CALIBRATED CLICK-THROUGH AUCTIONS

by

Dirk Bergemann, Paul Duetting,
Renato Paes Leme, and Song Zuo

COWLES FOUNDATION PAPER NO. 1771

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

2022

http://cowles.yale.edu/

Calibrated Click-Through Auctions

Dirk Bergemann
dirk.bergemann@yale.edu
Yale University

Paul Duetting
duetting@google.com
Google

Renato Paes Leme
renatoppl@google.com
Google

Song Zuo
szuo@google.com
Google

ABSTRACT

We analyze the optimal information design in a click-through auction with stochastic click-through rates and known valuations per click. The auctioneer takes as given the auction rule of the click-through auction, namely the generalized second-price auction. Yet, the auctioneer can design the information flow regarding the click-through rates among the bidders. We require that the information structure to be calibrated in the learning sense. With this constraint, the auction needs to rank the ads by a product of the value and a *calibrated* prediction of the click-through rates. The task of designing an optimal information structure is thus reduced to the task of designing an optimal calibrated prediction.

We show that in a symmetric setting with uncertainty about the click-through rates, the optimal information structure attains both social efficiency and surplus extraction. The optimal information structure requires private (rather than public) signals to the bidders. It also requires correlated (rather than independent) signals, even when the underlying uncertainty regarding the click-through rates is independent. Beyond symmetric settings, we show that the optimal information structure requires partial information disclosure, and achieves only partial surplus extraction.

CCS CONCEPTS

• **Theory of computation → Algorithmic game theory and mechanism design**.

KEYWORDS

second-price auction, stochastic click-through rates, revenue maximization

ACM Reference Format:
Dirk Bergemann, Paul Duetting, Renato Paes Leme, and Song Zuo. 2022. Calibrated Click-Through Auctions. In *Proceedings of the ACM Web Conference 2022 (WWW '22), April 25–29, 2022, Virtual Event, Lyon, France.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3485447.3512050

1 INTRODUCTION

In the world of digital advertising the allocation mechanism is frequently a pay-per-click auction. Hence, ad systems are a combination of an auction mechanism together with a machine learning model predicting click probabilities. The question of learning click-through rates has been analyzed extensively in terms of accuracy and scalability [e.g., 15, 24]. In this paper we take a different perspective and analyze how different machine learning models influence the performance of the auction in terms of revenue-efficiency trade-offs.

1.1 Our Approach and Results

Our focus in this paper is on how the click-through-rate model influences ranking and pricing, instead of focusing on accuracy alone. We approach this question by keeping the auction mechanics as simple as possible: a single-slot pay-per-click auction where click-through-rates are stochastic, correlated and bidder-dependent. What we will vary instead, will be the choice of the machine learning model. In other words, among different *calibrated predictors* for click-through rates, which of them leads to higher revenue?

We will borrow both the terminology and the technical tools from information design to tackle this question [e.g., 4]. In that context, we will refer to a click-through-rate model as an *information structure* or an *information policy*. We will enforce the constraint that the model is *calibrated* (in the sense of Foster and Vohra [13]). In other words, the expected click-through rate given a prediction of the model is equal to the prediction itself. Among all possible information policies, the complete information policy and the zero information policy are both leading examples, as well as extremal information policies. Under a complete information policy, the seller completely discloses all information to the bidders. It is thus as if the bidders were in a complete information environment. By contrast, in a zero information policy the seller does not disclose any information about the realized click-through rates. In consequence, each bidder acts as if the realized click-through rate is always equal to the ex ante expected click-through rate. These two extremal information policies have dramatically different payoff implications.

With stochastic click-through rates, social efficiency and revenue will depend on which information about the click-through rates is disclosed. With complete information, the resulting allocation is always efficient. But as the competitive position of each bidder can vary across the realized click-through rates, the resulting revenue of the seller can be low due to weak competition. By contrast, with zero information about the click-through rates, the resulting allocation will typically fail to be socially efficient. As the bidding behavior cannot reflect any information about the click-through rates, the socially relevant information fails to be reflected in the auction outcome. Yet, as the bidding reflects only the expected click-through rate, and hence only the mean of the click-through rate

WWW '22, April 25–29, 2022, Virtual Event, Lyon, France.

ACM ISBN 978-1-4503-9096-5/22/04.
https://doi.org/10.1145/3485447.3512050

(and not the higher moments of the click-through distribution), the resulting bids will have zero variance, and thus be more competitive.

Our first result, Proposition 3.1, shows that a statistically independent information structure can never improve the revenue over the no-disclosure information policy. Our second result, Corollary 3.2, shows that with any level of uncertainty in the click-through rates, the seller strictly prefers no-disclosure to full-disclosure. Thus, the seller always favors competition over information disclosure. We then ask whether there exists an improved information policy that can realign social efficiency with the revenue of the seller.

Our main result, Theorem 4.1, establishes that a calibrated and correlated information policy can completely align social efficiency and revenue maximization. In particular, when the common prior distribution over the click-through rates is symmetric across the bidders, we can then explicitly construct an information policy such that the socially efficient allocation is always realized and the competition between the bidder levels the information rent of the bidders to the ex-ante level. Given the symmetry of the common prior distribution, this implies that the bidders compete their residual surplus down to zero.

We also provide an explicit construction of the correlated information structure (or signalling scheme). Interestingly, the optimal information structure is an interior information structure; that is, it is neither zero nor complete information disclosure. The information structure balances two conditions that are necessary to attain the socially efficient allocation while maintaining competition: (i) it provides sufficient information to rank the alternative allocations according to social efficiency, and (ii) it limits the variance in the posterior beliefs of the competing bidders so that their equilibrium bids remain arbitrarily close to support competitive bids.

In the full version, we show that this result extends to a setting with two bidders under a weak notion of ex-ante symmetry: the socially efficient allocation and revenue maximization remain perfectly aligned as long as the expected click-through rate is equalized across bidders, even when the support of the ex-post realized click-through rates can vary across bidders.

Finally, in Theorem 5.1, we show that an interior information structure remains part of the optimal information design with stochastic click-through rates, even when the expected click-through rates across the bidders differ, and therefore one bidder is stronger from an ex-ante perspective. In particular, the optimal information structure releases less information about the winning bidder than about the losing bidder. This suggests that the optimal information structure in an asymmetric auction seeks to strengthen the weak bidder with additional information, relative to the strong bidder.

The optimality of a noisy information structure has significant implications in the world of digital advertising. Since the optimal information structure remains noisy, better click-through-rate predictions–achieved through improved learning–may not necessarily lead to better auction results. Thus, there might be limits to the returns of more elaborate machine learning algorithms to inform the prediction problem.

1.2 Discussion

Throughout this paper we maintain the simplifying assumption that the valuation of each bidder is known. Thus there is complete information regarding the value of each bidder. This setting is commonly adopted in the analysis of sponsored search auctions, see notably Edelman et al. [10] and Varian [25]. In our analysis, this means that we can assume truth-telling by the bidders, and thus the bids of the agents always equal their values. The complete information assumption allows us to focus the analysis entirely on the optimal information policy regarding click-through rates.

A significant next step would be to embed the current analysis into an environment with incomplete information regarding the values of the bidders. We would then be in a setting where both the auctioneer and bidders have private information. A general analysis with two-sided private information remains a wide-open issue, even in a setting with a single bidder. Currently, progress is being made only in specific settings, either binary actions and states or multiplicative separable settings (see Kolotilin et al. [20] and Candogan and Strack [7], respectively).

In our specific setting, we would need to augment the analysis with truth-telling constraints by the bidders. We expect that these incentive constraints would weaken the power of the information design without eliminating it completely. In particular, we expect that the bidders' surplus would then only be partially extracted.

We also maintain the simplifying assumption that there is only a single position to be allocated. This allows us to describe the distribution of the information among two bidders, the winning and losing bidder. With many positions, the resulting information design would have to balance additional constraints that may impose similar restrictions on surplus extraction as the introduction of private information discussed above. Yet, even with a single position, the click-through auction yields distinct outcomes from the second price auction. In particular, in the click-through auction, the revenue (and the bidders' net surplus) are impacted by the correlation between interim expected click-through rates and ex post realized click-through rates. By contrast, in the second price auction, it is only the interim expected rates that matter. It is the richer and subtle interaction between interim and ex post click-through rates that allows the click-through auction to align revenue and socially efficiency more strongly.

