Exploring Leximin Principle for Fair Core-Selecting Combinatorial Auctions: Payment Rule Design and Implementation

Hao Cheng[1] [*] , **Shufeng Kong**[2] [†] , **Yanchen Deng**[3] , **Caihua Liu**[4] , **Xiaohu Wu**[5] ,
Bo An[3] , **Chongjun Wang**[1]

[1]State Key Laboratory for Novel Software Technology, Nanjing University, China
[2]School of Software Engineering, Sun Yat-Sen University, China
[3]School of Computer Science and Engineering, Nanyang Technological University, Singapore
[4]School of Artificial Intelligence, Guilin University of Electronic Technology, China
[5]National Engineering Research Center of Mobile Network Technologies, Beijing University of Posts and Telecommunications, China
chengh@smail.nju.edu.cn, kongshf@mail.sysu.edu.cn, ycdeng@ntu.edu.sg, caihua.liu@guet.edu.cn,
xiaohu.wu@bupt.edu.cn, boan@ntu.edu.sg, chjwang@nju.edu.cn

Abstract

Core-selecting combinatorial auctions (CAs) restrict the auction result in the core such that no coalitions could improve their utilities by engaging in collusion. The minimum-revenue-core (MRC) rule is a widely used core-selecting payment rule to maximize the total utilities of all bidders. However, the MRC rule can suffer from severe unfairness since it ignores individuals' utilities. To address this limitation, we propose to explore the leximin principle to achieve fairness in core-selecting CAs since the leximin principle prefers to maximize the utility of the worst-off; the resulting bidder-leximin-optimal (BLO) payment rule is then theoretically analyzed and an effective algorithm is further provided to compute the BLO outcome. Moreover, we conduct extensive experiments to show that our algorithm returns fairer utility distributions and is faster than existing algorithms of core-selecting payment rules.

1 Introduction

Combinatorial auctions (CAs) have found applications in many real-world auction problems, including procurement auctions [Sandholm, 2007] and government spectrum auctions [Cramton, 2013; Ausubel and Baranov, 2017; Leyton-Brown et al., 2017]. CAs are attractive as they can produce efficient outcomes even when bidders have complex preferences on bundles of heterogeneous goods. A combinatorial auction mechanism often includes two components: 1) an allocation rule to specify the bundle of items assigned to each bidder; 2) a payment rule to specify the price a bidder should pay for his bundle. Note that an allocation rule is normally

obtained by solving the NP-hard winner determination problem [Rothkopf et al., 1998]. In this paper, we will focus on the payment rules' designs and implementations.

To address the truthfulness issue of bidding, where bidders are unwilling to give their real valuations of items but report false information to improve their utilities, the well-known Vickrey-Clarke-Groves (VCG) payment rule [Vickrey, 1961; Clarke, 1971; Groves, 1973] ensures that the optimal strategy for each bidder is to bid their true valuations of the items. However, the VCG payment rule still suffers from various other issues that prohibit it from real-world adoption, e.g., arbitrarily low revenue for the seller, which may cause collusion between the seller and some bidders [Ausubel et al., 2006].

The drawbacks of the VCG payment rule have motivated the development of core-selecting payment rules [Day and Raghavan, 2007], which provide a principled way to ensure that there are no opportunities for collusion between the seller and bidders. The core is a set of all possible payment outcomes where the utility of any coalition of bidders in the auction cannot be improved by engaging in collusion [Gillies, 1959]. However, the core can be exponentially large, and thus, selecting a suitable payment outcome from the core is a critical problem in core-selecting CAs. Many efforts have been devoted to solving this problem by designing some principles of selection [Day and Raghavan, 2007; Erdil and Klemperer, 2010; Day and Cramton, 2012; Lubin et al., 2015; Bünz et al., 2018]. Among them, the minimum-revenue-core (MRC) rule is particularly popular, where it achieves the largest total incentive by selecting the core outcome that maximizes the total utilities of all bidders.

However, it is observed that the MRC rule may cause severe unfairness, since some winners may be forced to sacrifice their utilities in the MRC rule. A detailed example to illustrate this unfair phenomenon caused by the MRC rule is provided in Example 1 of this paper, where, in the optimal MRC payment outcome, a winner's utility is 0 while he could have obtained a positive utility. More importantly, this zero-utility phenomenon is not rare since nearly 30% of winners are forced to get a utility of zero on average in our data gener-

---

[*]The work was done during a visiting scholar program at Nanyang Technological University.
[†]Corresponding author

ated by CATS [Leyton-Brown *et al.*, 2000], which simulates bidding behavior in many realistic economic environments.

To address the limitation of the MRC rule, we propose to explore the leximin principle to achieve fairness in core-selecting CAs and propose a novel bidder-leximin-optimal (BLO) payment rule that selects a core outcome based on the leximin principle which prefers to maximize the utility of the worst-off [Rawls, 2004]. Informally, the BLO rule maximizes the minimum utility, the second minimum utility and so forth among all the bidders, which achieves egalitarian fairness. Furthermore, we analyze the BLO outcome theoretically and provide an effective algorithm to compute it. We summarize our contributions[1] as follows:

1. We investigate the unfair phenomenon of the MRC rule and propose the BLO rule to achieve fairness in the core-selecting CAs.

2. We give a theoretical analysis of the BLO rule and show that the BLO outcome is unique and Pareto optimal. Moreover, we offer that the worst-case relative ratio between a bidder's utility (resp. the total utility of all bidders) in the BLO outcome and the maximum utility among all the core outcomes is $\frac{1}{|W|}$ (resp. $\frac{4}{|W|+2+\frac{|W| \mod 2}{|W|}}$), where $|W|$ is the winner size.

