THE LIMITS OF AUCTIONS UNDER *EX-ANTE* COLLUSION

RAVI JAGADEESAN, ILYA SEGAL, AND ALEX TORDJMAN

Abstract. We study revenue-maximizing auction design when bidders can form a bidding ring *ex ante*—i.e., before learning their values and before each bidder decides whether to participate in the auction. We show that bidders' *ex-post* private information about values does not preclude them from colluding efficiently, and as a result, the seller cannot achieve a higher revenue than from posting an optimal price and rationing the good. However, a posted-price mechanism leaves the seller vulnerable to *inefficient* collusion. We construct mechanisms that guarantee the non-collusive revenue of an optimal posted price against arbitrary collusion, even if a subset of bidders can form a bidding ring *ex ante* that can observe participants' values and commit to punish bidders who refuse to join. With symmetric bidders, this guarantee can be achieved by a mechanism that simply combines a posted price with a first-price knockout auction conducted by the seller on behalf of the bidders.

1. Introduction

Standard auctions are vulnerable to collusion. For example, in a second-price or ascending auction, members of a bidding ring may run a knockout auction to determine who bids in the final auction. This possibility lowers revenue by reducing price competition among bidders.[1] A large theoretical literature, initiated by Laffont and Martimort (1997) and Che and Kim (2006), has analyzed how to design auctions that are robust to short-term collusion schemes in which rings entice bidders who know their own private values to collude.

In practice, however, many bidding rings are long-term schemes, which bidders may enter before learning their values for a particular auction. For example, a bidding ring may form for repeated auctions, and adhere to similar collusion schemes across many auctions (Pesendorfer, 2000; Asker, 2010).

This paper examines the revenue limits of all auction formats when such long-term collusion schemes are possible. We capture the long-term impact of repeated

Date: May 6, 2025.

Jagadeesan: Department of Economics, Stanford University. Email: ravi.jagadeesan@gmail.com.
Segal: Department of Economics, Stanford University. Email: ilya.segal@stanford.edu.
Tordjman: Department of Economics, Stanford University. Email: tordjman@stanford.edu.

We thank Paul Milgrom, Mike Ostrovsky, Alex Teytelboym, and several seminar participants for helpful comments.

[1]See Graham and Marshall (1987), Mailath and Zemsky (1991) and McAfee and McMillan (1992).

collusion via a static model in which a bidding ring can be formed *ex ante*—i.e., before bidders learn their values and before they decide whether to participate in the auction. Intuitively, a bidding ring sets up a scheme by which bidders will collude in a sequence of many auctions with similar formats.

We consider the revenue-maximization problem of a seller with one good to sell to one of several bidders, who have independent private values. If a bidding ring is formed, it can dictate the messages that its members send to the seller, and then reallocate the good procured in the auction and implement (budget-balanced) side payments amongst its members. However, once bidders learn their values, the ring is subject to asymmetric information about values analogous to the asymmetric information problem faced by the seller (Laffont and Martimort, 1997; Che and Kim, 2006).

Our results show that *ex-ante* collusion sharply limits the revenue that any selling mechanism can obtain: the seller's optimal revenue guarantee under collusion is that of an optimal posted price. Under a posted price, the good is randomly rationed among the bidders who are willing to pay the price, so the seller is unable to extract any revenue from competition between bidders.

We develop our analysis in three parts. First, we show that the grand coalition of bidders is able to collude efficiently despite its members' private information, and if that happens, then a posted price is optimal to the seller. Second, we show that posting a price leaves the seller vulnerable to *inefficient* collusion. Last, we construct selling mechanisms that guarantee the (non-collusive) revenue of the optimal posted price against all forms of *ex-ante* collusion, including rings comprised of subsets of bidders, and including rings that can commit to punish bidders who refuse to join. When bidders are symmetric, our robust mechanism simply combines a posted price with a first-price knockout auction that the seller conducts on behalf of the bidders.

For the first part, we suppose that the grand coalition of bidders colludes according to a collusion scheme that is *ex-ante* Pareto-optimal subject to (1) the ring's asymmetric information constraints, and (2) the bidders' *ex-ante* participation constraints. Our first main result shows that for any selling mechanism, the revenue it delivers under any efficient collusion scheme is at most the (non-collusive) revenue of the optimal posted price. Moreover, we show that the optimal posted price achieves its non-collusive revenue under any efficient collusion scheme. Thus, a posted price is optimal.

To prove these results, we first show that the grand coalition, acting as a bidding ring, can overcome its asymmetric information and *ex-ante* participation constraints to collude in an *ex-post* efficient way. Indeed, it can do so by using budget-balanced

d'Aspremont and Gérard-Varet (1979) side payments to incentivize truthful revelation of values, plus lump-sum transfers to make joining the ring *ex ante* acceptable to all bidders.[2] Hence, under any efficient collusion scheme, the grand coalition effectively acts as a single agent that values the good at the highest of the bidders' values. The optimality of a posted price follows since a posted price is optimal for selling to a single agent (even among randomized selling mechanisms).

In contrast to the reserve for optimal auctions without collusion (Myerson, 1981), the optimal posted price depends on the number of bidders. If bidders are symmetric, the optimal posted price is strictly increasing in the number of bidders, and (with multiple bidders) strictly greater than the optimal reserve.[3] More generally, the posted price plays a much more important role than reserves do in the non-collusive model. For example, a standard auction without reserves yields zero revenue under efficient collusion regardless of the number of bidders, in contrast to the non-collusive case, in which adding a bidder is better than adding a reserve (Bulow and Klemperer, 1996).

Next, we turn to the case in which bidders may collude inefficiently. We show that in this case, posting a price is insufficient to guarantee protection against collusion. Intuitively, under a posted price, there are large gains from collusion that allocates the good efficiently rather than randomly. So, a collusion scheme in which bidders capture some of these efficiency gains but sometimes (inefficiently) abstain from buying the good can be *ex-ante* acceptable to all bidders while hurting the seller.

Rather than taking a stance on what forms of collusion are reasonable,[4] we seek to design a selling mechanism that is robust to all possible forms of collusion. While our first main result implies that it is impossible to guarantee more than the optimal non-collusive posted-price revenue, our second main result shows that it is in fact possible to guarantee this revenue with selling mechanisms that are robust to all kinds of *ex-ante* collusion by all coalitions of bidders.

One idea behind the construction is to use the logic of the Collusion-Proofness Principle (Laffont and Martimort, 1997): to deter collusion, the seller can implement efficient collusion on behalf of the bidders.[5] However, this approach may leave the

[2]Mailath and Zemsky (1991) first developed this construction for the context of second-price auctions.
[3]If the seller thinks there is a small chance that collusion will not occur, then they may prefer to run a standard auction with the optimal posted price as a reserve.
[4]In practice, collusion may be inefficient, especially when bidders are asymmetric (Asker, 2010).
[5]Thus, the seller allocates the good to the highest bidder if the good is sold. To deter collusion, the seller must allocate the revenue from the efficiency gain relative to the random rationing entailed by a posted price back to the bidders. In particular, some losing bidders receive payments from the seller.

seller vulnerable to collusion by rings that are smaller than the grand coalition.[6] For example, we show that with three bidders, if we attempted to construct a collusion-robust selling mechanism by combining a posted price with Mailath and Zemsky's (1991) efficient collusion scheme (based on an all-pay knockout auction), the seller would be vulnerable to collusion by smaller rings consisting of pairs of bidders.

To construct selling mechanisms that are robust to collusion by all rings, we combine a posted price with particular alternative efficient collusion schemes. We first focus on the symmetric case, for which we can develop a relatively simple mechanism. We show that combining the posted price with a first-price knockout auction[7] conducted by the seller whose revenue is remitted back to the bidders leads to a mechanism that is robust to *ex-ante* collusion by all rings. Intuitively, the first-price format works well for the seller's knockout auction as the first-price auction is relatively resistant to collusion (Robinson, 1985). While bidders in a first-price auction can still generally collude to lower revenue, the key to our mechanism is that they have no incentive to do so when the revenue is remitted back to them.

With asymmetric bidders, using a first-price knockout auction leads to inefficiency, and hence vulnerability to collusion. Our second mechanism is more complicated, but robust to collusion in general. We combine a posted price with an efficient collusion scheme whose side payments are based on a mechanism proposed by Safronov (2018). (Safronov (2018) sharpened d'Aspremont and Gérard-Varet's (1979) construction of an efficient budget-balanced mechanism to provide coalition-proofness.) However, we need to adapt Safronov's (2018) construction to provide a revenue guarantee, and to allow for nonparticipation in a way that protects the seller from collusion that stipulates that some bidders sometimes not participate in the selling mechanism.

The mechanisms we construct have several other desirable properties. First, they are robust to collusion even if members of a bidding ring have complete information about each others' types. Second, their outcomes are robust to collusion even by rings that can commit to punishing bidders who refuse to join the ring by sending minmax messages to the mechanism (as in Quesada (2005) and Dequiedt (2007)), and so may even make its members worse off but prevent defection by a threat of an even harsher punishment. Third, in addition to protecting the seller from collusion, non-colluding bidders' *ex-ante* utilities are also protected from the possibility of collusion by others

[6] Laffont and Martimort (1997, 2000) restricted attention to settings with two agents, where the only possible colluding coalition is the grand coalition.

[7] McAfee and McMillan (1992) showed that a first-price knockout auction with revenue shared equally among the bidders is an efficient collusion scheme when bidders are symmetric.

(as in Safronov (2018) and Csóka et al. (2024)). Last, the mechanisms are *ex-post* individually rational, so ensure participation security (in the sense of Dasgupta et al. (1979)) even if bidders are concerned about the possibility of others colluding.

The logic behind the robustness of our proposed mechanisms to collusion is based on four properties that the mechanisms share. First, they allocate the good efficiently in equilibrium. Second, in equilibrium, they charge the bidders the optimal posted price in aggregate if the good is sold regardless of who gets the good. Third, they charge at least as much out of equilibrium (e.g., when some bidders do not participate). These three properties are sufficient to protect the seller from collusion by the grand coalition if it cannot punish bidders who refuse to collude. Indeed, the three properties ensure that the grand coalition receives an efficient allocation (taking into account the posted price as a cost), and hence any deviation that alters the allocation would decrease some bidders' *ex ante* utilities. The second property ensures that any deviation that leaves the allocation unchanged cannot lower revenue.

However, these three properties are not sufficient to protect the seller from collusion by rings that may not include all bidders, or from rings that can commit to punishing bidders who refuse to join. Ensuring robustness to those forms of collusion relies on a fourth property of our mechanisms: that each bidder can guarantee their *ex-ante* expected equilibrium payoff by following their equilibrium strategies, regardless of how other bidders behave. This *guaranteed-utility* property ensures that no coalition can make any bidder obtain less than his equilibrium payoff. Hence, under any form of collusion, each bidder must receive at least their equilibrium payoff, implying that we only need to concern ourselves with collusion by the grand coalition.

The two mechanisms we construct satisfy the guaranteed-utility property. For the symmetric case, a rough intuition is that the first-price knockout format protects bidders from the behavior of others: other bidders only affect a bidder through their highest bid, and a bidder is compensated for losing the auction through his share of the knockout revenue. For the general case, we adapt Safronov's (2018) construction to ensure that the guaranteed-utility property holds even when some bidders choose not to participate; this requires appropriately defining participants' outcomes if some bidders do not participate, as well as adding particular lump-sum side transfers.

