Opaque Auctions: Implicit Signals and Value of Information in Ad Auctions

Stylianos Despotakis and Jungju Yu*

June 15, 2026

Abstract

Online advertising markets exhibit two characteristic features: opacity in the platform's allocation process and two-sided information asymmetry, whereby platforms observe user engagement data unknown to advertisers while advertisers possess private business valuations unknown to the platform. In standard auction settings with commitment power, a platform can implement publicly observable auction-design rules through which its private information is directly incorporated into advertisers' decisions. However, the platform's strategy is often unobservable to advertisers, and thus the platform lacks commitment, implying that its private information can only be transmitted implicitly through observable choices.

We develop a model of opaque auctions characterized by two-sided private information and the platform's lack of commitment to auction-design strategies. In this setting, the platform can set reserve prices and decide which advertisers to solicit, choices that may serve as implicit signals. First, we show that the reserve price cannot credibly convey the platform's private information, but selective bidder solicitation can: the platform optimally excludes some advertisers from the auction, and being invited serves as a credible signal of high advertiser-impression match value, inducing more aggressive bidding. When either the platform or the advertisers have full information, we show that being invited no longer carries information, and thus inviting all advertisers is optimal. Second, acquiring the advertisers' private information can harm the platform by undermining the credibility of the signaling mechanism, an effect absent with commitment. Third, the incentives of the platform and the advertisers to share their private information and acquire counter-party's information are aligned if and only if the market size is large enough. These results show that the value of information is determined by the interaction among the information environment, commitment power, and market size.

Keywords: Online advertising; Auction design; Information asymmetry; Value of information; Commitment; Transparency.

*Stylianos Despotakis: sdespota@cityu.edu.hk (City University of Hong Kong), Jungju Yu: yujungju@snu.ac.kr (Seoul National University). Despotakis acknowledges financial support from the Research Grants Council of Hong Kong (Project No. CityU 11509924).

1 Introduction

Online display ads are predominantly transacted programmatically, often through real-time auctions facilitated by ad exchanges or supply-side platforms.[1] When a user loads a webpage or opens an app, an ad impression becomes available. The exchange—used interchangeably with ad platform from here on—receives information about the impression, including user-level, context-level, device-level, and placement-level signals. It then constructs a bid request describing the impression and routes it to bidders, typically through demand-side platforms (DSPs). The DSPs evaluate the bid request on behalf of advertisers, estimate the value of the impression for each eligible campaign, and submit bids based on the information provided by the exchange, as well as their proprietary data. Thus, data are a key resource for ad exchanges and are essential to the functioning of online ad auction markets.

The exchange's payoff from ad auctions ultimately depends on bidders' ability to form accurate valuations of impressions and compete accordingly. Thus, the *economic value* of an ad exchange's data depends not only on the data it possesses, but also on whether that information can be communicated to bidders in a way that informs their bidding decisions.[2] However, in practice, how advertisers should interpret and incorporate exchange-provided signals into their bidding decisions is far from straightforward. This is because the exchange's impression-side information is proprietary and not directly available to advertisers in its raw form. In addition, advertisers typically observe only coarse descriptors of an impression, such as "female, mid-30s, metropolitan, sports enthusiast, professional," whose precise definitions, thresholds and predictive relevance may be unclear. As a result, it may be difficult for advertisers to translate the disclosed information into optimal bids.

Given these challenges, advertisers may resort to observable features of the auction environment configured by the ad exchange for indirect information about the impression. The reserve price is a natural example, as it may reflect the exchange's assessment of impression quality based on its private information on the impression side. Another institutionally important—but far less explored academically—example is the exchange's bidder selection through bid-request routing. Ad exchanges do not always send bid requests to all potential advertisers; instead, they often solicit

[1] According to eMarketer (2026), programmatic ads account for more than 90% of U.S. digital display ad spending.

[2] For instance, Ada et al. (2022) show that providing more granular context information in bid requests increased revenue per impression at a private ad exchange. This finding suggests that the way an ad exchange's information is communicated to bidders is economically consequential for both sides of the market.

only a subset of bidders for a given impression. Industry article (Google Authorized Buyers, 2026) documents that such routing decisions are shaped by technical and operational constraints, such as server capacity, callout quotas, and latency requirements.[3]

These auction-design choices, however, are not simple information disclosures. The reserve price and the bidder's solicitation status are observable to individual advertisers, but they are also strategic choices made by the exchange. Because advertisers do not observe the way that the exchange maps its private information into these choices, how advertisers should interpret these observable signals is not immediately clear, either. Nevertheless, these choices may serve as *implicit signals* through which platform information is transmitted to bidders, thus shaping their bidding decisions and, ultimately, the exchange's payoff.

This paper studies such an *opaque* information environment in an ad auction game and examines how exchange information can be transmitted implicitly through observable auction-design choices. Opacity has two key features in our setting. First, the ad exchange has private and partial information about advertisers' values for an impression. Second, the mapping from the exchange's private information to its observable auction-design choices—the reserve price and bidder selection—are unobservable to the advertisers, which means that the exchange cannot commit to any auction-design strategies. In this environment, we analyze whether and how any information relevant for the value of the impression can be credibly conveyed to the advertisers, and how that information affects the advertisers' bidding behavior and, in turn, the exchange's payoff. Our approach complements the literature on online ad auctions, which often assumes that the exchange can commit to an information-disclosure policy, or more generally a signal structure, that determines what bidders learn about an impression (e.g., De Corniere and De Nijs, 2016; Bergemann et al., 2022; Choi and Sayedi, 2024). In contrast, our approach focuses on information transmission without exchange's full commitment to auction-design strategies, and thus the exchange's design choices themselves may serve as implicit signals.

Understanding how platform information is transmitted in an opaque auction environment is critical for assessing the economic value of information. When information cannot be directly disclosed, its value depends on whether it can be credibly conveyed to bidders and translated into

[3] Bidder solicitation is also prevalent in government procurement, where contracting officers often have discretion over which firms to invite to bid (Elmaghraby, 2005; Coviello et al., 2018).

bidding decisions. Based on this observation, we examine the value of information for both the ad exchange and advertisers, including when each party benefits from acquiring the private information of the counterparty. We also compare these outcomes with a benchmark in which the exchange can commit to a bidder-solicitation strategy, allowing us to isolate the role of commitment in shaping equilibrium outcomes and information value.

We develop a stylized auction game to study the opaque information environment described above. A single ad impression is sold to one of $K \geq 2$ advertisers through a second-price auction.[4] Advertiser i's value for the impression is decomposed into two parts, i.e., $v_i = c_i \cdot \theta_i$, where c_i captures the expected conversion rate and θ_i captures the revenue conditional on conversion. Each advertiser's c_i and θ_i are drawn randomly and independently from a distribution. To reflect the platform's advantage in impression-side information and advertisers' advantage in business-side information, we assume that the platform privately observes c_i's of all advertisers, while each advertiser privately observes its own θ_i.

Before c_i's and θ_i's are realized, the platform announces n, the size of the invited bidder set. After these components are realized, the exchange configures the auction environment based on its private information by choosing which n advertisers to solicit and setting a reserve price r. Invited advertisers observe the reserve price, n, and their own invitation status, but not the platform's auction-design strategy. They then update their beliefs about their own c_i's, and hence their values v_i before submitting bids. The model therefore captures a setting where the platform lacks full commitment to an auction-design strategy, and bidders draw inferences from the observable auction environment.

We use perfect Bayesian equilibrium (PBE) as the equilibrium concept. As is common in signaling games, the model may admit multiple equilibria; we therefore focus on platform-optimal PBEs, namely those that maximize the platform's payoff among all PBEs.[5]

We first show that the reserve price cannot credibly transmit the platform's private information

[4] In a second-price auction, the highest bid wins if it exceeds the reserve price, and the winner pays the greater of the second-highest bid and the reserve price.

[5] Formally, a PBE is platform-optimal if no other PBE gives the platform a strictly higher payoff. This selection criterion is natural for our purpose because the paper studies the best equilibrium outcome that can be sustained for an informed platform when it cannot commit to a disclosure policy. It is conceptually related to sender-preferred equilibrium selection in strategic communication (e.g., Crawford and Sobel, 1982; Chen et al., 2008; Lipnowski and Ravid, 2020). It is also related in spirit to sender-optimal information design (e.g., Kamenica and Gentzkow, 2011; Bergemann and Morris, 2019), but differs because information design typically studies the sender's optimal committed signal structure rather than selecting among PBEs in a game without commitment.

in the platform-optimal equilibrium. This is because the platform will choose the reserve price that maximizes its expected payoff regardless of the realized c_i's due to the opaque environment, i.e., the unobservability of the platform's private information and reserve pricing strategy to the advertisers. Nevertheless, the platform's private information can be credibly transmitted through the platform's bidder solicitation. In particular, we identify a platform-optimal equilibrium in which the platform invites a subset of advertisers with n^* highest realized conversion rates c_i. Then, upon being solicited, each advertiser rationally updates its beliefs that its unobserved component c_i lies among the top n^* out of K i.i.d. draws and thus bids more aggressively. This signaling effect is stronger when the invited set is smaller, because solicitation becomes a more selective signal of advertiser-impression match quality. As a result, when there are sufficiently many potential advertisers, the platform optimally excludes one or more advertisers from the auction, even though doing so reduces the number of submitted bids. This result departs from those derived from standard auction models in that, when solicitation is informative, excluding some bidders can raise revenue by strengthening the signal conveyed to those who are invited.

Next, we compare the main model with several benchmark models of various information structures. In the first benchmark, the platform is fully informed while keeping the advertisers only partially informed: the only difference from the main model is that the platform now also has access to the advertiser-side component θ_i. In this benchmark, the platform's private information can no longer be transmitted through either the reserve price or bidder solicitation. The reason is that the platform cannot credibly commit to a bidder-selection strategy based on its private information. If advertisers believed that solicitation reflected the platform's private information c_i, the platform would have a profitable deviation to solicit bidders based on the commonly observed θ_i instead. Such a deviation would not change advertisers' posterior beliefs about c_i, while allowing the platform to select bidders with higher advertiser-side values and thereby induce higher bids. With no information conveyed through bidder selection, the invited advertisers may bid less aggressively, which sometimes leads to a less efficient auction outcome and lower payoff of the platform. Thus, additional information can hurt the platform by exacerbating its commitment problem and undermining the credibility of bidder solicitation as an implicit signal.

In contrast, when advertisers are fully informed about their own values, the platform's additional information unambiguously benefits the platform. To show this, we compare two benchmark models in which advertisers are fully informed: in one benchmark, the platform is only partially informed, and fully informed in the other. In this case, bidder selection has no informational role because bidders already know their exact values. The platform can therefore use its additional information to set reserve prices that are more finely tailored to realized bidder-specific values, allowing it to extract greater surplus.

Third, we study incentives to share private information, under the assumption that information—if disclosed—can be verified and is therefore truthful. Advertisers are always worse off from sharing their private information about θ_i: doing so moves the market toward a regime in which the platform can infer their exact values and extract their surplus via reserve pricing. In contrast, the platform can sometimes benefit from disclosing its private information about c_i's. When disclosure makes bidders better informed about their own values, the platform can raise reserves and intensify extraction, potentially increasing revenue. This asymmetry implies that, relative to advertisers, platforms may have stronger incentives to disclose and disseminate value-relevant information. From a policy perspective, this suggests that regulatory attention may be especially warranted on platform-side data practices, where incentives to share information can be stronger.

Finally, we examine the role of platform commitment by analyzing a modified model in which the platform can commit to a bidder-selection policy before realizing its private information. We focus on two natural and interpretable classes of monotone strategies. The first is a "top-n" rule, which was shown to be the equilibrium strategy without the platform's commitment. The second is a cutoff rule that invites all advertisers with $c_i \geq \tau$, where τ is optimally chosen. We show that, under the information structure of the main model with two-sided information asymmetry, commitment strictly increases platform profits and that the optimal cutoff policy dominates the optimal top-n policy. Commitment also affects information-sharing incentives. Under commitment, the platform no longer wishes to disclose its private information, in contrast to the no-commitment environment, where disclosure can be profitable in sufficiently thick markets. This suggests a regulatory implication: when evaluating policies governing the sharing of user-level information with advertisers or demand-side platforms, whether the ad exchange can commit to an auction design rule needs to be considered together as it may fundamentally affect its incentive for information sharing.

The rest of the paper proceeds as follows. The next subsection discusses related literature.

Section 2 introduces the main model and benchmark models. Section 3 characterizes the equilibrium and presents the main results. Section 4 analyzes benchmark information structures and studies the marginal value of information for the platform and advertisers. Section 5 examines the effects of commitment and how commitment and information jointly determine the value of information and the platform's incentives to disclose. Section 6 concludes.

1.1 Related Literature

Our paper relates to several strands of literature on bidder selection, signaling, and information disclosure in auctions and advertising markets.

First, our paper contributes to the literature on bidder selection and restricted participation in ad auctions. In the context of real-time bidding, Chakraborty et al. (2010) study ad exchanges that, due to latency and capacity constraints, call out only a subset of buyers for each impression and choose that subset to maximize expected welfare or revenue based on predicted bid distributions. Relatedly, Choi and Sayedi (2023) study restricted participation through private exchanges. Their focus is on the publisher's strategic decision of whether to sell through an open exchange, a private exchange, or both. Our paper instead studies the ad exchange's impression-level decision about which bidders to solicit. This decision is not only a strategic auction-design choice that determines who participates, but also a potential channel of credible information transmission: being included in the auction allows invited advertisers to infer the exchange's private assessment of their value for the impression.

Conceptually closest to our paper is Lauermann and Wolinsky (2017), who study bidder solicitation by an informed seller in a common-value first-price auction. They show that the act of being solicited can itself convey information and affect competition. In their model, the seller is informed about the quality of a good whose value is common to all bidders and chooses how many bidders to solicit. Thus, from the mere fact of being invited, a bidder can draw positive or negative inferences about the good's quality, depending on the seller's equilibrium solicitation strategy. Their focus is on how this inference mechanism affects competition and information aggregation. Our paper differs in both setting and mechanism. We study an online ad auction with bidder-specific private values, where the platform's private information concerns advertiser-specific match quality rather than a common object value. Accordingly, our focus is not on whether dispersed bidder signals about a common value are aggregated through the auction, but on whether bidder-specific information held by the ad exchange can be credibly communicated in an opaque auction environment. Whereas solicitation in their model may generate adverse selection about a common value, solicitation in our setting can certify bidder-specific relevance and thereby strengthen bidding incentives among selected advertisers.[6]

Related, Cai et al. (2007) study reserve-price signaling in an interdependent-values auction in which the seller has private information about object characteristics that are relevant to bidders' valuations. They show that the reserve price can serve as a credible signal: a seller with a higher private signal sets a higher reserve price than a seller with a lower signal. The mechanism relies on a type-dependent trade-off between a higher sale price conditional on sale and a lower probability of sale, which prevents low types from mimicking high types. In our setting, however, the platform's private information concerns bidder-specific impression match rather than a common vertical quality of the object. As a result, the reserve price cannot play the same signaling role. Instead, the platform's bidder-solicitation decision becomes the credible channel through which private information may be transmitted.

Finally, our paper relates to the literature on targeted advertising and consumer inference. Several papers study settings in which targeting itself becomes informative because consumers draw inferences from the fact that they are targeted. Shin and Yu (2021) show that targeted advertising can signal product value when advertisers have private information about consumers' valuations. Despotakis and Yu (2023) study multidimensional targeting with two-sided private information and an unobservable targeting strategy, showing how consumer private information can discipline firm behavior and facilitate credible communication. Ning et al. (2025) also analyze targeted advertising under two-sided private information and show how consumer private information affects strategic mistargeting and firm profits. Our paper shares with this literature the idea that selection can become a signaling device, but studies it in a different market interaction: an ad platform designs an auction environment for advertisers, rather than an advertiser targeting consumers. More

[6] Bidder participation is also an important design margin in other auction settings. In procurement, buyers often have discretion over which suppliers to invite, prequalify, or exclude from bidding, and such discretion can affect auction performance (e.g., Elmaghraby, 2005; Coviello et al., 2018). Related work on endogenous entry similarly shows that participation shapes bidding behavior and revenue (e.g., Levin and Smith, 1994; Bajari and Hortaçsu, 2003; Li and Zheng, 2009). Although our analysis is motivated by online ad auctions, these literatures suggest that our insights may extend to other auction environments in which participation is actively managed.

importantly, we use this mechanism to study how opacity and implicit signaling affect the value of information itself. In particular, we analyze when additional information benefits or hurts the platform and advertisers, and how these effects depend on the platform's ability to commit to auction-design strategies.