3. Finally, we propose an exact algorithm named WF-CGS-CR to compute the BLO outcome. Our WF-CGS-CR algorithm consists of the water-filling framework, a constraint generation search subroutine, and a constraint-reuse strategy. Given the oracle access for winner determination, it is effective in time polynomial in the number of winners. Extensive experiments show that our algorithm returns fairer utility distributions and is also faster than existing algorithms of core-selecting payment rules.

2 Related Work

As a practical mechanism, core-selecting CAs and the MRC rule were proposed in [Day and Raghavan, 2007] and used in several scenarios, including selling wireless spectra [Cramton, 2013] and online advertising [Goel *et al.*, 2015]. Then the quadratic rule [Day and Cramton, 2012] was proposed to select the MRC point nearest to the VCG point in l_2-distance. They also study some other reference points like all-zero points. Recently, some works design the core-selecting payment rule by approximating Bayesian equilibrium of different core outcomes [Lubin *et al.*, 2015; Bünz *et al.*, 2018].

Except for the incentives, fairness is an essential metric in the core-selecting mechanisms. [Lubin *et al.*, 2015] studied the fairness measured by the Gini coefficient and proposed a fractional rule that selects a point fraction to the VCG point performs best. However, their experiments show that relaxing the MRC constraint does not change the equilibrium properties; and there are no systematic shifts in efficiency, fairness, and aggregate incentives. There are many principles to define fairness. In this paper, we adopt the leximin principle in the

---

[1]Note that we will refer all proofs to the appendix. Our appendix and code are available at https://github.com/ha0cheng/Fair-BLO.

core-selecting CAs. The method is similar to the egalitarian core proposed by Arin [Arin and Inarra, 2001], which selects the utility distribution in the cooperative game via the leximin principle [Rawls, 2004]. The difference is that the total utility is not constant in core-selecting CAs.

As for the pricing algorithms, given access to the separation oracle of winner determination, one can use the Ellipsoid algorithm [Grötschel *et al.*, 1981] to optimize convex objectives (e.g., bidders' total utility) exactly over the core polytope in polynomial time. However, this approach is rather slow in practice, and Core Constraint Generation (CCG) algorithm [Day and Raghavan, 2007] was proposed as an effective heuristic algorithm to solve such an optimization problem, which is commonly used in core-selecting CAs [Day and Cramton, 2012; Cramton, 2013; Bünz *et al.*, 2015].

However, it is hard to formulate a global objective for the BLO rule, making it difficult to use the above algorithms directly. Then, Fast Core algorithm [Niazadeh *et al.*, 2022] was proposed to compute an approximate Pareto optimal outcome in the rich advertisement, which can be used to compute an approximate BLO core outcome. Instead of the approximate algorithm, we focus on the exact algorithm. Moreover, the experiments show that our exact algorithm is much faster than Fast Core.

3 Preliminaries

Combinatorial auctions are specifically designed for domains where bidders can have complex preferences over bundles of heterogeneous items. Formally, let $N = \{1, \ldots, n\}$ be a set of bidders and s be the seller. Let $M = \{a, b, c, \ldots\}$ be the set of items with $|M| = m$. Each bidder i has a valuation function $v_i : 2^M \to \mathbb{R}$ which specifies the bidder's valuation for every possible bundle of items $S \in 2^M$. Unless stated otherwise, bidders will only need to give their values for bundles in which they are interested, and all other bundles will be given values 0. Also, we assume that $v_i(\emptyset) = 0 \ (\forall i \in N)$.

A CA consists of two steps: step 1 computes an item allocation outcome, and step 2 computes an item payment outcome. An allocation rule specifies the bundle of items a_i assigned for each bidder i, where $a_i \in 2^M$ and $a_i \cap a_j = \emptyset \ (\forall i \neq j)$, and a payment rule specifies the price p_i that the bidder i should pay for his bundle a_i. Obviously, p_i must be 0 if a_i is empty.

Let π_i denote the utility of an auction participant $i \in N \cup \{s\}$. If i is the bidder, we assume that his utility is quasi-linear, i.e., $\pi_i = v_i(a_i) - p_i$. If i is the seller s, his utility is $\pi_s = \sum_{i \in N} p_i$, which is also known as the revenue. The *social welfare* is the summation of all participants' utility:

$$\pi_s + \sum_{i \in N} \pi_i = \sum_{i \in N} p_i + \sum_{i \in N} (v_i(a_i) - p_i) = \sum_{i \in N} v_i(a_i) \quad (1)$$

One can observe that the social welfare does not depend on the payment, and an allocation rule can be computed by maximizing the social welfare as follows:

$$\max_{\{a_i | a_i \subseteq M, i \in N, a_i \cap a_j = \emptyset, \forall i \neq j\}} \sum_{i \in N} v_i(a_i) \quad (2)$$

Note that the above optimization problem is also known as the *winner determination problem*. Let $w(\{v_i(\cdot)\}_{i \in N})$ be the

maximum social welfare given the valuation function profile $\{v_i(\cdot)\}_{i \in N}$. Once we have a social welfare maximizing allocation $\{a_i | i \in N, a_i \cap a_j = \emptyset, \forall i \neq j\}$ by solving the winner determination problem, in step 2, we need to decide the price p_i for each allocated item bundle $a_i (i \in N)$.

The simplest payment rule is the so-called *pay-as-bid* payment rule, where a winner $i(a_i \neq \emptyset)$ pays the price $p_i = v_i(a_i)$ for buying the item bundle a_i. However, the pay-as-bid payment rule suffers from the truthfulness issue, where bidders are unwilling to report the true valuation functions but strategically give false information instead. Therefore, in step 2, the famous VCG mechanism is proposed to offer bidders incentives to ensure that bidding truthfully is the dominant strategy. Unfortunately, VCG can lead to very low or even zero revenue, which is unstable and may cause collusion between the seller and some bidders.

