Rank-Guaranteed Auctions*

Wei He† Jiangtao Li‡ Weijie Zhong§

May 1, 2024

Abstract

We design a multi-item ascending auction that is "approximately" optimal, strategically simple, and robust to strategic and distributional uncertainties. Specifically, the auction is *rank-guaranteed*—ex-post revenue exceeds the maximal sum of the k^{th}-highest bidder's values when bidders play non-obviously dominated strategies. Moreover, under distributional uncertainty of valuations, the rank guarantee is asymptotically robustly optimal—it differs from the worst-case total surplus by at most $O(\frac{1}{N})$.

1 Introduction

The online advertising sector is a dynamic engine of economic activity, generating hundreds of billions of dollars annually through a simple mechanism: the auctioning of advertising "slots". Despite the critical role the auctions play, there's a surprising lack of theoretical groundwork to navigate the intricacies of auction design. This gap in knowledge stems from a unique challenge: "*not all slots are created equal*"—bidders often have complex, combinatorial preferences for different slots. For instance, YouTube intersperses promotional videos at regular intervals within longer content. Here, some advertisers might see value in the repetition of their ads, leveraging the complementarity, while others may fear overexposure could lead to negative perceptions akin to "spamming." Similarly, in the layout of Facebook's Marketplace, which presents content in two columns, preferences can vary widely. Some advertisers might seek to blend seamlessly with organic content by choosing single slots, whereas others might opt for an entire row to ensure their message is fully conveyed. Current methodologies oscillate between bespoke solutions designed for very specific preference frameworks—like the Generalized Second Price Auction (GSP) employed by Yahoo & Google, suited for singular demand

*This manuscript subsumes He et al. (2022).

†Department of Economics, The Chinese University of Hong Kong, hewei@cuhk.edu.hk
‡School of Economics, Singapore Management University, jtli@smu.edu.sg
§Graduate School of Business, Stanford University, weijie.zhong@stanford.edu

and vertically differentiated slots—and more universal models not primarily aimed at revenue optimization, such as the Vickrey-Clarke-Groves auction (VCG) employed by Meta.

This reality starkly contrasts with traditional auction theory, which often simplifies to an extent that overlooks the nuanced realities of the markets. Iconic theories, such as Myerson (1981), typically focus on the allocation of a *single item*, assuming *independent valuations* among participants, along with both the auctioneer and bidders possessing fully *Bayesian rationality*, underpinned by accurate, *shared prior beliefs*. However, none of these assumptions hold water in the complex scenarios described earlier, highlighting a clear disconnect. The development of a comprehensive theory for maximizing revenue through auction design in these more intricate and generalized environments remains a significant, unmet challenge.

In this study, we address the quadrilemma in multi-item auction design, achieving a *near-optimal* resolution: it is possible to simultaneously attain four pivotal objectives—optimal revenue, strategic simplicity, strategic robustness, and distributional robustness—through a simple auction mechanism, given certain approximations. Our focus is on scenarios where an auctioneer sells multiple items to several (potentially a large number of) strategic bidders, each with private valuations. We introduce a multi-item variant of the open ascending auction termed the (C)ombinatorial (As)sending (A)uction (CASA). Prior to the auction, the auctioneer curates a *menu* of item bundles for allocation. With the formal game theoretic form of the auction described in Section 2, this auction model distills down to two straightforward principles:

1. Bidders are allowed to place binding bids (increase prices) on any assortment of bundles from the menu, even if these selections overlap.
2. The auction concludes when bid prices stabilize, with the winning bids being those that maximize the total selling price.

Our findings reveal that CASA meets the quadrilemma's four criteria within certain approximate bounds:

- **Strategic Simplicity**: It is *obviously dominating* to bid up to the true value as long as there is a remaining surplus for a bidder given the current prices.
- **Optimal Revenue**: Any non-obviously dominated strategy profile yields an *ex-post* revenue that is *rank-guaranteed* — achieving the maximal revenue when each bundle within menu can be sold at the k^{th}-highest value among all bidders.
- **Distributional robustness**: Neither the auction format nor the revenue guarantee depends on a Bayesian prior on either the auctioneer's side or the bidders' side.

- **Strategic robustness**: CASA is rank-guaranteed even with irrational or collusive bidders.

Our second result is to quantify the k^{th}-guarantee using canonical robust optimality criteria. We took a worst-case analysis against distributional uncertainties, i.e. an adversarial nature choosing the joint distribution of values to minimize expected revenue against the mechanism (Carroll (2017)). We show that in the worst case, under a menu $\mathcal{M}$, the expected k^{th}-guarantee approximates the total surplus (the 1^{st}-guarantee) at the rate of $O\left(\frac{k|\mathcal{M}|}{N}\right)$, i.e. the k^{th}-guarantee asymptotically achieves full surplus extraction when the number of bidders is large relative to the menu size. We prove this by developing a novel statistical result that bounds the k^{th} largest order statistic of a given sample using a random element.

The k^{th}-guarantee we derive reveals a novel trade-off between *menu sufficiency* and *approximation efficiency*: a more complete menu achieves a higher benchmark total surplus but increases k and $|\mathcal{M}|$. Therefore, to close the approximation gap towards the various goals, the key exercise is to reduce the menu size while maintaining the allocation efficiency, leveraging further knowledge about the bidders' preferences. We focus on a specific type of *sufficient* menus that improves approximation efficiency "for free"—menus that achieve the same worst-case total surplus as the complete menu. Specifically, we show that when the bidder's preference exhibits canonical preference structures, without loss of the benchmark total surplus, the size of menus can be reduced to be polynomial in the number of items being auctioned and so is the convergence rate of revenue guarantee. The result is summarized in Table 1.

Preference	**Sufficient Menu**	k
Weak substitutability	Individual items	$O(M)$
Weak complementarity	Grand bundle	2
Partitional complementarity (with I different possible partitions)	Partitional bundles	$O(M)$ $O(M \times I)$
Homogenous goods (with I heterogenous types)	Menu of quantities	$O(M^2)$ $O(M^I)$

Table 1: Sufficient menus

1.1 Related literature

(Approximately) optimal auction design Beyond the simple environment studied in Myerson (1981) and Bulow and Klemperer (1996), solving for the exact optimal mechanism with confounding factors like multiple heterogeneous items, bounded distributional knowledge or bounded rationality is generally intractable. Various alternative optimality notions have been proposed to make progress (see surveys by Roughgarden (2015) and Hartline (2013)). Aggarwal and Hartline (2006) and Goldberg and Hartline (2001) obtained the "constant fraction" approximation in the auction of sponsored search and digital goods. Following a broader literature on robust mechanism design pioneered by Carroll (2017), authors studied "robustly" optimal auctions that maximize the distributional worst-case revenue. He and Li (2022) solved for the robustly optimal reserve price in SPA, Zhang (2021) solved for the robustly optimal dominance strategy mechanism for selling a single item. Bergemann et al. (2017) and Brooks and Du (2021) study auctions that maximize the informational worst-case revenue. To our knowledge, the rank-guarantee approximation is new, and has the unique feature of being more "adaptive" than alternative notions: the expected rank-guarantee varies with the underlying value distribution while the worst-case guarantee/maximal regret does not; the rank-guarantee increases with the population of bidders while the constant fraction guarantee does not. In Section 3, we apply the distributional robustness analysis similar to that of Carroll (2017) and show that the rank-guarantee has an appealing worst-case performance.

