$ is a marginal economy. If $M = N$, $E(M)$ is called the *main economy* (I will sometimes denote the main economy as $E(N)$). In the UCE auction there are $n+1$ economies: the main economy and the set of all marginal economies. Economy $E(M)$ has a corresponding price vector p_M that is a projection of p on $\mathbb{R}^{|M| \times |\Omega|}$ (denote a price vector that does not include buyer i as p_{-i}). p_M is a CE in economy $E(M)$ if there is a revenue maximizing allocation such that no buyer in M is left unsatisfied. The price vector p is a *universal competitive equilibrium* (UCE) if p_M is a CE price vector of $E(M)$ for every $M \in \mathbf{N}$. Hence, compared to the PD auction, the UCE auction puts additional requirements on reaching termination. Not only does p need to constitute a CE of the main economy, it also needs to constitute a CE of every marginal economy. Unlike the PD auction where it was sufficient to reach a CE price vector in the main economy, the UCE auction requires the price vector to be tested on every economy to see if it is possible to satisfy all buyers in the marginal economies as well.

Third, MUB is generalized by considering *universally minimally undersupplied buyers* (uMUB). Let the set of buyers K_M be a set of MUB in economy $E(M)$ (The same definition of MUB as in Section 2.3 apply). The set of buyers $K \supseteq K_M$ is uMUB if every buyer $i \in K$ is MUB in some economy $E(M)$. In other words, it is possible to choose the set of uMUB in two ways in the UCE auction. Either the set can be chosen from some economy $E(M)$, or it can be chosen to include MUB from many economies. As mentioned in Section 2.3, the UCE auction does not prescribe how this selection should be done. If it is desirable to obtain inactive buyers at an early stage, choosing the largest possible set of uMUB at every round might be a good strategy.

The UCE auction can be seen as a PD auction that is repeated over $n+1$ economies. By starting the auction in $E(N)$ and only selecting uMUB from the main economy, the UCE auction is identical to the PD auction. Once a CE price vector p is reached, the PD auction terminates whereas the UCE auction proceeds to the next element in $\mathbf{N}$. If p is a CE price vector in every economy, the UCE auction can terminate. If the price vector is updated in some economy, the procedure has to be repeated over all economies again.

By considering all marginal economies the UCE auction can implement VCG payments for any valuation profile. It can be shown [6] that if p is a UCE price vector, the payment schedule defined in Section 3 is equal to $p_i^{vcg} = p_i(X_i) - [\pi(p) - \pi(p_{-i})]$, where $\pi(\cdot)$ is the sellers revenue given p. The second part of the equation, $[\pi(p) - \pi(p_{-i})]$, can be considered as a discount on the final price. The discount is the Vickrey payoff and can be considered as buyer i's marginal product in the auction. The presence of buyer i increases seller revenue by $\pi(p) - \pi(p_{-i})$, and the final price is discounted with this value. Thus, if a buyer were to report some false valuation profile she could not reduce her final payment since $p_i^{vcg}(X_i)$ is the lowest price possible at which she can receive X_i (recall the Vickrey-auction where the lowest possible payment for the winning buyer is the second highest bid). As payments might deviate from the price vector seller revenue will not correspond to that explained in Section 2.3. Instead, seller revenue in the UCE auction is equal to $\sum_{i \in N} p_i^{vcg}(X_i)$.

In Mishra and Parkes [6] the UCE auction is actually called the Universal Quasi Competitive Equilibrium (uQCE)-invariant auction. This is to highlight the fact that in temporary allocations, prior to termination, the auction achieves quasi CE price vectors. I do not explain these quasi allocations here as it would complicate the formulation without adding any relevant information for this study. Thus I use the term UCE auction instead.

The UCE auction is an improvement of the PD auction in one sense: it achieves VCG payments under more general valuations. If the rules in Section 3 are also fulfilled, biding according to ones true valuation is a dominant strategy. Thus, the UCE auction induce truthful bidding in a more general setting than the PD auction. When an auction achieves truthful bidding, it is said to be strategy-proof. Using an auction that is strategy-proof as a mean of allocating goods is usually more desirable than using one that fails to be strategy-proof. Consider for example a government that wants to allocate a set of goods to firms in a market. If the implemented mechanism is strategy-proof, firms do not have to worry about their competitors bidding strategically [13]. The possibility of market manipulation is erased as all firms will bid according to their true valuation. The seller can maximize her revenue and be certain that the final allocation will give each item to the buyer that valued it the most.[7]

6 An example

In this section I will provide examples of how the combinatorial assignment problem can be solved. Given a set of buyers and a set of items, I will first show how the problem is solved using the PD auction. After that, I apply the UCE auction on the very same problem. For other illustrative examples of the combinatorial auction, see de Vries et al. [5] and Mishra and Parkes [6]. For ease of exposition, I will call a revenue maximizing allocation an allocation throughout this section.

There is a set of buyers $N = \{1, 2, 3\}$, a set of items $G = \{A, B\}$, and a null item $\emptyset$. Denote the number of buyers in a set of MUB as $|MUB|$. The profile of buyers valuations can be described by the 3 by 3 matrix

$$V = \begin{pmatrix} v_1(A) & v_1(B) & v_1(\{A,B\}) \\ v_2(A) & v_2(B) & v_2(\{A,B\}) \\ v_3(A) & v_3(B) & v_3(\{A,B\}) \end{pmatrix} =$$

[7] Of course, for this statement to be valid requires that firms fulfill all necessary assumptions such as monotonic preferences discussed in Section 2.

$$= \begin{pmatrix} 3 & 0 & 3 \\ 0 & 6 & 6 \\ 0 & 2 & 4 \end{pmatrix}.$$

Table 1 below describes how a solution to the combinatorial problem is solved using the PD auction. For every round, demanded bundles are highlighted with a parenthesis. The very right column states maximum seller revenue of the current round. In Round 0 all prices are zero, hence seller revenue can only be zero. Note that $\{1, 2\}$ can be satisfied since they demand A and B respectively. $\{1, 3\}$ can not be satisfied since buyer 3 only demands G. For the very same reason it is not possible to satisfy $\{2, 3\}$. Thus, it is possible to choose either of these sets as MUB. In this instance, choose $\{2, 3\}$.

From Section 4, I know that seller revenue increases with 1 when $|MUB| = 2$. Due to the increase in seller revenue in Round 1, it is still possible to satisfy $\{1, 2\}$ since $p_2(B) = 1$. Since neither buyer 2 nor buyer 3 has expanded their demand sets, $\{2, 3\}$ are still a set of MUB. Since buyer 1 is assigned A in one allocation, and buyer 3 is assigned G in another, $\{1, 3\}$ is a second set of MUB. Choose $\{2, 3\}$ as MUB again. Round 2 follows the same procedure as Round 1. This time however, choose $\{1, 3\}$ as MUB.

In Round 3, prices have increased for buyer 1 as well. Since maximum revenue in this round is 3, it is not possible to satisfy $\{1, 3\}$ or $\{2, 3\}$ in this round. Again, choose $\{1, 3\}$.

In Round 4, buyer 3 is inactive and demands the null item. At the current price vector there exists three allocations: $X^1 = [X_1, X_2, X_3] = [A, B, \emptyset]$, $X^2 = [A, \emptyset, B]$, and $X^3 = [\emptyset, \emptyset, G]$. Inspection of these allocations show that in X^1 no buyer is unsatisfied. Thus, it is not possible to find a set of MUB. A CE price vector has been reached and the auction can terminate. The final assignment is $X = [X_1, X_2, X_3] = [A, B, \emptyset]$ and seller revenue is $\sum_{i \in N} p(X_i) = 2 + 0 + 2 = 4$.

