Prefix-Based Collection Auction: A Mechanism against Market Power and Collusion

Taubman, Dmitriy

21 August 2025

Prefix-Based Collection Auction: A Mechanism against Market Power and Collusion

Dmitriy Taubman

Working Paper, August 21, 2025

Abstract

We introduce a new collection auction mechanism for selling multiple identical items to a single winner—the *Prefix-Based Collection Auction*. The auction restricts the winner to a prefix of their bids and imposes a payment rule based on both an internal prefix sum and an external second price. This dual structure offers strong protection against both market power and bidder collusion, while maintaining intuitive and truthful bidding behavior. The mechanism is robust, simple to implement, and has potential applications in art-collection markets, online advertising, and other environments where bundle demand is critical.

Note. This paper is a working paper. A version will also be available on arXiv.

Keywords: auctions; mechanism design; game theory; collusion resistance; market power; prefix structure; allocation rules.
JEL Codes: D44; C72; D47.

1 Introduction

In many auction environments, a single participant may acquire multiple items, forming a collection. However, standard mechanisms such as the Vickrey auction or uniform pricing do not prevent manipulation by dominant bidders or collusion among participants. To address this, we propose a novel mechanism that relies on the idea of a *prefix*—a sequence of consecutive bids by a single bidder—and combines two pricing safeguards: a prefix-based internal price and a second-best external price.

We propose a middle-ground solution: the *Collection Auction with prefix and monotonicity constraint*, where:

- Only a single bidder can win a contiguous prefix of the items;
- The bidder must submit a sequence of non-decreasing price bids—one for each unit;
- Competing bidders may also submit bids for any number of consecutive units;

- The winner is the bidder with the highest-value prefix, and payment follows a refined second-price rule:
  - If the sum of the winner's prefix excluding its last bid exceeds the value of the best competing prefix submitted by other bidders, the winner pays this internal "second price";
  - Otherwise, the winner pays the best external prefix value offered by other participants.

We analyze the trade-offs of this format versus classic bundle auctions and multi-winner auctions. We demonstrate that, under natural conditions, collection auctions with monotonic prefixes deliver better outcomes in terms of strategic robustness, resistance to collusion, and pricing discipline, particularly when one or more bidders seek the full collection but others still value parts of it.

This auction format is especially relevant for digital advertising markets. Advertisers often prefer to purchase entire collections of ad slots—for example, all banners on a single webpage—to achieve full visibility and avoid being shown alongside competitors. In such cases, the value derives not just from individual items, but from the entire bundle. Our mechanism is designed to accommodate this type of demand by allowing a single bidder to win a full prefix of slots and ensuring that the price they pay reflects both competition and exclusivity.

Finally, we emphasize the conceptual separation between allocation and pricing. We divide the mechanism into two components: the **Winner Determination Rule (WDR, allocation stage)** and the **Winner Payment Rule (WPR, payment stage)**. This separation makes the structure of the mechanism clearer and more modular, allowing allocation rules and payment rules to be analyzed independently and recombined in different designs. Unlike the traditional notion of the *Winner Determination Problem (WDP)*, which emphasizes computational difficulty, the concept of WDR highlights that allocation is not merely a problem to be solved, but a normative rule that shapes the overall properties of the mechanism.

2 Winner Determination Rule (WDR)

2.1 Setting and Definitions

The seller offers K identical indivisible items. Each bidder $i \in B$ submits a bid sequence (prefix):

$$b_{i,1},\ b_{i,2},\ \ldots,\ b_{i,K}.$$

The sequence must be non-decreasing:

$$b_{i,1} \leq b_{i,2} \leq \cdots \leq b_{i,K}.$$

The value of bidder i's prefix bid is

$$S_i(k) = \sum_{j=1}^{k} b_{i,j}. \tag{1}$$

The auction selects a single winning bidder i and a prefix length k, allocating items 1 to k to that bidder. Only one bidder wins, and the allocation must be a contiguous prefix from that bidder's sequence:

$$i^{=\arg\max_i \sum_{j=1}^{k} b_{i,j}}.(2)$$

In our prefix structure with a monotone bid allocation and appropriate conditions on the price function, only one agent can win the entire prefix (Appendix A). This structural constraint not only simplifies strategic behavior, but also provides strong protection against collusion and price manipulation.

