

Rank-Guaranteed Auctions*

Wei He[†] Jiangtao Li[‡] Weijie Zhong[§]

August 19, 2024

Abstract

We propose a combinatorial ascending auction that is “approximately” optimal, requiring minimal rationality to achieve this level of optimality, and is robust to strategic and distributional uncertainties. Specifically, the auction is *rank-guaranteed*, meaning that for any menu \mathcal{M} and any valuation profile, the ex-post revenue is guaranteed to be at least as high as the highest revenue achievable from feasible allocations, taking the $(|\mathcal{M}| + 1)^{th}$ -highest valuation for each bundle as the price. Our analysis highlights a crucial aspect of combinatorial auction design, namely, the design of menus. We provide simple and approximately optimal menus in various settings.

*We thank Mohammad Akbarpour, Martin Bichler, Tilman Börgers, Kim-Sau Chung, Rahul Deb, Jason Hartline, Paul Klemplerer, Fuhito Kojima, Andrew Komo, Paul Milgrom, Michael Ostrovsky, Antonio Penta, Ning Sun, Satoru Takahashi, and Andrzej Skrzypacz for helpful discussions. 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

1 Introduction

Auctions play a critical role in economic activities. For example, the online advertising sector generates trillions of dollars annually through the auctioning of advertising “slots”. The Federal Communications Commission (FCC) has collected over 200 billion dollars via auctioning radio spectrum. Despite the critical role these auctions play, there is 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 typically 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.” Likewise, in the realm of telecommunications, while larger service providers may pursue nationwide radio spectrum licenses to maximize their coverage, smaller providers often seek only regional licenses, prioritizing local markets over national presence.

The theoretical study of auctions, in contrast to reality, takes aggressive simplifications to an extent that overlooks the nuanced realities of the markets. The game theoretic approach focuses on the incentives of the bidders (incentive compatibility) and auctioneers (optimality), while limiting to very specific environments. Iconic theories, such as [Myerson \(1981\)](#), made progress by assuming a *single item, independent valuations* among bidders, full *Bayesian rationality*, and *shared prior beliefs*. None of these assumptions hold water in the complex scenarios described earlier, highlighting a clear disconnect. Conversely, the domain of operations research and computer science typically emphasizes the procedural aspects of communicating preferences and determining assignments under complete preferential complexity. This focus often neglects the incentives for bidders to reveal their preferences truthfully and for auctioneers to maximize revenue (see, e.g., a review by [Cramton et al. \(2006\)](#)). Therefore, the development of a comprehensive auction theory that simultaneously addresses both *incentives* and *complexity* represents a significant and unfulfilled challenge.

The goal of this study is to address the dilemma between incentives and complexity in auction design. This resonates with what [Carroll \(2019\)](#) speculates as the “future of economic design”:

“My expectation—and my hope—is that progress in mechanism design over the coming decades will come from developing more useful general models of preferences, information, and actions in complex environments, relatively free of structural assumptions; and developing conceptual tools to argue for why certain kinds of mechanisms will work well in such environments.”

To tackle the complexity of the practical environments, the same article then proposes:

“Making this progress toward more free-form models will require a shift in the criteria by which research in economic theory is evaluated.”

The methodology we develop in this paper hinges on the shift to a novel criterion of approximate optimality, termed the **rank guarantee**: the k^{th} -rank guarantee is the maximal ex-post revenue when each feasible bundle can be sold at the k^{th} -highest value among all bidders. With this approximation, we achieve a near-optimal resolution of the dilemma: it is possible to design a simple auction mechanism to respect the incentives—the auctioneer achieves a rank-guaranteed revenue as long as bidders avoid “obviously” bad strategies—while accommodating the full complexity—fully combinatorial preferences, without the need for any Bayesian prior.

In our model, 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)cending (A)uction (CASA). Prior to the auction, the auctioneer curates a *menu* of item bundles for allocation, denoted by \mathcal{M} . 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 resolves the dilemma within certain approximation bounds. First, we show that the outcome of CASA respects the incentives of both bidders and the auctioneer:

- **Bidder rationality**: All of our results apply to any *non-obviously dominated* strategy profile, a collection that excludes strategies that are *obviously dominated* in the sense that even in the most favorable case, they underperform some other strategies in their least favorable case (see [Li \(2017\)](#) and [Li and Dworzak \(2021\)](#)). In other words, we allow the bidders to bid fully strategically, while making a weak rationality assumption that all we know is that they avoid obviously bad choices.
- **Approximate auctioneer optimality**: In CASA, any non-obviously dominated strategy profile yields an *ex-post* revenue that is *rank-guaranteed* — achieving the maximal revenue when each bundle within the menu \mathcal{M} can be sold at the $(|\mathcal{M}| + 1)^{\text{th}}$ -highest value among all bidders. Since the 1^{st} -guarantee is the full ex-post trading surplus, the approximation we take is close in “ordinal distance” from the full surplus.

Second, we show how CASA maintains its performance in complex, unknown environments:

- **Prior-free:** The auction format, the selection of strategy profile, and the revenue guarantee does not depend on any Bayesian prior on either the auctioneer’s side or the bidders’ side, rendering the mechanism *prior-free* and our rank guarantee a prior-free approximation (see, e.g., Chapter 5 and 7 of [Hartline \(2013\)](#)).
- **Distributional robustness:** We quantify the rank-guarantee using canonical robust optimality criteria, i.e., the minimal ex-ante expectation when an adversarial nature chooses the joint distribution of values against the mechanism (see, e.g., [Carroll \(2017\)](#)). In the worst case, the revenue from CASA approximates the total surplus at the rate of $O\left(\frac{|\mathcal{M}|^2}{N}\right)$, i.e., CASA asymptotically achieves full surplus extraction when the number of bidders is large relative to the menu size.

To our knowledge, this paper is the first to introduce the rank-guarantee approximation. Compared to canonical approximation notions like constant-ratio guarantee and maxmin guarantee, the rank guarantee offers the added benefit of being easy to evaluate across various Bayesian and non-Bayesian environments. On the one hand, the rank guarantee provides an easily computable lower bound on the worst-case revenue, even when the underlying environment changes, such as when there are more bidders, when the menu changes, or when the auctioneer has additional information on the distribution of bidders’ valuations. Such adaptivity allows us to further simplify CASA by studying menu design. On the other hand, the lower bound performance is also straightforward to assess outside adversarial scenarios. For example, in the canonical setting where the values are independent and identically distributed, the rank guarantee is an appealing approximation when N is large, as all order statistics converge to the upper bound of the valuation support. More generally, beyond understanding the worst-case scenario, the rank guarantee can be useful to consider the potential upside of a mechanism in best-case scenarios.