Multi-item auctions Beyond the efficient Vickery auction, little clean results have been established regrading multi-item auctions with combinatorial preferences. Jehiel and Moldovanu (2001a) point out the vulnerability of efficiency under multidimensional bidder information. Ausubel and Milgrom (2002) point out the poor revenue performance and strategic vulnerability of the Vickery auction and propose simultaneous ascending auctions with package bidding (SAAPB). The multi-item auction design problem has also been extensively studied in the field of combinatorial auctions (see Cramton et al. (2006) for a survey). This literature mainly focuses on (approximately) efficient auction design and their computational complexity, which is orthogonal to our focus on revenue performance and strategic simplicity. CASA is a simpler variant of the SAAPB and the *Combinatorial Clock Acution* (CAA, see Ausubel et al. (2006) and Levin and Skrzypacz (2016)) in that bidders simply raise the prices of the bundles, as opposed to personalized prices in SAAPB and demand reporting in CCA. Our result overturned the seminal impossibility theorem of Roughgarden (2014) due to a different notion of simplicity: in the dynamic CASA, the *strategic form* game violates the simplicity requirement of Roughgarden (2014); meanwhile, it is straightforward to identify obviously optimal *behavioral* strategies in the *extensive form* game.

Implementation Under the extensive form of our auction format, we study outcomes that survive the elimination of obviously dominated strategies. This approach is a hybrid of the undominated-strategy implementation (Carroll (2014), Yamashita (2015), Jackson (1992), Börgers (1991)) and obviously strategy-proof implementation (Li (2017)).

The rest of the paper is organized as follows. In Section 2, we introduce the k^{th}-guarantee and the CASA auction. In Section 3, we bound the worst-case performance of k^{th}-guaranteed auctions under distributional uncertainties. In Section 4, we explore specific preference structures under which CASA with simple menus performs as well as the complete menu.

2 CASA and rank-gurantee

2.1 The environment

There is a set S of M items to be sold to N bidders. Each $b \subseteq S$ denotes a bundle of items. For any bidder n, let $\boldsymbol{v}^n = \{v_b^n\}_{b\subseteq S}$ denote the valuation vector of bidder n. We assume that $\forall b \neq \emptyset$, $v_b^n \in [\underline{v}, \bar{v}]$ with $\underline{v} \geq 0$ and normalize $v_\emptyset^n = 0$. Let $\boldsymbol{v} = (v_1, v_2, \ldots, v_N)$ denote the entire valuation profile. $\mathcal{M} \subset 2^S$ is a *menu* of (non-empty) bundles chosen by the auctioneer. While $\mathcal{M}$ is a choice variable of the auctioneer, for now we take it as exogenously given and defer the discussion of menu choice to Section 4. For our analysis, we assume $N \geq |\mathcal{M}| + 1$. Let

$$\mathcal{B}(\mathcal{M}) = \{X \subset \mathcal{M} | \forall b, b' \in X,\ b \cap b' = \emptyset\}$$

denote the set of *feasible* allocations of bundles within the menu, i.e. all collections consisting of non-overlapping bundles.

2.2 The auction format

We now define the extensive form of the *Combinatorial Ascending Auction* (CASA): Let $P \subset \mathbb{R}^+$ be a finite grid of feasible bids with grid size ϵ and $\max P > \bar{v}$.

(1) **Initialization**: $t = 0$. The *leading price* vector $\boldsymbol{p}^0 = (p_b^0)_{b\in\mathcal{M}} = \mathbf{0}$. The *leading bidder* vector $\boldsymbol{\phi}^0 = (\phi_b^0)_{b\in\mathcal{M}} = \mathbf{0}$. The *active bidder* set $\mathcal{N}^0 = \{1, \ldots, N\}$. Then, move on to period $t = 1$ bidding stage.

(2) **Bidding stage** t: An active bidder $n \in \mathcal{N}^{t-1}$ observes $(\boldsymbol{p}^{t-1}, \{b | \phi_b^{t-1} = n\})$ and chooses a bid $\{(b, p_b)\} \subset (\mathcal{M} \times P)$, subject to the constraint that[1]

- Leading bids are binding: if $\phi_b^{t-1} = n$, then the bid must includes b and $p_b \geq p_b^{t-1}$.
- Mnimum bid increment: if the bid includes b s.t. $\phi_b^{t-1} \neq n$, then $p_b > p_b^{t-1}$.

The auction states are updated as follows:

- Following a null bid $\emptyset$, $\boldsymbol{p}^t = \boldsymbol{p}^{t-1}$, $\boldsymbol{\phi}^t = \boldsymbol{\phi}^{t-1}$, $\mathcal{N}^t = \mathcal{N}^{-1} \setminus \{n\}$.

[1] The bidder selection rule and the observability of history is inconsequential for our analysis. For concreteness, consider the selection rule that $\mathcal{N}$ is cycled in increasing order. The observability of history minimized while being sufficient for determining a valid bid in order to maximally protect privacy.

- Following an active bid, $\phi_b^t = n$, $p_b^t = p_b$ for b included in the bid. $\phi_{b'}^t = \phi_{b'}^{t-1}$, $p_{b'}^t = p_{b'}^{t-1}$ for b' not included in the bid. $\mathcal{N}^t = \mathcal{N}^{t-1}$.

Then, move on to period $t+1$ bidding stage.

(3) **Resolution**: the auction ends in period $t+1$ when prices stay constant for $N+1$ consecutive periods. The auctioneer chooses a feasible allocation to maximize

$$\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})} \sum_{b\in\boldsymbol{b}} p_b^t.$$

Denote the maximizer by $\boldsymbol{b}^*$. Each bundle $b \in \boldsymbol{b}^*$ is allocated to ϕ_b^t, and charged a price p_b^t.