3 Winner Payment Rule (WPR) and Truthfulness

3.1 Payment Rule

Our goal is to define a payment rule such that truthful bidding becomes a dominant strategy for each bidder—meaning it maximizes their utility regardless of what others do.

Let:

- Bidder i win the auction with a prefix of k items;
- S_1 (internal prefix price) is the total bid submitted by the winner up to but not including their last winning item;
- S_2 (external alternative price) is the highest total bid that any other agent offered for the same collection of items.

We propose the following second-price-like payment rule:

$$\begin{aligned} S_1 &= \sum_{j=1}^{k-1} b_{i,j}, \\ S_2 &= \max_{m \neq i \sum_{j=1}^{k} b_{m,j}}, \\ P &= \max(S_1, S_2). \end{aligned} \tag{3}$$

Where P — Payment.

3.2 Bidding Format and Constraints

Each participant i submits a sequence of bids $\{b_{i,1}, b_{i,2}, \dots\}$, where $b_{i,j}$ is the price they are willing to pay for the j-th item in their collection. The auction enforces a maximum step constraint:

$$|b_{i,j+1} - b_{i,j}| \leq \Delta \qquad (\Delta > 0 \text{ small; e.g., linear or capped increment}).$$

Attempting to bid low and then suddenly jump would violate the Δ constraint or weaken their own prefix sum S_i. This condition ensures that no bidder can begin with small bids and then jump sharply to win an item at a disproportionately high price.

4 Why both S_1 and S_2 are necessary: mutual protection

- S_1 **protects against external collusion and manipulations by others.** Suppose other bidders try to help the winner by lowering their bids, hoping to reduce the winner's payment. This trick fails—the winner must still pay at least S_1, their own internal prefix sum. So even if $S_2 \downarrow$ due to collusion, the payment doesn't fall below S_1.
- S_2 **protects against internal manipulation by the winner.** The winner can try to reduce his contribution by making minimal bids to lower S_1. But even if $S_1 \downarrow$, the winner must pay at least the price offered by others and the payment doesn't fall below S_2.

Therefore, the final contribution is $P = \max(S_1, S_2)$. This rule ensures fairness and resistance to both collusion (reducing S_2 by other participants) and strategic manipulation (reducing S_1 by the winner).

5 What do S_1 and S_2 do separately

We conjecture that either S_1 or S_2 alone may be sufficient to guarantee *veracity*, the property that it is optimal for each agent to report her true valuation vector $v = (v_1, \ldots, v_k)$. This conjecture follows the same intuition as the classical proof of the Vickrey auction: a bidder cannot benefit from either overbidding ($b_k > v_k$) or underbidding ($b_k < v_k$) her valuation for any prefix of length $k \leq K$, since the payoff does not depend on the final marginal bid, but only on previous bids or competing alternatives.

One might ask whether the second-price rule S_1 preserves truthfulness, given that the winner pays the sum of their own bids for the first $K - 1$ items. After all, the payment appears to depend on the agent's own bid—seemingly contradicting the common principle that "the payment should not depend on the winner's own report."

Indeed, under S_1, the winner pays

$$S_1 = \sum_{j=1}^{k-1} b_{i,j},$$

i.e., their own prefix sum excluding the final bid that secures the K-th item. So yes—the payment depends on the winner's own bid. But only on the part of the bid that does not influence whether they win. The total bid determines whether the agent wins. The payment is then based on a prefix that had no marginal effect on the outcome.

Deviating from the true vector v can only lead to a worse outcome: either losing when she could have won, or winning at a loss. More precisely, if the agent wins k items, his utility is

$$u = \begin{cases} \sum_{\ell=1}^{k} v_\ell - P, & \text{if he wins,} \\ 0, & \text{if he loses,} \end{cases}$$

where P is determined by the rule $P = \max(S_1, S_2)$. Although a formal proof has yet to be established, this pricing rule preserves the basic idea of external and internal stability. It suggests that even implementing only one of the two components, either S_1 or S_2, may be sufficient to ensure incentive compatibility.

