

Bayesian Combinatorial Auctions

George Christodoulou¹, Annamária Kovács², and Michael Schapira³

¹ Max-Planck-Institut für Informatik, Saarbrücken, Germany.

`gchristo@mpi-inf.mpg.de`

² Institute for Computer Science, J. W. Goethe University, 60325 Frankfurt/Main, Germany. `panni@cs.uni-frankfurt.de`

³ The School of Computer Science and Engineering, The Hebrew University of Jerusalem, Israel. `mikesch@cs.huji.ac.il`. **

Abstract. We study the following Bayesian setting: m items are sold to n selfish bidders in m independent second-price auctions. Each bidder has a *private* valuation function that expresses complex preferences over *all* subsets of items. Bidders only have *beliefs* about the valuation functions of the other bidders, in the form of probability distributions. The objective is to allocate the items to the bidders in a way that provides a good approximation to the optimal social welfare value. We show that if bidders have submodular valuation functions, then every Bayesian Nash equilibrium of the resulting game provides a 2-approximation to the optimal social welfare. Moreover, we show that in the full-information game a pure Nash always exists and can be found in time that is polynomial in both m and n .

1 Introduction

Combinatorial Auctions. In a *combinatorial auction* m items $M = \{1, \dots, m\}$ are offered for sale to n bidders $N = \{1, \dots, n\}$. Each bidder i has a *valuation function* (or valuation, in short) v_i that assigns a non-negative real number to *every* subset of the items. v_i expresses i 's preferences over *bundles* of items. The value $v_i(S)$ can be thought of as specifying i 's maximum willingness to pay for S . Two standard assumptions are made on each v_i : $v_i(\emptyset) = 0$ (normalization), and $v_i(S) \leq v_i(T)$ for every two bundles $S \subseteq T$ (monotonicity). The objective is to find a partition of the items among the bidders S_1, \dots, S_n (where $S_i \cap S_j = \emptyset$ for all $i \neq j$) such that the social welfare $\sum_i v_i(S_i)$ is maximized.

The interplay between selfishness and computational optimization in combinatorial auctions is well-studied. Each of these aspects alone can be handled in a satisfactory way: The celebrated VCG mechanisms [19, 3, 10] motivate agents to truthfully report their private information, and optimize the social-welfare. The caveat is that this may take exponential time [15, 16] (in the natural parameters of the problem m and n). On the other hand, if we disregard strategic issues, it is possible to obtain good approximations to the optimal social-welfare

** Supported by a grant from the Israel Science Foundation.

in polynomial time, for restricted, yet very expressive special cases of combinatorial auctions (e.g., combinatorial auctions in which bidders have *submodular* valuations, known as combinatorial auctions with *submodular bidders* [13, 4–7, 20]). It is the combination of the two challenges that proves to be problematic.

Bayesian Combinatorial Auctions. In this paper we approach the problem of handling selfishness and computational hardness from an old-new perspective: Harsanyi [11] introduced *Bayesian games* as an elegant way of modeling selfishness in partial-information settings. In a Bayesian game, players do not know the private information of the other players, but only have *beliefs*, expressed by probability distributions over the different possible realizations of this private information. In combinatorial auctions this translates to probability distributions over the possible valuation functions of the other bidders. We are interested in maximizing the social-welfare in a way that is aligned with the interests of the different bidders. We ask the following question: Can we design an auction for which *any* Bayesian Nash equilibrium provides a good approximation to the optimal social-welfare? This question is, of course, an extension of the well known “*price of anarchy*” question [12, 17] to Bayesian settings.

Inspired by eBay, we study the simple auction in which the m items are sold in m independent second-price auctions. This auction induces a game in which a bidder’s strategy is the m -dimensional vector of bids he submits in the different single-item auctions, and his payoff is his value for the set of items he is allocated minus his payments. Unfortunately, in this general setting, some unnatural problems may arise: Consider the following simple example: $m = 1$, $n = 2$, and the bidders have complete information about each other. Let $v_1(1) = 1$ and $v_2(1) = 0$. Observe that the optimal social welfare is 1 (assign item 1 to bidder 1). Also observe that if bidder 1 bids 0 and bidder 2 bids 1 then this is a pure Nash equilibrium with a social-welfare value of 0. Therefore, the price of anarchy of this full-information game is infinity.

