INSTITUT NATIONAL DE LA STATISTIQUE ET DES ETUDES ECONOMIQUES
Série des Documents de Travail du CREST
(Centre de Recherche en Economie et Statistique)

n° 2007-25

The Ausubel-Milgrom Proxy Auction with Final Discounts*

L. LAMY [1]

* I am grateful above all to my Ph. D. advisor Philippe Jehiel for his continuous support. I would like to thank seminar participants in Paris-PSE and Vienna ESEM 2006 Conference. All errors are mine.
[1] Laboratoire d'Economie Industrielle, CREST-INSEE, 28 Rue des Saints-Pères, 75007 PARIS, France. laurent.lamy@ensae.fr

The Ausubel-Milgrom Proxy Auction with Final Discounts*

Laurent Lamy†

*I am grateful above all to my Ph.D. advisor Philippe Jehiel for his continuous support. I would like to thank seminar participants in Paris-PSE and Vienna ESEM 2006 Conference. All errors are mine.

†Laboratoire d'Economie Industrielle, CREST-INSEE, 28 rue des Saints-Pères 75007 Paris. e-mail: laurent.lamy@ensae.fr

Abstract

We slightly modify the Ausubel-Milgrom [3] Proxy Auction by adding a final stage which possibly induces some discounts relative to the final accepted bids of the 'original' auction. The proxy auction with final discounts is such that the outcome is a point in the bidder optimal frontier of the Core. Then truthful reporting is an equilibrium if and only if the Vickrey outcome is in the Core, a condition that is necessary but not sufficient in the original version of the proxy auction as illustrated by an example.

Keywords: Auctions, multi-unit auctions, Core, Vickrey implementation
JEL classification: D44, D45

Abstract

Nous modifions légèrement l'enchère proxy d'Ausubel-Milgrom [3] en ajoutant une étape supplémentaire qui correspond à des réductions au regard des prix finaux par rapport à la dynamique de l'enchère d'origine. L'enchère proxy avec des remises implémente une allocation dans la frontière des optima parétiens du Coeur. Reporter ces préférences de manière sincère est alors une stratégie d'équilibre si et seulement si l'allocation de Vickrey est dans le coeur, une condition qui est nécessaire mais pas suffisante dans la version d'origine du mécanisme, comme nous l'illustrons par un exemple.

Mots-clés: Enchères, Enchères multi-unitaires, Coeur, Allocation de Vickrey
Classification JEL: D44, D45

1 Introduction

Ausubel and Milgrom [3] (A&M henceforth) introduce an ascending proxy auction which is supposed to combine the advantages from Vickrey insofar as the efficient allocation is implemented in dominant strategy and from first price 'menu' auctions insofar as it is robust to shill bidding and losers' deviation. On the one hand, robustness to losers' deviation is satisfied very generally and follows from the fact that the final outcome of the proxy auction belongs to the Core relative to the reported preferences. On the other hand, truthful reporting is a dominant strategy only if the final outcome is the Vickrey outcome. Since the outcome of the proxy auction lies always in the Core, this condition is satisfied only if the Vickrey outcome lies in the Core. However, A&M derives a stronger sufficient condition: implementation of the efficient allocation in dominant strategy is obtained under a buyer-submodularity condition, which is equivalent to the condition that the Vickrey outcome lies in the Core for any set of bidders. Generically, it corresponds to the condition that goods are substitutes in assignment problems without allocative externalities.

A natural question raised by A&M's analysis is then to characterize for the ascending proxy auction the full set of preferences that implements the efficient outcome in truthful dominant strategy. The natural candidate are preferences such that the Vickrey outcome is in the Core. However, A&M do not clarify whether this condition is sufficient or not.

First, this note provides an example such that the Vickrey outcome is in the Core but where the Ausubel-Milgrom proxy auction does not lead to the Vickrey outcome. Second the main contribution of this note is to propose a slight modification of the auction by adding a final stage, where the seller awards some discounts to the final prices, such that the mechanism always leads to a bidder-optimal frontier outcome. Those final discounts are a kind of 'Vickreyfication' of the proxy auction giving better incentives for truthful reporting. For this modified auction, and more generally for the class of auctions leading to the bidder-optimal frontier, the set of preferences such that truthful reporting is a Nash Equilibrium is perfectly characterized: the Vickrey outcome must be in the Core. Without ambiguity, the final discount stage is an improvement from A&M perspective: it enlarges the set of bidders' preferences such that truthful reporting is an equilibrium whereas the modified auction keeps the desirable properties resulting from the 'Core membership' of the final outcome. Moreover, the truthful equilibrium of the proxy auction with final discounts can be implemented by a dynamic mechanism leaving some privacy about bidders' valuations. Finally, in the background of our analysis, we clarify the status of the final outcome of the 'original' Ausubel-Milgrom proxy auction: it implements a payoff in the weak bidder-optimal frontier relative to the reported preferences, i.e. some -but not all- bidders' payoff may be raised such that the outcome remains in the Core.

Another strand of the auction literature uses a seemingly different approach to reach dynamically the Vickrey allocation: Clock Auctions [8, 10, 6, 1, 2] are mimicking a Walrasian tâtonnement. Demange, Gale and Sotomayor [8] and Gul and Stacchetti [10] propose dynamic auction mechanisms with differentiated commodities which are converging to the smallest Walrasian prices and truthful reporting is hence an equilibrium if and only if those Walrasian prices coincide with Vickrey's. The Bikhchandani and Ostroy [5]'s linear programming formulation of the assignment model allows de Vries et al [7] to link the two approaches by interpreting the aforementioned clock auctions as a primal-dual algorithm and A&M's auction as a subgradient algorithm.

Ausubel's clock auctions ([1, 2]) are implementing the Vickrey outcome by means of an auctioneer whose announced prices are converging to a Walrasian equilibrium price vector. However, contrary to [8, 10], a 'clinching' rule disconnects the prices that are paid with the closing prices of the auctioneer making truthful reporting an equilibrium for a larger set of preferences. For general valuations and for a larger class of ascending price auctions, Mishra and Parkes [14] generalize the idea of using Walrasian prices to reveal preferences and then to implement the Vickrey payoffs via price discounts. Our final discount stage presents a similarity with such 'clinching' rules: the pricing rule gets closer to Vickrey's. We express our idea in the perspective of A&M's package auction. Nevertheless, as in [14], our discount stage applies more generally, e.g. also for auctions corresponding to primal-dual algorithms.