In words, CASA runs parallel ascending auctions for each bundle b. Then, the items are allocated to maximize the total price. Compared to other proposals like SAAPB (Ausubel and Milgrom (2002)) and CCA (Ausubel et al. (2006)), CASA shares the same open ascending feature but has a more straightforward rule of "pay-as-bid". CASA has the following properties:

2.3 The rank-guarantee of CASA

At any history, consider a non-leading bidder's choice between "to bid" or "to quit". Then, quitting the auction leads to a best outcome of zero payoff, while continuing guarantees a non-negative payoff if some leading prices are still below the bidder's valuations. Hence, it is "obviously optimal" to bid as opposed to quit. Analogous to the notion of obvious dominant strategy from Li (2017), we define a notion of obviously dominated strategies to derive our solution concept. Let $h = (t, (\mathcal{N}^0, \ldots, \mathcal{N}^{t-1}), (\boldsymbol{p}^0 \ldots, \boldsymbol{p}^{t-1}), (\boldsymbol{\phi}^0, \ldots, \boldsymbol{\phi}^{t-1}))$ denote a history of the game and H denote the set of histories. Let $I_n(\boldsymbol{p}, \boldsymbol{b}) = \{h \in H | \boldsymbol{p}^t = \boldsymbol{p}, \{b|\phi_b^{t-1} = n\} = \boldsymbol{b}\}$ denote bidder n's information set given observed prices $\boldsymbol{p}$ and n's leading bundles $\boldsymbol{b}$. Let $\mathcal{I}_n$ denote all information sets of n. Let $s_n : \mathcal{I}_n \to 2^{\mathcal{M}\times\mathbb{R}^+}$ denote bidder n's (pure behavioral) strategy and $u_n(\boldsymbol{s}, \boldsymbol{v}|h)$ denote the deterministic payoff to player n given value profile $\boldsymbol{v}$, strategy profile $\boldsymbol{s}$, conditional on the current history h and n bidding in period t.

Definition 1. *A bidding strategy $s_n : H \to 2^{\mathcal{M}\times P}$ is **obviously dominated** if there exists history h and s_n' that differs from s_n at the history s.t.*

$$\sup_{v_{-n}, s_{-n}, h\in I_n} u_n(\boldsymbol{s}, \boldsymbol{v}|h) \leq \inf_{v_{-n}, s_{-n}, h\in I_n} u_n(s_n', s_{-n}, \boldsymbol{v}|h);$$
$$\inf_{v_{-n}, s_{-n}, h\in I_n} u_n(\boldsymbol{s}, \boldsymbol{v}|h) < \sup_{v_{-n}, s_{-n}, h\in I_n} u_n(s_n', s_{-n}, \boldsymbol{v}|h).$$

The first inequality is identical to the definition of the obvious dominance relation in Li (2017), i.e. the best outcome under s_n is weakly worse than the worst outcome under s'_n. In addition, we require the dominated strategy to be non-equivalent in terms of the induced outcome to the dominating strategy. The second requirement guarantees that the set of non-obviously dominated strategies is non-empty. The earlier intuition then translates to:

Lemma 1. *If $\exists\ I_n \in \mathcal{I}_n$ s.t. $s_n(I_n) = \emptyset$ and*

- $\exists \boldsymbol{p}' \in P$ *and* $b \in \arg\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})} \sum_{b'\in\boldsymbol{b}} p'_{b'}$ *s.t.* $\boldsymbol{p}' \geq \boldsymbol{p}^t$ *and* $v_b^n > p'_b > p_b^t$,

then s_n is obviously dominated.

Proof. The payoff from s_n conditional on any $h \in I_n$ is always zero. We first discuss a special cases: $|\mathcal{N}^{t-1}| = 1$ and n is not a current leading bidder (otherwise $\emptyset$ is not a feasible bid). Then, the current prices must be all 0 and all other bidders have quit (and this is the unique consistent history). The auction ends with bid $s'_n(h) = (b, p'_b)$, leading to $u_n(\boldsymbol{s}', \boldsymbol{v}|h) > 0$.

Suppose $|\mathcal{N}^{t-1}| > 1$. Consider $s'_n(I_n) = (b, p'_b)$ and any other active bidder bidding to $p'_{b'}$ when it is his turn. Then, the auction ends with prices $\boldsymbol{p}'$, leading to $u_n(\boldsymbol{s}', \boldsymbol{v}|h) = v_b^n - p'_b > 0$, for all $h \in I_n$. *Q.E.D.*

Note that Lemma 1 suggests that our intuition is incomplete as we can not fully rule out strategies that quit when value is above some leading prices. To rule out such strategies, it further requires the existence of some scenario under which bidder n may be pivotal: a strictly profitable bid of n may ever be selected in the end. Let $S^n_{NOD}(P)$ denote the set of non-obviously dominated strategies of n given price grid P. Let $R(\boldsymbol{s}, \boldsymbol{v})$ denote the revenue to seller given value profile $\boldsymbol{v}$ and strategy profile $\boldsymbol{s}$. Define

$$\underline{R}_{CASA}(\boldsymbol{v}) := \varliminf_{\epsilon\to 0} \inf_{s_n \in S^n_{NOD}(P)} R(\boldsymbol{s}, \boldsymbol{v}).$$

$\underline{R}_{CASA}(\boldsymbol{v})$ is the worst-case ex-post revenue from CASA under non-obviously dominated strategies in the limit where grid P becomes dense.

Given the valuation profile $\boldsymbol{v}$ and menu $\mathcal{M}$, the k^{th}-guarantee is defined as the maximal revenue from feasible allocations within $\mathcal{M}$, taking the k^{th}-highest valuations for each bundle as the price.

Definition 2. *The k^{th}-guarantee given $\boldsymbol{v}$ and $\mathcal{M}$ is*

$$R^k_{\mathcal{M}}(\boldsymbol{v}) := \max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})} \sum_{b\in\boldsymbol{b}} v_b^{(k)},$$

where $v_b^{(k)}$ denotes the k^{th}-highest value of bundle b.

The key observation is that $\underline{R}_{CASA}(\boldsymbol{v})$ achieves the k^{th}-guarantee.

Theorem 1. *$\underline{R}_{CASA}(\boldsymbol{v}) \geq R^k_{\mathcal{M}}(\boldsymbol{v})$ for $k = |\mathcal{M}| + 1$.*

Proof. We claim that $\forall P$ with size ϵ,

$$\max_{\boldsymbol{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \boldsymbol{b}} p^t_b \geq \max_{\boldsymbol{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \boldsymbol{b}} (v^{(k)}_b - \epsilon).$$

Suppose for the purpose of contradiction that the statement is not true. Then, there exists strategy profile s, value profile $\boldsymbol{v}$ and $\delta > 0$ s.t. let $(t, \boldsymbol{p}^t, \boldsymbol{\phi}^t)$ be the outcome,

$$\max_{\boldsymbol{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \boldsymbol{b}} p^t_b < \max_{\boldsymbol{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \boldsymbol{b}} (v^{(k)}_b - \epsilon - \delta).$$