Related Literature. Our analysis builds on seminal analyses of collusion-proof mechanisms by Laffont and Martimort (1997, 2000) and Che and Kim (2006). Those papers considered *interim* collusion—i.e., where rings are formed after bidders learn their values and decide whether to participate in the auction. Che and Kim (2006) suggested

a mechanism that guarantees the seller's optimal revenue against any collective misbehavior that bidders might engage in after they have agreed to participate in the mechanism. Subsequent work by Pavlov (2008) and Che and Kim (2009) showed that collusion does sometimes limit revenue if bidders can form a ring *before* deciding whether to participate in the auction, but after learning their values. Instead, motivated by long-term collusion schemes, we suppose that rings are formed *ex ante*—i.e., *before* bidders learn their values and decide whether to participate in the auction. We show that collusion then limits the expected revenue that can be guaranteed to the optimal posted-price revenue. The optimal collusion-proof mechanisms that we construct also provide side benefits such as robustness to rings that can punish bidders, protection of non-colluding bidders from collusion, and *ex post* individual rationality.

Our conclusions are more similar to Dequiedt (2007), who studied optimal auctions in the presence of collusion with two bidders and two types. He considered a stronger form of interim collusion in which rings can not only collude on whether to participate in the auction, but can also commit to an arbitrary punishment of defectors who refuse to join the ring at the interim stage.[8] Dequiedt (2007) showed that the seller cannot improve upon a posted price. However, unlike our result, it is unclear if his result generalizes to more agents or types. Moreover, the mechanisms we construct are robust to the type of collusion considered by Dequiedt (2007).

In independent and contemporaneous work, Wu (2025) also studies auction design under *ex-ante* collusion. She develops results that are closely related to our analysis of efficient collusion.[9] She also constructs a selling mechanism by combining the optimal posted price with the efficient collusion scheme constructed by Mailath and Zemsky (1991), and shows that this mechanism is robust to collusion by the grand coalition. However, as we show in Section 4.2, this approach leads to a selling mechanism that is vulnerable to collusion by smaller coalitions. In contrast, our approach guarantees robustness to collusion by all rings of bidders; it also provides additional benefits not obtained by other approaches—such as robustness to rings that can punish bidders, protection of non-colluding bidders from collusion, and *ex-post* individual rationality.[10]

[8]Quesada (2005) studied auctions for procuring complements under a related form of collusion.
[9]Specifically, Wu (2025) shows that the optimal posted price provides the optimal revenue guarantee under worst-case efficient collusion by the grand coalition, if a ring's side payments only need to be budget-balanced *ex ante*. In contrast, we allow for any (i.e., non-worst-case) efficient collusion and restrict to the more standard case of *ex-post* budget-balanced side payments by rings, and show that the optimal posted price provides the optimal revenue.
[10]Wu (2025) also studies optimal collusion-proof mechanisms subject to the restriction of never making payments to bidders, which are beyond our scope.

Methodologically, our construction of mechanisms that are robust to collusion extends Safronov's (2018), Csóka et al.'s (2024), and Csóka et al.'s (2025) use of the guaranteed-utility property along several dimensions. First, we formalize a strong sense in which the guaranteed-utility property ensures robustness to collusion even by rings that may not include all bidders and can commit to punish agents who refuse to join.[11] Second, we show how the approach can be applied to provide a revenue guarantee instead of a balanced budget. Third, we obtain robustness to rings that can reallocate the good. Fourth, we adapt Safronov's (2018) construction to ensure robustness to collusion on participation and satisfy *ex-post* participation constraints. Fifth, we show that in the case of symmetric bidders, a much simpler mechanism, namely one based on a first-price knockout auction, also has the guaranteed-utility property and is therefore robust to collusion.

Paper Outline. Section 2 introduces our model. Section 3 shows that a posted price is optimal under efficient collusion. Section 4 illustrates difficulties in ensuring robustness to more general collusion. Section 5 formalizes robustness to collusion by rings of bidders that can commit to punish deviators. Section 6 constructs our collusion-robust mechanisms. Section 7 is a conclusion. The appendix presents omitted proofs.

2. Model

A seller has one unit of a good to sell to one of a finite set I of bidders i. Thus, the space of feasible allocations of the good is given by $\mathcal{Q} = \left\{ q \in \mathbb{R}^I_{\geq 0} \,\middle|\, \sum_{i \in I} q_i \leq 1 \right\}$. The bidders have independent, private values for the good. Bidder i's value θ_i is drawn from a continuous distribution F_i supported on a compact set $\Theta_i \subseteq \mathbb{R}_{\geq 0}$.

The game proceeds as follows. First, the seller announces a mechanism, which must include a "nonparticipation message" for each bidder. Second, an intermediary (or any bidder[12]) proposes a collusion scheme, whereby bidders can agree about how to manipulate their reports to the selling mechanism—including whether to send a nonparticipation message—as well as whether or how to reallocate the good and exchange side payments. Bidders commit to the collusion scheme before learning their values.

[11] Segal and Whinston (2012, Proposition 1), Safronov (2018, Theorem 2 and 3), Csóka et al. (2024, Proposition 2), and (Csóka et al., 2025, Section 3.4) formalized weaker senses in which the guaranteed-utility property ensures robustness to equilibrium selection and collusion.

[12] Unlike in Laffont and Martimort (1997) and Che and Kim (2006), it does not matter whether an intermediary proposes the collusion scheme or one of the bidders proposes it, as none of the bidders have learned their value at the time at which the collusion scheme is proposed.

For each bidder, the seller can set an arbitrary (measurable) message space M_i that includes a *nonparticipation message* NP_i. Letting $M = \times_{i\in I} M_i$ denote the space of message profiles, the seller also sets a (bounded, measurable) outcome function $(q,t) : M \to \mathcal{Q} \times \mathbb{R}^I$ such that $q_i(NP_i, m_{-i}) = t_i(NP_i, m_{-i}) = 0$ for all message profiles $m_{-i} \in M_{-i}$. Thus, the nonparticipation message must guarantee a bidder an allocation and payment of 0. To make the bidding ring's problem well-behaved, we impose the technical restriction that the space

$$\left\{ \left(\sum_{i\in I} q_i(m), \sum_{i\in I} t_i(m) \right) \middle| \, m \in M \right\}$$

of pairs of aggregate allocations and aggregate transfers obtainable by any message profile m be closed (Pavlov, 2008, Condition 2).[13]

A *selling mechanism* consists of message spaces M_i for each bidder, an outcome function $(q,t) : M \to \mathcal{Q} \times \mathbb{R}^I$, and strategies $\sigma_i : \Theta_i \to M_i$.[14] If all bidders follow the strategies prescribed by the selling mechanism, then bidder i's *ex ante* expected utility is $V_i(M, \sigma, q, t) = \mathbb{E}_{\tilde{\theta}}\left[\tilde{\theta}_i q_i(\sigma(\tilde{\theta})) - t_i(\sigma(\tilde{\theta}))\right]$ where we write $\sigma(\theta) = (\sigma_i(\theta_i))_{i\in I} \in M$.

For collusion schemes, by the Revelation Principle, it is without loss of generality to restrict attention to direct mechanisms (Laffont and Martimort, 1997). Rather than directly specifying the reallocation and side payments by a bidding ring, following Che and Kim (2006), we consider the allocations $\hat{q}(\theta)$ and transfers $\hat{t}(\theta)$ induced by the composite of the reallocation and side payments of the ring with the allocation and transfer of the submitted reports $\mu(\theta)$ to the selling mechanism.

Definition 1 (Laffont and Martimort, 1997; Che and Kim, 2006)**.** A *collusion scheme* for a selling mechanism (M, q, t, σ) consists of a report function $\mu : \Theta \to \Delta(M)$ and a direct mechanism[15] $(\hat{q}, \hat{t}) : \Theta \to \mathcal{Q} \times \mathbb{R}^I$ such that for all type profiles $\theta \in \Theta$, we have

$$\sum_{i\in I} \hat{q}_i(\theta) = \mathbb{E}_{\tilde{m}\sim\mu(\theta)}\left[\sum_{i\in I} q_i(\tilde{m}) \right] \qquad \text{(feasibility of reallocation of goods)} \tag{1}$$

$$\sum_{i\in I} \hat{t}_i(\theta) = \mathbb{E}_{\tilde{m}\sim\mu(\theta)}\left[\sum_{i\in I} t_i(\tilde{m}) \right] \qquad \text{(budget balance of side payments).} \tag{2}$$

In light of (2), expected revenue under a collusion scheme $(\mu, \hat{q}, \hat{t})$ is given by

$$\mathbb{E}_{\tilde{\theta}}\left[\mathbb{E}_{\tilde{m}\sim\mu(\tilde{\theta})}\left[\sum_{i\in I} t_i(\tilde{m}) \right] \right] = \mathbb{E}_{\tilde{\theta}}\left[\sum_{i\in I} \hat{t}_i(\tilde{\theta}) \right].$$

[13] Any outcome function can be approximated arbitrarily well by ones satisfying this restriction.
[14] The restriction to pure strategies is without loss of generality, as the message space can be expanded to include stochastic reports over pure messages. We specify strategies so a selling mechanism gives a well-defined non-collusive outside option to which a bidder can compare a proposed collusion scheme.
[15] The allocation and payment rules of direct mechanisms are required to be bounded and measurable.

If bidders collude, they are bound to send the message recommended by the bidding ring, and to follow through with the specified reallocation of the good and transfers.[16] In the interpretation of long-term collusion schemes in repeated auctions, bidders may be deterred from such deviations with the threat of punishments in future periods.

Our analysis of efficient collusion places two additional restrictions on the feasible forms of collusion. Imposing more constraints on collusion makes our first result on the limits of auctions stronger. (The mechanisms we construct in Section 6 are robust to more general forms of collusion that we introduce in Section 5.)

First, we suppose that all bidders must agree to a proposed collusion scheme for a bidding ring to be formed. Thus, bidders accept a proposed collusion scheme if and only if their resulting *ex ante* expected utility $U_i(\hat{q},\hat{t}) = \mathbb{E}_{\tilde{\theta}}\big[\tilde{\theta}_i\hat{q}_i(\tilde{\theta}) - \hat{t}_i(\tilde{\theta})\big]$ exceeds their *ex ante* expected utility $V_i(M,\sigma,q,t)$ under the selling mechanism.

Definition 2. A collusion scheme $(\mu,\hat{q},\hat{t})$ for a selling mechanism (M,q,t,σ) is *ex-ante acceptable* if for all bidders i, we have $U_i(\hat{q},\hat{t}) \geq V_i(M,\sigma,q,t)$.

Second, we suppose that after bidders learn their values, they (like the seller) may be subject to asymmetric information about others' values (Laffont and Martimort, 1997; Che and Kim, 2006). Thus, bidders can affect the message recommendations and reallocations by manipulating their type reports to the bidding ring. Hence, we restrict attentions to collusion schemes that are Bayesian incentive compatible.

Definition 3. A collusion scheme $(\mu,\hat{q},\hat{t})$ (for any selling mechanism) is *incentive-compatible* if for all bidders i, types $\theta_i \in \Theta_i$, and reports $\theta_i' \in \Theta_i$, we have

$$\mathbb{E}_{\tilde{\theta}_{-i}}\big[\theta_i\hat{q}_i(\theta_i,\tilde{\theta}_{-i}) - \hat{t}_i(\theta_i,\tilde{\theta}_{-i})\big] \geq \mathbb{E}_{\tilde{\theta}_{-i}}\big[\theta_i\hat{q}_i(\theta_i',\tilde{\theta}_{-i}) - \hat{t}_i(\theta_i',\tilde{\theta}_{-i})\big].$$

3. Optimality of a Posted Price under Efficient Collusion

To define efficient collusion, we compare collusion schemes $(\mu,\hat{q},\hat{t})$ based on the delivered profile of *ex-ante* utilities $(U_i(\mu,\hat{q},\hat{t}))_{i\in I}$. Our next definition introduces a notion of efficiency subject to the constraints that (1) the grand coalition must make collusion acceptable for each bidder at the *ex ante* stage, and (2) there is asymmetric information about the bidders' values *ex post*.