Our framework highlights a crucial aspect of combinatorial auction design, namely, the design of menus. Crucially, the rank 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 the rank $|\mathcal{M}|$. Therefore, to close the approximation gap, a 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 improve approximation efficiency “for free”—menus that achieve the same benchmark 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	Simple and Sufficient Menu	Rank k
Weak substitutability	Individual items	$O(M)$
Weak complementarity	Grand bundle	2
Partitional complementarity	Partitional bundles	$O(M)$
Homogeneous goods	Menu of quantities	$O(M^2)$

Table 1: Simple and sufficient menus

The remainder of the introduction reviews related literature. [Section 2](#) introduces the auction format of CASA and the notion of rank-guarantee, and shows that CASA achieves the rank-guarantee. [Section 3](#) bounds the worst-case performance of rank-guaranteed auctions under distributional uncertainties. [Section 4](#) explores specific preference structures where CASA with simple menus performs as well as the complete menu.

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 \(2017\)](#), various authors have studied “robustly” optimal auctions that maximize the distributional worst-case revenue.¹ Particularly, [He and Li \(2022\)](#), [Zhang \(2022\)](#), and [Suzdaltsev \(2022\)](#) study robust versions of the single-unit auction problem in the distribution robust framework where the auctioneer has non-Bayesian uncertainty about the joint distribution of the bidders’ valuations. In this paper, we propose the notion of rank-guarantee, a distinct notion of approximate optimality. 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 Vickrey auction, few theoretic results have been established regarding multi-item auctions with combinatorial preferences. [Jehiel](#)

¹This research direction complements the large body of papers that focus on the case in which the designer does not have reliable information about the agents’ hierarchies of beliefs about each other while assuming the knowledge of the payoff environment; see, for example, [Bergemann and Morris \(2005\)](#), [Chung and Ely \(2007\)](#), [Chen and Li \(2018\)](#), [Du \(2018\)](#), [Brooks and Du \(2021\)](#), [Yamashita and Zhu \(2022\)](#), and [Brooks and Du \(2023\)](#).

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 Vickrey 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 communication/computational complexity, which is orthogonal to our focus on revenue performance and bidder incentives. Compared to other proposals like SAAPB (Ausubel and Milgrom (2002)) and the Combinatorial Clock Auction (CCA, see Ausubel et al. (2006) and Levin and Skrzypacz (2016)), CASA uses a simpler “pay-as-bid” rule, as opposed to personalized prices in SAAPB and demand reporting in CCA. Importantly, we assume that the bidders are fully strategic instead of single-minded. The menu design problem we tackle in Section 4 is akin to Rothkopf et al. (1998), which seeks to make a combinatorial auction computationally tractable by restricting the menu. We achieved the same goal while allowing for strategic bidders and maintaining the approximate optimality of the revenue.

Implementation in strategies that are not obviously dominated We study outcomes when agents are rational in the sense of avoiding obviously dominated strategies. This solution concept draws from the idea of obvious strategy-proof mechanisms (see Li (2017)) and is systematically studied in Li and Dworzak (2021). As in Li and Dworzak (2021), we assume that agents avoid obviously dominated strategies, but refrain from making assumptions regarding how agents select among strategies that are not obviously dominated. This methodology aligns with the spirit of implementation in undominated strategies; see for example Carroll (2014), Börgers (1991), Jackson (1992), and Yamashita (2015).

2 CASA and rank-guarantee

2.1 The auction environment

There is a set S of M items to be sold to N bidders. Let $\mathcal{N} = \{1, 2, \dots, N\}$. We write $b \subseteq S$ to denote a generic bundle of items. Let $\mathbf{v}^n = \{v_b^n\}_{b \subseteq S}$ denote the valuation vector of bidder n , where v_b^n is bidder n 's valuation of bundle b . Valuations are normalized so that $v_\emptyset^n = 0$ and $v_b^n \in [\underline{v}, \bar{v}]$ ($\underline{v} \geq 0$) for all $b \neq \emptyset$. A generic valuation profile is denoted by $\mathbf{v} = (\mathbf{v}^1, \mathbf{v}^2, \dots, \mathbf{v}^N)$. Let $\mathcal{M} \subseteq 2^S$ denote a *menu* of bundles chosen by the auctioneer. While \mathcal{M} is a choice variable of the auctioneer, for now we take it as exogenously given;

we defer the discussion of menu design to [Section 4](#). Assume that $N \geq |\mathcal{M}| + 1$. Let

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

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

2.2 The Combinatorial Ascending Auction

We define the *Combinatorial Ascending Auction* (CASA) as follows. The auction has an iterative structure, with the “state of the auction” characterized by the identity of the leading bidder and the leading price for each bundle. Initially, the leading price for each bundle is zero and none of the bidders is a leading bidder for any bundle. Bidders take turns raising the bids on the bundles, which determines new leading bidders and leading prices. The process repeats itself until when there are no new bids on any bundle. At that point, the auction stops. The auctioneer chooses a feasible allocation to maximize revenue, taking the leading prices as the prices for the bundles. There is also an activity rule designed to ensure that bidding activity starts out high and declines during the auction as prices rise far enough to discourage some bidders from continuing.

Formally, let $P \subset \mathbb{R}^+$ be a finite grid of feasible bids with grid size ϵ and $\max P > \bar{v}$.

(1) **Initialization stage** $t = 0$. Define

- the *leading bidder vector* at stage 0: $\phi^0 = (\phi_b^0)_{b \in \mathcal{M}} = \mathbf{0}$,
- the *leading price vector* at stage 0: $\mathbf{p}^0 = (p_b^0)_{b \in \mathcal{M}} = \mathbf{0}$,
- the set of *active bidders* at stage 0: $\mathcal{N}^0 = \{1, 2, \dots, N\}$.

