Selling to Competitors *

Pëllumb Reshidi † João Thereze ‡

[Click Here for the Latest Version]

September 25, 2024

Abstract

A manufacturer seeks to license a product to downstream competitors with unknown productivities. She can design a mechanism to allocate licenses to one or multiple competitors. We identify the revenue-maximizing mechanism and show it can be implemented through an *interval auction*: the highest bidder is exclusively licensed if their bid is much higher than others, but multiple bidders are licensed otherwise. This mechanism does not allocate efficiently, and we characterize the distributions of buyer valuations that lead to over- or under-licensing. If buyers arrive over time, the seller may delay licensing, and we show that the seller only commits to exclusive contracts if she is less patient than the buyers.

*We thank Leeat Yariv, Roland Benabou, Alessandro Lizzeri, Curt Taylor, Giuseppe Lopomo, James Anton, and Udayan Vaidya for very helpful suggestions and feedback. We also thank audiences at Princeton University, Duke University, Fuqua School of Business, and the 35th Stony Brook International Conference on Game Theory for their helpful comments and discussions.

†Department of Economics, Duke University, Durham, NC 27710—pellumb.reshidi@duke.edu

‡Fuqua School of Business, Duke University, Durham, NC 27708—joao.thereze@duke.edu

1 Introduction

Consider BioNTech developing the next-generation mRNA vaccine and deciding whether to license this technology to two rival pharmaceutical companies: Pfizer and Johnson & Johnson (J&J). Each company privately understands the benefits the technology would bring to their vaccine development efforts, especially if they are the sole licensee. While BioNTech could license the technology to both Pfizer and J&J, doing so would reduce each company's competitive edge due to increased rivalry in the vaccine market. If BioNTech's objective is to maximize profits from this licensing, should it commit to exclusivity by auctioning the license to the highest bidder? Should it set a fixed price and offer the license to both companies? Or is there a better approach? This paper examines the revenue-maximizing mechanisms in such scenarios.

In many markets—such as information sales, franchise licensing, and government procurement—sellers face similar trade-offs. Should OpenAI allow only Apple to integrate ChatGPT, only Android, or both? Should Bloomberg provide proprietary market trend data exclusively to a top investment bank, to a hedge fund, or to anyone willing to pay? Should Intel sell its processors exclusively to Dell or offer them to multiple manufacturers? All these cases share three critical features: the seller is able to replicate the good, there are externalities between the buyers, and the buyers hold private information about their profits. Despite the ubiquity of these situations, the seller's optimal strategy remains largely unknown. We aim to bridge this gap.

We begin our analysis by modeling competitive profits, relevant when licenses are issued to more than one buyer, as a proportion of monopoly profits. In this setup, each buyer's profit depends on their private type. This framework effectively captures various market scenarios while maintaining analytical tractability. Moreover, it underscores a key assumption in this paper: while buyers have private information about their own valuations, the market structure, which captures the externalities among buyers, is common knowledge. As a result, the seller knows how to map buyers' valuations to outcomes when allocating the good to multiple buyers.

Our first result is identifying the optimal mechanism that maximizes the seller's profit. While the insights apply to any number of buyers, for simplicity, we explain them below using the case of two buyers. This mechanism allocates the good to a single buyer when their private valuation significantly exceeds that of the other and allocates it to both buyers otherwise. To implement the mechanism in dominant strategies, we introduce what we term an *interval auction*. In an interval auction, each buyer submits a bid, around which specific neighborhoods are defined. If a bid is below the neighborhood of the competitor's bid, that buyer is excluded and incurs no cost. If a bid falls within the neighborhood of the competitor's bid, both buyers are awarded the good and pay the lowest price consistent with them being in their competitor's neighborhood. Finally, if a bid exceeds the neighborhood of the competitor's bid, that bidder alone is awarded the good, paying a premium for it. Thus, in this auction, it is not only the highest bid that matters; the entire distribution of bids influences outcomes. If the bids cluster closely, multiple licenses are awarded; if they are widely dispersed, only the highest bidder receives a license. Despite the inherent complexities, our findings reveal that implementation is relatively simple, making it a viable option for real-world application.

Next, we characterize inefficiencies absent in traditional auctions. In our setup, the seller may either under or over-provide the good—selling to fewer or more buyers than would be optimal under symmetric information. In standard auctions with symmetric buyers, inefficiencies emerge only when virtual valuations are non-increasing or negative. As long as virtual valuations are monotonic, the auctioneer's most valuable bidder remains unchanged regardless of whether buyers have private information or not. However, when the auctioneer can sell to multiple buyers, the optimal allocation is governed by the ratio of valuations under symmetric information and by the ratio of virtual valuations under private information. It is this discrepancy between the two ratios that drives inefficiencies. We find that these inefficiencies are ubiquitous: the optimal mechanism is efficient if and only if the distribution of buyers' types belongs to the Pareto family. Importantly, we establish a link between the shape of the distribution of buyers' types and the nature of the inefficiency—whether the good is under or over-provided. Put simply, our result reveals that a policymaker can assess whether a good will be over or under-supplied in a market based on an understanding of the distribution of valuations without needing to know about market conduct or the magnitude of externalities.

We extend our analysis by adapting the baseline model to a dynamic framework where bidders arrive sequentially over time. In this setup, our key assumption is that once a seller grants a license to a buyer, they cannot revoke it later. Under the optimal dynamic mechanism, when the first buyer arrives, the seller either grants them a license if their type is high enough or asks them to wait for another buyer. When the seller faces multiple buyers, the profit-maximizing allocation aligns with that of the static model. It turns out that it is never optimal for the seller to promise future exclusivity to an early-arriving buyer. We find that in the dynamic environment, asymmetric information creates an additional inefficiency. It affects not only the number of licenses issued, as in the static case but also influences which initial buyers are or are not issued a license. Buyers who would receive licenses immediately under complete information might be asked to wait, or vice versa. Concretely, we observe that sellers tend to over-wait when they under-provide licenses and under-wait when they over-provide. We once again establish a direct link between these inefficiencies and the shape of the distribution of buyer valuations, which turns out to be the same condition as in the static model. A policymaker can narrow their concerns simply by understanding this distribution.

We then consider a scenario where the seller is less patient than the buyer, under the assumption that payments are made upfront. In this setup, the seller offers the initial buyer a contract that specifies future allocations for each period in which another buyer may arrive, with payments made upfront. Allocation cutoffs are proportional to those in the static model but become progressively more favorable to the initial buyer over time. There is a finite period in the future beyond which—if the other buyer has not yet arrived—the seller guarantees exclusivity to the current buyer. However, even in this case, the exclusive contract is offered only at a point sufficiently far in the future. In other words, we find that it is challenging to justify contracts that offer immediate exclusivity to buyers.

Finally, we extend the static model to accommodate more general profit functions with supermodular returns from exclusivity. We then extend the framework to account for interdependencies, where a buyer's profits, when winning with others, depend not only on their own type but also on the types of their competitors. This adjustment introduces correlations among buyers' valuations, shifting our analysis from a framework of independent private values to one that involves the complexities of common value auctions. We identify sufficient conditions on preferences and distributions, that ensure the optimality of our mechanism.

The remainder of the paper is organized as follows: In Section 2, we begin our analysis with the baseline model. In Section 3, building on insights from the previous section, we examine a dynamic version of the model. In Section 4, we extend the model to account for more general profit functions as well as interdependent valuations. We consider applications in Section 5, and conclude in Section 6.

1.1 Related Literature

Patent Licensing Our paper relates to a body of work on patent licensing in oligopolistic downstream industries (Kamien and Tauman, 1986; Katz and Shapiro, 1986; Kamien et al., 1992; Sen and Tauman, 2007; Li and Wang, 2010; Doganoglu and Inceoglu, 2014).[1] These papers conduct their analysis under 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. The role of informational asymmetry is taken seriously in later works such as Choi (2001), Poddar et al. (2002), and Sen (2005), which allow for asymmetric information but consider only a monopolistic buyer.[2] On the other hand, allowing for multiple buyers, Antelo and Sampayo (2017) studies a signaling problem, while Antelo and Sampayo (2024) studies a screening problem where the types of buyers can be high or low.[3] Both the earlier and more recent studies focus on identifying optimal licensing strategies within a range of mechanisms, such as determining the optimal fees, setting the optimal reservation price in a first-price auction, and establishing the optimal royalties. In contrast, our work identifies the optimal mechanism from the entire set of feasible options.

Within the licensing literature, of relevance for our initial static setup is the work of Schmitz (2002), who considers selling a license to two potential buyers. They characterize the profit-maximizing mechanism and identify potential inefficiencies that may arise from information asymmetries. Differently from this paper, we characterize precisely when such inefficiencies arise, allow for more than two buyers, study a dynamic version, and allow for general profit functions, including cases with interdependent types.

Mechanism Design with Externalities Our paper is related to the literature on mechanism design with externalities, particularly relating to the works of Jehiel et al. (1996) and Jehiel et al. (1999), which explore multidimensional settings with unknown market structures. In contrast, we model the market structure as a function of all buyers' types, which in turn reduces dimensionality and enhances tractability, allowing for a comprehensive characterization of the optimal mechanism. Additionally, our approach permits multiple sales of the good and considers externalities based on the

[1] For an early survey, see Kamien (1992).