A significant advance in our information design problem is to allow for multi-dimensional and private information in a strategic setting. By contrast, the most recent result in the design of optimal information structure requires one-dimensional, or equivalently symmetric, solutions to optimal design (see Kleiner et al. [19] and Bergemann et al. [3]). A general approach to optimal multi-dimensional information design in strategic settings is again a wide-open question. In the current context, we could and do make progress by insights specific to the auction setting.

1.3 Further Related Work

A recent strand of literature in algorithmic mechanism design has considered different aspects of signalling and targeting in ad auctions [e.g., 2, 3, 11, 14, 18]. The main difference to our work is that these works have focused on per-impression auctions, which are very different from the per-click auction considered here. An earlier literature [21, 23] has introduced the idea of "squashing" or "boosting" of click-through rates to optimize auction performance. These papers only consider expected click-through rates and do not insist on calibration. In our analysis, we clearly separate between the auction mechanism, the generalized second price auction, and the information design. The joint attainment of revenue and efficiency is due to the information design. This distinguishes our analysis from the full surplus extraction results of [8, 9]. There, the transfer payments in the mechanism are exactly chosen and tailored to the correlation of the signals to allow for surplus extraction. In particular, with positive probability the payments need to be arbitrarily larger to achieve the results. In the generalized second price auction, the payments are determined independent of the correlation in the signals. Despite these constraints, our positive results show how impactful the design of the information can be to guide the bids and the allocations in the auction.

We provide a more detailed discussion of this and additional related work in Appendix A.

2 MODEL

We will analyze the setting of click-through auctions where bidders are ranked by the product of their value (expressed as a maximum willingness-to-pay per click) and a calibrated prediction of the click-through rates. Our main goal will be to study how to engineer such calibrated prediction to achieve better revenue-efficiency trade-offs. In the language of information design, we will keep the auction format fixed and vary the information structure.

To allow us to focus on the information structure, we will work in the full information model where each bidder $i = 1, \dots, n$ has a fixed and known value $v_i \geq 0$ representing their willingness-to-pay for a click. Our central object of study will be the click-through rates (CTRs): before the auction, a vector $r = (r_1, \dots, r_n) \in [0,1]^n$ will be drawn from a joint prior distribution G, which is a multi-dimensional distribution that may display correlation across the bidders' CTRs.

The click-through rates are known by the auctioneer: typically, the platform is the one building a machine learning model to estimate them. The auctioneer must now decide on a score/signal $s = (s_1, \dots, s_n) \in [0,1]^n$ to rank the bidders. The design space will be to design a joint probability distribution ρ on pairs (r, s) such that the marginal on r is G:

$$\int_{r \in R} dG(r) = \int_{r \in R} \int_s d\rho(s, r), \quad \forall \text{ measurable } R \subseteq [0,1]^n.$$

This joint probability distribution will be referred as the *information structure*. For notational convenience, it will be useful to assume that both G and ρ are discrete distributions with finite support, and hence we can write $g(r)$ for the probability of a given vector r under distribution G and $x(r, s)$ as the probability of a pair (r, s) under distribution ρ. With that notation, the information structure is a function $x : [0,1]^n \times [0,1]^n \to [0,1]$ satisfying:

$$\sum_s x(r, s) = g(r), \quad \forall r \in [0,1]^n.$$

Auction Mechanics. Again for simplicity, we focus on the single slot setting where the goal of the auction is to select a single winner $i^* \in [n]$. The winner, assuming truthful bids, will be selected as the bidder having the largest $s_i v_i$, with a symmetric tie-breaking rule. The winner's cost per click is then:

$$p_{i^*} = \max_{j \neq i^*} \frac{v_j s_j}{s_{i^*}}.$$

The winner only pays when there is a click, which happens with probability r_{i^*}. Hence, the expected revenue from this per-click auction is $r_{i^*} p_{i^*}$.

Note that we are effectively operating a single-item, second-price auction in which truth-telling is a dominant strategy equilibrium. This is true even if the scores depend on the valuations, because in the complete information setting we consider a bidder cannot change signals by submitting a non-truthful bid. We therefore assume truth-telling by the bidders.

Calibration. The auctioneer is restricted to ranking with an estimator of the click-through rates. We require the information structure to be calibrated in the Foster-Vohra sense [13]. An information structure is called *calibrated* if the posterior, given any signal realization s'_i for bidder i, matches with the signal itself, i.e.,

$$\mathbb{E}[r_i | s_i = s'_i] = s'_i. \tag{1}$$

If the CTR and signal space is discrete, then we can write calibration as:

$$\sum_{(r,s); s_i = s'_i} x(r, s) \cdot (r_i - s'_i) = 0, \quad \forall i, s'_i.$$

There are two important examples of calibrated information structures:

- Full-disclosure: where $s_i = r_i$ almost surely.
- No-disclosure: where $s_i = \mathbb{E}[r_i]$ almost surely.

Since the calibration constraint is imposed on every bidder separately, it is possible to create information structures that combine disclosure and no-disclosure. For example, given two bidders, we can consider an information structure where bidder 1 receives only one signal, and the signal is equal to the ex-ante expectation of the click-through rate, thus $s_1 = \mathbb{E}[r_1]$; and bidder 2 receives as many signals as click-through rates, thus $s_2 = r_2$. This forms a calibrated information structure.

Whenever $s_i = \mathbb{E}[r_i]$ we will say that we *fully bundle* bidder i. Whenever $s_i = r_i$ we will say that we *unbundle* bidder i. If neither is the case we will say that we *partially bundle* the bidder.

Calibrated vs. Unbiased Estimator. The current notion of calibrated estimator is related but distinct to the notion of unbiased estimator. An unbiased estimator requires that $\mathbb{E}[s_i | r_i] = r_i$ whereas a calibrated prediction requires $\mathbb{E}[r_i | s_i] = s_i$.

When predictions are used to rank candidates in an auction, calibration is a more useful notion than unbiasedness. Suppose pCTR is the prediction and CTR is the true value of click-through-rates. Then, by using a calibrated model, ranking by pCTR $\cdot$ bid is the same as ranking by $\mathbb{E}[\text{CTR} \cdot \text{bid} \mid \text{pCTR}]$. This is why in practice pCTR's that are generated by a complex ML model are re-calibrated to satisfy the basic requirement that the pCTRs match with the unconditional expectation.

Independence and Correlation. A information structure is called *independent* if signals s_j for $j \neq i$ do not offer additional information on the expectation of r_i beyond s_i. Formally:

$$\mathbb{E}[r_i | s = s'] = \mathbb{E}[r_i | s_i = s'_i], \quad \forall i, \forall s. \tag{2}$$

Both full-disclosure and no-disclosure information structures are independent.

Whenever we do not assume independence, we will say that an information structure is *correlated.* Below, we give an example of a correlated and calibrated information structure. This will also serve as an example of how information structures will be illustrated throughout the paper. Consider the two-bidder setting where CTRs are $r \in \{(1/2, 1/2), (1/2, 1), (1, 1/2), (1, 1)\}$, each with probability $1/4$ (hence r_1 and r_2 are independent). We represent an information structure where rows correspond to pairs of CTRs (r_1, r_2) and columns correspond to pairs of signals (s_1, s_2). Each entry of the matrix will correspond to $x(r, s)$ which is the probability of the event that the CTRs are r and the signals are s. Table 1 shows the "*flipping the square*" structure, which we will discuss in detail in Section 4.1.