Definition 4. An *ex-ante* acceptable, incentive-compatible collusion scheme for a selling mechanism is *ex ante incentive-efficient* if it is *ex-ante* Pareto-optimal over all *ex-ante* acceptable, incentive-compatible collusion schemes for that selling mechanism.

[16] Laffont and Martimort (1997) and Che and Kim (2006) made the same assumption.

3.1. **Efficient Collusion.** To analyze auction design under *ex-ante* incentive-efficient collusion, we first show that the grand coalition can overcome its acceptability constraint and asymmetric information to achieve efficient collusive outcomes. Formally, we say that a collusion scheme $(\mu, \hat{q}, \hat{t})$ for a selling mechanism is *efficient* if it is *ex-ante* Pareto-optimal over all collusion schemes for the selling mechanism.

Proposition 1. *Every ex-ante incentive-efficient collusion scheme is efficient.*

To prove Proposition 1, we construct an efficient, *ex ante* acceptable, incentive-compatible collusion scheme by extending the arguments of Mailath and Zemsky (1991) to general selling mechanisms.[17] The formal argument is in the appendix.

Collusion being *ex-ante* is crucial to Proposition 1. At the interim stage, without having decided on a collusion scheme *ex ante*, the grand coalition cannot generally overcome acceptability constraints and asymmetric information to obtain efficient collusive outcomes (Laffont and Martimort, 1997; Che and Kim, 2006), even for selling mechanisms with nonparticipation messages (Pavlov, 2008; Che and Kim, 2009).

3.2. **Optimality of a Posted Price.** Our main result on auction design under *ex-ante* incentive-efficient collusion shows the optimality of a posted-price mechanism. A p-posted-price mechanism allocates the good randomly among those who are willing to pay p (if any), charging p to the bidder who receives the good. Each bidder only needs to express whether they would like to buy the good (i.e., send participation message P_i) or not (i.e., send nonparticipation message NP_i). Formally, the *p-posted-price mechanism* is defined by binary message spaces $M_i = \{P_i, NP_i\}$, allocation rule

$$q_i(m) = \begin{cases} 0 & \text{if } m_i = NP_i \\ \frac{1}{|\{i \,|\, m_i = P_i\}|} & \text{if } m_i = P_i \end{cases},$$

payment rule $t_i(m) = pq_i(m)$, and (dominant) strategies

$$\sigma_i(\theta_i) = \begin{cases} P_i & \text{if } \theta_i \geq p \\ NP_i & \text{if } \theta_i < p \end{cases}.$$

The revenue under the prescribed strategies is $\Pi(p) = p \Pr_{\tilde{\theta}}[\max_{i \in I} \tilde{\theta}_i \geq p]$.

Let $p^* \in \arg\max_p \Pi(p)$ be an optimal posted price. Our first main result shows that if bidders collude using any *ex-ante* incentive-efficient collusion scheme, posting a price of p^* is optimal. Formally, we bound the revenue of any mechanism under any *ex-ante*

[17] This scheme is *ex-ante* incentive-efficient by construction, so in particular an *ex-ante* incentive-efficient collusion scheme exists. Hence, Proposition 1, and Theorem 1, below are not vacuous.

incentive-efficient collusion scheme by $\Pi(p^*)$, and show that a p^*-posted-price mechanism yields revenue $\Pi(p^*)$ under any *ex-ante* incentive-efficient collusion scheme.

Theorem 1. *(a) For any selling mechanism, expected revenue under any ex-ante incentive-efficient collusion scheme is at most* $\Pi(p^*)$.

(b) All ex-ante incentive-efficient collusion schemes for the p^**-posted-price mechanism yield expected revenue* $\Pi(p^*)$.[18]

The intuition behind Theorem 1 is that since *ex-ante* incentive-efficient collusion allows the grand coalition to collude efficiently, it is as if the seller faces a single bidder whose value is the highest of the bidders' values. The ability of the grand coalition to send nonparticipation messages ensures *ex-post* individual rationality for the coalition. In turn, in a single-buyer setting with an individual rationality constraint, a posted-price is optimal, and (with continuous value distribution) provides a full implementation of the revenue-maximizing allocation.

3.3. **Proof of Theorem 1.** To prove Theorem 1, we use the following simple characterization of the messages that can be sent under efficient collusion schemes: these messages must maximize the surplus of the grand coalition, valuing the aggregate allocation at the highest bidder's value and accounting for aggregate transfers.

Lemma 1. *If* $(\mu, \hat{q}, \hat{t})$ *is an efficient collusion scheme for a selling mechanism* (M, q, t, σ)*, then the report function* μ *must satisfy*

$$\mu(\theta) \in \arg\max_{\lambda \in \Delta(M)} \left\{ \left(\max_{i \in I} \theta_i \right) \mathbb{E}_{\tilde{m} \sim \lambda} \left[\sum_{i \in I} q_i(\tilde{m}) \right] - \mathbb{E}_{\tilde{m} \sim \lambda} \left[\sum_{i \in I} t_i(\tilde{m}) \right] \right\} \quad \text{almost surely.} \tag{3}$$

Proof. As utility is transferable, efficiency of the collusion scheme implies that

$$\left(\mu(\theta), \hat{q}(\theta)\right) \in \arg\max_{(\lambda, q) \in \Delta(M) \times \mathcal{Q} \,\big|\, \sum_{i \in I} q_i = \mathbb{E}_{\tilde{m} \sim \lambda}\left[\sum_{i \in I} q_i(m)\right]} \left\{ \sum_{i \in I} \theta_i q_i - \mathbb{E}_{\tilde{m} \sim \lambda} \left[\sum_{i \in I} t_i(\tilde{m}) \right] \right\} \text{ almost surely.}$$

Substituting an optimal choice of q for any fixed λ yields (3), as the ring maximizes its total surplus by allocating $\mathbb{E}_{\tilde{m} \sim \lambda}\left[\sum_{i \in I} q_i(m)\right]$ to bidders with the highest value. □

To complete the proof of Theorem 1, we can then apply single-buyer arguments.

Proof of Theorem 1. By Proposition 1, it suffices to consider efficient collusion schemes.

First, consider any efficient collusion scheme $(\mu, \hat{q}, \hat{t})$ for any selling mechanism (M, q, t, σ). The presence of the message profile NP ensures the maximum in (3) is

[18] The proof shows more generally that for any posted price p, all *ex-ante* incentive-efficient collusion schemes for the p-posted-price mechanism yield expected revenue $\Pi(p)$.

nonnegative. Lemma 1 thus implies that the collusive aggregate outcome function $\left(\sum_{i\in I}\hat{q}_i(\mu(\theta)), \sum_{i\in I}\hat{t}_i(\mu(\theta))\right)$ defines a mechanism for selling to a single buyer with value $\max_{i\in I}\theta_i$ that is almost surely incentive compatible and individually rational. As a posted price is optimal for selling to one buyer (Myerson, 1981; Riley and Zeckhauser, 1983), revenue under the collusion scheme is at most $\Pi(p^*)$.

Now suppose that the mechanism (M, q, t, σ) is the p^*-posted-price mechanism. By Lemma 1 and the definition of posted-price mechanisms, we have

$$\mu(\theta) \in \arg\max_{\lambda\in\Delta(M)}\left\{\left(\max_{i\in I}\theta_i - p^*\right)\mathbb{E}_{\tilde{m}\sim\lambda}\left[\sum_{i\in I}q_i(\tilde{m})\right]\right\} \quad \text{almost surely.}$$

For $\max_{i\in I}\theta_i > p^*$ (resp. $\max_{i\in I}\theta_i < p^*$), this requires almost surely selecting a report distribution $\mu(\theta)$ that makes $\mathbb{E}_{\tilde{m}\sim\mu(\theta)}\left[\sum_{i\in I}q_i(\tilde{m})\right] = 1$ (resp. $\mathbb{E}_{\tilde{m}\sim\mu(\theta)}\left[\sum_{i\in I}q_i(\tilde{m})\right] = 0$). Hence, expected revenue under the collusion scheme is

$$\mathbb{E}_\theta\left[\mathbb{E}_{\tilde{m}\sim\mu(\theta)}\left[\sum_{i\in I}t_i(\tilde{m})\right]\right] = p^*\mathbb{E}_\theta\left[\mathbb{E}_{\tilde{m}\sim\mu(\theta)}\left[\sum_{i\in I}q_i(\tilde{m})\right]\right] = p^*\mathbb{E}_\theta\left[\mathbf{1}\left(\max_{i\in I}\theta_i \geq p^*\right)\right] = \Pi(p^*),$$

as claimed. □

3.4. **Properties of Optimal Posted Prices.** The role of a posted price is not analogous to the role of a reserve in standard auction theory. Indeed, unlike Myerson's (1981) optimal reserve, the optimal posted price depends on the number of bidders, and is strictly increasing in the number of bidders in the symmetric case.

Proposition 2. *If values are i.i.d from a distribution F supported on a closed interval with a continuous positive density, then any selection of optimal posted prices p_N^* for N bidders is strictly increasing in N.*

The intuition for this result is that (efficient) collusion effectively destroys competition between the bidders, forcing the seller to raise revenue in a less efficient way. It follows that the posted price exceeds the Myerson optimal reserve.

Corollary 1. *Under the hypotheses of Proposition 2, if F is regular and $N > 1$, then any optimal posted price is strictly higher than the Myerson (1981) optimal reserve.*

Proof. As shown by Myerson (1981), the optimal reserve is the optimal posted price for $N = 1$. Hence, the result follows from Proposition 2. □

More generally, the decrease in competition entailed by collusion means that a posted price in our model is more important than the usual reserve. For example, under efficient collusion, standard auctions without reserves raise zero revenue regardless

of the number of bidders, while a posted price raises positive revenue. Hence, unlike in Bulow and Klemperer (1996), competition between bidders does not substitute for the use of a better mechanism when bidders can collude *ex ante*.

4. Difficulties in Ensuring Robustness to Collusion

In this section, we provide three examples that illustrate issues with implementing a posted price in a way that is robust to collusion. Our first example shows that posted-price mechanisms on their own are vulnerable to inefficient collusion. Thus, a natural strategy to ensure robustness to collusion would be to combine the posted-price mechanism with an efficient collusion scheme. In our second example, we show that this strategy does not always work, as it can leave the seller vulnerable to collusion by smaller rings than the grand coalition. Our last example shows that the seller can also be left vulnerable to collusion by the grand coalition if a bidding ring can commit to punishing bidders who refuse to join by sending nonparticipation messages.

4.1. Vulnerability of Posted Prices to Inefficient Collusion.

It turns out that posted-price mechanisms leave the seller vulnerable to inefficient collusion, even via *ex-ante* acceptable, incentive-compatible collusion schemes.

Example 1. There are two bidders whose values are each uniformly distributed between 1 and 2. Consider the posted-price mechanism with optimal posted price $p^* = \frac{4}{3}$. Each bidder's *ex ante* utility is $V_i(M, q, t, \sigma) = \frac{4}{27}$, and expected revenue is $\Pi(p^*) = \frac{32}{27}$.

There is a way for bidders to collude to buy the good only if the higher bidder's value clears $\hat{p} = \frac{3}{2}$, thereby decreasing the probability the good is sold and hence lowering expected revenue. The collusion scheme we construct is incentive-compatible, and allocates the good efficiently among the bidders if the good is sold (as the lower bidder never participates). While it is (incentive-)inefficient, it improves the efficiency of allocation across bidders enough relative to rationing to obtain *ex ante* acceptability.