6 Conclusion

In a prefix bidding auction, only one bidder wins, and the allocation must be a continuous prefix of its bid sequence. This auction differs from standard multi-item auctions such as discriminatory auctions, single-price (uniform-price) auctions, direct fractional auctions (DFA), and even allocation rules based on a descending sequence of bids, where multiple winners are possible. The proposed prefix-based Collection Auction introduces a novel combination of incentive compatibility and collusion resistance, achieved by leveraging a structured bidding space and a second-price rule defined via the maximum of two alternative values—the internal prefix (S_1) and the external alternative (S_2).

This mechanism stands in contrast with a classic Vickrey-style auction for a full collection, where all items are sold as a single lot and the winner pays the second-highest bid. While such an auction ensures truthfulness, it remains vulnerable to collusion: bidders may coordinate to suppress the second-highest bid and reduce the payment of the winning coalition. In particular, multiple agents with overlapping interests in the collection can jointly underbid and still win—a strategy not penalized in the Vickrey framework.

By contrast, the prefix-based mechanism:

- Breaks the space of bids into incremental prefixes,
- Ensures that only one agent wins, eliminating competitive collusion,
- And applies a second-price rule using both internal and external benchmarks $P = \max(S_1, S_2)$, providing strong protection against both overbidding and underbidding.

Thus, the Collection Auction retains the desirable truthfulness property of Vickrey auctions while adding strategic robustness in environments with complex demand and the potential for collusion.

A Appendix. Theorem about singleton

1. Inequalities and notation

Consider several non-decreasing sequences $b_{i,1}, b_{i,2}, \ldots, b_{i,K}$, $i = 1, \ldots, N$. Define the partial sums

$$S_i(k) = \sum_{j=1}^{k} b_{i,j}, \quad k = 1, \ldots, K.$$

Define the global maximum prefix value as

$$F^{=\max_{i=1,\ldots,N;\, k=1,\ldots,K} S_i(k).}$$

A K-allocation $k(1), k(2), \ldots, k(N)$ is a "partitioning" of K in the sense that

$$\sum_{i=1}^{N} k(i) = K. \tag{A1}$$

Let $S_i(k(i))$ denote the value of partial sums of sequence i up to $k(i)$, and define the value of an allocation k as

$$F(k) = \sum_{i=1}^{N} S_i\Big(k(i)\Big). \tag{A2}$$

Theorem 1 (Singleton)**.** *If all $b_{i,\cdot}$ are non-decreasing sequences, then for any K-allocation k the inequality*

$$F(k) \le F$$

holds.

Proof. It is known that for non-decreasing sequences, averages of their prefixes also do not decrease:

$$\frac{S_i(k)}{k} \le \frac{S_i(k+1)}{k+1}. \tag{A3}$$

Substituting K in place of $k+1$ yields

$$\frac{S_i(k)}{k} \le \frac{S_i(K)}{K}. \tag{A4}$$

Because of (A4) and the bound $S_i(K) \le F$ we have

$$S_i(k) \le \frac{k}{K} F^{\cdot}\text{(A5)}$$

Now substituting (A5) into (A2), and using equality (A1), we get

$$F(k) = \sum_{i=1}^{N} S_i(k(i)) \le \sum_{i=1}^{N} \frac{k(i)}{K} F^{= \frac{1}{K}\left(\sum_{i=1}^{N} k(i)\right) F = 1 \cdot F \cdot}$$

End of proof. □ □

From the proven inequality $F(k) \le F$ it follows that the best allocation always could be the set with only one element—the unit set or singleton.

Acknowledgments

The author is grateful to Emmanuel Livshits for valuable discussions and for assistance in proving the main theorem.

References

1. Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1), 8–37.

2. Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 11(1), 17–33.

3. Groves, T. (1973). Incentives in teams. *Econometrica*, 41(4), 617–631.

4. Milgrom, P. R., & Weber, R. J. (1982). A Theory of Auctions and Competitive Bidding. *Econometrica*, 50(5), 1089–1122.

5. Ausubel, L. M., & Milgrom, P. R. (2006). The Lovely but Lonely Vickrey Auction. In Cramton, Shoham, & Steinberg (Eds.), *Combinatorial Auctions*. MIT Press.

6. Haeringer, G. (2021). *Market Design: Auctions and Matching*. MIT Press.

7. Taubman, D. (2024). Direct Fractional Auction: A Mechanism for Fractional Ownership of Indivisible Assets. arXiv:2411.11606.