

Optimal Revenue-Sharing Double Auctions with Applications to Ad Exchanges

Renato Gomes
Toulouse School of Economics
renato.gomes@tse-fr.eu

Vahab Mirrokni
Google Research, New York
mirrokni@google.com

ABSTRACT

E-commerce web-sites such as Ebay as well as advertising exchanges (AdX) such as DoubleClick's, RightMedia, or AdECN work as intermediaries who sell items (e.g. page-views) on behalf of a seller (e.g. a publisher) to buyers on the opposite side of the market (e.g., advertisers). These platforms often use fixed-percentage sharing schemes, according to which (i) the platform runs an auction amongst buyers, and (ii) gives the seller a constant-fraction (e.g., 80%) of the auction proceeds. In these settings, the platform faces asymmetric information regarding both the valuations of buyers for the item (as in a standard auction environment) as well as about the seller's opportunity cost of selling the item. Moreover, platforms often face intense competition from similar market places, and such competition is likely to favor auction rules that secure high payoffs to sellers. In such an environment, what selling mechanism should platforms employ? Our goal in this paper is to study optimal mechanism design in settings plagued by competition and two-sided asymmetric information, and identify conditions under which the current practice of employing constant cuts is indeed optimal.

In particular, we first show that for a large class of competition games, platforms behave in equilibrium as if they maximize a convex combination of seller's payoffs and platform's revenue, with weight α on the seller's payoffs (which is proxy for the intensity of competition in the market). We generalize the analysis of Myerson and Satterthwaite (1983), and derive the optimal direct-revelation mechanism for each α . As expected, the optimal mechanism applies a reserve price which is decreasing in α . Next, we present an indirect implementation based on "sharing schemes". We show that constant cuts are optimal if and only if the opportunity cost of the seller has a power-form distribution, and derive a simple formula for computing the optimal constant cut as a function of the sellers' distribution of opportunity costs, and the market competition proxy α . Finally, for completeness, we study the case of a seller's optimal auction with a fixed

profit for the platform, and derive the optimal direct and indirect implementations in this setting.

Categories and Subject Descriptors

H.m [Information Systems]: Miscellaneous; F.m [Theory of Computation]: Miscellaneous

1. INTRODUCTION

E-commerce web-sites such as Ebay as well as ad exchanges (AdX) such as Yahoo's RightMedia, Microsoft's AdECN, or Google's DoubleClick Ad Exchange sell items (e.g. page-views) on behalf of a seller (e.g. a publisher) by running an auction amongst a certain number of buyers (e.g., advertisers). In most cases, the e-commerce site or AdX platform then follow a fixed-percentage sharing scheme, according to which the seller is given a constant cut from the auction revenue, and the platform keeps the remaining proceeds from the auction. One typical mechanism employed by ad exchanges works as follows: (i) the publisher declares a reserve price p to be used in the auction, (ii) AdX runs a second-price auction with a reserve price $1.25p$, and collects 20% of the final revenue from the auction. The publisher, in turn, gets the remaining 80% of the revenue from the auction, making sure that it gets at least its declared reserve price of p .

In these settings, the platform faces asymmetric information regarding both the valuations of buyers for the item (as in a standard auction environment) as well as about the seller's opportunity cost of selling the good. For example, in the case of ad exchanges, the opportunity cost for the publisher (seller) is the profit that the publisher could obtain by sending the page-view to competing ad exchanges or by using the page-view to fulfill some guaranteed contract previously signed with advertisers. In such an environment, the optimal direct-revelation mechanism (following the Nobel-winning work by Myerson) is such that payments to the seller depend on a rather complex way on the distribution of valuations by buyers, and the distribution of the opportunity cost of the seller. This contrasts with the current practice of most Internet platforms, which adopt a simple fixed percentage or fixed fee models of revenue share. Moreover, in practice, platforms often face intense competition from similar market places. This is specially true in the case of advertising exchanges that play in a highly competitive market place with many other ad networks and ad exchanges. Such issues in these practical settings raise the following questions: what selling mechanism should platforms employ? In particular, is such a mechanism based on

a fixed-percentage sharing scheme? If so, what is the optimal fixed-percentage sharing scheme? Finally, how are sharing schemes affected by the degree of market competition? Our goal in this paper is to study optimal mechanism design in settings plagued by competition and two-sided asymmetric information, and identify conditions under which the current practice of employing constant cuts is indeed optimal.

1.1 Our Contributions.

Equilibrium Strategy for Competing Ad Exchanges. Competition among ad exchanges is likely to favor auction rules that secure high payoffs to sellers, so that sellers continuously trade through the ad exchange. Intuitively, competition has the effect of making ad exchanges internalize (at least partially) the sellers' payoffs in their objective functions. In this paper, we first show that for a large class of competition games between intermediaries, the equilibrium best responses of each intermediary maximizes a convex combination of seller's payoffs and intermediary's revenue. The weight assigned in the seller's profit, denoted by α , is a proxy for the degree of competition in the market - and depends on the primitives of the competition game between platforms. We refer to this maximization problem as the α -optimal problem.

Characterizing α -Optimal Mechanisms. As a first step toward understanding the optimal double auction in the presence of competing exchanges, we derive the optimal direct-revelation mechanism associated to each α -optimal problem. As expected, the α -optimal mechanism allocates the good to the buyer with the highest valuation if and only if the highest valuation is greater than some reserve price. This reserve price is decreasing in α : as the weight on the seller profits increase, the platform withholds the good less often, and makes greater transfers to the seller. In particular, when $\alpha \geq \frac{1}{2}$, the α -optimal mechanism coincides with the mechanism that the seller would run if he owned the platform (as in Myerson (1981)). To understand this result, notice that the platform's objective function puts weakly higher weight on the seller's payoff than on its own profits. Because of the seller's incentive constraints, by giving up one dollar in profits, the platform can add more than one dollar to the seller's payoff. As a consequence, the platform maximizes its objective by following the seller's optimal mechanism (that is, the Myerson auction).

On the other hand, when $\alpha < \frac{1}{2}$, the platform puts strictly more weight on its profits than on the seller's payoffs. As a consequence, it is willing to introduce distortions on the seller side of the market (therefore reducing the seller's payoffs) in order to generate higher profits. The allocative distortions generated by the seller's informational rents are proportional to

$$\frac{1 - 2 \cdot \alpha}{1 - \alpha},$$

which is decreasing in α . As expected, when $\alpha = 0$, the α -optimal mechanism coincides with the platform profit-maximizing mechanism derived in [8].

As a second step toward understanding α -optimal mechanisms, we study how to indirectly implement the α -optimal direct-revelation mechanism by means of "sharing schemes". Such sharing schemes work in the following manner: (i) the seller submits a reserve price r to the platform, (ii) the platform runs a second-price auction among sellers with reserve

price r , (iii) the platform gives a fraction $\theta^\alpha(r)$ of the auction proceeds to the seller, and keeps the rest. Such sharing schemes (in which the fraction $\theta^\alpha(r)$ is a function of the reported reserve price) implement the α -optimal direct-revelation mechanism for any distributions of the seller opportunity costs and the buyers' valuations that satisfy mild regularity conditions.