Formally, consider the collusion scheme with deterministic report function

$$\mu_i(\theta) = \begin{cases} P_i & \text{if } \theta_i \geq \hat{p} \text{ and } \theta_i \geq \theta_{-i} \\ NP_i & \text{otherwise} \end{cases},$$

allocation rule $\hat{q}_i(\theta) = q_i(\mu(\theta))$, and payment rule

$$\hat{t}_i(\theta) = p^* \hat{q}_i(\theta) + \Big((2 - \theta_i)(\hat{p} - p^*) + \tfrac{1}{2}(\theta_i - \hat{p})^2\Big)\mathbf{1}[\theta_i > \hat{p}] - \Big((2 - \theta_{-i})(\hat{p} - p^*) + \tfrac{1}{2}(\theta_{-i} - \hat{p})^2\Big)\mathbf{1}[\theta_{-i} > \hat{p}].$$

(1) and (2) hold by construction. The collusion scheme is incentive-compatible.[19] It delivers an *ex ante* utility of $U_i(\hat{q},\hat{t}) = \frac{1}{6} > \frac{4}{27} = V_i(M,q,t,\sigma)$ to each bidder, so is *ex ante* acceptable. But it leads to an expected revenue of $1 < \frac{32}{27} = \Pi(p^*)$.

In practice, collusion is not always efficient, especially with asymmetric bidders (Asker, 2010). And even so, a seller may want to protect themselves from the possibility of inefficient collusion. Hence, in the remainder of the paper, we seek to design mechanisms that guarantee the same revenue as the non-collusive revenue of the optimal posted price even in the presence of inefficient collusion.

4.2. **Vulnerability to Collusion by Smaller Rings.** Our overall strategy to construct mechanisms that are robust to collusion is to use the logic of the Collusion-Proofness Principle: to deter bidders from colluding, we combine the posted price with a specified efficient collusion scheme implemented on behalf of the bidders.

One natural approach would be to use the efficient collusion scheme proposed by Mailath and Zemsky (1991), as that collusion scheme is efficient with both symmetric and asymmetric bidders.[20] However, when there are more than two bidders, this approach actually leads to a mechanism that is not robust to all forms of collusion. The problem is that this efficient collusion scheme is not itself robust to collusion.

Indeed, Mailath and Zemsky's (1991) scheme asks bidders to bid for the right to buy the good for the posted price in an all-pay knockout auction whose revenue is remitted back to the bidders and divided equally among them. Hence, with three bidders, in the selling mechanism defined by the combination of a posted price with Mailath and Zemsky's (1991) collusion scheme, a pair of bidders can collude to make themselves both better off by running a private knockout auction to determine who participates in the seller's knockout auction, thereby saving the side payment the loser of the seller's knockout auction would have made to the third bidder.[21] Along the way, the ring can end up lowering the probability the good is sold, thereby lowering revenue.

Example 2. There are three bidders, whose values are each uniformly distributed between 0 and 1. The optimal posted price is $p^* = \frac{1}{\sqrt[3]{4}} \approx 0.63$.

Mailath and Zemsky's (1991) efficient collusion scheme would allocate the good to the highest bidder, provided that his value exceeds p^*. In the all-pay knockout

[19]Indeed, the allocation rule is monotone and expected payments satisfy an Envelope Formula.

[20]Graham and Marshall (1987) and McAfee and McMillan (1992) assumed bidders are symmetric. As Graham and Marshall's (1987) scheme is only budget-balanced *ex ante*, combining a posted price with it would leave the seller vulnerable even to efficient collusion by the grand coalition.

[21]This vulnerability is closely related to the vulnerability to collusion of the standard efficient partnership dissolution mechanism described by Cramton et al. (1987, page 624).

auction, the equilibrium bid of type θ_i is

$$g(\theta_i) = (\theta_i - p^*)^2 \frac{2\theta_i + p^*}{3} \mathbf{1}[\theta_i \geq p^*],$$

making bidder i's side payment be given by

$$g(\theta_i) - \frac{1}{2} \sum_{j \in I \setminus \{i\}} g(\theta_j).$$

Combining this collusion scheme with the posted-price mechanism gives a selling mechanism with message spaces $M_i = [0,1] \cup \{NP_i\}$, payment rule[22]

$$t_i(\theta) = p^* q_i(\theta) + g(\theta_i) - \frac{1}{2} \sum_{j \in I \setminus \{i\}} g(\theta_j),$$

and strategies $\sigma_i(\theta_i) = \theta_i$. For this mechanism, under truthful revelation, expected revenue is $\Pi(p^*) \approx 0.47$. Moreover, any *ex-ante* acceptable, incentive-compatible collusion scheme yields expected revenue $\Pi(p^*)$.

However, the mechanism leaves the seller vulnerable to collusion by rings consisting of pairs of bidders. Specifically, we construct a collusion scheme in which bidders 1 and 2 collude so that the lower-value bidder between them bids zero and the higher-value bidder reports truthfully provided that their value exceeds $\hat{p} = \frac{2}{3} > p^*$, and reports zero otherwise. This results in an efficient allocation of the good between bidders 1 and 2 and therefore they do not reallocate it. We construct collusive side payments that ensure incentive compatibility of the scheme for bidders 1 and 2, assuming that bidder 3 does not know about collusion and therefore reports truthfully.[23] Under this collusion scheme, bidders 1 and 2 each obtain an *ex-ante* expected utility of ≈ 0.0535, which exceeds their *ex-ante* expected utility $V_i(M, q, t, \sigma) \approx 0.0531$ from reporting truthfully. The collusion scheme lowers the expected revenue to $\approx 0.45 < \Pi(p^*)$.

Our constructions ensure robustness to collusion by all rings by combining the posted price with other collusion schemes that are themselves more robust to collusion. Rather than using Mailath and Zemsky's (1991) all-pay collusion scheme, we use McAfee and McMillan's (1992) first-price collusion scheme for the symmetric case, and side payments based on an idea of Safronov (2018) for the general case.

[22]The mechanism must also specify what happens when some agents do not participate, which is not material for this example, but important for the next one.

[23]The side payment from bidder 1 to bidder 2 is

$$\Big(\frac{2\theta_1^3 + \hat{p}^3}{3} - \theta_1^2 p^* - \frac{\theta_1 + \mathbf{1}[\theta_1 > \theta_2]}{2} g(\theta_1)\Big)\mathbf{1}[\theta_1 > \hat{p}] - \Big(\frac{2\theta_2^3 + \hat{p}^3}{3} - \theta_2^2 p^* - \frac{\theta_2 + \mathbf{1}[\theta_2 > \theta_1]}{2} g(\theta_2)\Big)\mathbf{1}[\theta_2 > \hat{p}].$$

4.3. **Vulnerability to Collusion by Rings that Commit to Punish.** Another issue is that even when we restrict attention to collusion that involves all the bidders, inefficient collusion may be enforced by the ring committing to punish bidders who refuse to join the ring, as in Quesada (2005) and Dequiedt (2007).

Example 3. Consider again the two-bidder setting of Example 1. Consider the direct mechanism that combines the optimal posted price $p^* = \frac{4}{3}$ with Mailath and Zemsky's (1991) efficient collusion scheme. This mechanism allocates to the highest bidder, provided that his value exceeds p^*. Its payment rule is given by

$$t_i(\theta) = p^* q_i(\theta) + \frac{1}{2}(\theta_i - p^*)^2 \mathbf{1}[\theta_i \geq p^*] - \frac{1}{2}(\theta_{-i} - p^*)^2 \mathbf{1}[\theta_{-i} \geq p^*].$$

Each agent's *ex-ante* utility under truthtelling is ≈ 0.173.

Now suppose that if either bidder sends a nonparticipation message, the seller implements the null outcome. Then, inefficient collusion can be sustained by one agent threatening the other agent with nonparticipation. For example, agents can sustain the (incentive-compatible) collusion scheme in which agent 2 always reports 1, agent 1 reports $\frac{5}{3}$ when $\theta_1 > \frac{25}{18}$ and reports 1 otherwise, and which has no reallocation or side transfers. This scheme gives agent 1 an *ex-ante* expected utility of $\approx 0.187 > 0.173$. The collusion scheme gives agent 2 a positive expected payment of $\frac{1}{54}$, and he is induced to participate in it by agent 1's threat of not participating in the selling mechanism—thereby canceling it—if collusion does not occur. Note that this collusion scheme reduces the expected revenue to $\frac{22}{27} < \Pi(p^*) = \frac{32}{27}$.

This example demonstrates, in particular, that to avoid the possibility that bidders force other bidders to collude using a threat of nonparticipation, outcomes following nonparticipation by some agents must be defined more carefully.

5. Strong Collusion Robustness

In this section, we introduce a desirideratum, termed *strong collusion robustness*, that requires robustness to a broad form of collusion by rings that can be comprised of arbitrary subsets of bidders that can commit to punish defectors in arbitrary ways. The mechanisms that we construct will satisfy this desideratum.

5.1. **Modeling Rings and Punishments.** We first define a concept of collusion by a ring $R \subseteq I$ that may not include all bidders. The ring can use any collusion scheme that only manipulates the reports of ring members, and involves reallocation of the

good and side payments only among ring members. (Bidders who are not in the ring follow the strategies prescribed by the selling mechanism.)

Definition 5. A collusion scheme $(\mu, \hat{q}, \hat{t})$ for selling mechanism (M, q, t, σ) is a *collusion scheme for ring* $R \subseteq I$ if there exists $\mu^R : \Theta^R \to \Delta(M^R)$ such that for all type profiles $\theta \in \Theta$, we have $\mu(\theta) = \mu^R(\theta_R) \times \sigma_{I \setminus R}(\theta_{I \setminus R})$,[24] and

$$\big(\hat{q}_i(\theta), \hat{t}_i(\theta)\big) = \mathbb{E}_{\tilde{m} \sim \mu(\theta)} \big[\big(q_i(\tilde{m}), t_i(\tilde{m})\big)\big] \quad \text{for all } i \in I \setminus R. \tag{4}$$

This notion of collusion allows that a ring involves only a subset of bidders. Moreover, a ring can observe its participants' types, so as not face incentive constraints among members (unlike in Example 2), and can commit to punishing a potential member i who refuses to join with an arbitrary reporting plan by the remaining ring members $j \in R \setminus \{i\}$ (the only restriction is that they cannot observe i's type).

We now define a notion of individual rationality at the *ex ante* stage for a collusion scheme for a ring $R \subseteq I$.[25]

Definition 6. A collusion scheme $(\mu, \hat{q}, \hat{t})$ for a ring $R \subseteq I$ is *individually rational* if for each $i \in R$ there exists a (measurable) punishment scheme $\nu^{i,R} : \Theta_{R \setminus \{i\}} \to \Delta(M_{R \setminus \{i\}})$ such that for all strategies $\sigma_i' : \Theta_i \to M_i$, we have

$$\mathbb{E}_{\tilde{\theta}}\big[\tilde{\theta}_i \hat{q}_i(\tilde{\theta}) - \hat{t}_i(\tilde{\theta})\big] \geq \mathbb{E}_{\tilde{\theta}}\big[\mathbb{E}_{\tilde{m}_{I \setminus \{i\}} \sim \nu^i(\tilde{\theta}_{I \setminus \{i\}})}\big[\tilde{\theta}_i q_i\big(\tilde{m}_{I \setminus \{i\}}, \sigma_i'(\tilde{\theta}_i)\big) - t_i\big(\tilde{m}_{I \setminus \{i\}}, \sigma_i'(\tilde{\theta}_i)\big)\big]\big], \tag{5}$$

where we write $\nu^i(\theta_{I \setminus \{i\}}) = \nu^{i,R}(\theta_{R \setminus \{i\}}) \times \sigma_{I \setminus R}(\theta_{I \setminus R})$.

In words, this says that a collusion scheme for a ring $R \subseteq I$ is individually rational if, at the *ex ante* stage, for any agent $i \in R$, the remaining ring members $j \in R \setminus \{i\}$ can ensure that i will find it in his interest to participate in the collusion scheme by resorting to punishments that uniformly hurt agent i in case he does not join the ring, that is for any strategy i might play in the selling mechanism without colluding.[26,27]

Note that a ring can induce all agents who are aware of it to join, as the best responses of bidders who are aware of the ring's existence would be individually

[24] By abuse of notation, we also let $\sigma_i(\theta_i)$ denote a point mass at $\sigma_i(\theta_i)$.

