Selling to Competitors

Pëllumb Reshidi [*] João Thereze [†]

[PRELIMINARY DRAFT. PLEASE DO NOT CIRCULATE.]
July 22, 2024

Abstract

A seller seeks to license the use of a product. She can design a mechanism to allocate licenses to one, two, or several downstream competitors with unknown productivities. How should she distribute the right to use her product? Under a set of conditions satisfied by standard competition models, we characterize the optimal mechanism: it can be implemented by an auction that licenses only the highest bidder when that bid is sufficiently higher than the rest but licenses multiple bidders otherwise. When profits depend only on a competitor's own type and the number of adversaries—but not the adversaries' productivities—we show that allocative inefficiencies are always present. Moreover, we show how the shape of virtual valuations generates under- or over-provision of licenses compared to the first-best.

[*] Department of Economics, Duke University, Durham, NC 27710—preshidi@princeton.edu
[†] Fuqua School of Business, Duke University, Durham, NC 27708—joao.thereze@duke.edu

1 Introduction

Consider an inventor contemplating whether to license her invention to two candidate firms, who compete in a downstream market. The benefit of the invention to their business is private information for each firm, but their gain is larger if they can adopt the innovation on their own. Although licensing the patent is freely replicable, licensing both firms will drive their expected profits down through competition. Should the inventor commit to scarcity and run an auction? Should she post a price for the license? This paper studies the revenue-maximizing mechanism for the seller in this and similar environments.

In multiple markets —e.g. the sale of information, the licensing of franchises, and government procurement— sellers face similar trade-offs as the inventor. In particular, there are three key elements that characterize the setting above: there are externalities among buyers, the seller can freely replicate the good to be allocated, and buyers have private information. Despite the ubiquity of such environnents, the optimal strategy for the seller remains unknown. We aim to bridge this gap.

First, we study a baseline model with two ex-ante symmetric buyers, in which their profits are determined exclusively by their private type when granted an exclusive license. To address the impact of competitive externalities, we model duopoly profits—relevant when licenses are issued to two firms—–as a proportion of monopoly profits. This basic setup is effective in representing various relevant market scenarios and ensures considerable analytical simplicity. Further, it highlights the key assumption in the paper: buyers have private information about their valuation, but the market structure is common knowledge, so the seller knows how to map private types to the valuation obtained when multiple buyers recieve the good. We later extend this model to embrace more intricate interdependencies.

Our first main result is that information asymmetries give rise to novel inefficiencies not encountered in traditional auction setups. In particular, we show that the seller might under-provide or over-provide the good—selling to fewer or more buyers than what is prescribed in the symmetric information allocation. These inefficiencies are ubiquitous: the optimal mechanism is efficient if and only if the distribution of buyer's types is in the Pareto family. Further, we link the shape of the initial distribution of buyers' types to the type of inefficiency—that is, whether the product is under/over-provided. In standard auctions with symmetric buyers, inefficiencies arise only if virtual valuations are either non-increasing or negative. In particular, as long as virtual valuations are monotonic, the most valuable bidder for the auctioneer is the same whether information is private or not. When the auctioneer can sell to both agents, however, their optimal strategy is determined by the ratio of valuations under symmetric information, but by the ratio of their *virtual* valuations when information is private. It is the gap between these two ratios that characterizes inefficiencies in our model.

Next, we identify the optimal mechanism that maximizes buyer profits and propose a direct implementation method. The optimal mechanism allocates the good to only one agent when his private valuation is much larger than the other one, and it allocates to both otherwise. To implement that mechanism in dominant strategies, we introduce what we term a *threshold auction*. In a threshold auction, each buyer submits a bid, and around these bids, specific neighborhoods are defined. If a bid is below the neighborhood of a competitor's bid, that buyer is excluded and incurs no cost. If a bid falls within the neighborhood of a competitor's bid, both buyers are awarded the good, and pay the lowest price that is consistent with them being in their competitor's neighborhood. Finally, if a bid exceeds the neighborhood of a competitor's bid, that bidder alone is awarded the good paying a premium for it. Thus, in such an auction, it is not solely the highest bid that is crucial; the entire distribution of bids influences the outcome. If the bids cluster closely, multiple licenses are awarded; if they are widely dispersed, only the highest bidder receives a license. Despite the complexities involved, we find the implementation process to be straightforward, which makes its real-world application feasible and practical.

; specifically, that profits depend not just on a buyer's own type but also on that of their competitor. This expanded framework requires additional assumptions, which we identify.

Finally, we extend the model introduced in the initial part of the paper and provide extra conditions under which the mechanism above is optimal. In this generalized model, when multiple buyers win the auction, their payoffs—their profits in the post-auction market—are influenced not only by their private types but also by the types of their competitors. This addition introduces correlations among buyers' valuations, moving our analysis away from a framework of independent private values to one that encompasses the complexities of common value auctions. A key factor aiding our analysis is the negative correlation between buyer types—whereby a higher competitor type leads to lower profits. This natural feature of competition mitigates the winner's curse when the good is allocated to a single agent. When the allocation awards both agents, the winner's curse poses difficulties for one to find the optimal mechanism. Nevertheless, we provide sufficient conditions, satisfied in several standard competitive models, that guarantee the optimality of our mechanism.

The rest of the paper is structured as follows. In section 2 we go through the baseline model, identify new inefficiencies that arise in this model, and characterize the revenue-maximizing mechanism and a natural dominant strategies implementation. In section 3.2 we expand the model to allow for correlated values and again identify the optimal mechanism as well as the implementation. We finish by considering an application which needs the more general framework in section 4, and conclude in section 5.

1.1 Related Literature

This paper relates to several strands of the literature.

Mechanism Design with Externalities First, it contributes to the literature on mechanism design with externalities. It closely relates with the works of Jehiel et al. (1996) and Jehiel et al. (1999), which delve into multidimensional settings where the market structure is unknown to the seller. In contrast, in our setup, the seller knows the market structure but lacks knowledge of the realized values of the types of buyers. In essence, we model the market structure as a function of all buyers' types, which in turn reduces the relevant dimensionality. While this distinction makes our setup more specialized, it enhances tractability and enables a comprehensive characterization of the optimal mechanism. Additionally, our approach diverges from this existing literature by permitting multiple sales of the good, as opposed to selling a single unit.

Optimal Licensing Our work also aligns with the existing body of work on optimal licensing, as seen in the studies by Kamien et al. (1992), Katz and Shapiro (1986), Jehiel and Moldovanu (2000), and Sen (2005); Sen and Tauman (2007). However, our approach diverges from these studies in two significant ways. First, the existing literature typically assumes no ex-ante uncertainty regarding the types of buyers. In contrast, in our setup, while the distribution of buyers' types is common knowledge, their realized values are private. Second, this body of work identifies optimal licensing within a constrained set of mechanisms, such as determining the optimal price to charge. In contrast, we identify the optimal mechanism with no such restrictions.

Auctions with Common Values Finally, our work also relates to the literature on auctions with common values, including the classic studies by Milgrom and Weber (1982) and Bulow and Klemperer (1994), as well as more recent approaches that identify the optimal mechanism under specific setups, such as Bergemann et al. (2020). Our work differs from this existing body of literature in two important ways. First, we allow for multiple goods to be sold, as opposed to a single good. Second, in our setup, buyers' types are negatively correlated. This negative correlation aids in characterizing the optimal mechanism.

2 The Setup

An auctioneer has an item to sell to N potential buyers, indexed in $\mathcal{N} = \{1, ..., N\}$. This item differs from typical commodities in two respects. First, it generates externalities: buyers' valuations of the product depend on who else purchases it. Second, the item can be replicated at zero cost—allowing the seller to sell multiple copies. Consider the 2^N possible subsets of $\mathcal{N}$, and let the k^{th} subset be denoted by $\mathcal{J}_k$. The cardinality of $\mathcal{J}_k$ is represented by n_k. For any subset $\mathcal{J}_k \subseteq \mathcal{N}$, the utility of buyer i when all members of $\mathcal{J}_k$ receive the item is given by

$$u(\theta_i, \mathcal{J}_k) = \theta^i \alpha_k^i.$$

We can think of α_k^i as the market share controlled by each buyer after the sale. For agents who do not purchase the item, their share is normalized to zero i.e., when $i \notin \mathcal{J}_k$, to $\alpha_k^i = 0$. Thus, utilities are characterized by a benefit from purchasing the good, θ_i, and a flexible market-share vector, α^i. By stacking the α^i vectors, we form a matrix A with dimensions $2^N \times N$. Our main assumption is that A is common knowledge, reflecting the market structure in the post-allocation stage, while each agent's taste for the good is private information. Each θ_i is assumed to be independently drawn from a regular distribution with cumulative distribution F with full support. Additionally, we assume $\alpha_k^i \in [0, 1]$ for all i and k. Utilities are assumed to be quasilinear in money.