Constant Cuts. As discussed earlier, in practice, platforms use constant sharing schemes, in which the fraction (or cut) assigned to the seller does not depend on his reported reserve price. Next, we turn to answering the following question: Under what conditions, can a constant sharing scheme implement the α -optimal direct-revelation mechanism?

Interestingly, we show that constant cuts are optimal if and only if the distribution of seller's opportunity cost, denoted by G , has the power form

$$G(v_0) = \left(\frac{v_0}{b}\right)^k,$$

where $k > 0$ and b is the superior limit of the support of the seller's opportunity cost distribution.

The result above can be understood in the light of monopolistic price theory. Intuitively, power-form distribution functions have a constant price elasticity of supply (which measures how many more percentage points of inventory sellers are willing to offer for one percentage increase in expected revenue). Namely, a power-form distribution with parameter k has a price elasticity of supply equal to k for all opportunity costs in the support. As it turns out, constant sharing schemes are optimal provided that the price elasticity of supply is constant.

In particular, we derive an easy-to-implement formula that relates the price elasticity of supply k and the degree of competition α to the constant seller's share:

$$\theta^{\alpha,k} = \frac{k}{k + \frac{1-2\alpha}{1-\alpha}}.$$

Data Analysis. As the price elasticity k increases, the seller's revenue share goes up. Intuitively, as the distribution of the seller's opportunity costs become concentrated on high values, the platform has to increase the seller's revenue share to make sure that the seller is willing to participate in the trading mechanism with high enough probability. In turn, for fixed k , the seller revenue share increases as the degree of competition α increases. Figure 1 represents the α -optimal constant sharing rule as a function of α for the following price elasticities of supply: $k = \{0.4, 0.5, 0.6, 1, 1.5\}$.

Furthermore, we explore a way of adjusting the design of the AdX mechanism to incorporate the above characterization, and apply it on real data sets. In particular, we use data on the distribution of publishers' opportunity costs to estimate (i) the polynomial distribution that better approximates its random generating process, and (ii) the objective function pursued by intermediaries as a function of the revenue share given to publishers. Our analysis suggests that a revenue share of 80% (which is the current practice of Google's AdEx) is consistent with assigning a weight of 47% to publishers' revenues on the platform's objective. See details in Section 4.1.

Constant Fee. Finally, we analyze an alternative model of competition between auction platforms. Namely, we assume that competition drives down expected profits to equal the quasi-fixed cost of intermediating each publisher's trans-

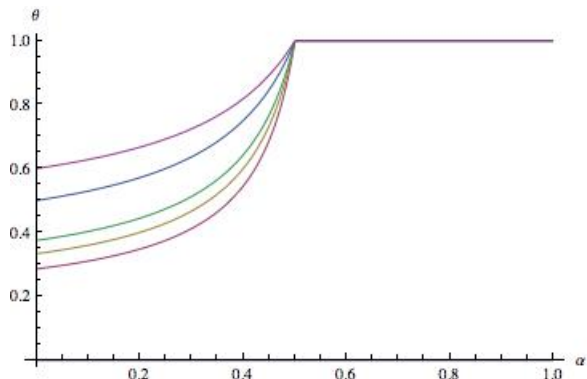


Figure 1: The α -optimal revenue sharing rules as a function of α when the price elasticity of supply equals $k = 0.4, 0.5, 0.6, 1$ and 1.5 (from lower to higher curves respectively).

action. We give a characterization of optimal double auctions in this setting, and present an indirect implementation in terms of sharing schemes.

1.2 Related Work

Double Auctions. Optimal double auctions have been widely studied in the economics literature and more recently CS literature. Following the Nobel-prize winning work of [7], [8] generalized the design of optimal auctions for two-sided settings, and characterized the platform-optimal double auction. This article generalizes the analysis of [8] by (i) studying settings in which the platform’s objective function is a convex combination of the seller’s profit and the platform’s profit, (ii) analyzing indirect implementation by sharing schemes, and (iii) providing a necessary and sufficient condition under which constant sharing schemes indirectly implement the optimal mechanism.

More recently, [1] studied multidimensional variants of the Bayesian optimal double auctions and present polynomial-time approximation algorithms for the problem. For the non-Bayesian (prior-free) setting, [2] studied revenue-maximizing double auctions when the auctioneer has no prior knowledge about bids. Unlike this paper, all the above work focus on the platform-optimal auctions, and do not consider other objective functions or fixed-percentage or fixed-cost auctions.

Competing Mechanisms. There is a large literature in economics studying competing mechanism design. The seminal work of [6] allows principals to post arbitrary trading mechanisms, but assumes that principals’ deviations do not account for their effects on the payoffs of agents who visit other principals (this is known as the large market assumption). In this setting, McAfee shows that there exists an equilibrium in which principals post second-price auctions. Dispensing with the large market assumption, but assuming that principals are only allowed to post second-price auc-

tions, [10] derive asymptotic reserve prices when the number of principals grows large.

Allowing for both arbitrary mechanisms and also dispensing with the large markets assumption, [9] provides a sufficient condition for the equilibrium outcome to be quasi-efficient (in the sense that inefficiencies are due to exclusion, but not to misallocations). Under this condition, sellers employ second-price auctions in equilibrium.

In contrast to the above cited papers, in our model, publishers choose which platform to patronize before knowing the realized opportunity cost of each impression. This assumption captures the notion that the publishers’ choice of platform is often a long-term choice, while the opportunity cost of each impression is a short-term contingency.

Other E-commerce Applications. Other than ad serving systems, the results of this paper apply to various online and offline retailers and e-commerce websites like Amazon and Ebay. For example, Ebay applies similar revenue-sharing auctions to the ones studied in this paper when it serves as a broker between a set of buyers and a seller. Roughly speaking, Ebay takes a 9% cut on each sale, referred to as *final value fee*, and also fixed fee for listing an item, referred to as an *insertion fee*¹. They also apply a convex cost function for the fixed fee as the number of listings, and a maximum of \$250 for the final value fee. A very recent paper by [4] studies EBay’s double auction problem, but their setting is different from this paper as they consider multiple sellers and one buyer, and explore approximately optimal pricing schemes for this setting.

2. MODEL

Before studying competition, we will introduce the model’s notation by considering the problem of a single platform populated by N buyers indexed by $i = 1, \dots, N$ and one seller (whose index is 0). The seller can obtain a price v_0 for his good by trading it outside of the platform (for example, by using the impression to fulfill guaranteed contracts). The value of v_0 is private information of the seller. From the perspective of the platform, v_0 is a draw from the distribution G with support $[a, b]$ (and density g).

Each buyer i has a valuation v_i for the good offered by the seller. The valuation of each buyer is his private information. From the perspective of the platform, the valuation of each buyer is an independent draw from the distribution F with support $[a, b]$ (and density f). As standard in mechanism design, we assume that F possesses a weakly increasing hazard rate, and that G possesses a weakly decreasing reverse hazard rate.