(2) **Bidding stage** $t \geq 1$. An active bidder $n \in \mathcal{N}^{t-1}$ observes $(\mathbf{p}^{t-1}, \{b \mid \phi_b^{t-1} = n\})$, and decides whether to quit, which bundles to bid on, and how much to bid.²

- Bidder n may choose to quit by submitting an empty bid only if $\{b \mid \phi_b^{t-1} = n\} = \emptyset$, i.e., bidder n is not a leading bidder for any bundle in stage $t - 1$. Quitting is irreversible, that is, if bidder n chooses to quit, then bidder n becomes an inactive bidder and does not participate in future bidding rounds. Update:

$$- \phi^t = \phi^{t-1}, \mathbf{p}^t = \mathbf{p}^{t-1}, \mathcal{N}^t = \mathcal{N}^{t-1} \setminus \{n\}.$$

- If bidder n chooses not to quit, then she submits a bid—a nonempty set of bundle-price pairs $\{(b, p_b)\} \subset \mathcal{M} \times P$, subject to the requirements that (1) Leading bids are binding: if bidder n is the leading bidder at some bundle in

²The bidder selection rule and the observability of history is inconsequential for our analysis. For concreteness, we consider the selection rule that active bidders are cycled in ascending order according to their indices, and the observability of history is minimized to maximally protect privacy.

stage $t - 1$, then she must include that bundle in her bid with a bid that is *weakly* higher than the current leading price for that bundle, and (2) Minimum bid increment: if bidder n would like to bid on some bundle for which she is not the leading bidder, then her bid for that bundle must be *strictly* higher than the current leading price for that bundle. Update:

- $\phi_b^t = n$ and $p_b^t = p_b$ for any bundle b included in her bid,
- $\phi_{b'}^t = \phi_{b'}^{t-1}$ and $p_{b'}^t = p_{b'}^{t-1}$ for any bundle b' not included in her bid,
- $\mathcal{N}^t = \mathcal{N}^{t-1}$.

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

(3) **Allocation.** The auction ends (in stage T) when the leading prices stay constant for N consecutive periods. The auctioneer chooses a feasible allocation to maximize

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

Denote the maximizer by \mathbf{b}^* . Each bundle $b \in \mathbf{b}^*$ is allocated to ϕ_b^T at the price p_b^T .

In words, the auction format of CASA runs parallel ascending auctions for each bundle $b \in \mathcal{M}$. Then, the items are allocated to maximize the total price. We discuss possible variants of CASA and its relation to existing auction formats in [Section 2.4](#).

2.3 The rank-guarantee of CASA

In this subsection, we study the strategic behavior of the bidders and establish the rank-guarantee property of CASA. We only assume minimal rationality on the part of the bidders—bidders are rational in the sense of not playing obviously dominated strategies. Formally we adopt the solution concept of implementation in strategies that are not obviously dominated (see [Li \(2017\)](#) and [Li and Dworzak \(2021\)](#)). We first sketch the intuition, and then provide the formal arguments.

At any history, consider a non-leading bidder’s choice as to whether to quit. Obviously, as quitting is irreversible, quitting the auction leads to a best possible outcome of a zero payoff. Suppose that there is some bundle for which the bidder’s valuation is higher than the current leading price for that bundle. Consider the following strategy where the bidder raises the price for this particular bundle and never revises her bid afterwards. Clearly, this continuing strategy guarantees a non-negative payoff for the bidder. Thus, at least for the purpose of deciding whether to quit, it is “obviously optimal” not to quit.

More formally, let $h = (t, (\mathcal{N}^0, \dots, \mathcal{N}^{t-1}), (\mathbf{p}^0, \dots, \mathbf{p}^{t-1}), (\phi^0, \dots, \phi^{t-1}))$ denote a history of the game in stage t , H_t the set of such histories in stage t , and $H = \cup_{t \geq 0} H_t$. Suppose that bidder n is the active bidder in some stage t and I_n is bidder n ’s information set. Then the observed prices \mathbf{p} and n ’s leading bundles \mathbf{b} are the same for all

$h \in I_n$. Let \mathcal{I}_n denote all information sets of n . Let $s_n : \mathcal{I}_n \rightarrow 2^{\mathcal{M} \times \mathbb{R}^+}$ denote bidder n 's (pure behavioral) strategy and $u_n(\mathbf{s}, \mathbf{v}^n | h)$ the payoff to bidder n given valuation vector \mathbf{v}^n , strategy profile \mathbf{s} , conditional on the current history h and n bidding in period t .

Definition 1. A bidding strategy $s_n : \mathcal{I}_n \rightarrow 2^{\mathcal{M} \times P}$ is obviously dominated if there exists s'_n such that at any earliest point of departure I_n between s_n and s'_n ,

$$\begin{aligned} \sup_{s_{-n}, h \in I_n} u_n(\mathbf{s}, \mathbf{v}^n | h) &\leq \inf_{s_{-n}, h \in I_n} u_n(s'_n, \mathbf{s}_{-n}, \mathbf{v}^n | h); \\ \inf_{s_{-n}, h \in I_n} u_n(\mathbf{s}, \mathbf{v}^n | h) &< \sup_{s_{-n}, h \in I_n} u_n(s'_n, \mathbf{s}_{-n}, \mathbf{v}^n | h). \end{aligned}$$

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 strategy that dominates it. The second requirement guarantees that the set of non-obviously dominated strategies is non-empty. The earlier intuition then translates to:

Lemma 1. If there exists an information set $I_n \in \mathcal{I}_n$ (with observed prices \mathbf{p}) such that $s_n(I_n) = \emptyset$ (i.e., bidder n quits) and

- $\exists \mathbf{p}' \in P$ and $b \in \arg \max_{b \in \mathcal{B}(\mathcal{M})} \sum_{b' \in \mathbf{b}} p'_b$ such that $\mathbf{p}' \geq \mathbf{p}$ and $v_b^n > p'_b > p_b$,

then s_n is obviously dominated.

Proof. To show that s_n is obviously dominated, we explicitly construct another strategy s'_n that obviously dominates it. Consider the information set $I_n \in \mathcal{I}_n$ (with observed prices \mathbf{p}) such that $s_n(I_n) = \emptyset$ and

- $\exists \mathbf{p}' \in P$ and $b \in \arg \max_{b \in \mathcal{B}(\mathcal{M})} \sum_{b' \in \mathbf{b}} p'_b$ such that $\mathbf{p}' \geq \mathbf{p}$ and $v_b^n > p'_b > p_b$.

Obviously, the payoff from s_n conditional on any $h \in I_n$ is zero. Let s'_n be the same as s_n before the information set I_n and let bidder n bid $s'_n(I_n) = (b, p'_b)$ and never revise her bid afterwards.

We first discuss a special case in which $|\mathcal{N}^{t-1}| = 1$ and n is not a current leading bidder (otherwise quitting the auction is not feasible for bidder n). Then, all the current prices must be 0 and all the other bidders have quit (as this is the unique consistent history). Following the strategy s'_n , the auction ends with bidder n bidding (b, p'_b) , leading to a positive payoff of $v_b^n - p'_b > 0$ for bidder n .

Next, we consider the case in which $|\mathcal{N}^{t-1}| > 1$. It is clear that bidder n will get a nonnegative payoff regardless of the value profiles and bidding strategies of other bidders. We show that the best possible payoff for bidder n following the strategy s'_n is positive.