This note is organized as follows. Section 2 introduces the assignment problem and the related Core concepts. Section 3 defines the algorithm which concisely characterizes the Ausubel-Milgrom proxy auction. The algorithm is the iteration of a mapping whose fixed points are Core outcomes. Section 4 gives an example where the final outcome is not in the bidder-optimal frontier of the Core. Section 5 concludes by proposing the final discount stage modification and characterizes the set of preferences that renders the truthful strategy an equilibrium.

2 Model and notation

There are N buyers (With a slight abuse of notation, N will represent the set as well as the number of buyers) indexed by $l = 1, \ldots, N$ and a seller designated by $l = 0$. For any set of buyers $S \subset N$, denote by S^* the set $S \cup \{0\}$. Denote by M the finite set of indivisible items to be auctioned. We define an allocation as an assignment of the items, denoted by $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1, \cdots, \mathcal{A}_N)$ where $\mathcal{A}_l \in M$ specifies the items acquired by agent l.

Denote by $\mathbf{A}$ the set of feasible allocations, i.e. such that $\bigcup_{l\in N^*}\mathcal{A}_l = M$ and $\mathcal{A}_l \cap \mathcal{A}_k = \emptyset$. Each buyer l has a valuation vector $\Pi_l = (\Pi_l(\mathcal{A}), \mathcal{A} \in \mathbf{A})$, where $\Pi_l(\mathcal{A}) \geq 0$ specifies the value of allocation $\mathcal{A}$ to bidder l. Denote by $\Pi := (\Pi_1, \ldots, \Pi_N)$ the vector of all buyers' types. As in A&M, we limit and simplify the set of preferences by the following assumptions.

First, we consider a private values framework where each agent is privately informed about his preferences and where one's valuation depends solely on his private signal and not on his opponents' signals. Hence, we exclude any informational externality. Second, we exclude any allocative externality: an agent's valuation on an assignment depends solely on the set of items that he acquires. Third, we consider that the seller is indifferent to the final allocation.[1] Fourth, we consider that a buyer obtains his lowest valuation when he acquires no item[2]. Fifth, we consider that agents are risk neutral.

To summarize, if bidder $l \in N$ pays a bid $b_l(\mathcal{A})$ such that the allocation $\mathcal{A}$ is chosen, then he earns a net payoff of $\Pi_l(\mathcal{A}) - b_l(\mathcal{A})$, where $\Pi_l(\mathcal{A})$ depends only on $\mathcal{A}_l$. On the other hand the seller's payoff is her revenue $\sum_{l\in N} b_l(\mathcal{A})$. The payoff of a bidder who acquires nothing and pays nothing is thus normalized to zero.

We define a (feasible) outcome as an $N+1$-uple $(\mathcal{A}, (b_l)_{l\in N})$ where $\mathcal{A}$ is the allocation chosen and b_l the price paid by buyer l to the seller. Equivalently, a vector of net payoffs $(\pi_l)_{l\in N^*}$, such that there exits an allocation $\mathcal{A}$ such that $\sum_{l\in N^*}\pi_l = \sum_{l\in N^*}\Pi_l(\mathcal{A})$,will be referred to as a (feasible) outcome. Given a set of players and of preferences, we face an allocation problem $\{N, (\Pi_l)_{l\in N}\}$. Our perspective is to implement the efficient allocation denoted $\mathcal{A}^*$, which maximizes the total welfare, i.e. $\mathcal{A}^* \in Arg\max_{\mathcal{A}\in\mathbf{A}}\{\sum_{l\in N}\Pi_l(\mathcal{A})\}$.

As a useful tool for the following analysis, we first characterize the coalitional form game (N^*, w) associated with the allocation problem $\{N, (\Pi_l)_{l\in N}\}$, where N^* is the set of players and w is the coalitional value function. For any coalition of buyers $S \subset N$, w is defined by the following expression:

$$w(S^*) = \max_{\mathcal{A}\in \mathrm{A}} \sum_{l\in S} \Pi_l(\mathcal{A}) \quad ; \quad w(S) = 0$$

In particular, it means that if the seller is not a member of the coalition, then the coalition obtains no items.

[1] Indeed, all the analysis can be extended to a framework with allocative externalities provided that non-purchasers are indifferent to the final assignment, i.e. they can 'escape to the moon' in Jehiel and Moldavonu's [12] terminology, and a non-neutral seller provided that she is not strategic. The generalization of the proxy auction with allocative externalities has been investigated by Ranger [15].

[2] This is a weaker form of the 'free disposal' assumption made in A&M and Ranger [15]. It is indeed sufficient in their analysis.

Then we define the set of core outcomes, denoted by $Core(N^*, w)$, related to this coalitional value function w:

$$Core(N^*, w) = \left\{ (\pi_l)_{l \in N^*} \mid (a): \sum_{l \in N^*} \pi_l = w(N^*);\ (b):\ (\forall S \subset N^*)\ w(S) \leq \sum_{l \in S} \pi_l \right\}$$

(a) is the feasibility condition meaning that a Core outcome implements the efficient allocation, whereas inequalities (b) mean that the payoffs are not blocked by any coalition S.

Remark 2.1 *The outcome resulting from a transfer of payoffs from a given buyer l to the seller remains in the Core if the initial outcome is in the Core and provided that π_l remains nonnegative. This comes from the fact that inequalities (b) when $0, l \in S$ are not altered, whereas such inequalities with $0 \notin S$ are always satisfied provided that $\pi_l \geq 0$, i.e. the individual rationality constraint is satisfied, and that inequalities (b) with $l \notin S$ are only strengthened. In particular, the outcome such that $\pi_0 = w(N^*)$ and $\pi_l = 0$ for all $l = 1 \dots N$ belongs to the Core which is thus non empty.*

Another specific outcome is the Vickrey outcome, denoted by $\pi^V := (\pi_l^V)_{l \in N^*}$, such that buyer l's payoff π_l^V equals to $w(N^*) - w(N^* \setminus \{l\})$ and the seller receives the revenue $\pi_0^V = w(N^*) - \sum_{l \in N} \pi_l^V$. A main issue in the analysis of the Ausubel-Milgrom proxy auction is whether the Vickrey outcome is a Core outcome. This is equivalent to the fact that the set of Pareto-optima from the perspective of the buyers is a singleton (A&M theorem 6). This set will be qualified as the bidder-optimal frontier of the core.

