

Optimal Information Disclosure in Classic Auctions[†]

By DIRK BERGEMANN, TIBOR HEUMANN, STEPHEN MORRIS,
CONSTANTINE SOROKIN, AND EYAL WINTER*

We characterize the revenue-maximizing information structure in the second-price auction. The seller faces a trade-off: more information improves the efficiency of the allocation but creates higher information rents for bidders. The information disclosure policy that maximizes the revenue of the seller is to fully reveal low values (where competition is high) but to pool high values (where competition is low). The size of the pool is determined by a critical quantile that is independent of the distribution of values and only dependent on the number of bidders. We discuss how this policy provides a rationale for conflation in digital advertising. (JEL D44, D82, D83, M37)

In digital advertising, the publisher sells impressions in real-time auctions to competing advertisers. As the publisher has detailed information about the viewers on the website, the publisher has a choice of how much information to reveal about the viewer to the competing bidders. Thus, the seller can influence the distribution of conditional expectations of the bidders for the value of the impressions. Motivated by this, we characterize the optimal information disclosure of a seller regarding the values of the bidders. More generally our analysis provides insights in environments where the seller's information is superior or supplemental in at least some dimensions to that of the bidders—for example, markets for experience and secondhand goods and markets for financial assets such insurance and private equity.

We consider a second-price auction where bidders' valuations are independently and symmetrically distributed but initially unknown to the bidders. The seller can

*Bergemann: Yale University (email: dirk.bergemann@yale.edu); Heumann: Pontificia Universidad Católica de Chile (email: tibor.heumann@uc.cl); Morris: Massachusetts Institute of Technology (email: semorris@mit.edu); Sorokin: Glasgow University and Higher School of Economics (email: constantine.sorokin@glasgow.ac.uk); Winter: Hebrew University of Jerusalem and Lancaster University (email: eyal.winter@mail.huji.ac.il). Andrzej Skrzypacz was coeditor for this article. We acknowledge financial support from NSF grants SES-2001208, SES-1948336, and SES-2049744 and ANID Fondecyt Iniciación 11200365. We thank three anonymous referees for helpful suggestions. We thank Steve Berry, Ozan Candogan, Roberto Corrao, Phil Haile, Sergiu Hart, Jason Hartline, Nicole Immorlica, Maxim Ivanov, Anton Kolotilin, Renato Paes Leme, Denis Shishkin, Rann Smorodinsky, Philipp Strack, Alex Wolitzky, and Song Zuo for informative discussions. We thank seminar audiences at Informs AMD, UCSD, PUC Chile, Harvard-MIT, Hebrew University, the China Meetings of the Econometric Society, and the Stony Brook International Conference on Game Theory for many helpful comments. This manuscript combines and subsumes the results in Bergemann, Heumann, and Morris (2021) and Sorokin and Winter (2021). These two papers shared the current Theorem 1, obtained via different proof strategies.

[†]Go to <https://doi.org/10.1257/aeri.20210504> to visit the article page for additional materials and author disclosure statement(s).

choose what information each bidder can learn about their own value. If the seller did not allow them to learn anything, then all bidders would bid their (common) expected value and the good would be randomly (and inefficiently) allocated among them. If the seller allowed bidders to learn their true value, then they would have a dominant strategy (under private values in a second-price auction) to bid their values. The good would be allocated efficiently to the bidder with the highest value. The revenue of the seller would equal the value of the efficient allocation minus the bidders' information rent. By permitting bidders to learn something but not everything about their values, the seller can trade off efficiency loss with information rent reduction. Our main result is a characterization of the optimal (among symmetric) information policies for the seller. In our analysis, we assume that the seller is unconstrained in the choice of information disclosure. In particular, the bidders do not possess any private information *a priori* that would constrain the information disclosure of the seller. This is an admittedly substantive assumption but one that helps us to frame the trade-off for the seller sharply.

Conditional on having a low value, a bidder is likely to be competing with other bidders and earn low information rents. But conditional on having a high value, a bidder is likely to win (facing weak competition) and thus can expect to win at a price significantly below his value, thus earning high information rents. Thus, the gains from concealing information will be highest when valuations are high. In the optimal policy, high values are pooled and low values are revealed. There is a critical threshold described by a quantile above which all valuations are bundled together (Theorem 1). The threshold is given by a quantile of the distribution that depends only on the number of bidders and not the distribution of valuations. The optimal quantile above which disclosure occurs is increasing in the number of participating bidders and goes toward 1 (i.e., full disclosure) as the number of bidders grows arbitrarily large. Thus, the information policy is influencing the distribution of bids, holding fixed the distribution of preferences among the bidders.

The optimal threshold is designed to keep a moderate level of competition at the top of the bid distribution. If the threshold is too low, more bidder surplus will be extracted but the expected value conditional on being above the threshold will be too low. If the threshold is too high, the expected value conditional on being above the threshold will be high but too little bidder surplus will be extracted. The optimal quantile keeps “around” two bidders at the top of the bid distribution—that is, with the same expected value and competing with each other. This gives just enough competition to extract the bidder surplus at low efficiency cost.

In the current setting with independent private values, the revenue equivalence result holds. Thus, all classic auction formats—such as first-price auction, second-price auction, and all-pay auction—generate the same expected revenue. Hence, while we formally study the second-price auction, the results presented here extend to all classic and revenue-equivalent auction formats.

We then consider a number of variations of the second-price auction and show how our insights remain robust and relevant in those settings. First, we consider the second-price auction with a given reserve price. We show that the optimal information policy maintains the earlier structure but now introduces a second interval of pooled values around the reserve price. We then extend the allocation problem to where K identical objects are allocated in a uniform $(K + 1)$ st-price auction. The

newly relevant $(K + 1)$ st-order statistics share the same curvature property as the second-order statistic and hence yield comparable optimal information structures. Finally, we consider objectives different from the revenue of the seller, such as the social surplus or the bidders' surplus. As all of these objectives can be represented as a convex combination of first- and second-order statistic, we obtain a complete characterization of the optimal information structure, now in terms of upper and lower censored information structures.