[2] There is also a literature that incorporates asymmetric information where the quality of the innovation/license is not fully known to the buyers (Zhang et al., 2016; Jeon, 2019; Wu et al., 2021).

[3] Differently from this body of work, Heywood et al. (2014) and Fan et al. (2018) consider a setup in which the seller is an active competitor in the market.

opponent's realized type, not just their identity.

Of relevance is also Jehiel and Moldovanu (2000), who study auctions with downstream interactions among buyers. Like our work, they model outcomes as a function of buyers' types, but unlike us, they focus on the sale of a single unit and focus their analysis on second-price, sealed-bid auctions.

Auctions with Common Values Our work also relates to the literature on auctions with common values, including classic studies by Milgrom and Weber (1982) and Bulow and Klemperer (1996), as well as more recent approaches that identify the optimal mechanism under specific setups, such as Bergemann et al. (2020).[4] Our work differs from this existing body of literature in two important ways. First, we allow for the sale of multiple goods. Second, in our setup, buyers' types are negatively correlated. As it turns out, this negative correlation aids in characterizing the optimal mechanism.

Multi Unit Auctions Our setup shares similarities with the literature on multi-unit auctions and bundling. In particular, the decision to offer two licenses to two different buyers, which reduces their individual payoff, rather than a license to one buyer is akin to the decision of selling goods to two or one buyer (Armstrong, 2000; Avery and Hendershott, 2000). We diverge from that setup in several ways. First, by assuming the market structure is known, and focusing on buyers' productivities as their types, we reduce the dimensionality of the type space—each bidder is no longer associated with different marginal values for each additional item. We also extend beyond the standard multi-unit auction approach by incorporating dynamics and by allowing for interdependent valuations of the goods.

2 The Setup

An auctioneer has an item to sell to N potential buyers, indexed in $\mathcal{N} = \{1, ..., N\}$. This item differs from standard commodities in two key ways. First, it generates externalities: buyers' valuations of the product depend on how many other buyers purchase it. Second, the item can be replicated at no cost—allowing the seller to sell to multiple

[4]This work differs from studies where the correlation lies on bidders' signals rather than directly in their valuations. Such a scenario was explored even by Myerson (1981), who illustrated that if bidders' private information is correlated, the seller can design a mechanism to extract the full surplus. Crémer and McLean (1985) demonstrated that Myerson's example has broad applicability, and subsequent research, including Crémer and McLean (1988), McAfee et al. (1989), and McAfee and Reny (1992), further established that this result holds under even more general conditions.

buyers. 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 $|k|$. 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 normalize the utilities of agents who do not purchase an item 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 vector of market coefficients, 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.[5] Each θ_i is assumed to be independently drawn from a regular distribution with cumulative distribution F that has full support. We make further assumptions about α values. In particular, we impose symmetry: if $i \in \mathcal{J}_k$, then $\alpha_k^i = \alpha_{|k|}$. Additionally, we assume $\alpha_k^i \in [0,1]$ for all i and k. Utilities are 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 $\mathcal{J}$ denote the set of all such subsets. Given this setup, the revelation principle applies, allowing us to focus on identifying the truthful direct revelation mechanism that maximizes revenue.

2.1 First Best Allocation

We start by establishing the revenue-maximizing allocation under symmetric information. If the principal knows the vector $\theta = (\theta^1, ..., \theta^N)$, she chooses transfers i and an allocation σ_k to solve:

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

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

It is clear that (IR) must hold with equality in any solution. Thus, the problem can be simplified to a simple accounting problem: the seller considers the maximal gross

[5]As we show later, the principal does not need to have this market knowledge. As long as the buyers know the market structure, the optimal mechanism can be implemented.

payoff that buyers can obtain across all possible groups and extracts all revenues. As usual, the revenue-maximizing allocation under symmetric information is also welfare-maximizing, so we call it the first-best allocation.

The principal's problem can be easily understood in the case where $N = 2$. Here, the payoff for being allocated the good alone is normalized to θ^i, and we define $\alpha\theta^i$ as the payoff when the good is allocated to both buyers. In this case, it is optimal to sell to i alone if $\frac{\theta^i}{\theta^j} \geq \frac{\alpha}{1-\alpha}$. By contrast, it is optimal to sell to both i and j if $\frac{\alpha}{1-\alpha} \geq \frac{\theta_i}{\theta_j} \geq \frac{1-\alpha}{\alpha}$. Importantly, the optimal allocation is driven by the ratio of valuations θ_i/θ_j.

2.2 Revenue-Maximization under Asymmetric Information

Next, we characterize the revenue-maximizing mechanism when the seller 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 feasible allocations, we can define the associated feasible assignment set as

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

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

Lemma 1. $\mathcal{Q} = \operatorname{co}\left\{\alpha_k : k \in \{1, \dots, 2^N\}\right\}$. *$\mathcal{Q}$ 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

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

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

Lemma 2. *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. While in standard auctions feasibility requires allocations to be located in the unit simplex, our feasibility condition imposes that trade probabilities are contained within the polytope $\mathcal{Q}$—which,in general, extends 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

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

The next proposition characterizes the allocation in the optimal mechanism when $N = 2$.

Proposition 1. *Let $N = 2$. In the optimal mechanism, allocations satisfy:*

$$(q^1, q^2)(\theta) = \begin{cases} (1, 0), & \text{if } \frac{v(\theta_1)}{v(\theta_2)} \geq \frac{\alpha}{1-\alpha} \\ (\alpha, \alpha), & \text{if } \frac{\alpha}{1-\alpha} \geq \frac{v(\theta_1)}{v(\theta_2)} \geq \frac{1-\alpha}{\alpha} \\ (0, 1), & \text{otherwise} \end{cases} \tag{1}$$

2.3 Inefficiencies

In traditional auction theory, asymmetric information can lead to inefficiencies in two primary ways. First, having non-increasing virtual valuations or heterogeneous agents could result in a scenario where an agent with a lower valuation wins the auction, causing an ex-post inefficient allocation. The second type of inefficiency arises if virtual values can be negative. If the realized virtual values are negative across all agents, the good remains unsold even if all agents value it more than the seller. In our setup, inefficiencies not found in traditional auctions emerge. To distinguish these from conventional suboptimal outcomes, we assume that all agents draw their types from the same regular distribution F and that virtual valuations are increasing. This eliminates the first inefficiency. Furthermore, by ensuring that virtual values are positive for all realizations, we eliminate the second inefficiency.

Assumption 1. *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:*

$$q^1(\theta) + q^2(\theta) \leq (\geq) q_f^1(\theta) + q_f^2(\theta) \text{ for all } \theta.$$

An allocation is efficient if equality holds above.

Let $h(\theta_i)$ represent the inverse hazard rate. Define $\lambda(\theta_i) \equiv h(\theta_i)\theta_i$.[6]

Proposition 2. *The profit-maximizing mechanism*

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

The above proposition implies that the profit-maximizing mechanism will prescribe the same allocation as the first-best outcome for any α and any realized type values if and only if the buyer's types follow a distribution within the Pareto family. While we know from previous work, such as Jehiel et al. (1996) and Schmitz (2002), that information asymmetries can lead a profit-maximizing monopolist to over-provide a good, our paper is, to the best of our knowledge, the first to specifically characterize

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

when such inefficiencies occur, based on the distribution of buyers' types. 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 1. The figure illustrates the profit-maximizing

Figure 1: Examples of Under and Overprovision

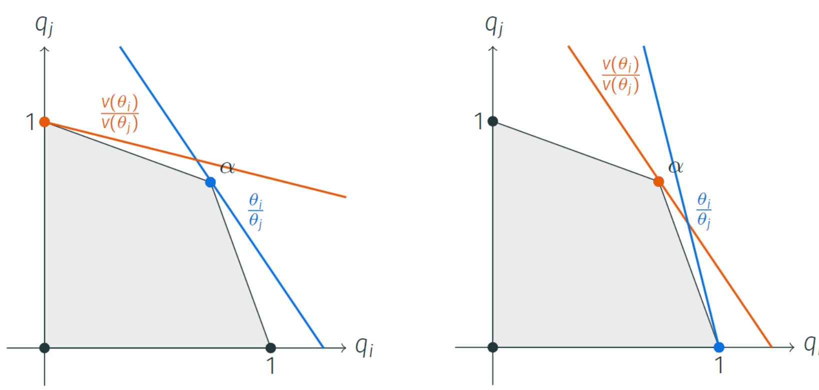

Notes: The figure above displays the profit-maximizing and the first-best allocations for different realized values of θ_i and θ_j. In the left panel, it is efficient to allocate the good to both agents, but it is profit-maximizing to allocate the good to agent i exclusively—*underprovision*. In the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

and the first-best allocations for different realized values of θ_i and θ_j. In the left panel, 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 right panel, 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 2 presents two such examples.

