Selling to Competitors *

Pëllumb Reshidi † João Thereze ‡

[Click Here for the Latest Version]

November 4, 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 how 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. 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 prevalence of these scenarios, 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 only on their private type and the number of competitors licensed. 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.[1]

Our first result identifies the optimal direct mechanism that maximizes the seller's profit. While our results hold for any number of buyers, we illustrate here the main intuition using the case of two buyers. This mechanism allocates the good to a single

[1] It is not essential for the seller to know the market structure; it suffices that buyers know it. We demonstrate that alternative implementations can achieve profit-maximizing outcomes even when the seller lacks knowledge of the market structure.

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 of the seller's allocation problem, our findings reveal that implementation is relatively simple, making it a viable option for real-world application.

Next, we characterize inefficiencies absent in standard 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 under symmetric information is governed by the ratio of valuations, whereas with private information, it is determined by the ratio of *virtual* valuations. 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—that holds for any level of externality. Put simply our result reveals that a policymaker can assess whether a good will be over or under-supplied in a market based solely on the distribution of valuations without needing to know about market conduct or the magnitude of externalities.

Focusing on the two-buyer case, 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 the other buyer. When the seller faces both 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 upon arrival. 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. Once again, we establish a direct link between these inefficiencies and the shape of the distribution of buyer valuations using the exact same conditions established 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 and flow from the buyer to the seller. In this setup, the seller offers the initial buyer a contract that specifies different allocation rules for each period. Revenue maximizing 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, in the two-buyer case, 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 shifts the analysis from independent to 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).[2] 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.[3] 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.[4] 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 static setup is the work of Schmitz (2002), who considers selling a license to two potential buyers. They determine the profit-maximizing mechanism and highlight that potential inefficiencies 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 relates to mechanism design literature with externalities, particularly Jehiel et al. (1996) and Jehiel et al. (1999), which study multidimensional settings with unknown market structures. In contrast, we model the market structure as a function of buyers' types, reducing dimensionality and improving tractability for a full characterization of the optimal mechanism. Additionally, our approach permits multiple sales of the good and considers externalities based on the opponent's realized type, not just their identity.

Relevant to our work is Dana Jr and Spier (1994), who examine mechanisms for

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

[3]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).

[4]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.

auctioning licenses for production rights to one or two producers. Similar to our findings, they show the optimal number of licenses is determined endogenously, with inefficiencies due to information asymmetries. However, in their setup, inefficiencies solely lead to underprovision, with monopolies being assigned more frequently than duopolies. In contrast, in our setup, inefficiencies can occur in either direction, and we precisely characterize when they arise. We further differ by identifying a dominant strategy implementation, reducing the seller's required information disclosure, and enabling portability to a dynamic framework without dynamic disclosure concerns. We then finally differ from this work by studying a dynamic setup.

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).[5] We differ from this body of literature by allowing for the sale of multiple goods.

Multi Unit Auctions Finally, 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.

[5]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.

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 buyers. Consider the 2^N possible subsets of $\mathcal{N}$, which we denote by $\mathcal{J}$, and let the ℓ^{th} subset be denoted by $\mathcal{J}_\ell$. The cardinality of $\mathcal{J}_\ell$ is represented by $|\ell|$. For any subset $\mathcal{J}_\ell \subseteq \mathcal{N}$, the payoff of buyer $i \in \mathcal{J}_\ell$, when only members of $\mathcal{J}_\ell$ are allocated the good is given by

$$u(\theta_i, \mathcal{J}_\ell) = \theta_i \alpha_{|\ell|}.$$

We normalize the payoffs of agents who are not allocated the good, $i \notin \mathcal{J}_\ell$, to zero—effectively excluding them from this market. Thus, utilities are characterized by a private benefit from purchasing the good θ_i and an externality coefficient determined by the number of winners $\alpha_{|\ell|}$. We assume that $\alpha_k \geq \alpha_{k+1}$ for any $k < N$. That is, as the size of the winning group grows, the payoffs of the winning members shrink. Let $a = \{\alpha_1, \alpha_2, \dots, \alpha_N\}$ represent the vector of externality coefficients. Our main assumption is that a is common knowledge, while each agent's taste for the good is private information.[6]

An allocation is a distribution on the family of subsets of $\mathcal{N}$, $\sigma \in \Delta\mathcal{J}$, and due to replicability, the auctioneer can supply any of these subsets. Given the setup, the revelation principle applies, allowing us to focus on identifying the truthful direct revelation mechanism that maximizes revenue.

Finally, each θ_i is assumed to be independently drawn from a distribution with a differentiable cumulative function F, with density f, and full support on $[\underline{\theta}, \overline{\theta}]$ for some $\underline{\theta} > 0$. Define the virtual valuation of a buyer in the usual way: $v(\theta) = \theta - \frac{1-F(\theta)}{f(\theta)}$. We maintain throughout the paper the standard assumption that the distribution F is regular, that is, $v(\theta)$ is a strictly increasing function.

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 r^i and an

[6]As we show later, the principal does not need to be the one who knows the magnitude of the externalities. The optimal mechanism can be implemented even if only the buyers know a.

allocation σ to solve:

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

$$\text{s.t.} \quad \theta_i \sum_{\ell : i \in \mathcal{J}_\ell} \sigma_\ell \alpha_{|\ell|} - r^i \geq 0 \qquad \text{for all } i = 1, \dots, N \tag{IR}$$

It is clear that (IR) must hold with equality in any solution. Thus, the problem can be simplified to an accounting problem: the seller considers the maximal gross 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 illustrated when $N = 2$, as shown in Figure 1. In this case, we normalize the payoff from being allocated the good exclusively to θ_i, while $\alpha\theta_i$ represents the payoff when the good is shared between both buyers. It is optimal to sell solely to buyer i if $\frac{\theta_i}{\theta_j} \geq \frac{\alpha}{1-\alpha}$. On the other hand, selling to both buyers i and j is optimal if $\frac{\alpha}{1-\alpha} \geq \frac{\theta_i}{\theta_j} \geq \frac{1-\alpha}{\alpha}$. Importantly, the optimal allocation is determined by the ratio of valuations θ_i/θ_j.

Figure 1: First Best Allocations

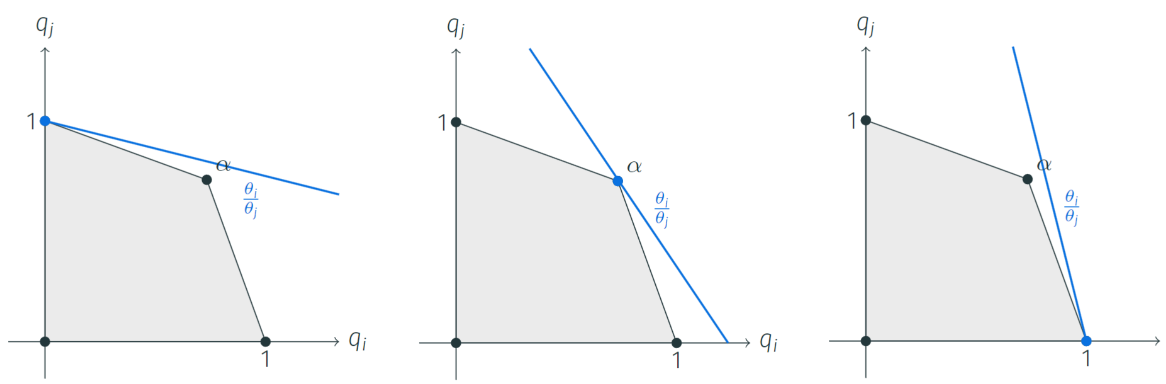

Notes: The figure above displays the first-best allocations for different realized values of θ_i and θ_j. In the left panel, it is efficient to allocate the good exclusively to agent j. In the middle panel, it is efficient to allocate the good to both agents, while in the right panel, it is efficient to allocate exclusively to agent i.