Our paper's motivation is the market for impressions in digital advertising. A large share of digital advertising, whether in search, display advertising, or social networks, is allocated by auction mechanisms. In digital advertising the auction forms a match between competing advertisers (the *bidders*) and a viewer. A match between viewer and advertiser creates an *impression* (or search result) on the publisher's website. We sketch in the final section a model of the market for impressions and describe its relationship to information design.

Literature.—Levin and Milgrom (2010) suggested that the idea of conflation (central in many commodity markets) by which similar but distinct products are treated as identical in order to make markets thicker or reduce cherry-picking, may be relevant for the design on online advertising markets. The optimal information structure we derive in Theorem 1 exactly determines when conflation should occur, in the upper interval, and when not, in the lower interval.

The paper relates to the literature studying optimal information disclosure in selling mechanisms. Milgrom and Weber (1982) established the “linkage principle,” which showed that public information in affiliated value auctions is revenue increasing. Board (2009) shows in a private value environment that the disclosure of information generates a trade-off between improving efficiency and decreasing competition. Bourreau, Caillaud, and de Nijs (2018) obtain a similar result with a specific model of differentiated advertisers and so do Hummel and McAfee (2016) by considering a specific form of public information considered earlier by Palfrey (1983), one that unbundles many symmetric objects.

Ganuza (2004) studies the optimal information disclosure in a second-price auction. The bidder's valuations are independent and private and determined by the quality of a match between a bidder's taste and the good's characteristic represented by Hotelling's model on a circle. The seller chooses the optimal public signal about the good's characteristic. The disclosure of information occurs through a *costly* public signal, and the equilibrium information provision is less than the social efficient level. Bergemann and Pesendorfer (2007) analyze the joint optimal design of auction and information structure. In particular, they allow for asymmetric information structures and personalized reserve prices. Here, we fix the selling mechanism to be a second-price auction and focus on the role of information to improve the revenue of the seller.

A recent strand of literature allows the information disclosure by the seller to depend on prior private information elicited from the bidders (see Eső and Szentes 2007; Li and Shi 2017). A critical issue is then whether the information disclosure should discriminate across agents and across types.

I. Model

There are N bidders who compete for an indivisible good in an auction. Bidder i 's value is denoted by $v_i \in \mathbb{R}_+$. We assume that the valuations are independently and identically distributed across bidders according to an absolutely continuous distribution, denoted by F . The assumption that F is absolutely continuous helps simplify some of the expressions, but all results go through essentially unchanged if we relax this assumption.

The seller can choose how much information each bidder will have about his own value v_i . An information structure is defined by

$$s_i : \mathbb{R}_+ \rightarrow \Delta \mathbb{R}_+,$$

where $s_i(v_i)$ is the signal observed by bidder i when his value is v_i . After observing s_i , the bidder forms his beliefs about his value. An agent's expected value is denoted by

$$w_i \triangleq \mathbb{E}[v_i | s_i].$$

We denote by G_i the distribution of expected valuations. The definition of the information structure implicitly imposes two restrictions. First, each bidder observes only information about his own value as s_i takes as an argument v_i only (instead of $(v_1, \dots, v_N)$). Second, there is no common source of randomization in the signals. Hence, the signals will be independently distributed across bidders. Finally, we additionally assume that the seller is restricted to symmetric information structures—that is, $s_i(\cdot) = s_j(\cdot)$.

The objective of the seller is to maximize revenue. Since bidders are competing in a second-price auction, it is a dominant strategy to bid their expected value. Hence, revenue is equal to the second-highest expected value across bidders. We denote the k th-highest value by $w_{(k)}$. The objective of the seller is to solve

$$(1) \quad R \triangleq \max_{\{s: \mathbb{R} \rightarrow \Delta \mathbb{R}\}} \mathbb{E}[w_{(2)}].$$

II. Optimal Information Structure

Since the expected revenue is equal to the expectation of second-highest value, the distribution of expected valuations generated by the signal is a sufficient statistic to compute the seller's expected revenue. Hence, instead of studying explicitly the information structure chosen by the seller, we frequently refer to the distribution of expected valuations generated by the signals, which we denote by G .

The second-order statistic of N symmetrically and independently distributed random variables is distributed according to

$$(2) \quad \Pr(w_{(2)} \leq t) = NG^{N-1}(t)(1 - G(t)) + G^N(t).$$

The expected revenue of the auctioneer is therefore

$$\mathbb{E}[w_{(2)}] = \int_0^\infty t d(NG^{N-1}(t)(1 - G(t)) + G^N(t)).$$

We now characterize the set of feasible distributions G .

By Blackwell (1951), Theorem 5, there exists a signal s that induces a distribution of expected valuations if and only if F is a mean-preserving spread of G . F is defined to be a mean-preserving spread of G if

$$(3) \quad \int_v^\infty F(t) dt \leq \int_v^\infty G(t) dt, \forall v \in \mathbb{R}_+,$$

with equality for $v = 0$. If F is a mean-preserving spread of G , we write $F \prec G$.

We can now express the seller's problem as maximizing revenue by choosing a distribution G subject to a mean-preserving constraint. The choice of the optimal information structure can be written as the following maximization problem:

$$(4) \quad R = \max_G \int_0^\infty t d(NG^{N-1}(t)(1 - G(t)) + G^N(t)),$$

subject to $F \prec G$.

This problem consists of maximizing over feasible distributions of expected valuations. However, the objective function is nonlinear in the probability (or density) of the optimization variable G . Moreover, the nonlinearity cannot be confined to be either concave or convex on G .