Definition 1 *The bidder-optimal frontier (respectively the weak bidder-optimal frontier) of the core is the set containing the elements $(\pi_l)_{0 \leq l \leq N^*} \in Core(N^*, w)$ such that there exists no outcome $(\pi'_l)_{0 \leq l \leq N^*} \in Core(N^*, w)$ with $\pi'_l \geq \pi_l$ for all $l = 1 \dots N$ and such that at least one inequality is strict (respectively with $\pi'_l > \pi_l$ for all $l = 1 \dots N$).*

On the one hand, the bidder-optimal frontier is a standard concept in the literature. In particular, Berheim and Whinston [4] have established that the outcomes of the coalitional-proof equilibria of the first price 'menu' auction coincide with the bidder-optimal frontier. On the other hand, we are not aware of any previous work in multi-unit auctions that focuses on the weak bidder-optimal frontier. Nevertheless, we show in next section that, in general, the Ausubel-Milgrom proxy auction ends in the weak bidder-optimal frontier.

A&M has introduced a condition that makes truthful reporting a dominant strategy in the ascending proxy auction: buyer-submodularity.

Definition 2 (Buyer submodularity) *The coalitional value function w is buyer-submodular if for any $l \in N$ and any coalitions S and S' satisfying $l \in S \subset S' \subset N$, we have*

$$w(S^*) - w(S^* \setminus \{l\}) \geq w(S'^*) - w(S'^* \setminus \{l\})$$

The term $w(S^*) - w(S^* \setminus \{l\})$ represents the surplus associated with the presence of bidder l in the coalition S^*. It is a kind of 'substitutes' condition: the bidders should be viewed as substitutes insofar as the surplus associated with the presence of a bidder is non-increasing with the set of competitors.

3 The Ausubel-Milgrom proxy auction

We first define the ϵ-Ausubel-Milgrom proxy auctions ($\epsilon > 0$). The Ausubel-Milgrom proxy auction is the mechanism defined by taking the limit $\epsilon \to 0$.[3] As in A&M, those sealed bid mechanisms can be viewed as the outcome of the ascending package auction where each bidder l has instructed a 'proxy agent' that bids on his behalf according to a straightforward bidding strategy parameterized by ϵ which reflects the increment used by the bidders to reduce their target profit. Indeed what is below referred to as the Ausubel-Milgrom proxy auction is a specific version of the family of proxy auctions defined by A&M. Considering the whole family would not modify the insights but introduce cumbersome notation. A&M slackens their analysis by considering a family of straightforward bidding strategies. Here we give more structure to the definition of the straightforward strategies used by proxy bidders. In A&M, the parameter ϵ corresponds to the upper bound on the bid increments. Here it reflects the increment used for target profit reductions. Furthermore, A&M restricts their analysis to the limit $\epsilon \to 0$, whereas all our results are valid for any increment ϵ.[4]

Definition 3 *The ϵ-Ausubel-Milgrom proxy auction (A^ϵ, b^ϵ) is the function mapping Π, the vector of reported valuations, into an outcome in $\mathbf{A} \times \mathbb{R}^N$*

[3]The outcome of ϵ-generalized ascending proxy auctions lies in the compact set of feasible and individually rational payoffs. So there exists a sequence $(\epsilon_n)_{n \in \mathbb{N}}$ such that the limit exists.

[4]Note also that the ascending proxy auction in A&M is not fully characterized: the final outcome still remains ambiguous (except under the buyer-submodularity condition). In their definition of the proxy auction, they consider the limit where the upper bound of bid increments are negligibly small $\epsilon \to 0$. For example, each bidder could use in its straightforward strategy a different increment ϵ_l such that $\epsilon_l = \lambda_l \cdot \epsilon$. Up to a normalization, each specific choice of $(\lambda_l)_{l=1,\cdots,N}$ defines a specific mechanism at the limit $\epsilon \to 0$. Implicitly, our analysis is restricted to a uniform increment across bidders, i.e. $\lambda_l = 1$. To convince himself that the Ausubel-Milgrom proxy auction is sensitive to the choice of the increments, it is left to the reader to check that, in our example in section 4, if $(\lambda_l)_{l=1,\cdots,N}$ is chosen such that λ_1 is sufficiently smaller than $\lambda_2 = \lambda_3 = \lambda_4$, then the Ausubel-Milgrom proxy auction (without final discounts) would converge to the Vickrey outcome unlike the case with uniform increments.

according to the following algorithm:

1 **Initialization** $a := (\emptyset, \ldots, \emptyset)$; $\widehat{\pi}_l = \max_{\mathcal{A} \in \mathbf{A}} \Pi_l(\mathcal{A})$; $y := 0$
2 **While** $y = 0$ **do**
3 $\quad y := 1$
4 $\quad$ **for** $l = 1$ *to* N **do**
5 $\qquad$ **if** $\Pi_l(a) < \widehat{\pi}_l$
6 $\qquad\quad$ **then** $\widehat{\pi}_l := \max\{0, \widehat{\pi}_l - \epsilon\}$; $y := 0$
7 $\qquad b_l(\mathcal{A}) = \max\{0, \Pi_l(\mathcal{A}) - \widehat{\pi}_l\}$ $(\forall l \in N, \forall \mathcal{A} \in \mathbf{A})$
8[5] $\quad a := Arg \max_{\mathcal{A} \in \mathbf{A}} \left(\sum_{l=1}^{N} b_l(\mathcal{A}) \right)$
9 $A^\epsilon(\Pi) := a$; $b_l^\epsilon(\Pi) := b_l(a)$

The function mapping $(a, (\widehat{\pi}_l)_{l \in N}) \in \mathbf{A} \times \mathbb{R}^{+N}$ *into itself according to the preceding loop 'while' (lines 2-8) is qualified as the* $T^\epsilon - mapping$*. The numbers* $\widehat{\pi}_l$ *are referred to as* target profits*. For a bidder* l*, if* $b_l(\mathcal{A}) = \Pi_l(\mathcal{A}) - \widehat{\pi}_l$*, then allocation* $\mathcal{A}$ *is referred to as a* target allocation *relative to bidder* l*.*

Note that the $N+1$-uples $(a, (\widehat{\pi}_l)_{l \in N})$ are not necessary (feasible) outcomes. That is the reason why we use the 'hatted' $\widehat{\pi}$ for target profits instead of π , which is the notation used for outcomes.