3.1 Core-selecting Payment Rules

The weaknesses of VCG payment rule have motivated the development of *core-selecting payment rules*, which provide a principled way to ensure that the revenue of the auction is high enough such that there are no opportunities for collusion.

Formally, let $S \subseteq N \cup \{s\}$ denote a coalition of participants. Then, if $s \notin S$, the social welfare of S must be 0 since the auction cannot happen without the seller's participation; If $s \in S$, the social welfare of S can be computed as follows:

$$\pi_s + \sum_{i \in S \setminus \{s\}} \pi_i = \sum_{i \in S \setminus \{s\}} p_i + \sum_{i \in S \setminus \{s\}} (v_i(a_i) - p_i) = \sum_{i \in S \setminus \{s\}} v_i(a_i)$$
(3)

The coalition characteristic function, $\mathcal{F} : 2^{N \cup \{s\}} \to \mathbb{R}$, which indicates the optimal social welfare of each coalition S is then defined as follows:

$$\mathcal{F}(S) = \begin{cases} w(\{v_i(\cdot)\}_{i \in S \setminus \{s\}}), \text{if } s \in S \\ 0, \text{if } s \notin S \end{cases}$$
(4)

The *core* is the set of all possible outcomes where the total utility of any coalition cannot be improved by engaging in collusion [Gillies, 1959]. The definition is given as follows:

Definition 1 (Core). *The core of a CA is a set of outcomes satisfying the following two conditions:*

- *Efficiency:*

$$\sum_{i \in N \cup \{s\}} \pi_i = \mathcal{F}(N \cup \{s\})$$
(5)

- *Coalitional rationality:*

$$\sum_{i \in S} \pi_i \geq \mathcal{F}(S), \forall S \subseteq N \cup \{s\}$$
(6)

Note that the efficiency condition requires that the allocation outcome must be optimal, namely, the allocations are obtained by solving the winner determination problem. It is possible that there is more than one optimal allocation, and we break ties by selecting a random optimal allocation $\{a_i^*\}_{i \in N}$. Denote by W the winner set, i.e., $W = \{i | a_i^* \neq \emptyset\}$. On the other hand, the coalitional rationality condition guarantees that the maximum social welfare that any coalition S

could get (i.e., RHS of the inequality) is not larger than what the coalition can already get under the current outcome (i.e., LHS of the inequality). *Finally, a core-selecting CA selects a social welfare maximizing allocation and picks a payment outcome in the core.*

If coalition S does not include the seller, then $\mathcal{F}(S) = 0$. Thus the corresponding constraints of coalitional rationality can be summarized as follows:

$$\pi_i \geq 0, \forall i \in N$$
(7)

If $s \in S$, substitute Eq.5 into Eqs.6 and we can arrange the constraints as follows:

$$\sum_{i \in N \setminus S} \pi_i \leq w(N) - w(S), \forall S \subseteq N$$
(8)

where $w(S)$ is the abbreviation of $w(\{v_i(\cdot)\}_{i \in S})$. Eqs.7 and Eqs.8 form the complete core constraint set, which includes $2^N + N$ constraints in total. We use the utility to analyze in the following discussion and each bidder's payment price is mapped by $p_i = v_i(a_i^*) - \pi_i$.

3.2 Minimum-Revenue-Core Payment Rule

Core-selecting CAs are not guaranteed to be truthful [Goeree and Lien, 2016]. To maximize the incentives of truthful bidding, the *Minimum-Revenue-Core rule* (MRC) is proposed as a fundamental principle to select the core outcomes [Day and Raghavan, 2007; Day and Milgrom, 2008; Day and Cramton, 2012]. The MRC rule maximizes the total utility of all bidders, which can be computed as follows:

$$\max_{\{\pi_i\}_{i \in N} \in U} \sum_{i \in N} \pi_i$$
(9)

where U is the core of the related CA.

Unfortunately, we observe that the MRC rule could suffer from severe unfairness of utility distribution, where some winners' utilities are forced to be zero. This unfair phenomenon largely dues to the fact that the MRC rule emphasizes the total utility and ignores individuals' utilities. We will illustrate this unfair utility distribution phenomenon in the following example.

Example 1. *Consider a CA with 5 bidders and 3 items. Let $N = \{1, 2, 3, 4, 5\}$ be the set of bidders and $M = \{a, b, c\}$ be the set of items. Table 1 gives the valuation functions of all bidders. By solving the winner determination problem, we have that the winners are bidders 1, 2, and 3, thus the maximum social welfare is that $w(N) = v_1(\{a\}) + v_2(\{b\}) + v_3(\{c\}) = 6$ according to Eq.1. Similarly, $w(\{3, 4, 5\}) = 4$ and the corresponding constraint is $\pi_1 + \pi_2 \leq 6 - 4 = 2$. The other constraints are obtained similarly. After removing the redundant constraints, the core of this example is shown as follows:*

$$Core = \begin{cases} \pi_1 \geq 0, \pi_2 \geq 0, \pi_3 \geq 0 \\ \pi_1 + \pi_2 \leq 2 \\ \pi_2 + \pi_3 \leq 2 \end{cases}$$
(10)

According to the MRC rule, the unique utility distribution is $\{2, 0, 2\}$. Bidder 2 must sacrifice all of his utility even

though he could have obtained a positive utility in other core outcomes (e.g., $\{1, 1, 1\}$). In this case, bidder 2 would prefer bidding untruthfully to get more utility. For example, he can bid the value 1 for the item b, which leads to a utility of 1 in the MRC rule.