2.2 Revenue-Maximization under Asymmetric Information

Next, we characterize the revenue-maximizing mechanism when the seller does not observe the realized profile of buyer types, θ. Our first observation is that we can change the space of allocations for each buyer 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 gross payoff for agent i:

$$\mathbb{E}_\sigma[u(\theta, \mathcal{J}_\ell)] = \theta_i \underbrace{\sum_{\ell : i \in \mathcal{J}_\ell} \sigma_\ell \alpha_{|\ell|}}_{q^i(\sigma)}$$

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

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

For any set $\mathcal{J}_\ell$, let $\mathbb{1}_\ell$ be a vector in which the $i - th$ entry is 1 if and only if $i \in \mathcal{J}_\ell$. It is clear that:

Lemma 1. $\mathcal{Q} = \text{co}\left\{ \alpha_{|\ell|} \mathbb{1}_\ell : \ell \in \{1, \dots, 2^N\} \right\}$. *$\mathcal{Q}$ is a convex polytope.*

This transformation allows us to identify conditions for implementability by using the standard Myersonian approach. For an assignment q^i define an expected assignment Q^i,

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

and

$$U^i(\theta_i) = \theta_i Q^i(\theta_i) - \underbrace{\int \left[\sum_\ell \sigma_\ell(\theta_i, \theta_{-i}) r^i_\ell(\theta_i, \theta_{-i}) \right] dF_{-i}(\theta_{-i})}_{R^i(\theta_i)}.$$

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 implementability, with the exception of feasibility. Unlike in standard auctions that confine allocation probabilities to the unit simplex (Myerson, 1981), our feasibility condition requires trade probabilities lie within the polytope $\mathcal{Q}$, which in general extends beyond the unit simplex. The problem of the principal then reduces to

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

Recall that the virtual valuation of a type θ_i agent is $v(\theta_i)$. 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 assignment in the optimal mechanism, as a consequence of the steps above. For each realization of buyers' private information, there is no loss of generality in reordering buyers so that $\theta_1 \geq \theta_2 \geq \ldots \geq \theta_N$. Define $\alpha_0 = 0$, with the interpretation that the seller can always choose to exclude all the buyers to obtain 0 revenues.

Proposition 1. *In the optimal mechanism, assignments satisfy:*

$$q^i(\theta) = \begin{cases} \alpha_{k^*(\theta)} & if \quad i \leq k^*(\theta) \\ 0 & otherwise \end{cases} \qquad k^*(\theta) \in \arg \max_{n \in \{0,1,\ldots,N\}} \alpha_n \sum_{i=1}^{k} v(\theta_i).$$

The optimal assignment can be reinterpreted as a deterministic allocation. The seller serves a set of size $k^*(\theta)$, consisting of the highest θ_i values. Therefore, she only needs to evaluate N possible outcomes that differ in the number of buyers. The profit-maximizing outcome depends on the externality coefficients, α, as well as the realized types of buyers, θ. Again, it is easy to illustrate these allocations for $N = 2$, where they satisfy

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

In the $N = 2$ case, if θ_1 is much larger than θ_2, the seller finds it optimal to sell to buyer one exclusively, or vice versa if θ_2 is much larger than θ_1. The intuition is as follows. When θ_1 is much larger than θ_2, the additional profit from including the second buyer, $\alpha v(\theta_2)$, is relatively small. However, by including the second buyer, the seller reduces the amount she can extract from the first buyer, from $v(\theta_1)$ to $\alpha v(\theta_1)$. Therefore, when θ_1 is much larger than θ_2, this tradeoff becomes unfavorable, leading the seller to exclude the second buyer. On the other hand, when θ_1 and θ_2 values are relatively close, the seller finds it optimal to sell to both. This is because $\alpha(\theta_1 + \theta_2)$ will be larger than $\max\{\theta_1, \theta_2\}$, which follows from $\alpha \geq \frac{1}{2}$.

2.3 Inefficiencies

In standard auctions with a single item for sale, asymmetric information can cause inefficiencies in two main ways. First, if agents are heterogeneous or if their virtual valuations are not increasing, it is possible that a bidder 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. Our setup introduces inefficiencies that do not occur in traditional auctions. Note that we assume that all agents draw their types from the same regular distribution F, which eliminates the first inefficiency. In this section, to distinguish our inefficiencies from the ones in conventional auctions, we further assume that virtual values are positive so that exclusion is never optimal in the standard setting.

Definition 1. *Let $k_f(\theta)$ be the size of the first-best optimal group of buyers when private information is θ. The revenue-maximizing mechanism under-(over-) provides if, for all θ:*

$$k^*(\theta) \leq (\geq)\, k_f(\theta).$$

An allocation is efficient if equality holds above.

Define $\lambda(\theta_i) \equiv \frac{f(\theta_i)}{1-F(\theta_i)}\theta_i$.[7]

[7]Which can be interpreted as the price-elasticity of demand. To see this, consider setting a price θ_i

Proposition 2. *Assume $v(\underline{\theta}) \geq 0$. The profit-maximizing mechanism*

- *Is **Efficient** for all values of a if and only if λ is constant — that is, F is in the Pareto family.*
- ***Under-provides** for all values of a if and only if λ is increasing.*
- ***Over-provides** for all values of a if and only 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 are distributed according to the Pareto family. While it is known 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 characterize *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, in contrast with the first-best outcome, the behavior of the principal is no longer dictated by the ratio of valuations θ_i/θ_j. Rather, the slope of the seller's 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 belongs to the Pareto family. To see that, assume $v(\theta) = (1 - \frac{1}{\lambda})\theta$, $\lambda > 0$. We then have:

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

Solving this differential equation yields the unique solution:

$$F(\theta) = 1 - c\theta^{-\lambda},$$

which is the CDF of a distribution in the Pareto family. For any other distribution, the two ratios highlighted above will differ at least for some realizations. We show two such examples in Figure 2. The figure illustrates the profit-maximizing 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,

---

to a buyer who draws his evaluation from $F(\theta)$. The probability that they will purchase—demand—will be $1 - F(\theta_i)$, leading to $\frac{\%\Delta Q}{\%\Delta P} = \frac{\frac{\frac{d(1-F(\theta_i))}{d\theta_i}}{1-F(\theta_i)}}{\frac{d\theta_i}{\frac{d\theta_i}{\theta_i}}} = \frac{f(\theta_i)}{1-F(\theta_i)}\theta_i$

Figure 2: Examples of Under and Overprovision

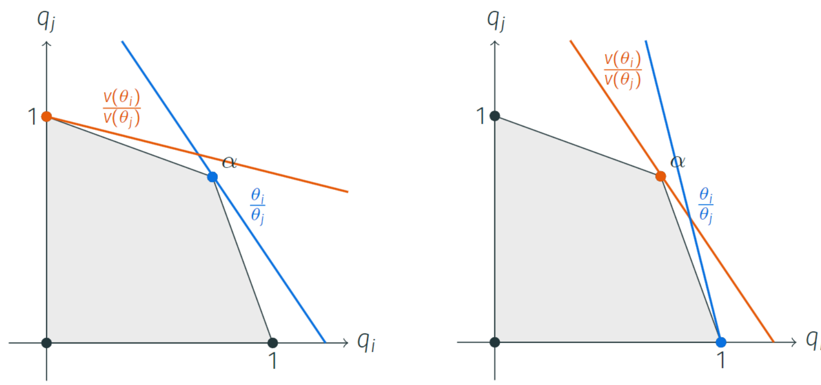

Notes: The figure above displays the profit-maximizing and the first-best allocations for different realized values of θ_i and θ_j. In the left panel, it is efficient to allocate the good to both agents, but it is profit-maximizing to allocate the good to agent i exclusively—*underprovision*. In the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

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 3 presents two such examples.