An allocation is a distribution over subsets of $\mathcal{N}$, and due to replicability, the auctioneer can supply any of these subsets. Let $\mathcal{J}$ represent the set of all such subsets. Given this setup, the revelation principle applies, allowing us to focus on finding the truthful direct revelation mechanism that maximizes revenue.

2.1 First Best Allocation

We start our analysis under the assumption of symmetric information. If the principal knows the vector $\theta = (\theta^1, ...)$, she chooses transfers τ^i and an allocation σ_k to solve:

$$\max_{\sigma \in \Delta\mathcal{J}, \{\tau^i\}_{i=1,\ldots,N}} \quad \sum_i \tau^i$$

$$\text{s.t.} \quad \theta^i \sum_k \sigma_k \alpha_k^i - \tau^i \geq 0 \qquad \text{for all } i = 1, ..., N \tag{IR}$$

It is clear that (IR) must hold with equality in any solution. Thus, the problem can be simplified to

$$\max_{\sigma \in \Delta\mathcal{J}} \quad \sum_i \theta^i \sum_k \sigma_k \alpha_k^i.$$

Because σ is a distribution, it must be that

$$\operatorname{supp} \sigma \subset \arg\max_k \left\{ \sum_i \theta^i \alpha_k^i \right\}.$$

Then, without loss of generality, we can then focus on deterministic allocations.

Lemma 1. *Let k be an optimal allocation for θ, and k' be optimal for $\theta' = \left(v^i, \theta_{-i}\right)$, with $v^i > \theta_i$. Then $\alpha_k^i \leq \alpha_{k'}^i$.*

In other words, keeping all other factors constant, as the realized benefit from purchase for agent i increases, the seller optimally allocates to a group of agents that boosts the market share of agent i. In many relevant scenarios, market shares α^i decrease in the number of competitors. To capture this, define, for all $i, n = 1, .., N$, $\overline{a}_n^i = \max_k \{\alpha_k^i : n_k = n\}$ and $\underline{a}_n^i = \min_k \{\alpha_k^i : n_k = n\}$.

Assumption 1. *Market shares decrease in the number of competitors: $\overline{a}_n^i \leq \underline{a}_{n-1}^i$ for all $i = 1, ..., N$ and all $n > 1$. Moreover, we normalize market shares to be zero for non-active agents: $i \notin \mathcal{J}_k \implies \alpha_k^i = 0$.*

Lemma 2. *Let Assumption 1 hold. Let k be an optimal allocation for θ such that $i \in \mathcal{J}_k$. If k' is optimal for $\theta' = \left(v^i, \theta_{-i}\right)$, with $v^i > \theta_i$, then: $n_k \geq n_{k'}$ and $i \in \mathcal{J}_{k'}$.*

In other words, with everything else remaining constant, as the realized benefit from the purchase for agent i increases, the seller optimally chooses to sell to a set of agents that is no larger. Next, we consider the case where the market share is determined solely by the number of competitors and is equal among all active agents. Define $I^n = \{k : n_k = n\}$.

Assumption 2. *For all subsets with the same number of agents, market shares are the same. For all* $k, k' \in I^n$, $i \in \mathcal{J}_k, i' \in \mathcal{J}_{k'}$, $\alpha_k^i = \alpha_{k'}^{i'} \equiv a_n$.

Let $\theta^{(m)}$ be the m-th largest element of θ, and let (m) be the individual associated with that valuation.[1]

Lemma 3. *Let Assumption 2 hold, and let k be an optimal allocation for* θ. *Then, generically,* $\mathcal{J}_k = \{(1), (2), ..., (n_k)\}$.

The lemma above implies that if, in the optimal allocation, k agents are active, they will be the k agents with the highest realized θ values. Define $f(n) = a_n \cdot n$. f is the size of the market share for subsets with n active agents under Assumption 2.

Lemma 4. *Let Assumption 1 and Assumption 2 hold.* $n \in \{1, ..., N\}$ *is the cardinality of an optimal allocation for some realization* θ *if and only if* $f(n) \geq \max_{z \leq n-1} f(z)$.

This lemma implies that if it is optimal to serve n agents, then the total market size must be larger than the market size in any possible allocation with fewer agents. The intuition for this is rather straightforward. If we can achieve a higher total market share with fewer agents, then we can select a subset of the agents with higher θ values and, thus, increase total transfers.

The above analysis characterizes the optimal allocation under complete information. The seller will allocate the goods to a subset of agents. These will be the agents with the highest realized θ values, and the seller will charge them exactly their utility values, thus receiving all surplus. The problem faced by the principal can be easily visualized for $N = 2$, which we do in Figure 1 below. In this case, we normalize the payoff of being allocated the good alone to θ^i. Furthermore, we let $\alpha\theta^i$ be the payoff of being allocated the good when one's competitor also receives it.

[1] Under continuous distributions, there is, generically, exactly one such individual.

Figure 1: First Best Allocation Examples

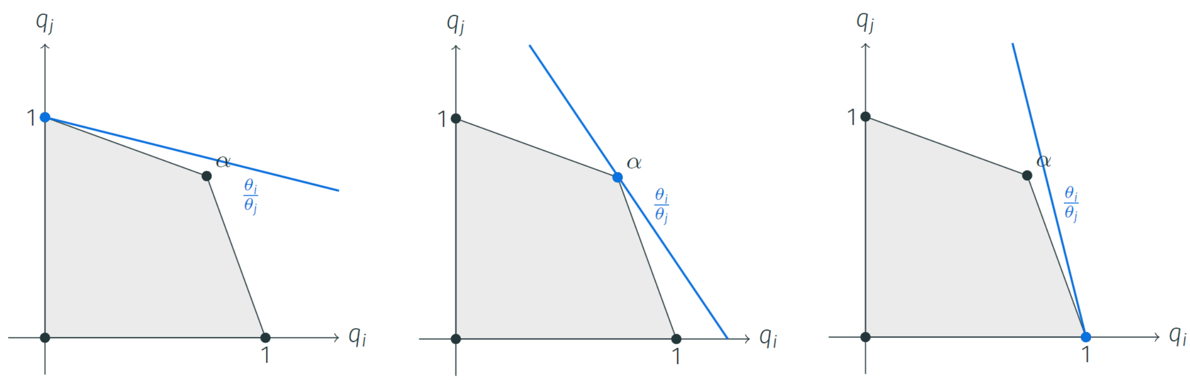

Notes: The figure above displays the first-best allocation for different ratios of θ_i/θ_j. In the example in the *left*, it is optimal to allocate the good exclusively to agent j. In the example in the *middle*, it is optimal to allocate the good to both agents. Finally, in the example displayed on the *right*, it is optimal to allocate only to agent i.

As displayed in the figure above, it is optimal to sell to i alone if

$$\theta_i \geq \alpha(\theta_i + \theta_j) \quad \rightarrow \quad \frac{\theta_i}{\theta_j} \geq \frac{\alpha}{1-\alpha}.$$

It is optimal to sell to both i and j if

$$\alpha(\theta_i + \theta_j) \geq \theta_i \quad \textit{and} \quad \alpha(\theta_i + \theta_j) \geq \theta_j \quad \rightarrow \frac{\alpha}{1-\alpha} \geq \frac{\theta_i}{\theta_j} \geq \frac{1-\alpha}{\alpha}.$$

And finally, it is optimal to sell to j alone if

$$\alpha(\theta_i + \theta_j) < \theta_j \quad \rightarrow \quad \frac{1-\alpha}{\alpha} > \frac{\theta_i}{\theta_j}.$$

The optimal allocation is driven by the ratio of valuations θ_i/θ_j, which in the figure above serves the role of an iso-profit curve. The principal thus attempts to increase the iso-profit curve as much as possible while still operating within the feasibility constraint, which, as we elaborate in the next section, is a polytope.