[25] For the special case of $R = I$, our definition is analogous to Csóka et al. (2024, page 14). By contrast, Csóka et al. (2025) does not allow punishments when modeling collusion by rings for $R \subsetneq I$, so, in effect, only allows for *ex ante* acceptable collusion by rings.

[26] The restriction to pure strategies σ_i' for each bidder i under punishment is without loss of generality.

[27] Every *ex ante* acceptable collusion scheme for a selling mechanism (M, q, t, σ) is an individually rational collusion scheme for the grand coalition for (M, q, t, σ), provided that σ is a Bayes–Nash equilibrium. (The punishment schemes for a bidder i can be taken to consist of equilibrium strategies for the other bidders.) Hence, individually rational collusion schemes (essentially) include the collusion schemes considered in Sections 2 and 3.

rational and therefore could be replicated within a larger collusion scheme. Therefore, the assumption that non-members of the ring play their non-collusive strategies in an (individually rational) collusion scheme for a ring is without loss of generality.

When sustained with punishments, individually rational collusion schemes need not raise all or even any of the ring members' payoffs. Thus, the rich form of potential punishments makes our notion of individually rational collusion very permissive.

5.2. **Strong Collusion Robustness.** It proves convenient for us to work with an even more permissive notion of individual rationality (and therefore even stronger notions of collusion and collusion robustness), which allows all agents (not just ring members) to join in punishing a defector, and constrains defectors to respond to the punishment with their prescribed strategies under the selling mechanism. This allows us to avoid any distinction between ring members and non-members in our arguments.

Definition 6′. A collusion scheme $(\mu, \hat{q}, \hat{t})$ for selling mechanism (M, q, t, σ) is *weakly individually rational* if for each $i \in I$ there exists a (measurable) punishment scheme $\nu^i : \Theta_{-i} \to \Delta(M_{-i})$ such that

$$(5') \qquad \mathbb{E}_{\tilde{\theta}}\big[\tilde{\theta}_i \hat{q}_i(\tilde{\theta}) - \hat{t}_i(\tilde{\theta})\big] \geq \mathbb{E}_{\tilde{\theta}}\big[\mathbb{E}_{\tilde{m}_{-i} \sim \nu^i(\tilde{\theta}_{-i})}\big[\tilde{\theta}_i q_i\big(\tilde{m}_{-i}, \sigma_i(\tilde{\theta}_i)\big) - t_i\big(\tilde{m}_{-i}, \sigma_i(\tilde{\theta}_i)\big)\big]\big].$$

As our terminology suggests, weakly individually rational collusion is even more permissive than individually rational collusion by any ring.

Lemma 2. *Any individually rational collusion scheme for a ring $R \subseteq I$ for a selling mechanism is a weakly individually rational collusion scheme for that mechanism.*

Proof. Consider a selling mechanism (M, q, t, σ). For each bidder $i \in R$, consider a punishment scheme $\nu^{i,R} : \Theta_{R \setminus \{i\}} \to \Delta(M_{R \setminus \{i\}})$ that satisfies (5) for all strategies σ_i'. Taking $\sigma_i' = \sigma_i$, we can see that (5′) holds for the punishment scheme $\nu^i(\theta_{I \setminus \{i\}}) = \nu^{i,R}(\theta_{R \setminus \{i\}}) \times \sigma_{I \setminus R}(\theta_{I \setminus R})$. For bidders $i \notin R$, (4) implies that (5′) holds for the punishment scheme $\nu^i(\theta_{I \setminus \{i\}})$ that is the marginal distribution of $\mu(\theta)$ on $M_{I \setminus \{i\}}$. □

Our concept of strong collusion-robustness requires that the outcome and revenue of a mechanism be robust to any weakly individually rational collusion scheme. As our notion of weakly individually rational collusion is extremely permissive, strong collusion robustness is an extremely strong concept.

Definition 7. A mechanism (M, q, t, σ) is *strongly collusion-robust* if all weakly individually rational collusion schemes

(i) give all bidders the same ex-ante utilities as if all bidders used strategy σ;

(ii) almost surely induce the same allocation as if all bidders used σ; and
(iii) yield the same expected revenue as if all bidders used strategy σ.

Note that due to Lemma 2, the definition implies that Properties (i)–(iii) must hold for all individually rational collusion schemes by arbitrary rings. Thus, our definition is stronger than the concept of collusion-proofness introduced by Csóka et al. (2024), which only requires robustness to individually rational collusion by the grand coalition, and hence than Safronov's (2018) weaker coalition-proofness concept.

Remark 1. Strong collusion-robustness also implies that Properties (i)–(iii) must hold for all *ex ante* acceptable collusion schemes. Indeed, every *ex ante* acceptable collusion scheme for a selling mechanism is a weakly individually rational collusion scheme for the grand coalition for that selling mechanism, in which the punishment schemes for a bidder consists of prescribed strategies for the other bidders. Hence, strong collusion-robustness ensures robustness to the form of collusion considered in Sections 2 and 3.

In addition to being more technically convenient to work with, strong collusion robustness allows us to unify collusion robustness and Bayes–Nash equilibrium. Moreover, the outcome and revenue of a strongly collusion-robust mechanism are guaranteed to be robust to equilibrium selection as well as to collusion. Hence, strong collusion-robustness also implies a form of full implementation.

Proposition 3. *If a mechanism (M, q, t, σ) is strongly collusion-robust, then σ is a Bayes–Nash equilibrium,*[28] *and Properties (i)–(iii) hold for all Bayes–Nash equilibria.*

Proof. To prove the first part, note that for each bidder i, any better response σ_i' to $\sigma_{I\setminus\{i\}}$ defines a weakly individually rational collusion scheme with deterministic report function $\mu(\theta) = (\sigma_i'(\theta_i), \sigma_{I\setminus\{i\}}(\theta_{I\setminus\{i\}}))$ in which no reallocation or side payments occur. Property (i) in Definition 7 implies that σ_i' cannot be a strictly better response.

For the second part, note that each Bayes–Nash equilibrium defines a weakly individually rational collusion scheme (with no reallocation or side payments), in which the punishment scheme for a bidder consists of equilibrium strategies for the other bidders. □

6. Collusion-Robust Selling Mechanisms

In this section, we construct two strongly collusion-robust mechanisms that provide the revenue of the optimal posted price. The first mechanism works in the case of symmetric bidders; the second mechanism is more complicated, but works in general.

[28] We consider *ex-ante* Bayes–Nash equilibrium throughout for technical convenience.

Both constructions rely on a set of sufficient conditions for a selling mechanism to be strongly collusion-robust and yield expected revenue $\Pi(p)$, for a given price p.

Proposition 4. *Let p be a price, and denote the bidders' maximal aggregate surplus net of price p by*

$$S^p(\theta) = \max_{q \in \mathcal{Q}} \Big\{ \sum_{i \in I} (\theta_i - p) q_i \Big\} = \max \Big\{ \max_{i \in I} \{\theta_i - p\}, 0 \Big\}.$$

If a selling mechanism (M, q, t, σ) is such that:

(I) for all type profiles $\theta \in \Theta$, we have

$$\sum_{i \in I} (\theta_i - p) q_i(\sigma(\theta)) = S^p(\theta);$$

(II) for all type profiles $\theta \in \Theta$, we have

$$\sum_{i \in I} t_i(\sigma(\theta)) - p \sum_{i \in I} q_i(\sigma(\theta)) = 0;$$

(III) for all message profiles $m \in M$, we have

$$\sum_{i \in I} t_i(m) - p \sum_{i \in I} q_i(m) \geq 0;$$

(IV) for all bidders i and message profiles $m_{-i} \in M_{-i}$, we have

$$\mathbb{E}_{\tilde{\theta}_i}\big[\tilde{\theta}_i q_i(\sigma_i(\tilde{\theta}_i), m_{-i}) - t_i(\sigma_i(\tilde{\theta}_i), m_{-i})\big] \geq \mathbb{E}_{\tilde{\theta}}\big[\tilde{\theta}_i q_i(\sigma(\tilde{\theta})) - t_i(\sigma(\tilde{\theta}))\big].$$

then it is strongly collusion-robust and yields expected revenue $\Pi(p)$ under strategies σ.

Proof. Consider any weakly individually rational collusion scheme $(\mu, \hat{q}, \hat{t})$ for (M, q, t, σ). Due to (IV), the collusive *ex-ante* expected utility of each bidder is bounded below by

$$\mathbb{E}_{\tilde{\theta}}[\tilde{\theta}_i \hat{q}_i(\tilde{\theta}) - \hat{t}_i(\tilde{\theta})] \geq \mathbb{E}_{\tilde{\theta}}\big[\tilde{\theta}_i q_i(\sigma(\tilde{\theta})) - t_i(\sigma(\tilde{\theta}))\big] \text{ for all } i \in I. \tag{6}$$

At the same time, using the collusion feasibility conditions (1) and (2), Condition (III), and Condition (II), for each type profile θ, we have

$$\sum_{i \in I} \hat{t}_i(\theta) - p \sum_{i \in I} \hat{q}_i(\theta) = \mathbb{E}_{\tilde{m} \sim \mu(\theta)}\Big[\sum_{i \in I} t_i(m) - p \sum_{i \in I} q_i(m) \Big] \geq 0 = \sum_{i \in I} t_i(\sigma(\theta)) - p \sum_{i \in I} q_i(\sigma(\theta)). \tag{7}$$

Adding (6) over all the bidders and adding the expectation of (7) yields

$$\sum_{i \in I} \mathbb{E}_{\tilde{\theta}}\big[(\tilde{\theta}_i - p) \hat{q}_i(\tilde{\theta})\big] \geq \mathbb{E}_{\tilde{\theta}}\big[S^p(\tilde{\theta})\big].$$

Together with Condition (I) and the definition of S, this implies that Property (ii) of Definition 7 holds, and that all the bounds must be tight. In particular, (6) must hold with equality for all bidders i, which implies Property (i) of Definition 7. Moreover, (7)

must hold with equality almost surely, which, together with Property (ii) of Definition 7, implies Property (iii) of Definition 7, i.e., the expected revenue under the collusion scheme is the same as under the prescribed strategies σ. Condition (I) and (II) then imply that expected revenue is $\Pi(p)$. □

Conditions (I) and (II) of the proposition require that if bidders follow the prescribed strategies, the aggregate allocation must be efficient for the bidders given price p, and the aggregate payment must be the corresponding posted price payment. Condition (III) is a revenue guarantee: it requires that revenue always be at least p times the aggregate allocation. Condition (IV) is a utility guarantee: it requires that if any bidder i follows his prescribed strategy, his *ex-ante* expected utility cannot be reduced by others deviating from their prescribed strategies (to any fixed message profile, and hence to any contingent message profile independent of bidder i's type).

Note that given Condition (II), Condition (III) can be interpreted as a case of Condition (IV) for a fictitious seller who values the good at p. With this interpretation, Condition (I) requires that the mechanism allocate the good efficiently, achieving the maximal available aggregate surplus of $S^p(\theta)$, and Condition (II) requires *ex-post* individual rationality for the fictitious seller.

A simple intuition behind Proposition 4 runs as follows: any weakly individually rational collusive deviation cannot lower any bidder's expected payoff by Condition (IV) and cannot reduce the fictitious seller's payoff by Condition (III), so cannot reduce the aggregate surplus of these agents. But as this aggregate surplus is already maximized by Condition (I), the deviation must yield a payoff-equivalent outcome for the agents. With continuous value distributions, such a deviation must almost surely not affect the allocation, and therefore cannot affect expected revenue.