Figure 2: Examples of Distributions leading to Under and Overprovision

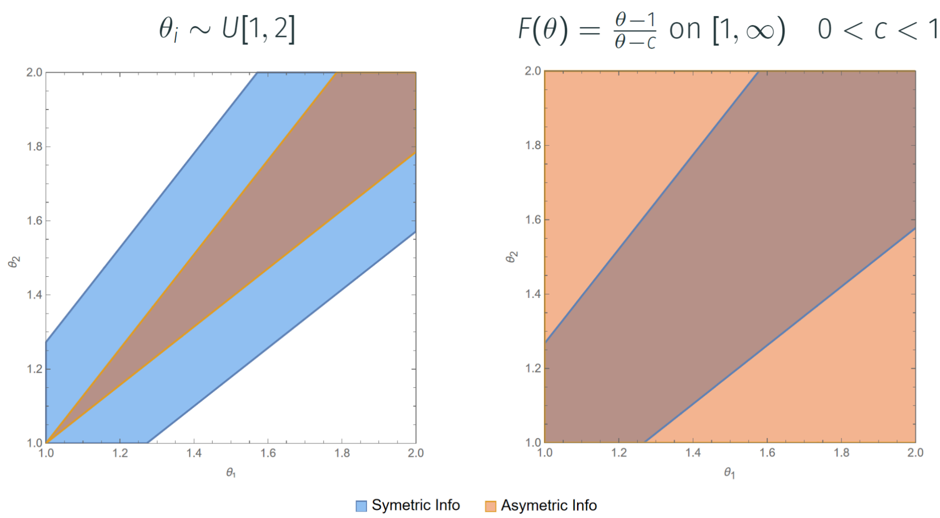

Notes: The figure above illustrates both the profit-maximizing and first-best allocations. The distributions used in each example are shown at the top of the graphs, with $\alpha = 0.56$.

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. On the left panel, the shaded orange region is contained 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 contained 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. Thus, there exists a nonempty set of distributions for which either outcome is possible.

We reiterate that these inefficiencies, whether they involve under- or over-provision of the good, are absent in standard auctions. To underscore that typical inefficiencies are not the drivers of these results, we have assumed that all agents draw their types from the same distribution F and that virtual values are positive and increasing. Under these assumptions, standard auctions do not exhibit inefficiencies. Yet, in this setup, over- or under-provision can occur.

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.

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

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

Proposition 3. *The optimal mechanism is implemented in dominant strategies by an interval 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: both agents are asked to submit bids. Assume, without loss, that $b_1 \geq b_2$. If $b_1 < \overline{\tau}(b_2)$, then allocate the good to both agents, who pay $\alpha\underline{\tau}(b_{-i})$ each. If $b_1 \geq \overline{\tau}(b_2)$, then allocate the good to the first bidder only. This bidder pays $\alpha\underline{\tau}(b_{-i}) + (1-\alpha)\overline{\tau}(b_{-i})$. We visualize the workings of this mechanism in Figure 3 below.

Figure 3: Interval Auction Implementation

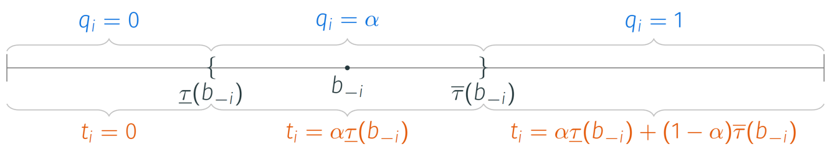

Notes: The figure above visualizes the profit-maximizing implementation via an interval auction. Around the bid of the opponent b_{-i} there is a n neighborhood $(\underline{\tau}(b_{-i}), \overline{\tau}(b_{-i}))$. If the agent's bid falls below this neighborhood, he is excluded and pays nothing $t_i = 0$. 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 $t_i = \alpha\underline{\tau}(b_{-i}) + (1-\alpha)\overline{\tau}(b_{-i})$.

In this implementation, for an agent to secure exclusive rights to the good, they must significantly outbid the other agent. Merely outbidding the other agent by a small margin results in both agents being allocated the good. Conversely, if an agent loses by only a small margin, both agents still receive the good. The agent is excluded only when their bid is substantially lower than their opponent's. 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.

2.5 Revenue Comparison

The revenue difference can be compared among various mechanisms, such as a posted price, a standard auction where the good is sold to a single buyer, and the optimal mechanism identified in this paper. This comparison can be formalized as follows. First, because virtual valuations are assumed to be positive, it can be shown that the optimal posted price involves setting the price at $\alpha\underline{\theta}$, where both agents purchase the product. Through standard manipulations of the virtual value function, this revenue can be expressed as

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

On the other hand, the revenue from a standard auction, where the designer commits to selling only one product, is determined by the expected value of the second-highest bid, which can be expressed as

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

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

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

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

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

Figure 4 illustrates this comparison. For any $\alpha \in (0.5, 1)$, interval auctions outperform either mechanism. Notably, as α approaches 0.5, the likelihood of selling to a single agent increases—the polytope discussed earlier converges to the unit simplex—causing the profits from interval auctions to converge to those of a regular auction. Conversely, as α approaches 1, the externalities from having two active firms diminish, leading profits to align with those from a posted price.

Therefore, when externalities are so severe that the seller would not consider selling to more than one buyer, using a traditional auction results in minimal or no revenue loss. On the opposite end, if there are no externalities, the seller maximizes profits by selling to all buyers, with the only optimization being the determination of the optimal price. Consequently, it is in the intermediate cases—where externalities are significant but not overwhelming—that our mechanism offers the greatest advantages.

3 Dynamic Model

We now consider a version of the model in which buyers arrive sequentially, so the seller also decides the timing of license concesssions. Time is discrete and runs indefinitely: $t \in \mathbb{N}$. At any time t, with probability λ, a buyer $i \in \{1, 2\}$ may arrive. Arrival times are independent between buyers. Buyers discount the future at rate δ, while the seller discounts the future at rate ρ, with $\rho \leq \delta$. We let $a^i \in \mathbb{N}$ denote the arrival time of buyer

Figure 4: Revenue Comparison

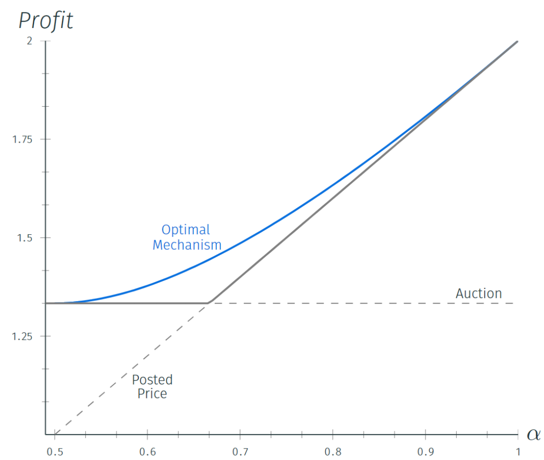

Notes: For different α values, the graph above compares the revenue from a posted price, a standard auction in which the good is sold to one buyer, and the optimal mechanism.

i—if the buyer has not arrived, denote $a^i = o$. A direct mechanism consists of functions $[q_t^i, r_t^i]_{i=\{1,2\},t\in\mathbb{N}}$: an allocation $q_t^i : \Theta^2 \times \mathbb{N}^2 \to [0,1]$ and a transfer $r_t^i : \Theta^2 \times \mathbb{N}^2 \to \mathbb{R}$ which specify, for every buyer i, time t, types $\theta = (\theta_i, \theta_{-i})$, and arrival times a^i, a^{-i}, a number between 0 and 1, and a value in the reals. Let $c_t^i = (q_t^i, r_t^i)$. We impose the following restrictions on mechanisms:

Definition 3. *A mechanism is permissible if it satisfies*

1. ***Feasibility**: for each t, $(q_t^1, q_t^2) \in \mathcal{Q}$;*

2. ***Consistency**: for $a^i > t$, $c_t^i = 0$ and $c_t^{-i}(\theta_i, \theta_{-i}, a^i, a^{-i}) = c_t^{-i}(\theta_i', \theta_{-i}, a', a^{-i})$ for all $\theta_{-i} \in \Theta$, $a' > t$;*

3. ***Irreversibility**: Let $t' > t \geq a^i$. Then, if $a^j > t'$, $q_{t'}^i \geq q_t^i$. If $a^j \leq t'$, then $q_{t'}^i \geq \alpha q_t^i$.*

The first condition is the same as in the static model and ensures that the assignments of the product to the agents are consistently represented in the allocation. The second condition restricts what can be offered when one agent has not yet arrived. In particular, it must be that if a buyer has not yet arrived, they cannot be allocated the good or be asked for any transfers. On the other hand, the allocation and transfer of the buyer who has arrived cannot depend on the type of buyers who have not yet arrived. The most significant restriction out of the three is irreversibility. Irreversibility implies that once the designer allocates a license to an agent, she cannot take it back.

This implies that the probability of being assigned a license cannot decrease over time. Thus, the only way a buyer's allocation can be reduced is if another buyer arrives and is also allocated a license with some probability.

We focus on cases in which agents arrive sequentially. We soon clarify that when buyers arrive simultaneously, the optimal mechanism is the one identified in the static model. Without loss of generality, say that agent 1 is the first to arrive. Because the problem of the principal effectively starts at that time—due to consistency—we normalize $a^i = 0$. We also assume all transfers r, from i happen at the time of arrival of agent i, which is without loss of optimality given the assumptions on discount rates. The payoff of an agent 1, who arrives first and at 0, is:

$$U^1(\theta_1) = \mathbb{E}_{\theta_2}\left[\sum_{j=0}^{\infty} \delta^j (1-\lambda)^j q_j^1(\theta_1) + \sum_{j=0}^{\infty} \lambda(1-\lambda)^j \sum_{k=j+1}^{\infty} \delta^k q_k^1(\theta_1, \theta_2, a^2 = j+1)\right]\theta_1 - r^1(\theta_1),$$

The first term in the parentheses takes into account the times t such that $a^2 > t$, that is buyer 2 has not yet arrived. In this case, we know that q^1 does not depend on θ^2 or a^2, by consistency, so we omit those variables. The second term takes into account the cases when buyer 2 arrives at time $j+1$.