The key step in our argument comes from a change of variables, rewriting the above in terms of the quantile q of the second-order statistic. We denote by $S_N(q)$ the cumulative distribution function of the quantile of the second-highest value:

$$S_N(q) \triangleq \Pr(G(w_{(2)}) \leq q).$$

We index by N to highlight the dependence on the number of bidders. We observe that $S_N(q)$ is given by

$$S_N(q) = Nq^{N-1}(1 - q) + q^N.$$

The distribution S_N , which is the quantile distribution of the second-order statistic $w_{(2)}$, is independent of the underlying distribution G . Just as the quantile of any random variable is uniformly distributed, the quantile of the second-order statistic of N symmetric independent random variables is distributed according to S_N for any underlying distribution G . Hence, the revenue can be computed by taking the expectation over quantiles using measure $S_N(q)$: the revenue given the quantile of the second-order statistic is G^{-1} . Thus, the maximization problem (4) can be transformed into

$$(5) \quad \max_{G^{-1}} \int_0^1 S'_N(q) G^{-1}(q) dq,$$

subject to $G^{-1} \prec F^{-1}$.

The corresponding constraint states that the seller can choose any distribution of expected valuations whose quantile function G^{-1} is a mean-preserving spread of the quantile function F^{-1} of the initial distribution of valuations. This uses a well-known property of the distribution function (see Shaked and Shanthikumar 2007; Chapter 3) stating that $F \prec G$ if and only if $G^{-1} \prec F^{-1}$. Hence, we have a linear (in G^{-1}) maximization problem subject to a majorization constraint, which will allow us to solve the problem with known methods.

THEOREM 1 (Optimal Information Structure):

1. *The unique optimal symmetric information structure is given by*

$$(6) \quad s(v_i) = \begin{cases} v_i, & \text{if } F(v_i) < q_N^*; \\ \mathbb{E}[v_i | F(v_i) \geq q_N^*], & \text{if } F(v_i) \geq q_N^*. \end{cases}$$

where the critical quantile $q_N^* \in [0, 1)$ is independent of F .

2. *The critical quantile satisfies $q_2^* = 0$; q_N^* is increasing in N ; $q_N^* \rightarrow 1$ as $N \rightarrow \infty$; and for each $N \geq 3$, q_N^* is the unique solution in $(0, 1)$ to the following polynomial equation of degree N :*

$$(7) \quad S'_N(q)(1 - q) = 1 - S_N(q).$$

Thus, the optimal information structure is to reveal the value of all those bidders who have a value lower than some threshold determined by a fixed quantile q_N^* and otherwise reveal no information beyond the fact that the value is above the threshold. The threshold in terms of the value is given by $F^{-1}(q_N^*)$, but the quantile q_N^* is independent of the distribution F of valuations.

The optimal information structure thus supports more competition at the top of the distribution at the expense of an efficient allocation. The information structure fails to distinguish in the allocation between any two valuations that are in the upper tail of the distribution $[F^{-1}(q_N^*), \infty)$. The benefit accrues through more competitive bids among the high-value bidders. Namely, if the second-highest bid is in the above upper interval, then its competitive bid matches exactly the bid of the winning bid and thus the information rent of the winning bidder is depressed considerably with a corresponding gain in the revenue for the seller.

To prove Theorem 1, we use Proposition 2 of Kleiner, Moldovanu, and Strack (2021). It gives necessary and sufficient conditions under which an extreme point G^{-1} in the set of mean-preserving distributions F^{-1} is optimal. The proposition uses the fact that the objective is a linear functional, and the constraint set is defined by majorization and monotonicity, which are indeed the defining conditions of our maximization problem (5).¹

¹ The proof of Theorem 1 in Sorokin and Winter (2021) was self-contained and did not refer to the verification result of Kleiner, Moldovanu, and Strack (2021).

PROPOSITION 1 (Kleiner, Moldovanu, and Strack 2021; Proposition 2): *Let G^{-1} be such that for some countable collection of intervals $\{[\underline{x}_i, \bar{x}_i) \mid i \in I\}$,*

$$G^{-1}(q) = \begin{cases} F^{-1}(q), & \text{if } q \notin \cup_{i \in I} [\underline{x}_i, \bar{x}_i); \\ \frac{\int_{\underline{x}_i}^{\bar{x}_i} F^{-1}(t) dt}{\bar{x}_i - \underline{x}_i}, & \text{if } q \in [\underline{x}_i, \bar{x}_i). \end{cases}$$

If $\text{conv} S_N$ is affine on $[\underline{x}_i, \bar{x}_i)$ for each $i \in I$ and if $\text{conv} S_N = S_N$ otherwise, then G^{-1} solves problem (5). Moreover, if F^{-1} is strictly increasing, the converse holds.

Here, $\text{conv} S_N$ is the convexification of S_N —that is, the largest convex function that is smaller than S_N . With this result we can prove our main result.

PROOF OF THEOREM 1:

The second derivative of the distribution S_N of the quantile of the second-order statistic is given by

$$S_N''(q) = q^{N-3}(N-1)N(N-2-q(N-1)).$$

Hence, $S_N(q)$ is concave if and only if

$$q \geq (N-2)/(N-1)$$

and convex otherwise. Thus, the convex hull of S_N for $N \geq 3$ is

$$\text{conv} S_N(q) = \begin{cases} S_N(q), & \text{if } q \leq q_N^*; \\ S_N'(q_N^*)(q - q_N^*) + S(q_N^*), & \text{otherwise,} \end{cases}$$

where q_N^* is defined as in (7) for $N \geq 3$. For $N = 2$, we have $\text{conv} S_2(q) = q$ and define $q_2^* = 0$. Now let G^{-1} be given by

$$(8) \quad G^{-1}(q) = \begin{cases} F^{-1}(q), & q < q_N^*; \\ \frac{\int_{q_N^*}^1 F^{-1}(t) dt}{1 - q_N^*}, & q \in [q_N^*, 1]. \end{cases}$$

Then G^{-1} satisfies all the assumptions of Proposition 1, so it is the unique optimal solution to (5). For all valuations below $F^{-1}(q_N^*)$, the distribution over expected valuations is the same as that of the real valuations. Hence, types below $F^{-1}(q_N^*)$ know their own values. On the other hand, for valuations above $F^{-1}(q_N^*)$, the distribution over expected valuations is a mass point at the expected value conditional on being above $F^{-1}(q_N^*)$. Hence, it is clear that this distribution is induced by information structure (6).