Evidently, $\exists b$ s.t. $p^t_b < v^{(k)}_b - \epsilon - \delta$. Raise each p^t_b to (the closest price below) $v^{(k)}_b - \delta$ until we find the first pivotal b, i.e. $\max_{\boldsymbol{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \boldsymbol{b}} p^t_b$ is strictly improved by raising p^t_b to $v^{(k)}_b - \delta$. Since $k = |\mathcal{M}| + 1$, there exists a player n that quits the auction before period t and $v^n_b \geq v^{(k)}_b$. Let h be the (on-path) history at which n quits (note that at this history, n must not be a leading bidder and $p_b \leq p^t_b \leq v^n_b - \delta - \epsilon$). Let $h \in I_n$. Then, $s_n(I_n) = \emptyset$. Let $\boldsymbol{p}'$ be the raised prices, then $\boldsymbol{p}'$ satisfies the condition in Lemma 1. Therefore, s_n is obviously dominated. *Q.E.D.*

As we have pointed out following Lemma 1, elimination of obviously dominated strategies do *not* guarantee the ex-post prices to be above the k^{th}-highest values. To establish Theorem 1, we prove in addition that the behaviors of non-pivotal bidders are inconsequential. With Theorem 1, we say that CASA is a rank-guaranteed auction format (under non-obviously dominated strategies). When we say another auction format is rank-guaranteed, the solution concept should also be specified. For instance, the SPA is second-guaranteed under dominant strategy equilibrium. The third-priced auction is third-guaranteed under undominated strategies.[2]

2.4 Discussions

Strategic simplicity: The two key features of CASA, the ascending auction format and the non-exclusivity rule made it strategically simple to decide "how to bid"—bid up as long as there is remaining surplus. Both features are necessary. The ascending format makes the "prices" transparent so that there is no uncertainty about prices. The non-exclusive bidding rule makes the allocation transparent.

However, the decision of "what to bid" remains strategically hard for the bidders. While the bidder can fully avoid the "exposure problem" by bidding only on the bundles

[2] Observe that in the third-priced auction, bidding below value is dominated by bidding the value.

he wants, he might strategically chooses to expose himself so that to manipulate the prices of the complementary bundles.

Robustness to collusion / irrationality: Since the argument of Theorem 1 solely relies on the existence of one strategic bidder with a value v_b^n above $v_b^{(k)}$ who quits the auction. The analysis extends to the case with non-strategic bidders easily. Firstly, when the number of non-strategic bidders is bounded, then Theorem 1 still holds when k is relaxed by the number of non-strategic bidders.

Proposition 1. *Suppose there are j non-strategic bidders, then $\underline{R}_{CASA}(\boldsymbol{v}) \geq R^k_{\mathcal{M}}(\boldsymbol{v})$ for $k = |\mathcal{M}| + 1 + j$.*

Proof. Observe that in the proof of Theorem 1, since $k = |\mathcal{M}| + 1 + j$, there exists at least one *strategic* player n that quits the auction before period t and $v_b^n \geq v_b^{(k)}$. The rest of the proof follows. *Q.E.D.*

Secondly, when the bidders form coalitions and they strategically maximize group payoffs[3], then Theorem 1 still holds when k is scaled by the coalition sizes.

Proposition 2. *Suppose bidders are partitioned into strategic coalitions $\{c_i\}_{i\in I}$, where the index is chosen such that $|c_i|$ decreases in i. Then, $\underline{R}_{CASA}(\boldsymbol{v}) \geq R^k_{\mathcal{M}}(\boldsymbol{v})$ for $k = \sum_{i\leq|\mathcal{M}|} |c_i| + 1$.*

Proof. Observe that in the proof of Theorem 1, since $k = \sum_{i\leq|\mathcal{M}|} |c_i| + 1$, there exists at least one coalition of players c that all quit the auction before period t and $\max_{n\in c}\{v_b^n\} \geq v_b^{(k)}$. Let h be the (on-path) history at which the last member n in the coalition quits (note that at this history, n must not be a leading bidder and $p_b \leq p_b^t \leq \max_n\{v_b^n\} - \delta - \epsilon$). Let $h \in I_n$. Then, $s_n(I_n) = \emptyset$. Obviously, quitting gives the entire group 0 payoff while bidding p'_b guarantees a non-negative payoff. Suppose all other bidders bid to p'_b when it is their turn, the auction ends with $\boldsymbol{p}'$ and the group obtains a strictly positive payoff. Therefore, s_n is obviously dominated for coalition c. *Q.E.D.*

The intuition behind the two extensions is exactly the strategic simplicity of "how to bid". The price of each bundle must be higher than the value of any *losing strategic* bidder or any *losing coalition group* as otherwise they will outbid the price. Of course, Proposition 2 has no bite when k is large compared to N; hence, it should be interpreted as the strategic robustness of CASA only in relatively thick markets. Nevertheless, CASA is also aligned with the philosophy of anti-collusion design even in thin markets (Klemperer (2002)), for the reason that CASA permits minimum transmission of information.[4]

[3] Each group can freely shift allocations within the group and maximize the total payoff.

[4] Anonymity prevents the reciprocity behavior documented in Cramton and Schwartz (2000). Theorem 1 holds when each bidder can only submit bid on one bundle at a time, which prevents the strategic communication documented in Jehiel and Moldovanu (2001b) and Grimm et al. (2003).

Efficiency: Note that it is *not* obviously dominated to bid strictly above the true valuation for a bundle.[5] Therefore, CASA might not satisfy ex-post IR; hence k^{th}-guaranteed revenue does not implie k^{th}-guaranteed surplus. Of course, CASA still satisfies ex-ante IR (assuming bidders having correct Bayesian priors) since quitting is always an option; hence, the ex-ante bounds we derive in this paper on the revenue of CASA also apply to surplus.

VCG & SAAPB: The celebrated pivot VCG mechanism is known to underperform the k^{th}-guarantee when bidder's preferences exhibits complementarity.[6] The SAAPB mechanism is indeed k^{th}-guaranteed under "straightforward bidding" strategies (Theorem 1 of Ausubel and Milgrom (2002)). However, whether SAAPB is k^{th}-guaranteed with fully strategic bidders is yet unknown to us. Nevertheless, our result suggests that the spirit of SAAPB, when carefully executed, leads to an appealing revenue guarantee.

3 Rank-guarantee in the worst case

Evidently, if bidders are i.i.d., then the k^{th}-guarantee is a quite appealing approximation when N is large, as all (fixed) order statistics converge to the upper bound of valuation support quickly.[7] In this section, we show that even in an adversarial scenario when bidders can be arbitrarily correlated and the joint distribution is chosen to minimize the revenue, the performance of the k^{th}-guaranteed auctions is still quite appealing.