When the second agent arrives, their utility at time a^2 is:

$$U^2(\theta_1, \theta_2, a^2) = \left[\sum_{j=0}^{\infty} \delta^j q^2_{a^2+j}(\theta_1, \theta_2, a^2)\right]\theta_2 - r^2(\theta_1, \theta_2),$$

The seller maximizes expected revenue, discounted by ρ, in the set of mechanisms that are available, incentive compatible and individually rational at the time of arrival.

Proposition 5. *Normalize the arrival time of buyer 1 to $t = 0$ and let a be the arrival time of buyer 2. There exists some $\hat{\theta} < \overline{\theta}$ such that*

For all $t < a$,

$$q_t^1(\theta_1) = \begin{cases} 1, & \text{if } \theta_1 \geq \hat{\theta} \\ 0, & \text{otherwise,} \end{cases} \tag{2}$$

Figure 5: Optimal Dynamic Mechanism

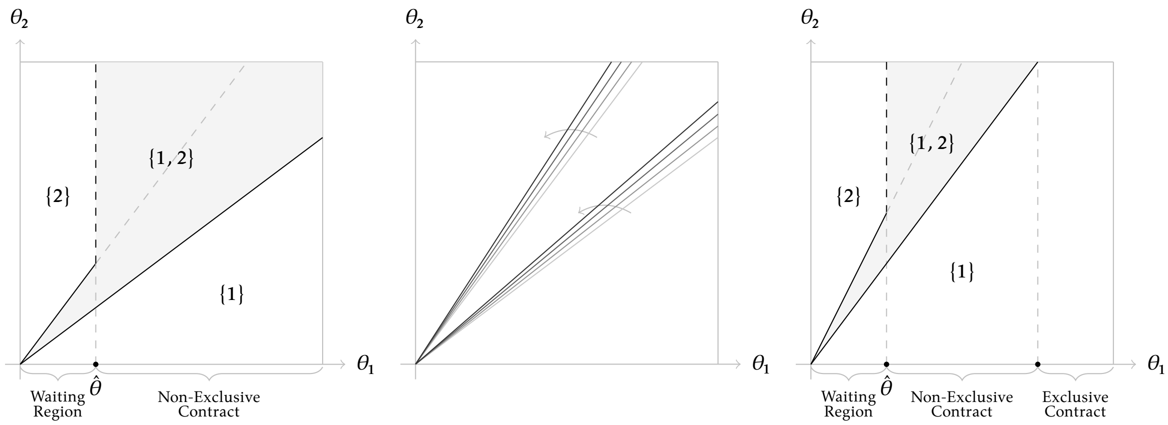

Notes: The figures above depict the revenue-maximizing dynamic mechanism. The left panel shows the $\hat{\theta}$ threshold: agents with types below this threshold are asked to wait, while those above it are immediately issued a contract. In the latter scenario, depending on the type of the subsequent agent, θ_2, the principal may choose to allocate to the first agent exclusively or to both agents but no longer exclusively to the second. The middle panel illustrates how these cutoffs evolve for later arrival dates of the second agent, becoming more favorable towards the first agent. The right panel presents the allocation regions for sufficiently delayed arrival dates of the second agent. As can be seen, an exclusive contract region emerges.

For all $t \geq a$

$$(q_t^1, q_t^2)(\theta_1, \theta_2, a) = \begin{cases} (1, 0), & \textit{if } \left(\frac{\delta}{\rho}\right)^a \frac{v(\theta_1)}{v(\theta_2)} \geq \frac{\alpha}{1-\alpha} \\ (\alpha, \alpha), & \textit{if } \frac{\alpha}{1-\alpha} \geq \left(\frac{\delta}{\rho}\right)^a \frac{v(\theta_1)}{v(\theta_2)} \geq \frac{1-\alpha}{\alpha} \\ \left(\alpha q_0^1(\theta_1), 1 - (1-\alpha) q_0^1(\theta_1)\right), & \textit{otherwise} \end{cases} \tag{3}$$

Figure 5 offers an illustration of the optimal mechanism. Based on the proposition above, the optimal dynamic mechanism can be conceptualized as a two-step process. In the first step, the decision is whether to issue a license to the buyer who has already arrived. The advantage of issuing a license immediately is that the buyer can be charged a higher price, as they will begin operating and generating profits before the second buyer arrives. However, the cost of issuing the license to the current buyer is the lost option value of having only the second buyer active when they arrive—due to irreversibility, this option is no longer available. Once the second buyer arrives, even if their type is significantly higher than that of the first buyer, it is no longer possible to revoke the first buyer's license, meaning that the principal can, at best, make the buyers share the market by issuing two licenses.

The second step of the mechanism is contingent on the decision made in the first step. If the decision was to wait, the principal then compares the types of buyers and decides whether to allocate the license to buyer 1, to buyer 2, or to both. The cutoffs for this decision are proportional to those from the static model but are adjusted by a factor of $(\delta/\rho)^a$. The higher the value of a, the lower this fraction becomes, thereby increasing the likelihood that the principal will either sell the license exclusively to buyer 1 or at least include him in the allocation. This is because the first buyer makes their payment in period 0, while the second buyer, if included, makes their payment upon arrival. Given the difference in discounting between the principal and the buyers, buyer 2's payment is reweighted, and this reweighting becomes more significant the later they arrive. Consequently, the seller finds it optimal to favor buyer 1 more as the arrival time of buyer 2 is delayed.[7] As it turns out, even if the initial decision is not to wait, the principal employs the same cutoffs in the second step, with the key difference being the absence of an upper cutoff—there is no longer a range of realized types where buyer 2 would receive an exclusive contract.

Finally, the proposition indicates that when a is sufficiently large, an exclusivity region emerges. In other words, the conditions eventually become so favorable for the first buyer, that even if the second buyer arrives with type $\overline{\theta}$, the principal still allocates the good exclusively to the first buyer. For any type of first buyer for whom the principal decides to allocate the license ($q_0^1 = 1$), this exclusivity region emerges in finite time. However, this exclusivity region only materializes if $\delta > \rho$, meaning the seller must be less patient than the buyers. If, instead, $\delta = \rho$, the seller never finds it optimal to write a contract that guarantees exclusivity to the current buyer.

The dynamic mechanism gives rise to two types of inefficiencies compared to the first best. First, when both agents have arrived, the designer may either over-allocate or under-allocate the license, similar to the static problem. Second, when buyer one arrives, the decision to grant him a license may also be inefficient. Compared to the first best, the designer might commit to allocating to buyer types that are too low or, conversely, fail to allocate to buyer types that would be chosen under the first-best outcome. We formally define these two inefficiencies below.

Definition 4. *Let $q_{f,t}^i$ represent the first-best allocations. We say that the allocation q_t in-*

[7]Note that if the seller is more patient than the buyer ($\rho > \delta$), then it would be optimal to postpone any payment as much as possible. With an infinite time horizon, the problem is no longer well-defined. However, if we were to assume that a deadline exists, e.g. the game ends after T periods, then all payments would be postponed to this $t = T$ period. Because both buyers would make payments at that period, the seller does not give preferential treatment to one of the buyers, thus, the optimal cutoffs from the static mechanism are preserved.

duced by a mechanism under- (over-)provides if:

$$q_t^1(\theta, a) + q_t^2(\theta, a) \leq (\geq) q_{f,t}^1(\theta, a) + q_{f,t}^2(\theta, a) \text{ for all } \theta \text{ and } t \geq a \ .$$

We say that the allocation is stringent (lenient) if:

$$q_t^1(\theta_1) \leq (\geq) q_{f,t}^1(\theta_1) \text{ for all } \theta_1 \text{ and } t < a.$$

Proposition 6. *Let λ be increasing (decreasing). Then, for all α, the allocation is stringent (lenient) and always under- (over-) provides the good.*

The proposition demonstrates that these inefficiencies are interconnected: the same conditions that lead to under-provision also imply that the designer becomes more stringent in granting licenses to the first buyer. Thus, the distribution of buyers in a market dictates whether there will be under- or over-provision of the good, as well as whether the principal will adopt a stringent or lenient approach when issuing initial contracts.

4 Model Generalizations

We started our analysis in Section 2 with a simplified model that allowed us to visualize the mechanism's inner workings and build intuition. One advantage of this initial setup was that it required no additional assumptions beyond those typically used in standard auction theory. While this model suffices for certain applications, its limitations are apparent. Notably, we assumed that when transitioning from exclusive control to sharing the market, the buyer's profits were merely scaled by a constant factor, $\alpha < 1$. In reality, a firm's profits in the presence of competition may not simply be a fixed fraction of what they earn as a monopolist. This fraction can vary depending on the buyer's type, and the effect of competition on profits often hinges on the characteristics of the competitor. For instance, if a competitor has significantly higher production costs, α should approach 1, as they represent little competitive pressure. On the other hand, a highly efficient competitor may dominate the market, leading to a much lower α. Thus, when the market is shared, the outcomes for agent i may depend not only on θ_i but also on the type of agent j, specifically θ_j.