Having explained how a combinatorial assignment problem is solved by the PD auction, I now proceed to apply the UCE auction to the very same problem. As there are three buyers, there is a total of four economies. For the rest of this section denote a marginal economy without buyer i as $E(N_{-i})$. To follow the structure of later sections, I will only choose uMUB from one economy in every round. As uMUB is selected from only one economy, it is more appropriate to call such a set MUB instead. I will start to look for a CE price vector of the main economy. Once a CE has been obtained, I examine wether the price vector achieves a CE in every marginal economy. This is in line with the sequential procedure suggested in Mishra and Parkes [6] for solving UCE auctions (see further Section 8.3).

The UCE auction is described in Table 2 below. The right most column show maximum revenue at the current price vector for every economy, starting with the main economy. Every marginal economy must be evaluated with respect to the maximum revenue that can be attained in that specific economy.

Since I start in the main economy, Round 0 to 3 is identical to the PD auction. Prices are updated accordingly, and maximum seller revenue is equal across all economies. Entering Round 4, a CE price vector has been obtained

Table 1: Example of a PD auction

Round		Buyer 1			Buyer 2			Buyer 3			π
	Bundle	A	B	AB	A	B	AB	A	B	AB	
	Value	3	0	3	0	6	6	0	2	4	
0	Price	0	0	0	0	0	0	0	0	0	0
	Surplus	(3)	0	(3)	0	(6)	(6)	0	2	(4)	
	$\{1, 3\}$ and $\{2, 3\}$ are MUB. Choose $\{2, 3\}$.										
1	Price	0	0	0	0	1	1	0	0	1	1
	Surplus	(3)	0	(3)	0	(5)	(5)	0	2	(3)	
	Same as Round 0. Choose $\{2, 3\}$.										
2	Price	0	0	0	0	2	2	0	0	2	2
	Surplus	(3)	0	(3)	3	(4)	(4)	0	(2)	(2)	
	Same as Round 0 and 1. Choose $\{1, 3\}$.										
3	Price	1	0	1	0	2	2	0	1	3	3
	Surplus	(2)	0	(2)	0	(4)	(4)	0	(1)	(1)	
	Same as previous rounds. Choose $\{1, 3\}$.										
4	Price	2	0	2	0	2	2	0	2	4	4
	Surplus	(1)	0	(1)	0	(4)	(4)	(0)	(0)	(0)	
	Buyer 3 is inactive. CE price vector reached.										

in the main economy. However, in economy $E(N_{-1})$, buyer 2 does not belong to any allocation. The only allocation in $E(N_{-1})$ that satisfies maximum seller revenue assigns G to buyer 3 (examine p_{-1} for this round). Thus, $\{2\}$ is a unique set of MUB in $E(N_{-1})$ (note that in all other economies, there is no set of MUB).

In Round 5, maximum seller revenue has increased by 1 in $E(N)$ and $E(N_{-3})$. The set of allocations has not changed in any economy since no buyer has expanded their demand set. Neither are the prices of buyer 2 sufficiently high to belong to some allocation. Hence, the current price vector is a CE in every economy except $E(N_{-1})$. Again, $\{2\}$ is a unique set of MUB.

In Round 6, seller revenue again increase by 1 in $E(N)$ and $E(N_{-3})$. The set of allocations has still not changed, but buyer 2's prices have now "caught up". In $E(N_{-1})$ it is possible to satisfy both buyers since one allocation assigns G to buyer 2 and the null item to buyer 3. Since $p_2(G) = 4$ and buyer 3 is inactive, the current price vector achieves a CE in $E(N_{-1})$. Also, p_M is a CE in every economy $E(M)$. A UCE price vector has been obtained and the auction terminates.

Though there are several allocations in Round 6, only one is valid as a final allocation. Any other allocation could not satisfy all buyers and can thus not be chosen. The final allocation is $X = [A, B, \emptyset]$ (same as in the PD auction). VCG payments are calculated as $p_i^{vcg} = p_i(X_i) - [\pi(p) - \pi(p_{-i})]$. Hence, $p_1^{vcg} =$

Table 2: Example of a UCE auction

Round		Buyer 1			Buyer 2			Buyer 3			$\pi(\cdot)$
	Bundle	A	B	AB	A	B	AB	A	B	AB	
	Value	3	0	3	0	6	6	0	2	4	
0	Price	0	0	0	0	0	0	0	0	0	0,0,0,0
	Surplus	(3)	0	(3)	0	(6)	(6)	0	2	(4)	
	$\{1,3\}$ and $\{2,3\}$ are MUB. Choose $\{2,3\}$.										
1	Price	0	0	0	0	1	1	0	0	1	1,1,1,1
	Surplus	(3)	0	(3)	0	(5)	(5)	0	2	(3)	
	Same as Round 0. Choose $\{2,3\}$.										
2	Price	0	0	0	0	2	2	0	0	2	2,2,2,2
	Surplus	(3)	0	(3)	3	(4)	(4)	0	(2)	(2)	
	Same as Round 0 and 1. Choose $\{1,3\}$.										
3	Price	1	0	1	0	2	2	0	1	3	3,3,3,3
	Surplus	(2)	0	(2)	0	(4)	(4)	0	(1)	(1)	
	Same as previous rounds. Choose $\{1,3\}$.										
4	Price	2	0	2	0	2	2	0	2	4	4,4,4,4
	Surplus	(1)	0	(1)	0	(4)	(4)	(0)	(0)	(0)	
	In $E(N_{-1})$ $\{2\}$ is uniquely MUB.										
5	Price	2	0	2	0	3	3	0	2	4	5,4,4,5
	Surplus	(1)	0	(1)	0	(3)	(3)	(0)	(0)	(0)	
	Same as Round 4.										
6	Price	2	0	2	0	4	4	0	2	4	6,4,4,6
	Surplus	(1)	0	(1)	0	(2)	(2)	(0)	(0)	(0)	
	p_M is a CE price vector in every economy. A UCE price vector has been obtained										

$2 - [4 - 6] = 2 - 2 = 0$, $p_2^{vcg} = 4 - [6 - 4] = 4 - 2 = 2$, $p_3^{vcg} = 0 - [6 - 6] = 0$. Seller revenue is $\sum_{i \in N} p_i^{vcg}(X_i) = 0 + 2 + 0 = 2$.

Both auctions yield the same final allocation. However, since final payments in the UCE auction must be VCG payments, prices are discounted. The PD auction does not require that prices in a CE are discounted if they fail to achieve VCG payments. For this reason, identical valuation profiles can lead to different seller revenue in the PD and the UCE auction. A second observation is that the UCE auction requires a higher number of iterations before termination. Once a CE price vector is achieved in the main economy, the PD auction terminates whereas the UCE auction must check if the price vector is a CE in every economy. In this example, this was not the case. In marginal economy $E(N_{-1})$ it was not possible to satisfy buyer 2 in any allocation. This was effectively a result of the current prices for buyer 2, but also for buyer 1. In the main economy, the low prices of buyer 2 could be "supported" by buyer 1 to attain a revenue maximizing allocation. However, in the absence of buyer 1, buyer 2's prices were not sufficiently high. Since buyer 3 would pay 4 for G, the low prices of buyer 2 could not satisfy any allocation in $E(N_{-1})$. Thus, prices for buyer 2 needed to "catch up" in order to achieve a UCE price vector.