It suffices to consider the case in which any other subsequent active bidder bids p'_b for some bundle b' whenever possible, the auction ends with prices \mathbf{p}' , leading to a positive payoff of $v_b^n - p'_b > 0$ for bidder n . *Q.E.D.*

Note that [Lemma 1](#) suggests that our intuition is incomplete as we cannot fully rule out strategies that quit when the bidder's value for some bundle is above the leading price for that bundle. To rule out all 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. Nevertheless, we show below that the non-pivotal bidders are inconsequential for our analysis.

Let $S_{NOD}^n(P)$ denote the set of non-obviously dominated strategies of n given price grid P . Let $R(\mathbf{s}, \mathbf{v})$ denote the revenue to seller given value profile \mathbf{v} and strategy profile \mathbf{s} . Define

$$\underline{R}_{CASA}(\mathbf{v}) := \lim_{\epsilon \rightarrow 0} \inf_{s_n \in S_{NOD}^n(P)} R(\mathbf{s}, \mathbf{v}),$$

where the first limit inf is taken over $\epsilon \rightarrow 0$ and all P with grid size ϵ . That is, $\underline{R}_{CASA}(\mathbf{v})$ 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 \mathbf{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 the menu \mathcal{M} and the value profile \mathbf{v} is*

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

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

Our key observation is that CASA achieves the k^{th} -guarantee as long as bidders avoid strategies that are obviously dominated. In other words, for CASA to achieve the k^{th} -guarantee, we only need minimal rationality on the part of the bidders.

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

Proof. Consider a grid P with any grid size $\epsilon > 0$. We show that as long as bidders avoid obviously dominated strategies, for any value profile \mathbf{v} ,

$$\max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} p_b^T \geq \max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} (v_b^{(k)} - \epsilon).$$

Suppose to the contrary, this is not true. Then, there exists strategy profile s (where s_n is not obviously dominated for each bidder n), value profile \mathbf{v} , and $\delta > 0$ such that

$$\max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} p_b^T < \max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} (v_b^{(k)} - \epsilon - \delta).$$

Evidently, there exists some bundle b such that $p_b^T < v_b^{(k)} - \epsilon - \delta$, or equivalently, $p_b^T + \epsilon < v_b^{(k)} - \delta$. For each such bundle b , raise the price of bundle b to (the closest price below) $v_b^{(k)} - \delta$ sequentially until we find the first pivotal bundle \tilde{b} when $\max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} p_b$ is strictly improved by such increase in prices.

Since $k = |\mathcal{M}| + 1$ and $p_{\tilde{b}}^T < v_{\tilde{b}}^{(k)} - \epsilon - \delta$, there exists a bidder n with $v_{\tilde{b}}^n \geq v_{\tilde{b}}^{(k)}$ that quits the auction before the auction ends. Let h be the (on-path) history at which n quits (note that at this history, n must not be a leading bidder and $p_{\tilde{b}}(h) \leq p_{\tilde{b}}^t \leq v_{\tilde{b}}^n - \delta - \epsilon$, where $p_{\tilde{b}}(h)$ is the price of bundle \tilde{b} at the history h). Let $h \in I_n$. Then, $s_n(I_n) = \emptyset$. Let \mathbf{p}' be the raised prices. Then $\tilde{b} \in \arg \max_{\mathbf{b}' \in \mathcal{B}(\mathcal{M})} \sum_{b' \in \mathbf{b}'} p_{b'}$, and \mathbf{p}' satisfies the condition in [Lemma 1](#). Therefore, s_n is obviously dominated. We arrive at a contradiction. *Q.E.D.*

As we have pointed out following [Lemma 1](#), elimination of obviously dominated strategies does *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 as long as we are comfortable assuming that bidders avoid strategies that are obviously dominated.

CASA and its rank-guarantee become extremely simple when $|\mathcal{M}| = 1$, where CASA reduces to the canonical English auction of the unique feasible bundle and the k^{th} -guarantee becomes the second highest value. In this special case, [Theorem 1](#) shares the insight from [Li \(2017\)](#) that English auction achieves the 2^{nd} -guarantee via an obviously strategy-proof outcome. To understand our general result, we provide an alternative interpretation of this special case. The English auction made it obvious that a bidder should avoid losing the auction when there is a remaining surplus. Therefore, it maximally utilizes the competition among *losers* of the auction, pushing the price to the highest value among them. Meanwhile, it does not further screen the *winner* at all. CASA is exactly the multi-item generalization of this philosophy: it maximally invokes competition among losers by making it obviously suboptimal to lose with a remaining surplus. Meanwhile, it does not “care” at all which winner gets which bundle and how much extra he pays. As a result, the auctioneer extracts at least as much surplus as that from the losers only, which is exactly the rank-guarantee.

2.4 Discussions

Alternative formats: As we have discussed above, CASA is an intuitive extension of the canonical single-item English auction. The idea of generalizing the English auction to accommodate multiple items has been extensively explored. However, crucial to the design of CASA is a set of unique properties that cope with the incentives of the players. In comparison to the simultaneous ascending auction (SAA, [Milgrom \(2000\)](#)), allowing for bidding on bundles (package bidding) has the advantage of mitigating the demand reduction problem or the exposure problem, leading to sufficient competition. Compared to the combinatorial variants of SAA like SAAPB ([Ausubel and Milgrom \(2002\)](#)) and CCA ([Ausubel et al. \(2006\)](#)), CASA uses a simpler "pay-as-bid" rule so that the bidders find it straightforward to determine the remaining surplus from the auction, leading to the obvious strategy-proof implementation.³ Compared to the majority of iterative combinatorial auctions (Chapter 2, [Cramton et al. \(2006\)](#)), which seek to replicate VCG (known not to be rank-guaranteed⁴) and maximize efficiency, CASA has much better revenue performance. The design of CASA accommodates fully strategic bidders, as opposed to single-minded or myopic bidders assumed in the cited papers.

Except for the several crucial features, the exact format of CASA can be flexibly tailored to the auctioneer's needs and practical considerations without losing its various desirable properties. Notably, the menu \mathcal{M} is a design variable of the auctioneer, which we examine in detail in Section 4. Additionally, instead of allowing bidders to bid on multiple bundles, we could restrict them to bid on a single bundle when it is their turn to move (if a bidder is already leading for some bundle when it is her turn to move, she could not bid on any other bundle but she would remain an active bidder). Such restriction prevents the type of collusive communication documented in [Grimm et al. \(2003\)](#); [Jehiel and Moldovanu \(2001b\)](#). Moreover, we could also allow the bidders to observe the entire history if that is considered desirable for transparency purposes.