2.2 Optimal Mechanism

Next, we characterize the profit-maximizing mechanism when the seller of the good does not observe the realized θ values of the buyers. Our first observation is that we can change the allocation space from $\Delta\mathcal{J}$ to an interval in $\mathbb{R}$. To see this, start with any allocation $\sigma \in \Delta\mathcal{J}$. This allocation leads to the following expected utility for agent i:

$$\mathbb{E}_\sigma[u(\theta, \mathcal{J}_k)] = \theta \underbrace{\sum_k \sigma_k \alpha_k^i}_{q^i(\sigma)}$$

We call $q^i(\sigma)$ an assignment. Let $q(\sigma)$ be the vector of assignments. Then, if $\Delta\mathcal{J}$ is the set of possible allocations, we can define the associated assignment set as

$$\mathcal{Q} = \left\{ q \in \mathbb{R}^N : \exists \sigma \in \Delta\mathcal{J}, q = q(\sigma) \right\}.$$

Define $\alpha_k = \left(\alpha_k^1 \alpha_k^2, \ldots, \alpha_k^N\right)$. It is clear that:

Lemma 5. $\mathcal{Q} = \operatorname{co}\left\{\alpha_k : k \in \{1, \ldots, 2^N\}\right\}$. *$\mathcal{Q}$ is a convex polytope.*

For an expected market share vector q define

$$Q^i(\theta^i) = \int q^i\left(\sigma\left(\theta^i, \theta_{-i}\right)\right) dF_{-i}(\theta_{-i}),$$

and

$$U^i(\theta^i) = \theta^i Q^i(\theta^i) - \underbrace{\int \left[\sum_k \sigma_k(\theta^i, \theta_{-i}) \tau_k^i(\theta^i, \theta_{-i})\right] dF_{-i}(\theta_{-i})}_{T^i(\theta^i)}$$

The expected utility of agent i, given their realized value θ^i, is the net gains minus the expected transfer.

Lemma 6. *An allocation σ is implementable if and only if the following conditions hold:*

1. ***Monotonicity:*** *Q^i is increasing for all i;*
2. ***Envelope Condition:*** $U^i(\theta^i) = U^i(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} Q^i(v)dv$;
3. ***Individual Rationality:*** $U^i(\theta^i) \geq 0$ *for all i, θ^i;*
4. ***Feasibility:*** $q(\sigma) \in \mathcal{Q}$.

This represents the usual set of conditions for auction implementability, with the exception of feasibility. In our context, feasibility requires that trade probabilities are contained within the polytope $\mathcal{Q}$—a polytope that may extend beyond the unit simplex. Conversely, in standard auctions, these probabilities must reside within the unit simplex, as detailed in (Myerson, 1981). The problem of the principal then reduces to

$$\max_{U^i, Q^i, q^i} \quad \int \sum_i \left(\theta^i Q^i(\theta^i) - U^i(\theta^i)\right) f(\theta) d\theta$$
$$\text{s.t.} \quad 1 - 4.$$

Define the virtual valuation of a type θ agent as: $v(\theta^i) = \theta^i - \frac{1-F_i(\theta^i)}{f_i(\theta^i)}$. Following the standard integration by parts approach, the problem of the principal becomes

$$\max_{q^i} \quad \int \sum_i v(\theta^i) q^i(\theta) f(\theta) d\theta$$
$$\text{s.t.} \quad 1 \text{ and } 4.$$

We henceforth focus on the case in which Assumptions 1 and 2 hold — that is, the number of competitors allocated the good alone determines the payoff of the buyer, and the more competitors share the good, the lower is each bidders' valuation.

2.3 Inefficiencies

In traditional auction theory, asymmetric information can lead to inefficiencies in two primary ways. First, if virtual valuations are non-increasing or if agents are heterogeneous. This could result in scenarios where an agent with a lower realized type submits a higher bid and thus wins the auction, causing an ex-post inefficient allocation. The second type of inefficiency arises if the virtual values can be negative. If the realized virtual values are negative across all agents, the good remains unsold even if all agents value the good positively, which is suboptimal. In our model, we identify a novel form of inefficiency. To differentiate from traditional suboptimal outcomes, we assume that all agents draw their types from the same regular distribution F, this eliminates the first inefficiency. Additionally, by ensuring that virtual values are positive for all realizations, we eliminate the second inefficiency.

Assumption 3. *v is increasing and $v(\underline{\theta}) \geq 0$.*

Definition 1. *Let q_f be the first-best allocation. We say that a mechanism inducing allocation q under- (over-) provides if:*

$$q^i(\theta) \leq (\geq) q^i_f(\theta) \text{ for all } \theta \text{ and } i.$$

An allocation is efficient if equality holds above.

Let $h(\theta_i)$ represent the inverse hazard rate. Define $\lambda(\theta_i) \equiv h(\theta_i)\theta_i$.[2]

Proposition 1. *Let assumptions 1, 2 and 3 hold. The profit-maximizing mechanism*

- ***Is Efficient*** *if and only if λ is constant — that is, F is in the Pareto family.*
- ***Under-provides*** *if λ is increasing.*
- ***Over-provides*** *if λ is decreasing.*

The proposition posits that the profit-maximizing mechanism will prescribe the same allocation as the first-best outcome for all realized values if and only if the buyer's types are distributed according to a family in the Pareto family. To build some intuition about this result, we once again go back to an $N = 2$ example. Note that, differently from the first-best outcome, the behavior of the principal, while similar, is no longer dictated by the ratio of valuations θ_i/θ_j. Rather, the iso-profit curve is now determined by the ratio of virtual valuations $v(\theta_i)/v(\theta_j)$. There is, of course, no reason for these two ratios to be the same, especially not for any realization of θ_i and θ_j. In particular, $\frac{\theta^1}{\theta^2} = \frac{v(\theta^1)}{v(\theta^2)}$ for all vectors θ if and only if v is linear. We complete the proof by showing that v is linear if and only if F_i belongs to the Pareto family. To see that, assume $v(\theta) = \lambda\theta$, $\lambda > 0$. We then have:

$$\theta - \frac{1 - F(\theta)}{f(\theta)} = \lambda\theta.$$

Solving this differential equation yields the unique solution:

$$F(\theta) = 1 + k\theta^{-\frac{1}{1-\lambda}}.$$

The only family of CDFs satisfying this equation is the Pareto family. For any other distribution, the two ratios highlighted above will differ at least for some realizations. We show two such examples in Figure 2 below.

[2] Which can be interpreted as the price-elasticity of demand.

Figure 2: Examples of Under and Overprovision

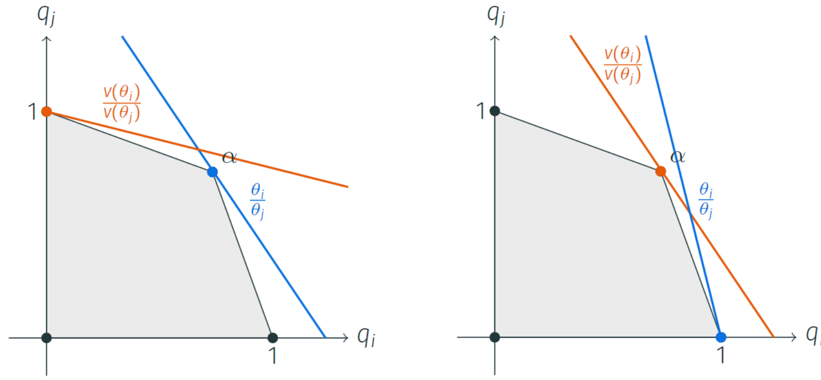

Notes: The figure above displays the profit-maximizing and the first-best allocations for different realized values of θ_i and θ_j. In the example in the *left*, it is efficient to allocate the good to both agents, but it is profit-maximizing to allocate the good to agent i exclusively—*underprovision*. In the example in the *right*, it is efficient to allocate exclusively to to agent i, but it is profit-maximizing to allocate to both—*overprovision*.

The figure above illustrates the profit-maximizing and the first-best allocations for different realized values of θ_i and θ_j. In the example on the *left*, when behavior is governed by the ratio of valuations θ_i/θ_j, it is efficient to allocate the good to both agents. However, in the case of asymmetric information, as previously discussed, behavior is driven by the ratio of valuations $v(\theta_i)/v(\theta_j)$, leading to the good being allocated exclusively to agent i as the profit-maximizing outcome. Consequently, the good is underprovided. In the example on the *right*, it is efficient to allocate the good exclusively to agent i, but profit maximization dictates allocating to both agents. Thus, the good is overprovided. The potential for overprovision and underprovision is not only theoretical; there exists a nonempty set of distributions for which either outcome is possible. Figure 3 below presents two such examples.