Table 1: A CA with 5 bidders and 3 items. Winning bids are marked with "*". The last three columns are utilities computed from the VCG, MRC, and our new BLO rule, respectively.

	bids					utilities		
	{a}	{b}	{c}	{a,b}	{b,c}	VCG	MRC	BLO
Bidder 1	2*					2	2	1
Bidder 2		2*				2	0	1
Bidder 3			2*			2	2	1
Bidder 4				2		0	0	0
Bidder 5					2	0	0	0

Extending this example to the general case, we have the following lemma:

Lemma 1. *Assuming that bidders bid truthfully, in the worst case, $|W| - 2$ winners may get zero utility in the MRC rule, even though there exists some core outcome that assigns them with positive utilities, where $|W|$ is the winner size.*

Moreover, we observe that the unfair phenomenon often happens in a wide range of domains and datasets when using the MRC rule. To address this issue, we will explore and propose a novel payment rule to achieve fairness.

4 Bidder-Leximin-Optimal Payment Rule

We propose to adopt the *leximin principle* to achieve fairness in core-selecting CAs, where the leximin principle prefers the utility distribution with a larger utility for the worst-off.

Definition 2 (Leximin Dominance). *Let $x, y \in \mathbb{R}^n$ be two real vectors. x leximin dominates y if the following holds:*

- *There exists some integer $0 \leq k \leq n - 1$ such that the k-smallest elements of both vectors are equal, and the $(k + 1)$-smallest element of x is larger than the $(k + 1)$-smallest element of y.*

We denote the leximin-dominant relation by "$\succ_{Lex}$", and then a *Bidder-Leximin-Optimal* (BLO) core outcome is defined as follows:

Definition 3 (Bidder-Leximin-Optimal). *Given a core U of a CA, a utility outcome $\pi \in U$ is bidder-leximin-optimal if and only if it is not leximin dominated by any utility outcome π' in the core, i.e.,*

$$\nexists \pi' \in U : \pi' \succ_{Lex} \pi \tag{11}$$

The BLO payment rule follows the leximin principle to give priority to those who are worst-off. It selects a core outcome that is BLO rather than a core outcome that maximizes the total utility of bidders. In order to illustrate the advantage of the BLO payment rule over the MRC payment rule, we consider the following example.

Example 2. *Let us recall the CA problem in example 1. The unique MRC utility outcome is $\pi = \{2, 0, 2\}$ for the winners. On the other hand, the BLO utility outcome is $\pi' = \{1, 1, 1\}$ under the core constraints in Eqs.10. One can observe that π' is a fairer utility distribution than π.*

4.1 Theoretical Analysis of the BLO Payment Rule

This section provides three theoretical results for the BLO payment rule. First, the following theorem establishes that the BLO outcome is unique and Pareto optimal.

Theorem 1. *The BLO outcome always exists, and it is unique and Pareto optimal, where no bidders can improve their utilities without harming other bidders' utilities.*

We then consider a lower bound of a bidder's utility in the BLO core outcome.

Theorem 2. *Assuming that bidders bid truthfully, given the BLO outcome $\{\pi_i^{BLO}\}_{i \in N}$, we have the following tight lower bound:*

$$\pi_i^{BLO} \geq \frac{1}{|W|} \pi_i^*$$

where π_i^* is the maximum utility for bidder i in the core.

The above lower bound shows that a bidder's utility is always larger than zero in the BLO outcome, unless it is zero in every core outcome's utility distribution. Regarding the total utility of the BLO outcome, it is clear that it cannot be better than that of an MRC outcome, because the total utility of an MRC outcome is maximum. The following result shows the ratio between the BLO outcome and the MRC outcome in terms of their total utility.

Theorem 3. *Assuming that bidders bid truthfully, given the BLO utility outcome $\{\pi_i^{BLO}\}_{i \in N}$, we have the following tight lower bound:*

$$\sum_{i \in N} \pi_i^{BLO} \geq \frac{4}{|W| + 2 + \frac{|W| \mod 2}{|W|}} \sum_{i \in N} \pi_i^{MRC}$$

where $\{\pi_i^{MRC}\}_{i \in N}$ is an MRC outcome.

One can observe that the relative ratio is 1 if the size of the winner set is 1 or 2, but the relative gap grows linearly with the winner set size $|W|$. Besides, the BLO outcome has a better guarantee than the general Pareto optimal outcome, which has a tight relative ratio guarantee of $\frac{1}{|W|-1}$ [2].

5 An Exact Algorithm for the BLO Payment Rule

In this section, we propose an exact algorithm, named WF-CGS-CR, to compute the BLO payment outcome. The WF-CGS-CR algorithm is a **water-filling (WF) algorithm** to increase all active winners' utilities uniformly in an interactive manner until reaching the BLO outcome, and a **constraint generation search (CGS) subroutine** is adopted to compute the maximum utility increment at each water-filling iteration. Moreover, we offer a **constraint-reuse (CR) strategy** to accelerate the CGS by setting a better initialization with the history of constraints. We will elaborate on each component of our WF-CGS-CR below.

Since core-selecting CAs involve solving the winner determination problem, existing works [Day and Raghavan, 2007; Day and Cramton, 2012] often assume that there is a so-called winner determination oracle to solve this problem automatically. We will follow such an assumption in this paper and introduce the concept as follows:

---

[2]More detail in the appendix.

Definition 4 ($\mathcal{WD}$). *Let $\mathcal{WD}$ be an oracle with this input-output relation:*
Input: *submitted bid profile $\{b_i(\cdot)\}_{i \in N}$, where $b_i : 2^M \to \mathbb{R}$ is bidder i's reported valuation function[3].*
Output: *a winner set W and maximum social welfare $w(\{b_i(\cdot)\}_{i \in N})$ under the submitted bids.*

In other words, we do not dig into the winner determination solving but only regard it as a black-box interface. Any state-of-the-art winner determination method can be used in this way. *Oracle complexity* is defined as the query time to the oracle in the worst case, which is the primary analysis aspect for the algorithms in this paper.