Robustness to collusion / irrationality: [Klemperer \(2002\)](#) identifies collusion as the first concern that "*really matters in auction design*". We illustrate below that our framework allows us to quantify the impact of collusive or irrational behavior on the performance of CASA. Recall that the key force that drives the revenue to the rank-guarantee is the competition among the losing bidders. Therefore, when there are non-strategic bidders, one can simply consider the competition among the subset of rational losing bidders.

Formally, 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.

³ The SAAPB is indeed k^{th} -guaranteed under non-strategic "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.

⁴ 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 $k \geq 2$.

Proposition 1. *Suppose there are j non-strategic bidders, then $\underline{R}_{CASA}(\mathbf{v}) \geq R_{\mathcal{M}}^k(\mathbf{v})$ 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.*

[Proposition 2](#) applies when a relatively small number of bidders play collusive/irrational strategies. Then, they reduce the rank of the guarantee by a small amount. Nevertheless, CASA is still rank-guaranteed.

What if *all* bidders are collusive? When the bidders form coalitions, and they strategically maximize group-level payoffs⁵, [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, $\underline{R}_{CASA}(\mathbf{v}) \geq R_{\mathcal{M}}^k(\mathbf{v})$ 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 $p_b \leq p_b^t \leq \max_{n \in c} \{v_b^n\} - \delta - \epsilon$). Let $h \in I_n$. Then, $s_n(I_n) = \emptyset$. Obviously, quitting gives the entire group zero payoff while bidding p_b' guarantees a non-negative payoff. Suppose all other bidders bid up to p_b' when it is their turn, the auction ends with \mathbf{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 competition among losing strategic bidders. 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, because CASA permits the minimum transmission of information. For example, anonymity prevents the reciprocity behavior documented in [Cramton and Schwartz \(2000\)](#).

Efficiency: In CASA, although bidders can fully avoid the exposure problem by bidding only on bundles with a positive surplus, they may strategically expose themselves. It is

⁵ Each group can freely shift allocations within the group and maximize the total payoff.

not obviously dominated to bid strictly above the true valuation for a bundle.⁶ Therefore, CASA might not satisfy ex-post IR; hence k^{th} -guaranteed revenue does not imply k^{th} -guaranteed surplus. Of course, CASA still satisfies ex-ante IR (assuming bidders have correct Bayesian priors) since quitting at the beginning is always an option; hence, the ex-ante bounds we derive in the following sections on the revenue of CASA also apply to surplus.

3 Rank-guarantee as a desideratum

Section 2 shows that CASA achieves the rank-guarantee as long as bidders avoid strategies that are obviously dominated. In this section, we explore the concept of rank-guarantee as a desideratum in auction design.

Clearly, if the auctioneer knows that the bidders' valuations are independent and identically distributed, then rank-guarantee is an appealing approximation when N is large, as all order statistics converge to the upper bound of the valuation support (at the rate of $\frac{1}{N}$). In what follows, we show that even when the auctioneer has non-Bayesian uncertainty about the joint distribution of the bidders' valuations and maximizes the revenue-guarantee (the worst-case expected revenue where the worst case is taken over all joint distributions that are perceived to be plausible), in many settings, rank-guarantee remains an appealing approximation.

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

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

where F_n is the marginal distribution of bidder n 's valuation. Thus, $\frac{1}{N} \sum F_n$ is the cumulative distribution function of the valuation of a uniformly randomly selected bidder in the population. We call \mathbb{F} an *ambiguity set*. Such an ambiguity set \mathbb{F} could come from the statistical estimation of F based on a "sanitized" dataset about valuations, that is, past bidders' 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 structures stemming from statistical inference.⁷

⁶ Consider, for example, the case with three bidders 1, 2, 3 and three items a, b, c . Bidder 1 only wants a , bidder 2 only wants b , and bidder 3 only wants the grand bundle $\{a, b, c\}$. The valuation of each bidder for the desired bundle is 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.

⁷ This setup covers a wide range of scenarios, as \mathbb{G} is completely general. \mathbb{G} could be a singleton set capturing the case in which the auctioneer has no uncertainty about the distribution of a representative bidder's valuation. \mathbb{G} could also be the set of distributions satisfying certain statistical properties (say moment conditions), capturing scenarios in which the

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

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

where ι , the assignment function satisfies $\iota(b) \neq \iota(b')$ for any $b \neq b'$. The ex-ante efficient surplus $V_{2^S}(F)$ with respect to the complete menu 2^S is denoted by $V^*(F)$ for simplicity.

Theorem 2. *For any menu \mathcal{M} ,*

$$\inf_{F \in \mathbb{F}} \mathbb{E}_F [R_{\mathcal{M}}^k(\mathbf{v})] \geq \inf_{F \in \mathbb{F}} V_{\mathcal{M}}(F) - \frac{(k-1)|\mathcal{M}|\bar{v}}{N}.$$

Proof. Let \mathbb{R} be a uniform random element of \mathcal{N} .

$$\begin{aligned} R_{\mathcal{M}}^k(\mathbf{v}) &= \max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} v_b^{(k)} \\ &\geq \max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \left(\sum_{b \in \mathbf{b}} v_b^{\mathbb{R}} - \sum_{b \in \mathbf{b}} \max \{v_b^{\mathbb{R}} - v_b^{(k)}, 0\} \right) \\ &\geq \max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \left(\sum_{b \in \mathbf{b}} v_b^{\mathbb{R}} \right) - \sum_{b \in \mathcal{M}} \max \{v_b^{\mathbb{R}} - v_b^{(k)}, 0\} \\ \implies \mathbb{E}_F [R_{\mathcal{M}}^k(\mathbf{v})] &\geq \mathbb{E}_F \left[\max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} v_b^{\mathbb{R}} \right] - \sum_{b \in \mathcal{M}} \mathbb{E}_F \left[v_b^{\mathbb{R}} - v_b^{(k)} \mid v_b^{\mathbb{R}} > v_b^{(k)} \right] \text{Prob}(v_b^{\mathbb{R}} > v_b^{(k)}) \\ &\geq \mathbb{E}_F \left[\max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} v_b^{\mathbb{R}} \right] - \sum_{b \in \mathcal{M}} \bar{v} \text{Prob}(v_b^{\mathbb{R}} > v_b^{(k)}) \\ &\geq \mathbb{E}_F \left[\max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} v_b^{\mathbb{R}} \right] - \frac{(k-1)|\mathcal{M}|\bar{v}}{N}. \end{aligned}$$