7 Purpose

The previous sections constructively outlined the reasoning behind the research hypotheses I intend to test:

- *Hypothesis 1) For any random profile of valuations, on average, the PD auction yields higher seller revenue than the UCE auction.*
- *Hypothesis 2) For any random profile of valuations, on average, the PD auction terminates within fewer iterations than the UCE auction.*

As should be clear from Section 5, achieving a strategy-proof auction is in some respect costly. In the UCE auction payments are discounted prices and could thus decrease seller revenue as compared to the PD auction. Since the UCE auction has to find a CE in every economy, it should generally require more iterations to achieve termination than in the PD auction. The research question is relevant as it investigates if, and to what degree, there exists a trade-off between implementing a strategy-proof auction *or* possibly achieving faster termination and higher revenue. It is generally accepted [9, 11, 2] that participation in an auction is costly as buyers must learn the mechanism and decide how to bid in every round. From this perspective, an auction that terminates as quick as possible is desirable. The seller in an auction wants to maximize her revenue and might also want to induce truthful bidding. Clearly, these objectives might be conflicting. From this perspective it is relevant to know the costs, if any, of implementing a strategy-proof auction rather than one that is not strategy-proof.

8 Constructing the simulation models

To test wether the PD auction yields higher revenue and terminates faster than the UCE auction I develop a simulation model of each auction. In this section I will describe the general construction of these models and discuss how I expect them to perform. Each auction is formulated as an algorithm that searches for an equilibrium in a way that is consistent with the theory. The core of these algorithms will be the CE algorithm which I will describe below. Then, following the pattern of previous sections, it will be convenient to describe the PD algorithm, after which I explain the UCE algorithm.

The simulations will be executed for different number of buyers. Buyers will always compete for exactly three items ($G = 3$). Keeping the number of items constant and increasing the number of buyers effectively shows how changes in the level of competition affects the performance of the auctions. Both algorithms initiates by drawing buyers valuations for the items in G from the uniform distribution with support $v = \{0, 1, \ldots, 25\}$. The value of bundle B to buyer i is $v_i(B) = \sum_{b \in B} v_i(b)$, i.e. valuations are additive. I allow valuations to be additive in order to speed up calculations. This should not inflict any bias in the results. If non-additivity (as described in Section 2) was allowed, both algorithms should increase revenue and number of iterations equally much. Every buyer and the seller has the objectives discussed in Section 2, and the buyers obey the rules defined in Section 3.

8.1 The CE algorithm

The CE algorithm can be described by Figure 1. As is clear from the name, the purpose of the algorithm is to find a CE price vector. Since MUB and uMUB drives the auction forward, the CE algorithms focus is on finding such sets. The algorithm uses the profile of buyers valuations and the current price vector as input. The input is used to calculate the demand set of every buyer, given the current price vector. The maximum achievable seller revenue for the current round is identified by some rule (different rules are used for the PD and the UCE algorithm, see below). When the algorithm initiates, the price vector is taken as input and then updates within the algorithm. Once the the demand sets and maximum revenue of the current round are known, the algorithm searches for a set of MUB. This procedure makes use of the following rule:

- *Rule 1) With $|G|$ items, the maximum number of buyers that could constitute a set of MUB is $|G| + 1$ (Propisition A.1).*

This rule gives an upper bound on how many buyers I can expect in a set of MUB. At the first round, buyers will typically only demand the full bundle G. Since maximum revenue is zero (all prices are zero), it will normally only be possible to satisfy one buyer simultaneously, implying that in any set of MUB there can only be two elements. The algorithm will for every round start to look for a set $|MUB| = 4$. $|MUB| = 4$ can only exist if there are at least *four* revenue maximizing allocations satisfying *three* buyers. Denote the set of active buyers in N as N_+. For for $\{i, j, k, l\} \in N_+$ to be MUB, it must be true that $\{D_i, D_j, D_k\} \in X^1$, $\{D_i, D_j, D_l\} \in X^2$, $\{D_i, D_k, D_l\} \in X^3$, and $\{D_j, D_k, D_l\} \in X^4$. Similarly, $|MUB| = 3$ can only exist if there are at least *three* revenue maximizing allocations that can satisfy *two* buyers. However, it can not not be possible to satisfy all three buyers simultaneously. Thus, for $\{i, j, k\} \in N_+$ to be MUB, it must be true that $\{D_i, D_j\} \in X^1$, $\{D_i, D_k\} \in X^2$, $\{D_j, D_k\} \in X^3$, and $\{D_i, D_j, D_k\} \notin X$ for any $X \in L(p)$. The process is identical for $|MUB| = 2$. For such a set to exist, there must be at least *two* allocations satisfying *one* of the buyers. $|MUB| = 1$ only requires that there is no revenue maximizing allocation that can satisfy some buyer (remember the definition in Section 2.2).

The process above formalized a general rule that controls the process of obtaining a set of MUB. Denote a set of active buyers K, where $|K| \in \{1, 2, 3, 4\}$.

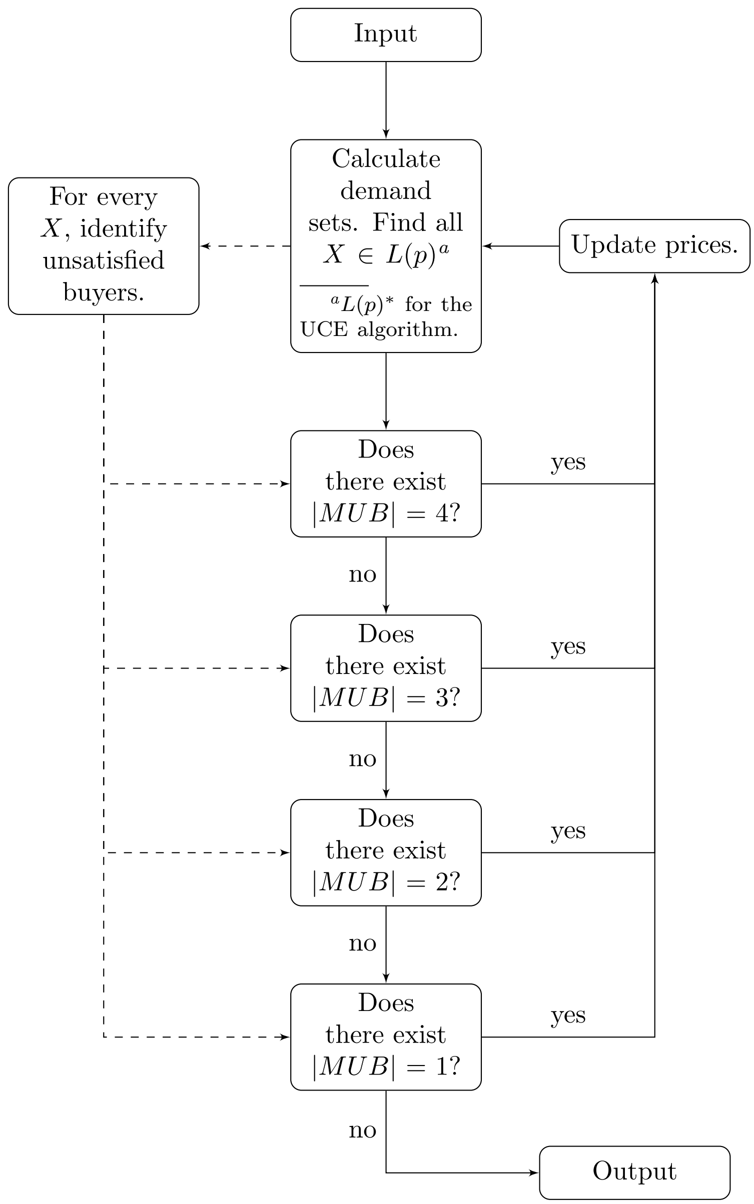

Figure 1: Flow chart of the CE algorithm.