Table 1: Flipping the square

r \ s	$(\frac{3}{4}-\epsilon, \frac{3}{4}-\epsilon)$	$(\frac{3}{4}-\epsilon, \frac{3}{4}+\epsilon)$	$(\frac{3}{4}+\epsilon, \frac{3}{4}-\epsilon)$	$(\frac{3}{4}+\epsilon, \frac{3}{4}+\epsilon)$
$(\frac{1}{2}, \frac{1}{2})$	ϵ	0	0	$\frac{1}{4}-\epsilon$
$(\frac{1}{2}, 1)$	0	$\frac{1}{4}$	0	0
$(1, \frac{1}{2})$	0	0	$\frac{1}{4}$	0
$(1, 1)$	$\frac{1}{4}-\epsilon$	0	0	ϵ

One can check that while the calibration constraints (equation (1)) hold, the independence condition (equation (2)) does not. So, this is a calibrated, correlated information structure.

Normalization and Symmetry. We note that it is without loss of generality to normalize all values to $v'_i = v$ while scaling CTRs and scores to $r'_i = v_i r_i / v$ and $s'_i = v_i s_i / v$. Both the allocations and expected payments remain identical after such normalization. We therefore generally assume that $v_1 = \ldots = v_n = 1$. A setting is *symmetric* if (after normalization) random variables $v_1 r_1, \ldots, v_n r_n$ are exchangeable. Whenever symmetry does not hold, we will say that the environment is *asymmetric.*

3 INDEPENDENT INFORMATION STRUCTURES

A first step in the analysis of optimal information design is a focus on independent signals. With independent signals, the signal s_i of each bidder i contains all the information about the CTR r_i that the auctioneer releases before the auction. Thus, bidder i could not learn anything more about the true CTR from any other bidder. In turn, the information that bidder i receives from the auctioneer is the maximal information that is available before the auction to place an informed bid.

3.1 Independent Signals and Two Bidders

We begin the analysis with two bidders and then generalize the insight to many bidders.

Proposition 3.1 (Independent and Calibrated Signalling). *In a two-bidder environment, the expected revenue of an independent and calibrated information structure cannot exceed the one from no-disclosure.*

Proof. We start by computing the expected revenue given signals s_1 and s_2. If $v_1 s_1 \geq v_2 s_2$, then the revenue can be written as:

$$\mathbb{E}\left[r_1 \cdot \frac{v_2 s_2}{s_1} \middle| s_1, s_2\right] = \mathbb{E}\left[r_1 | s_1, s_2\right] \cdot \frac{v_2 s_2}{s_1} = \mathbb{E}\left[r_1 | s_1\right] \cdot \frac{v_2 s_2}{s_1} = s_1 \cdot \frac{v_2 s_2}{s_1} = v_2 s_2,$$

where the second equality follows from the independence of the signaling scheme and the third equality follows from calibration.

Therefore, we can write the expected revenue as:

$$\text{Rev} = \mathbb{E}[\min_{i \in \{1,2\}} v_i s_i] \leq \min_{i \in \{1,2\}} \mathbb{E}[v_i s_i] = \min_{i \in \{1,2\}} v_i \, \mathbb{E}[s_i] = \min_{i \in \{1,2\}} v_i \, \mathbb{E}[r_i],$$

where the first inequality follows from Jensen's inequality and the concavity of the minimum. The last equality follows from calibration. Finally, note that $\min(v_1 \mathbb{E}[r_1], v_2 \mathbb{E}[r_2])$ is the revenue from no-disclosure. □

We have thus shown that full-disclosure can never revenue-dominate no-disclosure. We next show that generically full-disclosure is, in fact, strictly revenue-dominated by no-disclosure.

Corollary 3.2 (Zero vs. Complete Information Disclosure). *In the two-bidder environment, if both $v_1 r_1 > v_2 r_2$ and $v_1 r_1 < v_2 r_2$ occur with positive probability, then the revenue from no-disclosure strictly dominates the revenue from full-disclosure.*

Proof. If both events $v_1 r_1 > v_2 r_2$ and $v_1 r_1 < v_2 r_2$ occur with positive probability, then Jensen's inequality holds with a strict inequality $\mathbb{E}[\min(v_1 r_1, v_2 r_2)] < \min(\mathbb{E}[v_1 r_1], \mathbb{E}[v_2 r_2])$ in the previous proof. □

The argument suggests that the revenue dominance result does not extend to more than two bidders. With more than two bidders, the smaller of the two highest realizations determines the price, and the expectation of the smaller of the two highest is now larger than the unconditional expectation of the second highest click-through rate.

3.2 Independent Signals with Many Bidders

Indeed, the power of independent signalling is much improved in the presence of competition, and we now consider the case of more than two bidders, $n > 2$. We show by example that an independent symmetric-calibration signal can improve the revenue, and thus partial revelation is better than no- or full-disclosure.

Consider a symmetric three-bidder environment with $v_1 = v_2 = v_3 = 1$. For each bidder i let $r_i = 0$ with probability $2/3$, and $r_i = 1$ with probability $1/3$. Moreover, assume that r_1, r_2 and r_3 are independent. It is simple to check that full revelation has revenue $7/27$ and no-revelation has revenue $9/27$. This can be improved by the following signaling scheme with partial bundling:

Table 2: Partial bundling

r_i \ s_i	$s_i = 0$	$s_i = 4/9$
$r_i = 0$	$1/4$	$5/12$
$r_i = 1$	0	$1/3$

The revenue of partial disclosure is $4/9$ whenever at least two of the bidders have the high signal, which happens with probability $1-(1/4)^3-3(1/4)^2(3/4)=27/32$. Hence, the overall revenue of this partial disclosure scheme is $3/8$, which dominates both full-disclosure and no-disclosure.

Instead of trying to optimize for the optimal independent information structure, we will move to the more powerful model of correlated information structures in the next section.

The results here mirror earlier results by Board [6], who considers an ascending auction in a private value setting without click-through rates. In his analysis, he restricts attention to independent signals and establishes that with two bidders, the seller's revenue is smaller *with* than *without* information disclosure. He further shows that in a symmetric model, as the number of bidders become arbitrarily large, complete information disclosure eventually revenue-dominates zero information disclosure.

4 CORRELATED STRUCTURES IN SYMMETRIC ENVIRONMENTS

We obtain a much stronger result if we allow the information structure to be correlated across bidders. The information flow allows influence over the level of competition to some extent. This allows us to conflate auction items and restore some market thickness (see Levin and Milgrom [22]).

4.1 A First Example: Flipping the Square

To showcase the power of correlated information structures, we start with an example where the optimal structure is rather counterintuitive. Consider two bidders with values $v_1=v_2=1$ and independent click-through rates distributed uniformly in $\{\frac{1}{2},1\}$. Or rather: the vector r is uniformly distributed in $\{(\frac{1}{2},\frac{1}{2}),(\frac{1}{2},1),(1,\frac{1}{2}),(1,1)\}$. With an independent signaling scheme, the optimal information structure is no-disclosure, which yields revenue equal to $3/4$. With a correlated signaling scheme, one can obtain arbitrarily close to $7/8$ revenue, which is optimal since it corresponds to the welfare $\mathbb{E}[\max_i v_i r_i]$ of the optimal allocation.

The information structure is the "flipping the square" structure described in Table 1. To see that this information structure is calibrated, observe that $s_1=3/4-\epsilon$ with probability $1/2$. This probability event can be decomposed in two: with probability $1/4+\epsilon$ we output this signal with $r_1=1/2$, and with the remaining $1/4-\epsilon$ probability we have $r_1=1$. Hence:

$$\mathbb{E}[r_1|s_1=3/4-\epsilon]=\frac{(1/4+\epsilon)\cdot 1/2+(1/4-\epsilon)\cdot 1}{1/2}=3/4-\epsilon.$$

The counterintuitive nature of this mapping can best be seen when depicted as in Figure 1. We map the CTRs in $\{1/2,1\}$ to two values $\{3/4-\epsilon,3/4+\epsilon\}$ around the mean. The symmetric pairs, $(1/2,1/2)$ and $(1,1)$, are mapped with high probability into symmetric, but order-reversed pairs, $(3/4+\epsilon,3/4+\epsilon)$ and $(3/4-\epsilon,3/4-\epsilon)$, respectively. The calibration is nonetheless achieved by the off-diagonal pairs $(1/2,1)$ and $(1,1/2)$ that are mapped into order-preserving signals with probability 1, $(3/4-\epsilon,3/4+\epsilon)$ and $(3/4+\epsilon,3/4-\epsilon)$, respectively. The ϵ perturbation in the mapping of the diagonal pairs then achieves the ordering of the signals.