Denote by $\mathbf{v} \equiv (v_0, v_1, \dots, v_N)$ the profile of types (including the buyer and all sellers). Following the Revelation Principle, we will confine attention to truthful direct-revelation mechanisms. A direct-revelation mechanism consists of (i) an allocation rule $\{q_i(\mathbf{v})\}_{i=0,1,\dots,N}$ that maps \mathbf{v} into probabilities that each agent is assigned the good ($q_0(\mathbf{v})$ stands for the probability that the seller keeps the good), and (ii) a payment rule $\{p_i(\mathbf{v})\}_{i=0,1,\dots,N}$ that maps \mathbf{v} into payments for each buyer i and the seller. We denote a mechanism by $M = \{\{q_i(\mathbf{v})\}_{i=0,1,\dots,N}, \{p_i(\mathbf{v})\}_{i=0,1,\dots,N}\}$.

¹For details see <http://en.wikipedia.org/wiki/EBay>, and the charging table at http://pages.ebay.com/help/sell/fees.html#if_auction.

To recap, direct-revelation mechanisms display the following extensive form:

1. The seller is asked to report his opportunity cost v_0 and buyers are asked to report their valuations v_i ,
2. The platform delivers the good according to the allocation rule $\{q_i(\mathbf{v})\}_{i=0,1,\dots,N}$,
3. The platform charges $p_i(\mathbf{v})$ to each buyer and transfers $-p_0(\mathbf{v})$ to the seller.

An allocation is feasible if $\sum_{i=0}^N q_i(\mathbf{v}) = 1$ for all \mathbf{v} .

Denote by $Q_i(v_i) \equiv \mathbb{E}_{\mathbf{v}_{-i}}[q_i(\mathbf{v})]$ the interim probability that agent i is assigned the good and by $P_i(v_i) \equiv \mathbb{E}_{\mathbf{v}_{-i}}[p_i(\mathbf{v})]$ the interim payment of each agent i .

A mechanism is individually rational (IR) and incentive compatible (IC) for buyers if and only if for all $i \in \{1, \dots, N\}$

$$U_i(v) = Q_i(v) \cdot v - P_i(v) \geq \max \{Q_i(\hat{v}) \cdot v - P_i(\hat{v}), 0\}.$$

for all $v, \hat{v} \in [a, b]$.

A mechanism is individually rational and incentive compatible for the seller if and only if

$$U_0(v) = Q_0(v) \cdot v - P_0(v) \geq \max \{Q_0(\hat{v}) \cdot v - P_0(\hat{v}), v\}.$$

for all $v, \hat{v} \in [a, b]$.

It is standard to show that a mechanism is incentive compatible if and only if for all $i \in \{0, 1, \dots, N\}$, $Q_i(\cdot)$ is weakly increasing and

$$U_i(v) = U_i(a) + \int_a^v Q_i(\hat{v}) d\hat{v} = U_i(b) - \int_v^b Q_i(\hat{v}) d\hat{v}. \quad (1)$$

The platform's expected profits from mechanism M are then

$$\Pi(M) \equiv \sum_{i=1}^N \int_a^b P_i(v) dF(v) + \int_a^b P_0(v) dG(v).$$

The seller's ex-ante expected payoff from mechanism M is

$$\Gamma(M) \equiv \int_a^b U_0(v) dG(v).$$

As we will show in Section 2.1, the optimal best response strategy for the platform is to solve the following α -optimal problem: The α -optimal problem (denoted P^α) is to choose $\{q_i(\mathbf{v}), p_i(\mathbf{v})\}_{i=0,1,\dots,N}$ to

$$P^\alpha : \max \left\{ \alpha^k \cdot \Gamma(M) + (1 - \alpha^k) \cdot \Pi(M) \right\}, \text{ where } \alpha \in [0, 1], \quad (2)$$

subject to IR, IC, the feasibility constraint, and the platform's break-even constraint

$$\Pi \geq 0, \quad (3)$$

which states that the platform makes non-negative profits. As will be discussed in Section 2.1, the parameter α captures the intensity of competition in the market: in more competitive markets, the platform should assign a higher weight in its objective function to the sellers' payoffs. We refer to the mechanism that solves problem P^α as the α -optimal mechanism, and denote it by $\{q_i^\alpha(\mathbf{v}), p_i^\alpha(\mathbf{v})\}_{i=0,1,\dots,N}$.

2.1 Platforms' Competition and Best Responses

In this section, we study some general properties of competition games between K platforms, indexed by $k \in \{1, \dots, K\}$. It is often the case that publishers often have ex-ante preferences over which platform to join. These ex-ante preferences can be conveniently summarized by a vector (h^1, \dots, h^K) , where each component h^k describes the pure utility gain (or loss) that publishers associate with joining platform k .

Consider the following extensive-form competition game between K platforms.

1. Each platform k simultaneously post a mechanism M^k consisting of an allocation rule $\{q_i^k(\mathbf{v})\}_{i=0,1,\dots,N}$ and a payment rule $p_i^k(\mathbf{v})$ to each buyer and transfers $-p_0^k(\mathbf{v})$ to the seller.
2. Before observing her opportunity costs v_0 , but knowing her horizontal preferences (h^1, \dots, h^K) , the seller chooses the platform k that maximizes her ex-ante utility

$$\Gamma(M^k) + h^k.$$

All buyers join all platforms.

3. The buyers and the seller play the mechanism proposed at stage 1 by the platform selected by the seller.

Notice that the class of games described above accommodates any form of horizontal differentiation a la Hotelling [3].

Consider the set I of all pairs (Π, Γ) of expected payoffs to the platform and seller induced by all implementable mechanisms. The Pareto frontier of I , denoted by E , consists of all pairs (Π, Γ) such that there exists no implementable mechanism that leads to expected payoffs (Π', Γ') such that $\Pi' > \Pi$ and $\Gamma' \geq \Gamma$, or $\Pi' \geq \Pi$ and $\Gamma' > \Gamma$. If a pair (Π, Γ) belongs to the Pareto frontier E , we say that (Π, Γ) is Pareto efficient.

It is important to note that the set of implementable payoffs I is a convex set. To see why, take two implementable mechanisms M and M' that lead to expected payoffs (Π, Γ) and (Π', Γ') , respectively. Now consider the random mechanism $M'' \equiv \beta \cdot M + (1 - \beta) \cdot M'$. It follows directly from the incentive and rationality constraints that the mechanism M'' is implementable. Moreover, the expected payoffs of mechanism M'' are $(\beta \cdot \Pi + (1 - \beta) \cdot \Pi', \beta \cdot \Gamma + (1 - \beta) \cdot \Gamma')$.

The previous observation is key to prove the following result.

Proposition 1. Let $\{M^1, \dots, M^K\}$ be a Nash equilibrium of the competition game described above, and let (Π^k, Γ^k) be the expected payoff of mechanism M^k . Then there exists a vector $(\alpha^1, \dots, \alpha^K) \in [0, 1]^K$ such that each mechanism M^k solves

$$P^{\alpha^k} : \max_M \left\{ \alpha^k \cdot \Gamma(M) + (1 - \alpha^k) \cdot \Pi(M) \right\}.$$

Moreover, if the equilibrium is symmetric, $\alpha^k = \alpha$ for all k .