In the above scenario, the second bidder bid for (and eventually got) an object he was not interested in possessing. However, such a situation is unlikely to occur in real-life situations, especially if bidders are only partially informed and are therefore more inclined to avoid risks. To handle this, we adopt the well-known assumption in decision theory, that the players (bidders) are ex-post individually-rational. Informally, ex-post individual-rationality means that bidders play it safe, in the sense that a bidder will never make decisions that might (in some scenario) result in getting a negative payoff.

Our Results. Our main result is the following: We exhibit an auction (we refer to as a *Bayesian auction*) in which essentially all items are sold in a second-price auction. We prove that *any* mixed Nash equilibrium for this auction provides a good approximation to the optimal social welfare.

Theorem: If bidders are ex-post individually-rational, and have submodular valuation functions, then every (mixed) Bayesian Nash equilibrium of a Bayesian auction provides a 2-approximation to the optimal social welfare.

A bidder i is said to have a *submodular* valuation function if for all $S, T \subseteq M$ $v_i(S \cup T) + v_i(S \cap T) \leq v_i(S) + v_i(T)$. This definition of submodularity is known (see, e.g., [13]) to be equivalent to the following definition: A valuation function is submodular if for every two bundles $S \subseteq T$ that do not contain an item j it holds that

$$v_i(S \cup \{j\}) - v_i(S) \geq v_i(T \cup \{j\}) - v_i(T).$$

This last inequality has a natural interpretation as it implies that the bidders have decreasing marginal utilities (the marginal value of the item decreases as the number of items a bidder has increases). In fact, our theorem holds even if bidders have fractionally-subadditive valuation functions [6] (defined, and termed XOS, in [14]), a class of valuation functions that strictly contains all submodular ones [14, 13].

We stress that this “*Bayesian price of anarchy*” result is independent of the bidders’ beliefs. That is, the 2-approximation ratio is guaranteed for *any* common probability distribution (“common prior”) over the valuation functions (we do require the common prior to be the product of independent probability distributions). This suggests a middle-ground between the classical economic and the standard computer science approaches: Works in economics normally assume that the “input” is drawn from some *specific* probability distribution, and prove results that apply to that specific one. In contrast, computer scientists prefer a worst-case analysis that holds for every possible input. We require the assumption that the input is drawn from some (known) probability distribution but expect to obtain a good approximation ratio regardless of what it is. We note that an approach similar to ours was applied to selfish routing problems in [9, 8].

Open Question: Can a (mixed) Bayesian Nash equilibrium be computed in polynomial time?

Simple examples show that a *pure* Bayesian Nash is not guaranteed to exist in our Bayesian setting⁴. An interesting special case of the Bayesian game is the full-information game in which every bidder’s valuation function is known to all other bidders. We show, that in such full-information games a *pure* Nash equilibrium *always* exists. This is true even if the bidders are *not* ex-post individually rational. In fact, it is easy to show that even the optimal allocation of items can be achieved in a pure Nash equilibrium. So, while the price of anarchy in this game is infinity (without ex-post individual-rationality), the *price of stability* (the social welfare value of the *best* Nash equilibrium [1]) is 1. However, optimizing the social-welfare in combinatorial auctions with submodular bidders is known to be NP-hard [13]. Therefore, we are left with the following natural question: Can we find a pure Nash equilibrium that provides a good approximation to the optimal social welfare in polynomial time?

⁴ In fact, even a mixed Nash equilibrium is not guaranteed to exist unless one discretizes the strategy space (e.g., only allows bids that are multiples of some small $\epsilon > 0$).

We give the following answer for submodular bidders:

Theorem: If bidders have submodular valuation functions then a pure Nash equilibrium of the complete-information game that provides a 2-approximation to the optimal social welfare exists and can be computed in polynomial time.

We prove the theorem by showing that the approximation algorithm for maximizing social welfare in combinatorial auctions with submodular bidders, proposed by Lehmann et al. [13], can be used to compute the bids in a pure Nash equilibrium. We note that similar questions have been studied by Vetta in [18]⁵.