To check that q_N^* is strictly increasing in N , we define

$$\psi(q, N) \triangleq S_N'(q)(1 - q) - (1 - S_N(q)).$$

By definition, $\psi(q_N^*, N) = 0$. We now note that

$$\psi(q, N+1) - \psi(q, N) = N(q-1)^2(N(q-1)+1)q^{N-2},$$

so $\psi(q, N+1) - \psi(q, N) \geq 0$ if and only if $q \geq (N-1)/N$. As previously argued, $q_N^* < (N-2)/(N-1)$, so $q_N^* < (N-1)/N$, which implies that

$$(9) \quad \psi(q_N^*, N+1) < 0.$$

We also have that $\psi(0, N) = -1$ and $\psi(1 - \varepsilon, N) > 0$ for ε small enough, where the last part can be verified by noting that

$$\psi(1, N) = \frac{\partial \psi(1, N)}{\partial q} = 0 \quad \text{and} \quad \frac{\partial^2 \psi(1, N)}{\partial q^2} = N(N-1) > 0.$$

As previously argued, $\psi(q, N+1)$ has a unique root in $(0, 1)$, so (9) implies that $q_N^* < q_{N+1}^*$.

Finally, if N diverges to infinity and $\lim_{N \rightarrow \infty} q_N^* < 1$, then in the limit we would have that $S_N(q_N^*), S'_N(q_N^*) \rightarrow 0$. So (7) would not be satisfied. We thus must have that $\lim_{N \rightarrow \infty} q_N^* = 1$. ■

The gains from the optimal information structure will depend on the distribution of values, ranging from no gains to a revenue that can be arbitrarily larger than the expected revenue generated under complete information (see Bergemann et al. 2021). As the number of bidders grows large, the gains will remain substantive as long as the distribution of values has fat tails.

We restricted attention to the optimal symmetric information structure. While we do not have a general result showing that the optimal information structure is always symmetric, the symmetric information structure is indeed the unique optimal information structure when there are two or three bidders (see Board 2009 for the case $N = 2$ and Bergemann, Heumann, and Morris 2021 for the case $N = 3$).

Critical Quantile q^ .*—We now provide some intuition for the determination of the critical quantile. In particular, we show that it must be given by equation (7) as long as information takes the qualitative form given by Theorem 1. Suppose that we fix a quantile threshold q and write $v = F^{-1}(q)$ for the corresponding value. Now what happens to revenue if we decrease the threshold by dq ?

With probability $S'_N(q) dq$, the second-highest bid was not in the pooling zone before the decrease but is now after the decrease. The resulting revenue increases by

$$(10) \quad (\mathbb{E}[v_i | v_i \geq v] - v),$$

where the conditional expectation $\mathbb{E}[v_i | v_i \geq v]$ represents the value (and bid) in the pooling interval and v the value (and bid) before being included in the interval. This is the benefit of a marginally lower threshold. With probability $1 - S_N(q)$, the second-highest bid was already in the pooling zone before the decrease. The cost of a lower threshold is then a loss of revenue from the inframarginal bidders:

$$(11) \quad \frac{d\mathbb{E}[v_i | v_i \geq v]}{dq} = \frac{1}{1-q} (\mathbb{E}[v_i | v_i \geq v] - v).$$

Hence, the revenue gain is proportional to $S'_N(q)$ and the revenue loss is proportional to $(1 - S_N(q))/(1 - q)$. Equating gains and losses, we get (7), as the term

$(\mathbb{E}[v_i | v_i \geq v] - v) dq$ appears on both sides and hence cancels out. Furthermore, among the class of upper-pooling information structure, revenue is quasi-concave and single-peaked in the threshold q (see Sorokin and Winter 2021).

A notable implication of this marginal argument is that the critical quantile q^* is indeed independent of the distribution $F(v)$ of values, or indeed any other function of the distribution such as the inverse hazard rate $(1 - F(v))/f(v)$ that typically appears in the virtual utility. The absence of any distributional dependence becomes evident in the above marginal argument. The size of the pool is determined by the marginal probability gain and the inframarginal probability loss, each of which can be expressed solely in terms of the distribution $S_N(q)$ of the quantile of the second-order statistic. Moreover, the corresponding revenue gain or loss is proportional to $(\mathbb{E}[v_i | v_i \geq v] - v)$, as expressed by (10) and (11). Thus, the terms where we would expect the distributional information to enter drop out completely. The fact that the optimal quantile is independent of any distributional information regarding F is distinct from the fact that the quantile distribution of the second-order statistic is independent of F . The former is a consequence of the revenue optimization; the latter is a purely statistical property.

The optimal quantile threshold q^* is therefore independent of the distribution of values and only depends on the number of bidders. This qualitative result on the optimal information policy is exactly the opposite of the optimal reserve price policy, which is independent of the number of bidders but varies with the distribution.

Competition at the Top.—We can gain further intuition about the optimal policy by examining the degree of competition at the top of the bid distribution. We will show that the expected number of bidders above the threshold lies in a narrow range around two.

The number of bidders that have the highest expected value follows a binomial distribution, denoted conventionally by $B(N, 1 - q_N^*)$. The expected number of bidders at the top of the bid distribution is then given by²

$$\rho_N^* \triangleq N(1 - q_N^*).$$

Now $\rho_2^* = 2$, since $q_2^* = 0$. To characterize ρ_N^* for $N \geq 3$, we can substitute $q_N^* = 1 - \rho_N^*/N$ into (7) to get an N th-degree polynomial characterizing ρ_N^* :

$$(12) \quad \frac{\rho_N^*}{N} S_N' \left(1 - \frac{\rho_N^*}{N} \right) = 1 - S_N \left(1 - \frac{\rho_N^*}{N} \right).$$

One can numerically verify that ρ_N^* is decreasing in $N \geq 3$, with $\rho_3^* = 2.25$ and $\rho_N^* \downarrow \rho_\infty^* \approx 1.793$ as $N \rightarrow \infty$. The limit value ρ_∞^* can be solved analytically, as equation (12) reduces to

$$(\rho_\infty^*)^2 + \rho_\infty^* + 1 = e^{\rho_\infty^*},$$

²This is the expression for the mean of any binomial distribution (see, for example, Papoulis and Pillai 2002).

as $N \rightarrow \infty$, whose solution is $\rho_\infty^* \approx 1.793$. Furthermore, as $N \rightarrow \infty$, the binomial distribution $B(N, 1 - q_N^*)$ converges to a Poisson distribution with parameter ρ_∞^* by the Poisson limit theorem (see Papoulis and Pillai 2002). Hence, the optimal information policy always keeps “about” two bidders above the threshold, which is the key to extracting bidder surplus.