5.1 Water-filling Algorithm

Water-filling algorithm was first proposed in [Niazadeh *et al.*, 2022] to compute an approximate Pareto optimal outcome. Different from the original version, we develop the algorithm for the exact BLO outcome.

Given a utility distribution in the core, we have two kinds of winners: the active winner and the frozen winner. The active winner has the potential to increase his utility, while the frozen winner can not increase the utility further. Our water-filling algorithm starts from the all-zero utility distribution and the initial active winner set W (i.e., the original winner set). Then, at the t-th iteration, it does the following:

- Run the search subroutine to return the maximum utility increment $\Delta \pi^t$ and the frozen winner set W_f^t.

- Generate the next utility distribution by increasing each active winner's utility by $\Delta \pi^t$ and generate the next active winner set by removing W_f^t from the current active winner set W_a^t.

The iteration stops when the active winner set becomes empty, i.e., no one can increase his utility anymore. The pseudo-code is shown in Alg.1.

A search subroutine is adopted to find the maximum utility increment and the frozen winner set at each iteration. Formally, given that the current utility distribution $\{\pi_i^t\}_{i \in N}$ is in the core, the optimization problem for the search subroutine can be formulated as follows:

$$\max_{\{\Delta \pi | \{\pi_i^t + \Delta \pi\}_{i \in W_a^t} \in U\}} \Delta \pi \quad \text{(P0)}$$

where $\{\pi_i^t + \Delta \pi\}_{i \in W_a^t}$ means increasing the utilities of winners in W_a^t by $\Delta \pi$ while the other winners' utilities remain unchanged. After the utility increment, some winners' utilities can not increase anymore, i.e., the frozen winner set at this iteration. Then we have the following lemma:

Lemma 2. *Given that the results computed by the search subroutine are correct, the water-filling algorithm returns the BLO outcome.*

5.2 Constraint Generation Search Subroutine

A naive method to solve the problem (P0) is obtaining all the core constraints, but it needs to query the winner determination oracle exponential times according to Eqs.8. To avoid

---

[3]We regard the bidders' reported valuation functions as their true valuation functions in the computation part.

Algorithm 1 Water-filling algorithm for finding the exact BLO outcome

Input: bid profile $\{b_i(\cdot)\}_{i \in N}$.
Output: the BLO outcome.

1: Step $t \leftarrow 1$
2: $W, w(N) \leftarrow \mathcal{WD}(\{b_i(\cdot)\}_{i \in N})$
3: Initialize the utility distribution $\{\pi_i^t\}_{i \in N}$ as all-zero
4: Initialize the active winner set W_a^t as W
5: **while** W_a^t is not empty **do**
6:     Run the search subroutine to return the maximum utility increment $\Delta \pi^t$ and the frozen winner set W_f^t
7:     Generate the next utility distribution:

$$\pi_i^{t+1} = \begin{cases} \pi_i^t + \Delta \pi^t, \text{ if } i \in W_a^t \\ \pi_i^t, \text{ otherwise} \end{cases}$$

8:     Generate the next active winner set:

$$W_a^{t+1} = W_a^t \backslash W_f^t$$

9:     Step $t \leftarrow t + 1$
10: **return** $\{\pi_i^t\}_{i \in N}$

---

this, in this section, we first formulate the problem (P0) to an equivalent mathematical format and then solve it through the constraint generation search.

Optimization Problem Formulation
In the original format (P0), the restriction condition is to ensure the utility distribution is still in the core, in other words, satisfying all the core constraints. According to Eqs.8, the constraints can be written as:

$$\sum_{i \in N \backslash S} \pi_i^t + |W_a^t \backslash S| \Delta \pi \leq w(N) - w(S), \ \forall S \subseteq N \quad (12)$$

If $W_a^t \subseteq S$, then $|W_a^t \backslash S| = 0$, and the corresponding constraints can be ignored. Thus, we can arrange the remaining constraints as follows:

$$\Delta \pi \leq \frac{w(N) - w(S) - \sum_{i \in N \backslash S} \pi_i^t}{|W_a^t \backslash S|} \quad (13)$$

where $S \subseteq N$ and $W_a^t \not\subseteq S$. Every such constraint provides an upper bound for $\Delta \pi$. Denote by $\Delta \pi^S$ the upper bound corresponding to S, i.e., RHS in Eq.13. Thus, finding the maximum utility increment satisfying all the core constraints is equivalent to finding the minimum upper bound above. Problem (P0) is equivalent to the following optimization problem:

$$\min_{\{S | S \subseteq N, W_a^t \not\subseteq S\}} \frac{w(N) - w(S) - \sum_{i \in N \backslash S} \pi_i^t}{|W_a^t \backslash S|} \quad \text{(P1)}$$

Let S^* be one of the optimal coalitions for problem (P1), thus $\Delta \pi^t = \Delta \pi^{S^*}$. After the increment, the corresponding core constraint of S^* would be tight for the new utility distribution $\{\pi_i^{t+1}\}_{i \in N}$, thus we have

$$\sum_{i \in N \backslash S^*} \pi_i^{t+1} = w(N) - w(S^*) \quad (14)$$

According to Eq.14, bidders belonging to $N \backslash S^*$ cannot increase their utilities without violating this constraint, i.e., they are frozen. Therefore, the frozen winner set W_f^t is $W_a^t \backslash S^*$ at this iteration. Problem (P1) is actually a nesting optimization problem with a fractional format since $w(S)$ is corresponding to the winner determination problem. Thus it is intractable to solve directly, and we propose to use the constraint generation method that only considers the most useful constraints.