The figure displays the profit-maximizing and the first-best allocations for different values of θ_i and θ_j. The shaded blue/orange areas indicate the regions where the good is provided to both agents under the first-best/profit-maximizing allocation. In the example on the *left*, the shaded orange region is entirely within the shaded blue region, indicating that there are realizations of θ_i and θ_j for which both agents would receive the good under the first-best allocation, but only one agent receives it under the profit-maximizing allocation, leading to underprovision. Conversely, in the example on the *right*, the shaded blue region is entirely within the shaded orange region, indicating that there are realizations of θ_i and θ_j for which an agent would receive the good exclusively under the first-best allocation, but both agents receive it under the profit-maximizing allocation, leading to overprovision. Therefore, it is not difficult to find distributions of types for which either outcome occurs.

Figure 3: Examples of Distributions leading to Under and Overprovision

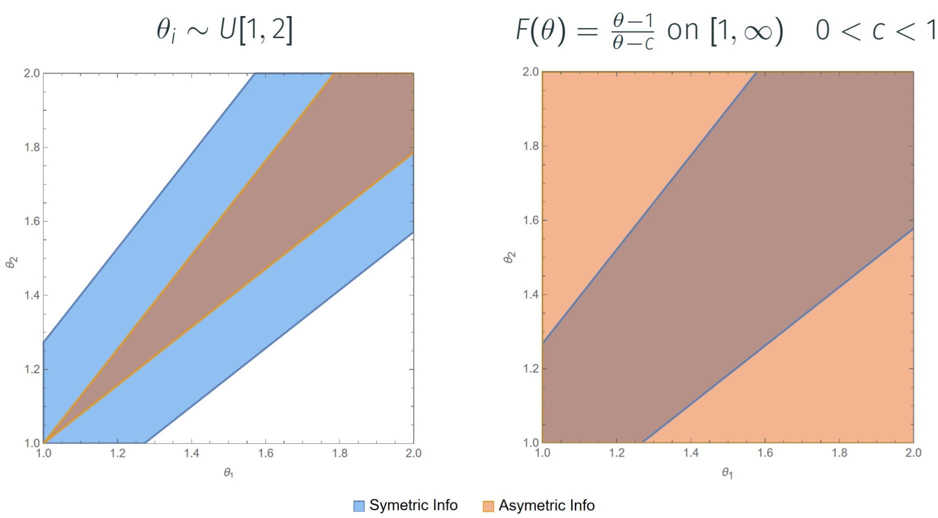

Notes: The figure above displays the profit-maximizing and the first-best allocations. The distributions utilized for each example are displayed on top of the graphs.

It is worth emphasizing once again that these inefficiencies, whether they involve under- or over-provision of the good, represent a novel inefficiency not present in standard auction examples. To underscore that typical inefficiencies are not driving these results, we have assumed that all agents draw their types from the same distribution F and that virtual values are positive for any realization. Under these two assumptions, standard auctions do not exhibit inefficiencies. Yet, in our setup, over- or under-provision may occur.

Figure 4: Relevant α Values

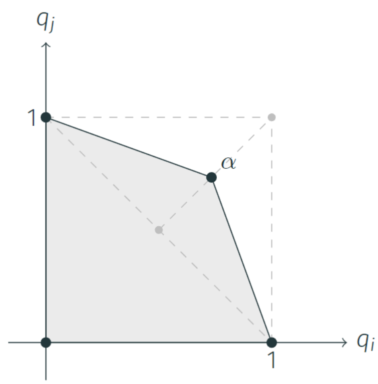

Notes: The figure above displays the relevant $\alpha \in (0.5, 1)$ values for which the problem we analyze is relevant. The value of β has been normalized to 1. If $\alpha \leq 0.5$, the principal never finds it optimal to allocate to two agents. At the other extreme, if $\alpha \geq 1$, there is no loss from providing to more agents, and therefore, the principal always provides to both.

2.4 Implementation

Next, we turn to the implementation of the optimal mechanism. In particular, we look for implementations that satisfy the following two desiderata:

1. Implements the optimal allocation *truthfully* and in dominant strategies.

2. Does not require payment from excluded agents.

Recall that for a type x, $g(x) \equiv \max\{\underline{\theta}, v^{-1}(av(x))\} < x$ is the lower-threshold for x: the type such that, if the other bidder has a valuation below $g(x)$, the other bidder is excluded. Accordingly, $g^{-1}(y)$ is the type for which y is the lower threshold.

Definition 2. *A* ***threshold auction****: for each bid b there exist thresholds* $\underline{\tau}(b) < b < \overline{\tau}(b)$ *such that*

$$q_i = \begin{cases} 1 & \text{if } b_i > \overline{\tau}(b_{-i}) \\ \alpha & \text{if } \overline{\tau}(b_{-i}) > b_i > \underline{\tau}(b_{-i}) \\ 0 & \text{otherwise} \end{cases} \qquad t_i = \begin{cases} \alpha\underline{\tau}(b_{-i}) + (1-\alpha)\overline{\tau}(b_{-i}) & \text{if } b_i > \overline{\tau}(b_{-i}) \\ \alpha\underline{\tau}(b_{-i}) & \text{if } \overline{\tau}(b_{-i}) > b_i > \underline{\tau}(b_{-i}) \\ 0 & \text{otherwise} \end{cases}$$

Proposition 2. *Under assumptions 1, 2 and 3, the optimal mechanism is implemented in dominant strategies by a threshold auction.*

In other words, the optimal allocation can be implemented truthfully and in dominant strategies without loss of revenue to the seller. The mechanism works as follows: ask both agents to bid. Assume, without loss, that $b_1 \geq b_2$. If $b_1 < g^{-1}(b_2)$, then allocate the good to both agents, who pay $\alpha g(b_{-i})$ each. If $b_1 \geq g^{-1}(b_2)$, then allocate the good to the first bidder only. This bidder pays $\alpha g(b_{-i}) + (\beta - \alpha) g^{-1}(b_{-i})$. We visualize the workings of this mechanism in Figure 5 below.

Figure 5: Threshold Auction Implementation

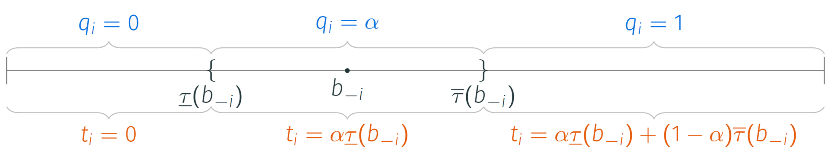

Notes: The figure above visualizes the profit-maximizing implementation via a threshold auction. Around the ibd of the opponent b_{-i} there is a n neighborhood $(\underline{\tau}(b_{-i}), \overline{\tau}(b_{-i}))$. If the agent's bid falls below this neighborhood, he is excluded and pays nothing $t_i = 0$. If his bid falls within this neighborhood, both agents are allocated the good and pay $t_i = \alpha\underline{\tau}(b_{-i})$. Finally, if an agent bid falls above this neighborhood, he is provided the good exclusively and pays $t_i = \alpha\underline{\tau}(b_{-i}) + (1-\alpha)\overline{\tau}(b_{-i})$.

Notice that in this mechanism, when both agents are allocated the product, the agent with the lowest bid pays more than the agent with the highest bid. Regardless, this does not imply incentives to increase their own bid, as their payment does not depend on their individual bid.

Figure 6: Revenue Comparison

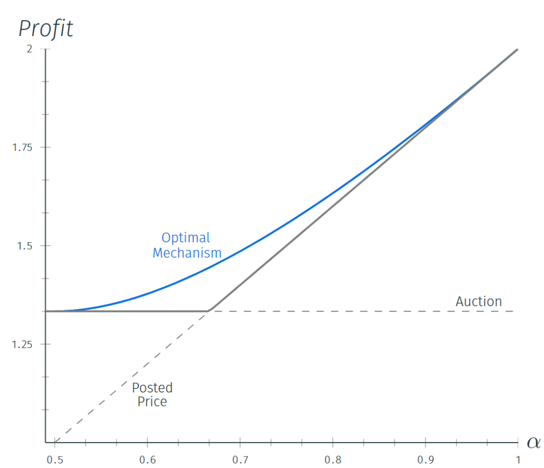

Notes: For different α values, the graph above compares the revenue from a posted price, a standard auction in which the good is sold to one buyer, and the optimal mechanism.