As in A&M, this algorithm could be interpreted as a dynamic auction with the original buyers represented by proxy bidders whose strategies are entirely determined by the reports Π_l to the proxy. We comment then the algorithm in this perspective. At each round of this relating dynamic auction algorithm (i.e. for each iteration of the loop 'while') and for each bidder l, there are two kind of allocations. First, there are the target allocations for which the profit equals to the target profit $\widehat{\pi}_l$ if this allocation is chosen by the seller. The target profit $\widehat{\pi}_l$ also corresponds to the highest profit that is conceivable at the current round for agent l (for any possible continuation of the algorithm). Second, there are the allocations for which the profit is less than the target profit and for which the bids are still null (line 7).

Given the current bids of the proxies, the algorithm chooses a current allocation which maximizes the auctioneer's payoff, i.e. that maximizes the sum of all agents bids (line 8).

The tricky part of the algorithm is the dynamic revision of the bids by the proxies (line 4-7). Two events may arise for bidder l when his is

[5] In the event of a tie between different allocations, we assume at least that the efficient allocation is chosen if available. This tie-breaking rule, that is not made in A&M, guarantees that the algorithm stops exactly when it reaches the Core. This is just a technical trick that guarantees that the weak bidder-optimal frontier are fixed points of the $T^\epsilon - mapping$. Otherwise, the algorithm would stop anyway when it reaches the interior of the Core.

asked to revise his bids. First, the current allocation chosen is a target allocation (if $\Pi_l(a) \geq \widehat{\pi}_l$), then he does not change his bids. Second, if the current allocation is not a target allocation (if $\Pi_l(a) < \widehat{\pi}_l$, line 5), then the buyer reduces his target profit by ϵ (line 6) (provided that his target profit remains positive, else he stops at his null target profit), which roughly corresponds to outbid ϵ on target allocations (line 7). When he reduces his target profit, it possibly corresponds to bid on new allocations: the target allocations sets are increasing sets as the algorithm goes along. For the other allocations the bids are still sticking to zero. Such a strategy is referred to as a *straightforward bidding strategy* in A&M. The algorithm stops when no agent outbids his previous bids (in the algorithm, the binary variable y reflects this information: $y = 1$ meaning that all agents have been inactive in the previous round.).

The rest of this section establishes the link between the final outcome of the proxy auction with Core outcomes. A&M has proved that the proxy auction terminates at a Core outcome which is roughly equivalent to the fact that fixed point of T^ϵ are in the Core. Hatfield and Milgrom [11] show that the Ausubel-Milgrom proxy auction is a cumulative offer processes that approaches the Core from above in term of bidders payoffs and ends at the bidder-optimal frontier under a 'substituability' condition. This condition enables them to apply fixed point theorems in lattices: the Core has then a lattice structure implying that the bidder-optimal frontier is a singleton and thus that the Vickrey outcome is in the Core. Nevertheless, the intuition of the tâtonnement from above is not restricted to such 'substitutes' preferences as will be stated in proposition [3.2] where we show that the final outcome always converges to the weak bidder optimal frontier relative to the reported preferences. Before, the following proposition characterizes the Core as the set of fixed points of the map T^ϵ, an approach which is similar to Echenique and Oviedo [9].

Proposition 3.1 *The set of fixed points of the $T^\epsilon - mapping$ are the Core outcomes.*

Proof 1 *The following lemma provides another expression for the revenue of the seller after the application of the $T^\epsilon - mapping$.*

Lemma 3.1 *The seller maximum revenue at the current allocation a that maximizes her revenue when the vector of target profit is $(\widehat{\pi}_l)_{l\in N}$ is given by:*

$$\widehat{\pi}_0 = \max_{a\in\mathbf{A}} \sum_{l=1}^{N} b_l(a) = \max_{S\subset N} \left\{ w(S^*) - \sum_{l\in S} \widehat{\pi}_l \right\} \tag{1}$$

Proof 2 *This is a careful rewriting by reversing the order of the 'max' operator. For more details, refer to A&M pp. 20. The seller's maximization* *program can be viewed as two stages: first choose the set of buyers S that will pay a strictly positive amount and obtain consequently a target allocation, second given this set of buyers choose the allocation that maximizes their contribution. For the target allocations of buyers in S there is no distortion between the bids and the preferences (they are all equal up to a constant) and the second stage then consists in choosing the allocation they jointly prefer, i.e. corresponding to the surplus $w(S^*)$.*

Suppose that $(a,(\widehat{\pi}_l)_{l\in N})$ is a fixed point of the $T^\epsilon-$mapping. For any fixed point, each bidder obtains his target profit, i.e. $b_l(a) = \Pi_l(a) - \widehat{\pi}_l$. Hence, for any other allocation a', we have:

$$\sum_{l\in N}\Pi_l(a) - \sum_{l\in N}\widehat{\pi}_l = \sum_{l\in N} b_l(a) \geq \sum_{l\in N} b_l(a') \geq \sum_{l\in N}\Pi_l(a') - \sum_{l\in N}\widehat{\pi}_l$$

The first inequality uses the definition of allocation a which maximizes the seller's revenue. The second inequality follows from the bidding strategy that corresponds to the target profits. The inequality between the two extreme implies that allocation a is the efficient allocation and the feasibility constraint is then satisfied. At a fixed point, target profit equals realized profit and equation (1) is satisfied with $\widehat{\pi}_l = \pi_l$ and the blocking constraints of the Core are thus fulfilled. Finally we have proved that fixed points of T^ϵ are Core outcomes.