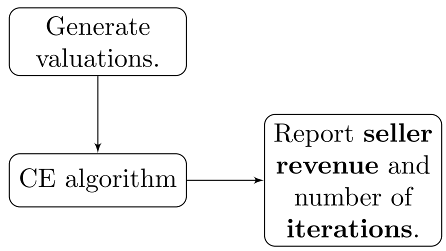

Figure 2: Flow chart of the PD algorithm.

For $|MUB| = |K|$ to exist, there must be at least $|K|$ revenue maximizing allocations satisfying K_{-i} for every $i \in K$, and there can be no allocation $X \in L(p)$ such that every buyer in K is satisfied. The CE algorithm starts at $|K| = 4$ and then work through every possible set of MUB in a descending order to $|K| = 1$. Whenever a set of MUB is identified, the algorithm cancels the search, update prices, and proceeds to the next round. If no set of MUB can be found, the algorithm terminates. The reported output in the CE algorithm is the CE price vector p, and the number of iterations to termination.

The CE algorithm follows the general procedure outlined in de Vries et al. [5] and Mishra and Parkes [6]. Collecting unsatisfied buyers for every $X \in L(p)$ is a simple method to obtain MUB candidates. Of course, every such candidate group has to be examined to make sure that there is no allocation satisfying all these buyers simultaneously. Due to the assumptions made on valuations, there will generally be $|MUB| = 2$ initially. As buyers will begin to demand bundles of only one item only when prices has reached a certain level, $|MUB| = 3, 4$ will generally only exist in later rounds of the algorithm. The CE algorithm selects candidate buyers with a queuing order q. For example when considering $|MUB| = 3$, the queuing order is $q = [\{1, 2, 3\}, \{1, 2, 4\}, \cdots, \{n-2, n-1, n\}]$. Thus, the algorithm will initially only select MUB from the "first" buyers in N. As a result, prices tend to increase for these buyers until they become inactive. At this stage the revenue maximizing allocations yield a sufficiently high level of revenue so that other buyers are not satisfied in any $X \in L(p)$. Each one of these buyers will hence constitute separate sets of $|MUB| = 1$. Prices will then adjust for one buyer at the time until that buyer "catches up" in prices. This will result in a rather large number of iterations before termination. As the number of buyers increases, this effect will be more and more apparent. The algorithm will however sort buyers contingent on their valuations so that the buyer that value G the least is placed first. In the light of the discussion in Section 2.3, this might somewhat reduce the number of iterations before termination.

8.2 The PD algorithm

The PD algorithm is described in Figure 2. As is clear from the figure, only a few amendments to the CE algorithm are necessary to obtain the PD algorithm.

To determine seller revenue in every round of the PD algorithm, the following rule is used:

- *Rule 2) From round t to round $t+1$, seller revenue increases by $|K^t| - 1$ (see Section 4).*

In the initializing round seller revenue is zero. For any following round where termination does not occur, there will be a set of MUB. This set determines the evolution of seller revenue from one round to the next. Rule 2 reduces calculations since it is not necessary to search through all demand sets and the price vector to obtain maximum revenue in a round. Instead, using Rule 2, the algorithm can directly test if a set of buyers can fulfill a revenue maximizing allocation or not. This in turn reveals wether these buyers are unsatisfied and should be considered as MUB or not. Once a CE price vector is obtained, the algorithm reports seller revenue and number of iterations and terminates. Due to Rule 2, it is not necessary to determine a final allocation.

8.3 The UCE algorithm

The UCE auction respects marginal economies and VCG payments, hence the algorithm requires a more dynamic structure than the PD algorithm. The UCE algorithm can be explained by Figure 3.

The algorithm initiates in the main economy where it searches for a CE price vector p_M. Once the price vector is obtained, the algorithm proceeds to check if p_M is a CE in $E(M)$ for every $M \neq N$. If so, a UCE price vector is reached. The algorithm proceeds to calculate VCG payments, report seller revenue and number of iterations, and terminates. If p_M is not a CE in some $E(M)$, the algorithm updates the price vector in that economy until a CE price vector p'_M is reached. The new price vector p'_M is then run for *every* economy. If p'_M is a CE in every economy, a UCE price vector is obtained. Otherwise, the price vector is again updated in the economy where it was not a CE, and is checked against every economy.

Sets of uMUB are chosen *only* from the economy that the algorithm is currently working in. From Section 5 it is clear that there might exist larger sets of uMUB other than that in some specific economy. This procedure might have a positive effect on the number of iterations required to reach termination.

It is not always possible to calculate seller revenue in the UCE algorithm with Rule 2. More specifically, whenever the algorithm changes from one economy to another, Rule 2 does not apply. Since the algorithm selected uMUB from one economy, the change in revenue in another economy might be different as the revenue maximizing allocations could be different. Whenever the price has been updated once in $E(M)$, Rule 2 can be used as long as the algorithm stays in $E(M)$. Whenever the algorithm goes from $E(M)$ to $E(M')$ (once a CE is reached in $E(M)$) a function calculates maximum seller revenue by using demand sets and the current price vector.

To calculate VCG payments it is necessary to select a final allocation. The UCE algorithm does so by searching , in a descending order, for groups of three,

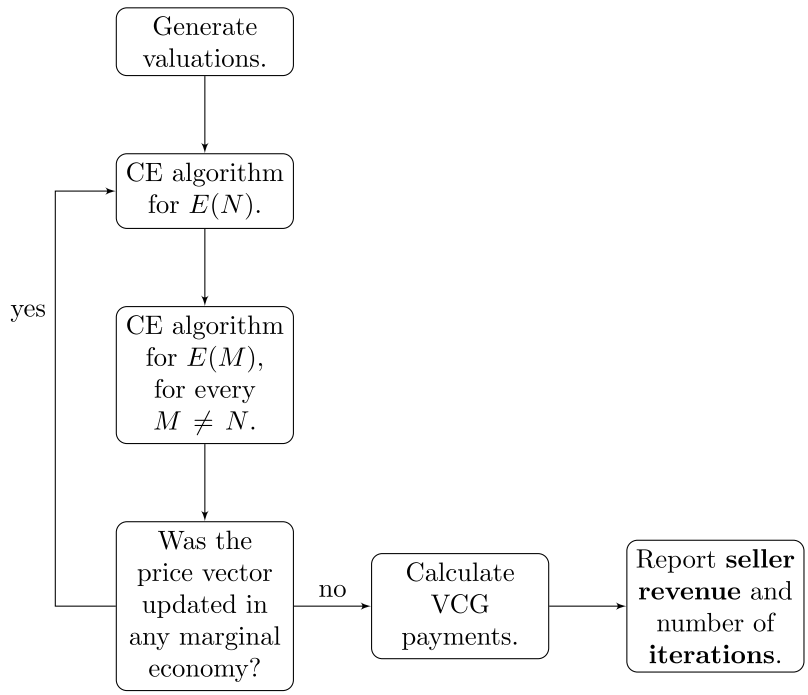

Figure 3: Flow chart of the UCE algorithm.

two or one buyers that satisfy maximum seller revenue. Once such a group of winners is identified, the algorithm can calculate VCG payments by examining maximum seller revenue for each marginal economy that does not include one winner.