Taken together, these streams of literature show that information disclosure, bidder participation, and strategic inference are central to advertising and auction markets. Our paper builds on these strands of research by studying an ad auction in which the platform's private information is implicitly transmitted through endogenous auction-design choices rather than direct disclosure. In this environment, auction-design choices—especially bidder solicitation—can affect revenue not only by determining participation, but also by shaping bidders' beliefs. This mechanism allows us to study how opacity and signaling affect the value of information for both the platform and advertisers.

2 Model

2.1 Main Model

We study a game between an ad exchange platform (hereafter, the *platform*) and $K \geq 2$ advertisers, in which the platform allocates a single ad impression through an auction. The platform can strategically determine which advertisers participate and set a reserve price, informed by its private knowledge of user conversion behavior. We now describe the primitives of the model.

Valuations. Each advertiser $i \in \{1, \ldots, K\}$ has a value v_i for the ad impression, representing the expected payoff from having its ad displayed. We decompose this value as

$$v_i = \theta_i \cdot c_i,$$

where $c_i \in [0, 1]$ is advertiser i's *conversion probability* (the probability that a user who views the ad takes a desired action, e.g. clicks, purchases, or signs up) and $\theta_i \in [0, 1]$ is advertiser i's *value per conversion* (the payoff conditional on such an action). The $2K$ random variables

$c_1, \ldots, c_K, \theta_1, \ldots, \theta_K$ are mutually independent, each drawn from the uniform distribution on $[0, 1]$.[7] In particular, advertisers are ex ante symmetric.

Information Structure. The platform and the advertisers possess different types of private information, giving rise to a two-sided information asymmetry. The platform observes the full vector of conversion probabilities $\mathbf{c} = (c_1, \ldots, c_K)$ but does not observe any advertiser's value per conversion. Each advertiser i privately observes its own value per conversion θ_i but does not observe its conversion probability c_i, nor does it observe any information about other advertisers. The number of advertisers K, the distributional assumptions, and the information structure described above are common knowledge among all players.

This information structure reflects institutional features of online advertising markets. Platforms possess rich user-level data (including browsing history, past ad engagement, and contextual signals such as device type, time of day, and ad placement) that enable accurate prediction of conversion probabilities. Individual advertisers, by contrast, observe only the outcomes of their own campaigns and lack the cross-platform data needed to estimate user-level conversion rates. Conversely, the value per conversion reflects proprietary business information (such as profit margins, customer lifetime value, inventory positions, and promotional strategies) that is better known to the advertiser than to the platform.

Platform's Decisions. The platform's strategy has two stages, reflecting the distinction between system-level policy choices and impression-level decisions.

In the first stage, before the realization of $\mathbf{c}$, the platform commits to the number of advertisers $n \in \{1, \ldots, K\}$ that will be invited to the auction.[8] The choice of n is made without knowledge of the realized conversion probabilities or values per conversion.

In the second stage, after observing $\mathbf{c}$, the platform makes two impression-level decisions. First,

[7] The uniform distribution assumption makes the analysis more tractable and allows us to derive closed-form solutions. This enables clear comparisons across information structures, allowing us to isolate the roles of information and commitment.

[8] This timing assumption reflects a salient feature of programmatic advertising markets, where the number of bid requests dispatched per impression is governed by system-level parameters, such as traffic-shaping configurations, computational budgets, and exchange policies, that are set in advance and applied uniformly across impressions. Similarly, platforms such as Google Ads determine the size of the candidate pool through algorithmic thresholds that do not vary on an impression-by-impression basis. The platform retains the flexibility to determine which advertisers to invite based on real-time signals, but the scale of participation is a policy-level commitment.

it selects a set of n advertisers to invite, according to a selection rule $\mathcal{S} : [0,1]^K \to \binom{[K]}{n}$ that maps the realized conversion probabilities to a subset of exactly n advertisers. Second, it sets a reserve price $r \geq 0$ that serves as the minimum acceptable bid; the reserve price may depend on $\mathbf{c}$ through a mapping $r : [0,1]^K \to \mathbb{R}_+$. Each invited advertiser $i \in \mathcal{S}(\mathbf{c})$ observes that it has been invited, the number of participants n, and the reserve price r. Importantly, it does not observe the platform's selection rule $\mathcal{S}$ or the conversion probabilities $\mathbf{c}$. Instead, it holds beliefs about the selection rule, denoted $\tilde{\mathcal{S}}$.

Auction Format. The impression is allocated through a second-price auction with reserve price r. Each invited advertiser $i \in \mathcal{S}(\mathbf{c})$ simultaneously submits a bid $b_i \geq 0$. The highest bidder wins the impression, with ties broken uniformly at random, and pays $\max\{r, b^{(2)}\}$, where $b^{(2)}$ denotes the second-highest bid among the invited advertisers (with $b^{(2)} \equiv 0$ if only one advertiser is invited). If all bids fall below r, the impression goes unallocated and no payments are made. Advertisers not in $\mathcal{S}(\mathbf{c})$ do not participate.

Payoffs. The platform's payoff equals the payment collected from the winning advertiser, or zero if the impression is unsold. Each advertiser i is risk-neutral: it receives a realized payoff of $v_i - p_i$ if it wins the impression and pays p_i, and zero otherwise. The platform seeks to maximize expected revenue, and each advertiser seeks to maximize expected payoff.

Game sequence. The game proceeds as follows:

1. The platform chooses the number of bidders $n \in \{1, \ldots, K\}$.[9]

2. The random variables $c_1, \ldots, c_K, \theta_1, \ldots, \theta_K$ are drawn independently from the uniform distribution on $[0,1]$. The platform observes $\mathbf{c}$, and each advertiser i privately observes θ_i.

3. The platform selects a set $\mathcal{S}(\mathbf{c}) \subseteq \{1, \ldots, K\}$ with $|\mathcal{S}(\mathbf{c})| = n$ and sets a reserve price $r \geq 0$. Both decisions may depend on the realized $\mathbf{c}$.

4. Each invited advertiser $i \in \mathcal{S}(\mathbf{c})$ observes (n, r) and submits a bid $b_i \geq 0$. Uninvited advertisers do not participate.

[9] If n is chosen after the realization of $\mathbf{c}$, then it can potentially act as a signal to the advertisers. Under this alternative timeline, the equilibrium structure becomes more complex, however our main results remain robust.

5. The highest bidder among those with $b_i \geq r$ wins the impression and pays $\max\{r, b^{(2)}\}$. If no bid meets the reserve, the impression goes unsold.

Strategies and Equilibrium. A strategy for the platform consists of three components: (i) a number of invitees $n \in \{1, \ldots, K\}$; (ii) a selection rule $\mathcal{S} : [0,1]^K \to \binom{[K]}{n}$ mapping the realized conversion probabilities to a set of exactly n advertisers; and (iii) a reserve-price rule $r : [0,1]^K \to \mathbb{R}_+$. A strategy for advertiser i maps each triple (θ_i, n, r) to a bid $b_i \geq 0$.

Because the platform's selection and reserve-price decisions may depend on $\mathbf{c}$, the observable auction design choices can be endogenously informative: an invited advertiser can rationally update its prior belief about its own conversion probability c_i based on the observed r and the anticipated selection rule $\tilde{\mathcal{S}}$. Although n is chosen before $\mathbf{c}$ is realized and therefore does not signal the realized conversion probabilities, it may affect the informativeness of invitation.[10] In equilibrium, these inferences must be consistent with the platform's actual behavior.

We analyze the game using the concept of perfect Bayesian equilibrium (PBE). A PBE consists of a strategy profile and a system of beliefs, specifying each invited advertiser's posterior distribution over its own conversion probability upon observing (n, r), such that: (i) the platform's choice of n maximizes ex ante expected revenue, and for each realization $\mathbf{c}$, its choices of $\mathcal{S}^*(\mathbf{c})$ and $r^*(\mathbf{c})$ maximize expected revenue given the advertisers' bidding strategy; (ii) for each (θ_i, n, r) and the advertisers' posterior beliefs about their own c_i, $g(\cdot \mid r, n, \tilde{\mathcal{S}})$, the bidding strategy β^* maximizes expected payoff; and (iii) invited bidders' posterior beliefs $g(\cdot \mid r, n, \tilde{\mathcal{S}})$ are derived from the platform's strategy via Bayes' rule wherever applicable. In general, the game admits multiple equilibria; we focus on the *platform-optimal* PBE, i.e., the equilibrium that maximizes the platform's expected revenue.

3 The platform-optimal equilibrium

It is well known that in a second-price auction, bidding one's expected value is a weakly dominant strategy, irrespective of beliefs about other bidders' strategies (Vickrey, 1961; Milgrom and Weber, 1982). Invited advertisers observe the reserve price r and the number of bidders n. They do not

[10] How n factors into the advertiser's belief updating is provided in Section 3.

observe the conversion probabilities c_i or the platform's selection rule $\mathcal{S}$, but hold beliefs $\tilde{\mathcal{S}}$ about the selection rule. In equilibrium, these beliefs are correct: $\tilde{\mathcal{S}} = \mathcal{S}^*$.

Since n is chosen before $\mathbf{c}$ is realized, it carries no information about the conversion probabilities. The reserve price r, however, is set after the platform observes $\mathbf{c}$ and could in principle convey information. Based on the fact that it has been invited, together with the observed r and the anticipated $\tilde{\mathcal{S}}$, each bidder i updates its belief about its own c_i. Two observations simplify this updating. First, in a second-price auction, only bidder i's own c_i is relevant for its bidding decision. Second, the ex ante symmetry of advertisers and the fact that all invited bidders receive the same signal (n, r) imply that they share a common posterior belief about their own c_i.

We denote this common posterior by $g(\cdot \mid r, n, \tilde{\mathcal{S}})$, a distribution over $[0, 1]$, and write $g(\cdot)$ for brevity when the conditioning variables are clear. Given this posterior, the optimal bid of advertiser i is its expected value for the impression:

$$b_i(\theta_i;\, r, n, \tilde{\mathcal{S}}) = \mathbb{E}[v_i \mid r, n, \tilde{\mathcal{S}}] = \bar{c}_g \cdot \theta_i, \tag{1}$$

where $\bar{c}_g := \mathbb{E}_g[c_i]$ is the posterior mean conversion probability. Although all bidders share the same belief about c_i, they submit different bids because their values per conversion θ_i are heterogeneous and privately observed.

From the platform's perspective, each advertiser's bid is a random variable with CDF $F_B(x) = x/\bar{c}_g$ on $[0, \bar{c}_g]$, since θ_i follows a standard uniform distribution. The platform's expected revenue, given n invited bidders and reserve price r, is therefore:

$$\begin{aligned} \pi(n, r;\, \bar{c}_g) &= r \cdot \Pr[X_{(n-1)} < r \le X_{(n)}] + \int_r^{\bar{c}_g} x \cdot f_{X_{(n-1)}}(x)\, dx \\ &= r\big(1 - F_B(r)^n\big) + \int_r^{\bar{c}_g} \big(1 - n\, F_B(x)^{n-1} + (n-1)\, F_B(x)^n\big)\, dx, \end{aligned} \tag{2}$$

where $X_{(k)}$ denotes the k-th smallest among n i.i.d. draws from F_B.[11]

The key feature of this revenue expression is the posterior mean $\bar{c}_g$, which governs the upper bound of the bid distribution. Without selection effects, the prior mean is $\mathbb{E}[c_i] = 1/2$. When bidders rationally account for the fact that they were selected, $\bar{c}_g$ exceeds $1/2$, leading to higher

[11]The first line holds by definition: the platform earns r if only the highest bid exceeds r, and earns the second-highest bid if both the highest and second-highest exceed r. The second line follows by integration by parts.

bids and greater platform revenue. Importantly, the revenue depends on the posterior only through its mean $\bar{c}_g$, not on the full distribution $g(\cdot)$. It also depends on the number of bidders n and on the anticipated selection strategy $\tilde{\mathcal{S}}$ (through $\bar{c}_g$), but not on the platform's actual strategy $\mathcal{S}$.

We now characterize the platform-optimal PBE. The analysis proceeds in three steps: we first derive the optimal reserve price for given beliefs, then show that the reserve price is uninformative in equilibrium, and finally determine the optimal selection rule and number of bidders.

The first result establishes that the platform's optimal reserve price takes a simple closed form, depending only on the posterior mean $\bar{c}_g$ and the number of bidders n.

Lemma 1 (Optimal reserve price)**.** *For a given number of bidders $n \geq 1$ and posterior mean $\bar{c}_g > 0$, the revenue function $\pi(n, r; \bar{c}_g)$ is strictly concave in r on $[0, \bar{c}_g]$, with unique maximizer $r^* = \bar{c}_g/2$. The maximized revenue is*

$$\pi^*(n; \bar{c}_g) = \bar{c}_g \cdot \alpha(n), \quad \textit{where} \quad \alpha(n) := \frac{(n-1) + \frac{1}{2^n}}{(n+1)}. \tag{3}$$

The optimal reserve is simply half the expected bid ceiling. This is a natural analogue of the classical Myerson optimal reserve for the uniform distribution, adapted to the setting where the bid distribution has support $[0, \bar{c}_g]$ rather than $[0, 1]$.

An important implication is that maximized revenue is linear in the posterior mean $\bar{c}_g$: any strategy that inflates bidders' beliefs about their conversion probability by a factor scales the platform's revenue by the same factor. The multiplier $\alpha(n) \in (0, 1)$ is strictly increasing in n, reflecting the competitive effect of additional bidders.

In our model, the reserve price is set after the platform observes $\mathbf{c}$ and can therefore serve as a signal about the conversion probabilities. The next result shows this cannot occur in the platform-optimal equilibrium.

Lemma 2 (No signaling through the reserve price)**.** *In the platform-optimal equilibrium, the reserve price is constant across all realizations of* $\mathbf{c}$. *The reserve price reveals no information about the conversion probabilities to invited advertisers.*

The intuition relies on the linearity established in Lemma 1. Suppose the platform tried to signal information through the reserve price, using different values of r for different realizations of $\mathbf{c}$.

Since the platform's actual revenue depends only on $(n, r, \bar{c}_g)$ and not on the realized $\mathbf{c}$ directly, the platform can freely choose any on-path reserve price regardless of the state. Optimality then requires that all on-path actions yield identical revenue; otherwise, the platform would deviate to the highest-revenue action in every state.

Now, signaling creates dispersion in posterior beliefs: some on-path actions induce beliefs above the unconditional mean $\bar{c}_g$, others below it. Because revenue is linear in beliefs, this amounts to a mean-preserving spread of the platform's payoff across actions. But the equal-revenue constraint pins the platform's payoff at the revenue achievable under the lowest posterior mean. Since the lowest posterior must be strictly below $\bar{c}_g$ whenever the signal is informative, the platform's revenue is strictly below what it would achieve by pooling. In other words, the platform would prefer to "hide" its information and keep bidders' beliefs at the unconditional level $\bar{c}_g$, which it can achieve by setting a constant reserve price.

With the reserve price pinned down, we turn to the platform's choice of which n advertisers to invite. Since the platform's revenue is linear in $\bar{c}_g$ (Lemma 1), the selection rule matters only through its effect on the posterior mean: the platform should select the rule that maximizes the expected conversion probability of invited bidders, conditional on selection.

Lemma 3 (Optimal selection rule)**.** *For a fixed number of bidders n, the selection rule that maximizes $\bar{c}_g$ is to invite the n advertisers with the highest conversion probabilities. Under this rule, the posterior mean is*

$$\bar{c}_g(n) = \frac{2K - n + 1}{2(K+1)}. \tag{4}$$

Moreover, the platform has no profitable deviation from this rule: since bids depend on the privately observed θ_i's (which are i.i.d. across all advertisers and independent of c_i), the platform's revenue is the same regardless of which n advertisers it actually selects, given fixed beliefs.

The first part of the result is straightforward: to maximize the posterior mean of a selected advertiser's c_i, the platform should select advertisers with the highest realized c_i's. The sum of the top n order statistics of any collection of numbers is maximized by choosing the n largest, and this pointwise optimality carries over to expectations.

The second part (incentive compatibility) follows from independence between the platform's private information ($\mathbf{c}$) and the advertisers' private information (θ_i). Because bids equal $\bar{c}_g \cdot \theta_i$, and the θ_i's are i.i.d. regardless of which advertisers are selected, the distribution of submitted bids (and hence the platform's revenue) is invariant to the actual selection. The platform is therefore indifferent across all selection rules of size n, making the top-n rule trivially sustainable.

Combining the three lemmas, the platform's optimization reduces to a single variable: the number of bidders n. The platform faces a fundamental trade-off. Inviting more bidders increases competition, captured by the increasing factor $\alpha(n)$, but dilutes the selection effect, as the posterior mean $\bar{c}_g(n) = (2K - n + 1)/(2(K+1))$ is decreasing in n. The optimal n^* balances these two forces.