For the wider class of fractionally-subadditive valuation functions, we provide a constructive way of finding a pure Nash that provides a 2-approximation via a simple and natural myopic procedure. This procedure is inspired by the greedy approximation-algorithm in [4]. Unfortunately, while this myopic procedure does compute a pure Nash equilibrium in polynomial time for some interesting (non-submodular) subcases, this is not true in general. We prove that the myopic procedure may take exponential time by exhibiting a non-trivial construction of an instance on which this can occur.

Open Question: Can a pure Nash equilibrium that obtains a 2-approximation to the optimal social-welfare be computed in polynomial time if bidders have fractionally-subadditive valuation functions?

The proofs omitted due to space limitations can be found in the full version of the paper.

2 Bayesian Price of Anarchy

In Subsection 2.1 we present the Bayesian setting we explore in this paper, which we term “Bayesian combinatorial auctions”. In Subsection 2.2 we exhibit our main result, which is that the Bayesian price of anarchy of Bayesian combinatorial auctions is 2.

2.1 The Setting - Bayesian Combinatorial Auctions

The Auction. m items are sold to n bidders in m independent second-price auctions (with some tie-breaking rule). A bidder’s strategy is a *bid-vector* $b_i \in R_{\geq 0}^m$ ($b_i(j)$ represents i ’s bid for item j).⁶ A (*pure*) *strategy profile* of all players

⁵ Vetta considers a general setting in which decisions are made by non-cooperative agents, and the utility functions are submodular. He proves that in this setting the price of anarchy is at most 2. The framework discussed there is such that the players’ pure strategies are *subsets* of a ground set (e.g., the items). This framework is not applicable to our auction, where the bids play a crucial role.

⁶ So as to have a finite (discrete) model, we assume that all acceptable bids are multiples of an arbitrary ϵ . Furthermore, they cannot exceed some maximum value B_{max} .

is an n -tuple $b = (b_1, \dots, b_n)$. We will use the notation $b = (b_i, b_{-i})$, to denote a strategy profile in which bidder i bids b_i and other bidders bid as in $b_{-i} = (b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$.

Given a strategy profile b , the items are allocated according to the second price rule, i.e., every object is sold to the highest bidder at a price equal to the second highest bid. For technical reasons, to be explained later, in each such second-price auction with negligible probability the item is randomly allocated to one of the bidders (who submitted a bid higher than zero). In this case, that bidder is charged his bid for the item.

For some fixed b we denote by $X_i(b)$ the set of items obtained by player i in the auction⁷. For a set $S \subseteq M$, let the sum of the (highest) bids be denoted by

$$Bids(S, b) = \sum_{j \in S} \max_k b_k(j),$$

$$Bids_{-i}(S, b_{-i}) = \sum_{j \in S} \max_{k \neq i} b_k(j),$$

and

$$Bids_i(S, b_i) = \sum_{j \in S} b_i(j).$$

The *utility* (payoff) of player i is then given by

$$u_i(b) = v_i(X_i(b)) - Bids_{-i}(X_i(b), b_{-i}).$$

We make two assumptions about the bidders: ex-post individual-rationality, and that the v_i s are *fractionally-subadditive*. A valuation is fractionally-subadditive if it is the pointwise maximum of a set of additive valuations: A valuation a_i is additive if for every $S \subseteq M$ $a_i(S) = \sum_{j \in S} a_i(\{j\})$. A valuation v_i is fractionally-subadditive if there are additive valuations $A = \{a_1, \dots, a_l\}$ such that for every $S \subseteq M$ $v_i(S) = \max_{a \in A} a(S)$. (We will call $a_k \in A$ a *maximizing additive valuation* for the set S if $v_i(S) = a_k(S)$.)

The class of fractionally-subadditive valuations is known to be strictly contained in the class of subadditive valuations and to strictly contain the class of submodular valuations [14, 13].