Figure 1: Depiction of the "flipping the square" structure (with ϵ-flows omitted).

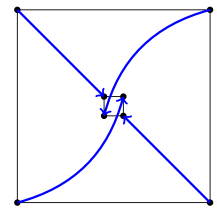

We notice the click-through signals s_i maintain the efficient ranking of the alternatives, and thus guarantee an efficient outcome in the auction. The revenue in the auction is given by:

$$\frac{1}{4}\cdot\frac{1}{2}+\frac{1}{4}\cdot\frac{\frac{3}{4}-\epsilon}{\frac{3}{4}+\epsilon}+\frac{1}{4}\cdot\frac{\frac{3}{4}-\epsilon}{\frac{3}{4}+\epsilon}+\frac{1}{4}=\frac{21-4\epsilon}{24+32\epsilon}=\frac{7}{8}-O(\epsilon),$$

which means that almost the entire surplus is extracted. The auction uses a *uniform* tie-breaking rule, thus allocating the object with equal probability if the signals are equal across the bidders.

4.2 A Second Example: Dispersion Along the Diagonal

The second example maintains symmetry across the bidders but has correlated click-through rates. The resulting information structure is more subtle, reflecting the need to balance information necessary to support an efficient allocation with information to support competition.

We present the construction, which we refer to as "dispersion along the diagonal," for a small number of signals. Our main result in this section (Theorem 4.1) builds on a generalization of this construction to more signals (see Lemma 4.2 and Figure 2).

Consider again two bidders with values $v_1=v_2=1$ and click-through rates either $(1/2,1)$ or $(1,1/2)$, with probability $1/2$ each. The CTRs are thus perfectly negatively correlated and the social surplus is 1. The revenue under full-disclosure would be $1/2$, and under no-disclosure it would be $3/4$. With no-disclosure, the price-per-click is always competitive, as $\mathbb{E}[r_1]/\mathbb{E}[r_2]=1$, but the auction fails to lead to the efficient allocation with probability $1/2$. With the following information flow, we attain a revenue of $0.79>3/4$: This

Table 3: Dispersion along the diagonal

r \ s	$(\frac{6}{10},\frac{6}{10})$	$(\frac{6}{10},\frac{3}{4})$	$(\frac{3}{4},\frac{6}{10})$	$(\frac{3}{4},\frac{15}{16})$	$(\frac{15}{16},\frac{3}{4})$	$(\frac{15}{16},\frac{15}{16})$
$(\frac{1}{2},1)$	$\frac{2}{15}$	$\frac{2}{5}$	0	$\frac{2}{5}$	0	$\frac{1}{15}$
$(1,\frac{1}{2})$	$\frac{2}{15}$	0	$\frac{2}{5}$	0	$\frac{2}{5}$	$\frac{1}{15}$

information flow lowers the probability of an inefficient allocation from $1/2$ to $1/5$ and attains an equilibrium price closer to 1. In particular,

$$\min\left|\frac{r_i}{r_j}\right|=\frac{1}{2}<\frac{4}{5}=\min\left|\frac{s_i}{s_j}\right|.$$

The information flow in this example generates some symmetric click-through *signals* in the absence of symmetric click-through *rates*. The symmetric signals in the presence of asymmetric rates create some inefficiency in the allocation. But the symmetric click-through rates create the basis for signals that are adjacent, in the sense that they are nearby, yet signal the correct ranking of the underlying click-through rates. If we increase the numbers of signals in the construction of the information flow, we can then reduce the revenue loss and bring it arbitrarily close to zero. This is the following content of Lemma 4.2.

4.3 Optimal Information Structure

We can now state and establish the first main result, showing that for any n-bidder symmetric environment it is possible to construct an information structure extracting revenue that is arbitrarily close to the optimal surplus.

Theorem 4.1 (Full Surplus Extraction in Symmetric Environments). *For every symmetric n-bidder environment, there exists a randomized and calibrated correlated information structure whose revenue is arbitrarily close to full surplus extraction.*

As a building block, we will consider the special case of a symmetric environment of two bidders where $v_1 = v_2 = v$ and

$$\Pr[r = (l, h)] = \Pr[r = (h, l)] = 1/2,$$

for two values $0 \leq l < h \leq 1$. We will then reduce the general symmetric case to a composition of information structures for pairs (h, l). The optimal information structure will be to disperse the signals along the diagonal as depicted in Figure 2. The following lemma, Lemma 4.2, is the technical heart of the paper.

Figure 2: Depiction of the "dispersion along the diagonal" structure for Lemma 4.2.

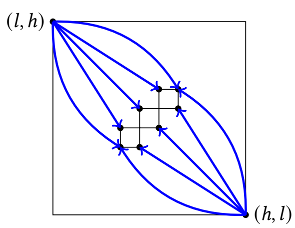

Lemma 4.2 (Dispersion Along the Diagonal). *Consider the symmetric setting of two bidders with normalized values $v_1 = v_2 = v$ where the click-through rate vector is either (l, h) or (h, l), each with probability $1/2$. Then, for every $\epsilon > 0$, there is a calibrated, correlated information structure with revenue $vh - \epsilon$.*

Proof of Lemma 4.2. To prove the lemma, we construct an information structure with a finite set of signals, $S \subset [0, 1]$. The key to this construction is to (i) properly select the signal set S, and (ii) come up with a discretized and calibrated information structure $x(r, s)$ for $r \in \{(l, h), (h, l)\}$ and $s \in S$ that achieves almost optimal revenue.

We consider the following construction with parameters $\delta > 0$ and $x_0 > 0$ to be determined later:

(1) Signal set $S = \{s_{-K}, \ldots, s_0, s_1, \ldots, s_K\}$, where $s_0 = (l + h)/2$, $s_k = s_0 \cdot (1 + \delta)^k$ for $-K \leq k \leq K$,
$$K = \left\lfloor \log_{(1+\delta)} \frac{2h}{l + h} \right\rfloor - 1;$$

(2) $x((l, h), (s_k, s_{k+1})) = x((h, l), (s_{k+1}, s_k)) = x_k$ for $-K \leq k \leq K - 1$, where
$$x_k = \frac{h - s_k}{s_k - l} \cdot x_{k-1} = x_0 \prod_{\kappa=1}^{k} \frac{h - s_\kappa}{s_\kappa - l}, \text{ when } 1 \leq k \leq K - 1,$$
$$x_k = \frac{s_{k+1} - l}{h - s_{k+1}} \cdot x_{k+1} = x_0 \prod_{\kappa=k}^{-1} \frac{s_{\kappa+1} - l}{h - s_{\kappa+1}}, \text{ when } -K \leq k \leq -1;$$

(3) $x((l, h), (s_{-K}, s_{-K})) = x((h, l), (s_{-K}, s_{-K})) = y$, where
$$y = \frac{s_{-K} - l}{l + h - 2s_{-K}} \cdot x_{-K};$$

(4) $x((l, h), (s_K, s_K)) = x((h, l), (s_K, s_K)) = z$, where
$$z = \frac{h - s_K}{2s_K - l - h} \cdot x_{K-1}.$$

In the rest of the proof, we first verify that the construction is a valid calibrated and correlated information structure, then show that by choosing a sufficiently small δ, the revenue is at least $h - \epsilon$.

Step 1: We verify that the signals are valid probabilities, i.e., $S \subset [0, 1]$.

For sufficiently small $\frac{h-l}{3(h+l)} > \delta > 0$, $K - 1 = \left\lfloor \log_{(1+\delta)} \frac{2h}{l+h} \right\rfloor - 1 \geq 1$. For all $-K \leq k \leq K$,

$$s_k \leq s_K = \tfrac{l+h}{2} \cdot (1 + \delta)^K \leq \tfrac{l+h}{2} \cdot \tfrac{2h}{l+h} \cdot \tfrac{1}{1+\delta} = \tfrac{h}{1+\delta} < h \leq 1;$$
$$s_k \geq s_{-K} = \tfrac{l+h}{2} \cdot (1 + \delta)^{-K} \geq \tfrac{l+h}{2} \cdot \tfrac{l+h}{2h} \cdot (1 + \delta) \geq \tfrac{4hl}{4h} \cdot (1 + \delta) > l \geq 0.$$

Therefore, $S \subset [0, 1]$ is a valid finite signal space.