Constraint Generation Search
Generally speaking, CGS starts from an upper bound of the utility increment, then reduces the upper bound iteratively through generated constraints until it is the minimum.

Formally, CGS initializes the upper bound $\Delta \bar{\pi}$ as ∞ and the frozen winner set as W_a^t. It checks whether this upper bound is the minimum by the following equation:

$$w(N) - \sum_{i \in N} \tilde{\pi}_i = w_B \tag{15}$$

where $\{\tilde{\pi}_i\}_{i \in N}$ is the utility distribution by setting the utility increment as $\Delta \bar{\pi}$ and w_B is the maximum social welfare under the truncated bid profile $\{\max(b_i(\cdot) - \tilde{\pi}_i, 0)\}_{i \in N}$. If Eq.15 is satisfied, the current upper bound $\Delta \bar{\pi}$ is the minimum, and CGS will return the result. Otherwise, a smaller upper bound is computed by

$$\Delta \bar{\pi} = \Delta \bar{\pi} - \frac{w_B - (w(N) - \sum_{i \in N} \tilde{\pi}_i)}{|\bar{W}_f|} \tag{16}$$

The corresponding frozen winner set is updated by

$$\bar{W}_f = W_a^t \backslash B \tag{17}$$

where B is the winner set under the truncated bid profile $\{\max(b_i(\cdot) - \tilde{\pi}_i, 0)\}_{i \in N}$. If the new frozen winner set satisfies $|\bar{W}_f| = 1$, then this upper bound is the minimum. Otherwise, CGS will start a new iteration with the smaller upper bound. The oracle complexity of CGS is given as follows:

Lemma 3. *Constraint generation search requires at most* $|W_a^t|$ *queries to the oracle* $\mathcal{WD}$ *to solve the optimization problem (P1).*

Assume that the iteration number of the water-filling algorithm is T. Naturally, since $W = W_a^1 \supset ... \supset W_a^T = \emptyset$, the query time of the CGS subroutine in the water-filling algorithm is at most $|W| + (|W| - 1) + ... + 1 = \frac{|W|(|W|+1)}{2}$.

5.3 Constraint-reuse Strategy
To avoid CGS starting from the naive upper bound ∞, we propose the constraint-reuse strategy to set a better initialization through the history of constraints.

Formally, we store the binary tuple $(W_f^S, \Delta \pi^S)$ in the stored constraint set $\mathbf{C}$, where W_f^S is the frozen winner set if S is an optimal coalition, i.e., $W_f^S = W_a^t \backslash S$. Then in each CGS subroutine, we first update each binary tuple $\mathbf{C}$ based on the new utility distribution and active winner set. The recurrence formula is given as follows:

$$\begin{cases} \acute{W}_f^S = W_f^S \cap W_a^t \\ \Delta \acute{\pi}^S = \frac{|W_f^S|}{|\acute{W}_f^S|}(\Delta \pi^S - \Delta \pi^{t-1}) \end{cases} \tag{18}$$

Algorithm 2 Constraint generation search with the constraint-reuse strategy

---

Input: bid profile $\{b_i(\cdot)\}_{i \in N}$, maximum social welfare $w(N)$, current utility distribution $\{\pi_i^t\}_{i \in N}$, active winner set W_a^t, last utility increment $\Delta \pi^{t-1}$, constraint set $\mathbf{C}$.
Output: the maximum utility increment $\Delta \pi^t$ and the frozen winner set W_f^t.

/*Constraint-reuse Strategy*/
1: Initialize $\Delta \bar{\pi} \leftarrow \infty$; $\bar{W}_f \leftarrow W_a^t$
2: **for** $(W_f^S, \Delta \pi^S)$ in $\mathbf{C}$ **do**
3: $\quad \mathbf{C} \leftarrow \mathbf{C} \backslash \{(W_f^S, \Delta \pi^S)\}$
4: $\quad$ **if** $W_f^S \cap W_a^t \neq \emptyset$ **then**
5: $\quad\quad \acute{W}_f^S \leftarrow W_f^S \cap W_a^t$; $\Delta \acute{\pi}^S = \frac{|W_f^S|}{|\acute{W}_f^S|}(\Delta \pi^S - \Delta \pi^{t-1})$
6: $\quad\quad$ **if** $\Delta \acute{\pi}^S < \Delta \bar{\pi}$ **then** $\Delta \bar{\pi} \leftarrow \Delta \acute{\pi}^S$; $\bar{W}_f \leftarrow \acute{W}_f^S$
7: $\quad\quad \mathbf{C} \leftarrow \mathbf{C} \cup \{(\acute{W}_f^S, \Delta \acute{\pi}^S)\}$
/*Constraint Generation Search*/
8: **while** True **do**
9: $\quad$ Generate the utility distribution by upper bound $\Delta \bar{\pi}$:

$$\tilde{\pi}_i = \begin{cases} \pi_i^t + \Delta \bar{\pi}, \text{ if } i \in W_a^t \\ \pi_i^t, \text{ otherwise} \end{cases}$$

10: $\quad B, w_B \leftarrow \mathcal{WD}(\{\max(b_i(\cdot) - \tilde{\pi}_i, 0)\}_{i \in N})$
11: $\quad$ **if** $w(N) - \sum_{i \in N} \tilde{\pi}_i = w_B$ **then** break
12: $\quad \Delta \bar{\pi} \leftarrow \Delta \bar{\pi} - \frac{w_B - (w(N) - \sum_{i \in N} \tilde{\pi}_i)}{|\bar{W}_f|}$; $\bar{W}_f \leftarrow W_a^t \backslash B$
13: $\quad \mathbf{C} \leftarrow \mathbf{C} \cup \{(\bar{W}_f, \Delta \bar{\pi})\}$
14: $\quad$ **if** $|\bar{W}_f| = 1$ **then** break
15: **return** $\Delta \bar{\pi}, \bar{W}_f$