2.5 Revenue Comparison

In this and other applications, it is possible to compare the revenue difference from, say, a posted price, a standard auction in which the good is sold to one buyer, and the optimal mechanism identified in this paper. Figure 6 offers such a comparison. It is clear that, for any $\alpha \in (0.5, 1)$ value, threshold auctions outperform either mechanism. Importantly, when α approaches 0.5, it becomes more and more likely that it will be optimal to sell to one agent—the polytope discussed above converges to the unit simplex—thus, the profit of threshold auctions converge to the profits from a regular auction. Conversely, as α approaches 1, the externalities from having two firms active are reduced; therefore, profits converge to those of the posted price.

The comparison between the posted price, standard auction and optimal mechanism revenues is actually tighter. To establish a formal result in that respect, consider the revenue of each of these mechanisms. First, because virtual valuations are assumed to be positive, it can be proved that the optimal posted price would entail setting a price of $\alpha\underline{\theta}$, in which both agents buy the product. By standard manipulations of the virtual value function, one can show that the revenue obtained can be rewritten as:

$$R^p = \alpha \mathbb{E}\left[v\left(\theta_{(1)}\right) + v\left(\theta_{(2)}\right)\right].$$

By its turn, the revenue obtained in a standard auction, in which the designer commits to sell only one product, is given by the expected value of the second highest, which can also be written as:

$$R^a = \mathbb{E}\left[v(\theta_{(1)}\right].$$

Thus, a constrained seller who chooses between these two mechanisms would obtain revenue:

$$R^c = \max\{R^p, R^a\}.$$

Now, consider the seller who chooses an optimal mechanism. We know the seller sells to the buyer with the highest realization if $v(\theta_{(1)}) \geq \alpha\left(v(\theta_{(1)}) + v(\theta_{(2)})\right)$. By the virtual-valuation representation of the seller's revenue, in that case the seller's revenue is exactly $v(\theta_{(1)})$. This simple logic establishes the following proposition, which states that the difference between the unconstrained and the constrained revenues is precisely quantified by a Jensen gap.

Proposition 3. *Under assumptions 1, 2 and 3, the difference between the optimal revenue, R and the revenue constrained to posted prices and standard auctions is:*

$$
\begin{aligned}
R - R^c = {} & \mathbb{E} \max \left\{ v(\theta_{(1)}), \alpha \left(v(\theta_{(1)}) + v(\theta_{(2)}) \right) \right\} \\
& - \max \left\{ \mathbb{E} \left[v(\theta_{(1)}) \right], \alpha \mathbb{E} \left[v(\theta_{(1)}) + v(\theta_{(2)}) \right] \right\}
\end{aligned}
$$

3 General Setup

We began our analysis in Section 2 with a simplified model that enabled us to visualize many of the inner workings of the mechanism and thus helped us develop intuition. This initial model was advantageous as it did not require any assumptions beyond those found in standard auction theory. Although this model is adequate for some applications, its limitations are clear. Notably, we assumed that when transitioning from being awarded the good exclusively to sharing the market with another, the buyer's profits were simply multiplied by constant $\alpha < 1$. However, often the extent to which competition reduces a firm's profit depends on the private information of one's competitor. For instance, if our competitor's production costs are extraordinarily high, then α should be near 1, as they pose little threat to our market dominance. Conversely, if our competitor is very efficient, they might capture most of the market for themselves, warranting a lower α. In other words, when the market is shared, we want outcomes for agent i to depend not only on θ_i but also on the type of agent j, namely θ_j. This section is dedicated to this goal.

We proceed in two steps. In Section 3.1, we begin by analyzing a setup in which the dependence is assumed to be multiplicative, while in Section 3.2, we analyze the most general model, allowing for any dependencies between agents' types and payoffs. In Table 1 below, we report the additional sufficient conditions needed for the analysis of these alternative models. As can be seen, we require one additional assumption for the multiplicative model, while to ensure that our analysis goes through for a general model, we need two additional assumptions. We derive and provide intuition for these assumptions in the sections below.

3.1 Multiplicative Model

We modify the setup in the following way. Let $N = 2$, we have that individual's types, θ are i.i.d. drawn from the same distribution F. If agent i, with type θ_i, is allocated

Table 1: Sufficient Conditions

Model	Win Alone	Win Together	Additional Assumptions
α Model	θ_i	$\alpha_i\theta_i$	$\varnothing$
Multiplicative	θ_i	$h(\theta_j)\theta_i$	$\frac{v'(\theta_i)}{v(\theta_i)} \geq \frac{[1-h(\theta_i)]'}{1-h(\theta_i)}$
General	θ_i	$g(\theta_i, \theta_j)$	$v'(\theta_i) \geq v_{g,1}(\theta_i, \theta_j) + v_{g,2}(\theta_j, \theta_i) \geq 0$ $1 - g_1(\theta_i, \theta_j) \geq 0$

Notes: The table above reports the additional assumptions ensuring the analysis goes through.

the product alone, his payoff is $\theta_i\beta$. On the other hand, if both agents are allocated the product, their payoffs are $\theta_i\alpha(\theta_{-i})$, for some decreasing function $\alpha(\theta_{-i})$, with $\beta \geq \alpha \geq \frac{\beta}{2}$.

In this setting, the allocation set can be enumerated as $k \in \{0, 1, 2, 3\}$, where $k \in \{1, 2\}$ means that agent i is allocated alone, $k = 0$ that nobody receives the good, and $k = 3$ that both agents are allocated. Then, define for each realization $\vartheta = (\theta_1, \theta_2)$:

$$\mathcal{P}(\vartheta) = \{q \in \mathbb{R}^2 : \exists \gamma \in \Delta\{0, 1, 2, 3\}, q_i = \gamma_i\beta + \gamma_3\alpha(\theta_{-i}), i = 1, 2\}$$

Just as before, it is easy to see that $\mathcal{P}(\vartheta)$ is a polytope, representing feasible expected payoffs given a type realization. Note that, in the space of expected payoff allocations, buyers' expected utilities can be written as:

$$U_i(\vartheta) = \theta_i q_i(\vartheta) - t(\vartheta)$$

We can then follow the argument in Bulow and Klemperer (1996), Lemma 3 to conclude that the principal solves:

$$\max_q \quad \mathbb{E}_\vartheta\left[\sum_i v(\theta_i)q_i(\vartheta)\right]$$

$$\text{s.t.} \quad q(\vartheta) \in \mathcal{P}(\vartheta)$$

$$\mathbb{E}_{\theta_{-i}}[q_i(\theta_i, \theta_{-i})] \text{ increasing in } \theta_i, \text{ for } i = 1, 2$$

where $v(\theta_i) = \theta_i - \frac{1-F(\theta_i)}{f(\theta_i)}$ is the standard virtual valuation. We start by relaxing the monotonicity assumption. Define $a_i(\vartheta) = \frac{\alpha(\theta_i)}{\beta-\alpha(\theta_{-i})}$. Note that a_i is decreasing in θ_i and in θ_{-i}. Then, it is easy to see that the optimal solution to the seller's problem, $\overline{q}$ solves:

$$\overline{q_i}(\vartheta) = \begin{cases} 0 & \text{if } v(\theta_{-i}) > a_{-i}(\vartheta)v(\theta_i), \\ \alpha(\theta_{-i}) & \text{if } \frac{v(\theta_i)}{a_i} \leq v(\theta_{-i}) < a_{-i}(\vartheta)v(\theta_i), \\ \beta & \text{if } v(\theta_{-i}) < \frac{v(\theta_i)}{a_i}. \end{cases}$$

We now find conditions under which $\overline{q}$ satisfies the monotonicity condition. Define $\overline{F} = F \circ v^{-1}$, $\overline{\alpha} = \alpha \circ v^{-1}$, and we have:

$$\mathbb{E}_{\theta_{-i}}[\overline{q}_i(\theta_i, \theta_{-i})] = \int_{\frac{v(\theta_i)}{a_i(\vartheta)}}^{a_{-i}(\vartheta)v(\theta_i)} \overline{\alpha}(z)\overline{f}(z)dz + \overline{F}\left(\frac{v(\theta_i)}{a_i(\vartheta}\right)\beta$$

Differentiating with respect to θ_i we obtain:

$$\overline{\alpha}(a_{-i}(\vartheta)v(\theta_i))\overline{f}(a_{-i}(\vartheta)v(\theta_i)) \qquad \underbrace{\left(\frac{da_{-i}(\vartheta)}{d\theta_i}v(\theta_i) + a_{-i}(\vartheta)v'(\theta_i)\right)}_{>0\text{: conditional on being the highest type you cannot lose}}$$

$$+\left(\beta - \overline{\alpha}\left(\frac{v(\theta_i)}{a_i(\vartheta)}\right)\right)\overline{f}\left(\frac{v(\theta_i)}{a_i(\vartheta)}\right) \qquad \underbrace{\left(\frac{v'(\theta_i)}{a_i(\vartheta)} - \frac{da_i(\vartheta)}{d\theta_i}\frac{v(\theta_i)}{a_i^2(\vartheta)}\right)}_{><0\text{? Conditional on being lowest you may lose}}$$

A sufficient condition for this expression ot be positive is that $\frac{v'(\theta_i)}{v(\theta_i)} \geq \frac{[\beta-\alpha(\theta_i)]'}{\beta-\alpha(\theta_i)}$, that is, virtual valuations are *more increasing* than $\beta - \alpha(\theta_i)$.