Bayesian Nash Equilibria. For all i , let V_i denote the finite set of possible valuations of player i . The set of possible valuation profiles of the players is then $V = V_1 \times \dots \times V_n$. There is a *known* probability distribution D over the valuations V (a *common prior*). D can be regarded as some market statistics that is known to all bidders (and to the auctioneer), and specifies their beliefs. We assume that $D = D_1 \times \dots \times D_n$ is the cartesian product of independent probability distributions D_i : any valuation profile $v = (v_1, \dots, v_n)$ occurs with probability $D(v) = \prod_{i=1}^n D_i(v_i)$, where $D_i(v_i)$ is the probability that bidder i

⁷ observe that $X_i(b)$ is actually a random variable, but since the event that not all items are sold in second-price auctions only occurs with very low probability we shall often refer to $X_i(b)$ as indicating a specific bundle of items.

has the valuation function v_i . Let $V_{-i} = \times_{k \neq i} V_k$, $D_{-i} = \times_{k \neq i} D_k$, and $v_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$.

A *bidding-function* B_i for player i is a function that assigns a bid-vector $b_i = B_i(v_i)$ to every valuation function $v_i \in V_i$. The reader may find it helpful to think of B_i as a suggestion made to player i by the auctioneer as to which bid to submit. An n -tuple of bidding-functions $B = (B_1, \dots, B_n)$ is a *Bayesian Nash equilibrium* if for every $i \in [n]$, and for every valuation function v_i , the bid $B_i(v_i)$ maximizes i 's expected utility given that his valuation function is v_i , and that the bid of every other bidder j is $B_j(v_j)$, where v_j is drawn from D_j . That is, a Bayesian Nash maximizes i 's expected payoff for any valuation function he may have, given his beliefs about the other bidders.

Bayesian Price of Anarchy. For a fixed valuation profile of the bidders $v = (v_1, \dots, v_n)$, the optimal social-welfare is $OPT(v) = \max_{S_1, \dots, S_n} \sum_i v_i(S_i)$, where the maximum is taken over all partitions of M into disjoint bundles S_1, \dots, S_n . For given D , the (expected) optimal social-welfare $SW(OPT)$ is the expectation $E[OPT(v)]$, where v is drawn from D . That is,

$$SW(OPT) = \sum_{v \in V} D(v) OPT(v)$$

Given a profile v , every pure strategy profile b induces a social-welfare value $SW(b) = \sum_{i \in [n]} v_i(X_i(b))$. For an n -tuple of bidding-functions $B = (B_1, \dots, B_n)$, we denote by $SW(B)$ the expected social welfare $E[SW(B_1(v_1), \dots, B_n(v_n))]$, where the $v = (v_1, \dots, v_n)$ is drawn from D :

$$SW(B) = \sum_{v \in V} D(v) SW(B(v)).$$

We are interested in Bayesian Nash equilibria B for which the ratio $\frac{SW(OPT)}{SW(B)}$ is small. The *Bayesian price of anarchy* of a game is

$$PoA = \max_{D, B \text{ Bayesian Nash}} \frac{SW(OPT)}{SW(B)},$$

that is the maximum of the expression $\frac{SW(OPT)}{SW(B)}$, taken over *all* probability distributions D , and *all* Bayesian Nash equilibria B (for these probability distributions). Intuitively, a Bayesian price of anarchy of α means that no matter what the bidders' beliefs are, every Bayesian Nash equilibrium provides an α -approximation to the optimal social-welfare.

Supporting Bids. Due to ex-post individual-rationality, a bidder will never submit a bid that might result in a negative payoff. Recall that the rules of the auction dictate that it is possible that a bidder get *any* subset of the items for which he submitted his non-zero bids and be charged the sum of these bids. Hence, a bidder's bid must uphold the following property:

Definition 1. A bid vector b_i is said to be a *supporting bid* given a valuation v_i , if for all $S \subseteq M$ $v_i(S) \geq \sum_{j \in S} b_i(j)$.

Recall that the event that not all items are sold in second-price auctions only occurs with negligible probability. Hence, for simplicity, we shall disregard this event and prove our theorem for the case that all items are sold in second-price auctions.