Let $\mathbb{G} \subset \Delta([\underline{v}, \bar{v}]^{2^S})$ be an arbitrary subset of distributions of valuations for all bundles. We interpret $\mathbb{G}$ as the auctioneer's estimation for a *representative bidder*'s valuation. Then, the joint distributions among all bidders' valuations that are considered possible by the auctioneer are

$$\mathbb{F} = \left\{ F \in \Delta([\underline{v}, \bar{v}]^{N \times 2^S}) \middle| \frac{1}{N} \sum F_n \in \mathbb{G} \right\},$$

where F_n is the marginal distribution of bidder n's valuation. $\frac{1}{N} \sum F_n$ is the CDF of the valuation of a uniformly randomly selected bidder in the population. We call $\mathbb{F}$ an *ambiguity set*. Such ambiguity set $\mathbb{F}$ could come from the statistical estimation of F based on a "sanitized" dataset about valuations i.e. past bidder valuations with identity information removed. The ambiguity set $\mathbb{F}$ captures the type of distributional uncertainty introduced by Carroll (2017), while further generalizing it to capture realistic knowledge

[5] Consider, for example, the case with three items a, b, c and three bidders $1, 2, 3$. Bidder 1 and 2 only wants a and b and bidder 3 only wants bundle $\{a, b, c\}$, respectively, all with valuation 1. By strategically bidding up item b even though bidder 1 gets zero value from it, bidder 1 can reduce the bid required for him to win item a, creating an exposure problem for 1.

[6] Imagine the case $\mathcal{M} = \{(a), (b), (a, b)\}$. $v_{(a,b)} = 1$ for all bidders. $v_a = v_b = 0$ for all bidders except for two, whose value for a and b are 1. The VCG revenue is 0, while the k^{th}-guarantee is 1 for any $k \geq 2$.

[7] At the rate of $\frac{1}{N}$.

structures stemming from statistical inference.[8]

$\forall F \in \Delta([\underline{v}, \bar{v}]^{N\times 2^S})$, define the ex-ante *efficient* surplus with menu $\mathcal{M}$

$$V_{\mathcal{M}}(F) := \mathbb{E}_F\left[\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})}\max_{\iota:\boldsymbol{b}\leftrightarrow\mathcal{N}}\sum_{b\in\boldsymbol{b}} v_b^{\iota(b)}\right].$$

The ex-ante efficient surplus with the complete menu $V_{2^S}(F)$ is denoted by $V^*(F)$.

Theorem 2. *$\forall\mathcal{M}$,*

$$\inf_{F\in\mathbb{F}}\mathbb{E}_F[R^k_{\mathcal{M}}(\boldsymbol{v})] \geq \inf_{F\in\mathbb{F}} V_{\mathcal{M}}(F) - \frac{k|\mathcal{M}|\bar{v}}{N}.$$

Proof. Let Ⓡ be a uniform random element of $\mathcal{N} = \{1,\dots,N\}$.

$$\begin{aligned}
R^k_{\mathcal{M}}(\boldsymbol{v}) &\geq \max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})}\sum_{b\in\boldsymbol{b}} v_b^{(k)} \\
&\geq \max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})}\sum_{b\in\boldsymbol{b}} v_b^{Ⓡ} - \sum_b (v_b^{Ⓡ} - v_b^{(k)})^+ \\
\implies \mathbb{E}_F[R^k_{\mathcal{M}}(\boldsymbol{v})] &\geq \mathbb{E}_F\left[\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})}\sum_{b\in\boldsymbol{b}} v_b^{Ⓡ}\right] - \sum_{b\in\boldsymbol{b}}\mathbb{E}_F\left[v_b^{Ⓡ} - v_b^{(k)}, v_b^{Ⓡ} > v_b^{(k)}\right] \\
&\geq \mathbb{E}_F\left[\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})}\sum_{b\in\boldsymbol{b}} v_b^{Ⓡ}\right] - \sum_{b\in\boldsymbol{b}}\bar{v}\mathrm{Prob}(v_b^{Ⓡ} > v_b^{(k)}) \\
&= \mathbb{E}_F\left[\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})}\sum_{b\in\boldsymbol{b}} v_b^{Ⓡ}\right] - \frac{k|\mathcal{M}|\bar{v}}{N}.
\end{aligned}$$

Then,

$$\begin{aligned}
\inf_{F\in\mathbb{F}}\mathbb{E}_F\left[R^k_{\mathcal{M}}(\boldsymbol{v})\right] &\geq \inf_{F\in F}\mathbb{E}_F\left[\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})}\sum_{b\in\boldsymbol{b}} v_b^{Ⓡ}\right] - \frac{k|\mathcal{M}|\bar{v}}{N} \\
&\geq \inf_{G\in\mathbb{G}}\mathbb{E}_G\left[\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})}\sum_{b\in\boldsymbol{b}} v_b\right] - \frac{k|\mathcal{M}|\bar{v}}{N} \\
&\geq \inf_{F\in\mathbb{F}} V_{\mathcal{M}}(F) - \frac{k|\mathcal{M}|\bar{v}}{N}.
\end{aligned}$$

For the last inequality, observe that $\forall G \in \mathbb{G}$, the joint distribution that all bidders identically distributed according to G is contained in $\mathbb{F}$ (in which case $V_{\mathcal{M}}(F) \leq \mathbb{E}_G\left[\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})}\sum_{b\in\boldsymbol{b}} v_b\right]$). Therefore,$\inf_{F\in\mathbb{F}} V_{\mathcal{M}}(F) \leq \inf_{G\in\mathbb{G}}\mathbb{E}_G\left[\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})}\sum_{b\in\boldsymbol{b}} v_b\right]$

Q.E.D.

Theorem 2 illustrates a key trade-off between *menu sufficiency* and *approximation*

[8] Carroll (2017) studies a multi-item single-buyer problem, where the unquantifiable uncertainty is about the correlation between items. In our setting, we assume the unquantifiable uncertainty is about the correlation between bidders, while among items there might or might not be uncertainty as $\mathbb{G}$ is completely general.

efficiency. $\mathcal{M}$ can be naively chosen to be the entire 2^S to guarantee full menu sufficiency i.e. $V_{\mathcal{M}}(F) = V^*(F)$. However, this leads to $|\mathcal{M}|$ to grow exponentially in M, causing both complex auction process and slow convergence. Vice versa, choosing a small menu achieves approximation efficiency but sacrifices allocation efficiency. Although such trade-off is generally non-trivial under general combinatorial preferences, we show in the next section that under canonical preference structures, menu sufficiency and approximation efficiency can often be achieved simultaneously.

Tightness of the bound:

Proposition 3. $\forall M, N, \mathcal{M}, k$, *exists* $\mathbb{G}$ *s.t.*

$$\inf_{F\in\mathbb{F}} \mathbb{E}_F[R^k_{\mathcal{M}}(\boldsymbol{v})] \leq \inf_{F\in\mathbb{F}} V_{\mathcal{M}}(F) - O\left(\tfrac{k}{N}\right).$$

Proof. See Appendix A.1. *Q.E.D.*

The coefficienct $k|\mathcal{M}|$ in Theorem 2 consists of two parts. The first k comes from the k^{th} highest value approximation. The second $|\mathcal{M}|$ comes from the total number of $|\mathcal{M}|$ bundles. In Proposition 3, we show that the dependence on k is tight.