The figure displays the profit-maximizing and the first-best allocations for different values of θ_i and θ_j. The shaded blue(orange) areas indicate the regions where the good is provided to both agents under the first-best(profit-maximizing) allocation. 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

Figure 3: Examples of Distributions leading to Under and Overprovision

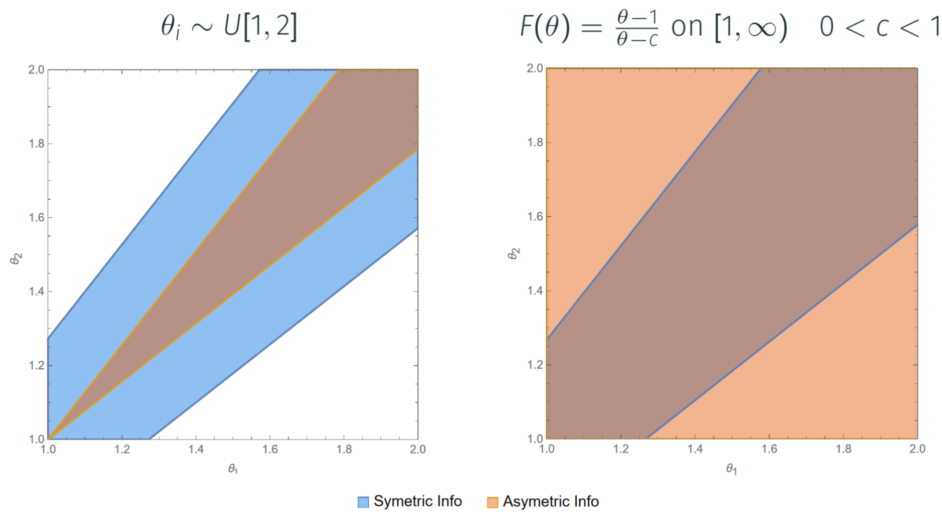

Notes: The figure above illustrate the profit-maximizing and first-best allocations. The shaded blue(orange) areas indicate the regions where the good is provided to both agents under the first-best(profit-maximizing) allocation. The distributions used in each example are shown at the top of the graphs, with $\alpha = 0.56$.

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.

What is additionally interesting is that these inefficiencies are entirely driven by the distribution of buyers' types. Proposition 2 directly links over- or under-provision to the distribution of buyer types without referencing market conduct (the size of externalities, α). Therefore, a policymaker concerned about over- or under-provision in a particular market can draw conclusions about potential inefficiencies simply by understanding the distribution of valuations—effectively by estimating demand.

2.4 Implementation

Next, we turn to the implementation of the optimal mechanism. In particular, we search for a protocol that (i) implements the optimal allocation *truthfully* and in dominant strategies and (ii) does not require payment from excluded agents. We define a class of auctions, which we dub interval auctions, that satisfy these two requirements.

Definition 2. *An **interval auction** has two stages. In the first stage, each potential buyer,* *i, submits a bid b_i. In the second stage, the seller defines a set of thresholds for each potential buyer, that does not depend on buyer i's bid: $\tau^i_N \leq \tau^i_{N-1} \leq \ldots \leq \tau^i_1$, such that the allocation and transfers as functions of buyers' i bid are:*

$$q_i = \alpha_k \mathbb{1}_{b_i \in (\tau_k, \tau_{k-1}]} \qquad\qquad r_i = \mathbb{1}_{b_i \in (\tau_k, \tau_{k-1}]} \sum_{j=k}^{N} (\alpha_k - \alpha_{k+1}) \tau_k$$

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 fact that the mechanism is implemented in dominant strategies plays a key role in making the auction rules fairly straightforward for the bidders. From the bidders' perspective, as long as the thresholds are monotonic and independent of their own bids, they would submit the same bids regardless of how these thresholds were constructed. That is to say, all the principal needs to convey to the potential buyers is that the thresholds she sets will not depend on their private bid. Thus, the principal does not need to burden the bidders with the precise construction of the thresholds, making the protocol straightforward from the bidder's point of view.

Once again, illustrating the implementation for $N = 2$ can help build intuition. When $N = 2$, the interval auction that implements the revenue-maximizing outcome is the following: for each bid b_i there exist thresholds $\underline{\tau}(b_{-i}) < b_{-i} < \overline{\tau}(b_{-i})$ such that

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

where $\underline{\tau}(b_{-i}) = v^{-1}\left(\frac{1-\alpha}{\alpha} v(b_{-i})\right)$, and $\overline{\tau}(b_{-i}) = v^{-1}\left(\frac{\alpha}{1-\alpha} v(b_{-i})\right)$.

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 4 below.

Figure 4: Interval Auction Implementation

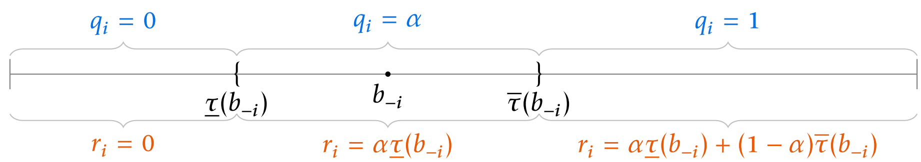

Notes: The figure above visualizes the profit-maximizing implementation via an interval auction. Around the bid of the opponent b_{-i} there is a n neighborhood $(\underline{\tau}(b_{-i}), \overline{\tau}(b_{-i}))$. If the agent's bid falls below this neighborhood, he is excluded and pays nothing $r_i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

In this implementation, for an agent to secure exclusive rights to the good, they must significantly outbid the other agent. Slightly outbidding the other agent 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.[8] Thus, unlike a standard auction, it is no longer the case that only the highest bid matters; instead, the entire distribution of bids determines the optimal allocation and transfers.

2.5 Revenue Comparison

We now focus on the case of $N = 2$ and develop some intuition on how the revenue of our mechanism compares with other allocation protocols. We compare our second-price mechanism with two benchmarks: a posted price that guarantees full-market coverage, and therefore ignores the costly externality; and a second-price auction, which guarantees exclusivity. Because virtual valuations are assumed to be positive, the seller can guarantee full participation by 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

[8]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.

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

Thus, a constrained seller who chooses between these two mechanisms would receive a revenue of

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

Now, consider the seller who chooses the 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. *The difference between the optimal revenue, R, and the revenue constrained to a full-participation posted price or a 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 5: Revenue Comparison

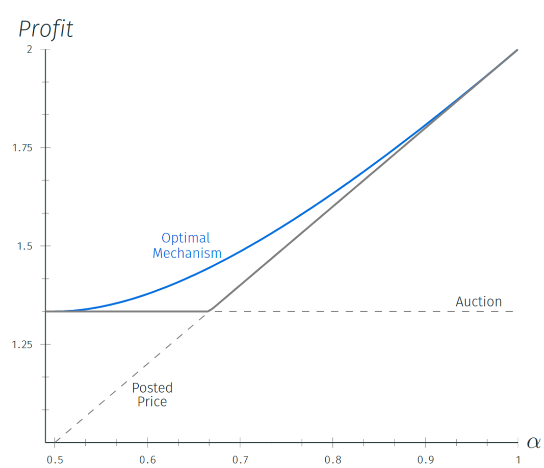

Notes: For different α values, the graph above compares the revenue from a full-participation posted price, a standard auction in which only one good is sold, and the optimal mechanism.