We proceed as follows. First, in Section 4.1, we retain the assumption that profits depend solely on a buyer's own type, but we analyze a scenario where the returns from exclusivity are captured by a general supermodular profit function. Next, we relax the assumption that profits depend only on a buyer's own type, allowing them to also depend on the types of other buyers. In Section 4.2, we examine a setup where this dependence is multiplicative. Finally, in Section 4.3, we explore the most general framework, permitting any dependencies between agents' types and payoffs. Table 1 below summarizes the additional sufficient conditions required for these alternative models. We derive these assumptions and provide the underlying intuition in the sections that follow.

Table 1: Sufficient Conditions

	Win Alone	**Win Together**	**Sufficient Conditions**
α **Model**	θ_i	$\alpha_i\theta_i$	$\varnothing$
A buyer's profit depends on own type only.			
Supermodular	θ_i	$u(\theta_i)$	$\max\left\{\frac{-u''(\theta_i)}{1-u'(\theta_i)}, \frac{u''(\theta_i)}{u'(\theta_i)}\right\} \leq \frac{v'(\theta_i)}{h(\theta_i)}$
A buyer's profit depends on both types.			
Multiplicative	θ_i	$g(\theta_j)\theta_i$	$\frac{v'(\theta_i)}{v(\theta_i)} \geq \frac{[1-g(\theta_i)]'}{1-g(\theta_i)}$
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.

4.1 General Supermodularity

Consider the following setting. If a buyer is allocated the good exclusively, their payoff is θ. If, however, the buyer shares the license, his payoff is $u(\theta) < \theta$, where u is an increasing function. Define $\Delta(\theta) = \theta - u(\theta)$. We assume that higher types gain more from exclusivity than lower types ($\Delta' > 0$)—that is, profits are supermodular. Let an allocation be represented by (x^1, x^2), where x^i is the probability of exclusive allocation to agent i. Assuming that the seller always finds it optimal to allocate to some agent, which we impose through conditions on virtual values, we have that $1 - x^1 - x^2$ is the probability that both agents share the license. Feasibility is equivalent to $x^1 + x^2 \leq 1$, with $x^i \in [0,1]$. We can write the expected payoff of agent i with type θ_i as:

$$U(\theta_i) = \mathbb{E}[\Delta(\theta_i)x^i(\theta) + u(\theta_i)\left(1 - x^j(\theta)\right) - r^i(\theta)],$$

where, as before, $\theta = (\theta_1, \theta_2)$. As usual, we denote by $x^i(\theta_k)$ the expectation of $x^i(\theta)$ given that the type of agent k is θ_k. For any θ_i, θ_i', BIC implies

$$U(\theta_i) - U(\theta_i') \geq (\Delta(\theta_i) - \Delta(\theta_i'))\, x^i(\theta_i') + (u(\theta_i) - u(\theta_i'))\,(1 - x^j(\theta_i'))$$

By switching the roles of θ_i and θ_i' we obtain two familiar conditions

$$U'(\theta_i) = \Delta'(\theta_i)x^i(\theta_i) + u'(\theta_i)(1 - x^j(\theta_i)), \tag{Envelope}$$

and

$$(\Delta(\theta_i) - \Delta(\theta_i'))\left(x^i(\theta_i) - x^i(\theta_i')\right) \geq (u(\theta_i') - u(\theta_i))\left(x^j(\theta_i') - x^j(\theta_i)\right). \tag{Monotonicity}$$

To satisfy the monotonicity condition above, it is sufficient that x^i is increasing with own-type θ_i and x^j is decreasing with other-type θ_i (or, equivalently, x^i is decreasing with type θ^j). Under the assumption that these conditions are met, after the usual integration by parts, the problem of the seller becomes

$$\max_{x^1+x^2\leq 1} \sum_{i=1,2} \mathbb{E}\left[\left(\Delta(\theta_i) - \Delta'(\theta_i)\frac{1-F(\theta_i)}{f(\theta_i)}\right)x^i(\theta) + \left(u(\theta_j) - u'(\theta_j)\frac{1-F(\theta_j)}{f(\theta_j)}\right)(1 - x^i(\theta))\right].$$

Without imposing any further constraints, the solution to this problem is $x^i = 1$ if and only if the two inequalities below hold

$$\begin{aligned}&\left(\Delta(\theta_i) - \Delta'(\theta_i)\frac{1-F(\theta_i)}{f(\theta_i)}\right) - \left(u(\theta_j) - u'(\theta_j)\frac{1-F(\theta_j}{f(\theta_j)}\right)\\ &\geq \left(\Delta(\theta_j) - \Delta'(\theta_j)\frac{1-F(\theta_j)}{f(\theta_j)}\right) - \left(u(\theta_i) - u'(\theta_i)\frac{1-F(\theta_i)}{f(\theta_i)}\right) \iff v(\theta_i) \geq v(\theta_j),\end{aligned}$$

and

$$\Delta(\theta_i) - \Delta'(\theta_i)\frac{1-F(\theta_i)}{f(\theta_i)} \geq u(\theta_j) - u'(\theta_j)\frac{1-F(\theta_j)}{f(\theta_j)}. \tag{4}$$

If the second inequality fails, then $x^i = x^j = 0$, implying both agents receive the license. We now find conditions to guarantee that this allocation satisfies monotonicity. Increasing virtual valuations implies that the first inequality is well-behaved. Let $h(\theta_i) = \frac{1-F(\theta_i)}{f(\theta_i)}$. Assume first that the problem is supermodular. In that case, our monotonicity condition requires x^i to be increasing in own-type, and therefore, that inequality (I) becomes more relaxed as θ_i grows. In other words, we want

$$\Delta'(\theta_i)\underbrace{(1-h'(\theta_i))}_{v'(\theta_i)>0}-\Delta''(\theta_i)h(\theta_i) \geq 0$$

At the same time, we want x^i decreasing in θ_j, which would follow from the inequality equation (4) becoming more strict as θ_j grows. This holds if

$$u'(\theta_j)v'(\theta_j) \geq u''(\theta_j)h(\theta_j).$$

Putting these two conditions together and noticing that $\Delta'' = -u''$, we obtain

$$\underbrace{-\Delta'(\theta_i)}_{<0 \text{ by supermodularity}} \frac{v'(\theta_i)}{h(\theta_i)} \leq u''(\theta_i) \leq u'(\theta_i)\frac{v'(\theta_i)}{h(\theta_i)}.$$

Notice that the sign of u'' is sufficient for one of these inequalities. A reasonable assumption is that $u'' > 0$ (equivalently, $\Delta'' < 0$), which states that although the gain from exclusivity increases with type, this increase is smaller and smaller as types grow. Under that assumption, the first inequality above holds immediately, and the second holds if $\frac{u''(\theta_i)}{u'(\theta_i)} \leq \frac{v'(\theta_i)}{h(\theta_i)}$. Of course, the inequalities automatically hold if $u'' = 0$, which is the case in our initial model. Alternatively, it could be that the gains from exclusivity not only increase but increase more as types grow. In this case, $\Delta'' > 0$ which implies that $u'' < 0$. Under this assumption, the second inequality above holds immediately, while the first holds if $\frac{-u''(\theta_i)}{1-u'(\theta_i)} \leq \frac{v'(\theta_i)}{h(\theta_i)}$.

4.2 Multiplicative Model

We modify the setup in the following way. We maintain the assumption that agents' types, θ, are i.i.d. drawn from the same distribution F. If agent i, with type θ_i, is allocated 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 represented as $k \in \{0, 1, 2, 3\}$, where $k \in \{1, 2\}$ indicates that agent i is allocated the good alone, $k = 0$ indicates that no one receives the good, and $k = 3$ indicates that both agents are allocated the good. For each realization $\vartheta = (\theta_1, \theta_2)$, we then define

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

Just as before, it is easy to see that, given a type realization, $\mathcal{P}(\vartheta)$ is a polytope, representing feasible expected payoffs. 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 determine conditions under which $\overline{q}$ satisfies the monotonicity requirement. Defining $\overline{F} = F \circ v^{-1}$, and $\overline{\alpha} = \alpha \circ v^{-1}$ leads to

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

Differentiating with respect to θ_i we obtain

$$\begin{aligned}
&\overline{\alpha}(a_{-i}(\vartheta)v(\theta_i))\overline{f}(a_{-i}(\vartheta)v(\theta_i)) \qquad \underbrace{\left(\frac{da_{-i}(\vartheta)}{d\theta_i}v(\theta_i)+a_{-i}(\vartheta)v'(\theta_i)\right)}_{>0\text{: conditional on being the highest type you cannot lose}} \\
&+\left(\beta-\overline{\alpha}\left(\frac{v(\theta_i)}{a_i(\vartheta)}\right)\right)\overline{f}\left(\frac{v(\theta_i)}{a_i(\vartheta)}\right) \qquad \underbrace{\left(\frac{v'(\theta_i)}{a_i(\vartheta)}-\frac{da_i(\vartheta)}{d\theta_i}\frac{v(\theta_i)}{a_i^2(\vartheta)}\right)}_{><0\text{? Conditional on being lowest you may lose}}.
\end{aligned}$$

A sufficient condition for this expression to 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 *increase fater* than $\beta-\alpha(\theta_i)$. Without loss we can normalize $\beta = 1$, obtaining the condition in Table 1.