Proposition 1 (Platform-optimal equilibrium)**.** *The platform-optimal PBE is unique (up to off-path beliefs and, when $\Pi(\cdot)$ has multiple maximizers, the choice among them) and is characterized as follows:*

1. *The platform invites the n^* advertisers with the highest conversion probabilities, where n^* maximizes*
$$\Pi(n) = \frac{2K - n + 1}{2(K+1)} \cdot \frac{(n-1) + \frac{1}{2^n}}{(n+1)} \tag{5}$$
over $n \in \{1, \dots, K\}$. The number of invitees n^* is independent of the realized $\mathbf{c}$.

2. *Each invited bidder's posterior mean conversion probability is $\bar{c}_g = \frac{2K-n^*+1}{2(K+1)}$.*

3. *The platform sets the reserve price to $r^* = \frac{2K-n^*+1}{4(K+1)}$, constant across all realizations of $\mathbf{c}$.*

4. *Each invited bidder bids $\beta^*(\theta_i) = \theta_i \cdot \frac{2K-n^*+1}{2(K+1)}$.*

Next, we focus on the optimal number of invited advertisers n^*. We start with a general Proposition 2 and then we identify the optimal n^* in Proposition 3.

Proposition 2 (Optimal number of bidders)**.** *Optimal number of bidders: For $K \leq 3$, it is $n^*(K) = K$ (i.e., the platform invites all advertisers), whereas for $K \geq 4$, it is $n^*(K) < K$ (i.e., the platform will not allow all advertisers to participate in the auction).*

The main takeaway of Proposition 2 is that for $K \geq 4$ the platform finds it optimal to select a strict subset of all advertisers to participate in the auction. This is because the platform can induce greater bids from the advertisers when it keeps the group of participating bidders more selectively—more precisely, providing the invited advertisers with the perception of belonging to a more selective group. More specifically, upon being invited, advertiser i updates belief about its c_i being among the top n^* advertisers' c_i's and thus increases its expected value from the auction (and its bid). So, the platform faces a tradeoff. Holding everything else constant, inviting an additional advertiser to the auction increases its revenue (i.e., the second highest bid) in expectation. However, keeping the group of bidders more selective improves the invited advertisers' inferences and, consequently, their bids. For a small enough $K(\leq 3)$, the benefit of including all advertisers in the auction outweighs the benefit of inducing positive inferences from advertisers, and therefore the platform will invite all advertisers to the auction by setting $n^*(K) = K$. However, for $K \geq 4$, the platform is better off inviting a strict subset of advertisers.

Next, we characterize the platform's optimal number of bidders $n^*(K)$ for a given number of advertisers. This analysis is a priori not trivial because n^* has to be a non-negative integer and thus cannot entirely rely on the first-order condition of the platform's profit function. Moreover, even if we start with the first-order condition, it does not offer a clear explicit formula for n^*. Nevertheless, the next lemma further characterizes $n^*(K)$.

Proposition 3 (Optimal n^*)**.** *Let $\{p_k\}_{k\geq 0}$ be the sequence defined by the formula:*

$$p_k = \begin{cases} \frac{1}{4}(k^2+3k-4) & \text{, if } k \geq 8 \text{ and } (k \bmod 4) \in \{0,1\}, \\ \left\lfloor \frac{1}{4}(k^2+3k) \right\rfloor & \text{, otherwise.} \end{cases}$$

The sequence $\{p_k\}_{k\geq 0}$ is strictly increasing. Then, for any integer $K \geq 1$, the optimal $n^*(K)$ is the unique integer k such that $p_{k-1} < K \leq p_k$.[12] *Using asymptotic notation, it is $n^*(K) \sim 2\sqrt{K}$.*

Given the optimal $n^*(K)$, Corollary 1 characterizes the platform's profit, the advertisers' surplus, and the social welfare. The platform's profit is the winner's payment, which is the maximum of the second highest bid and the reserve price. Note that the surplus of all advertisers except the winner of the auction is zero. Thus, the total advertiser surplus is the surplus of the winner, which is the highest value minus the payment. Thus, the social welfare is the expected value of the winner.

Corollary 1. *The expected equilibrium profit of the platform, the total surplus of the advertiser and*

[12] The first few terms of the sequence $\{n^*(K)\}_{K\geq 1}$ are as follows: $1, 2, 3, 3, 4, 4, 4, 5, 5, 5, 6, 6, 6, 7, 7, 7, 7, 8, 8, 8, 8, 9, \ldots$.

the social welfare are as follows:

$$\Pi^*(n^*,K) = \frac{(2K-n^*+1)\left(n^*-1+\frac{1}{2^{n^*}}\right)}{2(n^*+1)(K+1)}, \quad A^*(n^*,K) = \frac{(2K-n^*+1)\left(1-\frac{n^*}{2^{n^*+1}}-\frac{1}{2^{n^*}}\right)}{2(n^*+1)(K+1)},$$

$$SW^*(n^*,K) = \Pi^*(n^*,K) + A^*(n^*,K) = \frac{n^*(2K-n^*+1)\left(1-\frac{1}{2^{n^*+1}}\right)}{2(n^*+1)(K+1)}.$$

Provided with the social welfare, we can characterize the optimal social number of bidders n^{*s}, the details of which are relegated to Lemma 10 in Appendix A.1. The next corollary shows that the platform-optimal number of bidding advertisers is greater than the socially optimal number.

Corollary 2. *It is $n^{*s} \leq n^*$ (with a strict inequality for $K \geq 8$). In other words, the number of advertisers that maximizes social welfare is lower than the number of advertisers that maximizes revenue.*

The result that $n^{*s} \leq n^*$ shows a difference in how the platform and the social planner value competition. Both the platform's revenue and social welfare share the same signaling component: both are proportional to the posterior mean $\bar{c}_g(n)$, which is decreasing in n. Thus, both the platform and the social planner face the same cost of adding an additional bidder (diluting the informativeness of solicitation). However, they differ in the marginal benefit of competition. The platform's revenue depends heavily on the second-highest bid (the competitive pressure that determines the winner's payment), while social welfare depends on the highest bid, i.e., the value generated by the winner.

Adding a bidder has a larger marginal impact on the second-highest order statistic than on the first-highest. Formally, the competition multiplier $\alpha(n)$ (which governs revenue) increases faster in n than the efficiency multiplier $\frac{n}{n+1}\left(1-\frac{1}{2^{n+1}}\right)$ (which governs welfare).[13] Because the platform benefits more from each additional competitor, it optimally invites more bidders than the social planner would, tolerating a weaker signal in exchange for stronger competitive pressure on payments.

4 Value of information and other benchmark models

In this section, we analyze models of alternative information settings that will help us understand the value of information in various scenarios. We label each model by the information available to

[13] For large n, the marginal gain from competition for the platform is approximately $2/(n+1)^2$, whereas the marginal gain for social welfare is approximately $1/(n+1)^2$.

the platform and the advertisers, respectively.[14]

Benchmark 1 (Fully-informed platform, $P_{\text{ALL}}A_{\Theta}$)**.** The platform observes both $\mathbf{c} = (c_1, \ldots, c_K)$ and $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_K)$, while each advertiser i observes only θ_i (as in the main model). This benchmark allows us to assess the value to the platform of learning the advertisers' private information.

Benchmark 2 (Fully-informed advertisers, $P_C A_{\text{ALL}}$)**.** The platform observes $\mathbf{c}$ as in the main model, while each advertiser i observes both θ_i and c_i, and hence knows its value $v_i = \theta_i c_i$. This benchmark allows us to assess the value to advertisers of learning their own conversion probabilities. Because each advertiser knows its value, the auction operates as a standard second-price auction with independent private values, however the reserve price can depend on $\mathbf{c}$.

Benchmark 3 (Full information, $P_{\text{ALL}}A_{\text{ALL}}$)**.** The platform observes both $\mathbf{c}$ and $\boldsymbol{\theta}$, and each advertiser i observes both c_i and θ_i. This benchmark provides a full-information reference point in which all market participants are maximally informed.

4.1 Fully-informed platform and partially-informed advertisers

Lemma 4 (Model $P_{\text{ALL}}A_{\Theta}$)**.** *In a model where the platform knows both c_i's and θ_i's of all advertisers, the platform cannot convey any information about the advertisers' c_i's through its choice of n. It is weakly dominant for the platform to invite all K advertisers to the auction. Then, the platform sets the reserve price so that it extracts the full surplus of the advertiser with the highest expected value. Specifically,*

1. *The advertiser's equilibrium bidding function is $\beta^*_{P_{\text{ALL}}A_{\Theta}}(\theta_i) = \theta_i/2$;*

2. *The platform sets the reserve price to $r^*_{P_{\text{ALL}}A_{\Theta}} = \max_i \theta_i/2$;*

3. *The platform's profit is $\Pi^*_{P_{\text{ALL}}A_{\Theta}} = K/(2(K+1))$ and the advertiser surplus $A^*_{P_{\text{ALL}}A_{\Theta}} = 0$.*

In this scenario, both the platform and each advertiser know the advertiser's θ_i, whereas only the platform knows the advertiser's c_i. In contrast to the main model, the platform's equilibrium selection strategy cannot depend on the realized c_i's and thus cannot convey information about their c_i's through its observable choice of n bidders. This is because even if the advertisers believed that

[14] For instance, $P_C A_{\Theta}$ denotes the main model in which the platform observes $\mathbf{c}$ and each advertiser observes θ_i.

the platform's selection strategy depended on the c_i's, the platform has a profitable deviation to select advertisers with the highest commonly known θ_i's, not c_i's. In other words, once the platform knows θ_i, it cannot commit to not using this information to select the set of bidders in the auction, irrespective of what the advertisers believe about the platform's strategy.

To understand this result, let $\bar{c}$ be an invited bidder's posterior expectation for c_i. Then, its optimal bidding follows $\beta_i(\theta_i) = \bar{c}_g \cdot \theta_i$. This looks the same as Equation (1), but a key difference is that, in this model, θ_i is commonly observable to the platform and the advertiser i. Since $\bar{c}$ remains unchanged for any deviation by the platform, these bids are deterministic. Thus, the platform will induce maximum bids from the n bidders if and only if it selects n bidders with the highest θ_i values. Thus, in equilibrium, no information about c_i's can be credibly conveyed to the advertisers through inviting certain advertisers to the auction.

Given the deterministic bids, the advertiser with the highest θ_i will win the auction, and this advertiser's value (and the optimal bid) is known to the platform. Thus, the optimal strategy for the platform is to extract the full surplus of this advertiser by setting a reserve price equal to its value $\max_i \theta_i/2$ (or an infinitesimal amount below this level). Thus, the advertiser surplus is zero, and the platform's expected profit is $\mathbb{E}[\max_i \theta_i/2] = K/(2(K+1))$ as θ_i follows a standard uniform distribution on $[0,1]$.

We now compare the payoffs of this model with those of the main model.

Proposition 4. *When compared to the main model, it holds that* $A^*_{P_C A_\Theta} \geq A^*_{P_{\text{ALL}} A_\Theta}$ *for all* K *and* $\Pi^*_{P_C A_\Theta} \geq \Pi^*_{P_{\text{ALL}} A_\Theta}$ *for* $K \geq 8$. *Thus, the marginal value of the information about the advertisers'* θ_i *to the platform is negative for* K *large enough.*

Given the previous discussion, it is straightforward that the advertisers are worse off when the platform knows all compared to the main model with two-sided information asymmetry. More interestingly, additional information can backfire and harm the platform. Therefore, when there is a sufficiently high number of advertisers, it is better for the platform not to learn the advertisers' θ_i's even if obtaining this information is free. Thus, when the default environment is the main model denoted by $P_C A_\Theta$, the marginal value of the θ_i–information is negative for the platform if $K \geq 8$.

It can be shown that the platform's equilibrium payoff increases faster in K in the main model than in this benchmark model. In the main model, for K large enough, even if the platform left out some advertisers from the auction, the platform would still have enough number of bidders and enjoy the benefit from the bidders making positive inferences from being selected. Thus, for $K \geq 8$, the platform's payoff in the main model is greater than that in this model with the platform's access to additional information about θ_i.

4.2 Partially-informed platform and fully-informed advertisers

Next, we analyze a different model $P_C A_{\text{ALL}}$ in which each advertiser now knows both its c_i and θ_i. By comparing the equilibrium of this model with that of the main model, we can identify the marginal effect of the access of advertisers to additional information about their θ_i information.

Each advertiser now knows its own c_i. Thus, no information is transmitted from the platform to the advertisers, which implies that the platform wants to invite all K advertisers as bidders. This leads to another distinction that the optimal reserve price can now be a function of the realized c_i's.[15]

More specifically, in $P_C A_{\text{ALL}}$, we essentially have an auction with K advertisers with different upper bounds in their valuations; from the platform's perspective, advertiser i's valuation is uniformly distributed in $[0, c_i]$. In the proof of Lemma 5, we characterize the process through which the optimal reserve $r^*_{P_C A_{\text{ALL}}}$ is exactly pinned down for the general case with different realized values of c_i's. Subsequently, Proposition 5 compares the payoffs of this model with that of the main model.

Lemma 5 (Model $P_C A_{\text{ALL}}$)**.** *In equilibrium, the platform will invite all K advertisers and set the reserve price to $r^*_{P_C A_{\text{ALL}}}$, which is determined by the realized values of c_i's according to Algorithm 1 in Appendix A.1. Advertisers' bidding function is $\beta_{P_C A_{\text{ALL}}}(\theta_i, c_i) = \theta_i \cdot c_i$.*

Proposition 5. *Comparing the main model $P_C A_\Theta$ and model $P_C A_{ALL}$, we have $\Pi^*_{P_C A_\Theta} \leq \Pi^*_{P_C A_{\text{ALL}}}$ (for $K \geq 8$) and $A^*_{P_C A_\Theta} \leq A^*_{P_C A_{\text{ALL}}}$. Thus, for $K < 8$, the advertisers' access to additional information about c_i's benefits only the advertisers, whereas for $K \geq 8$, it benefits both the platform and advertisers.*

Proposition 5 implies that given the main model as the default setting—with the two-sided information asymmetry—the advertisers are always better off with additional information about

[15] In the main model, too, potentially the reserve price could also depend on c_i's as a signaling device, but we showed that the platform-optimal equilibrium involves a constant reserve price because signaling through the reserve is dominated by pooling.

c_i's. The effect is more nuanced for the platform, but nevertheless it is interesting that it can also benefit the platform. That is, the platform may sometimes want to provide information about c_i's to the advertisers even if it means that its access to information becomes inferior to that for the advertisers. Thus, information disadvantage does not necessarily lead to a worse situation for the platform.

This is because the advertisers' additional access to the c_i information comes with an economic tradeoff. On the one hand, it makes all advertisers fully informed about their valuations, which they bid. Also, it eliminates the credibility problem, and this enables the platform to invite all K advertisers. On the other hand, in the main model the platform excludes some advertisers, and this exclusion creates signaling value that causes each invited bidder to bid more aggressively than they would in the $P_C A_{\text{ALL}}$ model. As K increases, the former effect acting positively for the $P_C A_{\text{ALL}}$ model outweighs the latter negative effect, thus leading to an existence of a cutoff for K above which the platform's payoff is higher for any K above it.

4.3 Fully-informed platform and fully-informed advertisers

We now analyze the last benchmark model of this section, the complete information environment in which the advertiser's c_i and θ_i are known to both the advertiser and the platform. There is no room for the advertisers making inferences from any observation–including the reserve price r, the number of invited bidders n. The platform knows each advertiser's exact value to be $v_i = \theta_i \cdot c_i$. Thus, it will extract surplus of the advertiser with the highest value, i.e., $\max_i \theta_i \cdot c_i$, by setting the reserve price equal to this value.

Lemma 6 (Model $P_{\text{ALL}} A_{\text{ALL}}$)**.** *In equilibrium, the platform will set the reserve price equal to* $\max_i \theta_i \cdot c_i$*. The advertiser's bidding function is* $\beta_{P_{\text{ALL}} A_{\text{ALL}}}(\theta_i, c_i) = \theta_i \cdot c_i$*. The platform's equilibrium payoff is* $\Pi^*_{P_{\text{ALL}} A_{\text{ALL}}} = 1 - \sum_{k=0}^{K} \frac{K!}{(K-k)!(K+1)^{k+1}}$*, and the advertiser surplus* $A^*_{P_{\text{ALL}} A_{\text{ALL}}} = 0$.

The next proposition examines the value of marginal information for the platform and the advertisers when their counterpart knows all.