The proof of this proposition is an immediate application of the Support Hyperplane Theorem [5] (together with the fact that, in any equilibrium, the pair of expected payoffs is Pareto efficient). The result above formalizes the intuition

that competition among platforms has the effect of making each platform internalize (at least partially) the sellers' payoffs in its objective function.

Obviously, the vector of weights $(\alpha^1, \dots, \alpha^K)$ is endogenous to the equilibrium, and crucially depends on the horizontal preferences of the seller over platforms. Each α^k captures the equilibrium intensity of competition faced by platform k : in more competitive markets, each platform should assign a higher weight in its objective function to the sellers' payoffs.

Proposition 1 allows us to abstract from the details of the competition game played by firms, and derive properties of equilibrium mechanisms of any competition game. Specifically, in light of this proposition, we can characterize the equilibrium play of each firm by studying its respective P^{α^k} -problem. As stated earlier, for some arbitrary α , we refer to the mechanism that solves problem P^α as the α -optimal mechanism, and denote it by $\{q_i^\alpha(\mathbf{v}), p_i^\alpha(\mathbf{v})\}_{i=0,1,\dots,N}$.

Obviously, when $\alpha = 0$, the platform's problem is that of maximizing profits (as in Myerson and Satterthwaite (1983)). In turn, when $\alpha = 1$, the platform maximizes the seller's payoff subject to attaining weakly positive profits (which, as will be clear soon, is equivalent to the problem considered in Myerson (1981)). For $\alpha \in (0, 1)$, the platform behaves as if it were "partially owned" by sellers.

3. THE α -OPTIMAL MECHANISM

In this section, we derive the direct-revelation mechanism that solves each problem P^α .

Proposition 2. (The α -Optimal Direct-Revelation Mechanism) Let us choose indexes such that $v_i = \max_{j \in \{1, \dots, N\}} \{v_j\}$. The α -optimal direct-revelation mechanism is described below.

1. Let $\alpha \leq \frac{1}{2}$. Then $q_i^\alpha(v_i) = 1$ if

$$v_i - \frac{1 - F(v_i)}{f(v_i)} - v_0 - \frac{1 - 2 \cdot \alpha}{1 - \alpha} \cdot \frac{G(v_0)}{g(v_0)} \geq 0.$$

Otherwise, no sale occurs: $q_0^\alpha(v) = 1$.

2. Let $\alpha > \frac{1}{2}$. Then $q_i^\alpha(v_i) = 1$ if

$$v_i - \frac{1 - F(v_i)}{f(v_i)} - v_0 \geq 0.$$

Otherwise, no sale occurs: $q_0^\alpha(v) = 1$.

Proof. Applying the envelope formula (1), the platform's profits can be rewritten as

$$\int_{[a,b]^{N+1}} \left\{ \sum_{i=1}^N q_i(\mathbf{v}) \cdot \left(v_i - \frac{1 - F(v_i)}{f(v_i)} - v_0 - \frac{G(v_0)}{g(v_0)} \right) \right\} \cdot \prod_{i=1}^N dF(v_i) dG(v_0) - \sum_{i=1}^N U_i(a) - U_0(b). \quad (4)$$

In turn, the seller's payoff can be written as

$$\begin{aligned} \int_a^b U_0(v_0) dG(v_0) &= \\ U_0(b) - \int_a^b \int_{v_0}^b Q_0(\tilde{v}) g(v_0) d\tilde{v} dv_0 &= \\ U_0(b) - \int_a^b G(v_0) \cdot Q_0(v_0) dv_0 &= \\ U_0(b) - \int_a^b G(v_0) \cdot & \\ \int_a^b \dots \int_a^b \left(1 - \sum_{i=1}^N q_i(\mathbf{v}) \right) \left(\prod_{i=1}^N f(v_i) \right) \left(\prod_{i=1}^N dv_i \right) dv_0 &= \\ U_0(b) - b + \mathbb{E}(\tilde{v}_0) + & \\ \int_{[a,b]^{N+1}} \left\{ \sum_{i=1}^N q_i(\mathbf{v}) \cdot \frac{G(v_0)}{g(v_0)} \right\} \cdot \prod_{i=1}^N dF(v_i) dG(v_0). & \quad (5) \end{aligned}$$

The platform objective is then

$$\begin{aligned} \int_{[a,b]^{N+1}} \sum_{i=1}^N q_i(\mathbf{v}) \cdot & \\ \left((1 - \alpha) \cdot \left[v_i - \frac{1 - F(v_i)}{f(v_i)} - v_0 \right] - (1 - 2\alpha) \cdot \frac{G(v_0)}{g(v_0)} \right) \cdot & \\ \prod_{i=1}^N dF(v_i) dG(v_0) - & \\ (1 - \alpha) \cdot \left(\sum_{i=1}^N U_i(a) \right) + \alpha \cdot (\mathbb{E}(\tilde{v}_0) - b) + (2\alpha - 1) \cdot U_0(b). & \quad (6) \end{aligned}$$

If $\alpha \leq \frac{1}{2}$, the platform's objective is decreasing in $U_i(a)$ and $U_0(b)$. This implies that, at the α -optimum, the individual rationality constraints have to bind for every buyer of type a and every seller of type b : $U_i(a) = 0$ and $U_0(b) = b$. Maximizing the integral above point-wise leads to the following bang-bang solution: $q_i^\alpha(v_i) = 1$ for $i \in \{1, \dots, N\}$ if

$$v_i - \frac{1 - F(v_i)}{f(v_i)} - v_0 - \frac{1 - 2 \cdot \alpha}{1 - \alpha} \cdot \frac{G(v_0)}{g(v_0)} = \max_{j \in \{1, \dots, N\}} \left\{ v_j - \frac{1 - F(v_j)}{f(v_j)} - v_0 - \frac{1 - 2 \cdot \alpha}{1 - \alpha} \cdot \frac{G(v_0)}{g(v_0)}, 0 \right\}, \quad (7)$$

and $q_0^\alpha(v) = 1$ if there is no $j \in \{1, \dots, N\}$ that satisfies the equality above.