3.2 General Model

Let $N = 2$. Individuals have types θ i.i.d drawn from the same distribution $F \in \Delta\Theta$, and Θ is an interval of real numbers. If agent i is allocated the product alone, her value for the product is $\beta(\theta_i, \theta_{-i})$. If agents share the product, agent i's utility is $\alpha(\theta_i, \theta_{-i})$. Define $\gamma = (\beta, \alpha) \in \mathbb{R}^2$. A symmetric allocation is a triple of functions $\{q^i\}_{i=1,2}$, $q_\alpha : \Theta \times \Theta \to [0, 1]$, such that, q_α is symmetric and, for each realization $\theta, \nu \in \operatorname{supp} F$:

$$q^1(\theta_1, \theta_2) + q^2(\theta_1, \theta_2) + q_\alpha(\theta_1, \theta_2) \leq 1 \tag{F}$$

We interpret q^i as the probability that i is allocated alone, given θ_1, θ_2, and q_α to be the probability they are allocated together. We define $q_i = (q^i, q_\alpha)$. In a truthfully revealing direct mechanism, the expected utility of agent i with type θ is:

$$U_i(\theta) = \mathbb{E}\left[\gamma\left(\theta, \theta_{-i}\right) \cdot q_i\left(\theta, \theta_{-i}\right) - t\left(\theta, \theta_{-i}\right)\right]$$

We can then write the Bayesian incentive compatibility constraints as:

$$U_i(\theta) - U_i(\theta') \geq \mathbb{E}\left[\left(\gamma(\theta, \theta_{-i}) - \gamma(\theta', \theta_{-i})\right) \cdot q_i(\theta', \theta_{-i})\right],$$

for all θ, θ'. As usual, we say that an allocation is implementable if it satisfies Bayesian Incentive Constraints.

Lemma 7. *An allocation $\{q^1, q^2, q_\alpha\}$ is implementable only if:*

1. $U_i(\theta) = U_i(\underline{\theta}) + \int_0^\theta \mathbb{E}[\gamma'(\nu, \theta_{-i}) \cdot q_i(\nu, \theta_{-i})]d\nu$ *for all* $\theta \in \Theta$;
2. $\mathbb{E}\left[\left(\gamma(\theta, \theta_{-i}) - \gamma(\theta', \theta_{-i})\right) \cdot \left(q_i(\theta, \theta_{-i}) - q_i(\theta', \theta_{-i})\right)\right] \geq 0$ *for all* θ, θ'.

The lemma above provides necessary conditions for implementability, but these conditions are, in general, not sufficient. We now provide sufficient conditions.

Assumption 4. *Increasing differences: The difference between monopolist and duopolist profits is increasing in own-type:* $\beta'(\theta, \theta_{-i}) - \alpha'(\theta, \theta_{-i}) \geq 0$.

Proposition 4. *Let Assumption 4 hold. When q_i and $q_i + q_\alpha$ are increasing, condition 1 in Lemma 7 is sufficient for implementability.*

We can write expected transfers as $\mathbb{E}[\gamma(\theta, \theta_{-i}) \cdot q(\theta, \theta_{-i}) - U_i(\theta)]$. Using the usual integration by parts trick, we obtain that profits are:

$$\sum_i \int_\theta \mathbb{E}_{-i}\left[\left(\gamma(\theta, \theta_{-i}) - \frac{1 - F(\theta)}{f(\theta)}\gamma'(\theta, \theta_{-i})\right) \cdot q_i(\theta, \theta_{-i})\right] f(\theta) d\theta \tag{1}$$

Assumption 5. *We make the following assumptions on virtual valuations:*

Strong Regularity. *$v_\beta(\theta, \nu), v_\alpha(\theta, \nu)$ are increasing in θ for all ν.*

Virtual Gains. $v_\beta'(\theta, \nu) \geq v_\alpha'(\theta, \nu) + v_\alpha'(\nu, \theta) \geq \max\{0, v_{\beta,\nu}(\theta, \nu)\}$

Proposition 5. *Under Assumption 4 and Assumption 5, the revenue-maximizing mechanism has allocations:*

$$q_i(\theta, \theta_{-i}) = \begin{cases} 1 & \textit{if} \;\; v_\beta(\theta, \theta_{-i}) > \max\left\{v_\alpha(\theta, \theta_{-i}) + v_\alpha(\theta_{-i}, \theta), v_\beta(\theta_{-i}, \theta)\right\} \\ 0 & \textit{otherwise} \end{cases} \tag{2}$$

$$q_\alpha(\theta, \theta_{-i}) = \begin{cases} 1 & \textit{if} \;\; \max\{v_\beta(\theta, \theta_{-i}), v_\beta(\theta_{-i}, \theta)\} < v_\alpha(\theta, \theta_{-i}) + v_\alpha(\theta_{-i}, \theta) \\ 0 & \textit{otherwise} \end{cases} \tag{3}$$

Proposition 5 completes the characterization for the general model. Next, we explore applications that are now feasible but would have been unmanageable with the baseline model.

4 Applications

4.1 Application of the General Model: Hotelling Example

Consider a uniform distribution of consumers in the interval $[0, 1]$. Two potential franchisees are positioned at the ends. A franchisor, henceforth referred to as the principal, contemplates licensing a franchise to the franchisees, henceforth referred to as firms, positioned at 0, to the one positioned at 1, or to both of them. Each firm has private information regarding the quality of the products they will be able to offer. Let these qualities be uniformly distributed $q_j \sim U[\underline{q}, \overline{q}]$, with $j \in \{0, 1\}$, where j indicates their position in the interval. If a customer decides to purchase a good from a firm, say from firm $j = 0$, their utility will be $q_j - p_j - \delta x$, where p_j represents the price the firm charges, δ represents the travel costs, while x represents the consumer's position in the unit interval.

If the principal decides to license a franchise to only one firm, say $j = 0$, then this firm will be a monopolist. To find the profit-maximizing price, we first need to find the marginal consumer, the last consumer who justifies the travel cost. This will be the consumer positioned at $\tilde{x}$, where $\tilde{x} = \{x | q_j - p_j \delta x = 0\}$. The firm then maximizes $\max_{p_j} p_j \tilde{x}(p_j)$, and finds it optimal to charge $p_j^M = \frac{q_j}{2}$, where M represents their monopolistic status. The marginal consumer will thus be $\tilde{x}(p_j^M) = \frac{q_j}{2\delta}$, while the firms profits will be $\pi_j^M = \frac{q_j^2}{4\delta}$.

If the principal opts to grant franchises to both firms, then consumers compare the quality, price, and distance from each firm before deciding which one to buy from. To the buyers, this is the externality caused by providing two franchises. Although a franchise can be replicated at no cost, it intensifies competition, which may reduce profits by driving down the prices, leading to lower bids and potentially decreased profitability. With two active firms, the marginal client, the client indifferent from purchasing from $j = 0$ or $j = 1$, is

$$\tilde{x} = \left\{x \middle| q_0 - p_0 - \delta x = q_1 - p_1 - \delta(1-x)\right\} \qquad \rightarrow \qquad \tilde{x} = \frac{(q_0 - p_0) - (q_1 - p_1) + \delta}{2\delta},$$

Each firm then maximizes $\max_{p_j} p_j \tilde{x}(p_j, p_{-j})$, leading to the following optimal prices

$$p_0^D = \frac{q_0 - q_1 + 3\delta}{3}, \qquad p_1^D = \frac{q_1 - q_2 + 3\delta}{3}.$$

And duopoly profits of

$$\pi_0^D = \frac{(q_0 - q_1 + 3\delta)^2}{18\delta}, \qquad \pi_1^D = \frac{(q_1 - q_0 + 3\delta)^2}{18\delta}.$$

Importantly, note that the duopoly profits are not simply a fraction α of the monopoly profits, nor can they be expressed in a multiplicative form as a function of the competitors type q_{-j}. Thus, the machinery developed in section 3.2 is necessary to handle this example. It is trivial to check that for the right $\underline{q}$, $\overline{q}$, and δ parameters, all-sufficient conditions specified in section 3.2 are met. Thus, the principal can maximize expected profits by simply running a threshold auction.