Figure 5 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 from Figure 1 converges to the unit simplex—causing the profits from interval auctions to align with those of a regular auction for a single good. Conversely, as α approaches 1, the externalities from having two active firms diminish, leading profits to align with those from a posted price mechanism described above. Notably, observe that these two seemingly distinct problems—selling a single good through an auction and monopoly pricing—can be seen as specific instances of our framework. In one extreme, externalities are so significant that the seller avoids selling to multiple buyers; in the other, externalities are absent altogether. Our framework includes both cases and also maximizes profits for the intermediate scenarios.

The figure also highlights that when externalities are so severe that the seller would not consider selling to more than one buyer, using a traditional auction results in 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 most gains compared to more traditional alternatives.

3 Dynamic Model

In this section, we focus on the case in which there are only two agents: $N = 2$. We now consider a version of the model in which buyers arrive sequentially, so the seller also decides the timing of license concessions. Time is discrete and runs indefinitely: $t \in \mathbb{N}$. At any time t, with probability μ, 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 i—if the buyer does not arrive, 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 $a^j > t'$, $q^i_{t'} \geq q^i_t$. If $a^j \leq t'$, then $q^i_{t'} \geq \alpha q^i_t$;*

4. ***No money pumps**: $r^i \geq 0$.*

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 is irreversibility. Irreversibility implies that once the designer allocates a license to an agent, she cannot take it back. Therefore, 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. The last restriction makes sure that the seller and the buyer cannot benefit from a lending arrangement. This constraint is relevant when one party is strictly more patient than the other, which would allow for an infinite payoff to be achieved by the most patient party lending money to the least patient party.

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:

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

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 $t < a$,

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

For all $t \geq a$

$$(q^1_t, q^2_t)(\theta_1, \theta_2, a) = \begin{cases} (1, 0), & \text{if } \left(\frac{\delta}{\rho}\right)^a \frac{v(\theta_1)}{v(\theta_2)} \geq \frac{\alpha}{1-\alpha} \\ (\alpha, \alpha), & \text{if } \frac{\alpha}{1-\alpha} \geq \left(\frac{\delta}{\rho}\right)^a \frac{v(\theta_1)}{v(\theta_2)} \geq \frac{1-\alpha}{\alpha} \\ \left(\alpha q^1_0(\theta_1), 1 - (1-\alpha) q^1_0(\theta_1)\right), & \text{otherwise} \end{cases} \tag{3}$$

Figure 6 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 right away, as they will start operating and generating profits without delay. 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 higher this fraction becomes, thereby increasing the likelihood that the principal will either sell the license exclusively to the first buyer or at least include him in the allocation. This is because the first buyer makes

Figure 6: Optimal Dynamic Mechanism

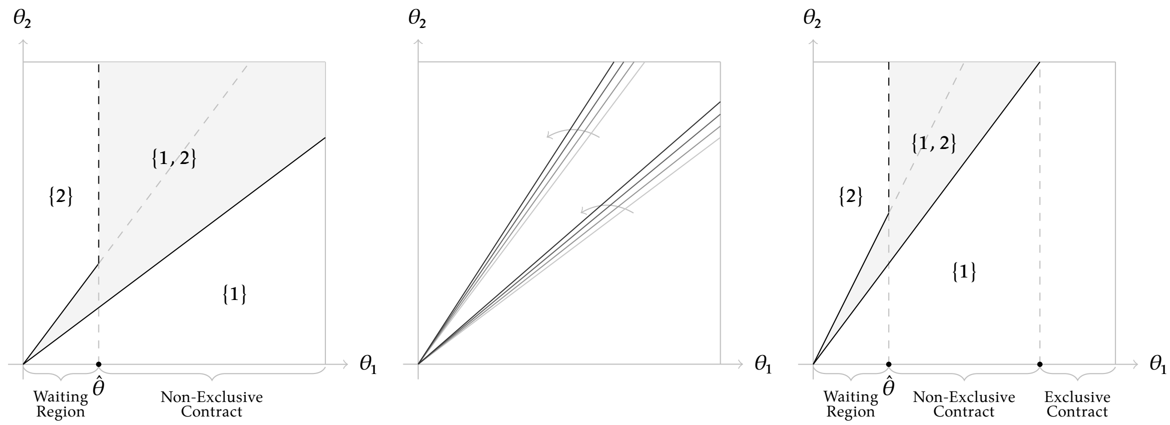

Notes: The figures above shows the revenue-maximizing dynamic mechanism. The left panel shows the $\hat{\theta}$ threshold: agents with types below this threshold are asked to wait, while those above it are immediately issued a contract. In the latter scenario, once the second buyer arrives, depending on their type, θ_2, the principal may choose to allocate to the first agent exclusively or to both agents but no longer exclusively to the second. For the case in which discount rates are not equal, 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.

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.[9] Consequently, the seller finds it optimal to favor the initial buyer more as the arrival time of the second buyer is delayed.[10] 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

[9] To build some intuition, consider this simplified scenario: suppose it is known that the second buyer arrives in period $t = 3$. For every dollar benefit the first buyer expects to receive in that period, he is willing to pay δ^3 immediately. In contrast, the buyer arriving in period $t = 3$ cannot pay immediately; their payment occurs only upon arrival. Consequently, each dollar of payment from this second buyer is worth, for the seller, ρ^3. Thus, in calculating the seller's net present profit, δ discounts the initial buyer's payment, while ρ applies to the second, resulting in the discrepancy.

[10] Note that if the seller is more patient than the buyer ($\rho > \delta$), it would be optimal to postpone any payment as much as possible. With an infinite time horizon, the problem is no longer well-defined. 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 period. Because both buyers would make payments at that period, the seller would not give preferential treatment to one of the buyers. Thus, the optimal cutoffs from the static mechanism would be preserved.

region emerges. In other words, the conditions eventually become so favorable for the first buyer that even if the second buyer arrives and has the highest type $\overline{\theta}$, the principal still allocates the good exclusively to the first buyer. This exclusivity region emerges in finite time for any type of first buyer for whom the principal decides to allocate the license ($q_0^1 = 1$). However, this exclusivity region only materializes if $\delta > \rho$, meaning the seller must be less patient than the buyers, and even then, after a sufficiently long period. 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 $q^i_{f,t}$ represent the first-best allocations. We say that the allocation q_t induced 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. We show the possible inefficiencies in Figure 7. Once again, the distribution of buyers in a market entirely 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. Therefore, a policymaker concerned about over- or under-provision or stringency or leniency in a particular market can draw conclusions about potential inefficiencies simply by understanding the distribution of valuations—effectively by estimating demand.

Figure 7: Dynamic Inefficiencies

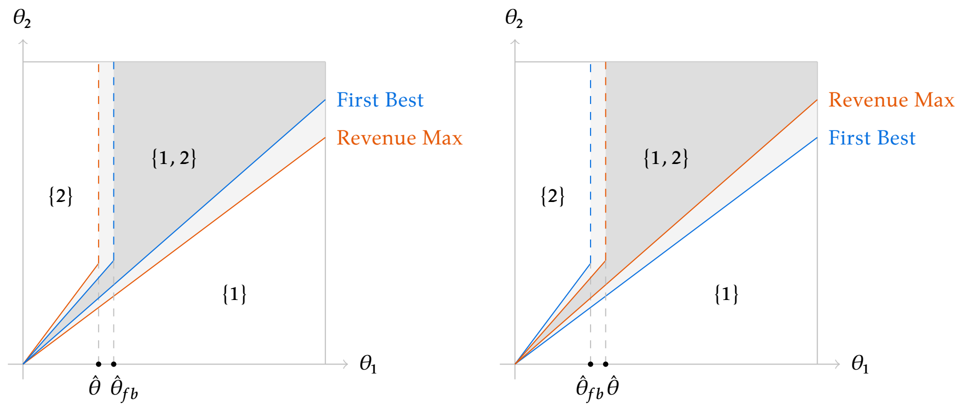

Notes: The left panel shows the stringent, underprovision case, which occurs when λ is increasing. The right panel illustrates the lenient, overprovision case, occurring when λ is decreasing.