The closest work that Proposition 4 builds upon is Safronov (2018), Csóka et al. (2024), and Csóka et al. (2025).[29] However, it differs from those papers in some important respects. First, to ensure a revenue guarantee (rather than budget balance), we introduce Conditions (II) and (III), which can be interpreted as a payoff guarantee for a fictitious seller who values the good at p. Second, to allow for nonparticipation messages (which are not required in the aforementioned papers), we state the guaranteed-utility conditions (III) and (IV) as inequalities rather than equalities.[30]

[29]It is also related to Segal and Whinston (2012, Proposition 1), in which the guaranteed-utility property was referred to as "nonharmful externalities."

[30]By changing off-path payments, any mechanism that satisfies these conditions can be converted to one of these conditions is satisfied and the other always holds with equality, but not in general one for which both conditions always hold with equality.

Third, we establish a stronger collusion-proofness property by allowing rings to commit to punish bidders that refuse to collude (unlike in Safronov (2018) and Csóka et al. (2025)), and by considering collusion by rings that are smaller than the grand coalition (unlike in Csóka et al. (2024)). Last, we show the mechanism remains robust even to collusion by rings that can reallocate the good.

In our constructions below, we ensure Conditions (I) and (II) by combining a posted price with an efficient, *ex-post* budget balanced collusion scheme. Obtaining Conditions (III) and (IV) in the case of nonparticipation is more subtle. Note, for example, that if nonparticipation by a bidder resulted in cancellation of the mechanism, Condition (IV) would fail—thereby enabling collusion using nonparticipation threats (as in Example 3), as well as creating undesirable nonparticipation equilibria. To avoid these issues, nonparticipation by some agents must allow the mechanism to proceed in a way that does not excessively harm participants. On the other hand, nonparticipation by some agents must not allow participants to obtain the good at the aggregate price below p—i.e., Condition (III) must be satisfied, for a violation of it would leave the seller vulnerable to collusion involving nonparticipation.

To thread the needle between these two considerations, our constructions stipulate that following nonparticipation by some agents, the mechanisms proceed with the others the same way as if the nonparticipants submitted the equilibrium messages of lowest types. As nonparticipants might have otherwise received payments from the seller, this approach relies on allowing inequality in Condition (III).

Finally, we show that there are mechanisms other than Safronov's (2018) that are strongly collusion-robust due to satisfying Conditions (I)–(IV). In fact, for the symmetric case, there is a natural indirect mechanism that works.

6.1. **The Symmetric Case.** Suppose that bidders are *symmetric* in that their value distributions are the same distribution $F_i = F$. We combine the p-posted-price mechanism with a first-price knockout auction for the ability to purchase the good for p whose revenue is distributed equally among losers of the knockout auction. The seller conducts the knockout auction on the behalf of the bidders.

We call this the *First-Price Knockout p-Posted Price (FPK-p) mechanism*. The message space is $M_i = [0, \bar{b}_i] \cup \{NP_i\}$, where $\bar{b}_i$ is sufficiently large. For interpretation, submitting bid b corresponds to a bid of $b-p$ in the first-price knockout auction. As in standard auctions with reserve p, if no one bids at least p, the good is not allocated; otherwise, the good is allocated to the highest bidder who bids at least b (breaking

ties uniformly). Formally, if all bidders submit messages $b_i \in [0, \overline{b}_i]$, the allocation is

$$q_i^p(b) = \begin{cases} 0 & \text{if } b_i < p \text{ or } b_i < \max_{j \in I \setminus \{i\}} b_j \\ \frac{1}{|\{j \mid b_j = b_i\}|} & \text{if } b_i \geq p \text{ and } b_i \geq \max_{j \in I \setminus \{i\}} b_j. \end{cases} \tag{8}$$

As for payments, the seller keeps p if the good is sold and returns any surplus revenue (which corresponds to the revenue of the knockout auction) back to the losing bidders. That is, payments if all bidders participate are given by

$$\begin{aligned} t_i^{\text{FPK-}p}(b) &= \underbrace{p q_i^p(b)}_{\text{posted price}} + \underbrace{(b-p) q_i^p(b)}_{\text{knockout payment}} - \frac{1}{N-1} \underbrace{\sum_{j \in I \setminus \{i\}} (b_j - p) q_j^p(b)}_{\text{knockout revenue}} \\ &= b q_i^p(b) - \frac{1}{N-1} \sum_{j \in I \setminus \{i\}} (b_j - p) q_j^p(b). \end{aligned}$$

When some agents do not participate, the mechanism runs for the participants the same way as if the nonparticipants j bid their lowest type $\underline{\theta}_j$.[31] Thus, for any nonparticipating coalition $S \subseteq I$, the participants' outcomes are

$$\left(q_i^p(NP_S, b_{I \setminus S}), t_i^{\text{FPK-}p}(NP_S, b_{I \setminus S})\right) = \left(q_i^p(\underline{\theta}_S, b_{I \setminus S}), t_i^{\text{FPK-}p}(\underline{\theta}_S, b_{I \setminus S})\right) \text{ for } i \in I \setminus S. \tag{9}$$

This mechanism has an important property that helps ensure Condition (IV) of Proposition 4: a bidder's allocation and payment both depend on the other agents' bids only through their highest bid, which limits their ability to hurt him. To satisfy Condition (IV), we aim to construct a strategy σ_i that makes bidder i's expected payoff independent of the others' highest bid.[32] By symmetry, the strategy σ_i is the same for all bidders, and if it is strictly increasing, all bidders using it yields the efficient allocation, implying Condition (I) of Proposition 4. As Conditions (II) and (III) are also easily verified, strong collusion robustness follows.

Instead of constructing the strategy σ_i from scratch, however, we make use of Proposition 3, which says that $\sigma = (\sigma_i)_{i \in I}$ must be a Bayes-Nash equilibrium. Thus, we can take advantage of McAfee and McMillan (1992, Theorem 3), who constructed

[31] Alternatively, we could allocate the knockout revenue among losing participants, rather than the seller keeping nonparticipants' shares of the knockout revenue. This alternative approach would also lead to a strongly collusion-robust mechanism.

[32] One caveat is that if the others' bid high enough to surely beat bidder i's bid, then raising the others' bids further would benefit bidder i due to his share of the knockout auction revenue. In this case, Condition (IV) can hold with strict inequality.

a Bayes–Nash equilibrium of the knockout auction in which type θ_i bids

$$h^p(\theta_i) = \frac{1}{F(\theta_i)^N}\frac{N-1}{N}\int_p^{\theta_i}(x-p)d(F(x)^N).$$

The corresponding strategy in the FPK-p mechanism is

$$\sigma_i^{\text{FPK-}p}(\theta_i) = \begin{cases}\theta_i & \text{if } \theta_i < p\\ p + h^p(\theta_i) & \text{if } \theta_i \geq p.\end{cases}$$

Then, verifying that the conditions of Proposition 4 are indeed satisfied yields the following robustness theorem (the complete proof is in the appendix).

Theorem 2. *With symmetric bidders, the FPK-p mechanism is strongly collusion-robust and yields expected revenue* $\Pi(p)$ *under strategy profile* $\sigma^{\text{FPK-}p}$.

6.2. **The General Case.** We next turn to the general case and construct a more complicated mechanism that is strongly collusion-robust with asymmetric bidders. We combine the posted price with a knockout mechanism whose side payments correspond to expected externalities when bidders submit reports sequentially. (The seller executes this knockout mechanism on behalf of the bidders.)

We call the mechanism the *Sequential Expected Externality Knockout p-Posted Price (SEEK-p) mechanism*. It is direct, so bidder i's message space is $M_i = \Theta_i \cup \{NP_i\}$ and his strategy is $\sigma(\theta_i) = \theta_i$. It uses the same allocation rule as the FPK-p mechanism.

We begin by describing payments when all bidders participate—i.e., $m = \theta \in \Theta$. Taking an arbitrary order of the bidders, we define the payments $t_i^{\text{SEEK-}p}(\theta)$ so that

$$\underbrace{\theta_i q_i^p(\theta) - t_i^{\text{SEEK-}p}(\theta)}_{\textit{ex post}\text{ utility}} = \underbrace{\mathbb{E}_{\tilde{\theta}_{\geq i}}\left[S^p\left(\theta_{<i}, \theta_i, \tilde{\theta}_{>i}\right) - S^p\left(\theta_{<i}, \tilde{\theta}_i, \tilde{\theta}_{>i}\right)\right]}_{\text{internalizing impact from } i \text{ reporting } \theta_i} + \gamma_i^p, \tag{10}$$

where the constants γ_i^p are chosen so that for each bidder i, we have that

$$t_i^{\text{SEEK-}p}(\underline{\theta}) = pq_i^p(\underline{\theta}). \tag{11}$$

The intuition for (10) is that bidders report sequentially (in an order that can be fixed arbitrarily, or randomized). Each bidder is charged an amount that makes him internalize the impact of his report on the aggregate surplus based on the preceding bidders' reported values but in expectation over the succeeding bidders' values. In addition, each bidder i receives an additional lump-sum payment γ_i^p so that for the lowest type profile $\underline{\theta}$, he just pays the posted price p for his purchase (if any).

We specify outcomes following nonparticipation by some agents analogously to (9): the mechanism proceeds for the participants the same way as if the nonparticipants

j reported their lowest types, which we denote by $\underline{\theta}_j$. Thus, for any nonparticipating coalition $S \subseteq I$, the participants' outcomes are

$$(12)\quad \left(q_i^p(NP_S, \theta_{I\setminus S}), t_i^{\text{SEEK-}p}(NP_S, \theta_{I\setminus S})\right) = \left(q_i^p(\underline{\theta}_S, \theta_{I\setminus S}), t_i^{\text{SEEK-}p}(\underline{\theta}_S, \theta_{I\setminus S})\right) \text{ for } i \in I\setminus S.$$

The construction of the payments is based on the ideas of Safronov (2018).[33] Specifically, Condition (II) of Proposition 4 can be verified when all bidders participate by adding (10) across agents and using a telescoping-sum property. (10) also ensures that when bidder i reports truthfully, his *ex-ante* expected utility equals γ_i^p for all possible type reports by the other bidders, which verifies Condition (IV) when all bidders participate.

A key difference from Safronov (2018) lies in ensuring Conditions (III) and (IV) when some bidders do not participate. Namely, we show that with the lump-sum payments γ_i^p constructed to satisfy (11), the lowest type of each agent *never* pays more than p for his purchase. Hence, when some agents do not participate, (12) and Condition (II) ensure that the participants pay at least p (per unit) for their aggregate purchase, making Condition (III) hold. (12) also ensures that Condition (IV) is preserved following nonparticipation by some agents.

Condition (I) is immediate. Thus, with all the conditions of Proposition 4 verified, we obtain the following robustness theorem (the complete proof is in the appendix).

Theorem 3. *The SEEK-p mechanism is strongly collusion-robust and yields expected revenue* $\Pi(p)$ *under truthtelling.*

6.3. ***Ex-Post* Individual Rationality.** In addition to offering robustness to *ex-ante* collusion, the FPK-p and SEEK-p mechanisms are *ex-post* individually rational—a property not shared by other collusion-proof mechanisms that have been proposed in the literature (e.g., Che and Kim (2006) and Pavlov (2008)).[34] Thus, it is secure for bidders to participate in these mechanisms (in the sense of Dasgupta et al. (1979)) even if bidders are concerned about others' colluding.

Formally, we say that a selling mechanism (M, q, t, σ) is *ex-post individually rational* if for all bidders i and type profiles $\theta \in \Theta$, we have $\theta_i q_i(\sigma(\theta)) - t_i(\sigma(\theta)) \geq 0$.

Proposition 5. *FPK-p and SEEK-p are ex post individually rational mechanisms.*

[33] While Safronov (2018) defines bilateral externality payments between agents, his total payments can be expressed similarly to (10).