Step 2: We verify that the parameters x_k, y, z are non-negative.

Since $s_k \in (l, h)$, by the construction of x_k for $k \neq 0$, $x_k/x_0 > 0$. For y and z, since $s_{-K} < s_0 = (l + h)/2 < s_K$,

$$y/x_{-K} = \tfrac{s_{-K}-l}{l+h-2s_{-K}} = \tfrac{(s_{-K}-l)/2}{(l+h)/2-s_{-K}} > 0,$$
$$z/x_{K-1} = \tfrac{h-s_K}{2s_K-l-h} = \tfrac{(h-s_K)/2}{s_K-(l+h)/2} > 0.$$

Therefore, as long as $x_0 > 0$, all probability terms are positive.

Step 3: We verify that we can choose x_0 such that:

$$\sum_s x((l, h), s) = \sum_s x((h, l), s) = 1/2$$

As we showed, the coefficients, x_k/x_0, y/x_0, and z/x_0 are all positive and fixed. Then with x_0 defined below,

$$x_0 = \tfrac{1}{2} \cdot \tfrac{1}{y/x_0 + z/x_0 + \sum_k x_k/x_0} > 0,$$

we have $\sum_s x((l, h), s) = \sum_s x((h, l), s) = y + z + \sum_k x_k = 1/2$.

Step 4: We check that the calibration constraints are satisfied.

For $-K + 1 \leq k \leq K - 1$, we verify the calibration constraint for sending signal s_k to bidder 1 as follows:

$$\begin{aligned} &\textstyle\sum_{s \in S} x((l, h), (s_k, s)) \cdot (l - s_k) + x((h, l), (s_k, s)) \cdot (h - s_k) \\ &= x_k \cdot (l - s_k) + x_{k-1} \cdot (h - s_k) \\ &= x_{k-1} \cdot \tfrac{h-s_k}{s_k-l} \cdot (l - s_k) + x_{k-1} \cdot (h - s_k) = 0. \end{aligned}$$

When sending signal s_K to bidder 1:

$$\begin{aligned}&\textstyle\sum_{s\in S} x((l,h),(s_K,s))\cdot(l-s_K)+x((h,l),(s_K,s))\cdot(h-s_K)\\ &= z\cdot(l-s_K)+x_{K-1}\cdot(h-s_K)+z\cdot(h-s_K)\\ &= \tfrac{h-s_K}{2s_K-l-h}\cdot x_{K-1}\cdot(l-s_K)+x_{K-1}\cdot(h-s_K)\\ &\qquad+\tfrac{h-s_K}{2s_K-l-h}\cdot x_{K-1}\cdot(h-s_K)=0.\end{aligned}$$

When sending signal s_{-K} to bidder 1:

$$\begin{aligned}&\textstyle\sum_{s\in S} x((l,h),(s_{-K},s))\cdot(l-s_{-K})+x((h,l),(s_{-K},s))\cdot(h-s_{-K})\\ &= x_{-K}\cdot(l-s_{-K})+y\cdot(h-s_{-K})+y\cdot(h-s_{-K})\\ &= x_{-K}\cdot(l-s_{-K})+\tfrac{s_{-K}-l}{l+h-2s_{-K}}\cdot x_{-K}\cdot(l-s_{-K})\\ &\qquad+\tfrac{s_{-K}-l}{l+h-2s_{-K}}\cdot x_{-K}\cdot(h-s_{-K})=0.\end{aligned}$$

As the construction is symmetric for bidder 1 and 2, we omit the verification of the calibration constraints for bidder 2.

Step 5: We bound the revenue.

Note that when $s\in\{(s_k,s_{k+1}),(s_{k+1},s_k)\}_{k=-K}^{K-1}$, the auction allocates the item efficiently and the auctioneer extracts almost all the surplus. More specifically, when $s=(s_k,s_{k+1})$ for CTR profile (l,h), or $s=(s_{k+1},s_k)$ for CTR profile (h,l), the conditional expected revenue is

$$h\cdot\tfrac{s_k}{s_{k+1}}=h/(1+\delta)>h-\epsilon/2,\ \text{when}\ \delta<\tfrac{\epsilon}{2h-\epsilon}.$$

Therefore, we remain to prove that the probability of not extracting revenue $h/(1+\delta)$ is sufficiently small, i.e., $y,z<\epsilon/8$.

Recall that $h\ge s_K\cdot(1+\delta)$, with sufficiently small $\delta<(h-l)/2(h+l)$,

$$\begin{aligned}z&=x_{K-1}\cdot\frac{h-s_K}{2s_K-l-h}\le x_{K-1}\cdot\frac{h-h/(1+\delta)}{2h/(1+\delta)-l-h}\\ &=x_{K-1}\cdot\frac{\delta h}{h-l-\delta\cdot(h+l)}<\delta\cdot\frac{2h}{h-l}\cdot x_{K-1}<\delta\cdot\frac{2h}{h-l},\end{aligned}$$

which is less than $\epsilon/8$ when $\delta<\frac{h-l}{16h}\cdot\epsilon$.

Similarly, $s_{-K}=(1+\delta)^{-K}\cdot(l+h)/2\le\frac{l+h}{2h}\cdot(1+\delta)\cdot(l+h)/2<(1+\delta)(h+3l)/4$, with sufficiently small $\delta<(h-l)/2(h+3l)$,

$$\begin{aligned}y&=x_{-K}\cdot\frac{s_{-K}-l}{l+h-2s_{-K}}\le x_{-K}\cdot\frac{(1+\delta)(h+3l)/4-l}{l+h-(1+\delta)(h+3l)/2}\\ &=x_{-K}\cdot\frac{h-l+\delta\cdot(h+3l)}{2h-l-2\delta\cdot(h+3l)}<\frac{3}{2}\cdot x_{-K}.\end{aligned}$$

It suffices to show $x_{-K}<\epsilon/12$ with a sufficiently small δ. Note that $s_{-K}<\cdots<s_{-1}<s_0=(l+h)/2<s_1<\cdots<s_K$, we then have $x_0>x_1>\cdots>x_{K-1}$ and $x_0>x_{-1}>\cdots>x_{-K}$. Then

$$1/2=y+z+\textstyle\sum_{k=-K}^{K-1}x_k>(K+1)\cdot x_{-K}.$$

Therefore, when $\delta<\frac{\epsilon}{6}\cdot\log\frac{2h}{l+h}$, x_{-K} can be bounded by $\epsilon/12$:

$$\begin{aligned}x_{-K}&<1/2(K+1)<\frac{1}{2\log_{(1+\delta)}(2h/(l+h))}\\ &=\frac{\log(1+\delta)}{2\log(2h/(l+h))}<\frac{\delta}{2\log(2h/(l+h))}.\end{aligned}$$

In summary, for any given $\epsilon>0$, we can conclude the proof with a sufficiently small δ:

$$\begin{aligned}\delta&<\frac{h-l}{3(h+l)}\implies K-1\ge1,\\ \delta&<\frac{\epsilon}{2h-\epsilon}\implies h/(1+\delta)\ge h-\epsilon/2,\\ \delta&<\min\left(\frac{h-l}{2(h+l)},\frac{h-l}{16h}\cdot\epsilon\right)\implies z<\epsilon/8,\\ \delta&<\min\left(\frac{h-l}{2(h+3l)},\frac{\epsilon}{6}\cdot\log\frac{2h}{l+h}\right)\implies y<\epsilon/8.\end{aligned}$$

□

Lemma 4.2 is stated for a very special case within the class of symmetric environments. Both bidders have value 1, each bidder has only one of *two possible* click-through rates, and the click-through rates are *perfectly negatively* correlated. However, the result can now be extended immediately to a general symmetric environment. The extension is based on two simple observations:

Corollary 4.3 (High-Low Pairing). *Consider n bidders with normalized values $v_i\equiv v$ and click-through rates uniformly distributed between two profiles where r and r' are such for two bidders $i,j\in[n]$ that we have: $r_i=r'_j>r_j=r'_i\ge r_k,r'_k$ for any $k\ne i,j$. Then, for any $\epsilon>0$ there is a calibrated, correlated information structure with revenue $vr_i-\epsilon$.*

The next lemma shows that information structures can be composed, in the sense that if we decompose a distribution of click-through rates and design an information structure for each of them, we can later compose them without loss in calibration.