4 Simple and sufficient menus

In this section, we explore preference structures under which there exists menus that are *sufficient*—the menu leads to full allocation efficiency—and *small*—menus size grows at polynomial rate as M grows.

Definition 3. *Menu* $\mathcal{M}$ *is* $\mathbb{G}$*-sufficient if:*

$$\inf_{G\in\mathbb{G}} \mathbb{E}_G\left[\max_{\boldsymbol{b}\in\mathcal{B}(2^S)} \sum_{b\in\boldsymbol{b}} v_b\right] = \inf_{G\in\mathbb{G}} \mathbb{E}_G\left[\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})} \sum_{b\in\boldsymbol{b}} v_b\right].$$

A menu $\mathcal{M}$ is sufficient if the worst-case surplus from allocating to (hypothetically) identical bidders with valuation distribution from $\mathbb{G}$ is the same as that under the complete menu 2^S. Importantly, sufficiency is defined w.r.t. the preference of a single bidder instead of all bidders. It is much weaker than assuming that restricting to allocations within $\mathcal{M}$ is without loss for ex-post efficiency.[9] Nevertheless, sufficiency guarantees full allocation efficiency:

[9] Consider for instance two items and two bidders, each valuing the sum of individual items more than bundles. Then, menu of individual items is "sufficient" per Definition 3, but not necessarily ex-post efficient when the two bidder's values are highly asymmetric.

Theorem 3. *If menu $\mathcal{M}$ is $\mathbb{G}$-sufficient:*

$$\inf_{F\in\mathbb{F}} \mathbb{E}_F \left[R^k_{\mathcal{M}}(\boldsymbol{v})\right] \geq \inf_{F\in\mathbb{F}} V^*(F) - \frac{k|\mathcal{M}|\bar{v}}{N}.$$

Proof. When $\mathcal{M}$ is $\mathbb{G}$-sufficient,

$$\begin{aligned}
\inf_{F\in\mathbb{F}} \mathbb{E}_F \left[R^k_{\mathcal{M}}(\boldsymbol{v})\right] &\geq \inf_{G\in\mathbb{G}} \mathbb{E}_G \left[\max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})} \sum_{b\in\boldsymbol{b}} v_b\right] - \frac{k|\mathcal{M}|\bar{v}}{N} \\
&= \inf_{G\in\mathbb{G}} \mathbb{E}_G \left[\max_{\boldsymbol{b}\in\mathcal{B}(2^S)} \sum_{b\in\boldsymbol{b}} v_b\right] - \frac{k|\mathcal{M}|\bar{v}}{N} \\
&\geq \inf_{F\in\mathbb{F}} V^*(F) - \frac{k|\mathcal{M}|\bar{v}}{N}.
\end{aligned}$$

Q.E.D.

A simple sufficient condition for $\mathbb{G}$-sufficiency is that $\max_{\boldsymbol{b}\in\mathcal{B}(2^S)} \sum_{b\in\boldsymbol{b}} v_b = \max_{\boldsymbol{b}\in\mathcal{B}(\mathcal{M})} \sum_{b\in\boldsymbol{b}} v_b$ holds ex-post, which allows us to convert preference assumptions into sufficiency. To simplify notation, let $\text{Supp}(\mathbb{G}) := \cup_{G\in\mathbb{G}}\text{Supp}(G)$.

Weak substitutability and itemized ascending auction

Definition 4. *Bidder preference exhibits* ***weak substitutability*** *if* $\forall \boldsymbol{v} \in \text{Supp}(\mathbb{G})$, $\forall b \subset S$,

$$\sum_{s\in b} v_{\{s\}} \geq v_b.$$

In words, a representative bidder finds the value of any bundle lower than the total value of all items in the bundle. Weak substitutability is a necessary condition for various substitutability notions studied in the literature.

Proposition 4. *If bidder preference exhibits weak substitutability, menu $\mathcal{M} = S$ is $\mathbb{G}$-sufficient and $k = M + 1$.*

When $\mathcal{M} = S$, the CASA reduces to a simple itemized ascending auction, where each item is auctioned individually and exclusively. Weak substitutability is one of the most widely used preference assumptions in the literature as it captures a natural diminishing return to scale. Our analysis shows that under such preference structure, CASA can exhibit extreme simplicity while achieving both allocation and approximation efficiency. Intriguingly, under weak substitutability, the canonical Vickery auction performs as well as CASA, despite its much worse performance under more general preference structures.

Proposition 5. *If bidder preference exhibits weak substitutability, the Vickery auction with truthful bidding achieves a revenue guarantee of $R^{M+1}_S(\boldsymbol{v})$.*

Proof. See Appendix A.2. *Q.E.D.*

An even more special case of weak substitutability is the sponsored search auction, where valuations of items are constant (and common) ratios of a one-dimensional private type. As is proved in Edelman et al. (2007), the clock auction version of *generalized second price* (GSP) auction is outcome equivalent to the Vickery auction; hence achieving the same $M+1^{th}$-guarantee.

Weak complementarity and SPA

Definition 5. *Bidder preference exhibits* ***weak complementarity*** *if* $\forall \boldsymbol{v} \in \mathrm{Supp}(\mathbb{G})$, $\forall \boldsymbol{b} \in \mathcal{B}(2^S)$,

$$\sum_{b\in\boldsymbol{b}} v_b \leq v_S.$$

In words, a representative bidder finds the value of the grand bundle higher than the total value of any feasible collection of bundles. Weak complementarity is a necessary condition for various complementarity notions studied in the literature.

Proposition 6. *If bidder preference exhibits weak substitutability, menu* $\mathcal{M} = \{S\}$ *is* $\mathbb{G}$*-sufficient and* $k = 2$.

When $\mathcal{M} = \{S\}$, the CASA reduces to a simple ascending auction for only the grand bundle. Evidently, in this case, the standard *second-priced auction* (SPA) is second-guaranteed and outcome-equivalent to CASA.

"Partitional" complementarity A hybrid case of complementarity and substitutability is the complementarity described by a partition $\mathcal{K}$ of S.