[34] If we required that selling mechanisms be *ex-post* individually rational instead of requiring that they include nonparticipation messages for each bidder, our theorems would persist unchanged.

7. Conclusion

This paper shows that if bidders can collude *ex ante* (i.e., before learning their values or deciding to participate), then the seller cannot obtain any revenue beyond a posted price. That is, *ex-ante* collusion leads to a complete unraveling of auctions.

If collusion is efficient, the seller can obtain the optimal posted-price revenue by posting an optimal price and letting bidders collude. However, posted-price mechanisms are vulnerable to inefficient collusion. We construct mechanisms that yield the optimal posted price revenue and are *strongly collusion-robust*: i.e., robust to collusion in which rings can observe their members' types, and even if it can force members to participate by committing to punish them for defections. In the symmetric case, strong collusion robustness can be obtained by simply combining the posted-price mechanism with a first-price knockout auction conducted by the seller on behalf of the bidders (with revenue remitted back to the bidders). In the general case, robustness is achieved by a direct mechanism with more complicated payments.

Conceptually, our main contribution is to delineate the limits of auctions when bidders can collude perfectly, even when they are subject to asymmetric information. In practice, it may be possible to achieve more revenue through auctions, but only due to legal or institutional features that limit bidders' ability to collude.

From a methodological perspective, our concept of strong collusion-robustness may be useful as a desideratum for market design beyond the question of revenue maximization under collusion. We leave the investigation of this topic to future research.

Appendix A. Omitted Proofs

A.1. **Proof of Proposition 1.** As utility is transferable, it suffices to prove that there exists an efficient, incentive-compatible, *ex ante* acceptable collusion scheme.

Let (M, q, t, σ) be a selling mechanism. Define a deterministic report function $\mu(\theta)$ and efficient collusive allocation rule $\hat{q}(\theta)$ by (an arbitrary measurable selection from)

$$\big(\hat{q}(\theta), \mu(\theta)\big) \in \underset{(q,m)\in\mathcal{Q}\times M \,|\, \sum_{i\in I} q_i = \sum_{i\in I} q_i(m)}{\arg\max} \left\{ \sum_{i\in I} \theta_i q_i - \sum_{i\in I} t_i(m) \right\},$$

which exists by our compactness assumption on the space of aggregate allocations and transfers. To construct the payment rule for the collusion scheme, we first define Vickrey–Clarke–Groves side payments by[35]

$$\hat{t}_i^{0,\mathrm{side}}(\theta) = \sum_{j\in I\setminus\{i\}} t_j(\mu(\theta)) - \sum_{j\in I\setminus\{i\}} \theta_j \hat{q}_j(\theta).$$

We can then define a (collusive) payment rule by

$$\hat{t}_i(\theta) = t_i(\mu(\theta)) + \mathbb{E}_{\tilde{\theta}_{-i}}\big[\hat{t}_i^{0,\mathrm{side}}(\theta_i, \tilde{\theta}_{-i})\big] - \frac{1}{N-1} \sum_{j\in I\setminus\{i\}} \mathbb{E}_{\tilde{\theta}_{-j}}\big[\hat{t}_j^{0,\mathrm{side}}(\theta_j, \tilde{\theta}_{-j})\big].$$

Conditions (1) and (2) hold by construction, so we have defined a collusion scheme. It is incentive-compatible by the arguments of d'Aspremont and Gérard-Varet (1979).[36]

The collusion scheme is efficient by construction. In particular, we have that $\sum_{i\in I} U_i(\hat{q}, \hat{t}) \geq \sum_{i\in I} V_i(M, q, t, \sigma)$. Hence, the collusion scheme can be made *ex ante* acceptable by adding budget-balanced lump sum side transfers among the bidders.

A.2. **Proof of Proposition 2.** Denote the support of F by $[\underline{\theta}, \overline{\theta}]$. By construction, we have $p_N^* \in \arg\max_p r(p, N)$, where $r(p, N) = p(1 - F(p)^N)$. First, note $p_N^* \in (\underline{\theta}, \overline{\theta})$

[35] These payments are based on the externalities that i exerts on other bidders through the (collusive) allocation rule $\hat{q}(\theta)$ and payments $t(\mu(\theta))$ to the seller.

[36] We give the argument for sake of completeness. Defining payment rule $\hat{t}_i^0(\theta) = t_i(\mu(\theta)) + \hat{t}_i^{0,\mathrm{side}}(\theta)$, the mechanism $(\hat{q}, \hat{t}^0)$ is a Groves (1973) scheme. By construction, for all bidders i and types θ_i, we have

$$\mathbb{E}_{\tilde{\theta}_{-i}}[\hat{t}_i(\theta)] = \mathbb{E}_{\tilde{\theta}_{-i}}\big[\hat{t}_i^0(\theta_i, \tilde{\theta}_{-i})\big] - \frac{1}{N-1} \sum_{j\in I\setminus\{i\}} \mathbb{E}_{\tilde{\theta}}\big[\hat{t}_j^{0,\mathrm{side}}(\tilde{\theta})\big].$$

It follows that for all bidders i, types θ_i, and reports θ_i', we have

$$\begin{aligned}
\mathbb{E}_{\tilde{\theta}_{-i}}\big[\theta_i \hat{q}_i(\theta_i, \tilde{\theta}_{-i}) - \hat{t}_i(\theta_i, \tilde{\theta}_{-i})\big] &= \mathbb{E}_{\tilde{\theta}_{-i}}\big[\theta_i \hat{q}_i(\theta_i, \tilde{\theta}_{-i}) - \hat{t}_i^0(\theta_i, \tilde{\theta}_{-i})\big] + \frac{1}{N-1} \sum_{j\in I\setminus\{i\}} \mathbb{E}_{\tilde{\theta}}\big[\hat{t}_j^{0,\mathrm{side}}(\tilde{\theta})\big] \\
&\geq \mathbb{E}_{\tilde{\theta}_{-i}}\big[\theta_i \hat{q}_i(\theta_i', \tilde{\theta}_{-i}) - \hat{t}_i^0(\theta_i', \tilde{\theta}_{-i})\big] + \frac{1}{N-1} \sum_{j\in I\setminus\{i\}} \mathbb{E}_{\tilde{\theta}}\big[\hat{t}_j^{0,\mathrm{side}}(\tilde{\theta})\big] \\
&= \mathbb{E}_{\tilde{\theta}_{-i}}\big[\theta_i \hat{q}_i(\theta_i', \tilde{\theta}_{-i}) - \hat{t}_i(\theta_i', \tilde{\theta}_{-i})\big],
\end{aligned}$$

where the inequality follows by the incentive compatibility of Groves's (1973) mechanisms.

for all $N > 1$. Indeed, this follows from the facts that $r(\overline{\theta}, N) = 0$, that $r(\underline{\theta}, N) = \underline{\theta} \geq 0$, and that $\frac{\partial r(p,N)}{\partial p}\big|_{p=\underline{\theta}} = 1$ (for $N > 1$). Second, note that for $p < \overline{\theta}$ and $N > 1$, we have

$$\frac{\partial^2 \log r(p,N)}{\partial N \partial p} = \frac{\partial}{\partial p} \frac{\log(N) F(p)^N}{1 - F(p)^N} > 0.$$

Thus, $\log r(p,N)$ has increasing marginal returns for $p < \overline{\theta}$ in the sense of Edlin and Shannon (1998). The result follows by Edlin and Shannon (1998, Theorem 1).

A.3. **Proof of Theorem 2.** We verify the conditions of Proposition 4. For notation, let $\overline{\theta}_i = \max \Theta_i$. Condition (I) holds as h^p is strictly increasing on $[p, \overline{\theta}_i]$.

To verify Condition (II), note that

$$\sum_{i \in I} t_i^{\text{FPK-}p}(b) = p \sum_{i \in I} q_i^p(b) \tag{A.1}$$

holds for all bid profiles $b \in \times_{i \in I} [0, \overline{b}_i]$ by construction.

To verify Condition (III), note that given any message $m = (NP_S, b_{I \setminus S})$, by construction, we have

$$t_i^{\text{FPK-}p}(m) - p q_i^p(m) = 0 \geq t^{\text{FPK-}p}(\underline{\theta}_S, b_{I \setminus S}) - p q_i^p(\underline{\theta}_S, b_{I \setminus S})$$

for all $i \in S$. By the construction of participants' outcomes if some bidders choose not to participate in (9), and using (A.1) for the bid profile $(\underline{\theta}_S, b_{I \setminus S})$, we have

$$\sum_{i \in I} t_i^{\text{FPK-}p}(m) - p \sum_{i \in I} q_i^p(m) \geq \sum_{i \in I} t_i^{\text{FPK-}p}(\underline{\theta}_S, b_{I \setminus S}) - p \sum_{i \in I} q_i^p(\underline{\theta}_S, b_{I \setminus S}) = 0.$$

To verify Condition (IV), consider any message profile $m_{-i} = (NP_S, m_{I \setminus \{i\} \setminus S}) \in M_{-i}$. Let $z = \max\left\{0, \max_{j \in I \setminus \{i\} \setminus S} \{b_j\}\right\}$ denote the highest bid of the other bidders. Let ψ be the largest value such that $\psi + h^p(\psi) \leq z$, so $\sigma_i^{\text{FPK-}p}(\theta_i) > z$ if and only if $\theta_i > \psi$.

We evaluate bidder i's *ex-ante* utility from following σ_i if others' send message profile m_{-i} by first calculating i's expected knockout payment in this situation: note

$$\begin{aligned}
\int_{\psi}^{\overline{\theta}_i} h^p(x) dF(x) &= \int_{\psi}^{\overline{\theta}_i} \int_p^x \frac{(N-1)(y-p)}{N F(x)^N} d(F(y)^N) dF(x) \\
&= \int_p^{\overline{\theta}_i} \int_{\max\{\psi, y\}}^{\overline{\theta}_i} \frac{(N-1)(y-p)}{N F(x)^N} dF(x) d(F(y)^N) \\
&= \int_p^{\overline{\theta}_i} \frac{y-p}{N} \left(\frac{1}{F(\max\{\psi, y\})^{N-1}} - 1 \right) d(F(y)^N) \\
&= \int_p^{\overline{\theta}_i} \frac{y-p}{N F(\max\{\psi, y\})^{N-1}} - \int_p^{\overline{\theta}_i} \frac{y-p}{N} d(F(y)^N). \qquad \text{(A.2)}
\end{aligned}$$

where the second equality is by Fubini's Theorem and the third is by evaluating the inner integral. Breaking the first integral in the last line into two parts yields that

$$\begin{aligned}\int_p^{\bar\theta_i} \frac{y-p}{NF(\max\{\psi,y\})^{N-1}} &= \int_p^{\psi} \frac{y-p}{NF(\psi)^{N-1}} d(F(y)^N) + \int_\psi^{\bar\theta_i} \frac{y-p}{NF(y)^{N-1}} d(F(y)^N) \\ &= F(\psi)\frac{h(\psi)}{N-1} + \int_\psi^{\bar\theta_i} (y-p)dF(y). \end{aligned} \tag{A.3}$$

Note also that as $F(y)^N$ is the distribution of the highest bidder value, we have that

$$\begin{aligned}\int_p^{\bar\theta_i} \frac{y-p}{N} d(F(y)^N) &= \frac{1}{N}\mathbb{E}_{\tilde\theta}\Big[\max\Big\{\max_{j\in I}\{\theta_j - p\}, 0\Big\}\Big] \\ &= \frac{1}{N}\mathbb{E}_{\tilde\theta}[S^p(\tilde\theta)] = V_i(M, q^p, t^{\text{FPK-}p}, \sigma^{\text{FPK-}p}), \end{aligned} \tag{A.4}$$

where the second equality is by the definition of S^p and the third by symmetry, (I), and (II). Combining (A.2), (A.3) and (A.4) yields that

$$\begin{aligned} V_i(M, q^p, t^{\text{FPK-}p}, \sigma^{\text{FPK-}p}) &= \int_\psi^{\bar\theta_i} \big(x - p - h^p(x)\big) dF(x) + F(\psi)\frac{h(\psi)}{N-1} \\ &= \int_\psi^{\bar\theta_i} \big(x - p - h^p(x)\big) dF(x) + \int_{\underline\theta_i}^{\psi} \frac{h(\psi)}{N-1} dF(x) \\ &\le \int_\psi^{\bar\theta_i} \big(x - p - h^p(x)\big) dF(x) + \int_{\underline\theta_i}^{\psi} \frac{z-p}{N-1} dF(x) \\ &= \mathbb{E}_{\tilde\theta_i}\Big[\tilde\theta_i q_i^p\big(\sigma^{\text{FPK-}p}(\tilde\theta_i), m_{I\setminus\{i\}}\big) - t_i^{\text{FPK-}p}\big(\sigma^{\text{FPK-}p}(\tilde\theta_i), m_{I\setminus\{i\}}\big)\Big], \end{aligned}$$

where the inequality holds due to the definition of ψ, and the last equality holds due to the definitions of q_i^p, $t_i^{\text{FPK-}p}$, z, and ψ. Hence, the theorem follows from Proposition 4.