Comparison to Bayesian Persuasion.—A brief contrast with a comparable Bayesian persuasion model may be instructive. In the continuous action and continuous state version as analyzed by Dworczak and Martini (2019), for example, the objective function is given by a nonlinear evaluation $u(x)$ of an outcome x and density $dG(x)$, thus a functional linear in probability:

$$\max_G \int_0^1 u(x) dG(x), \text{ subject to } F \prec G.$$

Our original maximization problem (4) did not take this form, as it was nonlinear in the density $dG(t)$. However, we reformulated the problem to one that is linear in the new optimization variable G^{-1} , changing the direction of the majorization constraint. Thus, we do not restrict G to be a mean-preserving contraction of F anymore as it is common in Bayesian persuasion; rather, we require that G^{-1} must be a mean-preserving spread of F^{-1} . For this problem, the convexification of S_N was key to identifying the optimal information structure. The information structure (6) that emerges here is sometimes referred to as “upper censorship” in the Bayesian persuasion literature, as it pools all the states above a threshold and reveals all the states below the threshold (Alonso and Câmara 2016; Kolotilin, Mylovannov, and Zapechelnyuk 2022).³

III. Variations of the Second-Price Auction

We gave a description of the optimal information structure in the standard second-price auction. We now explore the nature of the optimal information structure in significant variations of the second-price auction. First, we consider the optimal information structure in the presence of a reserve price r . Second, we consider a generalization of the second-price auction—namely, the $(K + 1)$ st-price auction where each bidder has unit demand but the seller can offer K homogenous units at a uniform price equal to the $K + 1$ highest bid. And third, we investigate the nature of the optimal information structure when the objective of the auction is different from the revenue of the seller—for example, the surplus of the bidders or the social surplus.

Reserve Price.—We now consider the second-price auction with a common reserve price $r > 0$ for all bidders. We keep the rest of the model as presented in Section I unchanged. We then derive the optimal information structure *given* the

³There are a number of recent papers that use the ironing technique as introduced by Mussa and Rosen (1978) and Myerson (1979) to solve problems in optimal pricing and optimal gerrymandering—see, for example, Loertscher and Muir (2022); Ashlagi, Monachou, and Nikzad (2021); Dworczak, Kominers, and Akbarpour (2021); Kang (2020); and Kolotilin and Wolitzky (2020), respectively. Kleiner, Moldovanu, and Strack (2021) provide an elegant formulation of problems where majorization is added to the ironing problem that we can apply directly.

auction format of a second-price auction with reserve price r . In particular, we do not attempt to solve jointly for the optimal reserve price and information structure.

The seller's problem regarding the choice of the information structure is now given by

$$(13) \quad R \triangleq \max_{\{s: \mathbb{R} \rightarrow \Delta \mathbb{R}\}} \left\{ \mathbb{P}(w_{(1)} \geq r, w_{(2)} < r) r + \mathbb{P}(w_{(2)} \geq r) \mathbb{E}[w_{(2)} | w_{(2)} \geq r] \right\}.$$

In other words, if $w_{(1)} < r$, then the object is not sold; if $w_{(2)} < r$ and $w_{(1)} \geq r$, then the object is sold to the bidder with the highest expected value at price r ; if $w_{(2)} \geq r$, then the object is sold to the bidder with the highest expected value at price $w_{(2)}$.

As before, the seller's problem is an optimization over the feasible distributions G of expected valuations. We denote the left limit of G at $v = r$ as

$$q_r \triangleq \lim_{v \uparrow r} G(v).$$

We can now write the seller's problem as follows:

$$R = \max_{G^{-1}} \left\{ N q_r^{N-1} (1 - q_r) r + \int_{q_r}^1 S'_N(q) G^{-1}(q) dq \right\},$$

subject to $G^{-1} \prec F^{-1}$.

This expresses the seller's problem as an optimization problem over quantiles G^{-1} , subject to G^{-1} being a mean-preserving spread of F^{-1} .

In the standard second-price auction, a positive reserve price r provides an upper bound on the information rent and thus a lower bound on the revenue that the seller can get in any sale of the object. The reserve price r maintains this beneficial property in the presence of information design. But the reserve price r gives the seller now an additional reason to pool the values of some bidders. Namely, the seller can create an intermediate pool of values with an expectation equal to the reserve price r . The new pool around r allows the seller to simultaneously increase the probability of a sale from low and intermediate values and maintain an upper bound on the information rent of the high-value bidders.

We denote the highest value in the support of the distribution G by $\bar{v}_G$.

PROPOSITION 2 (Optimal Information Structure with Reserve Price): *Given a reserve price r , an optimal distribution of expected valuations is given by*

$$(14) \quad G^{-1}(q) = \begin{cases} F^{-1}(q), & \text{if } q \in [0, q_1) \cup (q_2, q_3]; \\ r, & \text{if } q \in (q_1, q_2]; \\ \bar{v}_G, & \text{if } q \in (q_3, 1]; \end{cases}$$

for some quantiles $q_1 \leq q_2 \leq q_3$ and $F^{-1}(q_1) \leq r \leq F^{-1}(q_2) \leq F^{-1}(q_3) \leq \bar{v}_G$.

The proposition states that there will be two full-disclosure regions $[0, q_1) \cup [q_2, q_3)$, one pooling interval at r and a second pooling interval at $\bar{v}_G$ (which is the highest value in the support of expected valuations). The inequalities for the threshold quantiles are weak, $q_1 \leq q_2 \leq q_3$, and thus some intervals may not be present in the optimal information structure.⁴

The main novelty is the introduction of a pooling region at r . The reason for the additional pooling region is that the seller's expected revenue is discontinuous in the expected values around the threshold, so bidders who are initially marginally below the reserve price r are included in the interval $[q_1, q_2)$, so that they bid the reserve price. The full-disclosure region $[0, q_1)$ contains all values that do not buy the object.