Lemma 4.4 (Signal Composition). *Let $\mathcal{G}'$ and $\mathcal{G}''$ be distributions over click-through rate profiles r of n bidders and let $\mathcal{F}'$ and $\mathcal{F}''$ be corresponding calibrated information structures given by joint distributions over vector pairs (r,s) such that the r-marginals are $\mathcal{G}'$ and $\mathcal{G}''$ respectively.*

Let $\mathcal{G}$ be the distribution obtained by sampling from $\mathcal{G}'$ with probability λ, and $\mathcal{G}''$ with probability $1-\lambda$. Define a distribution $\mathcal{F}$ similarly. Then, $\mathcal{F}$ is a calibrated information structure for $\mathcal{G}$ and

$$\textsc{Rev}(\mathcal{F})=\lambda\textsc{Rev}(\mathcal{F}')+(1-\lambda)\textsc{Rev}(\mathcal{F}'').$$

Combining the previous lemmas, we can prove Theorem 4.1:

Proof of Theorem 4.1. Consider an n-bidder symmetric environment and assume for simplicity that the distribution over CTRs is discrete. For every profile of CTRs where bidder i has the highest CTR r_i and bidder j has the second highest CTR r_j (breaking ties lexicographically), we can pair with a profile where the CTRs of i and j are reversed. This leads to a decomposition of the original distribution of CTRs into distributions with support two of the form studied in Corollary 4.3. The results follow from applying Corollary 4.3 together with the composition technique in Lemma 4.4. □

It turns out that the idea behind the construction can be generalized to work under the weaker requirement of *equal means*. We refer to the full version for a discussion of this generalization.

5 CORRELATED STRUCTURES BEYOND SYMMETRIC ENVIRONMENTS

In this section, we pursue a more limited objective for general asymmetric environments. To this end we consider an environment in which there are two bidders with values $v_1 = v_2 = 1$, and two possible click-through rate configurations, namely (r_1, r_2) and (r_1', r_2'). Without loss of generality, we can always label the identities of the bidders and the click-through rates so that:

$$r_1 \geq r_1', r_2, r_2'.$$

The prior probability of the pair (r_1, r_2) is p, the other pair (r_1', r_2') has the complementary probability $1-p$. Let $\mu_1 = p \cdot r_1 + (1-p) \cdot r_1'$ and $\mu_2 = p \cdot r_2 + (1-p) \cdot r_2'$ be the expected click-through rates of bidder 1 and bidder 2, respectively.

5.1 The Disclosure Lattice

Different information structures offer different levels of informativeness, and in fact form a *lattice* (see Figure 3). No disclosure is the minimal, and full disclosure the maximal policy. Together, they form the set of *extremal* policies. We refer to every policy that is not extremal as *moderate*, and to any policy that does not consist of a combination of full- or no-disclosure as *interior*.

Figure 3: Disclosure lattice. The label XY denotes the policies of bidder 1 and bidder 2. F = full, P = partial, and N = no.

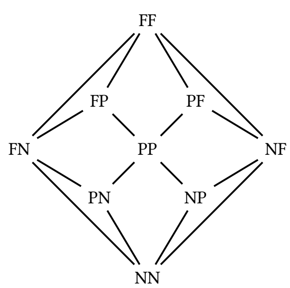

5.2 Extremal Structures are Dominated

The central result of this section establishes that optimal information design remains a powerful instrument to increase revenue, irrespective of the joint distribution of the click-through rates.

Theorem 5.1 (Moderate Information Structures). *There always exists a moderate information structure that strictly dominates any extremal information structure.*

Our proof of this result, which we defer to the full version, proceeds by distinguishing the following cases:

Uniform Winner: $r_1' \geq r_2'$	Variable Winner: $r_1' < r_2'$
Congruent Loser: $r_2 \geq r_2'$	Incongruent Loser: $r_2 < r_2'$
Weak Competition: $r_2 \leq \mu_1$	Strong Competition: $r_2 > \mu_1$

5.3 Moderate vs. Interior Structures

We highlight two building blocks in the proof of Theorem 5.1. Our first result is for the uniform winner congruent loser, and weak competition case. For this setting, we can identify the uniquely optimal information structure. Namely, it is optimal to bundle the click-through rates of the winner and to unbundle the click-through rates of the loser. The bundling of the click-through rates of the winner generates more competitive prices, and hence higher revenue for the auctioneer. Thus, the optimal information structure is moderate, but not interior.

Theorem 5.2 (Uniform Winner, Congruent Loser, Weak Competition). *The optimal information structure for the uniform winner, congruent loser, and weak competition case leaves the loser unbundled and fully bundles the winner.*

We provide the proof of this result, which relies on Chebyshev's sum inequality and nicely demonstrates the gist of our results for the asymmetric setting in Appendix C.

For the next result we stay with the uniform winner setting, but now flip the ranking of click-through rates across click-through realizations. Thus, we consider the case of the incongruent loser. We can verify that the competition is guaranteed to be weak in the above sense. Now, we can show that an interior information structure will always be the optimal information structure. In contrast, the moderate information structure that was optimal in the congruent setting can be shown to perform worse than either of the extremal information structures.

Proposition 5.3 (Uniform Winner and Incongruent Loser). *With a uniform winner and incongruent loser, there is always an interior information structure that is revenue-improving over all exterior information structures. In particular, partially bundling the winner* and *the loser is revenue-improving, relative to all exterior information structures.*

These two cases demonstrate how rich the optimal structures in the asymmetric setting can be, but they also emphasize that the same guiding principles that governed the optimal policy in symmetric settings are in play. Namely, as in the symmetric setting, the optimal schemes that we identify are not extremal, maintain efficiency, and seek to strengthen competition through calibrated signals.

6 CONCLUSION

The disclosure policy of the auctioneer regarding the click-through rates influences the distribution of bids, holding fixed the distribution of preferences among the bidders. By disclosing less, the auctioneer in effect bundles certain features, or as Levin and Milgrom [22] suggest, conflates features of the viewer. The process of conflation influences the thickness or thinness of the market—in other words, the strength of the competition.

In our analysis, conflation was achieved by bundling the information regarding click-through rates in an optimal manner. The optimal information structure conveys just enough information so that the resulting bidding process ranks the bidders according to the true social value of each bidder. At the same time, the information is released only partially to maintain bids as close as possible to a perfectly competitive level. We show that this requires that the information is provided to each bidder as private information at an individual level, rather than as public information on a market level.