4 Model Generalizations

In this section, we generalize the optimal mechanism beyond the simple environment considered so far. While our initial model suffices for certain applications, its limitations are apparent. Notably, we assumed (1) that when transitioning from exclusive control to sharing the market, the buyer's profits were merely scaled by a constant factor, $\alpha < 1$; and (2) that a firm's profits do not depend on the type of its competitor, even when they share the market. In general, firms' profits under competition may not simply be a scaled-down version of their monopoly profits. Furthermore, firms' profits in the presence of competition may depend not only on their own characteristics but also on the characteristics of their competitor. Below, we provide a general model that relaxes both of these assumptions at the cost of additional restrictions on the preferences and on the distribution of types. Throughout this section, we maintain the assumption that there are two buyers $N = 2$.

As in the previous sections, the payoff of buyer i who receives the good alone is θ_i, drawn independently from a distribution F. Buyer i's virtual valuation is then $v(\theta_i) = \theta_i - \frac{1-F(\theta_i)}{f(\theta_i)}$, which is assumed to be strictly increasing. When both buyers are allocated the good, the payoff of buyer i is lower than his monopoly payoff. Moreover, his duopoly payoff increases with his own type and decreases with his competitor's type. Formally, When buyer i, with private information θ_i shares the good with another buyer, with type θ_{-i}, his payoff is given by $\alpha(\theta_i, \theta_{-i}) < \theta_i$. We assume α is differentiable, strictly increasing in the first entry—$\alpha_1 > 0$ —and decreasing in the second entry—$\alpha_2 \leq 0$.

We define an allocation as a triple (q^1, q^2, q_α) such that q^i is the probability that buyer i is allocated the good alone, and q_α is the probability that both buyers receive the good. Naturally, $q^1 + q^2 + q_\alpha \leq 1$. The common feature of the static and dynamic mechanisms we obtained in the previous sections is that the optimal allocations have a specific structure: as a buyer's type increases, the buyer receives more exclusivity. We refer to such allocations as threshold allocations.

Definition 5. *An allocation for agent i is a* threshold allocation *if, for each type θ_{-i}, there exist thresholds, $\underline{\tau} \leq \overline{\tau}$ such that:*

$$q_\alpha(\theta_i, \theta_{-i}) = \mathbb{1}_{\underline{\tau} \leq \theta_i < \overline{\tau}}, \qquad \textit{and} \qquad q^i(\theta_i, \theta_{-i}) = \mathbb{1}_{\theta_i \geq \overline{\tau}}.$$

Below, we provide sufficient conditions for the allocation induced by an optimal mechanism to be a threshold allocation. We define the duopoly virtual value of a bidder with type θ_i to be:

$$v_\alpha(\theta_i, \theta_{-i}) = \alpha(\theta_i, \theta_{-i}) - \alpha_1(\theta_i, \theta_{-i}) \frac{1 - F(\theta_i)}{f(\theta_i)}.$$

Assumption 1. *Preferences and distributions satisfy the following:*

1. ***Increasing differences.*** *The difference between monopoly and duopoly payoffs is increasing in own-type: $1 \geq \alpha_1(\theta_i, \theta_{-i})$;*

2. ***Strong Regulatiry.*** *$v_\alpha(\cdot, \theta_{-i})$ is increasing for all $\theta_{-i} \in \Theta$*

3. ***Virtual Gains.*** *$v'(\theta_i) \geq v_{\alpha,1}(\theta_i, \theta_{-i}) + v_{\alpha,2}(\theta_{-i}, \theta_i) \geq 0$*

The first assumption establishes that monopoly payoffs grow faster with a buyer's type than his duopoly payoffs. The second assumption generalizes the usual regularity of virtual values to settings with correlated values. It is in line with an assumption found in Bulow and Klemperer (1996). Finally, the third condition specifies that, in virtual space, as the valuation of a buyer increases, his monopoly virtual value grows faster than the joint virtual value when both agents share the good. At the same time, this duopoly joint virtual value also grows with an agent type. The next result shows that these assumptions are sufficient to generalize our results.

Proposition 7. *Under Assumption 1, the revenue-maximizing mechanism implements threshold allocations given by:*

$$q^i(\theta_i, \theta_{-i}) = \begin{cases} 1 & \textit{if } v(\theta_i) > \max\{v_\alpha(\theta_i, \theta_{-i}) + v_\alpha(\theta_{-i}, \theta_i), v(\theta_{-i})\} \\ 0 & \textit{otherwise} \end{cases},$$

and

$$q_\alpha(\theta_i, \theta_{-i}) = \begin{cases} 1 & \textit{if } \max\{v(\theta_i), v(\theta_{-i})\} < v_\alpha(\theta_i, \theta_{-i}) + v_\alpha(\theta_{-i}, \theta_i) \\ 0 & \textit{otherwise.} \end{cases}$$

Proposition 7 completes the characterization for the general model. In some environments, the full generality of the model above may not be necesairy. For example, it may be that the payoff of a buyer when sharing the good does not depend on his competitor's type, yet this payoff is not a fraction of his monopoly profits. In the table below, we provide sufficient conditions for a mechanism analogous to the one in Proposition 7 to be optimal, provided that the returns from exclusivity are increasing in the buyer's type—we call this the supermodular model. We also provide conditions for the case in which duopoly prices are a scaled version of monopoly profits, but the scale may depend on the competitor's type—which we call the multiplicative model. Table 1 below summarizes the sufficient conditions required for these alternative models, which follow from Proposition 7. In Section 5.2, we explore an application that is feasible under this general framework 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

Table 1: Sufficient Conditions

	Win Alone	Win Together	Sufficient Conditions
α **Model**	θ_i	$\alpha\theta_i$	$\varnothing$
A buyer's profit depends on own type only.			
Supermodular	θ_i	$\alpha(\theta_i)$	$\max\left\{\frac{-\alpha''(\theta_i)}{1-\alpha'(\theta_i)}, \frac{\alpha''(\theta_i)}{\alpha'(\theta_i)}\right\}\frac{1-F(\theta_i)}{f(\theta_i)} \leq v'(\theta_i)$ $\alpha'(\theta_i) \leq 1$
A buyer's profit depends on both types.			
Multiplicative	θ_i	$\alpha(\theta_j)\theta_i$	$\frac{v'(\theta_i)}{v(\theta_i)} \geq -\frac{\alpha'(\theta_i)}{\alpha(\theta_i)}$ $\frac{1}{2} \leq \alpha(\theta_i) \leq 1$
General	θ_i	$\alpha(\theta_i, \theta_j)$	$v'(\theta_i) \geq v_{\alpha,1}(\theta_i, \theta_j) + v_{\alpha,2}(\theta_j, \theta_i) \geq 0$ $\alpha_1(\theta_i, \theta_j) \leq 1$

Notes: The table reports the additional assumptions ensuring that the optimal mechanism allocates to the group of buyers with maximal joint virtual values. The optimal mechanism also implements threshold allocations. These assumptions are in addition to increasing virtual values: $v'(\theta_i) > 0$.

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 the proportion of informed investors is $\eta = \frac{2}{N}$. We solve for the equilibrium in the financial market given η.