This further implies that

$$\begin{aligned} \inf_{F \in \mathbb{F}} \mathbb{E}_F [R_{\mathcal{M}}^k(\mathbf{v})] &\geq \inf_{F \in \mathbb{F}} \mathbb{E}_F \left[\max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} v_b^{\mathbb{R}} \right] - \frac{(k-1)|\mathcal{M}|\bar{v}}{N} \\ &= \inf_{G \in \mathbb{G}} \mathbb{E}_G \left[\max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} v_b \right] - \frac{(k-1)|\mathcal{M}|\bar{v}}{N} \\ &\geq \inf_{F \in \mathbb{F}} V_{\mathcal{M}}(F) - \frac{(k-1)|\mathcal{M}|\bar{v}}{N}, \end{aligned}$$

where the equality follows from the definition of the set \mathbb{F} . For the last inequality, observe that for any $G \in \mathbb{G}$, the joint distribution where each bidder's value is distributed according to G and all bidders' values are maximally positively correlated is contained in

auctioneer also has some non-Bayesian uncertainty about the distribution of a representative bidder's valuation.

\mathbb{F} . Therefore,

$$\inf_{F \in \mathbb{F}} V_{\mathcal{M}}(F) \leq \inf_{G \in \mathbb{G}} \mathbb{E}_G \left[\max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} v_b \right].$$

Q.E.D.

Theorem 2 highlights a key trade-off between *menu sufficiency* and *approximation efficiency*. Evidently, presenting the bidders with a larger menu has the potential of increasing allocation efficiency. Particularly, \mathcal{M} can be naively chosen to be the complete menu 2^S to guarantee full menu sufficiency, i.e., $V_{\mathcal{M}}(F) = V^*(F)$. However, this leads to $|\mathcal{M}|$ growing exponentially in M , causing both complex auction process and slow convergence. On the other hand, 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.

3.1 Discussions

Alternative ambiguity sets: The proof of **Theorem 2** goes through for a general \mathbb{F} if the following equality holds.

$$\inf_{F \in \mathbb{F}} \mathbb{E}_F \left[\max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} v_b^{\textcircled{R}} \right] = \inf_{F \in \mathbb{F}} V_{\mathcal{M}}(F).$$

That is, in the worst case, randomly selecting bidders performs as well as optimally selecting bidders. Therefore, it is straightforward that **Theorem 2** holds under an alternative ambiguity set:

$$\widehat{\mathbb{F}} = \left\{ F \in \Delta([\underline{v}, \bar{v}]^{N \times 2^S}) \mid \forall n, F_n \in \mathbb{G} \right\}.$$

Under $\widehat{\mathbb{F}}$, the auctioneer knows that the marginal distributions of bidders are contained in \mathbb{G} , but nothing beyond that. This ambiguity set exactly captures the correlational uncertainty studied in [Carroll \(2017\)](#), [He and Li \(2022\)](#), [Zhang \(2022\)](#), and [Suzdaltsev \(2022\)](#).

Distributions with unbounded support: Suppose the support of distributions in \mathbb{G} has unbounded support, i.e., $\bar{v} = \infty$, then the bound derived in **Theorem 2** has no bite. However, observe that

$$\mathbb{E}_F \left[v_b^{\textcircled{R}} \mid v_b^{\textcircled{R}} > v_b^{(k)} \right] \text{Prob}(v_b^{\textcircled{R}} > v_b^{(k)}) \leq \int_{Q_{\bar{F}_b}^{(k-1)/N}}^{\infty} v_b d\bar{F}(v),$$

where $\bar{F} = \frac{F_n}{N}$ and $Q_{\bar{F}_b}^{(k-1)/N}$ is the top $\frac{k-1}{N}$ quantile of the marginal of \bar{F} for v_b . To see why the inequality holds, the LHS is the constrained expectation of v_b given \bar{F} on *some* event of probability at most $\frac{k-1}{N}$. The RHS is the constrained expectation of v_b given \bar{F} on the probability- $\frac{k-1}{N}$ event that maximizes its value. Therefore,

$$\inf_{F \in \mathbb{F}} \mathbb{E}_F[R_{\mathcal{M}}^k(\mathbf{v})] \geq \inf_{F \in \mathbb{F}} V_{\mathcal{M}}(F) - \sup_{G \in \mathbb{G}} \sum_{b \in \mathcal{M}} \int_{Q_{G_b}^{(k-1)/N}}^{\infty} v_b dG(v).$$

For any G , the concentration inequality implies $\int_{Q_{G_b}^{(k-1)/N}}^{\infty} v_b dG(v) \leq \frac{\mathbb{E}_G[v_b^2]}{Q_{G_b}^{(k-1)/N}}$. If G has unbounded support, so is the quantile $Q_{G_b}^{(k-1)/N}$. As a result, when distributions in \mathbb{G} have unbounded support but uniformly bounded second moment, the worst-case rank-guarantee converges to the full surplus at the rate of $O\left(\frac{|\mathcal{M}|}{\inf_{G \in \mathbb{G}} Q_{G_b}^{(k-1)/N}}\right)$.

Tightness of the bound: The coefficient $k|\mathcal{M}|$ in [Theorem 2](#) consists of two parts. The coefficient k comes from the k^{th} highest value approximation. The coefficient $|\mathcal{M}|$ comes from the total number of bundles in the menu \mathcal{M} . [Proposition 3](#) below shows that the dependence on k is tight.

Proposition 3. *For any M, N, \mathcal{M}, k , there exists some \mathbb{G} such that*

$$\inf_{F \in \mathbb{F}} \mathbb{E}_F[R_{\mathcal{M}}^k(\mathbf{v})] \leq \inf_{F \in \mathbb{F}} V_{\mathcal{M}}(F) - O\left(\frac{k}{N}\right).$$

Proof. See [Appendix A.1](#).

Q.E.D.

The dependence of the approximation gap on $|\mathcal{M}|$, however, might not be tight as we take a very loose upper bound in the proof of [Theorem 2](#): we bound the revenue loss from the allocated bundles (up to M of them) by the revenue loss from all bundles ($|\mathcal{M}|$ of them). While we speculate that the coefficient $|\mathcal{M}|$ can be improved, a formal proof is yet unknown to us.

4 Menu design and simple menus

In this section, we examine several canonical classes of preference structures where there exist menus that are both *sufficient*, ensuring full allocation efficiency, and *simple*, with menu size growing at a polynomial rate as M increases.