Definition 6. *Let* $\mathcal{K}$ *be a partition of* S*. Bidder preference exhibits* $\mathcal{K}$***-partitioned complementarity*** *if* $\forall \boldsymbol{v} \in \mathrm{Supp}(\mathbb{G})$,

$$\begin{aligned}
&\forall b \in \mathcal{K}, \forall \textit{partition } \kappa \textit{ of } b,\ \sum_{b'\in\kappa} v'_b \leq v_b;\\
&\forall b' \subset S,\ \sum_{b\in\mathcal{K}} v_{b\cap b'} \geq v_{b'}.
\end{aligned}$$

Preference exhibits $\mathcal{K}$-partitioned complementarity structure means there is weak complementarity within each $b \in \mathcal{K}$ and weak substitutability across each $b \in \mathcal{K}$. In this case, $\mathcal{M} = \mathcal{K}$ is $\mathbb{G}$-sufficient and $k = |\mathcal{K}| + 1$.

The result can be easily extended to the case with multiple possible partitions $\{\mathcal{K}_i\}_{i=1}^{I}$, where I is bounded. In this case $\mathcal{M} = \cup_{i\in I}\mathcal{K}_i$ and $k \sim Poly(M)$. Such partitional preference structure arises when there is clear synergy between "nearby" bundles. Think about land auctions, for example. There are finitely many possible partitions that are determined by the major divisions of lands by rivers, highways, or railroads. If two distinct lands are segregated by those divisions, then there is substitutability among them. In such cases, our theory guarantees the performance of CASA with the partitional menu.

Homogeneous goods and quantity-CASA

Definition 7. *The goods are* ***homogeneous*** *if there exists* $u : \mathbb{N} \to [\underline{v}, \bar{v}]$ *s.t.* $\forall \boldsymbol{v} \in \mathrm{Supp}(\mathbb{G})$, $\forall b \in S$,

$$v_b = u(|b|).$$

With homogeneous goods, a representative bidder's valuation for any bundle only depends on the size of the bundle. Note that the dependence of u on $|b|$ is arbitrary. We do not even require monotonicity. In this case, we redefine the notion of feasible allocations to $\mathcal{B} : (\mathcal{M}) = \{X \subset \mathcal{M} | \sum_{b \in X} |b| \leq M\}$, i.e. an allocation is feasible as long as the total number of items is below M.

Proposition 7. *If goods are homogeneous, menu* $\mathcal{M} = \cup_{l \in \{1,\dots,M\}} \{b_l^1, \dots, b_l^{\lfloor \frac{M}{l} \rfloor}\}$*, where* b_l^j*'s are distinct bundles of size* l *is* $\mathbb{G}$*-sufficient and* $k \leq \frac{M^2+M}{2}$*.*

In this case, CASA simply auctions $\lfloor \frac{M}{l} \rfloor$ copies of each quantity level $l \leq M$ via individual ascending auctions. Like the discussion in partitional complementarity, there may finitely many types of homogeneous goods. As long as the number of types I is bounded, the menu consists of all combinations of $\lfloor \frac{M}{l} \rfloor$ copies of each type is sufficient and of size $Poly(M)$. Such preference structure is typical in examples like the spectrum auctions. Different frequencies are almost physically homogeneous, except that "middle" frequencies might be of different value from "boundary" frequencies.

5 Concluding remarks

In this paper, we design an auction format (CASA) that guarantees an approximately optimal ex-post revenue. In addition, we show that CASA is strategically simple for bidders and robust to distributional and strategic uncertainties under certain approximations. In practice, these approximation gaps may become non-negligible, rendering the deployment of CASA challenging.

- *Menu size:* the revenue performance of CASA as well as its strategic robustness crucially hinges on the rank k (menu size) being small relative to the number of bidders. In the online advertising examples we introduce, the complete menu is small enough that a handful of bidders may be sufficient to make CASA an appealing design. However, other interesting auctions may suffer the large menu problem (e.g. the land auctions) or the thin market problem (e.g. the route auctions of rideshare apps) or both (e.g. the spectrum auctions), rendering the guarantee underpowered.

In the latter cases, the menu sufficiency-approximation efficiency tradeoff becomes eminent. Our theory suggests the importance of preference estimation in those settings. Finding a simple sufficient menu still keeps the revenue guarantee appealing and CASA directly applicable. Even in settings like the spectrum auction where our theory has no bite, menu design may still be a cost-effective way to promote competition and improve the revenue performance of existing auctions.

- *Proxy bidding:* Since CASA makes "how to bid" but not "what to bid" straightforward, a full proxy-bidding version of CASA is still unknown to us, making deploying CASA in settings that require fast resolution of auctions challenging. We would like to argue that this can be mitigated by providing simple "copilot" features thanks to the significant progress in AI algorithms. If the bidder were to use the copilot, he simply need to provide values for the desired bundles, dictating the decision of "how to bid" but let the AI recommend "what to bid". The bidding process can even take a hybrid design where bidders may simply do full proxy bidding using the AI provided by the platform, or train their own complete bidding algorithm, or use a mixture of the two: fine-tune the AI at key decision nodes. The transparent nature of CASA makes "truthfulness" even easier to maintain than traditional truthful mechanisms like the VCG.

References

Aggarwal, G. and Hartline, J. D. (2006). Knapsack auctions. In *SODA*, volume 6, pages 1083–1092.

Ausubel, L. M., Cramton, P., and Milgrom, P. (2006). The clock-proxy auction: A practical combinatorial auction design. *Handbook of spectrum auction design*, pages 120–140.

Ausubel, L. M. and Milgrom, P. R. (2002). Ascending auctions with package bidding. *The BE Journal of Theoretical Economics*, 1(1):20011001.

Bergemann, D., Brooks, B., and Morris, S. (2017). First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143.

Börgers, T. (1991). Undominated strategies and coordination in normalform games. *Social Choice and Welfare*, 8:65–78.

Brooks, B. and Du, S. (2021). Optimal auction design with common values: An informationally robust approach. *Econometrica*, 89(3):1313–1360.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1):180–194.

Carroll, G. (2014). A complexity result for undominated-strategy implementation. Technical report, Working Paper.

Carroll, G. (2017). Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488.

Cramton, P. and Schwartz, J. A. (2000). Collusive bidding: Lessons from the fcc spectrum auctions. *Journal of regulatory Economics*, 17(3):229–252.

Cramton, P. C., Shoham, Y., Steinberg, R., and Smith, V. L. (2006). *Combinatorial auctions*, volume 1. MIT press Cambridge.

Edelman, B., Ostrovsky, M., and Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1):242–259.

Goldberg, A. V. and Hartline, J. D. (2001). Competitive auctions for multiple digital goods. In *European Symposium on Algorithms*, pages 416–427. Springer.

Grimm, V., Riedel, F., and Wolfstetter, E. (2003). Low price equilibrium in multi-unit auctions: the gsm spectrum auction in germany. *International journal of industrial organization*, 21(10):1557–1569.

Hartline, J. D. (2013). Mechanism design and approximation. *Book draft. October*, 122(1).

He, W. and Li, J. (2022). Correlation-robust auction design. *Journal of Economic Theory*, 200:105403.