Now suppose that $(a,(\widehat{\pi}_l)_{l\in N})$ is a Core outcome. Then we have both: $\sum_{l\in N} b_l(a) = w(N^*) - \sum_{l\in N}\widehat{\pi}_l$ and $\sum_{l\in N} b_l(a) \geq w(S^*) - \sum_{l\in S}\widehat{\pi}_l$. Finally, the maximum in equation (1) is reached for the coalition N, i.e. the entire set of bidders receive their target profit. Hence, due to our tie-breaking rule restriction (see footnote 5), the efficient allocation is chosen and each agent obtains his target profit. Finally, it means that no agent reduces his target profit and then that the outcome is a fixed point

The ϵ-Ausubel-Milgrom proxy auction corresponds to the iteration of the mapping T^ϵ with the initial target profits being equal to $\max_{\mathcal{A}}\Pi_l(\mathcal{A})$ for each bidder l (line 1). For any increment ϵ, the target profit of each bidder is decreasing along the path of the algorithm and is also decreasing from less than ϵ at each step. Therefore, when the algorithm stops, target profits are distant of at most ϵ from the weak bidder-optimal frontier. We say that the final outcome $\epsilon-$*approximates* the weak bidder-optimal frontier according to the following definition.

Definition 4 *An outcome $(a,(\pi_l)_{l\in N})$ is said to $\epsilon-$approximate a set K if there is an outcome $(a',(\pi'_l)_{l\in N})$ in K such that $a = a'$ and $|\pi_l - \pi'_l| \leq \epsilon$ for any $l \in N$.*

Proposition 3.2 *In the ϵ-Ausubel-Milgrom proxy auction, given the reported preferences, the final outcome $\epsilon-$approximates the weak bidder optimal frontier of the Core. Consequently, the Ausubel-Milgrom proxy auction ends in the weak bidder-optimal frontier.*

Proof 3 *Suppose that the final outcome $(a, (\pi_l)_{l \in N})$ does not $\epsilon-$approximates the weak bidder optimal frontier of the Core. Suppose additionally that the outcome $(a, (\pi_l + \epsilon)_{l \in N})$ is not in the Core. From standard convex analysis (see in Rockafellar [16]), the Core is a polyhedral convex set and there exists a hyperplane separating the Core and the singleton $(a, (\pi_l + \epsilon)_{l \in N})$. Thus there is a point in the interval $[(\pi_l)_{l \in N^*}, (\pi_l + \epsilon)_{l \in N^*}]$ which belongs to the bidder optimal frontier raising a contradiction with $(a, (\pi_l)_{l \in N})$ $\epsilon-$approximating the weak bidder optimal frontier. Finally we have proved that the outcome $(a, (\pi_l + \epsilon)_{l \in N})$ is in the Core. Then as pointed by remark [2.1], the whole cube $\{(a, (x_l)_{l \in N}) | \pi_l \leq x_l \leq \pi_l + \epsilon\}$ is included in the Core. It means that, in the previous round of the ϵ-Ausubel-Milgrom proxy auction, the state of the algorithm was necessary in this cube, which raises a contradiction with proposition [3.1] which states that the algorithm stops when it reaches a Core outcome.*

4 An Illustrative Example

The following example with 4 bidders and 2 identical items for sale without externalities has several aims.[6] First, it illustrates the dynamics of the algorithm. Second, it gives an example where the final outcome is not in the bidder-optimal frontier but only in the weak bidder-optimal frontier of the Core. Third, it illustrates our proposal to add a final discount stage to the Ausubel-Milgrom proxy auction. Thanks to the modification, truthful reporting is an equilibrium which leads to the single-valued bidder-optimal frontier, hence the Vickrey outcome. On the contrary, thruthful reporting is not an equilibrium in the 'original' proxy auction in this example.

Bidder 1 is valuing 100 the first item and 0 an additional item, bidder 2 and 3 are identical and are valuing any additional item 60. For the moment, the bidders have substitutes preferences. Let us introduce an additional

[6] Four is the minimal number of bidders such that Vickrey is in the Core (relative to the whole coalition) and w is not buyer-submodular. The buyer submodularity condition is equivalent to Vickrey being in the Core for each possible coalition (A&M theorem 7). Then for less than two bidders, the equivalence between buyer-submodularity and Vickrey is the Core is tautologic. For three bidders, Vickrey in the Core implies $w(N^*) - w(N^* \setminus \{i\}) \leq w(N^* \setminus \{j\}) - w(N^* \setminus \{i, j\})$ (otherwise $N^* \setminus \{i, j\}$ would be a blocking coalition). Given that $w(\{0\}) = 0$, the remaining inequalities to obtain the buyer-submodularity are of the kind: $w(\{i, j\}^*) \leq w(\{i\}^*) + w(\{j\}^*)$. Those inequalities are always satisfied without externalities. Thus, for three bidders, Vickrey in the Core implies w buyer-submodular. Obviously, two is also the minimal number of items for w not being buyer-submodular without externalities.

bidder 4 who has complement preferences: he values the bundle of the two items 100, but values 0 a single item. Bidder 4 suffers from complementarity. Note that bidder 4 is neutral from a Vickrey implementation point of view: he does not change the efficient allocation which is to assign the items either to the couple $\{1,2\}$ or to $\{1,3\}$ and to make both purchasers pay the amount of 60. Bidder 4 is also neutral from a Core allocation point of view: he does not modify the structure of the Core.

$$Core(\{1,2,3,4\}) = \left\{ (\widehat{\pi}_l)_{0 \leq l \leq N} \mid \sum_{l=0}^{N} \widehat{\pi}_l = 160; \quad \pi_1 \in [0,40] \ ; \ \pi_2 = \pi_3 = \pi_4 = 0 \right\}$$

If bidder 4 were absent, then we could apply A&M's results since the buyer-submodularity would be satisfied. Consequently, the Vickrey outcome would be in the Core and truthful reporting would be a Nash equilibrium in the proxy auction. Nevertheless, the mere presence of bidder 4 disturbs the dynamics of the auction. The final outcome is no longer the Vickrey outcome.