In turn, if $\alpha > \frac{1}{2}$, the platform's objective is decreasing in $U_i(a)$ and increasing in $U_0(b)$. This implies that, at the α -optimum, $U_i(a) = 0$, and $U_0(b)$ is set to satisfy the break-even constraint with equality:

$$\begin{aligned} U_0(b) &= \\ \int_{[a,b]^{N+1}} \sum_{i=1}^N q_i(\mathbf{v}) & \\ \cdot \left(v_i - \frac{1 - F(v_i)}{f(v_i)} - v_0 - \frac{G(v_0)}{g(v_0)} \right) & \\ \cdot \prod_{i=1}^N dF(v_i) dG(v_0). & \end{aligned}$$

Plugging (8) into the objective (6) leads to

$$\begin{aligned} & \int_{[a,b]^{N+1}} \sum_{i=1}^N q_i(\mathbf{v}) \\ & \cdot \alpha \cdot \left[v_i - \frac{1 - F(v_i)}{f(v_i)} - v_0 \right] \cdot \prod_{i=1}^N dF(v_i) dG(v_0) \\ & + \alpha \cdot (\mathbb{E}(\tilde{v}_0) - b). \end{aligned} \quad (8)$$

Maximizing the integral above point-wise leads to the following bang-bang solution: $q_i^\alpha(v_i) = 1$ for $i \in \{1, \dots, N\}$ if

$$v_i - \frac{1 - F(v_i)}{f(v_i)} - v_0 = \max_{j \in \{1, \dots, N\}} \left\{ v_j - \frac{1 - F(v_j)}{f(v_j)} - v_0, 0 \right\},$$

and $q_0^\alpha(v) = 1$ if there is no $j \in \{1, \dots, N\}$ that satisfies the equality above.

Because F has an increasing hazard rate and G has a decreasing reverse hazard rate, it follows that $q_i^\alpha(v_i, \mathbf{v}_{-i})$ is weakly increasing in v_i for every $\alpha \in [0, 1]$ (i.e., the α -optimal mechanism is implementable). Q.E.D.

When $\alpha \geq \frac{1}{2}$, the α -optimal mechanism coincides with the mechanism that the seller would run if he owned the platform (derived in Myerson (1981)). To understand this result, notice that the platform's objective function puts weakly higher weight on the seller's payoff than in its own profits. Because of the seller's incentive constraints, by giving up one dollar in profits, the platform can add more than on dollar to the seller's payoff. As a consequence, the platform maximizes its objective by following the seller's optimal mechanism (that is, the Myerson auction).

When $\alpha < \frac{1}{2}$, the platform puts strictly more weight on its profits than on the seller's payoffs. As a consequence, it is willing to introduce distortions on the seller side of the market (therefore reducing the seller's payoffs) in order to generate higher profits. The allocative distortions generated by the seller's informational rents are captured by the term

$$\frac{1 - 2 \cdot \alpha}{1 - \alpha} \cdot \frac{G(v_0)}{g(v_0)},$$

which is decreasing in α . As expected, when $\alpha = 0$, the α -optimal mechanism coincides with the platform profit-maximizing mechanism derived in Myerson in Satterthwaite (1983).

4. INDIRECT IMPLEMENTATION: SHARING RULES AND CONSTANT CUTS

The direct-revelation mechanism derived above offers little insight on the actual practice of ad-exchanges, as the payment functions that implement the optimal allocations are not very intuitive. In order to shed light on the actual mechanisms used by advertising platforms, we will now study one natural (and widely adopted) indirect implementation: Revenue-sharing mechanisms. They work as follows:

1. The seller is asked to report a reserve price \hat{r} .
2. The platform runs a second-price auction with reserve price \hat{r} .
3. The platform gives to the seller a fraction $\theta(\hat{r})$ of the proceeds of the auction (we call $\theta(\cdot)$ the *sharing rule*).

Before proceeding, let us define the function

$$h(\alpha) = \begin{cases} \frac{1-2\cdot\alpha}{1-\alpha} & \text{if } \alpha \leq \frac{1}{2} \\ 0 & \text{if } \alpha > \frac{1}{2}. \end{cases}$$

Note that the function $h(\alpha)$ is continuous at $\alpha = \frac{1}{2}$.

The next proposition shows how to define the sharing rule $\theta^\alpha(\cdot)$ that implements the α -optimal mechanism identified in Proposition 2. To this end, let us define the function $r^\alpha(\cdot)$ such that

$$r^\alpha(v_0) - \frac{1 - F(r^\alpha(v_0))}{f(r^\alpha(v_0))} = v_0 + h(\alpha) \cdot \frac{G(v_0)}{g(v_0)},$$

and the inverse $v_0^\alpha(r)$ according to

$$r - \frac{1 - F(r)}{f(r)} = v_0^\alpha(r) + h(\alpha) \cdot \frac{G(v_0^\alpha(r))}{g(v_0^\alpha(r))}.$$

Define the *minimum reserve price* \underline{r} according to $\underline{r} - \frac{1 - F(\underline{r})}{f(\underline{r})} = a$ (if this equation does not admit a solution in $[a, b]$, set $\underline{r} = a$).

Denote by $R(r, N)$ the expected revenue of a second-price auction with N bidders and reserve price r :

$$R(r, N) \equiv N \cdot \int_r^b \left\{ v - \frac{1 - F(v)}{f(v)} \right\} f(v) \cdot F^{N-1}(v) dv.$$

Proposition 3. (α -Optimal Revenue-Sharing Mechanism) Consider the revenue-sharing mechanism described by the sharing rule

$$\theta^\alpha(\hat{r}) = \frac{b - v_0^\alpha(\hat{r}) \cdot F^N(\hat{r}) - \int_{v_0^\alpha(\hat{r})}^b F^N(r(\hat{v})) d\hat{v}}{R(\hat{r}, N)} \quad \text{for } \hat{r} \in [\underline{r}, b] \quad (9)$$

and $\theta^\alpha(\hat{r}) = 0$ for $\hat{r} < \underline{r}$. This mechanism indirectly implements the α -optimal direct-revelation mechanism.

Proof. By the Envelope formula (1), incentive compatibility implies that the expected payment to a seller with value v_0 under the α -optimal mechanism is given by

$$b - v_0 \cdot F^N(r^\alpha(v_0)) - \int_{v_0}^b F^N(r^\alpha(\hat{v})) d\hat{v}. \quad (10)$$

Rather than asking sellers to report v_0 , the revenue-sharing mechanism considered here requires that sellers report some reserve price r . Let us posit that the seller equilibrium reserve price strategy is given by $r^\alpha(v_0)$. We will derive the expected payments induced by incentive compatibility under this strategy, and then argue that submitting reserve prices according to $r^\alpha(v_0)$ is indeed profit-maximizing for the seller.

Because F has an increasing hazard rate, G has a decreasing reverse hazard rate, and $h(\cdot) \geq 0$, we know that the function $r^\alpha(v_0)$ is strictly increasing. Therefore we can rewrite the expected payments of a seller with value v_0 (as implied by IC) in terms of his submitted reserve price $\hat{r} = r^\alpha(v_0)$:

$$b - v_0^\alpha(\hat{r}) \cdot F^N(\hat{r}) - \int_{v_0^\alpha(\hat{r})}^b F^N(r^\alpha(\hat{v})) d\hat{v}.$$

We will define the sharing rule $\theta^\alpha(\cdot)$ such that the expected revenue that the seller obtains from the second-price auction run at stage 2 equals the expected payments implied by incentive compatibility. This is equivalent to

$$\theta^\alpha(\hat{r}) \cdot R(\hat{r}, N) = b - v_0^\alpha(\hat{r}) \cdot F^N(\hat{r}) - \int_{v_0^\alpha(\hat{r})}^b F^N(r^\alpha(\hat{v})) d\hat{v}.$$

After rearranging, we get the formula (9).