4.3 General Model

Once again, we maintain the assumption that agents' types, θ, are 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 $\theta, \nu \in \operatorname{supp} F$

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

Given θ_1 and θ_2, we interpret q^i as the probability that agent i is allocated the good alone, and q_α as the probability that both agents are allocated the good together. 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 3. *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. $\mathbb{E}\left[(\gamma(\theta, \theta_{-i}) - \gamma(\theta', \theta_{-i})) \cdot (q_i(\theta, \theta_{-i}) - q_i(\theta', \theta_{-i}))\right] \geq 0$ *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 2. *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 7. *Let Assumption 2 hold. When* q_i *and* $q_i + q_\alpha$ *are increasing, condition 1 in Lemma 3 is sufficient for implementability.*

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

$$\sum_i \int_\theta \mathbb{E}_{-i}\left[\left(\gamma(\theta, \theta_{-i}) - \frac{1 - F(\theta)}{f(\theta)}\gamma'(\theta, \theta_{-i})\right) \cdot q_i(\theta, \theta_{-i})\right] f(\theta)d\theta. \tag{5}$$

Assumption 3. *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 8. *Under Assumption 2 and Assumption 3, the revenue-maximizing mechanism has allocations:*

$$q_i(\theta, \theta_{-i}) = \begin{cases} 1 & \text{if } v_\beta(\theta, \theta_{-i}) > \max\left\{v_\alpha(\theta, \theta_{-i}) + v_\alpha(\theta_{-i}, \theta), v_\beta(\theta_{-i}, \theta)\right\} \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

$$q_\alpha(\theta, \theta_{-i}) = \begin{cases} 1 & \text{if } \max\{v_\beta(\theta, \theta_{-i}), v_\beta(\theta_{-i}, \theta)\} < v_\alpha(\theta, \theta_{-i}) + v_\alpha(\theta_{-i}, \theta) \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

Proposition 8 completes the characterization for the general model. In Section **??**, we explore an application that is feasible with the general model but would have been unmanageable with the baseline model.

5 Applications

5.1 Selling Information in Financial Markets

We consider the market for one risky security with payoff $v \in \{0, 1\}$. Trade happens at time 0, and the payoff of the asset is revealed at time 1. There are $N > 2$ traders in the market: $N - 2$ being liquidity traders and 2 rational investors. Our trading protocol is inspired by Glosten and Milgrom (1985). At time 0, perfectly competitive market makers publicly post a price at which they stand ready to buy (bid, b) and sell (ask, a) the security. Subsequently, each trader interested in buying or selling is randomly matched with a market maker, and they trade 1 unit of the security at the posted price.

The payoff of a trader with marginal utility of wealth θ who buys one unit of the asset at the ask price is $\theta(v - a)$. If the same buyer were to sell the asset at the bid price, the payoff would be $\theta(b - v)$. We assume that the marginal utility of wealth, θ, is private information and symmetrically distributed according to a continuous distribution F. Rational investors trade to maximize their expected payoff. Liquidity traders trade randomly: for simplicity, we assume that liquidity traders are always willing to trade, and buy or sell with the same probability.

At time 0, all traders and market makers share a common prior assigning equal probability to $v \in \{0, 1\}$. Before prices are posted, an information seller (the principal) who is fully informed about the value of the security can sell that information to one or both of the rational traders. If only one of the rational investors is informed, it will be optimal for the uninformed rational investor to not trade, so the proportion of informed investors on the pool of traders is $\eta = \frac{1}{N-1}$. If both rational investors are informed, then all traders are active in the market and that fraction is $\eta = \frac{2}{N}$. We solve for the equilibrium in the financial market given α.

Market makers are competitive, but they are aware of adverse selection, which will give rise to a bid-ask spread in equilibrium. For example, upon observing a buying demand, a market maker knows that there is a probability that they are observing an informed trader, which implies that the asset value is 1. To protect themselves against that possibility, they raise their ask price. In equilibrium we must have:

$$a = \mathbb{E}[v|\text{buy}] = \frac{1+\eta}{2} \qquad\qquad b = \mathbb{E}[v|\text{sell}] = \frac{1-\eta}{2}$$

Therefore, if there is only one informed investor, her payoff is:

$$\pi^M(\theta) = \frac{1}{2}\frac{N-2}{N-1}\theta \propto \theta,$$

and the payoff of an informed investors when both are informed is:

$$\pi^D(\theta) = \frac{N-1}{N}\pi^M(\theta).$$

Therefore, this application fits into our baseline model, and the revenue-maximizing allocation of information follows Proposition 1.

5.2 Horizontally Differentiated Products

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.

Figure 6: Hotelling Application

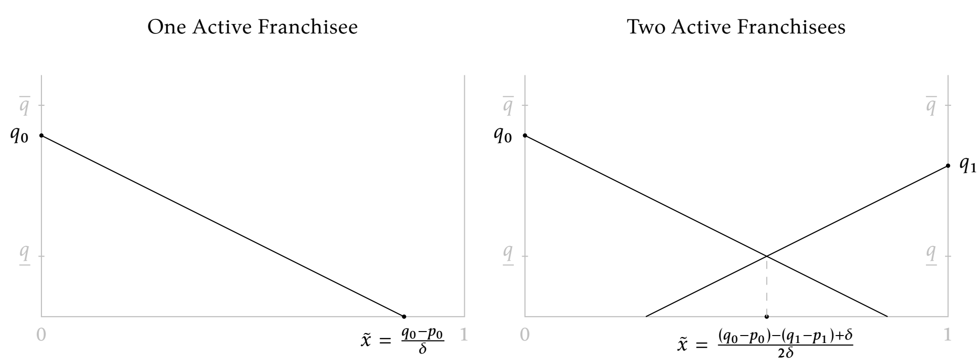

Notes: The figure above displays...

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

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

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 4 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 4 are met. Thus, the principal can maximize expected profits by simply running an interval auction.

6 Conclusions

This paper explores the optimal licensing strategy for a seller facing downstream competitors with private information, relevant in various market contexts, such as franchise operations, patent licensing, and information sales, to name a few. Stemming from asymmetric information, we characterize inefficiencies that do not arise in conventional auctions, leading to scenarios where the seller may either over- or under-provide the good. We link these inefficiencies to the distribution of buyer valuations. We propose an *interval auction* as the revenue-maximizing mechanism, where the allocation decision is based not only on the highest bid but on the distribution of all bids. If bids are closely clustered, the mechanism favors selling to multiple bidders; if bids are widely dispersed, exclusive licensing to the highest bidder becomes optimal. In a dynamic setting where buyers arrive sequentially, we analyze the timing of licensing decisions and show that a seller may delay licensing or issue a license immediately to an available buyer. Furthermore, we show that the decision to offer an exclusive license depends on the seller's relative patience compared to the buyers; exclusivity is promised only when the seller is less patient. Lastly, we explore sufficient conditions that allow for a framework where buyer valuations are interdependent.

7 Appendix

7.1 Unknown Market Structure: Reported Interval Auction

We now consider an implementation of the optimal mechanism identified in Section 2.2 when the seller does not know the market structure, namely α, but the buyers do. Consider the following *reported interval auction*. Both bidders are asked to submit intervals $[l_i, h_i]$. From these intervals, the principal calculates

$$\hat{\alpha}_i = \frac{\sqrt{v(h_i)}}{\sqrt{v(h_i)} + \sqrt{v(l_i)}}, \qquad \hat{\theta}_i = v^{-1}\left(\frac{\hat{\alpha}_i}{1-\hat{\alpha}_i} v(l_i)\right)$$

Where $\hat{\alpha}_i$ and $\hat{\theta}_i$ represent the principal's estimated α and θ_i respectively, based on the interval submitted by bidder i, while γ is some constant. The allocation rules are as follows: with equal probability, the principal considers $\hat{\theta}_i$ and $[l_{-i}, h_{-i}]$ or $\hat{\theta}_j$ and $[l_i, h_i]$ to determine the allocation. Consider without loss that $\hat{\theta}_i$ and $[l_{-i}, h_{-i}]$ were chosen.

1. If $\hat{\theta}_i$ falls in $[l_{-i}, h_{-i}]$, both bidders are allocated the good and pay $\hat{\alpha}_{-i} l_{-i}$.

2. If $\hat{\theta}_i > h_{-i}$, the good is allocated to bidder i exclusively who pays $\hat{\alpha}_{-i} l_{-i} + (1-\hat{\alpha}_{-i}) h_{-i}$, while bidder $-i$ pays nothing.

3. If $\hat{\theta}_i < l_{-i}$, the good is allocated to bidder $-i$ exclusively who pays $\hat{\alpha}_i l_i + (1-\hat{\alpha}_i) h_i$, while bidder i pays nothing.