2.2 Bayesian Price of Anarchy of 2

This subsection exhibits our main result. For ease of exposition we prove the theorem regarding the Bayesian price of anarchy for pure Bayesian Nash equilibria. The (more complicated) proof for mixed Nash is omitted due to space limitations and appears in the full version of the paper. The proof of the theorem exploits the fractional subadditivity of the valuations via the following lemma:

Lemma 1. *Let S be a set of items, and a_i be a maximizing additive valuation of player i for this set, restricted so that $a_i = 0$ on $M \setminus S$. If i bids according to a_i , while all the others bid according to any pure profile b_{-i} , then*

$$u_i(a_i, b_{-i}) \geq v_i(S) - \text{Bids}_{-i}(S, b_{-i}).$$

Proof. Let $X_i := X_i(a_i, b_{-i})$ be the set of items that player i is going to get. Note that if i wins any item $j \notin S$ then the maximum bid on this j was 0. Thus we can assume w.l.o.g. that $X_i \subseteq S$. Moreover, $a_i(j) - \text{Bids}_{-i}(\{j\}, b_{-i}) \leq 0$ holds for every non-obtained item $j \in S - X_i$. Therefore, we have

$$\begin{aligned} u_i(a, b_{-i}) &= v_i(X_i) - \text{Bids}_{-i}(X_i, b_{-i}) \\ &\geq \sum_{j \in X_i} a_i(j) - \text{Bids}_{-i}(X_i, b_{-i}) \\ &\geq \sum_{j \in S} a_i(j) - \text{Bids}_{-i}(S, b_{-i}) \\ &= v_i(S) - \text{Bids}_{-i}(S, b_{-i}). \end{aligned}$$

Theorem 1. *Let D be a distribution over fractionally-subadditive valuations of the bidders. If $B = (B_1, \dots, B_n)$ is a Bayesian Nash, such that each B_i maps every valuation function v_i to a supporting bid (given v_i) then $\frac{SW(OPT)}{SW(B)} \leq 2$.*

Proof. Let $v = (v_1, \dots, v_n)$ be a fixed valuation profile. We denote by $O^v = (O_1^v, \dots, O_n^v)$ the optimum allocation with respect to profile v .

Now for every player i , let a_i denote the maximizing additive valuation for the set O_i^v , (in particular, $a_i(j) = 0$ if $j \notin O_i^v$). For all i , we consider a_i as an alternative strategy to $B_i(v_i)$.

Let us fix a bidder i . Let w_{-i} be an arbitrary valuation profile of all bidders except for i . We introduce the short notation

$$X_i^{w_{-i}} \stackrel{\text{def}}{=} X_i(B_i(v_i), B_{-i}(w_{-i})).$$

Furthermore, for any $S \subseteq M$ we will use

$$Bids_{-i}^{w_{-i}}(S) \stackrel{\text{def}}{=} Bids_{-i}(S, B_{-i}(w_{-i})),$$

resp.

$$Bids^w(S) \stackrel{\text{def}}{=} Bids(S, B(w)),$$

where $w = (w_i, w_{-i})$ is a valuation profile.

Since B is a Bayesian Nash, the strategy $B_i(v_i)$ provides higher expected utility to player i than the strategy a_i :

$$\sum_{w_{-i} \in V_{-i}} D(w_{-i}) u_i(B_i(v_i), B_{-i}(w_{-i})) \geq \sum_{w_{-i} \in V_{-i}} D(w_{-i}) u_i(a_i, B_{-i}(w_{-i})).$$

The utility values on the left-hand-side are

$$u_i(B_i(v_i), B_{-i}(w_{-i})) = v_i(X_i^{w_{-i}}) - Bids_{-i}^{w_{-i}}(X_i^{w_{-i}}) \leq v_i(X_i^{w_{-i}}).$$

On the right-hand-side, applying Lemma 1 yields

$$u_i(a_i, B_{-i}(w_{-i})) \geq v_i(O_i^v) - Bids_{-i}^{w_{-i}}(O_i^v).$$

By merging the inequalities above, we get

$$\begin{aligned} \sum_{w_{-i} \in V_{-i}} D(w_{-i}) v_i(X_i^{w_{-i}}) &\geq \sum_{w_{-i} \in V_{-i}} D(w_{-i}) [v_i(O_i^v) - Bids_{-i}^{w_{-i}}(O_i^v)] \\ &= v_i(O_i^v) \sum_{w_{-i} \in V_{-i}} D(w_{-i}) - \sum_{w_{-i} \in V_{-i}} D(w_{-i}) Bids_{-i}^{w_{-i}}(O_i^v) \\ &= v_i(O_i^v) \cdot 1 - \sum_{w \in V} D(w) Bids_{-i}^{w_{-i}}(O_i^v) \\ &\geq v_i(O_i^v) - \sum_{w \in V} D(w) Bids^w(O_i^v). \end{aligned}$$