We will now argue that submitting reserve prices according to $r^\alpha(v_0)$ is profit-maximizing for the seller. To see why, notice that the seller problem can be written as

$$\max_{\hat{r}} v_0 \cdot F^N(\hat{r}) + \theta^\alpha(\hat{r}) \cdot R(\hat{r}, N).$$

By construction, the selection $r^\alpha(v_0)$ satisfies the envelope formula (1). Because the seller's objective satisfies strictly increasing differences, we can use the Constraint Simplification Theorem (Milgrom (2004), page 105) to conclude that $r^\alpha(v_0)$ maximizes the seller's payoff among all reserve prices in the range $[\underline{r}, b]$ if and only if $r^\alpha(\cdot)$ is weakly increasing. As argued before, because F has an increasing hazard rate, G has a decreasing reverse hazard rate, and $h(\cdot) \geq 0$, we know that the function $r^\alpha(v_0)$ is indeed strictly increasing. Because it is clearly suboptimal to submit a reserve price lower than \underline{r} (since $\theta^\alpha(\hat{r}) = 0$ for $\hat{r} < \underline{r}$) or greater than b , we conclude that submitting reserve prices according to $r^\alpha(v_0)$ is profit-maximizing for the seller. Q.E.D.

An important class of revenue-sharing mechanisms is described by *constant sharing rules*, where $\theta^\alpha(\cdot)$ is constant across all possible reserve prices. Such "constant-sharing schemes" are fairly simple to describe and are widely adopted in practice by Ad Exchanges and internet auction sites.

The next proposition provides a necessary and sufficient condition for the α -optimal mechanism to employ a constant sharing rule.

Proposition 4. (α -Optimal Constant Sharing Rule)

The α -optimal direct-revelation mechanism can be implemented by a revenue-sharing mechanism that employs a constant sharing rule if and only if

$$G(v_0) = \left(\frac{v_0}{b}\right)^k,$$

where $k > 0$. In this case,

$$\theta^\alpha(r) = \frac{k}{k + h(\alpha)} \quad \text{for all } r \in [\underline{r}, b],$$

and $\theta^\alpha(\hat{r}) = 0$ for $\hat{r} < \underline{r}$.

Before going through the proof, we present the following example confirming the above result for the case of uniform distribution and $\alpha = 0$, i.e., the platform-optimal mechanism. Interestingly, under uniform distribution, the platform-optimal mechanism can be indirectly implemented by a fifty-fifty revenue sharing rule.

Example 1. Assume $\alpha = 1$, and $F = G \sim U[0, b]$, in which case $r(v_0) = \frac{b}{2} + v_0$ and $r^{-1}(\hat{r}) = \hat{r} - \frac{b}{2}$. Then the seller's revenue under some sharing rule $\lambda(\hat{r})$ (which a priori is not necessarily constant) is given by

$$\begin{aligned} \lambda(\hat{r}) \cdot N \cdot \int_{\hat{r}}^b \left\{ v - \frac{1 - F(v)}{f(v)} \right\} f(v) \cdot F^{N-1}(v) dv &= \\ \lambda(\hat{r}) \cdot N \cdot \int_{\hat{r}}^b \{2v - b\} \cdot \frac{1}{b} \cdot \left(\frac{v}{b}\right)^{N-1} dv &= \\ \lambda(\hat{r}) \cdot \frac{N}{b^N} \cdot \int_{\hat{r}}^b \{2v^N - bv^{N-1}\} dv &= \lambda(\hat{r}) \cdot \frac{N}{b^N} \cdot \left[\frac{2v^{N+1}}{N+1} - \frac{bv^N}{N} \right]_{\hat{r}}^b \end{aligned}$$

$$= \lambda(\hat{r}) \cdot \frac{N}{b^N} \cdot \left[b^{N+1} \cdot \frac{N-1}{N(N+1)} - \hat{r}^{N+1} \cdot \frac{2}{N+1} + b\hat{r}^N \cdot \frac{1}{N} \right]. \quad (11)$$

In turn, incentive compatibility (and the envelope formula) deliver another expression for the seller's revenue:

$$\begin{aligned} b - r^{-1}(\hat{r}) \cdot F^N(\hat{r}) - \int_{r^{-1}(\hat{r})}^b F^N(r(\hat{v})) d\hat{v} &= \\ b - \left(\hat{r} - \frac{b}{2}\right) \cdot \left(\frac{\hat{r}}{b}\right)^N - \int_{\hat{r} - \frac{b}{2}}^{\frac{b}{2}} \left(\frac{\frac{b}{2} + \tilde{v}_0}{b}\right)^N d\tilde{v}_0 - \int_{\frac{b}{2}}^b d\hat{v} &= \\ \frac{b}{2} - \left(\hat{r} - \frac{b}{2}\right) \cdot \left(\frac{\hat{r}}{b}\right)^N - \frac{1}{b^N} \left[\left(\frac{b}{2} + \tilde{v}_0\right)^{N+1} \cdot \frac{1}{N+1} \right]_{\hat{r} - \frac{b}{2}}^{\frac{b}{2}} &= \\ \frac{b}{2} + \frac{\hat{r}^{N+1}}{b^N} \left(1 - \frac{N}{N+1}\right) + b \frac{\hat{r}^N}{b^N} \cdot \frac{1}{2} - \frac{b^{N+1}}{b^N} \frac{1}{N+1} &= \\ \frac{N}{b^N} \cdot \left[b^{N+1} \cdot \left(\frac{1}{2N} - \frac{1}{N(N+1)}\right) - \hat{r}^{N+1} \frac{1}{N+1} + b\hat{r}^N \cdot \frac{1}{2N} \right] &= \\ \frac{N}{b^N} \cdot \left[b^{N+1} \cdot \frac{N-1}{2N(N+1)} - \hat{r}^{N+1} \frac{1}{N+1} + b\hat{r}^N \cdot \frac{1}{2N} \right]. & \quad (12) \end{aligned}$$

Comparing (11) and (12) leads to $\lambda(\hat{r}) = \frac{1}{2}$ for all $\hat{r} \in [\frac{b}{2}, b]$. This shows that the only revenue sharing auction that implements the platform-optimal mechanism shares the revenue from the auction evenly between the platform and the seller.

The proof below establishes that constant sharing rules implement the α -optimal mechanisms if and only if the distribution of the seller's outside options has a power-form distribution.