Let us detail a little the dynamic of the auction which could be decomposed into three distinct stages which correspond to modifications in the target allocation set for some agents. Just to fix ideas, consider that the bid increment ϵ equals to 1 (but any smaller increment does not modify the insights). Moreover, in case of ties (line 8 of the algorithm), say that the allocations which favour the bidders with the smaller indexes are chosen. To simplify the presentation, we just focus on five possible assignments: the optimal assignments $\{1,2\}$ and $\{1,3\}$, and the assignments $\{2\}$, $\{3\}$, $\{4\}$ which give both items to the related single bidder. Indeed those are the only relevant assignments in the auction algorithm. Other assignments, such as $\{1\}$ or $\{2,3\}$, are omitted in the following analysis of the target allocation and received bids.

Stage 1: from round 1 to round 80 Initially, the target allocations of bidder 1 are $\{1,2\}$ and $\{1,3\}$. The target allocation of bidder 2 (respectively 3 and 4) is $\{2\}$ (respectively $\{3\}$ and $\{4\}$). Each bidder adds the bid increment 1 on his target allocations except when the chosen allocation is one of those target allocations, i.e. three times out of four he actively bids whereas one time out of four his target allocation is chosen. After round 80, the submitted bids are those shown in table 1.

Stage 2: from round 81 to round 130 Now there is a change in the target allocation set of bidders 2 (respectively 3): he is willing to bid for assignment $\{1,2\}$ (respectively $\{1,3\}$). Note that at the beginning of that stage, bidder 1 has reached his Vickrey payoff and that consequently, if he were now strategic, he should better not raise any incremental bid. Nevertheless his proxy bidder will do so. The reason is that he outbids when $\{4\}$ is the selected assignment. The precise dynamic of the selected allocation

	{1,2}	{1,3}	{2}	{3}	{4}
1	60	60	0	0	0
2	0	0	60	0	0
3	0	0	0	60	0
4	0	0	0	0	60

Table 1: Bids after round 80

follows a five period-cycle. Table 2 follows the five first rounds of this second stage. For each round and for each (relevant) assignment, the first row reports the set of agents that is raising an incremental bid on this assignment. For example, at round 83, bidder 1 is raising his submitted bids for assignment $\{1,2\}$ and $\{1,3\}$, whereas bidder 2 (respectively 3) is raising his bid on allocation $\{1,2\}$ (respectively $\{1,3\}$). The second row reports for each assignment the sum of submitted bids and the last column the revenue maximizing assignment that is selected. For example, at round 82, the allocation $\{4\}$ is the single revenue maximizing allocation and is therefore selected.

Round	{1,2}	{1,3}	{4}	Selected
81		3	4	
	60	61	61	{1,3}
82	2		4	
	61	61	62	{4}
83	1,2	1,3		
	63	63	62	{1,2}
84		3	4	
	63	64	63	{1,3}
84	2		4	
	64	64	64	{1,2}

Table 2: Incremental Bids for rounds 81-85

Then this stage ends at round 130 where the corresponding bids are reported in table 3. The striking point is that bidder 1 has continued to participate actively in the mechanism and has overbidden 10 above the Vickrey outcome bid. He has overbidden as if bidder 4 were a ‘serious’ opponent against which he should fight.

Stage 3: from round 131 to round 170 Now bidder 4’s target profit is null and he quits the auction. One of bidder 1’s target allocation is always selected such that from now on, bidder 1 does not raise bids anymore. The auction terminate in a ‘duel’ between agent 2 and 3. The final assignment is either $\{1,2\}$ or $\{1,3\}$. The final revenue is 130. All bids are reported in

	{1,2}	{1,3}	{2}	{3}	{4}
1	70	70	0	0	0
2	30	0	90	0	0
3	0	30	0	90	0
4	0	0	0	0	100

Table 3: Bids after round 130

table 4. Note that if bidder 1 were able to reduce slightly his submitted bids, he would do so without modifying the payoffs of the other bidders but only reducing the revenue of the seller.

	{1,2}	{1,3}	{2}	{3}	{4}
1	70	70	0	0	0
2	60	0	120	0	0
3	0	60	0	120	0
4	0	0	0	0	100

Table 4: Bids after round 170

This suggest to add a stage to the generalized proxy auction where the auctioneer reduce incrementally the bid of some bidders such that the remaining allocation still stay in the Core. This stage will be referred to as the final discount stage. In our case, it is clear that such a discount for the winner among 2 and 3 is impossible since the payoffs would be driven out of the Core. On the other hand, the final discount stage will imply a reduction in bidder 1's price: he will pay only 60. Indeed, 60 corresponds to the amount that he should pay to internalize the externality imposed on his opponents. Somehow clumsily, he has bid above 60 because bidder 1's proxy, at the beginning, bid as if he should internalize the externality imposed only on his opponents $\{4\}$, an externality which is stronger than the one he imposes on the bigger set of opponents $\{2, 3, 4\}$. This is exactly those events that the buyer-submodularity condition avoids.

	{1,2}	{1,3}	{2}	{3}	{4}
1	60^+	60^+	0	0	0
2	60	0	120	0	0
3	0	60	0	120	0
4	0	0	0	0	100

Table 5: Bids after the discount to bidder 1

Indeed, it can be proved that the payoff of bidder l in the Ausubel-

Milgrom proxy auction is at least $\min_{S\subset N} w(S^*) - w(S^* \setminus l)$.[7] If $\min_{S\subset N} w(S^*) - w(S^* \setminus l) \geq \pi_l^V$, then truthful reporting leads to the Vickrey payoffs and is thus a best response for bidder l. The buyer-submodularity condition is a stronger one. In our example, we have chosen precisely that $\min_{S\subset N} w(S^*) - w(S^* \setminus l) = w(\{1,4\}) - w(\{4\}) = 0 < \pi_l^V = 40$.

Next section defines properly the final discount stage. Unambiguously, such an additional stage gives better incentives for truthful reporting.

5 The Ausubel-Milgrom generalized proxy auction with final discounts

The final discount stage that we propose for the ϵ-Ausubel-Milgrom proxy auction intervenes when the previous mechanism has stopped: this stage does not modify the final allocation but only the prices that will be paid at the end. Instead of the final accepted bids for the final allocation, the bidders will pay final discounted bids. Iteratively, for all winning bidders, bids on all allocations are uniformly discounted provided that the final allocation remains the revenue maximizing allocation. It should be emphasized that the final discount stage is defined only from the final allocation and the entire set of submitted bids in the last round. Thus it is defined without any reference to the reported preferences to the proxy bidder. Hence, such a discount stage could be implemented in the dynamic version of the ascending package auction. Actually, the Ausubel-Milgrom generalized proxy auction with final discounts falls within the class of ascending price auctions introduced by Mishra and Parkes [14].