Proposition 6. *It holds that* $\Pi^*_{P_C A_{\text{ALL}}} \leq \Pi^*_{P_{\text{ALL}} A_{\text{ALL}}}$ *and* $A^*_{P_{\text{ALL}} A_{\Theta}} = A^*_{P_{\text{ALL}} A_{\text{ALL}}} = 0$*. That is, if the advertisers know all, the platform is better off with additional information about* $\boldsymbol{\theta}$*. On the* *other hand, if the platform knows all, the advertisers are indifferent whether they have additional information about c_i.*

The proposition shows that the marginal value of information can be meaningfully different depending on how much information the counterpart has. When advertisers know all about their values, the platform is better off having additional information about θ. Recall that the same wasn't true in Proposition 4 when the platform faced advertisers that know their θ_i's only. Likewise, the result about the value of information to the advertisers is not the same as before. Here, facing an ad platform that knows all, the advertisers are indifferent whether they have access to additional information, and thus the value of additional information about c_i's is zero. However, in Proposition 5 when the platform only knows the advertisers' c_i's, the value of information about c_i's to the advertisers was positive. Thus, these results together imply that each party's strategic decision for acquiring additional information and accessing economic value of information, the overall information structure in the market—i.e., what other parties know and what they don't—should also be considered. Also, the debate and investigation for how to regulate data transparency—e.g., which type of information and whose information to be made public and accessible to other stakeholders—must take into account the current status of information structure and the stakeholders' strategic incentives to acquire new information or to share their private information.

4.4 Information acquisition and information disclosure

The next result is based on analysis of a game where the platform and advertisers could choose the set of information to which they have access. More specifically, the platform can choose between c_i's and all information, whereas the advertisers can choose between θ_i's and all information.

Corollary 3 (Information acquisition)**.** *In a game in which both the platform and the advertisers can choose its information set between (C, ALL) and (Θ, ALL), respectively, pure-strategy Nash equilibrium is identified as follows: For $K \geq 8$, $(ALL,\ ALL)$ is the unique equilibrium. For $K < 8$, $(ALL,\ ALL)$ and (ALL, Θ) are equilibria. In any equilibrium, the advertiser surplus is always zero.*

Thus, in any equilibrium of this game, the platform will obtain all information and secure positive surplus, whereas the advertiser surplus will be fully extracted.

Next, we analyze a game in which each party simply decides whether to reveal its private information. So, the platform chooses whether to share its C information with advertisers and the advertisers Θ information with the platform.

Corollary 4 (Information disclosure)**.** *In a game where the platform and the advertisers decide whether to share their private information with the other party, for $K \geq 8$ the unique equilibrium is the platform sharing its information but the advertisers don't, i.e., $P_C A_{ALL}$. For $K < 8$, neither the platform nor the advertisers share their private information, thus leading to the environment $P_C A_\Theta$.*

It is direct from the previous corollary that the advertisers do not want to share their private information with the platform because, if they did, that would make the platform all-knowing and their surplus will then be fully extracted. Given that the advertisers will keep their information private, Proposition 5 guides us what the platform's equilibrium strategy will be. For $K \geq 8$, the platform will share information about c_i's, which mitigates information asymmetry about it, which makes the auction outcome more efficient. Moreover, it allows the platform to invite all K advertisers and thus potentially make the bids more competitive. Otherwise, if $K < 8$, then the platform will choose not to share its private information and rely on selective invitation to the auction as a credible signal and consequently induce higher bids from the invited advertisers.

Unlike the information acquisition game analyzed in the previous corollary, the information disclosure game results in equilibrium with positive surplus for the advertisers. This indicates that advertisers will be more invested in keeping their data private than the platform will be about its data. This result provides implications for the regulators in that their efforts to limit data-sharing activities without user consent could be more concentrated on the behaviors of the platforms than on the advertisers.

5 Impact of commitment in ad auctions

In this section, we investigate a model in which the platform can commit to a strategy as a mapping from a realization of $\mathbf{c} = (c_1, c_2, ..., c_K)$ to a set of bidders participating in the ad auction and a reserve price.

Benchmark 4 (Platform commitment, $P_C^{\text{Com}} A_\Theta$)**.** The information structure is identical to the main model: the platform observes $\mathbf{c}$ and each advertiser i observes θ_i. However, the platform can publicly commit, before the realization of $(\mathbf{c}, \boldsymbol{\theta})$, to a selection rule and a reserve-price rule, both of which may condition on the subsequent realization of $\mathbf{c}$.[16] Because the committed rules are publicly known, invited advertisers can directly compute their posterior beliefs about c_i without strategic inference, eliminating the signaling effects present in the main model. This benchmark allows us to understand the impact of commitment for the platform.

Unlike in the main model, the platform's committed strategy is publicly observable, so the platform cannot unknowingly deviate to an alternative strategy. We aim to characterize the equilibrium in this game and then compare it with the equilibrium of the main model without commitment, which will help us understand the impact of commitment in the platform's strategy for choosing the bidders based on the information it can access.

5.1 Analysis of commitment equilibria

The set of feasible strategies in this case can be larger than in the case without commitment because a strategy can be any mapping from the realized value of c_i's without having to worry about credibility. Nevertheless, it must be true that, all else equal, the platform prefers an advertiser with a higher c_i to participate in the auction. Thus, the optimal selection strategy of the platform must be monotone in the sense that an advertiser with a higher c_i is more likely to become a bidder.

Note that the equilibrium selection strategy of the non-commitment case in Section 3 is a type of threshold strategy, where the threshold is determined by the rank order of realized c_i's, i.e., a relative threshold. There exists another type of cutoff strategy where the threshold is an absolute level τ. For example, $\tau = 0$ results in all K advertisers participating in the auction with probability 1; $\tau = 1$ in no bidder with probability 1. Otherwise, $\tau \in (0, 1)$ results in a distribution over the number of bidders and the distribution over c_i conditional on being selected as a bidder.

We restrict our attention to these two classes of threshold strategies: one cutoff strategy based on rank order of realized c_i's and another based on an absolute cutoff τ.[17] The former, we refer to

[16] To make the game sequence under commitment more explicit, steps 1 and 3 in the game sequence of Section 2.1 are combined together as step 1, and the platform can commit on the functions $\mathcal{S}(\mathbf{c})$ and $r(\mathbf{c})$. Note also that in benchmark $P_C^{\text{COM}} A_\Theta$ the number of invited bidders n is not necessarily fixed and can depend on the realized $\mathbf{c}$.

[17] In general, a selection strategy can take a more complicated form. For instance, it can combine these two classes of strategies for different sets of realizations of c_i's. The analysis of general strategies is beyond the scope of this research and is reserved for future research.

as "top n-strategy," has already been analyzed in Section 3.

In the latter class of strategy, which we call a threshold strategy with a cutoff τ, upon being included as a bidder, the advertiser's posterior distribution for its unobserved c_i is a uniform distribution on $[\tau, 1]$. Thus, each bidder's optimal bidding is $\beta(\theta_i) = \theta_i \cdot (1+\tau)/2$. The probability that the number of bidders will be n is precisely $\binom{K}{n} \cdot (1-\tau)^n \cdot \tau^{(K-n)}$. This provides us with the distribution of the number of bidders in the auction, which allows us to compute the platform's expected payoff from the auction. The platform will then choose the cutoff τ to maximize its expected payoff.

Lemma 7. *Given a threshold strategy with the absolute cutoff τ^*,*

1. *The advertisers' equilibrium bidding strategy is $\beta(\theta_i) = \theta_i \cdot (1+\tau^*)/2$.*
2. *The platform's optimal reserve price is $r^* = (1+\tau^*)/4$.*
3. *The platform's expected payoff is $\Pi^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau^*) = \frac{1+\tau^*}{1-\tau^*}\left[\frac{K}{K+1} - \frac{1+\tau^*}{2} + \frac{1}{K+1}\left(\frac{1+\tau^*}{2}\right)^{K+1}\right]$.*
4. *The advertiser surplus is $A^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau^*) = \frac{1+\tau^* - \left(\frac{1+\tau^*}{2}\right)^{K+1}(2+K(1-\tau^*))}{2(K+1)(1-\tau^*)}$.*
5. *The social welfare is $SW^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau^*) = \frac{1+\tau^*}{2}\left[1 - \frac{1}{2}\left(\frac{1+\tau^*}{2}\right)^K + \frac{\left(\frac{1+\tau^*}{2}\right)^{K+1} - 1}{(K+1)(1-\tau^*)}\right]$.*

The equilibrium cutoff level τ^* satisfies the first-order condition $\partial \Pi^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau)/\partial\tau = 0$ when evaluated at $\tau = \tau^*$.

By the virtue of commitment, it is straightforward that the platform is weakly better off with commitment than without commitment. Given that the "top n advertisers" based on realized c_i's is the platform-optimal strategy in an environment without commitment, the equilibrium payoff of the platform in the commitment case (among all possible commitment strategies) must be weakly higher than the "top n advertisers" strategy. However, whether the platform will obtain a higher payoff in an equilibrium within a particular class of strategies—particularly, an absolute cutoff strategy—compared to the optimal "top n-strategy" is not so trivial. Nevertheless, the result below shows that the absolute cutoff-strategy equilibrium with commitment is always better for the platform than the optimal top n-strategy (which also is the equilibrium without commitment).

Lemma 8. *Consider the following two selection strategies for the platform:*

- *A threshold-τ strategy: The platform invites all advertisers with $c_i \geq \tau$.*

- *A top n-strategy: The platform invites the advertisers with the top n highest values of c_i.*

Let $\tau^* \in [0,1]$ be the optimal threshold for the threshold-τ strategy and $n^* \in [1, K]$ be the optimal integer for the top-n strategy. When there is commitment, the platform will always choose the threshold-τ^* strategy rather than the top-n^* strategy.

5.2 Comparison of commitment value vs. information value

Finally, we compare the value of commitment with the value of a better information structure to the platform. From the previous analysis provided with no commitment, we know that the platform is better off sharing its private information about c_i than acquiring additional information about θ_i (for large enough K).

Lemma 9. *The ability of the platform to commit on a selection rule leads to higher revenue than its ability to truthfully communicate the values of c_i to the advertisers. In other words, $\Pi^*_{P_C^{\text{COM}} A_\Theta} \geq \Pi^*_{P_C A_{\text{ALL}}}$.*[18]

This result finds that the value of commitment to the platform can be substantial, and it can be more valuable than sharing its private information with the advertisers, which results in a better information environment without commitment. This result is more interesting because in the (C, ALL) environment, the advertisers' c_i information is public, and the platform is indifferent whether it has commitment or not. And, we show that under commitment, the platform is weakly better off facing partially-informed advertisers than fully-informed advertisers. This is because there exists a reserve-pricing strategy to which the platform can commit that reveals $\mathbf{c}$ precisely, and the information environment becomes (C, ALL). This result is in contrast to the non-commitment case in Proposition 5 which shows that the platform can be better off facing all-knowing advertisers for thick markets. In other words, the platform would only want to share its private information if it did not have the commitment.

[18] Note that Lemma 9 relies on the platform using an unrestricted commitment strategy. $\Pi^*_{P_C^{\text{COM}} A_\Theta}$ is the optimal revenue over all commitment strategies, which can be higher than $\Pi^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau^*)$, the optimal within the threshold class. In fact, as we can see in Appendix A.2, for large enough K, $\Pi^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau^*)$ is lower than $\Pi^*_{P_C A_{\text{ALL}}}$.

6 Conclusion

This paper studies ad auctions in an opaque information environment characterized by two-sided information asymmetry and the platform's lack of commitment power. In this environment, we examine how the platform's private information can be credibly transmitted to the advertisers through the platform's auction designs as observable signals. Moreover, we investigate how this mechanism affects the value of information for the platform and the advertisers.

Our analysis produces several core findings. First, selective invitation to the auction functions as an endogenous signaling mechanism. Because advertisers can rationally infer that their conversion probability is likely to be high conditional on being invited, a more exclusive group of bidders leads to more aggressive bidding. This creates a tradeoff for the platform between competition (more bidders) and inference (higher individual bids), and we show that for markets with four or more advertisers, the platform optimally excludes some advertisers from the auction. Second, the value of information to each party depends on the overall information structure, i.e. on what the other side of the market knows. Surprisingly, the platform can be harmed by learning the advertisers' private valuations, because doing so undermines the credibility of its selective invitation as a signal about conversion probabilities. Conversely, when advertisers gain access to their own conversion probabilities, both sides of the market can benefit, as the resulting transparency enables more efficient allocation and eliminates the need for exclusionary signaling. Third, commitment power is a strong strategic instrument for the platform. The ability to commit to a selection rule, even within the simple class of absolute threshold strategies, strictly dominates the platform-optimal equilibrium without commitment, and its value exceeds that of voluntarily sharing conversion probability data with advertisers.

These findings have several implications for practice and for the ongoing regulatory debate surrounding online advertising markets. For ad platforms, our results formalize the economic logic behind common practices such as traffic shaping and quality-score thresholds: by curating which advertisers see a given impression, platforms shape advertisers' beliefs and their willingness to pay. Our analysis suggests that the design of invitation policies should be viewed as a core element of auction design that interacts with information provision and reserve pricing. The result that commitment power is highly valuable also provides a strategic rationale for platforms to invest in transparent, rule-based allocation mechanisms, because credible commitment to a selection rule can dominate opaque, discretionary practices in terms of revenue.

For advertisers, our results highlight the asymmetric stakes of data transparency. Advertisers are consistently better off keeping their private valuations confidential, whereas gaining access to their own conversion data is always (weakly) beneficial. This asymmetry suggests that industry initiatives aimed at increasing advertiser access to platform-side engagement data, such as conversion-probability estimates or quality scores, may benefit the entire market, including the platform itself when the number of competing advertisers is sufficiently large. By contrast, advertiser-side data sharing with platforms, e.g. disclosing profit margins or customer lifetime values, tends to erode advertiser surplus.

For policymakers and regulators, our analysis offers a unique perspective on data transparency mandates in online advertising. A blunt requirement that all information be shared among all parties does not necessarily improve welfare; indeed, full transparency enables the platform to extract the entire advertiser surplus. Rather, the welfare consequences of transparency depend on which information is disclosed and to whom. Our information disclosure game shows that in equilibrium, the platform may voluntarily share conversion data while advertisers rationally withhold their valuations, and this asymmetric outcome can support positive surplus for both sides. These findings suggest that regulatory efforts might be most productively directed at facilitating platform-to-advertiser data flows, such as requiring platforms to share engagement metrics or predicted conversion rates, rather than mandating blanket transparency or restricting all forms of data use indiscriminately. Furthermore, the finding that opacity in the platform's selection process has real economic costs, relative to a commitment benchmark, lends support to policies that promote auditability and rule-based governance of algorithmic ad allocation.

Our analysis has several limitations that point to natural directions for future research. We study a single ad impression in isolation, abstracting from dynamic considerations such as learning, reputation, and budget constraints that may arise in repeated environments. Under commitment, we focus on two natural classes of selection strategies—top-n and absolute-threshold rules—and show that the threshold rule dominates within this class, while a full characterization of the optimal commitment mechanism lies beyond the paper's scope. For tractability, we also assume uniform and independent distributions, which limit the heterogeneity captured by the model; richer distributional structures, including correlation between conversion probabilities and valuations, may generate additional insights. Finally, the model treats advertisers symmetrically, abstracting from differences in sophistication, data access, and strategic awareness.

Despite these limitations, this paper contributes to our understanding of opaque ad auctions by showing that auction-design decisions—especially bidder selection—can serve as credible channels of information transmission. The findings also highlight the subtlety of information value: whether information benefits the platform or advertisers depends on the surrounding information structure, the platform's commitment ability, and the thickness of the advertiser side of the market.

References

Ada, S., Abou Nabout, N., and Feit, E. M. (2022). Context information can increase revenue in online display advertising auctions: Evidence from a policy change. *Journal of Marketing Research*, 59(5):1040–1058.

Bajari, P. and Hortaçsu, A. (2003). The winner's curse, reserve prices, and endogenous entry: Empirical insights from ebay auctions. *RAND Journal of Economics*, 34(2):329–355.

Bergemann, D., Heumann, T., Morris, S., Sorokin, C., and Winter, E. (2022). Optimal information disclosure in classic auctions. *American Economic Review: Insights*, 4(3):371–388.

Bergemann, D. and Morris, S. (2019). Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95.

Cai, H., Riley, J., and Ye, L. (2007). Reserve price signaling. *Journal of Economic Theory*, 135(1):253–268.

Chakraborty, T., Even-Dar, E., Guha, S., Mansour, Y., and Muthukrishnan, S. (2010). Selective call out and real time bidding. In *International Workshop on Internet and Network Economics*, pages 145–157. Springer.

Chen, Y., Kartik, N., and Sobel, J. (2008). Selecting cheap-talk equilibria. *Econometrica*, 76(1):117–136.

Choi, W. J. and Sayedi, A. (2023). Open and private exchanges in display advertising. *Marketing Science*, 42(3):451–475.

Choi, W. J. and Sayedi, A. (2024). Agency market power and information disclosure in online advertising. *Marketing Science*, 43(6):1279–1298.

Coviello, D., Guglielmo, A., and Spagnolo, G. (2018). The effect of discretion on procurement performance. *Management Science*, 64(2):715–738.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451.