Above all, we have introduced the three components of the WF-CGS-CR algorithm and the runtime guarantee is shown as follows:

Theorem 4. *The WF-CGS-CR algorithm requires at most* $\frac{|W|(|W|+1)}{2} + 1$ *queries to oracle* $\mathcal{WD}$ *and an additional time complexity* $O(|W|^5)$ *to obtain the BLO outcome, where* $|W|$ *is the winner size.*

By adopting this formula, we transfer each binary tuple $(W_f^S, \Delta \pi^S)$ to a new binary tuple $(\acute{W}_f^S, \Delta \acute{\pi}^S)$ before CGS. Eventually, the minimum upper bound in the constraint set is computed and serves as the initial upper bound for CGS. Besides, each new binary tuple is stored in the constraint set if $\acute{W}_f^S$ is not empty. The complete pseudo-code for the CGS subroutine is shown in Alg.2. Note that this strategy would increase some extra computation complexity that is minor since the winner determination oracle is the main time-consuming part.

6 Experiments
In the experiments, we use the CATS (Combinatorial Auction Test Suite) [Leyton-Brown *et al.*, 2000] as the CA instance generator. Same as [Bünz *et al.*, 2015], we choose six representative CATS distributions: Arbitrary, Decay(L4),

Algorithm	Arbitrary	Decay(L4)	Matching	Paths	Regions	Scheduling
VCG	12.49 (4.52)	1.83 (0.87)	**0.75** (0.03)	**2.50** (0.18)	2.35 (0.45)	**0.83** (0.36)
MRC	22.09 (10.30)	4.84 (2.35)	2.12 (0.62)	5.70 (1.54)	4.42 (1.40)	2.37 (1.93)
MRC-VCG	23.32 (10.86)	4.84 (2.21)	2.23 (0.63)	5.86 (1.51)	4.76 (1.48)	2.90 (2.49)
MRC-Zero	23.48 (10.91)	5.17 (2.74)	2.23 (0.64)	5.90 (1.55)	4.87 (1.47)	2.99 (2.79)
Fast Core	33.26 (15.24)	6.80 (3.19)	7.67 (1.00)	21.46 (3.07)	7.23 (2.54)	11.48 (8.9)
WF-CGS	9.67 (4.85)	3.57 (1.83)	2.98 (0.42)	10.37 (1.56)	2.37 (1.13)	2.36 (0.97)
WF-CGS-CR	**7.15** (3.64)	**1.81** (0.93)	1.22 (0.18)	3.25 (0.36)	**1.55** (0.54)	1.58 (1.17)

Table 3: Results for the run time. Average run times (in seconds) over 50 CA instances are provided, with standard error in the parentheses. Each CA instance is generated by CATS, with 64 goods and 1000 bids. The lowest average run time in each column is marked in **bold**.

Matching, Paths, Regions, and Scheduling. For each distribution, we generate 50 CA instances with 64 goods and 1000 bids. All the experiments were run on a high-performance computer with a 3.10GHz Intel core and 16GB of RAM. Besides, the winner determination problem is transformed into a format of integer programming, solved by the solver CPLEX 20.1.0 [Manual, 1987]. The following payment rules are implemented for the comparison:

- *VCG*: the VCG payment rule, where each winner's utility is $\pi_i^{VCG} = w(N) - w(N \setminus \{i\})$.

- *MRC*: the MRC payment rule that selects one core outcome with the maximum total utility of bidders [Day and Raghavan, 2007].

- *MRC-VCG*: a quadratic payment rule that selects the MRC point nearest to the VCG point in l_2-distance [Day and Cramton, 2012].

- *MRC-Zero*: a quadratic payment rule that selects the MRC point nearest to the all-zero point in l_2-distance [Day and Cramton, 2012].

- *Fast Core*: the approximate BLO payment rule implemented through the water-filling algorithm with the binary search [Niazadeh *et al.*, 2022].

We use the state-of-the-art Core Constraint Generation (CCG) algorithm to compute all the MRC outcomes above [Day and Raghavan, 2007]. Unfortunately, it has no theoretical guarantee for the query time; in other words, its oracle complexity may be exponential. Instead, the oracle complexity of Fast Core is $O(|W| \log \frac{|W|}{\epsilon})$, where ϵ is the parameter representing the approximate gap. Same as [Niazadeh *et al.*, 2022], ϵ is set as 0.01 in our experiments.

6.1 Experimental Results

The auction results are shown in Table 2. We can see that core-selecting payment rules enhance the seller's revenue compared to the VCG payment rule. Among the core-selecting payment rules, the BLO rule achieves a fairer utility distribution with a tiny reduction of the total utility (i.e., from 173.65 to 147.17). Moreover, the MRC rules produce an unfair utility distribution, with nearly 30% winners obtaining zero utility, which impairs the incentives for bidders. Instead, the BLO rule has only 1.69% zero-utility winners on average, thus providing a guarantee for the worst-off. Fast Core achieves approximate fairness, but there are still 12.71% zero-utility winners.

Payment rule	Revenue	Total utility	Min ↑	Std↓	Zero ratio↓
VCG	11661.54	513.04	5.03	13.63	1.69%
MRC-VCG	12000.93	173.65	0.01	10.36	31.63%
MRC-Zero	12000.93	173.65	0.05	9.54	28.57%
Fast Core	12046.95	127.63	0.83	5.94	12.71%
BLO	12027.41	147.17	1.18	6.77	1.69%

Table 2: Results for the auction. The measurements include the revenue, the total utility, the minimum utility (Min), the standard deviation (Std), and the percentage of the zero-utility winners among all winners (Zero ratio). All results are averages over 300 CA instances, with 50 instances for each distribution. Note that ↑ indicates larger values are better, and ↓ indicates smaller values are better in terms of fairness.