The UCE algorithm follows the sequential approach suggested in Mishra and Parkes [6]. By ordering the elements in $\mathbf{N}$, starting with $E(N)$, one economy at a time is considered. In this study, the procedure on one hand simplifies the construction of the algorithm since the differences in calculations between economies are minimal. On the other hand, it does not fully utilize the concept of uMUB since only buyers from one economy is chosen.

9 Results

In this section I will present the results from the simulations. The results are intended to answer wether the PD auction on average yields higher seller revenue and requires fewer iterations than the UCE auction. As will be explained further below, the results suggest that discrepancy in revenue increases in the number of buyers, i.e. with higher levels of competition in the economy, whereas the number of iterations tends to be the same over both auctions.

I run simulations for 2 to 10 buyers with 1.000 repetitions for every auction and every number of buyers (totally 18.000 repetitions). Of course, results become more reliable with more repetitions. The algorithms described above are essentially based on running tests over permutations of buyers. Thus, increasing the number of buyers drastically increases necessary calculations for each iteration. Also, more buyers generally requires more iterations before termination. For example running a PD auction over 2 buyers takes approximately 0.052 seconds whereas a PD auction for 3 buyers takes 0.243 seconds. Considering the time it takes to run an auction for larger number of buyers, I chose 1.000 repetitions as a suitable number. A summary of the results are given in Table 3.

Table 3: Summary statistics

	UCE			PD		
Buyers	Max	Min	Mean	Max	Min	Mean
	Revenue					
2	56	0	24.542	59	1	24.768
3	68	9	37.604	69	10	37.610
4	70	13	45.188	70	18	45.812
5	70	22	50.531	69	22	50.849
6	73	28	54.028	72	31	54.936
7	73	32	56.569	73	33	57.521
8	72	34	58.716	74	31	59.627
9	74	42	59.983	74	46	61.331
10	74	42	61.570	74	43	62.727
	Iterations					
2	56	0	24.542	59	1	24.768
3	111	10	55.373	118	13	55.336
4	162	22	89.204	156	31	89.561
5	207	54	125.575	209	46	123.435
6	270	81	161.360	256	82	161.396
7	306	95	197.633	284	96	196.111
8	351	141	233.573	356	94	232.241
9	382	163	269.115	382	163	269.474
10	554	189	308.647	426	186	305.423

9.1 Average revenue

The general pattern of average revenue with 2-10 buyers can best be understood by Figure 4. Revenue tends to increase at a diminishing rate for both the PD and the UCE auction. This is a reasonable result as the competitive effects when introducing one more buyer ought to be higher over initially smaller number of buyers. In the case of 2 buyers, there will be a relatively high number of allocations that satisfy both buyers in the early rounds of the auction. Introducing a third buyer decreases the likelihood of finding an allocation that satisfy all buyers early in the auction and tends to drive final prices and revenue higher. When there are 4 buyers, a final allocation must assign the null item to at least one buyer. This implies that at least one buyer must have $p = v$ across all bundles, i.e., the buyer is inactive. For 4 buyers and above there can be at most three active buyers in any final allocation. Thus, all other buyers must be made inactive by raising prices equal to valuation.

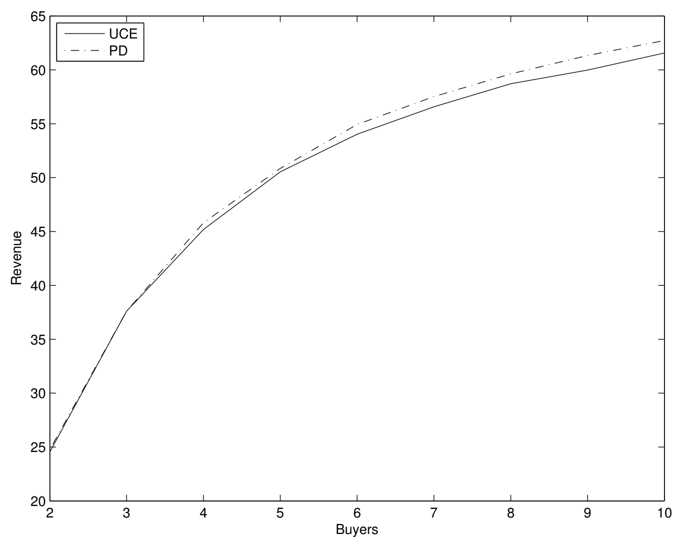

Figure 4: Average seller revenue in the PD and the UCE auction.

Increasing the number of buyers results in a more homogeneous valuation profile in the sense that the distribution of valuations approaches the uniform. As the valuation profile becomes more homogeneous the effect on revenue of the marginal buyer decreases. In other words, making one buyer inactive has more effect on revenue for smaller number of buyers. With higher number of buyers there does not need to be significant relative changes in the price vector to make many buyers inactive. As the number of buyers increases, it becomes more likely that valuations of the marginal buyer is "contained" within the valuation profile of the other buyers. Based on these results, I can now turn to research hypothesis 1:

- *Hypothesis 1) For any random profile of valuations, on average, the PD auction yields higher seller revenue than the UCE auction.*

Comparing average revenue of the PD and the UCE auction in Figure 4, there seems to be a clear tendency of divergence when the number of buyers increases.

To test wether this can be statistically supported, I apply a two-sided t-test allowing unequal variance, and calculate confidence intervals (CI) for each number of buyers. The results are displayed in Table 4. The CI uses the PD auction as a base. Thus, a strictly positive CI implies that, for some significance level, the PD auction yields a higher revenue than the UCE auction. As CIs including zero are not relevant I do not report these here. The significance level for calculating CIs is 1%.

Table 4: Comparison of average revenue.

Buyers	2	3	4	5	6	7	8	9	10
t-statistic	0.479	0.013	1.500	0.879	2.755	3.293	3.376	5.553	5.185
p-value	0.632	0.989	0.133	0.379	0.006	0.001	0.000	0.000	0.000
CI	-	-	-	-	[0.058, 1.757]	[0.206, 1.697]	[0.215, 1.606]	[0.722, 1.973]	[0.581, 1.732]

For 2 to 5 buyers I can not reject the null of equal means at any relevant significance level (given the sample sizes I would not consider t-tests with p-values above 5% to give any reliable indications on differences in revenue). Going from 5 to 6 buyers significantly improves p-values and for 6 buyers and above I can reject the null at significance levels below 1%. Examining relevant CIs, there does seem to be a tendency of revenue divergence as the number of buyers increases. The CIs tend to be higher, and the span of every CI narrower. The statistical tests imply that with low numbers of buyers the PD and the UCE auction yield similar revenue; and with higher numbers of buyers the PD auction yield increasingly higher revenue than the UCE auction. The magnitude of this difference is rather small, implying that a random draw of a profile of valuations for e.g. 9 buyers will yield a difference in revenue of only about 1 to 2. A difference in revenue means that the PD auction does not achieve VCG payments and that final prices in the UCE auction are discounted.

It does not seem plausible that VCG payments would be less frequent in the PD auction when the level of competition increases. This would imply that the submodularity condition is more likely to be fulfilled for lower numbers of buyers. Since the valuation profile becomes more homogeneous in the level of competition it would rather be reasonable to assume that the submodularity condition was more likely to be fulfilled when there are many buyers present.