Note that if $-i$ submits interval $[\underline{\tau}(\theta_{-i}), \overline{\tau}(\theta_{-i})]$, then

$$\hat{\alpha}_{-i} = \frac{\sqrt{v(\overline{\tau}(\theta_{-i}))}}{\sqrt{v(\overline{\tau}(\theta_{-i}))} + \sqrt{v(\underline{\tau}(\theta_{-i}))}} = \alpha,$$

$$\hat{\theta}_{-i} = v^{-1}\left(\frac{\hat{\alpha}_i}{1-\hat{\alpha}_i} v(\underline{\tau}(\theta_{-i}))\right)\gamma + v^{-1}\left(\frac{1-\hat{\alpha}_i}{\hat{\alpha}_i} v(\overline{\tau}(\theta_{-i}))\right)(1-\gamma) = \theta_{-i}.$$

Proposition 9. *Under assumptions 1, 2, and 3, the optimal mechanism is implemented as a Bayesian Nash equilibrium by a reported interval auction.*

Proof of Proposition 9

Assume that bidder $-i$ reports truthfully, so $[\underline{\tau}(\theta_{-i}), \overline{\tau}(\theta_{-i})]$. Start by assuming that bidder i reports $[\underline{\tau}(\theta_i), \overline{\tau}(\theta_i)]$. We will see if there is ever a profitable deviation from this report. First, note that deviations that do not change the allocations do not change payoffs either. Thus, from such deviations, the bidder stands to gain nothing. In what follows, we now consider deviations that may change the allocation.

Case 1: $\theta_i < \underline{\tau}(\theta_{-i})$

Deviation 1.1 Consider misreporting the interval upwards so that both bidders are allocated the good. The utility of bidder i would have been 0 but now is $\alpha\theta_i - \alpha\underline{\tau}(\theta_{-i}) < 0$. This deviation hurts the bidder.

Deviation 1.2 Consider misreporting the interval upwards so that bidder i is allocated the good exclusively. The utility of bidder i would have been 0 but now is $\theta_i - \alpha\underline{\tau}(\theta_{-i}) - (1-\alpha)\overline{\tau}(\theta_{-i}) < 0$. The inequality follows by substituting $\overline{\tau}(\theta_{-i})$ with $\underline{\tau}(\theta_{-i})$ and once again realizing that $\theta_i < \underline{\tau}(\theta_j)$. This deviation hurts the bidder.

Case 2: $\underline{\tau}(\theta_j) < \theta_i < \overline{\tau}(\theta_j)$

Deviation 2.1 Consider misreporting the interval downwards so that bidder i is excluded. The utility of bidder i would have been $\alpha\theta_i - \alpha\underline{\tau}(\theta_j)$ but now is 0. This deviation hurts the bidder.

Deviation 2.2 Consider misreporting the interval upwards so that bidder i is allocated the good exclusively. The utility of bidder i would have been $\alpha\theta_i - \alpha\underline{\tau}(\theta_j)$ but is now $\theta_i - \alpha\underline{\tau}(\theta_j) - (1-\alpha_{-i})\overline{\tau}(\theta_j)$. The bidder changed his utility by $(1-\alpha)\theta_i - (1-\alpha)\overline{\tau}(\theta_j) < 0$. This deviation hurts the bidder.

Case 3: $\overline{\tau}(\theta_j) < \theta_i$

Deviation 3.1 Consider misreporting the interval downwards so that both bidders are allocated the good. The utility of bidder i would have been $\theta_i - \alpha\underline{\tau}(\theta_j) - (1-\alpha_{-i})\overline{\tau}(\theta_j)$ but now is $\alpha\theta_i - \alpha\underline{\tau}(\theta_j)$. The bidder changed his utility by $(1-\alpha)\overline{\tau}(\theta_j) - (1-\alpha)\theta_i < 0$. This deviation hurts the bidder.

Deviation 3.2 Consider misreporting the interval downwards so that bidder i is excluded. The utility of bidder i would have been $\theta_i - \alpha\underline{\tau}(\theta_j) - (1-\alpha_{-i})\overline{\tau}(\theta_j)$ but now is 0. The bidder changed his utility by $-\theta_i + \alpha\underline{\tau}(\theta_j) + (1-\alpha_{-i})\overline{\tau}(\theta_j) < 0$. This deviation hurts the bidder.

There are either deviations that do not have any impact or deviations that hurt bidder i. Because bidder i does not know the interval reported by bidder $-i$, in expectation, any deviation will hurt him. Thus, the bidder does not find it beneficial to deviate from reporting $[\underline{\tau}(\theta_i), \overline{\tau}(\theta_i)]$. ■

7.2 Proofs

Proof of Proposition 2

We can parameterize the problem by $\alpha, \beta > 0$, 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. Recall, $\alpha \leq \beta$, 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 $\mathcal{Q}$: 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 1, 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

$$\alpha(\theta^1 + \theta^2) \geq \beta\theta^1 \iff \frac{\alpha}{\beta - \alpha} \geq \frac{\theta^1}{\theta^2}. \tag{8}$$

By contrast, the optimal mechanism serves both agents only if

$$\alpha\left(v(\theta^1) + v(\theta^2)\right) \geq \beta v(\theta^1) \iff \frac{\alpha}{\beta - \alpha} \geq \frac{v(\theta^1)}{v(\theta^2)}. \tag{9}$$

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, we provide the proof of under-provision. The proof for over-provision is symmetric. Note that $v(x) = \left(1 - \frac{1-F(x)}{f(x)x}\right)x = (1 - \frac{1}{\lambda(x)})x$. Without loss of generality, let $\theta_1 \geq \theta_2$. First, assume that λ is increasing. Then, if the first best allocates to only one agent:

$$\frac{v(\theta_1)}{v(\theta_2)} = \frac{1 - \frac{1}{\lambda(\theta_1)}}{1 - \frac{1}{\lambda(\theta_2)}}\frac{\theta_1}{\theta_2} \geq \frac{\theta_1}{\theta_2} \geq \frac{\alpha}{1-\alpha},$$

and the revenue-maximizing mechanism also allocates to only one agent. Thus, the revenue-maximizing mechanism can only under-provide.

For the converse, assume the revenue-maximizing mechanism under-provides for all $\alpha \in [0, 1]$. Fix $\theta_1 > \theta_2$ and let α be such that:

$$\frac{\alpha}{1-\alpha} = \frac{\theta_1}{\theta_2} \leq \frac{1 - \frac{1}{\lambda(\theta_1)}}{1 - \frac{1}{\lambda(\theta_2)}}\frac{\theta_1}{\theta_2},$$

where the inequality follows from the assumption. We thus have $\lambda(\theta_1) \geq \lambda(\theta_2)$. Because $\theta_1 > \theta_2$ were arbitrary, the result follows. ■

Proof of Proposition 3

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:

$$t_\beta - t_\alpha \leq g^{-1}(\theta_2)(\beta - \alpha). \tag{10}$$

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

$$t_\beta \leq g^{-1}(\theta_2)\beta. \tag{11}$$

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 10, we obtain an expression for the difference in payments:

$$t_\beta - t_\alpha = g^{-1}(\theta_2)(\beta - \alpha). \tag{12}$$

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 Proposition 5

Define x^1 and x^2 as follows

$$U^1(\theta_1) = \mathbb{E}_{\theta_2}\left[\underbrace{\sum_{j=0}^{\infty} \delta^j (1-\lambda)^j q_j^1(\theta_1) + \sum_{j=0}^{\infty} \lambda(1-\lambda)^j \sum_{k=j+1}^{\infty} \delta^k q_k^1(\theta_1, \theta_2, a^2 = j+1)}_{x^1}\right]\theta_1 - r^1(\theta_1),$$

$$U^2(\theta_1,\theta_2,a^2) = \left[\underbrace{\sum_{j=0}^{\infty} \delta^j q^2_{a^2+j}(\theta_1,\theta_2,a^2)}_{x^2}\right]\theta_2 - r^2(\theta_1,\theta_2).$$

$$\mathbb{E}[r^1(\theta_1,\theta_2) + \sum_{j=0}^{\infty} \rho^{j+1}\lambda(1-\lambda)^j r^2(\theta_1,\theta_2,a^2 = j+1)]$$

Using integration by parts, the seller maximizes:

$$\mathbb{E}\left[v(\theta_1)x_1(\theta_1,\theta_2) + \rho\sum_{i=0}^{\infty}\rho^i\lambda(1-\lambda)^i v(\theta_2)x_2(\theta_1,\theta_2,a^2 = i+1)\right]$$

$$= \mathbb{E}\sum_{j=0}^{\infty}\delta^j(1-\lambda)^j q^1_j(\theta_i)v(\theta_1) + \sum_{j=0}^{\infty}\lambda(1-\lambda)^j\sum_{k=j+1}\delta^k q^1_k(\theta_i,a^2 = j+1)v(\theta_1) +$$
$$\sum_{j=0}^{\infty}\rho^{j+1}\lambda(1-\lambda)^j\sum_{k=j+1}^{\infty}\frac{\delta^k}{\delta^{j+1}}q^2_k(\theta,a^2 = j+1)v(\theta_2)$$

So if we fix any j, $a^2 = j+1$ and any time $k > j+1$ we have that the seller solves, given an irreversibility constraint q:

$$\max_{q^1\geq q}\lambda(1-\lambda)^j\delta^k\left(q^1_k(\theta,a^2 = j+1)v(\theta_1) + \frac{\rho^{j+1}}{\delta^{j+1}}q^2_k(\theta,a^2 = j+1)v(\theta_2)\right).$$