He, W., Li, J., and Zhong, W. (2022). Order statistics of large samples: theory and an application to robust auction design. Technical report, Technical report, Mimeo.

Jackson, M. O. (1992). Implementation in undominated strategies: A look at bounded mechanisms. *The Review of Economic Studies*, 59(4):757–775.

Jehiel, P. and Moldovanu, B. (2001a). Efficient design with interdependent valuations. *Econometrica*, 69(5):1237–1259.

Jehiel, P. and Moldovanu, B. (2001b). The european umts/imt-2000 licence auctions. *Imt-2000 Licence Auctions (May 2001).*

Klemperer, P. (2002). What really matters in auction design. *Journal of economic perspectives*, 16(1):169–189.

Levin, J. and Skrzypacz, A. (2016). Properties of the combinatorial clock auction. *American Economic Review*, 106(9):2528–2551.

Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–3287.

Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–71.

Roughgarden, T. (2014). Barriers to near-optimal equilibria. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 71–80. IEEE.

Roughgarden, T. (2015). Approximately optimal mechanism design: Motivation, examples, and lessons learned. *ACM SIGecom Exchanges*, 13(2):4–20.

Yamashita, T. (2015). Implementation in weakly undominated strategies, with applications to auctions and bilateral trade. *Review of Economic Studies*, 82(3):1223–1246.

Zhang, W. (2021). Correlation robustly optimal auctions. mimeo.

A Omitted Proofs

A.1 Proof of Proposition 3

Proof. Pick arbitrary $b \in \mathcal{M}$. Let $v_{b'} = \mathbf{1}_{b'=b} \cdot U[0,1]$, i.e. b is the only valuable bundle and its value is uniformly distributed on $[0,1]$. Let $G \sim \boldsymbol{v}$ and $\mathbb{G} = \{G\}$. Then, $\forall F \in \mathbb{F}$, $V_{\mathcal{M}}(F) \geq 0.5$. Define F^* as follows: uniformly randomly pick $k-1$ bidders and their values for b are identical and distributed according to $U[1-\frac{k-1}{N}, 1]$. For the remaining bidders, their values for b are identical and distributed according to $U[0, 1-\frac{k-1}{N}]$. It is easy to verify that $F^* \in \mathbb{F}$ and

$$\mathbb{E}_{F^*}[R^k_{\mathcal{M}}(\boldsymbol{v})] = \mathbb{E}_{U[0,1-\frac{k-1}{N}]}[x] = \frac{1}{2} - \frac{1}{2}\frac{k-1}{N} \leq \inf_{F \in \mathbb{F}} V_{\mathcal{M}}(F) - O\left(\frac{k}{N}\right).$$

Q.E.D.

A.2 Proof of Proposition 5

Proof. We slightly abuse notation and represent an allocation by a vector of sets $\boldsymbol{b} = (b_1, \ldots, b_N)$, where $b_n \cap b_{n'} = \emptyset$. b_n is the bundle allocated to bidder n. Let $\mathcal{B}_N$ denote the set of all feasible allocations with N bidders. Let $\boldsymbol{b}^*(\boldsymbol{v})$ denote the efficient allocation.

We establish a lower bound of the revenue-guarantee of the VCG mechanism by constructing, for each n, an allocation $\boldsymbol{b}^n \in \mathcal{B}_{N-1}$ of the objects to the bidders other than bidder n. Clearly, for any such profile $\boldsymbol{b}^n$,

$$\begin{aligned} R_{VCG}(\boldsymbol{v}) &= \sum_{n=1}^{N} \left(\sup_{\boldsymbol{b} \in \mathcal{B}_{N-1}} \sum_{n' \neq n} v^{n'}_{b_{n'}} - \sum_{n' \neq n} v^{n'}_{b^*_{n'}(\boldsymbol{v})} \right) \\ &\geq \sum_{n=1}^{N} \left(\sum_{n' \neq n} v^{n'}_{b^n_{n'}} - \sum_{n' \neq n} v^{n'}_{b^*_{n'}(\boldsymbol{v})} \right) \end{aligned} \tag{1}$$

For each n, we construct allocation $\boldsymbol{b}^n \in \mathcal{B}_{N-1}$ via the following algorithm:

Algorithm. Bundle $b^n_{n'} = \emptyset$ for all n'. Set $O = b^*_n(\boldsymbol{v})$.

(1). For each $n' \neq n$:

If $b^*_{n'}(\boldsymbol{v}) \neq \emptyset$, set $b^n_{n'} = b^*_{n'}(\boldsymbol{v})$.

Let $\bar{N} = \{n' : b^n_{n'} = \emptyset, n' \neq n\}$.

(2). If $O \neq \emptyset$, pick $o \in O$.

Set $b^n_{n'} = \{o\}$ for some $n' \in \arg\max_{n'' \in \bar{N}} v_{j''}(\{o\})$.

Update $O \leftarrow O \setminus \{o\}$ and $\bar{N} \leftarrow \bar{N} \setminus \{n'\}$.

(3). Repeat (2) until $O = \emptyset$.

(4). Return allocation $\boldsymbol{b}^n = (b^n_1, b^n_2, \ldots, b^n_{n-1}, b^n_{n+1}, \ldots, b^n_N)$.

In words, if an object is allocated to a bidder other than bidder n under $\boldsymbol{b}^*(\boldsymbol{v})$, the object is still allocated to that bidder. We then iteratively pick an object o that is allocated to bidder n under $\boldsymbol{b}^*(\boldsymbol{v})$, and allocate the object to the bidder n' whose value for the object $v^{n'}_{\{o\}}$ is the highest among all the bidders who are not allocated any object yet. For each $o \in b^*_n$, define n_o to be the index n' such that $b^n_{n'} = \{o\}$.

It follows from Equation (1) that

$$\begin{aligned} R_{VCG}(\boldsymbol{v}) &\geq \sum_{n=1}^{N} \left(\sum_{n' \neq n} v^{n'}_{b^n_{n'}} - \sum_{n' \neq i} v^{n'}_{b^*_{n'}(\boldsymbol{v})} \right) \\ &= \sum_{n=1}^{N} \sum_{o \in b^*_n} v^{n_o}_{\{o\}} \end{aligned}$$

$$\geq \sum_{o=1}^{M} v_{\{o\}}^{(M+1)} = R_{\mathcal{M}}^{M+1}(\boldsymbol{v}),$$

The second inequality follows from the construction of $\boldsymbol{b}^i$: when an object $o \in b_i^*$ is being allocated, it is allocated to the bidder n' whose value for the object $v_{\{o\}}^{n'}$ is the highest among all the bidders who are not allocated any object yet. Since each iteration assigns at least one good to one bidder and there are at most M goods, we have $v_{\{o\}}^{n'_o}$ must be at least the $(M+1)$-th highest value among all $v_{\{o\}}^{n}$. *Q.E.D.*