Proof of Proposition 4. The expression (9) can be rewritten as

$$\theta^\alpha(r) \cdot R(r, N) = b - v_0^\alpha(r) \cdot F^N(r) - \int_{v_0^\alpha(r)}^b F^N(r^\alpha(\hat{v})) d\hat{v}.$$

Differentiating with respect to r leads to the following linear differential equation.

$$\begin{aligned} (\theta^\alpha)'(r) \cdot R(r, N) + \theta^\alpha(r) \cdot \frac{\partial R}{\partial r}(r, N) &= \\ -v_0^\alpha(r) \cdot N \cdot F^{N-1}(r) \cdot f(r). \end{aligned}$$

Therefore, $(\theta^\alpha)'(r) = 0$ for all r if and only if

$$\theta^\alpha(r) \cdot \frac{\partial R}{\partial r}(r, N) = -v_0^\alpha(r) \cdot N \cdot F^{N-1}(r) \cdot f(r).$$

Note that at the α -optimal mechanism

$$\begin{aligned} \frac{\partial R}{\partial r}(r, N) &= \\ -N \cdot \left\{ r - \frac{1 - F(r)}{f(r)} \right\} f(r) \cdot F^{N-1}(r) &= \\ -N \cdot \left\{ v_0^\alpha(r) + h(\alpha) \cdot \frac{G(v_0^\alpha(r))}{g(v_0^\alpha(r))} \right\} f(r) \cdot F^{N-1}(r). \end{aligned}$$

Therefore, $(\theta^\alpha)'(r) = 0$ for all r if and only if

$$\left(v_0^\alpha(r) + h(\alpha) \cdot \frac{G(v_0^\alpha(r))}{g(v_0^\alpha(r))} \right) \theta^\alpha(r) = v_0^\alpha(r),$$

which can be rewritten as

$$\frac{1 - \theta^\alpha(r)}{\theta^\alpha(r)} = h(\alpha) \cdot \frac{G(v_0^\alpha(r))}{v_0^\alpha(r) \cdot g(v_0^\alpha(r))}.$$

The function $v_0^\alpha(r)$ is strictly increasing in r . Therefore, the expression above is constant for every r if and only if $\frac{xg(x)}{G(x)}$ is constant, what leads to $G(v_0) = \left(\frac{v_0}{b}\right)^k$. Finally, note that if $G(v_0) = \left(\frac{v_0}{b}\right)^k$, then

$$\frac{1 - \theta^\alpha(r)}{\theta^\alpha(r)} = \frac{h(\alpha)}{k},$$

what leads to $\theta^\alpha(r) = \frac{k}{k+h(\alpha)}$ for all $r \in [x, b]$. Q.E.D.

The result above can be understood in the light of monopolistic price theory. Intuitively, polynomial distribution functions have a constant price elasticity of supply (which measures how many more percentage points of inventory sellers are willing to offer for one percentage increase in expected revenue). Namely, a distribution of the form $G(v_0) = \left(\frac{v_0}{b}\right)^k$ has a price elasticity of supply equal to k for all opportunity costs v_0 . As it turns out, constant sharing schemes are optimal provided that the price elasticity of supply is constant.

As the price elasticity k increases, the seller's revenue share goes up. Intuitively, as the distribution of the seller's opportunity costs become concentrated on high values, the platform has to increase the seller's revenue share to make sure that the seller is willing to participate in the trading mechanism with high enough probability.

In turn, for fixed k , the seller revenue share increases as the weight α on the seller's payoffs increases.

The next section applies the results above to data in order to estimate the objective function that rationalizes "constant-sharing schemes" commonly used in practice.

4.1 Data Analysis: Estimating a Platform's Objective Function

In this section we bring the theory developed above to data. We apply the results above to answer the following questions: (i) Given an objective function in mind for the platform, what is the best constant cut sharing rule to use? and (ii) What objective function rationalizes the revenue-sharing rule used by an Ad Exchange (AdX platform), which uses a constant sharing rule that assigns (for example) 80% of the auction revenue to sellers, and 20% to the platform?

To answer the above questions, a main methodological challenge is the estimation of the distribution of opportunity costs of publishers (as data on the value of guaranteed contracts signed by publishers is not available). In order to overcome this difficulty, we will make the assumption that the distribution of opportunity costs of publishers coincides with the distribution of revenue in the Ad Exchange. This assumption captures the idea that publishers should be indifferent between selling impressions in the Ad Exchange and selling impressions through guaranteed contracts (as otherwise, publishers are expected to sell all of their inventory through one of these channels).

Under this assumption, we can use the following procedure to estimate the platform-optimal mechanism:

- get the empirical distribution of revenues obtained by a given publisher on a particular ad slot,
- remove the observation where AdX platform was not able to sell the publisher's ad slot,
- estimate the power distribution that best fits the empirical revenue distribution, with coefficient k^*
- use this estimate to pin down the weight that the AdX platform gives to sellers' payoffs, denoted α^* , by solving the equation

$$0.8 = \frac{k^*}{k^* + h(\alpha^*)}.$$

In order to estimate the polynomial distribution that best fits the data, we use the maximum likelihood method. This method consists in choosing the distribution parameter k^* that maximizes the theoretical probability that the realized sample was generated by a power distribution with parameter k^* .

We applied the technique described above to five ad slots, and obtained the following results.

#obs	#sales	bid support	k^*	α^*
710,195	515,509	[0,20.35]	0.42	47.22%
42,037	37,152	[0,7.99]	0.65	45.51%
29,813	21,919	[0,12]	0.51	46.59%
78,324	65,582	[0,27.07]	0.39	47.43%
6,059	3,616	[0,54.38]	0.29	48.11%

Our estimate of α^* is roughly constant in the five data sets considered. Interestingly, our estimate suggests that an AdX platform that uses 20% as a constant cut puts roughly 46% of weight on the seller's payoffs when designing the auction rules of AdX. Also given a desirable objective function, and k^* for each publisher, the ad exchange can decide about the declared fixed percentage if it needs to negotiate them.

5. SELLER-OPTIMAL MECHANISM WITH A MINIMUM PROFIT CONDITION

In the previous sections, we studied the mechanism design problem of an Ad Exchange that wishes to maximize a convex combinations of its own profits and the seller's payoffs. This formulation captures in reduced form competition between Ad Exchanges: as more Ad Exchanges are present in the market, the higher is the weight assigned to the sellers' payoff in the Ad Exchanges objective function.

In this section, we follow an alternative approach to study competition between Ad Exchanges. Namely, we assume that Ad Exchanges maximize the sellers' payoff subject to attaining a minimal profit level (which should be used to cover operating costs, for example). Such reduced-form formulation is the outcome of a Bertrand competition game between Ad Exchanges (under the assumption that sellers can multi-home and have non-binding capacity constraints).

Formally, the Ad Exchange problem is to choose a mechanism $\{q_i(\mathbf{v}), p_i(\mathbf{v})\}_{i=0,1,\dots,N}$ to maximize the seller's profit subject to attaining a minimum profit level of π .

$$\max \int_a^b U_0(v) dG(v), \quad (13)$$

subject to IR, IC, the feasibility constraint, and

$$\sum_{i=1}^N \int_a^b P_i(v) dF(v) + \int_a^b P_0(v) dG(v) \geq \pi.$$

The next proposition characterizes the solution to this problem.