When the irreversibility constraint does not bind, we have:

$$q^1_t(\theta_1,\theta_2,a^2) = 1 \iff \frac{v(\theta_1)}{v(\theta_2)} \geq \left(\frac{\rho}{\delta}\right)^{a^2}\frac{\alpha}{1-\alpha},$$

$$q^1_t(\theta_1,\theta_2,a^2) = \alpha \iff \left(\frac{\rho}{\delta}\right)^{a^2}\frac{1-\alpha}{\alpha} \leq \frac{v(\theta_1)}{v(\theta_2)} \leq \left(\frac{\rho}{\delta}\right)^{a^2}\frac{\alpha}{1-\alpha},$$

and

$$q^1_t(\theta_1,\theta_2,a^2) = \alpha q \iff \left(\frac{\rho}{\delta}\right)^{a^2}\frac{1-\alpha}{\alpha} \geq \frac{v(\theta_1)}{v(\theta_2)}.$$

It is clear that, because of discounting, the seller has incentives to frontload the solo allocation of the good for agent 1, $q^1(\theta)$ before the arrival of agent 2. By irreversibility, that allocation cannot decrease until agent 2 arrives, so it is without loss of generality to consider an $q_t^1(\theta) = q$. The profit of the seller is then:

$$\frac{1}{1-\delta(1-\lambda)} q v(\theta_1) + \sum_{j=0}^{\infty} \frac{\lambda}{1-\delta}(1-\lambda)^j$$
$$\mathbb{E}_{\theta_2} \max\left\{\delta^{j+1} v(\theta_1), \alpha\left(\delta^{j+1} v(\theta_1) + \rho^{j+1} v(\theta_2)\right), q\alpha\left(\delta^{j+1} v(\theta_1) + \rho^{j+1} v(\theta_2)\right) + (1-q)\rho^{j+1} v(\theta_2)\right\}$$

This function is affine in q. To see that, fix any θ_2. Note that the third term in the max is the maximum of the three for $q = 0$ if and only if it is also the maximum for $q = 1$. In other words, for a fixed θ_2, either the max does not change with q, in which case the expression above is affine in q; or the max changes linearly in q, so the expression above is again affine in q. Thus, once one takes expectation in θ_2, the expression above is still affine in q. Therefore, $q \in \{0, 1\}$.

Having established the possible values of q, we next show that $q^w(\theta_1)$ is monotonically decreasing. It is optimal to choose $q = 0$ instead of $q = 1$ if

$$\begin{aligned}\sum_{j=0}^{\infty} \frac{\lambda}{1-\delta}(1-\lambda)^j \ \mathbb{E}_{\theta_2} \Big[&\max\left\{\delta^{j+1} v(\theta_1), \alpha\left(\delta^{j+1} v(\theta_1) + \rho^{j+1} v(\theta_2)\right), \rho^{j+1} v(\theta_2)\right\} \\ &- \max\left\{\delta^{j+1} v(\theta_1), \alpha\left(\delta^{j+1} v(\theta_1) + \rho^{j+1} v(\theta_2)\right)\right\}\Big] \\ &- \frac{1}{1-\delta(1-\lambda)} v(\theta_1) \geq 0\end{aligned}$$

At $\theta_1 = \overline{\theta}$, the difference within the expectation operator is 0, implying that the whole term is negative. Since the inequality does not hold the optimal decision is $q = 1$. At $\theta_1 = \underline{\theta}$, the inequality cannot be generally signed. Notice that if $v(\underline{\theta}) \leq 0$, then we obtain that it is always optimal to wait for $\theta_1 = \underline{\theta}$. Thus, for high enough θ_1 the optimal q surely becomes 1, but might be 0 or 1 for low values of θ_1. Fix a θ_2 value, and note that the difference between the max expressions weakly decreases as θ_1 increases, while the $v(\theta_1)$ term outside also decreases, leading to a decrease of the total expression. Since this holds for any θ_2 value, it also holds in expectation. Thus, q is either 1 to begin with, or goes from 0 to 1 as θ_1 increases. ■

Proof of Proposition 6

We can establish the inefficiencies related to the dynamic revenue-maximizing mechanism. To do that, start by noticing that, in that mechanism, for a fixed θ_1, the designer waits if and only if:

$$\frac{1-\delta(1-\lambda)}{1-\delta}\mathbb{E}_{a,\theta_2}\left[\max\left\{\delta^a, \alpha\left(\delta^a+\rho^a\frac{v(\theta_2)}{v(\theta_1)}\right), \rho^a\frac{v(\theta_2)}{v(\theta_1)}\right\} - \max\left\{\delta^a, \alpha\left(\delta^a+\rho^a\frac{v(\theta_2)}{v(\theta_1)}\right)\right\}\right] \geq 1,$$

whereas in the first best, the designer waits if and only if:

$$\frac{1-\delta(1-\lambda)}{1-\delta}\mathbb{E}_{a,\theta_2}\left[\max\left\{\delta^a, \alpha\left(\delta^a+\rho^a\frac{\theta_2}{\theta_1}\right), \rho^a\frac{\theta_2}{\theta_1}\right\} - \max\left\{\delta^a, \alpha\left(\delta^a+\rho^a\frac{\theta_2}{\theta_1}\right)\right\}\right] \geq 1.$$

Notice first that the expression inside the expectation on the left-hand side of the inequalities above is different from zero only if $\theta_2 > \theta_1$. We focus on that case from now on. Consider the function $H(\phi) = \max\{\delta^a, \alpha(\delta^a + \rho^a\phi), \rho^a\phi\} - \max\{\delta^a, \alpha(\delta^a + \rho^a\phi)\}$. This function is increasing in ϕ.

If $\frac{v(\theta_2)}{v(\theta_1)} \geq \frac{\theta_2}{\theta_1}$ for $\theta_2 > \theta_1$, then, conditional on θ_1, the distribution over $(\frac{v(\theta_2)}{v(\theta_1)}, a)$ first-order stochastically dominate the distribution over $(\frac{\theta_2}{\theta_1}, a)$. Hence, for all θ_1's such that the designer chooses to wait in the first-best, she also chooses to wait in the second-best. Notice that this is the same condition as underprovision for all α values (namely, λ is decreasing). ■

Proof of Lemma 3

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

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

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):

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

■

Proof of Proposition 7

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

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

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 2, 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 8

We solve the relaxed problem of maximizing profits subject to condition 1 in Lemma 3 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:

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

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:

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

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

Antelo, M. and Sampayo, A. (2017). On the number of licenses with signalling. *The Manchester School*, 85(6):635–660.

Antelo, M. and Sampayo, A. (2024). Licensing of a new technology by an outside and uninformed licensor. *Journal of Economics*, 142(2):111–162.

Armstrong, M. (2000). Optimal multi-object auctions. *The Review of Economic Studies*, 67(3):455–481.

Avery, C. and Hendershott, T. (2000). Bundling and optimal auctions of multiple products. *The Review of Economic Studies*, 67(3):483–497.

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. (1996). Auctions versus negotiations. *The American Economic Review*, 86(1):180–194.

Choi, J. P. (2001). Technology transfer with moral hazard. *International Journal of Industrial Organization*, 19(1-2):249–266.

Crémer, J. and McLean, R. P. (1985). Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent. *Econometrica*, 53(2):345–361.

Crémer, J. and McLean, R. P. (1988). Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, pages 1247–1257.

Doganoglu, T. and Inceoglu, F. (2014). Licensing of a drastic innovation with product differentiation. *The Manchester School*, 82(3):296–321.

Fan, C., Jun, B. H., and Wolfstetter, E. G. (2018). Optimal licensing under incomplete information: the case of the inside patent holder. *Economic Theory*, 66(4):979–1005.

Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of financial economics*, 14(1):71–100.

Heywood, J. S., Li, J., and Ye, G. (2014). Per unit vs. ad valorem royalties under asymmetric information. *International Journal of Industrial Organization*, 37:38–46.

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.

Jeon, H. (2019). Licensing and information disclosure under asymmetric information. *European Journal of Operational Research*, 276(1):314–330.

Kamien, M. I. (1992). Patent licensing. *Handbook of game theory with economic applications*, 1:331–354.

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

Kamien, M. I. and Tauman, Y. (1986). Fees versus royalties and the private value of a patent. *The Quarterly Journal of Economics*, 101(3):471–491.

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

Li, C. and Wang, J. (2010). Licensing a vertical product innovation. *Economic Record*, 86(275):517–527.

McAfee, R. P., McMillan, J., and Reny, P. J. (1989). Extracting the surplus in the common-value auction. *Econometrica: Journal of the Econometric Society*, pages 1451–1459.

McAfee, R. P. and Reny, P. J. (1992). Correlated information and mecanism design. *Econometrica: Journal of the Econometric Society*, pages 395–421.

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.

Poddar, S., Sinha, U. B., et al. (2002). The role of fixed fee and royalty in patent licensing. Technical report, Citeseer.

Schmitz, P. W. (2002). On Monopolistic Licensing Strategies under Asymmetric Information. *Journal of Economic Theory*, 106(1):177–189.

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.

Wu, C.-T., Peng, C.-H., and Tsai, T.-S. (2021). Signaling in technology licensing with a downstream oligopoly. *Review of Industrial Organization*, 58(4):531–559.

Zhang, H., Wang, X., Qing, P., and Hong, X. (2016). Optimal licensing of uncertain patents in a differentiated stackelberg duopolistic competition market. *International Review of Economics & Finance*, 45:215–229.