The observed differences in average revenue might be explained by statistical properties of the auctions that arise due to the way valuations are drawn. Since the relative difference in revenue between the auctions is rather small, it can only be small differences in revenue between the main and the marginal economies. As the number of buyers increases the profile of valuations become more homogeneous, and thus it becomes more likely that 3 buyers will be assigned one item each. This in turn implies that there will be more strictly positive prices in the final allocation. In the UCE auction a price will be discounted if the revenue without that particular buyer would have been lower. This occurs with a positive probability and thus when the number of non-zero prices increase, there is a higher probability that at least some price will be discounted. At the same time, however, there will be smaller relative differences between one buyers valuations to that of others. This will have an opposing effect on differences in revenue, as a marginal buyer contributes less and less to seller revenue when the number of buyers increase. The results from Table 4 implies that the first effect of increasing the probability of higher numbers of strictly positive prices outweighs the probability that revenue in a marginal economy is the same as that of the main economy.

The reasoning above suggests that differences in average revenue would not increase continuously as the number of buyers increased. Rather there would be some critical number (or region of numbers) of buyers after which the probability that there would always be three buyers that paid strictly positive prices in the final allocation would be close to 1. At this point, adding new buyers would only move the profile of valuations closer to the uniform distribution which would reduce the difference between revenue in the main and marginal economies. Thus it would be possible that above some number of buyers the difference in revenue would begin to decrease. Changing the number of items, or the distribution of valuations, would affect when and if such a level of competition would be reached.

9.2 Average number of iterations

Average number of iterations for the PD and the UCE auction can be seen in Figure 5. From a graphical analysis iterations seem to increase in a fairly linear manner. Since at most three buyers can be active in a final allocation, all other buyers need to be made inactive by raising their prices equal to valuation. As the construction of the algorithms will tend to force a subgroup of buyers to compete amongst themselves until they become inactive, all other buyers will at this stage constitute sets of $|MUB| = 1$. The price vector will then be increased for one buyer at a time until they become inactive or a final allocation can be reached. Price paths like this will be the usual case and thus adding one more buyer requires that the process is repeated for this buyer too. Hence the linear pattern in Figure 5 seems reasonable. I now turn to research hypothesis 2:

- *Hypothesis 2) For any random profile of valuations, on average, the PD auction terminates within fewer iterations than the UCE auction.*

Examining Figure 5, there does only seem to be any differences in number of iterations. To investigate this observation further, I apply the same statistical tests as for Hypothesis 1. The results are displayed in Table 5. I can find no support that there is a significant difference in the number of iterations at a 5 % level or below. Though the p-values are quite close to the 5 % level for 5 and 10 buyers, there is no apparent difference between these levels of competition to others. It should for example not be possible to more easily obtain sets of MUB that contain more elements for these number of buyers. Hence, it seems likely that these results are caused by outliers.

Since both auctions start in the main economy, this means that most of the times the CE price vector in the main economy is also a CE in all marginal economies. This might be explained by the fact that there can be at most 3 active buyers in a final allocation. Hence, as all other buyers are made inactive it might be trivial to satisfy all buyers in the marginal economies. The main economy consists of all buyers and thus there might just be a few, if any, changes that has to be made for the initial CE price vector in the main economy to be a UCE price vector. This reasoning is supported by the rather small differences that were obtained for average revenue.

Table 5: Comparison of average number of iterations.

Buyers	2	3	4	5	6	7	8	9	10
t-statistic	0.479	-0.049	0.368	-1.896	0.027	-1.121	-0.862	0.216	-1.816
p-value	0.632	0.960	0.712	0.058	0.972	0.262	0.388	0.829	0.069

As uMUB were chosen from only one economy at the time, this could have had a positive effect on the number of iterations before termination. The results from the simulations suggests that by always picking uMUB from a specific economy, the UCE auction will on average not yield more iterations than the PD auction. Since the UCE algorithm did not consider the maximum sizes of the sets of uMUB and still did not give higher numbers of iterations, picking the maximum size sets might even result in the UCE algorithm terminating within fewer iterations.

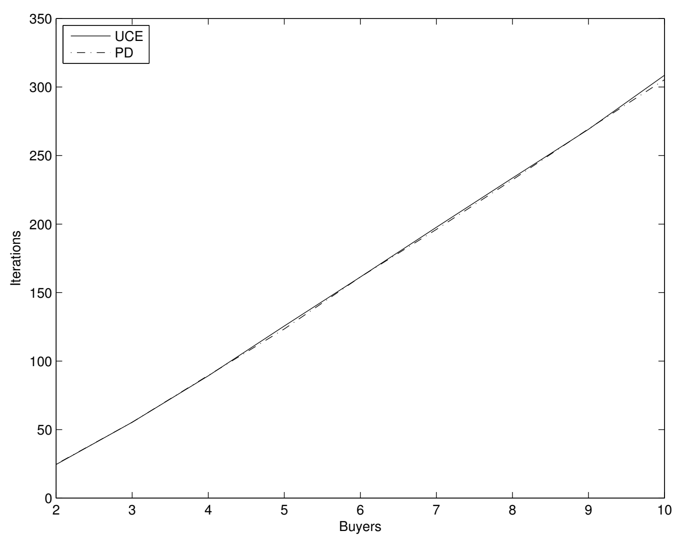

Figure 5: Average number of iterations in the PD and the UCE auction.

A last observation related to my results is that when there are 2 buyers, the number of iterations is identical to seller revenue (see Table 3). This follows directly from Rule 2 and how adjusted buyers are chosen. In the first round both buyers will either be satisfied, or both buyers will be in the unique set of MUB. For any round where the auction does not terminate, seller revenue increases by 1 in the main economy. Thus, both buyers will belong to the unique set of MUB (and uMUB) until an allocation is found that satisfy both buyers. This occurs at the T'th iteration where seller revenue is equal to T. In the marginal economies, buyers are trivially satisfied in every round. Since both buyers see a price increase in every round prior to termination, revenue increases in increments of 1 for every round and maximum revenue is the same across all economies in every round. Thus, for any given valuation profile with 2 buyers, revenue and number of iterations in the PD and the UCE auction should always be the same. For the case of 2 buyers there are more items than buyers and thus it is possible that seller revenue is zero (this could also happen if there are 3 buyers). Hence, a seller that wants to be certain to raise revenue when there are more or equal number of items and buyers must specify strictly positive reservation values.

10 Discussion

This study presented two different combinatorial auctions where one, the UCE auction, yielded VCG payments under more general settings than the other, the PD auction. It was argued that since VCG payments required that buyers payments sometimes were discounted, the UCE auction would on average yield lower seller revenue than the PD auction. A second argument directly related to fulfilling VCG payments was that the UCE auction would on average require a higher number of iterations since every marginal economy needed to terminate in a CE price vector.

To test wether these arguments were valid, I constructed simulation models of both auctions. The results showed that as the number of buyers increased, seller revenue diverged though the absolute difference remained small. I claimed that this was caused by the fact that as competition increased, it became more probable that a buyer only received one item, rather than a bundle containing many items. Hence, more buyers would pay a strictly positive price in a final allocation. Since this effectively meant that more prices were subjected to a possible discount, this might explain the divergence in revenue. This argument also suggested that a continuously increasing divergence was unlikely and that above some number of buyers, revenue would begin to converge as the valuation profile became more homogeneous. There was no apparent difference between the auctions with regards to the number of iterations before termination.

My results suggests that with regards to the properties studied here, there are only small, nonetheless significant, differences between the PD and the UCE auction. Thus the choice of which auction to implement is contingent on what the auctioneer wants to achieve. If the auctioneer is strictly concerned with maximizing revenue, the PD auction might be preferable. However, if buyers reporting truthfully is important, implementing the UCE auction will only yield small differences on achievable revenue. Concerning the time aspect, there does not seem to be a difference between the auctions, hence this aspect might be ignored when deciding on which auction to execute.