Here the expected highest bids $\sum_{w_{-i} \in V_{-i}} D(w_{-i}) Bids_{-i}^{w_{-i}}(O_i^v)$, and $\sum_{w \in V} D(w) Bids_{-i}^{w_{-i}}(O_i^v)$ are equal, because D is independent for all bidders. Finally, $Bids_{-i}^{w_{-i}}(O_i^v) \leq Bids^w(O_i^v)$ obviously holds, since in the latter case we consider maximum bids over a larger set of players. We obtained

$$v_i(O_i^v) \leq \sum_{w_{-i} \in V_{-i}} D(w_{-i}) v_i(X_i^{w_{-i}}) + \sum_{w \in V} D(w) Bids^w(O_i^v).$$

We sum over all i , and then take the expectation over all valuations $v = (v_1, \dots, v_n)$ on both sides:

$$\begin{aligned} \sum_{v \in V} D(v) \sum_{i \in [n]} v_i(O_i^v) &\leq \sum_{v \in V} D(v) \sum_{i \in [n]} \sum_{w_{-i} \in V_{-i}} D(w_{-i}) v_i(X_i^{w_{-i}}) \\ &\quad + \sum_{v \in V} D(v) \sum_{i \in [n]} \sum_{w \in V} D(w) Bids^w(O_i^v). \end{aligned}$$

Note that $\sum_{v \in V} D(v) \sum_{i \in [n]} v_i(O_i^v) = SW(OPT)$. Furthermore, we claim that both summands on the right-hand-side are at most $SW(B)$, so that $SW(OPT) \leq 2SW(B)$, which will conclude the proof. The first summand is

$$\begin{aligned}
& \sum_{i \in [n]} \sum_{v_i \in V_i} D(v_i) \sum_{v_{-i} \in V_{-i}} D(v_{-i}) \sum_{w_{-i} \in V_{-i}} D(w_{-i}) v_i(X_i^{w_{-i}}) \\
&= \sum_{i \in [n]} \sum_{v_i \in V_i} D(v_i) \sum_{w_{-i} \in V_{-i}} D(w_{-i}) v_i(X_i^{w_{-i}}) \sum_{v_{-i} \in V_{-i}} D(v_{-i}) \\
&= \sum_{i \in [n]} \sum_{v_i \in V_i} \sum_{w_{-i} \in V_{-i}} D(v_i) D(w_{-i}) v_i(X_i^{w_{-i}}) \cdot 1 \\
&= \sum_{i \in [n]} \sum_{v \in V} D(v) v_i(X_i^{v_{-i}}) \\
&= \sum_{v \in V} D(v) \sum_{i \in [n]} v_i(X_i(B(v))) = SW(B).
\end{aligned}$$

Finally, the second summand is

$$\begin{aligned}
\sum_{v \in V} D(v) \sum_{w \in V} D(w) \sum_{i \in [n]} Bids^w(O_i^v) &= \sum_{v \in V} D(v) \sum_{w \in V} D(w) Bids^w(M) \\
&= \sum_{w \in V} D(w) Bids^w(M) \sum_{v \in V} D(v) \\
&= \sum_{w \in V} D(w) Bids^w(M) \cdot 1 \\
&= \sum_{w \in V} D(w) \sum_{i \in [n]} Bids_i(X_i(B(w)), B_i(w_i)) \\
&\leq \sum_{w \in V} D(w) \sum_{i \in [n]} w_i(X_i(B(w))) = SW(B).
\end{aligned}$$

The last inequality holds since for all i , the $B_i(w_i)$ contains supporting bids for any set of items including the obtained set $X_i(B(w))$.

A simple example shows that even in the full-information setting, this Bayesian price of anarchy result is tight.

3 Computing Pure Nash Equilibria

In this section we consider the following full-information game: The m items are sold to n bidders with fractionally-subadditive valuation functions in m independent second-price auctions. The players' valuation functions are assumed to be common knowledge.

In Subsection 3.1, we show that a pure Nash that provides a good approximation to the social welfare always exists in such games and provide a constructive way of finding one. In fact, we also prove that the price of stability [1] is 1, i.e. the optimum can always be achieved in a Nash equilibrium.