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:

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

Thus, if there is only one informed investor, her payoff is

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

On the other hand, if there are two informed investors, their payoffs will be

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

If we renormalize investor types to be $\tilde{\theta} = \frac{1}{2}\frac{N-2}{N-1}\theta$, the above payoffs fit into our baseline model, with $\alpha = \frac{N-1}{N}$. The revenue-maximizing allocation of information follows Proposition 1.

5.2 Horizontally Differentiated Products

Consider a uniform distribution of consumers on 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. The principal can issue an exclusive license to the firm located at 0, an exclusive license to the firm located at 1, or issue licenses to both of them. Each firm has private information about 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 firm j, 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 8: Hotelling Application

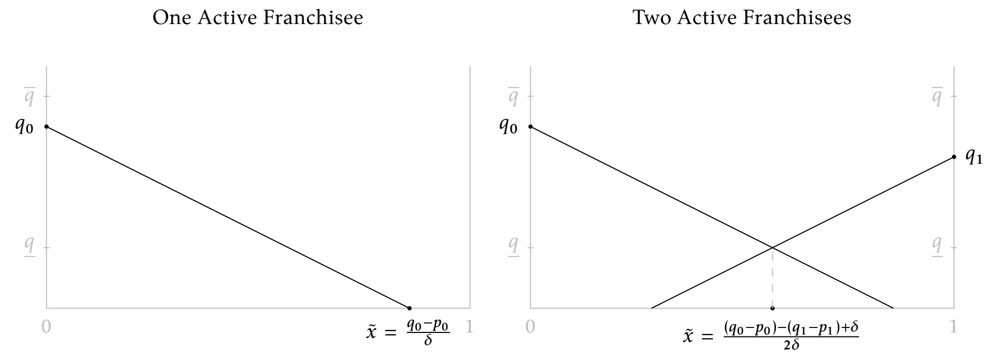

Notes: The left panel shows the case where an exclusive franchise license is granted to the firm at position 0, along with the corresponding marginal consumer. The right panel displays the case where franchise licenses are issued to both sellers, along with the associated marginal consumer.

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 who is indifferent between traveling and purchasing the good or staying home and receiving 0 utility. 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 $p_j^M = \frac{q_j}{2}$, with M representing their monopolistic status. The marginal consumer will thus be $\tilde{x}(p_j^M) = \frac{q_j}{2\delta}$, while the firms profits will be $\pi_j^M = \frac{q_j^2}{4\delta}$.

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

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

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

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

And duopoly profits of

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

Importantly, note that the duopoly profits are not simply a fraction α of the monopoly profits, nor can they be expressed in a multiplicative form as a function of the competitors type q_{-j}. Thus, the machinery developed in Section 4 is necessary to handle this example. It is straightforward to verify that, with the appropriate $\underline{q}$, $\overline{q}$, 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 examines the optimal licensing strategy for a seller dealing with downstream competitors who hold private information, applicable in numerous market contexts such as franchise operations, patent licensing, and information sales, to name a few. We characterize conditions under which inefficiencies—absent in conventional auctions—arise due to asymmetric information, leading the seller to either over- or under-supply the good. We link these inefficiencies to the distribution of buyer valuations, emphasizing that a policymaker only needs to estimate demand to assess whether over- or under-provision may occur. We propose an *interval auction* as a dominant strategy implementation of the revenue-maximizing mechanism, where the allocation decision depends not only on the highest bid but on the overall bid distribution. When bids are closely clustered, the mechanism favors selling to multiple bidders; when bids are widely dispersed, exclusive licensing to the highest bidder is optimal.

In a dynamic setting where buyers arrive sequentially, we analyze the timing of licensing decisions, demonstrating that a seller may choose to delay licensing or issue it immediately to an available buyer. We show that the decision to offer a promise of exclusivity depends on the seller's relative patience compared to the buyers; a promise of exclusivity is issued only if the seller is less patient and enough time has passed since the arrival of the first buyer. We also uncover additional inefficiencies in the dynamic setup related to the seller's inclination to license to the first arriving buyer. We once again link these inefficiencies to the same conditions identified in the static model.

Lastly, we explore sufficient conditions that allow for more general profit functions and a framework where buyer valuations are interdependent.

7 Appendix

7.1 Proofs

Proof of Proposition 1

The problem of the principal is:

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

We follow the standard approach to solve a relaxed problem by ignoring condition 1 (monotonicity), so we maximize the seller's problem conditional to feasibility alone. We then check that the solution to the relaxed problem does satisfy monotonicity. We first transform the problem back to allocation space:

$$\max_{\sigma \in \Delta \mathcal{J}} \quad \int \sum_\ell \sigma_\ell \left(\sum_i v(\theta_i) \alpha_{|\ell|} \mathbb{1}_{i \in \mathcal{J}_\ell} \right) \sigma_\ell$$

For each θ, the choice of the distribution σ_ℓ must satisfy:

$$\operatorname{supp} \sigma \in \arg\max_\ell \sum_i v(\theta_i) \alpha_{|\ell|} \mathbb{1}_{i \in \mathcal{J}_\ell}.$$

Because v is monotonic, by ordering the types we obtain the result in the proposition. Moreover, monotonicity is guaranteed because as θ_i grows, they can only be included in more exclusive groups, since $\alpha_|\ell| > \alpha_|\ell'|$ for $|\ell| < |\ell'|$. ■

Proof of Proposition 2

Assume $\lambda(x)$ increasing. We want to prove that the revenue-maximizing mechanism does not over-provide the good. To see that, fix an arbitrary realization of types and assume, without loss of generality, $\theta_1 \geq \theta_2 \geq \ldots \geq \theta_N$. For that realization, let the first-best number of buyers served be $k_f(\theta) = k$. Then, for every $k' > k$ we have:

$$\alpha_k \sum_{i=1}^{k} \theta_i \geq \alpha_{k'} \sum_{i=1}^{k'} \theta_i,$$

which can be reordered to obtain:

$$\frac{\alpha_{k'}}{\alpha_k - \alpha_{k'}} \leq \sum_{i=1}^{k} \frac{\theta_i}{\sum_{j=k+1}^{k'} \theta_j}.$$

Thus:

$$\begin{aligned}\frac{\alpha_{k'}}{\alpha_k - \alpha_{k'}} &\leq \sum_{i=1}^{k} \frac{\theta_i}{\sum_{j=k+1}^{k'} \theta_j} \leq \sum_{i=1}^{k} \frac{\left(1 - \frac{1}{\lambda(\theta_i)}\right)}{\left(1 - \frac{1}{\lambda(\theta_k)}\right)} \frac{\theta_i}{\sum_{j=k+1}^{k'} \theta_j} \\ &\leq \sum_{i=1}^{k} \frac{\left(1 - \frac{1}{\lambda(\theta_i)}\right)\theta_i}{\sum_{j=k+1}^{k'} \left(1 - \frac{1}{\lambda(\theta_j)}\right)\theta_j} = \sum_{i=1}^{k} \frac{v(\theta_i)}{\sum_{j=k+1}^{k'} v(\theta_j)},\end{aligned}$$

where the first ineequality follows from optimality of k (over k'). The second and third inequalities follow from the fact that $1 - \frac{1}{\lambda(\theta_i)} \geq 1 - \frac{1}{\lambda(\theta_j)}$ for all $j \geq i$ because λ is increasing. By rewriting the implied inequality:

$$\alpha_k \sum_{i=1}^{k} v(\theta_i) \geq \alpha_{k'} \sum_{i=1}^{k'} v(\theta_i),$$

so, under asymmetric information, the designer prefers to allocate to k over any larger number k' of buyers. Thus, the number of allocated buyers is at-most efficient.