Of course, my simulations and the results that can be inferred from these are limited by the number of items, and the number of buyers. Due to time constraints I could only run simulations to a certain extent. Thus the conclusions that I have presented should in no way be interpreted as applying for *every* setting where an auctioneer wants to allocate items by conducting a combinatorial auction. Rather my results point to the fact that under certain conditions, there seems to differences in how much seller revenue the PD and the UCE auction raise, and this might be explained by reasons mentioned above.

Combinatorial auctions are essentially intricate systems that can be studied from many different perspectives. As such there were many aspects, both theoretical and practical, that due to the limited time available to finish this study were not covered. The environment in which the PD and the UCE auction were described do not allow for e.g. seller reservation values or some types of buyer valuations such as substitute values. To be able to make normative statements about when and how these auctions can be implemented in real-world situations requires further research.

Allowing buyers to express valuations over bundles of items can have both positive and negative effects on the general performance of an auction. The positive effects can be related to the valuation profiles of buyers and their information. First, both auction formats discussed in this study respect some characteristics of buyers valuations that are likely to be observed in some markets. It is allowed that items are perfect complements so that a buyer might have a strictly positive valuation only for a bundle and value each individual item within this bundle to zero. Complementaries could for example exist in spectrum auctions where the possibility to control two geographically adjacent licenses is worth more than the sum of the individual licenses. The PD and UCE auction allow for such possibilities. If buyers were restricted to submit independent bids for both licenses, this could lead to inefficiencies. At the risk of winning only one license, buyers might either refrain from bidding, or they might be forced to bid above their valuation in order to secure the synergies. In this sense, the combinatorial auction increases efficiency. Second, the combinatorial auction does not necessarily require that all buyers reveal their true willingness to pay. Since preferences are private information, buyers might be reluctant to state their willingness to pay. Unlike one-shot auctions (e.g. the Vickrey auction), ascending auctions does not necessarily require that buyers reveal all their private information. A second observation related to the information structure of these auctions is that buyers are not required to express their valuation of every bundle at every round. Buyers only have to report the currently utility maximizing bundles and can remain silent about their ranking over other bundles. The possibility to express preferences over bundles rather than items and not being forced to completely reveal true valuations are two features that makes the PD and the UCE auction attractive as allocation mechanisms.

Though both auctions discussed in this study do account for some types of valuations, these are far from completely covering those that could occur in reality. It is easy to see that, for example, the monotonicity assumption could break down for a variety of reason. Say that a contractor is participating in a bidding process of receiving building contracts in three different geographical areas, A, B, and C. Area A and B are neighbor regions but area C is far away from both these regions. The contractor estimates that it is possible to extract 200 each in rent from every region. The necessary machines for construction costs 300. Since A and B are adjacent it is possible to use the same machines in both areas but area C, being located far away from A and B, requires a second set of machines. Thus the contractor would value the bundle $\{A, B\}$ to 200+200-300=100. However, due to the geographic distance, the bundle $\{A, B, C\}$ would be worth 200+200+200-300-300=0. The contractor would fail to satisfy the assumption of monotonicity and hence neither the PD nor the UCE auction would be applicable. Though this is a greatly simplifying example, it stresses the fact that when applying a certain combinatorial auction, the required assumptions need to be carefully evaluated against the actual situation.

Another problem with allowing bids over bundles of items is the increased complexity required to handle such a bidding procedure. When there are G items there is a total of 2^{G-1} bundles that needs to be evaluated with respect to possible revenue and reported demand sets. As the number of bundles increase exponentially in the number of items, the computational requirements rather quickly become vast. On behalf of the seller, determining an efficient allocation can thus be rather demanding. Also, for the buyers to bid optimally they must be able to understand the rules of the auction, and appropriately estimate their valuations. In this sense, combinatorial auctions are costly for the seller and the buyers since participation is encompassed by transaction costs.

10.1 Future research

There is still a lot to learn about combinatorial allocation problems. This study could for example be improved by considering higher number of buyers and conducting more repetitions. It would be interesting to see if the results presented here holds in an experimental setting. Also, with the rather large number of combinatorial auctions available, a general comparison of performance should be conducted for families of combinatorial auctions.

Theoretical research should aim at relaxing constraints to make combinatorial auctions more general and hence more applicable in the field. For the same reason, rules for e.g. iteration minimization should be examined. It would be interesting to explore the submodularity condition in a combinatorial setting further by examining under which general conditions, with respect to number of buyers and the profile of valuations, it will hold.

References

[1] Lawrence Ausubel M. An efficient dynamic auction for heterogeneous commodities. *Working paper, University of Maryland*, 2002.

[2] Patrick Bajari, Stephanie Houghton, and Steve Tadelis. Bidding for incomplete contracts: An empirical analysis. Working Paper 12051, National Bureau of Economic Research, February 2006.

[3] Sushil Bikhchandani and Joseph Ostroy M. The package assignment model. *Journal of Economic Theory*, 107(2):377 – 406, 2002.

[4] Peter Crampton, Yoav Shoham, and Richard Steinberg. *Combinatorial Auctions*. MIT Press, 2006.

[5] Sven de Vries, James Schummer, and Rakesh. Vohra V. On ascending vickrey auctions for heterogeneous objects. *Journal of Economic Theory*, 132(1):95–118, January 2007.

[6] Mishra Debasis and David Parkes C. Ascending price vickrey auctions for general valuations. *Journal of Economic Theory*, 132(1):335 – 366, 2007.

[7] Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94:863–872, 1986.

[8] Rafael Epstein, Lysette Henríquez, Jaime Catalán, Gabriel Y. Weintraub, and Cristián Martínez. A combinational auction improves school meals in chile. *Interfaces*, 32(6), November 2002.

[9] Shmuel Gal, Michael Landsberger, and Arkadi Nemirovski. Participation in auctions. *Games and Economic Behavior*, 60:75–103, 2007.

[10] Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982.

[11] William F. Samuelson. Competitive bidding with entry costs. *Economics Letters*, 17(1-2):53–57, 1985.

[12] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):pp. 8–37, 1961.

[13] Xia Zhou, Sorabh Gandhi, Subhash Suri, and Haitao Zheng. ebay in the sky: Strategy-proof wireless spectrum auctions. In *Proceedings of the 14th ACM international conference on Mobile computing and networking*, MobiCom '08, pages 2–13. ACM, 2008.

A Appendix

Here I will make a more technical summary of the PD and the UCE auction. I start with defining the common elements in the auctions. Everything mentioned here should be seen as an extension to Section 2. The same definitions of buyers, seller, items, bundles and valuations apply here. Without distorting the general results of the PD and the UCE auction, I will not be as explicit in my definitions as de Vries et al. [5], and Mishra and Parkes [6]. The interested reader should consult these papers for a complete description of the auctions.

A combinatorial auction seeks to find an efficient allocation of bundles over the buyers in an economy. An allocation X is a feasible vector of bundles on buyers. X assigns a bundle, possibly $\emptyset$, to *every* buyer in the economy. For any X it is required that $X_i \cap X_j = \emptyset \ \forall i, j \in N$. Hence, any X can, unlike items in G, assign the null item to multiple buyers. Due to the assumption of monotonicity, any X will have $\bigcup_{i \in N} X_i = G$. The set of all feasible allocations is $\mathbf{X}$. A feasible allocation X is said to be *efficient* if there does not exist another allocation $Y \in \mathbf{X}$ such that $\sum_{i \in N} v_i(Y) > \sum_{i \in N} v_i(X)$.