A.4. **Proof of Theorem 3.** We prove the theorem by verifying the conditions of Proposition 4 in turn. Condition (I) holds by construction.

For notation, we let $u_i^{\text{SEEK-}p}(\theta_i, m_{I\setminus\{i\}}) = \theta_i q_i^p(\theta_i, m_{I\setminus\{i\}}) - t_i^{\text{SEEK-}p}(\theta_i, m_{I\setminus\{i\}})$.

To verify Condition (II), note that the construction of $t_i^{\text{SEEK-}p}$ in (10) implies

$$\begin{aligned}\sum_{i\in I} u_i^{\text{SEEK-}p}(\theta) &= \sum_{i\in I} \mathbb{E}_{\tilde\theta_{\ge i}}\big[S^p\big(\theta_{<i}, \theta_i, \tilde\theta_{>i}\big) - S^p\big(\theta_{<i}, \tilde\theta_i, \tilde\theta_{>i}\big)\big] + \sum_{i\in I}\gamma_i^p \\ &= S^p(\theta) - \mathbb{E}_{\tilde\theta}\big[S^p(\tilde\theta)\big] + \sum_{i\in I}\gamma_i^p = \sum_{i\in I}(\theta_i - p)q_i^p(\theta) - \mathbb{E}_{\tilde\theta}\big[S^p(\tilde\theta)\big] + \sum_{i\in I}\gamma_i^p, \end{aligned}$$

where the second equality uses a telescoping property of the sum involving expectations of differences of S^p, and the third is by the definition of q_i^p. It follows that

$$p\sum_{i\in I} q_i^p(\theta) - \sum_{i\in I} t_i^{\text{SEEK-}p}(\theta) = \sum_{i\in I}\big(u_i^{\text{SEEK-}p}(\theta) - (\theta_i - p)q_i^p(\theta)\big) = -\mathbb{E}_{\tilde\theta}\big[S^p(\tilde\theta)\big] + \sum_{i\in I}\gamma_i^p.$$

Thus, the left-hand side is constant in θ. By the construction of the lump-sum payments γ_i^p in (11), it equals 0 for $\theta = \underline{\theta}$. Hence, it equals 0 for all type profiles θ.

To verify Condition (III), note that as $S^p(\theta)$ is a submodular function, the construction of $t^{\text{SEEK-}p}$ in (10) implies

$$u_i^{\text{SEEK-}p}(\underline{\theta}_i, \theta_{I\setminus\{i\}}) \geq u_i^{\text{SEEK-}p}(\underline{\theta}) = (\underline{\theta}_i - p)q_i^p(\underline{\theta}) \geq (\underline{\theta}_i - p)q_i^p(\underline{\theta}_i, \theta_{I\setminus\{i\}}), \tag{A.5}$$

where the first equality is by the construction of the lump-sum payments γ_i^p in (11), the second inequality is by the definition of the allocation rule q^p in (8). Rearranging terms yields that $t_i^{\text{SEEK-}p}(\underline{\theta}_i, \theta_{I\setminus\{i\}}) - pq_i^p(\underline{\theta}_i, \theta_{I\setminus\{i\}}) \leq 0$. It follows that given any message $m = (NP_S, \theta_{I\setminus S})$, we have

$$t_i^{\text{SEEK-}p}(m) - pq_i^p(m) = 0 \geq t_i^{\text{SEEK-}p}(\underline{\theta}_S, \theta_{I\setminus S}) - pq_i^p(\underline{\theta}_S, \theta_{I\setminus S}) \quad \text{for all } i \in S.$$

Using (II) for the message $(\underline{\theta}_S, \theta_{I\setminus S})$ and the construction of participants' outcomes if some bidders choose not to participate in (9) and (12), we have

$$\sum_{i\in I} t_i^{\text{SEEK-}p}(m) - p\sum_{i\in I} q_i^p(m) \geq \sum_{i\in I} t_i^{\text{SEEK-}p}(\underline{\theta}_S, \theta_{I\setminus S}) - p\sum_{i\in I} q_i^p(\underline{\theta}_S, \theta_{I\setminus S}) = 0.$$

To verify Condition (IV), note that for all type profiles $\theta_{I\setminus\{i\}} \in \Theta_{I\setminus\{i\}}$, (10) implies

$$\mathbb{E}_{\tilde{\theta}_i}\left[u_i^{\text{SEEK-}p}\left(\tilde{\theta}_i, \theta_{I\setminus\{i\}}\right)\right] = \mathbb{E}_{\tilde{\theta}_i, \tilde{\theta}'_{>i}}\left[S^p\left(\theta_{<i}, \tilde{\theta}_i, \tilde{\theta}'_{>i}\right)\right] - \mathbb{E}_{\tilde{\theta}'_{\geq i}}\left[S^p\left(\theta_{<i}, \tilde{\theta}'_{>=i}\right)\right] + \gamma_i^p = \gamma_i^p.$$

Since, in particular, taking expectations over $\theta_{I\setminus\{i\}}$ yields $V_i(M, q^p, t^{\text{SEEK-}p}, \sigma) = \gamma_i^p$, (IV) holds with equality for messages $m_{I\setminus\{i\}} \in \Theta_{I\setminus\{i\}}$. If bidders $S \subseteq I \setminus \{i\}$ do not participate, then by (12) we have

$$\mathbb{E}_{\tilde{\theta}_i}\left[u_i^{\text{SEEK-}p}\left(\tilde{\theta}_i, NP_S, \theta_{I\setminus\{i\}\setminus S}\right)\right] = \mathbb{E}_{\tilde{\theta}_i}\left[u_i^{\text{SEEK-}p}\left(\tilde{\theta}_i, \underline{\theta}_S, \theta_{I\setminus\{i\}\setminus S}\right)\right] = V_i\left(M, q^p, t^{\text{SEEK-}p}, \sigma\right),$$

so (IV) still holds. Hence, the theorem follows from Proposition 4.

A.5. **Proof of Proposition 5.** For the FPK-p mechanisn, as $h^p(\theta_i) \leq \max\{\theta_i - p, 0\}$ holds for all types θ_i by construction, we have

$$\begin{aligned}\theta_i q_i^p(\sigma^{\text{FPK-}p}(\theta)) - t_i^{\text{FPK-}p}(\sigma^{\text{FPK-}p}(\theta)) &\geq \left(\theta_i - p - h^p(\theta_i)\right)q_i^p(\sigma^{\text{FPK-}p}(\theta)) \\ &\geq \left(\theta_i - p - \max\{\theta_i - p, 0\}\right)q_i^p(\sigma^{\text{FPK-}p}(\theta)) \geq 0,\end{aligned}$$

where the last inequality holds as $q_i^p(\sigma^{\text{FPK-}p}(\theta)) = 0$ for $\theta_i < p$.

For the SEEK-p mechanism, note that as $S^p(\theta)$ is nondecreasing in θ_i, (10) implies

$$u_i^{\text{SEEK-}p}(\theta) \geq u_i^{\text{SEEK-}p}(\underline{\theta}_i, \theta_{I\setminus\{i\}}) \geq (\underline{\theta}_i - p)q_i^p(\underline{\theta}_i, \theta_{I\setminus\{i\}}) \geq 0,$$

where the second inequality is (A.5), and the last is by definition of q_i^p.

References

Asker, J. (2010). A study of the internal organization of a bidding cartel. *American Economic Review 100*(3), 724–762.

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

Che, Y.-K. and J. Kim (2006). Robustly collusion-proof implementation. *Econometrica 74*(4), 1063–1107.

Che, Y.-K. and J. Kim (2009). Optimal collusion-proof auctions. *Journal of Economic Theory 144*(2), 565–603.

Cramton, P., R. Gibbons, and P. Klemperer (1987). Dissolving a partnership efficiently. *Econometrica 55*(3), 615–632.

Csóka, E., H. Liu, A. Rodivilov, and A. Teytelboym (2024). A collusion-proof efficient dynamic mechanism. Working paper.

Csóka, E., A. Pongrácz, and A. Rodivilov (2025). Guaranteed utility equilibrium. Working paper.

Dasgupta, P., P. Hammond, and E. Maskin (1979). The implementation of social choice rules: Some general results on incentive compatibility. *Review of Economic Studies 46*(2), 185–216.

d'Aspremont, C. and L.-A. Gérard-Varet (1979). Incentives and incomplete information. *Journal of Public Economics 11*(1), 25–45.

Dequiedt, V. (2007). Efficient collusion in optimal auctions. *Journal of Economic Theory 136*(1), 302–323.

Edlin, A. S. and C. Shannon (1998). Strict monotonicity in comparative statics. *Journal of Economic Theory 81*(1), 201–219.

Graham, D. A. and R. C. Marshall (1987). Collusive bidder behavior at single-object second-price and English auctions. *Journal of Political Economy 95*(6), 1217–1239.

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

Laffont, J.-J. and D. Martimort (1997). Collusion under asymmetric information. *Econometrica 65*(4), 875–911.

Laffont, J.-J. and D. Martimort (2000). Mechanism design with collusion and correlation. *Econometrica 68*(2), 309–342.

Mailath, G. J. and P. Zemsky (1991). Collusion in second price auctions with heterogeneous bidders. *Games and Economic Behavior 3*(4), 467–486.

McAfee, R. P. and J. McMillan (1992). Bidding rings. *American Economic Review 82*(3), 579–599.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research 6*(1), 58–73.

Pavlov, G. (2008). Auction design in the presence of collusion. *Theoretical Economics 3*(3), 383–429.

Pesendorfer, M. (2000). A study of collusion in first-price auctions. *Review of Economic Studies 67*(3), 381–411.

Quesada, L. (2005). Collusion as an informed principal problem. Working paper, Universidad Torcuato Di Tella.

Riley, J. and R. Zeckhauser (1983). Optimal selling strategies: When to haggle, when to hold firm. *Quarterly Journal of Economics 98*(2), 267–289.

Robinson, M. S. (1985). Collusion and the choice of auction. *RAND Journal of Economics 16*(1), 141–145.

Safronov, M. (2018). Coalition-proof full efficient implementation. *Journal of Economic Theory 177*, 659–677.

Segal, I. and M. Whinston (2012). Property rights. In R. Gibbons and J. Roberts (Eds.), *The Handbook of Organizational Economics*, pp. 100–158. Princeton University Press.

Wu, K. (2025). Posted price and a knockout auction: Revenue maximization under the threat of collusion. Working paper, University of Chicago.