In Subsection 3.2 we show that if bidders have submodular valuation functions then such a pure Nash can be reached in polynomial time.

3.1 Fractionally-Subadditive Valuation Functions

Despite the fact that (as shown in the Introduction) some Nash equilibria may fail to provide good approximation to the social-welfare, we present a constructive way for finding a pure Nash that yields a 2-approximation. We introduce a natural procedure we call the POTENTIAL PROCEDURE which always reaches such an equilibrium. The POTENTIAL PROCEDURE is a simple myopic procedure for fractionally-subadditive bidders⁸.

For every i let $A_i = \{a_1^i, \dots, a_{m_i}^i\}$ be a set of additive valuations such that for every $S \subseteq M$ $v_i(S) = \max_{a \in A_i} a(S)$. Recall that since v_i is fractionally subadditive such A_i must exist. Informally, the POTENTIAL PROCEDURE simply starts with some arbitrary supporting bids (corresponding to some maximizing additive valuations) and let players best-reply, one by one, to the bids of other players by switching to new supporting bids.

The Procedure:

1. Initialize $b_i^*(j) \leftarrow 0, S_i \leftarrow \emptyset, r_j \leftarrow 0$, for $i = 1, \dots, n$ and $j = 1, \dots, m$.
2. While there is a bidder i such that $S_i \neq \arg \max_{S \subseteq M} (v_i(S) - \sum_{j \in (S \setminus S_i)} r_j)$:
 - (a) Let $T = \arg \max_{S \subseteq M} (v_i(S) - \sum_{j \in (S \setminus S_i)} r_j)$. Let $a \in A_i$ be such that $v_i(T) = a(T)$.
 - (b) Set $b_i^*(j) \leftarrow 0$ and $r_j \leftarrow 0$ for all $j \in S_i$.
 - (c) Set $b_i^*(j) \leftarrow a(j)$ for all $j \in T$. Set $r_j \leftarrow a(j)$ for all $j \in T$.
 - (d) Set $S_i \leftarrow T$.
 - (e) For all $k \neq i$ set $S_k \leftarrow S_k \setminus S_i$, and set $b_k^*(j) \leftarrow 0$ for all $j \in S_i$.
3. Output $b^* = (b_1^*, \dots, b_n^*)$.

Observe that in the definition we do not require $a(j) \geq r_j$ in lines (a)–(c) (i.e., that r_j increase). However, this follows from the fact that T maximizes $v_i(S) - \sum_{j \in (S \setminus S_i)} r_j$. We use a potential-function argument and the fractional-subadditivity of the bidders to show that the POTENTIAL PROCEDURE eventually converges to a “good” pure Nash.

Theorem 2. *If the valuation functions are fractionally subadditive, then the POTENTIAL PROCEDURE converges to a pure Nash equilibrium that provides a 2-approximation to the optimal social-welfare.*

⁸ This procedure requires bidder i to be able to determine which bundle he would prefer most, given a vector of per-item payments $r = (r_1, \dots, r_m)$. That is, to declare a bundle S for which $v_i(S) - \sum_{j \in S} r_j$ is maximized. This type of query is called a *demand query* and is very common in combinatorial auctions literature (see, e.g., [2, 4, 6, 7]).

We note that the proof shows that even the optimal social-welfare can be obtained in a pure Nash equilibrium and so the price of stability is 1. So, we have a natural procedure, that is essentially a best response sequence of the players, that leads to a pure Nash equilibrium. But, how long will it take the POTENTIAL PROCEDURE to converge? A non-trivial construction shows that unfortunately the worst case running time is exponential in n and m .

Theorem 3. *There is an instance with 2 bidders, each with a fractionally-subadditive valuation function, on which the POTENTIAL PROCEDURE converges after $\Omega(2^m)$ steps.*

Theorem 3 leaves us with two interesting open questions: First, will the POTENTIAL PROCEDURE converge in polynomial time if the valuation functions are submodular? Second, does the POTENTIAL PROCEDURE always run in time that is polynomial in the size of the sets of additive valuations that underlie every fractionally-subadditive valuation? An affirmative answer to this question would imply that the POTENTIAL PROCEDURE runs in polynomial time if the bidders have fractionally-subadditive valuations encoded in a bidding language (see [14, 13, 4, 5]). We note that in the instance in the proof of Theorem 3 the size of the sets of additive valuations was exponential n and m .