Definition 5 (The Final Discount Stage) *The final discount stage is the function mapping $(\mathcal{A}^*, (b_l(\mathcal{A}))_{1\leq l\leq N, \mathcal{A}\in\mathbf{A}}$, i.e. the final allocation $\mathcal{A}^*$ and the entire set of submitted bids, into the outcome $(\mathcal{A}^*, (b_l(\mathcal{A}^*))_{1\leq l\leq N}$ where the final accepted bids are discounted according to the following algorithm:*

for $l = 1$ *to* N **do**
 if $\mathcal{A}_l^* \neq \emptyset$
  then do $b_l(\mathcal{A}) := \max\{0, b_l(\mathcal{A}) - d\}$
  where d is the largest number such that
  $\mathcal{A}^* \in Arg\max_{\mathcal{A}\in\mathbf{A}} \left(\max\{0, b_l(\mathcal{A}) - d\} + \sum_{k=1,k\neq l}^{N} b_k(\mathcal{A})\right)$

Remark 5.1 *The final discount stage has not been presented in a symmetric way: a bidder may find strictly advantageous to be selected at the first iterations of the discount stage. A proper randomization makes the discount stage symmetric.*

[7]It is left to the reader to check that it is proved implicitly in the proof of A&M's Theorem 8.

The final discount stage could also be interpreted relative to the target profits: the discount d on submitted bids of bidder l is equivalent to the increase in d of the target profit of bidder l. After a discount, the efficient allocation must remain the profit maximizing allocation: in the perspective of lemma [3.1], it means that $w(N^*) - \sum_{l \in N} \pi_l = \max_{S \subset N} w(S^*) - \sum_{l \in S} \pi_l$.

The preceding definition is illustrated in Figure [1a] and [1b]. Those figures show the dynamic of the Ausubel-Milgrom proxy auction in a two dimensional target profit space of two selected bidders [8]. Dashed lines depict those dynamics starting from the initial target profit π^{Max} where each bidder obtains his most preferred allocation with a null bid and stoping at a point in the weak bidder-optimal frontier. In both cases, the final outcome of the Ausubel-Milgrom proxy auction is not in the bidder-optimal frontier and the discount stage is an active stage illustrating how the proposal concretely modify the outcome of the action. Fig [1a] is such that the Vickrey payoff is in the Core. In this case, the 'original' proxy auction does not implement truthfully the efficient assignement whereas the proxy auction with final discounts does.

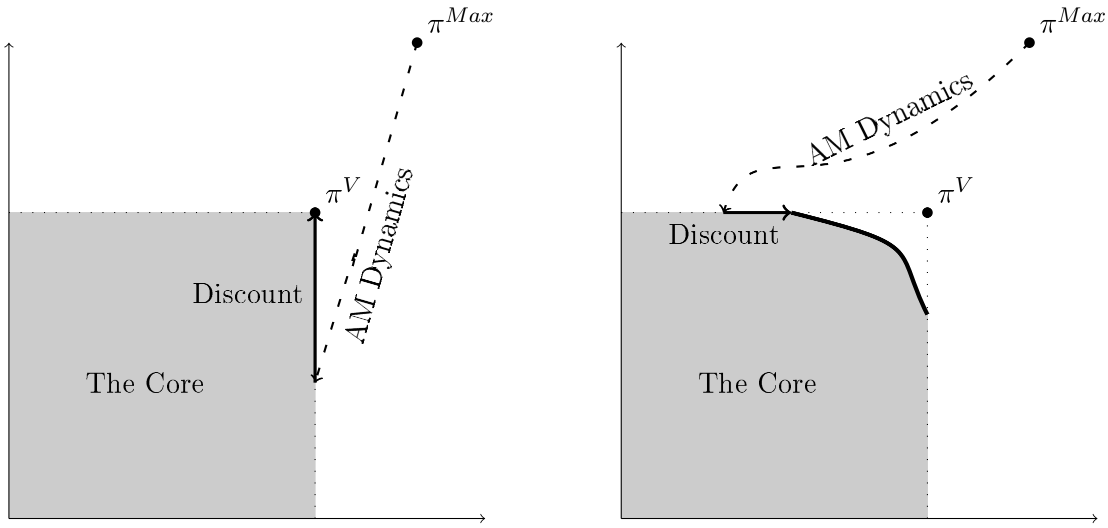

Fig 1a: Vickrey in the Core

Fig 1b: Vickrey outside the Core

We then prove that, from any point in the Core, the final discount stage leads to a point in the bidder-optimal frontier. This is proved in two steps. First, a bid discount iteration is such that the payoffs remain in the Core. Second, we show that the final payoffs $(\pi_l)_{l \in N}$ do not lie below the bidder-optimal frontier. Otherwise, there is a bidder l such that his bids could be discounted until a target profit $\pi_l^* > \pi_l$ and such that the final allocation is

[8]With only two bidders, w is always buyer-submodular and the discounts have no effect on the outcomes since the Vickrey outcome is directly implemented. From 3 bidders, the discounts can possibly modify the final outcome: consider section 4's example and suppress bidder 3 as an example.

still the revenue maximizing one. Then, when it was bidder l's turn in the final discount stage algorithm, bidder l's target profit should have been raised until π_l^*, which raises a contradiction since it has been raised only until π_l. Since the final allocation of the original ϵ-Ausubel-Milgrom proxy auction is in the Core, the final discount stage applied to the ϵ-Ausubel-Milgrom proxy auction leads to an outcome in the bidder-optimal frontier as stated in the following proposition.