5 Conclusions

This paper investigates the optimal mechanism for selling a replicable good with externalities relevant to a variety of scenarios—franchise operations, patent licensing, and information sales, to name a few. Stemming from asymmetric information, we uncover unique inefficiencies that do not arise in conventional auctions, that can lead a seller to either overprovide or underprovide the good. These inefficiencies are closely tied to the initial distribution of buyers. We identify both the optimal mechanism and a straightforward implementation strategy in practical settings. We propose a *threshold auction*, where the decision to sell is influenced not just by the highest bid but also by the overall distribution of bids. In essence, if bids are closely grouped, it is advantageous to sell the good to multiple bidders; conversely, if bids are widely dispersed, selling exclusively to the highest bidder maximizes profits.

6 Appendix

6.1 Proofs

Proof of Lemma 1

By optimality of k':

$$v^i\alpha^i_{k'} + \underbrace{\sum_{j\neq i}\theta^j\alpha^j_{k'}}_{u_{k'}} \geq v^i\alpha^i_k + \underbrace{\sum_{j\neq i}\theta^j\alpha^j_k}_{u_k}.$$

Similarly, by optimality of k

$$\theta^i\alpha^i_{k'} + u_{k'} \leq \theta^i\alpha^i_k + u_k.$$

Summing up the two inequalities leads to

$$(v^i - \theta^i)\left(\alpha^i_{k'} - \alpha^i_k\right) \geq 0.$$

Because $v^i > \theta^i$, we conclude the result. ■

Proof of Lemma 2

The proof follows from Lemma 1. ■

Proof of Lemma 3

Let $\mathcal{J}_k$ be an optimal allocation. The auctioneer's profits are:

$$a_{n_k}\sum_{i\in\mathcal{J}_k}\theta^i \leq a_{n_k}\sum_{i\in 1,\ldots,n_k}\theta^{(i)},$$

with equality if and only if $\mathcal{J}_k = \{(1),\ldots,(m)\}$. The result follows. ■

Proof of Lemma 4

First, assume $f(n) < \max_{z\leq n-1} f(z)$. By this assumption, there exists $m \leq n-1$ with $f(n) < f(m)$. From Lemma 3 (Assumption 2), we know that the optimal way for the principal to allocate to n (m) is to serve the buyers with the top n (m) valuations. We then have that profits satisfy:

$$a(n)\sum_{i=1}^{n}\theta^{(i)} = f(n)\frac{\sum_{i=1}^{n}\theta^{(i)}}{n} < f(m)\frac{\sum_{i=1}^{m}\theta^{(i)}}{m} = a(m)\sum_{i=1}^{m}\theta^{(i)},$$

where the inequality comes from $f(m) > f(n)$ and the fact that the average of the highest n types is smaller than the average of the highest $m > n$ types. We have then proved that the auctioneer can profit from reducing the number of agents in the allocation.

Conversely, assume $f(n) \geq \max_{z\leq n-1} f(z)$. Consider the profile $(\underbrace{\theta, ...\theta}_{\text{n times}}, \overbrace{0, ..., 0}^{\text{(N-n) times}})$. Thus, for any $m < n$:

$$a(m)\sum_{i=1}^{m}\theta^{(i)} = \theta a(m)m < \theta a(n)n = a(n)\sum_{i=1}^{n}\theta^{(i)},$$

where the inequality follows from the assumption on $f(n)$. Thus, the principal chooses no allocation with less than n agents. At the same time, for $m > n$, the principal's profits are:

$$a(m)\sum_{i=1}^{m}\theta^{(i)} = \theta a(m)n < \theta a(n)n = a(n)\sum_{i=1}^{n}\theta^{(i)},$$

where the inequality follows from Assumption 1. Therefore, the optimal allocation includes exactly n agents. ■

Proof of Lemma 5

It follows from the definition that each $q(\sigma)$ is a convex combination of α_k. Because Q is the convex hull of a finite set of points, it is a polytope by the vertex description of polytopes.

Proof of Proposition 1

Given Assumption 2, we can parameterize the problem by α, $\beta > 0$, where α is the payoff the multiplier when bot agents are served, while β is the multiplier when they are the only ones receiving the product. By Assumption 1, $\alpha \leq \beta$, and the payoff of not receiving the product is zero.

By ignoring constraint 1, the problem of the principal is a linear programming problem, which can be solved by an extreme point of the polytope Q: that is, by a degenerate allocation. Moreover, the problem can be solved realization by realization.

Fix θ and assume, without loss of generality, $\theta^1 \geq \theta^2$. By Assumption 3, virtual valuations are positive, so at least one agent is served, thus any allocation includes buyer 1. Then, in the first best allocation—under symmetric information—the principal serves both agents only if

$$\alpha(\theta^1 + \theta^2) \geq \beta\theta^1 \iff \frac{\alpha}{\beta - \alpha} \geq \frac{\theta^1}{\theta^2}. \tag{4}$$

By contrast, the optimal mechanism serves both agents only if

$$\alpha\left(v(\theta^1) + v(\theta^2)\right) \geq \beta v(\theta^1) \iff \frac{\alpha}{\beta - \alpha} \geq \frac{v(\theta^1)}{v(\theta^2)}. \tag{5}$$

Therefore, the allocation is efficient for all vectors θ if and only if $\frac{\theta^1}{\theta^2} = \frac{v(\theta^1)}{v(\theta^2)}$, for all $\theta^2 \leq \theta^1$, which happens if and only if v is linear. We complete the proof by showing that v is linear if and only if F_i is the Pareto distribution. To see that, assume $v(\theta) = \lambda\theta$, $\lambda > 0$. We then have:

$$\theta - \frac{1 - F(\theta)}{f(\theta)} = \lambda\theta.$$

Solving this differential equation yields the unique solution:

$$F(\theta) = 1 + k\theta^{-\frac{1}{1-\lambda}}.$$

The only family of CDFs satisfying this equation is the Pareto family.

Next, define $f(x) = \frac{v(x)}{x}$. f is decreasing if and only if $v'(x)x - v(x) < 0$, which is guaranteed by strict concavity of v, given $v \geq 0$. Similarly, f is increasing if v is strictly convex.

Assume first v is strictly concave. Consider a realization θ such that $\frac{\theta^1}{\theta^2} = \frac{\alpha}{\beta-\alpha}$ and $\theta^1 > \theta^2$. By the argument above, $f(\theta^1) < f(\theta^2)$ which implies:

$$\frac{\theta^1}{\theta^2} = \frac{\alpha}{\beta - \alpha} > \frac{v(\theta^1)}{v(\theta^2)}.$$

Moreover, for small enough ε we have, by continuity of v:

$$\frac{\theta^1 + \varepsilon}{\theta^2} > \frac{\alpha}{\beta - \alpha} > \frac{v(\theta^1 + \varepsilon)}{v(\theta^2)}.$$

By inequalities 4 and 5, this implies that the optimal mechanism provides the good for both agents, whereas the first best allocation only provides it for agent 1. The argument for under-provision is symmetric. ■

Proof of Proposition 2

Start with any mechanism that implements the optimal allocation and charges $t_\alpha(\theta_{-i})$ in case the agent shares, and $t_\beta(\theta_{-i})$ if the agent does not share. It is clear that, conditional on an allocation, bids cannot depend on own-type under Dominant-Strategy implementation. In what follows I omit the argument of t_α, t_β whenever possible.