De Corniere, A. and De Nijs, R. (2016). Online advertising and privacy. *The RAND Journal of Economics*, 47(1):48–72.

Despotakis, S. and Yu, J. (2023). Multidimensional targeting and consumer response. *Management Science*, 69(8):4518–4540.

Elmaghraby, W. J. (2005). The effect of asymmetric bidder size on an auction's performance: Are more bidders always better? *Management Science*, 51(12):1763–1776.

eMarketer (2026). Faq on programmatic advertising: Keeping up with automated ad buying. https://www.emarketer.com/content/faq-on-programmatic-advertising--keeping-up-with-automated-ad-buying. Accessed June 14, 2026.

Google Authorized Buyers (2026). Callout quota system. https://developers.google.com/authorized-buyers/rtb/callout-quota-system. Google for Developers, last updated April 14, 2026; accessed June 4, 2026.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Lauermann, S. and Wolinsky, A. (2017). Bidder solicitation, adverse selection, and the failure of competition. *American Economic Review*, 107(6):1399–1429.

Levin, D. and Smith, J. L. (1994). Equilibrium in auctions with entry. *American Economic Review*, 84(3):585–599.

Li, T. and Zheng, X. (2009). Entry and competition effects in first-price auctions: Theory and evidence from procurement auctions. *Review of Economic Studies*, 76(4):1397–1429.

Lipnowski, E. and Ravid, D. (2020). Cheap talk with transparent motives. *Econometrica*, 88(4):1631–1660.

Milgrom, P. R. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122.

Ning, Z. E., Shin, J., and Yu, J. (2025). Targeted advertising as implicit recommendation: Strategic mistargeting and personal data opt-out. *Marketing Science*, 44(2):390–410.

Shin, J. and Yu, J. (2021). Targeted advertising and consumer inference. *Marketing Science*, 40(5):900–922.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37.

A Appendix

A.1 Analysis and Proofs

Proof of Lemma 1

Substituting $F_B(x) = x/\bar{c}_g$ into Equation (2) and performing the change of variable $\rho = r/\bar{c}_g$, $t = x/\bar{c}_g$:

$$\begin{aligned}
\pi &= \bar{c}_g \left[\rho(1-\rho^n) + \int_\rho^1 \left(1 - n\,t^{n-1} + (n-1)\,t^n\right) dt\right] \\
&= \bar{c}_g \left[\rho - \rho^{n+1} + (1-\rho) - (1-\rho^n) + \frac{n-1}{n+1}(1-\rho^{n+1})\right] \\
&= \bar{c}_g \left[\frac{n-1}{n+1} + \rho^n - \frac{2n}{n+1}\,\rho^{n+1}\right].
\end{aligned}$$

Differentiating with respect to ρ:

$$\frac{\partial \pi}{\partial \rho} = \bar{c}_g \cdot n\,\rho^{n-1}(1-2\rho).$$

For $\rho \in (0,1)$, the unique zero is $\rho^* = 1/2$. The second derivative $\partial^2\pi/\partial\rho^2\big|_{\rho=1/2} = -\bar{c}_g \cdot n/2^{n-2} < 0$ confirms a strict global maximum on $[0,1]$. Evaluating at $\rho^* = 1/2$:

$$\pi^* = \bar{c}_g \left[\frac{n-1}{n+1} + \frac{1}{2^n} - \frac{n}{(n+1)\,2^n}\right] = \bar{c}_g \cdot \frac{(n-1)\,2^n + 1}{(n+1)\,2^n} = \bar{c}_g \cdot \alpha(n).$$

To verify that $\alpha(n)$ is strictly increasing, note that $\alpha(n) = 1 - 2/(n+1) + 1/((n+1)\,2^n)$. As n increases, both $2/(n+1)$ and $1/((n+1)\,2^n)$ decrease, but the former effect dominates. Formally, $\alpha(n+1) - \alpha(n) = \frac{2^{n+2}-n-3}{(n+1)(n+2)2^{n+1}} > 0$ for all $n \geq 1$. ■

Proof of Lemma 2

Suppose, toward a contradiction, that a platform-optimal equilibrium has $M \geq 2$ distinct on-path reserve prices $r_1, \ldots, r_M$, inducing posterior means $\bar{c}_{g1}, \ldots, \bar{c}_{gM}$.

Step 1 (Equal revenue across on-path actions). The platform's revenue from choosing r_k is $\pi(n, r_k;\, \bar{c}_{gk})$, which depends on $(n, r_k, \bar{c}_{gk})$ but not on the realized $\mathbf{c}$. This is because bids $b_i = \bar{c}_{gk} \cdot \theta_i$ are determined by the induced beliefs (through the signal r_k) and the types θ_i, which are i.i.d.

Uniform$[0,1]$ regardless of the selection. Since the platform can choose any on-path r_k for any realization of $\mathbf{c}$, optimality requires all on-path actions to yield the same revenue:

$$\pi(n, r_1; \bar{c}_{g1}) = \cdots = \pi(n, r_M; \bar{c}_{gM}) =: \pi^E.$$

Step 2 (Belief constraint). Let $\bar{c}_g$ denote the unconditional posterior mean of c_i for a selected advertiser (absent any signal from r), and define $w_k := \Pr[r = r_k \mid i \text{ selected}] > 0$ with $\sum_k w_k = 1$. It is $\bar{c}_g = \sum_{k=1}^{M} w_k\, \bar{c}_{gk}$. Since the signal is informative ($M \geq 2$), the posterior means are not all equal, so $\min_k \bar{c}_{gk} < \bar{c}_g$.

Step 3 (Revenue bound). By Lemma 1, $\pi^E = \pi(n, r_k; \bar{c}_{gk}) \leq \bar{c}_{gk} \cdot \alpha(n)$ for every k. Taking the minimum:

$$\pi^E \leq \big(\min_k \bar{c}_{gk}\big) \cdot \alpha(n) < \bar{c}_g \cdot \alpha(n).$$

Step 4 (Pooling dominates). The pooling equilibrium with constant $r = \bar{c}_g/2$ yields revenue $\bar{c}_g \cdot \alpha(n) > \pi^E$. It is supportable by assigning pessimistic off-path beliefs: for any off-path $r' \neq \bar{c}_g/2$, set $\bar{c}_g^{\text{off}}(r') = \varepsilon$ for arbitrarily small $\varepsilon > 0$, giving the platform revenue at most $\varepsilon \cdot \alpha(n) < \bar{c}_g \cdot \alpha(n)$. This contradicts the assumption that the signaling equilibrium is platform-optimal. ■

Proof of Lemma 3

By Lemma 2, the reserve price is constant, so the platform's equilibrium revenue is $\pi^* = \bar{c}_g(\mathcal{S}) \cdot \alpha(n)$, where $\bar{c}_g(\mathcal{S}) = E[c_i \mid i \text{ selected by } \mathcal{S}]$. Since $\alpha(n) > 0$ is fixed, the optimal selection rule maximizes $\bar{c}_g(\mathcal{S})$.

By exchangeability (all c_i are i.i.d. and the selection rule is measurable with respect to $\mathbf{c}$), the ex ante probability that advertiser i is selected is n/K for any symmetric rule. Therefore:

$$\bar{c}_g(\mathcal{S}) = \frac{E[c_i \cdot \mathbf{1}\{i \in \mathcal{S}(\mathbf{c})\}]}{\Pr[i \in \mathcal{S}(\mathbf{c})]} = \frac{K}{n} \cdot E[c_i \cdot \mathbf{1}\{i \in \mathcal{S}(\mathbf{c})\}].$$

Summing over all advertisers:

$$n\, \bar{c}_g(\mathcal{S}) = E\left[\sum_{i \in \mathcal{S}(\mathbf{c})} c_i\right].$$

Since exactly n advertisers are selected for each $\mathbf{c}$, the right-hand side is maximized pointwise by choosing the n largest c_i's. Hence the top-n rule uniquely maximizes $\bar{c}_g(\mathcal{S})$.

Under the top-n rule, a selected advertiser's c_i is equally likely to be any of the top n order statistics of K i.i.d. Uniform$[0,1]$ random variables. The expected value of the k-th order statistic of K standard uniforms is $k/(K+1)$, so:

$$\bar{c}_g(n) = \frac{1}{n}\sum_{k=K-n+1}^{K}\frac{k}{K+1} = \frac{1}{n(K+1)}\cdot\frac{n(2K-n+1)}{2} = \frac{2K-n+1}{2(K+1)}.$$

For incentive compatibility: conditional on fixed beliefs, each bidder bids $\bar{c}_g \cdot \theta_i$, where $\theta_i \sim$ Uniform$[0,1]$ independently of c_i and identically across all advertisers. The joint distribution of submitted bids, and therefore the platform's revenue, is invariant to which n advertisers are selected. The platform is thus indifferent among all size-n selection rules, and in particular has no profitable deviation from the top-n rule. ■

Proof of Proposition 1

Existence. We verify that the described strategy profile and beliefs constitute a PBE. Given the platform's strategy (top-n^* selection, reserve $r^* = \bar{c}_g(n^*)/2$), an invited advertiser rationally infers that its c_i is among the top n^* out of K, yielding posterior mean $\bar{c}_g(n^*) = (2K - n^* + 1)/(2(K+1))$ by Lemma 3. Truth-telling gives the optimal bid $\beta^*(\theta_i) = \bar{c}_g(n^*)\cdot\theta_i$. The resulting bids are i.i.d. Uniform$[0, \bar{c}_g(n^*)]$; Lemma 1 confirms $r^* = \bar{c}_g(n^*)/2$ is optimal. The platform chose n^* to maximize $\Pi(n)$ by construction.

Off-path beliefs support the equilibrium as follows. For any off-path reserve price $r' \neq r^*$, assign $\bar{c}_g^{\text{off}}(r') = \varepsilon$ (arbitrarily small), making such deviations unprofitable.

Optimality. Any PBE must satisfy: the reserve price is at most as good as the optimal $r^* = \bar{c}_g/2$ (Lemma 1); information revelation through r is weakly suboptimal (Lemma 2); the selection rule achieves posterior mean at most $\bar{c}_g(n)$ (Lemma 3). Hence no PBE achieves revenue exceeding $\max_n \Pi(n) = \Pi(n^*)$.

Uniqueness. Suppose a PBE achieves revenue Π^*. By Lemma 1, it must use reserve $r = \bar{c}_g/2$ (unique optimizer for given beliefs), so the revenue is $\bar{c}_g \cdot \alpha(n)$. Setting this equal to $\Pi^* = \bar{c}_g(n^*)\cdot\alpha(n^*)$ and noting that $\bar{c}_g \leq \bar{c}_g(n)$ (by Lemma 3, top-n is the best achievable belief), we require $\bar{c}_g(n) \cdot \alpha(n) \geq \Pi^*$, i.e., $\Pi(n) \geq \Pi(n^*)$, which forces $n = n^*$ when the maximizer is unique. Given $n = n^*$, the revenue Π^* is achieved only if $\bar{c}_g = \bar{c}_g(n^*)$ (since $\bar{c}_g \leq \bar{c}_g(n^*)$ and $\bar{c}_g \cdot \alpha(n^*) = \Pi^*$ requires equality), which by Lemma 3 requires the top-n^* selection rule. The reserve is then $r^* = \bar{c}_g(n^*)/2$. Hence the equilibrium is unique up to off-path beliefs. ■

Proof of Proposition 2

This result is a direct consequence of Proposition 3, which characterizes the optimal n^*. For $K \leq 3$, it is easy to verify that $n^*(K) = K$, since $\{p_k\}_{1 \leq k \leq 3} = \{1, 2, 4\}$. For $K \geq 4$, it is $p_{n^*-1} < K \leq p_{K-1}$. Since the sequence $\{p_k\}_{k \geq 0}$ is strictly increasing, it follows that $n^* < K$, for $K \geq 4$. ■

Proof of Corollary 1

From Proposition 1, we have that the platform's expected revenue is

$$\Pi^*(n^*, K) = \bar{c}_g(n) \cdot \alpha(n^*) = \frac{(2K - n^* + 1)\left(n^* - 1 + \frac{1}{2^{n^*}}\right)}{2(n^* + 1)(K + 1)}.$$

Next, the Social Welfare $SW^*(n^*, K)$ is the expected value of the winning bid when it exceeds the reserve price. For the bid distribution $F_B(x) = x/\bar{c}_g$, we have that

$$SW^*(n^*, K) = \int_{\bar{c}_g/2}^{\bar{c}_g} x \cdot n^* F_B(x)^{n^*-1} f_B(x) dx = \bar{c}_g \int_{1/2}^{1} n^* y^{n^*} dy = \bar{c}_g \frac{n^*}{n^* + 1}\left(1 - \frac{1}{2^{n^*+1}}\right).$$

Substituting $\bar{c}_g$ gives the formula for $SW^*(n^*, K)$.

Finally, the total advertiser surplus $A^*(n^*, K)$ is $SW^*(n^*, K) - \Pi^*(n^*, K)$. Subtracting the two expressions gives us the formula for $A^*(n^*, K)$. ■

Proof of Proposition 3

For $K \leq 12$, we can easily verify numerically that the given formula for p_k gives the correct values for n^* (i.e., $\{n^*(K)\}_{1 \leq K \leq 12} = \{1, 2, 3, 3, 4, 4, 4, 5, 5, 5, 6, 6\}$). Therefore, let's assume that $k \geq 6$ to determine the optimal values for $K \geq 13$. For $k \geq 6$, the formula for p_k can be simplified as follows:

$$p_k = \begin{cases} \frac{1}{4}(k^2+3k-4) & \text{, if } k \geq 8 \text{ and } (k \bmod 4) \in \{0,1\}, \\ \frac{1}{4}(k^2+3k-2) & \text{, if } k \geq 6 \text{ and } (k \bmod 4) \in \{2,3\}. \end{cases}$$

Since $\{n^*(K)\}_{K \geq 1}$ is a non-decreasing sequence, it is enough to prove that for every $k \geq 6$, if $K = p_k$ then the optimal n^* is equal to k, while if $K = p_k + 1$ then the optimal n^* is equal to $k+1$. In other words, we want to prove the following two cases: (i) $n^*(p_k) = k$, (ii) $n^*(p_k+1) = k+1$.

Let's start with case (i). Because the function $\Pi(K,n) = \frac{(2K-n+1)\cdot\left(n-1+\frac{1}{2^n}\right)}{2(n+1)(K+1)}$ is concave in n, it is sufficient to prove the following two sub-cases: (i.a) $\Pi(p_k,k) > \Pi(p_k,k+1)$, (i.b) $\Pi(p_k,k) > \Pi(p_k,k-1)$.

Let's start with sub-case (i.a). We further consider two sub-sub-cases: (i.a.1) $(k \bmod 4) \in \{0,1\}$, (i.a.2) $(k \bmod 4) \in \{2,3\}$.

Let's start with sub-sub-case (i.a.1). The inequality is equivalent to $\Pi\left(\frac{1}{4}(k^2+3k-4),k\right) > \Pi\left(\frac{1}{4}(k^2+3k-4),k+1\right)$. It is

$$\Pi\left(\frac{1}{4}(k^2+3k-4),k\right) - \Pi\left(\frac{1}{4}(k^2+3k-4),k+1\right) = \frac{\left(k^3+4k^2+3k+2^{k+3}-4\right)}{2^{k+1}k(k+1)(k+2)(k+3)} > 0,$$

for any $k \geq 1$, therefore sub-sub-case (i.a.1) holds.

Let's continue with sub-sub-case (i.a.2). The inequality is equivalent to $\Pi\left(\frac{1}{4}(k^2+3k-2),k\right) > \Pi\left(\frac{1}{4}(k^2+3k-2),k+1\right)$. It is $\Pi\left(\frac{1}{4}(k^2+3k-2),k\right) - \Pi\left(\frac{1}{4}(k^2+3k-2),k+1\right) = \frac{1}{2^{k+1}(k+2)} > 0$, for any $k \geq 1$, therefore sub-sub-case (i.a.2) holds.

Now let's continue with sub-case (i.b). Similarly, we consider two sub-sub-cases: (i.b.1) $(k \bmod 4) \in \{0,1\}$, (i.b.2) $(k \bmod 4) \in \{2,3\}$.

For sub-sub-case (i.b.1), we have that

$$\Pi\left(\frac{1}{4}(k^2+3k-4),k\right) - \Pi\left(\frac{1}{4}(k^2+3k-4),k-1\right) = \frac{\left(-k^2-3k+2^{k+2}-4\right)}{2^k k(k+1)(k+3)} > 0,$$

for any $k \geq 2$.

For sub-sub-case (i.b.2), we have that

$$\Pi\left(\frac{1}{4}(k^2+3k-2),k\right)-\Pi\left(\frac{1}{4}(k^2+3k-2),k-1\right)=\frac{\left(-k^2-2k+2^{k+2}-4\right)}{2^k k(k+1)(k+2)}>0,$$

for any $k \geq 1$.