For the converse, assume that the optimal mechanism underprovides for all α. Consider $\alpha_1 = 1$, $\alpha_2 = \alpha$ and $\alpha_3 = ... = \alpha_N = 0$, so at most two agents are served. Let $\theta_1 > \theta_2$ be the (arbitrary) highest valuations. Choose α 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 that the optimal mechanism underprovides for all α. We thus have $\lambda(\theta_1) \geq \lambda(\theta_2)$. Because $\theta_1 > \theta_2$ were arbitrary, the result follows. A symmetric argument works for λ decreasing.

By definition, the optimal mechanism is efficient if and only if it over- and underprovides, so λ is a constant. That is, for all x:

$$\lambda(x) = \frac{f(x)}{1 - F(x)} = \lambda,$$

for some real λ. Solving this differential equation generates the unique solution:

$$F(x) = 1 + c\theta^{-\lambda},$$

which is a member of the Pareto family. ■

Proof of Proposition 3

Interval Structure We first prove the following: for any bidder i, and for a fixed realization of other bidders, there exists an interval structure $\tau_1 < \tau_2 < ... < \tau_n$, with associated natural numbers $m_1 = N + 1 > m_2 > ... > m_n$, such that if $\theta_i \in [\tau_\ell, \tau_{\ell+1})$, $\ell > 1$, bidder i is allocated the good in a group with m_ℓ other buyers. If $\theta_i \in [\tau_1, \tau_2)$, then bidder i is excluded. We assume without loss of generality that buyer i is indexed as N, and that the other bidders' types are ordered in decreasing order: $\theta_1 > \theta_2 > ... > \theta_{N-1}$.

For that, notice that for any group containing k individuals including buyer N, the seller's revenue for that group, is a linear function of $v_n \equiv v(\theta_N)$. Formally:

$$g_k(v_N) = \alpha_k \sum_{j=1}^{k-1} v(\theta_j) + \alpha_k v_N.$$

If the group contains k individuals but excludes buyer N, then the revenue does not depend on v_N:

$$g_{-k}(v_N) = \alpha_k \sum_{j=1}^{k} v(\theta_j).$$

Because each g_{-k} does not depend on θ_N, we can define $g_{N+1} \equiv \max_k g_{-k}(v_N)$ as the maximal revenue obtained by selling to a group without agent N.

The revenue of the seller as a function of v_N can be calculated as:

$$g(v_N) = \max_{k \in \{1,...,N+1\}} \{g_k(v_N), g_{-k}(v_N)\}.$$

As a maximum of finite linear functions, g is linear by parts, convex with finite non-differentiable points. Order the points of non-differentiability of g: $\hat{\tau}_1 < \hat{\tau}_2 < ... < \hat{\tau}_n$. For any point $\hat{\tau}_i$ such that g is non-differentiable, define $m_i = \min\{k : g(\hat{\tau}_i) = g_k(\hat{\tau}_i)\}$. We argue that $m_i > m_{i+1}$ for any i. Indeed, notice that g is convex, so the slope of g at $\hat{\tau}_i$ is lower than the slope of g at $\hat{\tau}_{i+1}$. Because $g(\hat{\tau}_i) = g_{m_i}(\hat{\tau}_i)$ has slope α_{m_i}, and $g(\hat{\tau}_{i+1}) = g_{m_{i+1}}(\hat{\tau}_{i+1})$ has slope $\alpha_{m_{i+1}}$, we conclude $\alpha_{m_{i+1}} > \alpha_{m_i}$—with the interpretation that $\alpha_{N+1} = 0$. This implies $m_{i+1} < m_i$.

Thus, if $v_N \in [\hat{\tau}_i, \hat{\tau}_{i+1})$, the seller maximizes revenue by choosing a group of size m_i that includes buyer N, if $m_i < N+1$, or excludes the buyer N if $m_i = N+1$. Because v is monotonic, we can define $\tau_i = v^{-1}(\hat{\tau}_i)$ without changing the optimality of the seller's choice. Moreover, $m_1 > m_2 > ... > m_n$.

Implementation For any possible realization of the other N+1 agents' types, the procedure above provides us with a set of thresholds for agent i, $\tau_1 < \tau_2 < ... < \tau_n$, with associated group sizes $m_1 = N+1 > m_2 > ... > m_n$. For all bids b_i such that $\tau_k < b_i \leq \tau_{k+1}$ define:

$$r_k = \sum_{j=2}^{k} \left(\alpha_{m_j} - \alpha_{m_{j-1}}\right)\tau_j.$$

Let's show that the buyer has no incentive to deviate. Consider first a buyer with $\theta \in [\tau_\ell, \tau_{\ell+1}]$ who considers deviating to a group of size $m_k > m_\ell$ (that is, $k < \ell$)—notice that the deviation for $m_k = m_1 = N+1$ already includes the case of deviating to exclusion. Then the difference between the payoff of truthful revelation and of the deviation is:

$$\alpha_{m_\ell}\theta - \sum_{j=2}^{\ell}(\alpha_{m_j} - \alpha_{m_{j-1}})\tau_j - \left(\alpha_{m_k}\theta - \sum_{j=2}^{k}(\alpha_{m_j} - \alpha_{m_{j-1}})\tau_j\right) = \left(\alpha_{m_\ell} - \alpha_{m_k}\right)\theta - \sum_{j=k+1}^{\ell}(\alpha_{m_j} - \alpha_{m_{j-1}})\tau_j$$
$$\geq \left(\alpha_{m_\ell} - \alpha_{m_k}\right)\tau_\ell - \sum_{j=k+1}^{\ell}(\alpha_{m_j} - \alpha_{m_{j-1}})\tau_\ell = 0,$$

where the inequality follows from $\theta \in [\tau_\ell, \tau_{\ell+1}]$, $\alpha_{m_\ell} > \alpha_{m_k}$, and monotonicity of the thresholds, and the last equality follows from the telescopic summation. Thus, no such deviation would benefit the agent. Consider now $m_k < m_\ell$—that is, $k > \ell$. Then:

$$\alpha_{m_\ell}\theta - \sum_{j=2}^{\ell}(\alpha_{m_j} - \alpha_{m_{j-1}})\tau_j - \left(\alpha_{m_k}\theta - \sum_{j=2}^{k}(\alpha_{m_j} - \alpha_{m_{j-1}})\tau_j\right) = \left(\alpha_{m_\ell} - \alpha_{m_k}\right)\theta + \sum_{j=\ell+1}^{k}(\alpha_{m_j} - \alpha_{m_{j-1}})\tau_j$$

$$\geq \left(\alpha_{m_\ell} - \alpha_{m_k}\right)\tau_{\ell+1} - \sum_{j=k+1}^{\ell}(\alpha_{m_j} - \alpha_{m_{j-1}})\tau_{\ell+1} = 0,$$

confirming that the agent has no incentive to deviate. Because this payment scheme implements the optimal allocation and the agent with the lowest type receives zero rents, revenue-equivalence implies that this scheme implements the optimal mechanism.