Case 1 $\theta_2 \leq g^{-1}(\theta_2) < \theta_1$. Agent 1 is allocated the good alone. There is clearly no benefit in deviating to a higher bid, as that does not change either the allocation or the payment. So consider a deviation to a lower bid that makes the seller allocate the goods to both. Then, it must be the case that:

$$\theta_1\beta - t_\beta \geq \theta_1\alpha - t_\alpha \iff t_\beta - t_\alpha \leq \theta_1(\beta - \alpha).$$

Because this has to hold for all θ_1 in this set, we have the first constraint:

$$t_\beta - t_\alpha \leq g^{-1}(\theta_2)(\beta - \alpha). \tag{6}$$

A similar argument holds for deviations that exclude agent 1. Because under exclusion there are no payments, we have:

$$t_\beta \leq g^{-1}(\theta_2)\beta. \tag{7}$$

Case 2 $g(\theta_2) \leq \theta_1 < g^{-1}(\theta_2)$. In this case, both agents get the product. The deviation to higher types is avoided if:

$$\theta_1\alpha - t_\alpha \geq \theta_1\beta - t_\beta \iff t_\beta - t_\alpha \geq \theta_1(\beta - \alpha).$$

For this second constraint to hold for any θ_1 in this set, we have: $t_\beta - t_\alpha \geq g^{-1}(\theta_2)(\beta - \alpha)$. Combining this equality with 6, we obtain an expression for the difference in payments:

$$t_\beta - t_\alpha = g^{-1}(\theta_2)(\beta - \alpha). \tag{8}$$

Conversely, the downward deviation is avoided if $t_\alpha \leq \theta_1\alpha$, which is satisfied for all θ_1 in this set only if $t_\alpha \leq g(\theta_2)\alpha$.

Case 3 $\theta_1 < g(\theta_2)$. We are finally in the case in which 1 is excluded. For this to be optimal we have $t_\beta \geq \theta_1\beta$ and $t_\alpha \geq \theta_1\alpha$, which imply, respectively, $t_\beta \geq g(\theta_2)\beta$, and $t_\alpha \geq g(\theta_2)\alpha$. The last inequality pins down t_α given the discussion in the previous paragraph. We have then proved that t_α and $t_\beta - t_\alpha$ are pinned down by dominant-strategy ICs, and satisfy the mechanism in the statement. Therefore, this mechanism not only implements the optimal allocation in dominant strategies, but it also is the only one to do so conditional on the excluded agent not paying anything.

Proof of Lemma 7

By switching the order of θ and θ' in the BIC inequality above and putting the two together we obtain:

$$\mathbb{E}\left[(\gamma(\theta,\theta_{-i})-\gamma(\theta',\theta_{-i}))\cdot q_i(\theta',\theta_{-i})\right] \leq U_i(\theta)-U_i(\theta') \leq \mathbb{E}\left[(\gamma(\theta,\theta_{-i})-\gamma(\theta',\theta_{-i}))\cdot q_i(\theta,\theta_{-i})\right]$$

Divide all three terms by $\theta-\theta'$ and take the limit as $\theta' \to \theta$ to obtain condition (1). By combining the first and second inequality, we obtain condition (2):

$$\mathbb{E}\left[(\gamma(\theta,\theta_{-i})-\gamma(\theta',\theta_{-i}))\cdot(q_i(\theta,\theta_{-i})-q_i(\theta',\theta_{-i}))\right] \geq 0.$$

■

Proof of Proposition 4

Start by writing $\overline{\gamma} = (\beta-\alpha,\alpha)$, and $\overline{q}_i = q_i, q_i + q_\alpha$. Assume first $\theta > \theta'$. Then:

$$\begin{aligned} U(\theta)-U(\theta') &= \int_{\theta'}^{\theta} \mathbb{E}\left[\gamma'(\nu,\theta_{-i})\cdot q_i(\nu,\theta_{-i})\right] d\nu \\ = \mathbb{E}\left[\int_{\theta'}^{\theta} \gamma'(\nu,\theta_{-i})\cdot q_i(\nu,\theta_{-i}) d\nu\right] &= \mathbb{E}\left[\int_{\theta'}^{\theta} \overline{\gamma}'(\nu,\theta_{-i})\cdot \overline{q}_i(\nu,\theta_{-i}) d\nu\right] \\ &\geq \mathbb{E}\left[\int_{\theta'}^{\theta} \overline{\gamma}'(\nu,\theta_{-i}) d\nu \cdot \overline{q}_i(\theta',\theta_{-i})\right] \\ &= \mathbb{E}\left[(\gamma(\theta,\theta_{-i})-\gamma(\theta',\theta_{-i}))\cdot q_i(\theta',\theta_{-i})\right], \end{aligned}$$

where the first equality comes from condition 1, the second equality switches the order of integration, the third equality rewrites the integrand using the definitions of $\overline{\gamma}$ and $\overline{q}$, and the inequality uses the fact that, by Assumption 4, both entris of $\overline{\gamma}'$ are positive and, by the statement of the result, both entries of $\overline{q}_i$ are increasing.

The symmetric argument holds for $\theta' > \theta$, so we proved that BIC is satisfied. ■

Proof of Proposition 5

We solve the relaxed problem of maximizing profits subject to condition 1 in Lemma 7 and feasibility, (F). By usual arguments, the solution to that relaxed problem is the one above. We next show that the solution above satisfies incentive compatibility.

We start by proving q_i is increasing in θ for any θ_{-i}. Fix θ. By the the third inequality on virtual gains there is a threshold in the opponent's type space, call it x, such that the allocation rule $q_i(\theta, \theta_{-i})$ is one if and only if $\theta_{-i} < x$. That threshold satisfies:

$$v_\beta(\theta, x) = v_\alpha(\theta, x) + v_\alpha(x, \theta)$$

By total differentiation, we obtain:

$$\underbrace{\Big(v'_\beta(\theta, x) - (v'_\alpha(\theta, x) + v_{\alpha,\nu}(x, \theta))\Big)}_{>0 \text{ by virtual gains,inequality 1}} d\theta = -\underbrace{\Big(v_{\beta,\nu}(\theta, x) - (v_{\alpha,\nu}(\theta, x) + v'_\alpha(x, \theta))\Big)}_{<0 \text{ by virtual gains, inequality 2}} dx.$$

Thus, the threshold x is increasing with θ. Then, if $q_i(\theta, \theta_{-i}) = 1$, and $\theta' > \theta$, it must be that $q_i(\theta', \theta_{-i}) = 1$. Thus, q_i is increasing, as we wanted to prove.

We now show that $q_\alpha + q_i$ is increasing. Again fix any θ. Once more, using virtual gains it is easy to see that there is a threshold in the adversaries' type space, $y > \theta > x$ such that $q_\alpha + q_i = 1$ if and only if $\theta_{-i} < y$. y is defined by:

$$v_\beta(y, \theta) = v_\alpha(\theta, y) + v_\alpha(y, \theta).$$

Using total differentiation again:

$$\underbrace{\Big(v'_\beta(y, \theta) - v_{\alpha,\nu}(\theta, y) - v'_\alpha(y, \theta)\Big)}_{>0 \text{ by virtual gains, inequality 1}} dy = -\underbrace{\Big(v_{\beta,\nu}(y, \theta) - v'_\alpha(\theta, y) - v_{\alpha,\nu}(y, \theta)\Big)}_{<0 \text{ by virtual gains, inequality 2}} d\theta$$

Again, the threshold y grows. So if $q_i(\theta, \theta_{-i}) + q_\alpha(\theta, \theta_{-i}) = 1$, the same holds for $\theta' > \theta$, which guarantees that this sum is increasing.

We have now proved q_i and $q_i + q_\alpha$ are increasing, and we are thus in the conditions of Proposition 2.

■

References

Bergemann, D., Brooks, B., and Morris, S. (2020). Countering the winner's curse: Optimal auction design in a common value model. *Theoretical Economics*, 15(4):1399–1434.

Bulow, J. I. and Klemperer, P. D. (1994). Auctions vs. negotiations.

Jehiel, P. and Moldovanu, B. (2000). Auctions with downstream interaction among buyers. *Rand journal of economics*, pages 768–791.

Jehiel, P., Moldovanu, B., and Stacchetti, E. (1996). How (not) to sell nuclear weapons. *The American Economic Review*, pages 814–829.

Jehiel, P., Moldovanu, B., and Stacchetti, E. (1999). Multidimensional mechanism design for auctions with externalities. *Journal of economic theory*, 85(2):258–293.

Kamien, M. I., Oren, S. S., and Tauman, Y. (1992). Optimal licensing of cost-reducing innovation. *Journal of Mathematical Economics*, 21(5):483–508.

Katz, M. L. and Shapiro, C. (1986). Technology adoption in the presence of network externalities. *Journal of political economy*, 94(4):822–841.

Milgrom, P. R. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research*, 6(1):58–73.

Sen, D. (2005). On the coexistence of different licensing schemes. *International Review of Economics & Finance*, 14(4):393–413.

Sen, D. and Tauman, Y. (2007). General licensing schemes for a cost-reducing innovation. *Games and Economic Behavior*, 59(1):163–186.