PROOF OF PROPOSITION 2:

We fix an optimal information structure G^* and denote the left limit of G^* at r by q_r^* . We can then write the seller's problem as follows:

$$(15) \quad R = \max_{G^{-1}} N(q_r^*)^{N-1} (1 - q_r^*) r + \int_{q_r^*}^1 S'_N(q) G^{-1}(q) dq,$$

subject to

$$(16) \quad G^{-1} \prec F^{-1} \text{ and } G^{-1}(q_r^*) = r.$$

Here we added the constraint that the probability that a bidder's expected value is below the reserve price is q_r^* . By construction this constraint is satisfied by the optimal mechanism.

We now note that at q_r^* , constraint (16) is binding—that is,

$$(17) \quad \int_{q_r^*}^1 G^{*-1}(t) dt = \int_{q_r^*}^1 F^{-1}(t) dt.$$

Otherwise, one could increase the probability that a bidder's value is above r while keeping the distribution of values above quantile q_r^* the same as in G^{*-1} . In other words, we could find $\hat{G}^{-1}(q)$, with $\hat{q}_r < q_r^*$ and $\hat{G}^{-1}(q) = G^{*-1}(q)$ for all $q \geq q_r^*$. This would generate higher revenue, which is a contradiction to G^{*-1} being optimal.

Since $G^{*-1}(q) < r$ for all $q < q_r^*$ and (17) is satisfied, we can assume without loss of generality that $G^{*-1}(q) = F^{-1}(q)$ for all $q < q_r^*$. To verify this, note that we can always modify G^{*-1} in this way and it would continue to be non-decreasing and satisfy $G^{*-1} \prec F^{-1}$. Hence, we focus on finding the properties of G^{*-1} for quantiles above q_r^* . We can write constraint (16) as follows:

$$(18) \quad \int_q^1 G^{-1}(q) dq \leq \int_q^1 F^{-1}(q) dq \text{ for all } q \geq q_r^*$$

and with equality for $q = q_r^*$ and $G^{-1}(q_r^*) = r$.

⁴This is largely a function of the given level of the reserve price r . Already with the uniform distribution on the unit interval—depending on r —any of the following three configurations can arise: $q_1 < q_2 < q_3 = q_N^*$, $q_1 < q_2 = q_3 = q_N^*$, or $q_N^* < q_1 = q_2 = q_3$.

That is, we can write the constraint completely in terms of quantiles above q_r^* . So (18) is a majorization on a restricted domain $[q_r^*, 1]$ plus a constraint determining $G^{-1}(q_r^*) = r$.

We now define q_c^* implicitly as follows:

$$\frac{\int_{q_r^*}^{q_c^*} F^{-1}(t) dt}{q_c^* - q_r^*} \triangleq r$$

and define the following information structure:

$$\tilde{F}^{-1}(q) = \begin{cases} r & q \in [q_r^*, q_c^*]; \\ F^{-1}(q) & q \notin [q_r^*, q_c^*]. \end{cases}$$

That is, $\tilde{F}$ pools values in an interval $[F^{-1}(q_r^*), F^{-1}(q_c^*)]$ so that they induce an expected value r (and offer full disclosure otherwise).

We now note that a quantile function satisfying $G^{-1}(q) = F^{-1}(q)$ for all $q \leq q_r^*$ satisfies (18) if and only if $G^{-1} \prec \tilde{F}^{-1}$. Hence, we can write (15) as a maximization problem subject to a majorization constraint (as in (5)) with the following modifications: the optimization is over the domain $[q_r^*, 1]$, and it is using distribution $\tilde{F}$ instead of F . We can thus find a solution following the same procedure as in Theorem 1 (and we would obtain exactly the same problem if we were to rescale the domain $[q_r^*, 1]$ and replace F with $\tilde{F}$). We recover (14) by setting

$$(q_1, q_2, q_3) = \begin{cases} (q_r^*, q_r^*, q_r^*), & \text{if } q_N^* \leq q_r^*; \\ (q_r^*, q_c^*, q_c^*), & \text{if } q_r^* \leq q_N^* < q_c^*; \\ (q_r^*, q_c^*, q_N^*), & \text{if } q_c^* \leq q_N^*. \end{cases}$$

Note that q_r^* and q_c^* also depend on N , but we did not add the additional subindices to simplify the notation. ■

$(K + 1)$ st-Price Auction.—An important generalization of the second-price auction is the sale of K identical and homogeneous objects when each of the N bidders has unit demand. The seller then offers the objects at a uniform price equal to the $K + 1$ highest bid (assuming $K < N$). Each bidder still has a dominant strategy to bid truthfully, and the revenue can be described by the $(K + 1)$ st-order statistic.

We are now interested in characterizing the optimal information structure for the $(K + 1)$ st-price auction. The seller's revenue is

$$R = \max_{\{s: \mathbb{R} \rightarrow \Delta \mathbb{R}\}} K \cdot \mathbb{E}[w_{(K+1)}].$$

The distribution of the $(K + 1)$ st-order statistic is given by

$$\mathbb{P}(w_{(K+1)} \leq v) = \sum_{k=0}^K \binom{N}{k} G^{N-k}(v) (1 - G(v))^k.$$

The optimal choice of the information structure can then proceed as before except that we are tracking the $(K + 1)$ st- rather than the second-order statistic. With the same change in variables, from values to quantiles, we can then write the revenue as follows:

$$R = \int_0^1 \left(\sum_{k=0}^K \binom{N}{k} q^{N-k(v)} (1-q)^k \right) G^{-1}(q) dq.$$

Fortunately, the order statistic for all $K \geq 2$ share the same qualitative structure in terms of the curvature. Namely, the K th-order statistic expressed in quantile q changes its curvature only once, and from convex to concave. It follows that we obtain a similar result to Theorem 1 for the $(K + 1)$ st-price auction:

COROLLARY 1 ($(K + 1)$ st-Price Auction): *The unique optimal symmetric information structure is given by*

$$(19) \quad s(v_i) = \begin{cases} v_i, & \text{if } F(v_i) < q_{N,K}^*; \\ \mathbb{E}[v_i | F(v_i) \geq q_{N,K}^*], & \text{if } F(v_i) \geq q_{N,K}^*, \end{cases}$$

where the critical quantile $q_{N,K}^* \in [0, 1)$ is independent of F .