REFERENCES

[1] I. Arieli and Y. Babichenko. 2019. Private Bayesian Persuasion. *Journal of Economic Theory* 182 (2019), 185–217.
[2] A. Badanidiyuru, K. Bhawalkar, and H. Xu. 2019. Targeting and Signaling in Ad Auctions. In *Proceedings of ACM-SIAM SODA'19*. 2545—2563.
[3] Dirk Bergemann, Tibor Heumann, Stephen Morris, Constantine Sorokin, and Eyal Winter. 2022. Optimal Information Disclosure in Auctions. *American Economic Review: Insights* 4 (2022), forthcoming.
[4] D. Bergemann and S. Morris. 2019. Information Design: A Unified Perspective. *Journal of Economic Literature* 57 (2019), 44–95.
[5] D. Bergemann and M. Pesendorfer. 2007. Information Structures in Optimal Auctions. *Journal of Economic Theory* 137 (2007), 580–609.
[6] S. Board. 2009. Revealing Information in Auctions: The Allocation Effect. *Economic Theory* 38 (2009), 125–135.
[7] O. Candogan and P. Strack. 2021. Optimal Disclosure of Information to a Privately Informed Receiver. (2021).
[8] J. Cremer and R. McLean. 1985. Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands are Interdependent. *Econometrica* 53 (1985), 345–361.
[9] J. Cremer and R. McLean. 1988. Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions. *Econometrica* 56 (1988), 1247–1258.
[10] B. Edelman, M. Ostrovsky, and M. Schwartz. 2007. Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review* 97 (2007), 242–259.
[11] Y. Emek, M. Feldman, I. Gamzu, R. Paes Leme, and M. Tennenholtz. 2012. Signaling Schemes for Revenue Maximization. In *Proceedings of ACM EC'13*. 92–109.
[12] P. Eső and B. Szentes. 2007. Optimal Information Disclosure in Auctions. *Review of Economic Studies* 74 (2007), 705–731.
[13] D. Foster and R. Vohra. 1997. Calibrated Learning and Correlated Equilibrium. *Games and Economic Behavior* 21 (1997), 40–55.
[14] Hu Fu, Patrick Jordan, Mohammad Mahdian, Uri Nadav, Inbal Talgam-Cohen, and Sergei Vassilvitskii. 2012. Ad auctions with data. In *International Symposium on Algorithmic Game Theory*. Springer, 168–179.
[15] Thore Graepel, Joaquin Quiñonero Candela, Thomas Borchert, and Ralf Herbrich. 2010. Web-scale Bayesian click-through rate prediction for sponsored search advertising in Microsoft's bing search engine. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*. 13–20.
[16] G. H. Hardy, J. E. Littlewood, and G. Pólya. 1988. *Inequalities*. Cambridge University Press, Cambridge, UK.
[17] J. D. Hartline, A. Johnsen, D. Nekipelov, and O. Zoeter. 2019. Dashboard mechanisms for online marketplaces. In *Proceedings of ACM EC'19*. 591–592.
[18] Patrick Hummel and R Preston McAfee. 2016. When does improved targeting increase revenue? *ACM Transactions on Economics and Computation (TEAC)* 5, 1 (2016), 1–29.
[19] A. Kleiner, B. Moldovanu, and P. Strack. 2020. Extreme Points and Majorization: Economic Applications. *Econometrica* 89 (2020), 1557–1593. Issue 4.
[20] A. Kolotilin, M. Li, T. Mylovanov, and A. Zapechelnyuk. 2017. Persuasion of Privately Informed Receiver. *Econometrica* 85 (6) (2017), 1949–1964.
[21] S. Lahaie and D. Pennock. 2007. Revenue Analysis of a Family of Ranking Rules for Keyword Auctions. In *Proceedings of ACM EC'07*. 50–56.
[22] J. Levin and P. Milgrom. 2010. Online Advertising: Heterogeneity and Conflation in Market Design. *American Economic Review Papers and Proceedings* 100 (2010), 603–607.
[23] M. Mahdian and M. Sundararajan. 2015. *Robust Mechanisms for Sponsored Search*. Technical Report. Google Research.
[24] H. B. McMahan, G. Holt, D. Sculley, M. Young, D. Ebner, J. Grady, L. Nie, T. Phillips, E. Davydov, D. Golovin, S. Chikkerur, D. Liu, M. Wattenberg, A. M. Hrafnkelsson, T. Boulos, and J. Kubica. 2013. Ad click prediction: a view from the trenches. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. 1222–1230.
[25] H.R. Varian. 2007. Position Auctions. *International Journal of Industrial Organization* 25 (2007), 1163–1178.

A DETAILED LITERATURE REVIEW

Bergemann and Pesendorfer [5] and Eső and Szentes [12] are among the first to investigate the design of optimal information structure in an auction setting. They consider an auction environment with n bidders and independent private values where the seller jointly optimizes auction and information policies. A number of contributions have recently analyzed the optimal information structure for a given mechanism or auction format. Badanidiyuru et al. [2], Emek et al. [11], Fu et al. [14] consider a second-price auction where the information about the valuation of bidder i is partially shared between the bidder and the platform. Emek et al. [11] study the computational problem of the revenue-optimal public information structure. Badanidiyuru et al. [2] prove that in order to guarantee a constant approximation to the optimal revenue in this case, the size of the information structure needs to be exponentially large. Fu et al. [14] provide examples where partial information disclosure may yield higher expected revenue than both full- and no-disclosure in second price auctions with reserves. Hummel and McAfee [18] study a wide range of position auctions, including both GSP and VCG auctions, with or without reserve prices. They consider different orders between realizing values and submitting bids, which are essentially equivalent to full-disclosure (if realize-value-then-bid) and no-disclosure (if bid-then-realize-value).

Given the randomness of the click-through rates an earlier literature noticed that a systematic modification of the click-through rates may positively impact the revenue of the click-through auction. The adjusted click-through rates in Lahaie and Pennock [21] and Mahdian and Sundararajan [23] are not calibrated, as the true click-through rate is "squashed" or "boosted" for the purpose of ranking. This suggests a natural distinction between the analysis of optimal calibrated and non-calibrated information structures.

In the current setting, we explicitly allow for private information disclosure, rather than public information disclosure. By contrast, in most of the preceding literature the information disclosure was either public, as in Arieli and Babichenko [1], or independent across bidders, as in Bergemann and Pesendorfer [5]. Here, we are allow for, and importantly show the optimality of private and correlated information structures. The role of correlated signals for the revenue maximizing mechanism has been observed earlier in Cremer and McLean [8, 9]. They establish that correlation in the private values among bidders can be used in an optimal mechanism to extract the full surplus. While our results also highlight the increased power of correlated signals relative to independent signals, the setting and arguments differ substantially. In Cremer and McLean [8, 9], the auctioneer is free to choose the optimal mechanism, while we take the generalized second-price auction as given. We choose the information structure so that the signals are sufficiently informative, yet yield competitive interim expectations. Thus, the correlation of the signals achieves very different objectives in these two settings, and accordingly the construction differs significantly. In particular, Cremer and McLean [8, 9] take as given the signals and then design the optimal transfer function. We take as given the transfer function, namely the payment rules, and then design the signals to maintain competition.

Hartline et al. [17] suggest a class of dashboard mechanism that shares some of the themes with our contribution in a very distinct setting. They offer dashboards as an instrument for platforms to offer bid recommendation in mechanism when truthtelling does not form an equilibrium strategy. The leading examples are first price auction and all pay auctions. The dashboard converts the initial information of the bidders into bid recommendation and then asks which allocation and bid recommendations lead to obedient behavior by the bidders.

B PROOFS OMITTED FROM SECTION 4

Proof of Corollary 4.3. Treat bidders i and j as the high/low pair in Lemma 4.2 and do full-disclosure for any other bidder. The signal is still calibrated and only i and j win the item since their signals will be above the signals of any other bidder. Hence, the revenue bound in Lemma 4.2 still holds. □

Proof of Lemma 4.4. The r-marginal of $\mathcal{F}$ is clearly $\mathcal{G}$ and it holds that $\text{Rev}(\mathcal{F}) = \lambda\text{Rev}(\mathcal{F}') + (1-\lambda)\text{Rev}(\mathcal{F}'')$. The only non-trivial part is to check that $\mathcal{F}$ is calibrated, which we do below:

$$\begin{aligned}\mathop{\mathbb{E}}_{\mathcal{F}}[r_i|s_i = s_i'] &= \frac{\mathbb{E}_{\mathcal{F}}[r_i \mathbf{1}\{s_i - s_i'\}]}{\Pr_{\mathcal{F}}(s_i = s_i')}\\ &= \frac{\lambda\,\mathbb{E}_{\mathcal{F}'}[r_i \mathbf{1}\{s_i - s_i'\}] + (1-\lambda)\,\mathbb{E}_{\mathcal{F}''}[r_i \mathbf{1}\{s_i - s_i'\}]}{\lambda\Pr_{\mathcal{F}'}(s_i = s_i') + (1-\lambda)\Pr_{\mathcal{F}''}(s_i = s_i')}\\ &= \frac{\lambda\,\mathbb{E}_{\mathcal{F}'}[s_i' \mathbf{1}\{s_i - s_i'\}] + (1-\lambda)\,\mathbb{E}_{\mathcal{F}''}[s_i' \mathbf{1}\{s_i - s_i'\}]}{\lambda\Pr_{\mathcal{F}'}(s_i = s_i') + (1-\lambda)\Pr_{\mathcal{F}''}(s_i = s_i')} = s_i'\end{aligned}$$

This concludes the proof. □

C PROOF OF THEOREM 5.2

We provide the proof of Theorem 5.2. The main tool in the proof is the following lemma.