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

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

In words, a menu \mathcal{M} is \mathbb{G} -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, \mathbb{G} -sufficiency is defined with respect to 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.⁸ Nevertheless, sufficiency guarantees full allocation efficiency:

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

$$\inf_{F \in \mathbb{F}} \mathbb{E}_F [R_{\mathcal{M}}^k(\mathbf{v})] \geq \inf_{F \in \mathbb{F}} V^*(F) - \frac{(k-1)|\mathcal{M}|\bar{v}}{N}.$$

Proof.

$$\begin{aligned} \inf_{F \in \mathbb{F}} \mathbb{E}_F [R_{\mathcal{M}}^k(\mathbf{v})] &\geq \inf_{G \in \mathbb{G}} \mathbb{E}_G \left[\max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} v_b \right] - \frac{(k-1)|\mathcal{M}|\bar{v}}{N} \\ &= \inf_{G \in \mathbb{G}} \mathbb{E}_G \left[\max_{\mathbf{b} \in \mathcal{B}(2^S)} \sum_{b \in \mathbf{b}} v_b \right] - \frac{(k-1)|\mathcal{M}|\bar{v}}{N} \\ &\geq \inf_{F \in \mathbb{F}} V^*(F) - \frac{(k-1)|\mathcal{M}|\bar{v}}{N}, \end{aligned}$$

where the equality follows from the \mathbb{G} -sufficiency of menu \mathcal{M} , and the two inequalities have been established in the proof of Theorem 2. *Q.E.D.*

A simple sufficient condition for the \mathbb{G} -sufficiency of menu \mathcal{M} is that

$$\max_{\mathbf{b} \in \mathcal{B}(2^S)} \sum_{b \in \mathbf{b}} v_b = \max_{\mathbf{b} \in \mathcal{B}(\mathcal{M})} \sum_{b \in \mathbf{b}} v_b \tag{1}$$

holds ex-post, i.e. $\forall \mathbf{v} \in \text{Supp}(\mathbb{G}) := \cup_{G \in \mathbb{G}} \text{Supp}(G)$. This condition allows us to convert combinatorial preferences into sufficiency. Theorem 3, as well as Equation (1), states that one only needs to verify the sufficiency of a menu based on individual bidder's preference, as opposed to the distribution of valuations among all bidders. This is a consequence of the robustness concern. Recall from our analysis in Section 3 that the adversarial nature minimizes the rank-guarantee by making all losing bidders identical. Then, in the worst-case, a \mathbb{G} -sufficient menu performs as good as the complete menu. In the alternative cases where losing bidders are asymmetric, even though the \mathbb{G} -sufficient menu underperforms the complete menu, the extra surplus from asymmetry leads to an even higher rank-guarantee. With Theorem 3 and Equation (1), we derive simple sufficient menus for several canonical preference structures.

⁸ Consider for instance two items and two bidders, where for each bidder the sum of the value for each individual item is more than her value for the grand bundle. 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.

Weak substitutability and itemized ascending auction

Definition 4. We say that bidder preferences exhibit **weak substitutability** if for any $\mathbf{v} \in \text{Supp}(\mathbb{G})$ and $b \subseteq S$,

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

In words, a representative bidder finds the value of any bundle weakly lower than the sum of her value for each item in the bundle. Weak substitutability is a necessary condition for various substitutability notions studied in the literature.

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

When $\mathcal{M} = S$, CASA reduces to a simple itemized ascending auction, where the allocation is determined jointly in the end. Weak substitutability is one of the most widely studied preference assumptions in the literature as it captures a natural diminishing return to scale. Our analysis shows that under such preference structures, CASA exhibits 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 preferences exhibit weak substitutability, then the Vickery auction achieves a revenue guarantee of $R_S^{M+1}(\mathbf{v})$.*

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 shown 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 rank-guarantee.

Weak complementarity and the second-price auction

Definition 5. Bidder preferences exhibit **weak complementarity** if for any $\mathbf{v} \in \text{Supp}(\mathbb{G})$ and $\mathbf{b} \in \mathcal{B}(2^S)$,

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

In words, a representative bidder finds the value of the grand bundle weakly 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 preferences exhibit weak substitutability, then the menu $\mathcal{M} = \{S\}$ is \mathbb{G} -sufficient and $k = 2$.*

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

“Partitional” complementarity A hybrid case of substitutability and complementarity is the partitional complementarity which we define below, described by a partition \mathcal{K} of S .

Definition 6. *Let \mathcal{K} be a partition of S . Bidder preferences exhibit \mathcal{K} -partitional complementarity if for any $\mathbf{v} \in \text{Supp}(\mathbb{G})$,*

$$\begin{aligned} &\text{for any } b \in \mathcal{K} \text{ and partition } \kappa \text{ of } b, \sum_{b' \in \kappa} v_{b'} \leq v_b; \\ &\text{for any } b' \subseteq S, \sum_{b \in \mathcal{K}} v_{b \cap b'} \geq v_{b'}. \end{aligned}$$

In words, \mathcal{K} -partitional complementarity structure means there is weak complementarity within each $b \in \mathcal{K}$ and weak substitutability across each $b \in \mathcal{K}$.

Proposition 7. *If bidder preferences exhibit \mathcal{K} -partitional complementarity, then the menu $\mathcal{M} = \mathcal{K}$ is \mathbb{G} -sufficient and $k = |\mathcal{K}| + 1$.*

In some cases, the auctioneer may understand that bidder preferences exhibits partitional complementarity, but does not know the exact partition. Proposition 7 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 $k \sim \text{Poly}(M)$. Such partitional complementarity 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 $u : \mathbb{N} \rightarrow [\underline{v}, \bar{v}]$ such that for any $\mathbf{v} \in \text{Supp}(\mathbb{G})$ and $b \in S$,*

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

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 $\mathcal{B} : (\mathcal{M}) = \{X \subset \mathcal{M} \mid \sum_{b \in X} |b| \leq M\}$, i.e., an allocation is feasible as long as the total number of items being allocated is below M .

Proposition 8. *If goods are homogeneous, then the menu $\mathcal{M} = \cup_{l \in \{1, \dots, M\}} \{b_l^1, \dots, b_l^{\lfloor \frac{M}{l} \rfloor}\}$ is \mathbb{G} -sufficient and $k \leq \frac{M^2 + M}{2}$, where $\{b_l^j\}$ are distinct bundles of size l .*

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 be 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 $\text{Poly}(M)$. 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 of CASA that guarantees an approximately optimal ex-post revenue. To achieve this, we only need to assume minimal rationality on the part of the bidders. In addition, we show that CASA is robust to distributional and strategic uncertainties under certain approximations. In practice, however, these approximation gaps may become non-negligible, rendering the deployment of CASA challenging.