This concludes case (i). Let's continue with case (ii). As before, we consider two sub-cases: (ii.a) $\Pi(p_k+1,k+1) > \Pi(p_k+1,k+2)$, (ii.b) $\Pi(p_k+1,k+1) > \Pi(p_k+1,k)$. And for each sub-case, we consider two sub-sub-cases based on the remainder of k when divided by 4.

For sub-sub-case (ii.a.1), we have

$$\Pi\left(\frac{1}{4}(k^2+3k-4)+1,k+1\right)-\Pi\left(\frac{1}{4}(k^2+3k-4)+1,k+2\right)=\frac{\left(k^3+5k^2+\left(2^{k+4}+6\right)k+2^{k+4}+4\right)}{2^{k+2}(k+2)(k+3)\left(k^2+3k+4\right)}>0,$$

for any $k \geq 1$.

For sub-sub-case (ii.a.2), we have

$$\Pi\left(\frac{1}{4}(k^2+3k-2)+1,k+1\right)-\Pi\left(\frac{1}{4}(k^2+3k-2)+1,k+2\right)=\frac{\left(k^3+5k^2+\left(2^{k+4}+8\right)k+12\right)}{2^{k+2}(k+2)(k+3)\left(k^2+3k+6\right)}>0,$$

for any $k \geq 1$.

For sub-sub-case (ii.b.1), we have $\Pi\left(\frac{1}{4}(k^2+3k-4)+1,k+1\right)-\Pi\left(\frac{1}{4}(k^2+3k-4)+1,k\right)=\frac{\left(-k^3-4k^2-7k+2^{k+3}-8\right)}{2^{k+1}(k+1)(k+2)(k^2+3k+4)}>0$, for any $k \geq 6$.

For sub-sub-case (ii.b.2), we have $\Pi\left(\frac{1}{4}(k^2+3k-2)+1,k+1\right)-\Pi\left(\frac{1}{4}(k^2+3k-2)+1,k\right)=\frac{\left(-k^3-4k^2-9k+2^{k+4}-14\right)}{2^{k+1}(k+1)(k+2)(k^2+3k+6)}>0$, for any $k \geq 1$.

This concludes case (ii) and therefore the proof. ∎

Proof of Corollary 2

For this corollary, we need the following lemma that characterizes the socially optimal number of bidders n^{*s}.

Lemma 10 (Socially optimal n^{*s}). *Let $\{p_k^s\}_{k\geq 0}$ be the sequence that starts with $\{0,1,3,6,10,16,23,31,41,51,63,75\}$ and for $k \geq 12$, it is defined as $p_k^s=\frac{k^2+3k-2}{2}$. The sequence $\{p_k^s\}_{k\geq 0}$ is strictly increasing. Then, for any integer $K \geq 1$, the socially optimal $n^{*s}(K)$ is the unique integer k such that $p_{k-1}^s < K \leq p_k^s$.*[19]

[19]The first few terms of the sequence $\{n^{*s}(K)\}_{K\geq 1}$ are as follows: $1,2,2,3,3,3,4,4,4,4,5,5,5,5,5,5,6,6,6,6,6,\ldots$.

Proof. For $K \leq 88$, we can easily verify numerically that the given values for p_k^s give the correct values for n^{*s} (i.e., $\{n^{*s}(K)\}_{1\leq K\leq 88} = \{1,2,2,3,3,3,4,4,4,4,5,\dots\}$). Therefore, let's assume that $k \geq 12$ to determine the optimal values for $K \geq 89$.

Since $\{n^{*s}(K)\}_{K\geq 1}$ is a non-decreasing sequence, it is enough to prove that for every $k \geq 12$, if $K = p_k^s$ then the optimal n^{*s} is equal to k, while if $K = p_k^s + 1$ then the optimal n^{*s} is equal to $k+1$. In other words, we want to prove the following two cases: (i) $n^{*s}(p_k^s) = k$, (ii) $n^*(p_k^s + 1) = k+1$.

Let's start with case (i). Because the function $SW(K,n) = \frac{n(2K-n+1)\left(1-\frac{1}{2^{n+1}}\right)}{2(n+1)(K+1)}$ is concave in n, it is sufficient to prove the following two sub-cases: (i.a) $SW(p_k^s, k) > SW(p_k^s, k+1)$, (i.b) $SW(p_k^s, k) > SW(p_k^s, k-1)$.

Let's start with sub-case (i.a). It is

$$SW\left(\frac{k^2+3k-2}{2}, k\right) - SW\left(\frac{k^2+3k-2}{2}, k+1\right) = \frac{-k^4-4k^3-3k^2+2k+2^{k+3}-2}{2^{k+2}k(k+1)(k+2)(k+3)} > 0,$$

for any $k \geq 12$, therefore sub-case (i.a) holds.

Let's continue with sub-case (i.b). It is

$$SW\left(\frac{k^2+3k-2}{2}, k\right) - SW\left(\frac{k^2+3k-2}{2}, k-1\right) = \frac{k^3+2k^2-k+2^{k+2}-4}{2^{k+1}k(k+1)(k+3)} > 0,$$

for any $k \geq 1$, therefore sub-case (i.b) holds.

This concludes case (i). Let's continue with case (ii). As before, we consider two sub-cases: (ii.a) $SW(p_k^s+1, k+1) > SW(p_k^s+1, k+2)$, (ii.b) $SW(p_k^s+1, k+1) > SW(p_k^s+1, k)$.

For sub-case (ii.a), we have

$$SW\left(\frac{k^2+3k-2}{2}+1, k+1\right) - SW\left(\frac{k^2+3k-2}{2}+1, k+2\right) = \frac{-k^3-4k^2-3k+2^{k+4}-2}{2^{k+3}(k+1)(k+2)(k+3)} > 0,$$

for any $k \geq 1$, therefore sub-case (ii.a) holds.

For sub-case (ii.b), we have

$$SW\left(\frac{k^2+3k-2}{2}+1, k+1\right) - SW\left(\frac{k^2+3k-2}{2}+1, k\right) = \frac{k}{2^{k+2}(k+2)} > 0,$$

for any $k \geq 1$, therefore sub-case (ii.b) holds.

This concludes case (ii) and therefore the proof. $\square$

The corollary is now a consequence of the fact that the sequence p_k^s grows faster than the sequence p_k, and more specifically it holds that $p_k^s \geq p_k$ for every $k \geq 1$. This means that $n^*(K) \geq n^{*s}(K)$ for every $K \geq 1$. This completes the proof of the corollary. $\blacksquare$

Corollary 5. *The social welfare is non-monotonic with respect to K. More specifically, it is $SW_{P_C A_\Theta}(K) > SW_{P_C A_\Theta}(K+1)$ for $K \in \{p_k\}_{k \geq 7} = \{17, 21, 26, 32, 38, 44, 51, 59, \dots\}$, and $SW_{P_C A_\Theta}(K) < SW_{P_C A_\Theta}(K+1)$ otherwise. In other words, there are cases where adding an advertiser reduces the social welfare.*[20]

Proof. We want to prove that (i) for $K \in \{p_k\}_{k \geq 7}$, $SW_{P_C A_\Theta}(K) > SW_{P_C A_\Theta}(K+1)$, and (ii) for $K \notin \{p_k\}_{k \geq 7}$, $SW_{P_C A_\Theta}(K) < SW_{P_C A_\Theta}(K+1)$.

Let's start with case (i). Let $SW(K, n^*) = \frac{n^*(2K - n^* + 1)\left(1 - \frac{1}{2^{n^*+1}}\right)}{2(n^*+1)(K+1)}$ and suppose that $K = p_k$ for some $k \geq 7$. From Proposition 3, we know that $n^*(K) = k$ and $n^*(K+1) = k+1$. Therefore, we need to prove that $SW(p_k, k) > SW(p_k + 1, k + 1)$. We consider two sub-cases: (i.1) (k mod 4) $\in \{0, 1\}$, (i.2) (k mod 4) $\in \{2, 3\}$.

For sub-case (i.1), we have that

$$SW(p_k, k) - SW(p_k + 1, k + 1) = \frac{\left(-k^5 - 6k^4 + 2\left(2^{k+1} - 7\right)k^3 - 2\left(2^{k+2} + 3\right)k^2 - 3\left(5\ 2^{k+2} - 9\right)k - 2^{k+6} + 32\right)}{2^{k+2}(k+1)(k+2)(k+3)\left(k^2 + 3k + 4\right)} > 0,$$

for any $k \geq 7$.

For sub-case (i.2), we have that

$$SW(p_k, k) - SW(p_k + 1, k + 1) = \frac{\left(-k^4 - 3k^3 + \left(2^{k+2} - 7\right)k^2 - 5\left(2^{k+2} - 1\right)k - 2^{k+3} + 2\right)}{2^{k+2}(k+1)(k+2)\left(k^2 + 3k + 6\right)} > 0,$$

for any $k \geq 7$.

Now let's continue with case (ii). For $K < 17$, we can easily verify numerically that the inequality holds. Therefore, let's consider the case $K > 17$. Since $K \notin \{p_k\}_{k \geq 7}$, from Proposition 3 we know that $n^*(K) = n^*(K+1)$. Therefore, it is sufficient to prove that $SW(K, n) < SW(K+1, n)$, for

[20] Intuitively, this happens because being in the top n out of K advertisers is better than being in the top $n+1$ out of $K+1$ advertisers, for sufficiently large K. The first time this happens is during the transition from $K = 17$ to $K = 18$ where $n^*(K)$ goes from 7 to 8, and it happens again every time there is a transition after that, i.e. for $K \in \{p_k\}_{k \geq 7}$.

any $n \geq 1$. It is $SW(K,n) - SW(K+1,n) = -\frac{(2^{n+1}-1)n}{2^{n+2}(K^2+3K+2)} < 0$, for any $n \geq 1$. This concludes the proof. $\square$

Proof of Lemma 4

In model $P_{\text{ALL}}A_\Theta$, the platform observes both the conversion probabilities $\mathbf{c}$ and the values per conversion $\boldsymbol{\theta}$.

For Part 1, suppose the platform attempts to use a selection strategy that depends on $\mathbf{c}$, thereby inducing a posterior belief $\bar{c}$ for the invited advertisers. The advertisers' bidding strategy would then be $\beta_i(\theta_i) = \bar{c} \cdot \theta_i$. However, since the platform knows the exact realization of $\boldsymbol{\theta}$, the bids are deterministic from the platform's perspective. To maximize its revenue, the platform's optimal deviation is to always select the advertisers with the highest θ_i values, regardless of their actual c_i values. Because the platform cannot commit to ignoring $\boldsymbol{\theta}$, any selection strategy that depends on $\mathbf{c}$ is not credible. In equilibrium, the advertisers rationally anticipate this, so their posterior belief about c_i remains equal to their prior expectation: $\bar{c} = \mathbb{E}[c_i] = 1/2$. Consequently, each advertiser's optimal bid is $\beta^*_{P_{\text{ALL}}A_\Theta}(\theta_i) = \theta_i/2$.

For Part 2, because the platform knows all bids exactly, it can maximize its revenue by setting a reserve price exactly equal to the highest bid among all advertisers. Thus, the optimal reserve price is $r^*_{P_{\text{ALL}}A_\Theta} = \max_i \theta_i/2$.

For Part 3, because the reserve price is exactly equal to the winning bidder's expected valuation, the platform extracts the full surplus of the winning advertiser. This leaves the winning advertiser, and all losing advertisers, with a payoff of zero. Thus, the total expected advertiser surplus is $A^*_{P_{\text{ALL}}A_\Theta} = 0$.

The platform's expected profit is the expected value of the reserve price: $\Pi^*_{P_{\text{ALL}}A_\Theta} = \mathbb{E}[\max_i \theta_i/2]$. Since $\theta_i \sim U[0,1]$ are drawn independently, the expected value of the highest order statistic among K draws is $\mathbb{E}[\max_i \theta_i] = \frac{K}{K+1}$. Therefore, the platform's expected profit is $\Pi^*_{P_{\text{ALL}}A_\Theta} = \frac{K}{2(K+1)}$. This completes the proof. $\blacksquare$

Proof of Proposition 4

For the advertiser surplus, since $A^*_{P_{\text{ALL}}A_\Theta} = 0$, it follows that $A^*_{P_C A_\Theta} \geq A^*_{P_{\text{ALL}}A_\Theta}$ for all $K \geq 1$.

Next, we compare the platform's expected profit. We want to show that $\Pi^*_{P_C A_\Theta} \geq \Pi^*_{P_{\text{ALL}} A_\Theta}$ for $K \geq 8$. From Corollary 1, the platform's profit in the main model is:

$$\Pi^*_{P_C A_\Theta} = \Pi^*(n^*, K) = \frac{(2K - n^* + 1)\left(n^* - 1 + 2^{-n^*}\right)}{2(n^* + 1)(K + 1)}.$$

From Lemma 4, the profit in the benchmark model is $\Pi^*_{P_{\text{ALL}} A_\Theta} = \frac{K}{2(K+1)}$.

Because n^* is chosen to maximize $\Pi^*(n, K)$ in the main model, we have $\Pi^*(n^*, K) \geq \Pi^*(n, K)$ for any integer $n \in [1, K]$. Thus, to prove $\Pi^*_{P_C A_\Theta} \geq \Pi^*_{P_{\text{ALL}} A_\Theta}$, it is sufficient to show that there exists some specific n such that $\Pi^*(n, K) \geq \Pi^*_{P_{\text{ALL}} A_\Theta}$ for $K \geq 8$.

Let us evaluate the platform's profit in the main model if it were to choose $n = 5$. Substituting $n = 5$ into the profit function, we get:

$$\Pi^*(5, K) = \frac{(2K - 5 + 1)\left(5 - 1 + 2^{-5}\right)}{2(5 + 1)(K + 1)} = \frac{(2K - 4)\left(4 + \frac{1}{32}\right)}{12(K + 1)} = \frac{2(K - 2)\left(\frac{129}{32}\right)}{12(K + 1)} = \frac{129(K - 2)}{192(K + 1)}.$$

We compare this to the benchmark profit, which can be rewritten with a common denominator as $\Pi^*_{P_{\text{ALL}} A_\Theta} = \frac{96K}{192(K+1)}$. The inequality $\Pi^*(5, K) \geq \Pi^*_{P_{\text{ALL}} A_\Theta}$ is equivalent to:

$$\frac{129(K - 2)}{192(K + 1)} \geq \frac{96K}{192(K + 1)} \Leftrightarrow K \geq \frac{86}{11} \approx 7.818.$$

Because K is an integer, this inequality holds strictly for all $K \geq 8$. Since the optimal profit $\Pi^*_{P_C A_\Theta} = \Pi^*(n^*, K)$ must be at least as large as $\Pi^*(5, K)$, it follows that $\Pi^*_{P_C A_\Theta} > \Pi^*_{P_{\text{ALL}} A_\Theta}$ for all $K \geq 8$. This completes the proof. ■

Proof of Lemma 5

Algorithm 1 Find the optimal reserve price in model $P_C A_{\text{ALL}}$.

1: Order the advertisers by their maximum valuation c_i in descending order: without loss of generality, let's assume that $c_1 \geq c_2 \ldots \geq c_K$.

2: For every integer $k \in \{1, 2, \ldots, K\}$ set the reserve price to $r = \frac{\sum_{i=1}^{k} c_i}{2 \cdot k}$ and calculate the expected revenue $R(k)$.

3: Find the integer k^* that maximizes $R(k)$. The optimal reserve price is $r^*_{P_C A_{\text{ALL}}} = \frac{\sum_{i=1}^{k^*} c_i}{2 \cdot k^*}$.

Given $c_1, c_2, \ldots, c_K \in [0, 1]$, let a_1 be the highest value among them, a_2 the second highest value, and so on, such that $a_1 \geq a_2 \geq \cdots \geq a_K$. We also define $a_{K+1} = 0$.

For a given reserve price $r \in [0, a_1)$, let n be the unique integer such that $a_{n+1} \leq r < a_n$. If the platform sets the reserve price r, then its expected revenue is denoted by $R_n(a_1, \ldots, a_K, r)$.

$$R_n(a_1, \ldots, a_K, r) := r\left(1 - \frac{r^n}{\prod_{i=1}^n a_i}\right) + \sum_{k=1}^{n-1}\left[t - \frac{1-k}{k+1}\frac{t^{k+1}}{\prod_{j=1}^k a_j} - \frac{t^k}{k}\frac{\sum_{j=1}^k a_j}{\prod_{j=1}^k a_j}\right]_{t=a_{k+1}}^{t=a_k}$$
$$+ \left[t - \frac{1-n}{n+1}\frac{t^{n+1}}{\prod_{j=1}^n a_j} - \frac{t^n}{n}\frac{\sum_{j=1}^n a_j}{\prod_{j=1}^n a_j}\right]_{t=r}^{t=a_n}.$$

It can be shown (by taking the derivative of R_n with respect to r) that this function has a local maximum at $r = \frac{\sum_{j=1}^n a_j}{2n}$. Therefore, if this value falls within the valid range for n, i.e., if $a_{n+1} \leq \frac{\sum_{j=1}^n a_j}{2n} < a_n$, then this value is a candidate for the optimal reserve price.