Lemma C.1 (Chebyshev's sum inequality [16]). *Given two sequences $a_1 \geq a_2 \geq \ldots a_n \geq 0$ and $b_1 \geq b_2 \geq \ldots b_n \geq 0$ that are monotone in the same direction, and a set of non-negative weights $w_i \geq 0$ (not necessarily monotone), then:*

$$\left(\sum_i w_i a_i b_i\right) \cdot \left(\sum_i w_i\right) \geq \left(\sum_i w_i a_i\right) \cdot \left(\sum_i w_i b_i\right).$$

If $\{a_i\}$ and $\{b_i\}$ sequences are monotone in different directions (one increasing and one decreasing), the inequality holds in the opposite direction.

(Aside: The probabilistic interpretation of Lemma C.1 is that if A and B are two positively-correlated random variables, then $\mathbb{E}[AB] \geq \mathbb{E}[A] \cdot \mathbb{E}[B]$.)

Proof of Theorem 5.2. We will start with a generic solution $x(r,s)$ and show that using two applications of Chebyshev's sum inequality (Lemma C.1) we can bound it with respect to

$$p \cdot r_1 \cdot \frac{r_2}{\mu_1} + (1-p) \cdot r_1' \cdot \frac{r_2'}{\mu_1}, \tag{3}$$

which is the revenue obtained from leaving the loser unbundled and bundling the winner. We will proceed in three steps.

Step 1: Bounding the revenue. Consider any information structure defined by $x(r,s)$. For each signal $s = (s_1, s_2)$, the revenue in the event that bidder 1 wins is $(x(r,s)r_1 + x(r',s)r_1') \cdot s_2/s_1$. If bidder 2 wins, the revenue is $(x(r,s)r_2 + x(r',s)r_2') \cdot s_1/s_2$. In the second case, observe that since $s_2 \geq s_1$ we have $\frac{s_1}{s_2} < 1 < \frac{s_2}{s_1}$, and since we are in the uniform winner case, we know that $x(r,s)r_1 + x(r',s)r_1' \geq x(r,s)r_2 + x(r',s)r_2'$. Hence, we can bound:

$$(x(r,s)r_2 + x(r',s)r_2') \cdot \frac{s_1}{s_2} \leq (x(r,s)r_1 + x(r',s)r_1') \cdot \frac{s_2}{s_1},$$

and write:

$$\textsc{Rev} \leq \sum_s (x(r,s)r_1 + x(r',s)r_1') \cdot \frac{s_2}{s_1}.$$

Step 2: Unbundling the loser. Substituting s_2 by the calibration constraint we obtain:

$$\begin{aligned}\textsc{Rev} \leq \sum_{s_2}\sum_{s_1} &\left(x(r,(s_1,s_2))\frac{r_1}{s_1} + x(r',(s_1,s_2))\frac{r_1'}{s_1}\right)\\ &\cdot \frac{\sum_{s_1} x(r,(s_1,s_2))r_2 + x(r',(s_1,s_2))r_2'}{\sum_{s_1} x(r,(s_1,s_2) + x(r',(s_1,s_2))}\\ \leq \sum_{s_2}\sum_{s_1} &\left(x(r,(s_1,s_2))\frac{r_1 r_2}{s_1} + x(r',(s_1,s_2))\frac{r_1' r_2'}{s_1}\right),\end{aligned}$$

where the second inequality follows from Chebyshev's sum inequality with

$$\{a_i\}_i = \left(\frac{r_1}{s_1^1},\ldots,\frac{r_1}{s_1^n},\frac{r_1'}{s_1^1},\ldots,\frac{r_1'}{s_1^n}\right), \quad \{b_i\}_i = (r_2,\ldots,r_2,r_2',\ldots,r_2').$$

Congruence implies that the sequence $\{b_i\}$ is sorted. The sequence $\{a_i\}$ is sorted because $r_1' \leq s_1^1 \leq \ldots \leq s_1 \leq r_1$, and hence $\frac{r_1}{s_1} \geq 1 \geq \frac{r_1'}{s_1}$. Now that the loser is unbundled, there is no longer any need to keep track of s_2. To simplify notation we will define:

$$\tilde{x}(r,s_1) = \sum_{s_2} x(r,(s_1,s_2))$$

and re-write our current bound on the objective as:

$$\textsc{Rev} \leq \sum_{s_1} \frac{\tilde{x}(r,s_1)r_1r_2 + \tilde{x}(r',s_1)r_1'r_2'}{s_1}.$$

Step 3: Bundling the winner. We will replace s_1 according to the calibration constraint in the expression above and replace:

$$\lambda(s_1) = \frac{\tilde{x}(r,s_1)}{\tilde{x}(r,s_1)+\tilde{x}(r',s_1)} \quad \text{and} \quad \lambda'(s_1) = \frac{\tilde{x}(r',s_1)}{\tilde{x}(r,s_1)+\tilde{x}(r',s_1)}.$$

We obtain:

$$\begin{aligned}\textsc{Rev} &\leq \sum_j (\tilde{x}(r,s_1)r_1r_2 + \tilde{x}(r',s_1)r_1'r_2') \cdot \frac{\tilde{x}(r,s_1)+\tilde{x}(r',s_1)}{r_1\tilde{x}(r,s_1) + r_1'\tilde{x}(r',s_1)}\\ &= \sum_j \frac{\lambda(s_1)r_1r_2 + \lambda'(s_1)r_1'r_2'}{\lambda(s_1)r_1 + \lambda'(s_1)r_1'} \cdot (\tilde{x}(r,s_1)+\tilde{x}(r',s_1)).\end{aligned}$$

Now we can apply Chebyshev's sum inequality one more time with:

$$a(s_1) = \frac{\lambda(s_1)r_1r_2 + \lambda'(s_1)r_1'r_2'}{\lambda(s_1)r_1 + \lambda'(s_1)r_1'}, \quad b(s_1) = \lambda(s_1)r_1 + \lambda'(s_1)r_1',$$
$$w(s_1) = \tilde{x}(r,s_1)+\tilde{x}(r',s_1).$$

Note that we can reorder signals s_1 in any order we wish. Let us reorder the signals so that $\lambda(s_1)$ is increasing. This immediately implies that $b(s_1)$ is increasing. Namely, $b(s_1) = \lambda(s_1)r_1 + (1 - \lambda(s_1))r_1' = \lambda(s_1)(r_1 - r_1') + r_1'$. For $a(s_1)$, we can take the derivative in λ,

$$\frac{d}{d\lambda}a(s_1) = \frac{r_1 \cdot r_1' \cdot (r_2 - r_2')}{(r_1\lambda - r_1'(1-\lambda))^2} > 0.$$

Thus, $a(s_1)$ is also increasing in λ, allowing us to apply the inequality to obtain:

$$\sum_{s_1} w(s_1)a(s_1) \leq \left(\sum_{s_1} w(s_1)a(s_1)b(s_1)\right)\frac{\sum_{s_1} w(s_1)}{\sum_{s_1} w(s_1)b(s_1)},$$

which translates to (since $\sum_{s_1} w(s_1) = 1$):

$$\begin{aligned}\textsc{Rev} &\leq \frac{\sum_{s_1}\tilde{x}(r,s_1)r_1r_2 + \tilde{x}(r',s_1)r_1'r_2'}{\sum_{s_1}\tilde{x}(r,s_1)r_1 + \tilde{x}(r',s_1)r_1'}\\ &= \sum_{s_1}\tilde{x}(r,s_1)r_1\frac{r_2}{\mu_1} + \tilde{x}(r',s_1)r_1'\frac{r_2'}{\mu_1} = pr_1\frac{r_2}{\mu_1} + (1-p)r_1'\frac{r_2'}{\mu_1}.\end{aligned}$$

This is the revenue obtained by bundling the winner and unbundling the loser as desired. □