3.2 Submodular Valuation Functions

In this subsection, we focus on submodular valuation functions. We show that one can find, in polynomial time, a pure Nash equilibrium that also satisfies the premises of Theorem 1. The procedure we present exploits the algorithm due to Lehmann et al. [13]. This procedure, which we will call the MARGINAL-VALUE PROCEDURE, therefore provides a 2-approximation to the optimal social welfare.

The Procedure:

1. Fix an arbitrary order on the items. W.l.o.g. let this order be $1, \dots, m$.
2. Initialize $S_i \leftarrow \emptyset$, and $r_j \leftarrow 0$, for $i = 1, \dots, n$, and $j = 1, \dots, m$.
3. For each item $j = 1, \dots, m$:
 - (a) Let $i = \arg \max_{t \in N} v_t(S_t \cup \{j\}) - v_t(S_t)$. Set $S_i \leftarrow S_i \cup \{j\}$.
 - (b) Set $r_j \leftarrow \max_{t \in N} v_t(S_t \cup \{j\}) - v_t(S_t)$.
4. For every bidder i set $b_i^*(j) \leftarrow r_j$ for all $j \in S_i$ and $b_i^*(j) \leftarrow 0$ for all $j \notin S_i$.
5. Output $b^* = (b_1^*, \dots, b_n^*)$.

Observe, that the resulting n -tuple of bid-vectors b^* is such that for each item, only one bidder offers a non-zero bid for that item. This is due to the fact that we are dealing with a complete-information setting (intuitively, if a bidder does not win an item he might as well not bid on it). Also notice that the MARGINAL-VALUE PROCEDURE only requires m rounds and so ends in polynomial time.

Theorem 4. *If the valuation functions are submodular then a pure Nash equilibrium that provides a 2-approximation to the optimal social-welfare can be computed in polynomial time.*

Acknowledgements

We thank Noam Nisan for enriching conversations that lead to the writing of this paper.

References

1. E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *FOCS 04*, pages 295–304, 2004.
2. Liad Blumrosen and Noam Nisan. On the computational power of iterative auctions I: demand queries. Discussion paper no. 381, The Center for the Study of Rationality, The Hebrew University. An extended abstract in EC'05 contained preliminary results.
3. E. H. Clarke. Multipart pricing of public goods. *Public Choice*, pages 17–33, 1971.
4. Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *The 37th ACM symposium on theory of computing (STOC)*, 2005.
5. Shahar Dobzinski and Michael Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In SODA'06.
6. Uriel Feige. On maximizing welfare where the utility functions are subadditive. In STOC'06.
7. Uriel Feige and Jan Vondrak. The allocation problem with submodular utility functions. Manuscript, 2006.
8. Martin Gairing, Burkhard Monien, and Karsten Tiemann. Selfish routing with incomplete information. In *Proceedings of the 17th Annual ACM Symposium on Parallel Algorithms (SPAA)*, pages 203–212, 2005.
9. Dinesh Garg and Yadati Narahari. Price of anarchy of network routing games with incomplete information. In *Internet and Network Economics, First International Workshop (WINE)*, pages 1066–1075, 2005.
10. T. Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.
11. J. C. Harsanyi. Games with incomplete information played by 'bayesian' players, parts i ii and iii. *Management science*, 14, 1967-68.
12. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.
13. Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *ACM conference on electronic commerce*, 2001.
14. Noam Nisan. Bidding and allocation in combinatorial auctions. In *ACM Conference on Electronic Commerce*, 2000.
15. Noam Nisan and Amir Ronen. Algorithmic mechanism design. In *STOC*, 1999.
16. Noam Nisan and Amir Ronen. Computationally feasible vcg-based mechanisms. In *ACM Conference on Electronic Commerce*, 2000.
17. T. Roughgarden and Eva Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236 – 259, 2002.
18. Adrian Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *FOCS*. IEEE Computer Society, 2002.
19. W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.
20. Jan Vondrak. Optimal approximation for the submodular welfare problem in the value oracle model. In STOC'08.