Given $a_1, \ldots, a_K$, the platform's overall expected revenue for any reserve price r is a piecewise function:

$$R(a_1, \ldots, a_K, r) := \begin{cases} 0 & \text{, if } r \geq a_1, \\ R_1(a_1, \ldots, a_K, r) & \text{, if } a_2 \leq r < a_1, \\ R_2(a_1, \ldots, a_K, r) & \text{, if } a_3 \leq r < a_2, \\ \vdots & \vdots \\ R_K(a_1, \ldots, a_K, r) & \text{, if } a_{K+1} \leq r < a_K. \end{cases}$$

Let $r^*(a_1, \ldots, a_K)$ be the optimal reserve price, i.e., the value of r that maximizes $R(a_1, \ldots, a_K, r)$. To find $r^*(a_1, \ldots, a_K)$ algorithmically, we can check all the possible local maxima, $r_n^* = \frac{\sum_{j=1}^n a_j}{2n}$ for $n \in \{1, 2, \ldots, K\}$. For each n, we check if r_n^* is a valid candidate (i.e., if $a_{n+1} \leq r_n^* < a_n$). We then evaluate the revenue $R(a_1, \ldots, a_K, r_n^*)$ for all valid candidates and select the one that maximizes this revenue.

We now want to calculate the expected value of the optimal revenue,

$$R^*(a_1, \ldots, a_K) := R(a_1, \ldots, a_K, r^*(a_1, \ldots, a_K)),$$

assuming that the initial values $c_1, \ldots, c_K$ are independent uniform random draws from $[0, 1]$. In this case, the sorted values $a_1, \ldots, a_K$ are the order statistics of K i.i.d. Uniform(0,1) random variables.

Their joint probability density function over the region $1 > a_1 > a_2 > \cdots > a_K > 0$ is $K!$. The expected optimal revenue is therefore given by the integral:

$$\Pi_{P_C A_{\text{ALL}}} = K! \int_0^1 \int_0^{a_1} \int_0^{a_2} \cdots \int_0^{a_{K-1}} R^*(a_1, \ldots, a_K)\, da_K\, da_{K-1} \ldots da_1.$$

Similarly, for a reserve price r such that $a_{n+1} \leq r < a_n$, the expected social welfare is equal to

$$SW_n(a_1, \ldots, a_K, r) := \left(a_1 - \sum_{k=1}^{n} \frac{a_k^{k+1} - a_{k+1}^{k+1}}{(k+1)\prod_{j=1}^{k} a_j} \right) - \frac{n r^{n+1}}{(n+1)\prod_{i=1}^{n} a_i}.$$

The overall expected social welfare for any reserve price r is:

$$SW(a_1, \ldots, a_K, r) := \begin{cases} 0 & \text{, if } r \geq a_1, \\ SW_1(a_1, \ldots, a_K, r) & \text{, if } a_2 \leq r < a_1, \\ SW_2(a_1, \ldots, a_K, r) & \text{, if } a_3 \leq r < a_2, \\ \vdots & \vdots \\ SW_K(a_1, \ldots, a_K, r) & \text{, if } a_{K+1} \leq r < a_K. \end{cases}$$

For the optimal reserve price r (optimal in terms of revenue), the expected social welfare is

$$SW^*(a_1, \ldots, a_K) := SW(a_1, \ldots, a_K, r^*(a_1, \ldots, a_K)).$$

Finally, assuming that the initial values $c_1, \ldots, c_K$ are independent uniform random draws from $[0, 1]$, the expected social welfare is

$$SW_{P_C A_{\text{ALL}}} = K! \int_0^1 \int_0^{a_1} \int_0^{a_2} \cdots \int_0^{a_{K-1}} SW^*(a_1, \ldots, a_K)\, da_K\, da_{K-1} \ldots da_1.$$

The expected advertiser's welfare can then be calculated as $A_{P_C A_{\text{ALL}}} = SW_{P_C A_{\text{ALL}}} - \Pi_{P_C A_{\text{ALL}}}$. ■

Proof of Proposition 5

The proof proceeds by combining the exact evaluation of the expected payoffs for small K with an asymptotic analysis for large K.

Small K**.** For small values of K, we apply the algorithm described in the proof of Lemma 5. By evaluating the piecewise integrals over the order statistics $a_1, \ldots, a_K$, we obtain the exact expected platform revenue $\Pi^*_{P_C A_{\text{ALL}}}$ and expected social welfare $SW^*_{P_C A_{\text{ALL}}}$, which also gives the advertiser surplus $A^*_{P_C A_{\text{ALL}}} = SW^*_{P_C A_{\text{ALL}}} - \Pi^*_{P_C A_{\text{ALL}}}$. Comparing these exact numerical values with the closed-form expressions for $\Pi^*_{P_C A_\Theta}$ and $A^*_{P_C A_\Theta}$ from Corollary 1, we verify that $\Pi^*_{P_C A_\Theta} \geq \Pi^*_{P_C A_{\text{ALL}}}$ for $K < 8$, and $\Pi^*_{P_C A_\Theta} < \Pi^*_{P_C A_{\text{ALL}}}$ for $K \geq 8$. Furthermore, $A^*_{P_C A_\Theta} \leq A^*_{P_C A_{\text{ALL}}}$ holds for all evaluated K. We provide all the numerical results in Table 1 and Table 3 of Appendix A.2.

Large K**.** To prove that this relationship holds for all sufficiently large K, we analyze the asymptotic behavior of the payoffs as $K \to \infty$.

In the main model $P_C A_\Theta$, we substitute the optimal number of bidders $n^* \sim 2\sqrt{K}$ (from Proposition 3) into the formulas from Corollary 1. For the platform's revenue, we have:

$$\begin{aligned}
\Pi^*_{P_C A_\Theta} &= \frac{(2K - n^* + 1)(n^* - 1 + 2^{-n^*})}{2(n^* + 1)(K + 1)} \\
&\approx \left(1 - \frac{n^*}{2K}\right)\left(1 - \frac{2}{n^*}\right) \approx \left(1 - \frac{1}{\sqrt{K}}\right)\left(1 - \frac{1}{\sqrt{K}}\right) = 1 - \frac{2}{\sqrt{K}} + o\left(\frac{1}{\sqrt{K}}\right).
\end{aligned}$$

Similarly, the expected social welfare expands to $SW^*_{P_C A_\Theta} \approx 1 - \frac{1.5}{\sqrt{K}}$, giving an asymptotic advertiser surplus of $A^*_{P_C A_\Theta} = SW^*_{P_C A_\Theta} - \Pi^*_{P_C A_\Theta} \approx \frac{0.5}{\sqrt{K}}$.

In model $P_C A_{\text{ALL}}$, each advertiser bids its true valuation $V_i = \theta_i c_i$. Since θ_i and c_i are independent standard uniform variables, the CDF of V_i is $F(v) = v - v \ln v$ for $v \in (0, 1]$. To understand the upper tail, let $v = 1 - x$ for small $x > 0$. Using the Taylor expansion $\ln(1 - x) = -x - x^2/2 + O(x^3)$, the tail probability is:

$$1 - F(1 - x) = 1 - (1 - x) + (1 - x)\ln(1 - x) \approx x + (1 - x)\left(-x - \frac{x^2}{2}\right) \approx \frac{x^2}{2}.$$

Let $X_i = 1 - V_i$. The approximation $\Pr(X_i < x) \approx x^2/2$ implies that X_i is distributed approximately as $\sqrt{2Y_i}$, where $Y_i \sim U[0, 1]$. The platform's revenue $\Pi^*_{P_C A_{\text{ALL}}}$ is bounded below by the expected second-highest valuation (which corresponds to setting a reserve price $r = 0$). The second-highest valuation $V_{(K-1)}$ corresponds to the second-lowest value of X_i, denoted $X_{(2)} \approx \sqrt{2Y_{(2)}}$.

Because $Y_{(2)}$ is the second-order statistic of K standard uniform variables, it follows a Beta distribution $Y_{(2)} \sim \text{Beta}(2, K-1)$. We can compute its expected square root using the Gamma function:

$$\mathbb{E}\left[X_{(2)}\right] \approx \sqrt{2}\frac{\Gamma(2.5)\Gamma(K+1)}{\Gamma(2)\Gamma(K+1.5)} \approx \sqrt{2}\left(1.5 \times 0.5 \times \sqrt{\pi}\right) K^{-0.5} = \frac{3}{4}\sqrt{2\pi}\frac{1}{\sqrt{K}} \approx \frac{1.88}{\sqrt{K}}.$$

Therefore, the platform's expected revenue in $P_C A_{\text{ALL}}$ is bounded below by:

$$\Pi^*_{P_C A_{\text{ALL}}} \geq \mathbb{E}\left[V_{(K-1)}\right] \approx 1 - \frac{1.88}{\sqrt{K}}.$$

Comparing the two models, since $1 - \frac{1.88}{\sqrt{K}} > 1 - \frac{2}{\sqrt{K}}$, it follows that $\Pi^*_{P_C A_{\text{ALL}}} > \Pi^*_{P_C A_\Theta}$ for sufficiently large K.

By a similar extreme-value argument, the expected highest valuation $V_{(K)}$ corresponds to $X_{(1)} \approx \sqrt{2Y_{(1)}}$, where $Y_{(1)} \sim \text{Beta}(1, K)$. Its expectation is:

$$\mathbb{E}\left[X_{(1)}\right] \approx \sqrt{2}\frac{\Gamma(1.5)\Gamma(K+1)}{\Gamma(1)\Gamma(K+1.5)} \approx \sqrt{2}\left(0.5 \times \sqrt{\pi}\right) K^{-0.5} = \frac{1}{2}\sqrt{2\pi}\frac{1}{\sqrt{K}} \approx \frac{1.25}{\sqrt{K}}.$$

The expected social welfare is $SW^*_{P_C A_{\text{ALL}}} \approx 1 - \frac{1.25}{\sqrt{K}}$. The advertiser surplus is then:

$$A^*_{P_C A_{\text{ALL}}} = SW^*_{P_C A_{\text{ALL}}} - \Pi^*_{P_C A_{\text{ALL}}} \approx \left(1 - \frac{1.25}{\sqrt{K}}\right) - \left(1 - \frac{1.88}{\sqrt{K}}\right) = \frac{0.63}{\sqrt{K}}.$$

Since $\frac{0.63}{\sqrt{K}} > \frac{0.5}{\sqrt{K}}$, the advertiser surplus is strictly greater in $P_C A_{\text{ALL}}$ than in $P_C A_\Theta$ asymptotically. This concludes the proof. ■

Proof of Lemma 6

In model $P_{\text{ALL}} A_{\text{ALL}}$, both the platform and the advertisers observe the full realization of $\mathbf{c}$ and $\boldsymbol{\theta}$. Consequently, the exact valuation of each advertiser i, $v_i = \theta_i c_i$, is common knowledge.

Since the advertisers know their own valuations exactly, it is a weakly dominant strategy for each advertiser to bid its true valuation: $\beta(\theta_i, c_i) = \theta_i c_i$.

Because the platform knows all the bids in advance, it maximizes its revenue by setting the reserve price equal to the highest valuation among all advertisers. Thus, $r^* = \max_i(\theta_i c_i)$. By setting this reserve price, the platform extracts the entire surplus from the winning advertiser, leaving all advertisers with a payoff of zero. Hence, $A^*_{P_{\text{ALL}}A_{\text{ALL}}} = 0$.

The platform's expected payoff is the expected value of the highest valuation: $\Pi^*_{P_{\text{ALL}}A_{\text{ALL}}} = \mathbb{E}[\max_i V_i]$, where $V_i = \theta_i c_i$. Since θ_i and c_i are independent standard uniform random variables, the cumulative distribution function (CDF) of V_i is $F_V(v) = v - v \ln v$ for $v \in (0, 1]$. The CDF of the maximum order statistic $V_{(K)}$ is $(F_V(v))^K$. The expected value can be computed using the integral of the survival function:

$$\Pi^*_{P_{\text{ALL}}A_{\text{ALL}}} = \int_0^1 \left(1 - (v - v \ln v)^K\right) \, dv = 1 - \int_0^1 v^K (1 - \ln v)^K \, dv.$$

To evaluate the integral, we use the substitution $x = -\ln v$, which gives $v = e^{-x}$ and $dv = -e^{-x} dx$. The limits of integration change from $v \in (0, 1]$ to $x \in (\infty, 0]$. The integral becomes:

$$\int_0^\infty e^{-Kx} (1 + x)^K e^{-x} \, dx = \int_0^\infty e^{-(K+1)x} (1 + x)^K \, dx.$$

Using the binomial theorem, we expand $(1 + x)^K = \sum_{k=0}^K \binom{K}{k} x^k$. Substituting this into the integral, we get:

$$\sum_{k=0}^K \binom{K}{k} \int_0^\infty x^k e^{-(K+1)x} \, dx.$$

Recall that the Gamma function gives $\int_0^\infty x^k e^{-ax} \, dx = \frac{k!}{a^{k+1}}$. Applying this formula with $a = K + 1$, we get:

$$\sum_{k=0}^K \frac{K!}{k!(K - k)!} \frac{k!}{(K + 1)^{k+1}} = \sum_{k=0}^K \frac{K!}{(K - k)!(K + 1)^{k+1}}.$$

Substituting this back into the expectation gives the platform's equilibrium payoff:

$$\Pi^*_{P_{\text{ALL}}A_{\text{ALL}}} = 1 - \sum_{k=0}^K \frac{K!}{(K - k)!(K + 1)^{k+1}}.$$

This completes the proof. ■

Proof of Proposition 6

First, from Lemma 4 and Lemma 6, we get that $A^*_{P_{\text{ALL}}A_{\Theta}} = A^*_{P_{\text{ALL}}A_{\text{ALL}}} = 0$.

Next, we prove $\Pi^*_{P_C A_{\text{ALL}}} \leq \Pi^*_{P_{\text{ALL}}A_{\text{ALL}}}$. In model $P_C A_{\text{ALL}}$, the platform only knows $\mathbf{c}$ and sets a reserve price based on this partial information. The winning advertiser pays either the second-highest bid or the reserve price, which leaves a strictly positive information rent for the advertisers ($A^*_{P_C A_{\text{ALL}}} > 0$). The platform's revenue $\Pi^*_{P_C A_{\text{ALL}}}$ is strictly less than the total expected social welfare in $P_C A_{\text{ALL}}$, denoted $SW^*_{P_C A_{\text{ALL}}}$.

Furthermore, the maximum possible social welfare in any allocation mechanism is achieved when the impression is always allocated to the advertiser with the highest actual valuation $v_i = \theta_i c_i$. Therefore, $SW^*_{P_C A_{\text{ALL}}} \leq \mathbb{E}[\max_i(\theta_i c_i)]$.

In model $P_{\text{ALL}} A_{\text{ALL}}$, the platform extracts exactly this maximum possible social welfare as revenue, meaning $\Pi^*_{P_{\text{ALL}} A_{\text{ALL}}} = \mathbb{E}[\max_i(\theta_i c_i)]$.

Combining these inequalities, we get $\Pi^*_{P_C A_{\text{ALL}}} < SW^*_{P_C A_{\text{ALL}}} \leq \Pi^*_{P_{\text{ALL}} A_{\text{ALL}}}$, which implies $\Pi^*_{P_C A_{\text{ALL}}} \leq \Pi^*_{P_{\text{ALL}} A_{\text{ALL}}}$. This completes the proof. ■

Proof of Corollary 3

We analyze the normal-form game where the platform chooses its information set $S_P \in \{C, \text{ALL}\}$ and the advertisers choose their information set $S_A \in \{\Theta, \text{ALL}\}$. The payoffs (u_P, u_A) for each strategy profile correspond to the expected platform revenue and advertiser surplus in the respective benchmark models:

- (C, Θ) results in model $P_C A_\Theta$ with payoffs $(\Pi^*_{P_C A_\Theta}, A^*_{P_C A_\Theta})$.
- (C, ALL) results in model $P_C A_{\text{ALL}}$ with payoffs $(\Pi^*_{P_C A_{\text{ALL}}}, A^*_{P_C A_{\text{ALL}}})$.
- (ALL, Θ) results in model $P_{\text{ALL}} A_\Theta$ with payoffs $(\Pi^*_{P_{\text{ALL}} A_\Theta}, A^*_{P_{\text{ALL}} A_\Theta} = 0)$.
- (ALL, ALL) results in model $P_{\text{ALL}} A_{\text{ALL}}$ with payoffs $(\Pi^*_{P_{\text{ALL}} A_{\text{ALL}}}, A^*_{P_{\text{ALL}} A_{\text{ALL}}} = 0)$.