The average run time results are shown in Table 3, where WF-CGS represents our algorithm without the constraint-reuse strategy. Note that we include the run time for computing VCG prices for the MRC rules since it serves as the initial point in the CCG algorithm [Day and Raghavan, 2007]. As we can see, the WF-CGS-CR algorithm performs the best among all the core-selecting algorithms, which is even close to the VCG algorithm. Moreover, the constraint-reuse strategy is effective in our experiments, especially in the distribution of Paths.

7 Conclusion

In this paper, we propose the BLO payment rule to address the unfair issue with the MRC rule. The BLO payment rule adopts the leximin principle to select the core outcome, prioritizing the worst-off bidders. Furthermore, we theoretically analyze the BLO rule and propose the WF-CGS-CR algorithm to compute the exact BLO outcome. Our algorithm has the quadratic oracle complexity and achieves the best run time in the experiments, even close to VCG.

Compared with the MRC rule, the BLO rule has the following advantages: (1) the BLO outcome is unique, avoiding the problem of core outcome selection; (2) the BLO outcome achieves leximin fairness and thus avoids the unfair phenomenon caused by the MRC rule; (3) given the winner determination oracle access, the BLO outcome is computed effectively in time polynomial in the winner size. Therefore, the BLO rule could serve as an alternative payment rule for the core-selecting CAs.

Acknowledgements

References

[Arin and Inarra, 2001] Javier Arin and Elena Inarra. Egalitarian solutions in the core. *International Journal of Game Theory*, 30(2):187–193, 2001.

[Ausubel and Baranov, 2017] Lawrence M. Ausubel and Oleg Baranov. A practical guide to the combinatorial clock auction. *The Economic Journal*, 127(605):334–350, 2017.

[Ausubel et al., 2006] Lawrence M Ausubel, Paul Milgrom, et al. The lovely but lonely vickrey auction. *Combinatorial Auctions*, 17:22–26, 2006.

[Bünz et al., 2015] Benedikt Bünz, Sven Seuken, and Benjamin Lubin. A faster core constraint generation algorithm for combinatorial auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 827–834, 2015.

[Bünz et al., 2018] Benedikt Bünz, Benjamin Lubin, and Sven Seuken. Designing core-selecting payment rules: A computational search approach. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 109–109, 2018.

[Clarke, 1971] Edward H Clarke. Multipart pricing of public goods. *Public Choice*, pages 17–33, 1971.

[Cramton, 2013] Peter Cramton. Spectrum auction design. *Review of Industrial Organization*, 42(2):161–190, 2013.

[Day and Cramton, 2012] Robert W Day and Peter Cramton. Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research*, 60(3):588–603, 2012.

[Day and Milgrom, 2008] Robert Day and Paul Milgrom. Core-selecting package auctions. *International Journal of Game Theory*, 36(3):393–407, 2008.

[Day and Raghavan, 2007] Robert W Day and Subramanian Raghavan. Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science*, 53(9):1389–1406, 2007.

[Erdil and Klemperer, 2010] Aytek Erdil and Paul Klemperer. A new payment rule for core-selecting package auctions. *Journal of the European Economic Association*, 8(2-3):537–547, 2010.

[Gillies, 1959] Donald B Gillies. Solutions to general non-zero-sum games. *Contributions to the Theory of Games*, 4:47–85, 1959.

[Goel et al., 2015] Gagan Goel, Mohammad Reza Khani, and Renato Paes Leme. Core-competitive auctions. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 149–166, 2015.

[Goeree and Lien, 2016] Jacob K Goeree and Yuanchuan Lien. On the impossibility of core-selecting auctions. *Theoretical Economics*, 11(1):41–52, 2016.

[Grötschel et al., 1981] Martin Grötschel, László Lovász, and Alexander Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.

[Groves, 1973] Theodore Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.

[Leyton-Brown et al., 2000] Kevin Leyton-Brown, Mark Pearson, and Yoav Shoham. Towards a universal test suite for combinatorial auction algorithms. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pages 66–76, 2000.

[Leyton-Brown et al., 2017] Kevin Leyton-Brown, Paul Milgrom, and Ilya Segal. Economics and computer science of a radio spectrum reallocation. *Proceedings of the National Academy of Sciences*, 114(28):7202–7209, 2017.

[Lubin et al., 2015] Benjamin Lubin, Benedikt Bünz, and Sven Seuken. New core-selecting payment rules with better fairness and incentive properties. In *The Third Conference on Auctions, Market Mechanisms and Their Applications*. ACM, 2015.

[Manual, 1987] CPLEX User's Manual. Ibm ilog cplex optimization studio. *Version*, 12(1987-2018):1, 1987.

[Niazadeh et al., 2022] Rad Niazadeh, Jason Hartline, Nicole Immorlica, Mohammad Reza Khani, and Brendan Lucier. Fast core pricing for rich advertising auctions. *Operations Research*, 70(1):223–240, 2022.

[Rawls, 2004] John Rawls. A theory of justice. In *Ethics*, pages 229–234. Routledge, 2004.

[Rothkopf et al., 1998] Michael H Rothkopf, Aleksandar Pekeč, and Ronald M Harstad. Computationally manageable combinational auctions. *Management science*, 44(8):1131–1147, 1998.

[Sandholm, 2007] Tuomas Sandholm. Expressive commerce and its application to sourcing: How we conducted $35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–45, 2007.

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