Proposition 5.1 *In the ϵ-Ausubel-Milgrom proxy auction with final discounts, given the reported preferences, the final outcome is in the bidder optimal frontier of the Core.*

Proof 4 *It remains to prove the two aforementioned steps. A bid discount is just a transfer between the seller and one bidder. Both outcomes share the same final allocation $\mathcal{A}^*$. Consequently, condition (a) in the definition of the Core remains unchanged. Moreover the discounts are such that $\mathcal{A}^*$ remains the profit maximizing allocation relative to the target profit, then due to lemma [3.1], inequalities (b) are still satisfied. In a nutshell, the outcome after a discount remains in the Core provided that the initial allocation is in the Core.*
Then it remains to prove that if the final outcome is strictly below the bidder-optimal frontier, or equivalently if the bids of a bidder, say l, could be discounted until a target profit π_l^*, then it should have been discounted previously until that level in the algorithm, which raises a contradiction. Denote by $(\pi_l)_{l\in N}$ (respectively $(\pi_l')_{l\in N}$) the target profits at the end of the final discount stage (respectively just before bidder l's turn in the final discount algorithm). For both target profit vectors, $\mathcal{A}^*$ is the profit maximizing allocation. Moreover, $\pi_l \geq \pi_l'$ since the continuation of the algorithm involves only discounts in bids, or equivalently increases in the target profits. We have assumed above that the target profit vector $(\pi_1, \cdots, \pi_{l-1}, \pi_l^*, \pi_{l+1}, \cdots, \pi_N)$ is such that $\mathcal{A}^*$ is still profit maximizing:[9]

$$w(N^*) - \sum_{k\in N, k\neq l} \pi_k - \pi_l^* = \max_{S\subset N} \left\{ w(S^*) - \sum_{k\in S, k\neq l} \pi_k - \pi_l^* \cdot I[l \in S] \right\}.$$

Since $\pi_l \geq \pi_l'$, we have:

$$\max_{S\subset N} \left\{ w(S^*) - \sum_{k\in S, k\neq l} \pi_k - \pi_l^* \cdot I[l \in S] \right\} + \sum_{k\in N, k\neq l} \pi_k \geq$$
$$\max_{S\subset N} \left\{ w(S^*) - \sum_{k\in S, k\neq l} \pi_k' - \pi_l^* \cdot I[l \in S] \right\} + \sum_{k\in N, k\neq l} \pi_k'.$$

We conclude that:

$$w(N^*) - \sum_{k\in N, k\neq l} \pi_k' - \pi_l^* \cdot I[l \in S] \geq \max_{S\subset N} \left\{ w(S^*) - \sum_{k\in S, k\neq l} \pi_k' - \pi_l^* \cdot I[l \in S] \right\}.$$

Consequently, for the target profit vector $(\pi_1', \cdots, \pi_{l-1}', \pi_l^*, \pi_{l+1}', \cdots, \pi_N')$, $\mathcal{A}^*$ is the profit maximizing allocation. Thus we have raised a contradiction with the course of the final discount stage.

If the Vickrey outcome is in the Core, then the ϵ-Ausubel-Milgrom proxy auction with final discounts implements the Vickrey outcome under truth-

[9] Denote by $I[l \in S]$ the indicator function equal to 1 if $l \in S$ and else equal to 0.

ful reporting since the bidder-optimal frontier is then a singleton coinciding with the Vickrey payoffs (see A&M theorem 6). Hence truthful reporting is a Nash equilibrium strategy profile if the Vickrey outcome is in the Core. The converse statement results from the fact that if his opponents are truthful, a bidder can guarantee himself his Vickrey outcome by one of his best response's report: it corresponds to report the Vickrey profit as a target profit. Thus we have proved the following corollary which is indeed true for any mechanism that implements an outcome in the bidder-optimal frontier relative to the reported preferences.

Corollary 5.2 *Truthful reporting is a Nash equilibrium strategy profile if and only if the Vickrey outcome is in the Core. Then the $\epsilon-$Ausubel-Milgrom proxy auction with final discounts leads to the Vickrey outcome.*

In A&M, the truthful Nash Equilibrium is obtained for the mechanism by taking the limit $\epsilon \to 0$. Otherwise, the mechanism ends generically in the interior of the Core thus not at the Vickrey outcome. On the other hand, note that our corresponding result for the proxy auction with final discounts is true for any increment ϵ.

More generally, even if the Vickrey outcome is not in Core, then the final stage brings the final outcome unambiguously closer to the Vickrey outcome. According to Milgrom's [13] terminology, the Ausubel-Milgrom proxy auction with final discounts is a core-selecting auction that provides optimal incentives.

References

[1] L. Ausubel. An efficient ascending-bid auction for multiple objects. *Amer. Econ. Rev.*, 94(5):1452–1475, 2004.

[2] L. Ausubel. An efficient dynamic auction for heterogenous commodities. *Amer. Econ. Rev.*, 96(3):602–629, 2006.

[3] L. Ausubel and P. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1), 2002.

[4] D. B. Bernheim and M. Whinston. Menu auctions, resource allocation and economic influence. *Quaterly Journal of Economics*, 101(1):1–31, 1986.

[5] S. Bikhchandani and J. Ostroy. The package assignment model. *Journal of Economic Theory*, 107(2):377–406, 2002.

[6] S. Bikhchandani and J. Ostroy. Ascending price vickrey auctions. *Games and Economic Behavior*, 55:215–241, 2006.

[7] S. de Vries, J. Schummer, and R. Vohra. On ascending vickrey auctions for heterogenous objects. *Journal of Economic Theory*, 132:95–118, 2007.

[8] G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, August 1986.

[9] F. Echenique and J. Oviedo. Core many-to-one matchings by fixed-point methods. *Journal of Economic Theory*, 115:358–376, 2004.

[10] F. Gul and E. Stacchetti. The english auction with differentiated commodities. *Journal of Economic Theory*, 92:66–95, 2000.

[11] J. W. Hatfield and P. Milgrom. Matching with contracts. *American Economic Review*, 95(4):913–935, September 2005.

[12] P. Jehiel and B. Moldovanu. Strategic nonparticipation. *RAND J. Econ.*, 27(1):84–98, 1996.

[13] P. Milgrom. Incentives in core-selecting auctions. *mimeo*, October 2006.

[14] D. Mishra and D. Parkes. Ascending price vickrey auctions for general valuations. *Journal of Economic Theory*, 132:335–366, 2007.

[15] M. Ranger. The generalized ascending proxy auction in the presence of externalities. 2005.

[16] T. Rockafellar. *Convex Analysis.* Princeton University Press, 1970.