Similar to Theorem 1, the critical threshold is determined independent of the underlying distribution, but dependent on the number of bidders and now the number of objects.

Alternative Welfare Objectives.—We analyzed how information disclosure can maximize the revenue of the seller. But we might be interested how different objective functions would influence the optimal disclosure policy. For example, we might ask what is the social-surplus-maximizing information disclosure policy or what is the bidder optimal information disclosure policy, or conversely what is the revenue-minimizing or bidder-surplus-minimizing information policy?

In the second-price auction, these objectives can all be described by weighted sums of the first-order and second-order statistic. For example, the sum of the bidders' surplus is the expectation of the difference of the first-order statistic and the second-order statistic. It turns out that the difference of first- and second-order statistic has exactly the reverse curvature properties than the second-order statistic—namely, first concave and then convex. An implication of our earlier argument is then that the bidders' optimal information structure is a lower censored information structure.

COROLLARY 2 (Bidder Optimal Information Structure): *The bidder optimal symmetric information structure is given by*

$$(20) \quad s(v_i) = \begin{cases} \mathbb{E}[v_i | F(v_i) \leq q], & \text{if } F(v_i) < o_N^*; \\ v_i, & \text{if } F(v_i) \geq o_N^*, \end{cases}$$

where the critical quantile $o_N^* = (N - 2)/(N - 1) \in [0, 1)$ is independent of F .

In fact, we can describe the entire range of payoff outcomes for the bidders and the seller by computing the optimal information policy for the weighted sum of the expectation of the first-order statistic and the second-order statistic. The objective is now to maximize the expectation of a linear combination of the first-order and second-order statistic:

$$(21) \quad m\mathbb{E}[(1 - |\lambda|)w_{(2)} + \lambda w_{(1)}],$$

where $\lambda \in [-1, 1]$ and $m \in \{-1, 1\}$. The problem we have studied so far (revenue maximization) corresponds to $(m, \lambda) = (1, 0)$. If we wanted to maximize total surplus, the objective corresponds to $(m, \lambda) = (1, 1)$. Bidders' surplus maximization corresponds to $(m, \lambda) = (-1, -1/2)$.

We can find the optimal information structure following the same steps as in Section II. Throughout this section we assume that $N \geq 3$ (the case $N = 2$ follows after some qualifications). Define

$$W_N(q) \triangleq (1 - |\lambda|)S_N(q) + \lambda q^N,$$

and write the optimization problem as follows:

$$\max_{G^{-1}} \int_0^1 W'_N(q) G^{-1}(q) dq \quad \text{subject to} \quad G^{-1} \prec F^{-1}.$$

The second derivative of W gives us the curvature of the objective function. We can verify that if (i) $\lambda \geq 1/2$ and $m = 1$, then W_N is convex, (ii) $\lambda \geq 1/2$ and $m = -1$, then W_N is concave, (iii) $\lambda \leq 1/2$ and $m = 1$, then W_N is convex for small q and concave for large q , and (iv) $\lambda \leq 1/2$ and $m = -1$, then W_N is concave for small q and convex for large q . The shape of W_N determines the optimal information structure by finding the convex hull of $W_N(q)$.

COROLLARY 3 (Bidder Optimal Information Structure): *If $m = 1$, the optimal information structure is upper censorship, furthermore, is full disclosure if and only if $\lambda \geq 1/2$. If $m = -1$, the optimal information structure is lower censorship, furthermore, is no disclosure if $\lambda \geq 1$.*

Thus, any linear combination of revenue and social surplus is maximized by an upper or lower censored information structure. Figure 1 illustrates the set of feasible pairs of bidders' surplus and revenue. Any point can be attained by an information structure that solves (21) for some (λ, m) and the slope of the curve is given by $-\lambda/(1 - |\lambda| + \lambda)$. Furthermore, any payoff pair that is attainable by some information structure must be inside these two curves. The red curve illustrates the set of payoff pairs that can be attained by upper censored ($m = 1$). The maximum revenue is attained by an upper censored information policy. The blue curve illustrates the set of payoff pairs that can be attained by lower censorship ($m = -1$). The bidder optimal information structure is a lower censored information structure. Social surplus is maximized under complete information; the bidder surplus is *minimized* under no information (in this case, bidder surplus is zero). Complete information




FIGURE 1. EXPECTED REVENUE AND BIDDERS' SURPLUS ACROSS INFORMATION STRUCTURE, $N = 5$ AND UNIFORM DISTRIBUTION OF VALUES

and no information are the two information structures that are both upper and lower censorship.

IV. The Market for Impressions and Information Design

We studied the revenue-maximizing information structure in a second-price auction with and without reserve price. In the introduction, we mentioned the market for impressions in digital advertising as our main motivation. We conclude with a brief discussion how we may translate our current results to a market for impressions with two-sided information. This translation recasts the optimal information design as bidding mechanisms in the world of digital advertising. The choice of the optimal information structure can then be interpreted in terms of the information policy of the publisher who matches the viewer with the advertisers.

The seller (a *publisher* or *intermediary*) uses an auction platform to sell the attention (“eyeball”) of the viewer to competing advertisers. The viewer is thus the object of the auction. The viewers are typically heterogeneous in many attributes: their demographic characteristics, their preferences, their (past) shopping behavior, their browsing history, and other aspects, observable and unobservable. The advertisers display a corresponding degree of heterogeneity in their willingness to pay for a match between their advertisement and a specific viewer. The private (and the social) value of any particular match is jointly determined by the viewer’s attributes and the advertiser’s preferences for those attributes. In the presence of this heterogeneity on both sides of the match, viewer and advertiser, internet advertising has moved toward targeted advertising to join the information. The auction can therefore support highly targeted advertising that may increase the social efficiency in the match formation between viewer and advertiser.