At any instance of the auction there is a price vector $p \in R_+^{|N| \times |\Omega|}$, for which buyer i has an associated demand set:

$$D_i(p) = \{B \in \Omega : v_i(B) - p_i(B) \leq v_i(T) - p_i(T) \ \forall \ T \in \Omega\}.$$

The supply set of the seller in the PD auction under the same price vector p is denoted:

$$L(p) = \{X \in \mathbf{X} : \sum_{i \in N} p_i(X_i) \geq \sum_{i \in N} p_i(Y_i) \ \forall \ Y \in \mathbf{X}, \text{and } X_i \in D_i(p) \cup \{\emptyset\}\}$$

The supply set in the UCE auction under the same price vector p is denoted:

$$L^*(p) = \{X \in \mathbf{X} : \sum_{i \in N} p_i(X_i) \geq \sum_{i \in N} p_i(Y_i) \ \forall \ Y \in \mathbf{X}, \\ \text{and } X_i \in D_i(p) \cup \{\emptyset\}\} \ \ \forall M \subseteq N$$

From the definition above it is fully possible that X does not correspond to the demand set of buyer i. The demand set and supply set capture the standard notion that buyers and sellers are maximizing utility/revenue. Of course, every buyer can have several bundles in her demand set, and a seller can have several allocations in her supply set.

A.1 Prices

Adjustment towards equilibrium in the combinatorial auction is completely channeled through a price path. A price path is defined as a function $P : T \rightarrow$

$R_+^{|N|\times|\Omega|}$, where T is the set of rounds in the auction. $P(t)$ is a mapping from every round $t \in T$, to the corresponding price vector p. This type of price path is *non-linear* and *non-anonymous*. Allowing a non-linear price path means that it is possible to have $p(B)+p(T) \leq p(B+T)$ for $B, T \in \Omega$ (much like non-linear valuations). A non-anonymous price path maintains individual prices over Ω. In other words, it is fully possible that at any point in the auction, two buyers can face different prices for the very same bundle. $P(t)$ is non-decreasing over time and thus it can be called an ascending price path. Furthermore, P is single so that it is only possible to have one price path.[8]

A.2 Minimally undersupplied buyers

At price vector p, buyers might demand the same items. Also, the seller could have several revenue maximizing allocations. When these problems occur, it might not be possible to fulfill the demand of some buyers simultaneously. As long as $\emptyset \notin D_i(p)$, a buyer is said to be active. Denote the set of active buyers in an economy as N_+. I will call buyer i unsatisfied if $X_i \notin D_i(p)$.[9] Denote the set of unsatisfied buyers $K \subseteq N_+$ at price vector p and allocation $X \in L(p)$ as $U(K, X, p)$. The total number of unsatisfied buyers at some $U(K, X, p)$ is $|U(K, X, p)|$. Find an allocation $X^* \in L(p)$ such that $|U(K, X^*, p)| \leq |U(K, X, p)|$ for all $X \in L(p)$. A set of buyers $K \subseteq N_+$ are said to be *undersupplied* at price vector p if and only if $|U(K, X^*, p)| > 0$. The coalition $K \subseteq N_+$ is a set of *minimally undersupplied buyers* (MUB) if $|U(K_{-i}, X, p)|$=0 for every $i \in K$. In other words, for every $K' \subset K$ it must be possible to find an allocation $X \in L(p)$ such that no buyer in K' is unsatisfied. Note that even though some buyers are undersupplied in $X \in L(p)$ this does not necessarily imply that there exists a set of MUB. As long as there is *some* $X \in L(p)$ that satisfy all active buyers, there is no set of MUB.

I will now prove a useful proposition related to the concept of MUB that will be used in the simulation models in Section 8.

Proposition A.1. *If there are $|G|$ items, the maximum number of buyers that could constitute a set of MUB is $|G|+1$.*

Proof. The proof follows readily from the definition of MUB. Consider only the active buyers $N_+ \subseteq N$ at current price vector p. Assume that $|N_+| \geq |G|+1$. Note that for some set of buyers $K \subset N_+$, where $|K| = |G|$, an allocation X that satisfy every buyer in K must assign *exactly one item* in G to every buyer. Now include one more active buyer k, and call this set of buyers $K+ = \{K, k\}$. In the set $K+$ it is no longer possible to satisfy all buyers since there are too few items. For the set $K+$ to be MUB, it must be possible to find some allocations $X \in L(p)$ that satisfy $K+_{-i}$, for every $i \in K+$. This is only possible if every $K+_{-i}$ allocates *one* item from G to every buyer. Generally $|G|+a$ buyers, where $a \geq 2$, can never constitute a set of MUB since at least a buyers will always be unsatisfied at any $X \in L(p)$. □

[8] Other auctions allow for multiple price paths, see e.g. Ausbel [1].

[9] A buyer that is unsatisfied must by definition be active.

A.2.1 Universally minimally undersupplied buyers

At price vector p, K is a set of *universally minimally undersupplied buyers* (uMUB) if $K \cap M_+$ is a set of MUB in $E(M)$ for some $M \in \mathbf{N}$. Identifying MUB in this instance is done with $L^*(p)$, rather than $L(p)$. uMUB is a generalization of MUB as it considers not only MUB in the main economy, but also sets of MUB in the marginal economies.

A.3 The submodularity condition

If $V(M \cup \{i\}) - V(M) \geq V(M' \cup \{i\}) - V(M')$ holds for every $M \subseteq M' \subseteq N$ and all $i \in N$, the submodularity condition is fulfilled. Thus, the submodularity condition implies that buyer i's contribution to the social surplus is greater for smaller coalitions than for larger coalitions. In Section 6, the profile of valuations does not fulfill the submoularity condition.

A.4 An ascending price combinatorial auction

The following definition describes the PD auction as an algorithm using previously defined notions.

Definition A.1. *The price path starts at $P(0)$ and all prices are set to zero. The auction ends at $P(T)$ where final allocation and payments are decided. At any round t, where $0 \leq t \leq T$, the procedure is as follows:*

- Step 1) Every buyer i reports her demand set D_i, given the current price vector p^t.
- Step 2) From the reported demand, the seller identifies every $X \in L(p^t)$.
- Step 3) Using the information in 1) and 2), a set of MUB, K^t is chosen. If there is no MUB, go to step 6.
- Step 4) Every buyer $i \in K^t$ see a price increase on bundles in their demand set. If $B \in D_i(p^t)$, then $p_i^{t+1}(B) = p_i^t(B) + 1$. For all other buyers $p_i^{t+1}(B) = p_i^t(B)$.
- Step 5) Proceed to round $t+1$ by repeating from step 1.
- Step 6) The auction terminates. Bundles are allocated according to $X \in L(p^T)$ where X_i is the assignment to buyer i. The payment for every buyer is $p_i^T(X_i)$.

Clearly, the demand set will only change from round t to $t+1$ for buyer $i \in K^t$. Every buyer $j \notin K^t$ will be unaffected by the price adjustments in terms of demanded bundles.

To obtain the UCE auction, replace $L(p)$ with $L^*(p)$ throughout Definition A.1, and replace Step 3 as

- Step 3.1) Using the information in 1) and 2), a set of uMUB, K^t is chosen. If there is no uMUB, go to step 6.

Also, replace step 6 with

- Step 6.1) The auction terminates. Bundles are allocated according to $X \in L(p^T)$ where X_i is the assignment to buyer i. The payment for every buyer is $p_i^T(X_i) + [\pi(p^T) - \pi(p_{-i}^T]$.