- *Thin markets:* 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 keeps the revenue guarantee appealing and CASA directly applicable. Even in settings with a large number of items and a small number of bidders where our theory has little bite, menu design may still be a cost-effective way to promote competition and improve the revenue performance of existing auctions.

- *Proxy bidding:* While CASA simplifies the bidding process by clarifying “whether to quit,” the complexity of determining “which bundles to bid on” and “how much

to bid” remains unresolved. A *truthful* and full proxy-bidding version of CASA is not yet known to us. This makes the deployment of CASA challenging in environments that require fast resolutions of auctions. Nevertheless, we propose that advancements in AI could mitigate this by introducing ”copilot” features that assist bidders in decision-making. By integrating AI as the bidding proxy, bidders would only need to specify values for desired bundles, with the AI advising on bid placement. This could evolve into a hybrid model where bidders either rely fully on platform-provided AI, develop their own bidding algorithms, or use a combination of both strategies.

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. and Morris, S. (2005). Robust mechanism design. *Econometrica*, 73:1771–1813.
- 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.
- Brooks, B. and Du, S. (2023). On the structure of informationally robust optimal mechanisms. *Available at SSRN 3663721*.
- 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.
- Carroll, G. (2019). Design for weakly structured environments. *The Future of Economic Design: The Continuing Development of a Field as Envisioned by Its Researchers*, pages 27–33.
- Chen, Y.-C. and Li, J. (2018). Revisiting the foundations of dominant-strategy mechanisms. *Journal of Economic Theory*, 178:294–317.
- Chung, K.-S. and Ely, J. C. (2007). Foundations of dominant-strategy mechanisms. *Review of Economic Studies*, 74(2):447–476.
- 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.
- Du, S. (2018). Robust mechanisms under common valuation. *Econometrica*, 86(5):1569–1588.
- 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/int-2000 licence auctions. *Int-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, J. and Dworzak, P. (2021). Are simple mechanisms optimal when agents are unsophisticated? In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 685–686.
- Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–3287.

- Milgrom, P. (2000). Putting auction theory to work: The simultaneous ascending auction. *Journal of political economy*, 108(2):245–272.
- Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–71.
- Rothkopf, M. H., Pekeč, A., and Harstad, R. M. (1998). Computationally manageable combinatorial auctions. *Management science*, 44(8):1131–1147.
- Roughgarden, T. (2015). Approximately optimal mechanism design: Motivation, examples, and lessons learned. *ACM SIGecom Exchanges*, 13(2):4–20.
- Suzdaltsev, A. (2022). Distributionally robust pricing in independent private value auctions. *Journal of Economic Theory*, 206:105555.
- Yamashita, T. (2015). Implementation in weakly undominated strategies, with applications to auctions and bilateral trade. *Review of Economic Studies*, 82(3):1223–1246.
- Yamashita, T. and Zhu, S. (2022). On the foundations of ex post incentive-compatible mechanisms. *American Economic Journal: Microeconomics*, 14(4):494–514.
- Zhang, W. (2022). Correlation-robust optimal auctions.

A Omitted Proofs

A.1 Proof of Proposition 3

Proof. Pick an arbitrary bundle $b \in \mathcal{M}$. Let $v_b = \mathbf{1}_{b'=b} \cdot U[0, 1]$; that is, b is the only valuable bundle and its value is uniformly distributed on $[0, 1]$. Let G denote such a distribution and $\mathbb{G} = \{G\}$. Then, $V_{\mathcal{M}}(F) \geq \frac{1}{2}$ for any $F \in \mathbb{F}$. 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 straightforward to verify that $F^* \in \mathbb{F}$ and

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

Q.E.D.

A.2 Proof of Proposition 5

Proof. We slightly abuse notation and represent an allocation by a vector of sets $\mathbf{b} = (b_1, b_2, \dots, b_N)$, where $b_n \cap b_{n'} = \emptyset$ and b_n is the bundle allocated to bidder n . Let \mathcal{B}_N denote the set of all feasible allocations with N bidders. Let $\mathbf{b}^*(\mathbf{v})$ denote the efficient allocation.

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

$$\begin{aligned} R_{VCG}(\mathbf{v}) &= \sum_{n=1}^N \left(\sup_{\mathbf{b} \in \mathcal{B}_{N-1}} \sum_{n' \neq n} v_{b_{n'}}^{n'} - \sum_{n' \neq n} v_{b_{n'}^*(\mathbf{v})}^{n'} \right) \\ &\geq \sum_{n=1}^N \left(\sum_{n' \neq n} v_{b_{n'}^n}^{n'} - \sum_{n' \neq n} v_{b_{n'}^*(\mathbf{v})}^{n'} \right). \end{aligned} \quad (2)$$

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

Algorithm. Bundle $b_{n'}^n = \emptyset$ for all n' . Set $O = b_n^*(\mathbf{v})$.

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

If $b_{n'}^*(\mathbf{v}) \neq \emptyset$, set $b_{n'}^n = b_{n'}^*(\mathbf{v})$.

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

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

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

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

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

(4). Return allocation $\mathbf{b}^n = (b_1^n, b_2^n, \dots, b_{n-1}^n, b_{n+1}^n, \dots, b_N^n)$.

In words, if an object is allocated to a bidder other than bidder n under $\mathbf{b}^*(\mathbf{v})$, then the object is still allocated to that bidder. We then iteratively pick an object o that is allocated to bidder n under $\mathbf{b}^*(\mathbf{v})$, and allocate the object 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. For each $o \in b_n^*$, define n_o to be the index n' such that $b_{n'}^{n_o} = \{o\}$.

It follows from Equation (2) that

$$\begin{aligned}
R_{VCG}(\mathbf{v}) &\geq \sum_{n=1}^N \left(\sum_{n' \neq n} v_{b_{n'}^n}^{n'} - \sum_{n' \neq n} v_{b_{n'}^*(\mathbf{v})}^{n'} \right) \\
&= \sum_{n=1}^N \sum_{o \in b_n^*} v_{\{o\}}^{n_o} \\
&\geq \sum_{o=1}^M v_{\{o\}}^{(M+1)} \\
&= R_{\mathcal{M}}^{M+1}(\mathbf{v}).
\end{aligned} \tag{3}$$

The first equality holds since (a) when $b_{n'}^*(\mathbf{v}) \neq \emptyset$, $b_{n'}^n = b_{n'}^*(\mathbf{v})$, and (b) when $b_{n'}^*(\mathbf{v}) = \emptyset$, $b_{n'}^n$ is either $\{o\}$ for some $o \in O$, or \emptyset otherwise. The second inequality follows from the construction of \mathbf{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.*