The last step is to notice that we can relabel the thresholds to reflect the number of agents in their groups to obtain, for a certain subset of $\mathcal{N}$, $\tau_{n_1} < \tau_{n_2} < \ldots$, with $n_{j-1} > n_j$. Finally, if there is a natural number $n_{j-1} > n > n_j$, we define $\tau_n = \tau_{n_{j-1}}$. ■

Proof of Proposition 5

Define x^1 and x^2 as follows

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

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

The revenue of the seller is:

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

Using integration by parts, the seller maximizes:

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

$$=\mathbb{E}\sum_{j=0}^{\infty}\delta^j(1-\mu)^j q_j^1(\theta_i)v(\theta_1)+\sum_{j=0}^{\infty}\mu(1-\mu)^j\sum_{k=j+1}\delta^k q_k^1(\theta_i,a^2=j+1)v(\theta_1)+$$

$$\sum_{j=0}^{\infty}\rho^{j+1}\mu(1-\mu)^j\sum_{k=j+1}^{\infty}\frac{\delta^k}{\delta^{j+1}}q_k^2(\theta,a^2=j+1)v(\theta_2)$$

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:

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

When the irreversibility constraint does not bind, we have:

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

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

and when it binds:

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

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:

$$\frac{1}{1-\delta(1-\mu)}qv(\theta_1)+\sum_{j=0}^{\infty}\frac{\mu}{1-\delta}(1-\mu)^j$$

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

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

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

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

Start by noticing that, in the optimal mechanism, for a fixed θ_1, the designer waits if and only if:

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

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

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

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 Proposition 7

Define $\gamma(\theta_i, \theta_{-i}) = (1, \alpha(\theta_i, \theta_{-i}))$. Recall that an allocation is a triple of functions $\{q^i\}_{i=1,2}$, $q_\alpha : \Theta \times \Theta \to [0, 1]$, such that for each realization $\theta_1, \theta_2 \in \Theta$:

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

Define $q_i = (q^i, q_\alpha)$. In a truthfully revealing direct mechanism, the expected utility of agent i with type θ_i is

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

Following the usual approach for quasilinear mechanism design, we can then write the Bayesian incentive compatibility constraints as

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

for all θ_i, θ_i'. As usual, we say that an allocation is implementable if it satisfies Bayesian Incentive Constraints. The following lemma establishes necessary conditions for implementability.

Lemma 3. *An allocation* $\{q^1, q^2, q_\alpha\}$ *is implementable only if:*

(E) $U_i(\theta_i) = U_i(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} \mathbb{E}[\frac{\partial \gamma}{\partial \nu}(\nu, \theta_{-i}) \cdot q_i(\nu, \theta_{-i})] d\nu$ *for all* $\theta_i \in \Theta$.

Proof of Lemma 3

This proof closely follows Lemma 3 in Bulow and Klemperer (1996). By switching the order of θ_i and θ_i' in the BIC inequality above and putting the two together we obtain:

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

Divide all three terms by $\theta - \theta'$ and take the limit as $\theta' \to \theta$ to obtain condition (E).
■

The lemma shows that (E) is a necessary condition for implementability. We show now that, given the Increasing differences assumption, monotonicity of q^i and $q^i + q_\alpha$ in θ_i makes condition (E) sufficient for implementability.

Lemma 4. *When* q^i *and* $q^i + q_\alpha$ *are increasing, condition (E) in Lemma 3 is sufficient for implementability.*

Proof of Lemma 4

We assume q^i and $q^i + q_\alpha$ are increasing and prove that IC is satisfied. Start by writing $\overline{\gamma} = (\beta - \alpha, \alpha)$, and $\overline{q}_i = q^i, q^i + q_\alpha$. Assume first $\theta_i > \theta_i'$. Then:

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

where the first equality comes from condition (E), 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 the increasing differences assumption, both entries of $\overline{\gamma}'$ are positive and both entries of $\overline{q}_i$ are monotonic. The symmetric argument holds for $\theta' > \theta$, so we proved that BIC is satisfied. ■

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

$$\sum_i \int_\theta \mathbb{E}_{-i}\left[\left(\gamma(\theta_i, \theta_{-i}) - \frac{1-F(\theta_i)}{f(\theta_i)} \frac{\partial \gamma}{\partial \theta_i}(\theta_i, \theta_{-i})\right) \cdot q_i(\theta_i, \theta_{-i})\right] f(\theta_i) d\theta_i. \tag{4}$$

We solve the relaxed problem of maximizing profits subject to 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_i for any θ_{-i}. Fix θ_i. Let T be the set of opponent's types, θ_{-i}, such that $q^i(\theta_i, \theta_{-i}) = 1$. That is,

$$T = \{\theta_{-i} : v(\theta_i) \geq \max\{v_\alpha(\theta_{-i}, \theta_i) + v_\alpha(\theta_i, \theta_{-i}), v(\theta_{-i})\}\}.$$

Let x be defined by:

$$v(\theta_i) = v_\alpha(\theta_i, x) + v_\alpha(x, \theta_i),$$

and define $\hat{x} = \min\{\theta_i, x\}$. We now prove $T = [\underline{\theta}, \hat{x}]$. Indeed, let $\theta_{-i} \leq \hat{x}$. By the second inequality in the virtual gains assumption, $h(z) \equiv v(\theta_i) - v_\alpha(\theta_i, z) - v_\alpha(z, \theta_i)$ is a decreasing function of z, so $\theta_{-i} \leq x$ implies

$$v(\theta_i) \geq v_\alpha(\theta_i, x) + v_\alpha(x, \theta_i).$$

Moreover, if $\theta_{-i} \leq \theta_i$, then $v(\theta_i) \geq v(\theta_{-i})$. Therefore, $\theta_{-i} \leq \hat{x}$ implies $\theta_{-i} \in T$. The converse follows from the same argument.

We now show prove that $\hat{x}$ increases in θ_i. For that, it is sufficient to show that x increases in θ_i. Indeed, by total differentiation:

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

Therefore, x is increasing. Thus, if $q^i(\theta_i, \theta_{-i}) = 1$, and $\theta_i' > \theta_i$, then $q^i(\theta_i', \theta_{-i}) = 1$, and we proved that q^i is monotonic in θ_i.

We now show that $q_\alpha + q_i$ is increasing. Again fix any θ_i, and let T be the set of θ_{-i}

such that $q_\alpha + q_i = 1$, that is:

$$T = \{\theta_{-i} : \max\{v(\theta_i), v_\alpha(\theta_{-i}, \theta_i) + v_\alpha(\theta_i, \theta_{-i})\} \geq v(\theta_{-i})\}.$$

Let y be such that:

$$v(y) = v_\alpha(\theta_i, y) + v_\alpha(y, \theta_i),$$

and define $\hat{y} = \max\{\theta_i, y\}$. We now prove $T = [\underline{\theta}, \hat{y}]$. Indeed, let $\theta_{-i} \leq \hat{y}$. Then, if $\hat{y} = \theta_i$, we have $v(\theta_{-i}) \leq v(\theta_i)$, and thus $q^i + q_\alpha = 1$. If $\hat{y} = y$, then notice that the function $h(z) \equiv v(z) - v_\alpha(\theta_i, z) - v_\alpha(z, \theta_i)$ is increasing in z by the first inequality in the virtual gains assumption. Thus, $\theta_{-i} \leq y$ implies that $v(\theta_{-i}) \leq v_\alpha(\theta_i, \theta_{-i}) + v_\alpha(y, \theta_{-i})$. Therefore, $\theta_{-i} \in T$. The converse result follows from the same logic.

We now prove that $\hat{y}$ is increasing in θ_i. For that, it is enough to show that y is increasing. By total differentiation of the equation that defines y:

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

Again, the threshold y grows. So if $q_i(\theta_i, \theta_{-i}) + q_\alpha(\theta_i, \theta_{-i}) = 1$, the same holds for $\theta_i' > \theta_i$, which guarantees that $q^i + q_\alpha$ is increasing in θ_i. We have now proved q_i and $q_i + q_\alpha$ are increasing in θ_i, and we are thus in the conditions of Lemma 4. ■

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.

Dana Jr, J. D. and Spier, K. E. (1994). Designing a private industry: Government auctions with endogenous market structure. *Journal of public Economics*, 53(1):127–147.

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.