Proposition 5. (The π -Optimal Direct-Revelation Mechanism) Let us choose indexes such that $v_i = \max_{j \in \{1, \dots, N\}} \{v_j\}$. The π -optimal direct-revelation mechanism sets $q_i^\pi(v_i) = 1$ if and only if

$$v_i - \frac{1 - F(v_i)}{f(v_i)} - v_0 - \lambda(\pi) \cdot \frac{G(v_0)}{g(v_0)} \geq 0,$$

and sets $q_0^\pi(v) = 1$ otherwise. The function $\lambda : [0, \pi^*] \rightarrow [0, 1]$ is strictly increasing in π , and satisfies $\lambda(0) = 0$ and $\lambda(\pi^*) = 1$, where

$$\pi^* = \int_a^b \int_{r(v_0)}^b \left\{ v - \frac{1 - F(v)}{f(v)} - v_0 - \frac{G(v_0)}{g(v_0)} \right\} f(v) \cdot F^{N-1}(v) g(v_0) dv dv_0,$$

with $r(v_0) - \frac{1 - F(r(v_0))}{f(r(v_0))} = v_0 + \frac{G(v_0)}{g(v_0)}$. For each π , $\lambda(\pi)$ solves

$$\pi = \int_a^b \int_{s(v_0, \lambda(\pi))}^b \left\{ v - \frac{1 - F(v)}{f(v)} - v_0 - \lambda(\pi) \cdot \frac{G(v_0)}{g(v_0)} \right\} f(v) \cdot F^{N-1}(v) g(v_0) dv dv_0, \quad (14)$$

with $s(v_0, \lambda) - \frac{1 - F(s(v_0, \lambda))}{f(s(v_0, \lambda))} = v_0 + \lambda \cdot \frac{G(v_0)}{g(v_0)}$.

Proof. The platform's profits can be rewritten as

$$\int_a^b \dots \int_a^b \sum_{i=1}^N q_i(\mathbf{v}) \cdot \left(v_i - \frac{1 - F(v_i)}{f(v_i)} - v_0 - \frac{G(v_0)}{g(v_0)} \right) \cdot g(v_0) \left(\prod_{i=1}^N dF(v_i) \right) - U_i(a) + [b - U_0(b)].$$

In turn, the platform's objective (which is the seller's payoff) rewrites

$$U_0(b) - b + \mathbb{E}(\tilde{v}_0) - \int_a^b \dots \int_a^b \left\{ \sum_{i=1}^N q_i(\mathbf{v}) \cdot \frac{G(v_0)}{g(v_0)} \right\} \cdot g(v_0) \left(\prod_{i=1}^N dF(v_i) \right).$$

Expressing the platform's constrained maximization problem in Lagrangian form leads to

$$\begin{aligned} & \int_a^b \dots \int_a^b \sum_{i=1}^N q_i(\mathbf{v}) \\ & \cdot \left(\mu \cdot v_i - \mu \cdot \frac{1 - F(v_i)}{f(v_i)} - \mu \cdot v_0 - (1 + \mu) \cdot \frac{G(v_0)}{g(v_0)} \right) \\ & \cdot g(v_0) \left(\prod_{i=1}^N dF(v_i) \right) \end{aligned}$$

$$- \mu \cdot U_i(a) + \mu \cdot [b - U_0(b)] + U_0(b) - b + \mathbb{E}(\tilde{v}_0),$$

where $\mu > 0$. It is immediate that at the optimum $U_i(a) = 0$ and $U_0(b) = b$. Maximizing the integral above point-wise leads to the following bang-bang solution: $q_i^\pi(v_i) = 1$ if

$$\mu \cdot v_i - \mu \cdot \frac{1 - F(v_i)}{f(v_i)} - (1 + \mu) \cdot \left(v_0 - \frac{G(v_0)}{g(v_0)} \right) \geq 0,$$

and $q_0^\pi(v) = 1$ otherwise. Defining $\lambda \equiv \frac{\mu}{1+\mu}$ leads to the statement in the proposition. Finally, since the profit constraint is always binding, the value of λ is given by (14). Q.E.D.

In order to obtain insight on the actual practice of ad-exchanges, the next proposition derives the Revenue-sharing mechanisms that implement the direct-revelation mechanism described above. Define the function $v_0(r, \lambda)$ according to

$$r - \frac{1 - F(r)}{f(r)} = v_0(r, \lambda) + \lambda \cdot \frac{G(v_0(r, \lambda))}{g(v_0(r, \lambda))}.$$

Then the following result is true. The proof is identical to that of Proposition 4.1, and is therefore omitted.

Proposition 6. (π -Optimal Indirect Implementation) The following trading procedure indirectly implements the π -optimal mechanism:

1. The seller is asked to report a reserve price \hat{r} .
2. The platform runs a second-price auction with reserve price \hat{r} .
3. The platform gives to the seller a fraction $\theta^\pi(\hat{r})$ of the proceeds of the auction, where the sharing rule $\theta^\pi(\hat{r})$ satisfies

$$\theta^\pi(\hat{r}) = \frac{b - v_0(\hat{r}, \lambda(\pi)) \cdot F^N(\hat{r}) - \int_{v_0(\hat{r}, \lambda(\pi))}^b F^N(r(\hat{v})) d\hat{v}}{R(\hat{r}, N)}$$

for $\hat{r} \in [\underline{r}, b]$

Interestingly, constant sharing schemes are not able to implement the π -optimal mechanism. This suggests that platforms might need to adopt non-constant sharing schemes if competition in the ad exchange market increases, therefore approaching Bertrand competition, and operational costs per transaction are bounded by some π greater than zero.

6. FUTURE RESEARCH

This paper studies the optimality of constant sharing schemes in settings plagued by two-sided asymmetric information and competition. We provide necessary and sufficient conditions under which the optimal mechanism can be implemented by constant cuts for the seller, and analyze how such cuts are affected by the degree of competition in the market.

One fruitful direction of future research pertains the relative performance of constant sharing mechanisms when the necessary and sufficient conditions identified in this paper fail. In such settings, what fraction of the revenue associated with the optimal mechanism does constant sharing schemes achieve? The widespread use of constant sharing rules renders this an important question with first-order practical relevance.

Acknowledgements. We thank our colleagues Gagan Goel, Eyal Manor, and Martin Pal, for very helpful discussions and comments on the paper.

7. REFERENCES

- [1] X. Deng, P. Goldberg, B. Tang, and J. Zhang. Revenue maximization in a bayesian double auction market. In *ISAAC*, 2011.
- [2] K. Deshmukh, A. V. Goldberg, J. D. Hartline, and A. R. Karlin. Truthful and competitive double auctions. In *ESA*, pages 361–373, 2002.
- [3] H. Hotelling. Stability in competition. *The Economic Journal*, 39(153), 1929.
- [4] K. Jain and C. A. Wilkens. Ebay market intermediation problem. 2012.
- [5] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [6] R. P. McAfee. Mechanism Design by Competing Sellers. *Econometrica*, pages 1281–1312, 1993.
- [7] R. Myerson. Optimal auction design. *Mathematics of operations research*, pages 58–73, 1981.
- [8] R. Myerson and M. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economics Theory (JET)*, pages 265–281, 1983.
- [9] M. Pai. Competing auctioneers. 2012.
- [10] M. Peters and S. Severinov. Competition among sellers who offer auctions instead of prices. *Journal of Economics Theory (JET)*, 75:141–179, 1997.