Formally, the viewer may have attributes $x \in X \subset \mathbb{R}^J$ distributed according to F_x . Each advertiser i has a preference for the attributes described by $y \in Y \subset \mathbb{R}^J$, distributed according to F_y , identically and independently distributed across advertisers.

An impression is a match between an advertiser and a viewer. The value v_i of advertiser i from attracting a viewer is determined by a function $u : X \times Y \rightarrow \mathbb{R}_+$, such that the value is determined by the joint value of attributes and preferences: $v_i \triangleq u(x, y)$. The distribution of characteristics (x, y) and the value function u induce a distribution of the bidder i 's value v_i , which generates a distribution of values as denoted earlier by F .

We can then analyze bidding algorithms in which the publisher commits to complement the advertiser's information with a signal regarding the match quality and to set the advertiser optimal bid. In turn, the advertiser submits his preference y (and thus a description of the attributes he cares about). The central aspect of the bidding algorithms is that the publisher complements the advertiser's private information y with information about the viewer's attribute x that is unknown to the advertiser. Bergemann, Heumann, and Morris (2021) provide two algorithms in which the optimal information structure under Theorem 1 can either be exactly or approximately implemented while satisfying the incentive compatibility conditions of the privately informed bidder.

REFERENCES

- Alonso, Ricardo, and Odilon Câmara. 2016. "Persuading Voters." *American Economic Review* 106 (11): 3590–605.
- Ashlagi, Itai, Faidra Monachou, and Afshin Nikzad. 2021. "Optimal Dynamic Allocation: Simplicity through Information Design." In *EC '21: Proceedings of the 22nd ACM Conference on Economics and Computation*, 101–102. New York: Association for Computing Machinery.
- Bergemann, Dirk, Tibor Heumann, and Stephen Morris. 2021. "Selling Impressions: Efficiency vs. Competition." Cowles Foundation Discussion Paper 2291.
- Bergemann, Dirk, Tibor Heumann, Stephen Morris, Constantine Sorokin, and Eyal Winter. 2021. "Selling Impressions: Efficiency vs. Competition." Cowles Foundation Discussion Paper 2300.
- Bergemann, Dirk, and Martin Pesendorfer. 2007. "Information Structures in Optimal Auctions." *Journal of Economic Theory* 137 (1): 580–609.
- Blackwell, David. 1951. "Comparison of Experiments." In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, ed. Jerzy Neyman, 93–102. Oakland, CA: University of California Press.
- Board, Simon. 2009. "Revealing Information in Auctions: The Allocation Effect." *Economic Theory* 38 (1): 125–35.
- Bourreau, Marc, Bernard Caillaud, and Romain de Nijis. 2018. "The Value of Consumer Data in Online Advertising." *Review of Network Economics* 16 (3): 269–89.
- Dworzak, Piotr, Scott Duke Kominers, and Mohammad Akbarpour. 2021. "Redistribution through Markets." *Econometrica* 89 (4): 1665–98.
- Dworzak, Piotr, and Giorgio Martini. 2019. "The Simple Economics of Optimal Persuasion." *Journal of Political Economy* 127 (5): 1993–2048.
- Eső, Péter, and Balázs Szentes. 2007. "Optimal Information Disclosure in Auctions and the Handicap Auction." *Review of Economic Studies* 74 (3): 705–31.
- Ganuzza, Juan-José. 2004. "Ignorance Promotes Competition: An Auction Model with Endogenous Private Valuations." *RAND Journal of Economics* 35 (3): 583–98.
- Hummel, Patrick, and R. Preston McAfee. 2016. "When Does Improved Targeting Increase Revenue?" *ACM Transactions on Economics and Computation* 5 (1): 1–29.
- Kang, Zi Yang. 2020. "Markets for Goods with Externalities." Unpublished.
- Kleiner, Andreas, Benny Moldovanu, and Philipp Strack. 2021. "Extreme Points and Majorization: Economic Applications." *Econometrica* 89 (4): 1557–93.

- Kolotilin, Anton, Timofiy Mylovanov, and Andriy Zapechelnyuk.** 2022. "Censorship as Optimal Persuasion." *Theoretical Economics* 17 (2): 561–85.
- Kolotilin, Anton, and Alexander Wolitzky.** 2020. "The Economics of Partisan Gerrymandering." University of New South Wales (UNSW) Economics Working Paper 2020-12.
- Levin, Jonathan, and Paul Milgrom.** 2010. "Online Advertising: Heterogeneity and Conflation in Market Design." *American Economic Review* 100 (2): 603–07.
- Li, Hao, and Xianwen Shi.** 2017. "Discriminatory Information Disclosure." *American Economic Review* 107 (11): 3363–85.
- Loertscher, Simon, and Ellen V. Muir.** 2022. "Monopoly Pricing, Optimal Randomization, and Resale." *Journal of Political Economy* 130 (3): 566–635.
- Milgrom, Paul R., and Robert J. Weber.** 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (5): 1089–122.
- Mussa, Michael, and Sherwin Rosen.** 1978. "Monopoly and Product Quality." *Journal of Economic Theory* 18 (2): 301–17.
- Myerson, Roger B.** 1979. "Incentive Compatibility and the Bargaining Problem." *Econometrica* 47 (1): 61–73.
- Palfrey, Thomas R.** 1983. "Bundling Decisions by a Multiproduct Monopolist with Incomplete Information." *Econometrica* 51 (2): 463–83.
- Papoulis, Athanasios, and S. Unnikrishna Pillai.** 2002. *Probability, Random Variables, and Stochastic Processes*. New York: McGraw-Hill.
- Shaked, Moshe, and J. George Shanthikumar.** 2007. *Stochastic Orders*. New York: Springer.
- Sorokin, Constantine, and Eyal Winter.** 2021. "Pure Information Design in Classic Auctions." Hebrew University of Jerusalem Discussion Paper 742.