We first determine the advertisers' best responses. If the platform plays C, the advertisers choose between $A^*_{P_C A_\Theta}$ and $A^*_{P_C A_{\text{ALL}}}$. By Proposition 5, $A^*_{P_C A_\Theta} \leq A^*_{P_C A_{\text{ALL}}}$, so the advertisers weakly prefer ALL. If the platform plays ALL, the advertisers receive a surplus of 0 regardless of their choice, making them indifferent between Θ and ALL.

Next, we determine the platform's best responses. If the advertisers play ALL, the platform chooses between $\Pi^*_{P_C A_{\text{ALL}}}$ and $\Pi^*_{P_{\text{ALL}} A_{\text{ALL}}}$. By Proposition 6, $\Pi^*_{P_C A_{\text{ALL}}} \leq \Pi^*_{P_{\text{ALL}} A_{\text{ALL}}}$, so the platform strictly prefers ALL. If the advertisers play Θ, the platform chooses between $\Pi^*_{P_C A_\Theta}$ and $\Pi^*_{P_{\text{ALL}} A_\Theta}$. By Proposition 4, $\Pi^*_{P_C A_\Theta} \geq \Pi^*_{P_{\text{ALL}} A_\Theta}$ for $K \geq 8$, and $\Pi^*_{P_C A_\Theta} < \Pi^*_{P_{\text{ALL}} A_\Theta}$ for $K < 8$.

Now we find the pure-strategy Nash equilibria:

- For $K \geq 8$: The platform prefers C in response to Θ, but the advertisers prefer ALL in response to C. Thus, (C, Θ) is not an equilibrium. If the profile is (ALL, Θ), the platform deviates to C. If the profile is (C, ALL), the platform deviates to ALL. For (ALL, ALL), the platform strictly prefers ALL (since $\Pi^*_{P_{\text{ALL}} A_{\text{ALL}}} \geq \Pi^*_{P_C A_{\text{ALL}}}$) and the advertisers are indifferent (both give 0). Thus, (ALL, ALL) is the unique equilibrium.

- For $K < 8$: The platform strictly prefers ALL regardless of the advertisers' strategy (since $\Pi^*_{P_{\text{ALL}} A_\Theta} > \Pi^*_{P_C A_\Theta}$ and $\Pi^*_{P_{\text{ALL}} A_{\text{ALL}}} \geq \Pi^*_{P_C A_{\text{ALL}}}$). Thus, ALL is a strictly dominant strategy for the platform. Given that the platform plays ALL, the advertisers earn 0 either way and are indifferent. Therefore, both (ALL, ALL) and (ALL, Θ) are equilibria.

In all identified equilibria, the platform plays ALL, which results in the advertisers' surplus being fully extracted ($A = 0$). This completes the proof. ■

Proof of Corollary 4

We analyze the normal-form game where the platform chooses whether to share its information $\mathbf{c}$ with the advertisers ($D_P \in \{\text{Share}, \text{Don't Share}\}$), and the advertisers choose whether to share their information $\boldsymbol{\theta}$ with the platform ($D_A \in \{\text{Share}, \text{Don't Share}\}$). The payoffs (u_P, u_A) for each strategy profile are:

- $(\text{Don't Share}, \text{Don't Share})$ results in model $P_C A_\Theta$ with payoffs $(\Pi^*_{P_C A_\Theta}, A^*_{P_C A_\Theta})$.
- $(\text{Share}, \text{Don't Share})$ results in model $P_C A_{\text{ALL}}$ with payoffs $(\Pi^*_{P_C A_{\text{ALL}}}, A^*_{P_C A_{\text{ALL}}})$.
- $(\text{Don't Share}, \text{Share})$ results in model $P_{\text{ALL}} A_\Theta$ with payoffs $(\Pi^*_{P_{\text{ALL}} A_\Theta}, 0)$.
- $(\text{Share}, \text{Share})$ results in model $P_{\text{ALL}} A_{\text{ALL}}$ with payoffs $(\Pi^*_{P_{\text{ALL}} A_{\text{ALL}}}, 0)$.

We first evaluate the advertisers' optimal strategy. If the advertisers choose Share, their surplus is fully extracted by the platform, resulting in a payoff of 0. If they choose Don't Share, they earn either $A^*_{P_C A_\Theta} > 0$ or $A^*_{P_C A_{\text{ALL}}} > 0$. Therefore, Don't Share is a strictly dominant strategy for the advertisers.

Knowing that the advertisers will play Don't Share, the platform's decision reduces to choosing between Share (getting $\Pi^*_{P_C A_{\text{ALL}}}$) and Don't Share (getting $\Pi^*_{P_C A_\Theta}$).

By Proposition 5, we know that $\Pi^*_{P_C A_{\text{ALL}}} > \Pi^*_{P_C A_\Theta}$ for $K \geq 8$, and $\Pi^*_{P_C A_{\text{ALL}}} \leq \Pi^*_{P_C A_\Theta}$ for $K < 8$.

- For $K \geq 8$: The platform maximizes its payoff by choosing Share. The unique equilibrium is $(\text{Share}, \text{Don't Share})$, which corresponds to the environment $P_C A_{\text{ALL}}$.
- For $K < 8$: The platform maximizes its payoff by choosing Don't Share. The unique equilibrium is $(\text{Don't Share}, \text{Don't Share})$, which corresponds to the environment $P_C A_\Theta$.

This completes the proof. ■

Proof of Lemma 7

We analyze the platform's expected payoffs for an arbitrary absolute cutoff $\tau \in [0, 1]$. An advertiser is invited to the auction if and only if its conversion probability satisfies $c_i \geq \tau$. Since c_i is drawn from a standard uniform distribution, the probability of any given advertiser being invited is $p = 1 - \tau$. The number of participating bidders n therefore follows a Binomial distribution, $n \sim \text{Binomial}(K, p)$.

For Part 1, conditional on being invited, an advertiser's posterior belief about its own conversion probability is a uniform distribution on $[\tau, 1]$. The expected conversion probability is thus $\bar{c} = \frac{1+\tau}{2}$. In a second-price auction, it is a weakly dominant strategy for each invited advertiser to bid its expected valuation:

$$\beta(\theta_i) = \theta_i \cdot \mathbb{E}[c_i \mid c_i \geq \tau] = \theta_i \frac{1+\tau}{2} = \theta_i \bar{c}.$$

For Part 2, from the platform's perspective, the bids of the n participating advertisers are independent and uniformly distributed on $[0, \bar{c}]$. The optimal reserve price for uniformly distributed bids on $[0, \bar{c}]$ is independent of the number of bidders n, and is given by $r^* = \bar{c}/2$. Substituting $\bar{c}$, we get $r^* = \frac{1+\tau}{4}$.

For Part 3, we compute the platform's expected payoff. Conditional on n bidders participating, the expected revenue is given by the formula derived in Corollary 1:

$$\Pi^*(n) = \bar{c}\frac{n-1+2^{-n}}{n+1} = \bar{c}\left(1 - \frac{2}{n+1} + \frac{2^{-n}}{n+1}\right).$$

To find the unconditional expected revenue $\Pi^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau)$, we take the expectation over $n \sim$ Binomial(K, p). We use the following identities for binomial sums:

$$\mathbb{E}\left[\frac{1}{n+1}\right] = \sum_{n=0}^{K}\binom{K}{n}p^n(1-p)^{K-n}\frac{1}{n+1} = \frac{1-(1-p)^{K+1}}{p(K+1)},$$
$$\mathbb{E}\left[\frac{2^{-n}}{n+1}\right] = \sum_{n=0}^{K}\binom{K}{n}p^n(1-p)^{K-n}\frac{2^{-n}}{n+1} = \frac{2}{p(K+1)}\left(\left(1-\frac{p}{2}\right)^{K+1} - (1-p)^{K+1}\right).$$

Noting that $1 - p/2 = 1 - \frac{1-\tau}{2} = \frac{1+\tau}{2} = \bar{c}$, the second sum simplifies to $\frac{2}{p(K+1)}\left(\bar{c}^{K+1} - (1-p)^{K+1}\right)$. Substituting these into the expectation, we get:

$$\Pi^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau) = \bar{c}\left(1 - \frac{2-2(1-p)^{K+1}}{p(K+1)} + \frac{2\bar{c}^{K+1} - 2(1-p)^{K+1}}{p(K+1)}\right) = \bar{c}\left(1 - \frac{2}{p(K+1)} + \frac{2\bar{c}^{K+1}}{p(K+1)}\right)$$
$$= \bar{c}\left(\frac{p(K+1) - 2 + 2\bar{c}^{K+1}}{p(K+1)}\right).$$

Since $p = 2(1-\bar{c})$, we have $p(K+1) - 2 = 2(K+1) - 2\bar{c}(K+1) - 2 = 2K - 2\bar{c}(K+1)$. Thus:

$$\Pi^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau) = \bar{c}\left(\frac{2K - 2\bar{c}(K+1) + 2\bar{c}^{K+1}}{p(K+1)}\right) = \frac{2\bar{c}}{p}\left[\frac{K}{K+1} - \bar{c} + \frac{1}{K+1}\bar{c}^{K+1}\right].$$

Substituting $\bar{c} = \frac{1+\tau}{2}$ and $p = 1 - \tau$ gives the exact formula in the lemma.

For Part 5, conditional on n bidders, the expected social welfare from Corollary 1 is $SW^*(n) = \bar{c}\frac{n}{n+1}(1 - 2^{-(n+1)})$. Taking the expectation over n:

$$SW^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau) = \bar{c}\,\mathbb{E}\left[1 - \frac{1}{n+1} - 2^{-(n+1)} + \frac{2^{-(n+1)}}{n+1}\right].$$

Using the previous sums and $\mathbb{E}[2^{-(n+1)}] = \frac{1}{2}\mathbb{E}[(1/2)^n] = \frac{1}{2}(1-p/2)^K = \frac{1}{2}\bar{c}^K$, we get:

$$
\begin{aligned}
SW^{\text{TH}}_{P_C^{\text{Com}} A_\Theta}(\tau) &= \bar{c}\left(1 - \frac{1-(1-p)^{K+1}}{p(K+1)} - \frac{1}{2}\bar{c}^K + \frac{\bar{c}^{K+1}-(1-p)^{K+1}}{p(K+1)}\right) \\
&= \bar{c}\left(1 - \frac{1}{2}\bar{c}^K + \frac{\bar{c}^{K+1}-1}{p(K+1)}\right).
\end{aligned}
$$

Substituting $\bar{c} = \frac{1+\tau}{2}$ and $p = 1-\tau$ gives the social welfare formula.

For Part 4, the advertiser surplus is the difference between social welfare and platform revenue:

$$
\begin{aligned}
A^{\text{TH}}_{P_C^{\text{Com}} A_\Theta}(\tau) &= SW^{\text{TH}}_{P_C^{\text{Com}} A_\Theta}(\tau) - \Pi^{\text{TH}}_{P_C^{\text{Com}} A_\Theta}(\tau) = \bar{c}\left(1 - \frac{1}{2}\bar{c}^K + \frac{\bar{c}^{K+1}-1}{p(K+1)} - 1 + \frac{2}{p(K+1)} - \frac{2\bar{c}^{K+1}}{p(K+1)}\right) \\
&= \bar{c}\left(-\frac{1}{2}\bar{c}^K + \frac{1-\bar{c}^{K+1}}{p(K+1)}\right) = \frac{-p(K+1)\bar{c}^{K+1} + 2\bar{c} - 2\bar{c}^{K+2}}{2p(K+1)} \\
&= \frac{2\bar{c} - \bar{c}^{K+1}(p(K+1)+2\bar{c})}{2p(K+1)}.
\end{aligned}
$$

Since $2\bar{c} = 2-p$, we have $p(K+1) + 2\bar{c} = pK + p + 2 - p = pK + 2$. Substituting this back:

$$
A^{\text{TH}}_{P_C^{\text{Com}} A_\Theta}(\tau) = \frac{2\bar{c} - \bar{c}^{K+1}(2+pK)}{2p(K+1)}.
$$

Substituting $\bar{c} = \frac{1+\tau}{2}$ and $p = 1-\tau$ gives the formula for advertiser surplus in the lemma.

Finally, because the platform has commitment power, it will choose the cutoff τ^* that maximizes its expected payoff $\Pi^{\text{TH}}_{P_C^{\text{Com}} A_\Theta}(\tau)$. Since the payoff function is differentiable on $(0,1)$, the optimal interior cutoff τ^* satisfies the first-order condition $\partial\Pi^{\text{TH}}_{P_C^{\text{Com}} A_\Theta}(\tau)/\partial\tau = 0$ evaluated at $\tau = \tau^*$. ■

Proof of Lemma 8

The revenue under the top-n strategy is $\Pi^{\text{TOP}}(n) = \frac{(2K-n+1)\left(n-1+\frac{1}{2^n}\right)}{2(n+1)(K+1)}$ and n^* is the optimal integer $n \in [1, K]$ that maximizes $\Pi^{\text{TOP}}(n)$. The revenue under the threshold-τ strategy is $\Pi^{\text{TH}}(\tau) = \frac{1+\tau}{1-\tau}\left[\frac{K}{K+1} - \frac{1+\tau}{2} + \frac{1}{K+1}\left(\frac{1+\tau}{2}\right)^{K+1}\right]$ and τ^* is the optimal real $\tau \in [0,1]$ that maximizes $\Pi^{\text{TH}}(\tau)$. We want to prove that $\Pi^{\text{TH}}(\tau^*) \geq \Pi^{\text{TOP}}(n^*)$.

We define the functions $f(x) = \frac{x}{(K+1)(1-x)}\left(K - (K+1)x + \left(\frac{1}{2^{2(1-x)}}\right)^{K+1}\right)$ and $g(x) = \frac{x}{(K+1)(1-x)}\left(K - (K+1)x + x^{K+1}\right)$ for $x \in [1/2, 1]$. Notice that $\Pi^{\text{TOP}}(n) = f\left(1 - \frac{n+1}{2(K+1)}\right)$ and $\Pi^{\text{TH}}(\tau) = g\left(\frac{1+\tau}{2}\right)$.

Since $x \geq \frac{1}{2^{2(1-x)}}$ for any $x \in [1/2, 1]$, it holds that $g(x) \geq f(x)$ for any $x \in [1/2, 1]$. Let

$x_g^* \in [1/2, 1]$ be the value of x that maximizes $g(x)$ and $x_f^* \in [1/2, 1]$ be the value of x that maximizes $f(x)$. It is $x_g^* = \frac{1+\tau^*}{2}$, therefore

$$\Pi^{\text{TH}}(\tau^*) = g\left(\frac{1+\tau^*}{2}\right) = g(x_g^*) \geq g(x_f^*) \geq f(x_f^*) \geq f\left(1 - \frac{n^*+1}{2(K+1)}\right) = \Pi^{\text{TOP}}(n^*). \quad \blacksquare$$

Proof of Lemma 9

We show that $\Pi^*_{P_C^{\text{COM}} A_\Theta} \geq \Pi^*_{P_C A_{\text{ALL}}}$ by constructing a feasible commitment strategy whose expected revenue is arbitrarily close to $\Pi^*_{P_C A_{\text{ALL}}}$; since the platform can do at least as well as any feasible strategy, the inequality follows.

Under commitment, the platform publicly announces (before the realization of $\mathbf{c}$) a selection rule $\mathcal{S}(\mathbf{c})$ and a reserve-price rule $r(\mathbf{c})$. Since these are publicly known, each invited advertiser i can compute its posterior belief about c_i from the observed r and the publicly announced rules.

Consider the following commitment strategy. Let $r^*_{P_C A_{\text{ALL}}} : [0,1]^K \to \mathbb{R}_+$ denote the optimal reserve-price function in model $P_C A_{\text{ALL}}$ (as characterized in Algorithm 1), and let $\phi : [0,1]^K \to [0,1)$ be an injection (e.g., the decimal-interleaving map that interleaves the decimal digits of $c_1, \ldots, c_K$). For each integer $m \geq 1$, we define:

- Selection rule: $\mathcal{S}(\mathbf{c}) = \{1, \ldots, K\}$ (invite all advertisers).
- Reserve-price rule: $r_m(\mathbf{c}) = \frac{\lfloor r^*_{P_C A_{\text{ALL}}}(\mathbf{c}) \cdot 10^m \rfloor}{10^m} + \frac{1}{10^m} \cdot \phi(\mathbf{c})$.

Since $\lfloor r^*_{P_C A_{\text{ALL}}}(\mathbf{c}) \cdot 10^m \rfloor$ is a non-negative integer and $\phi(\mathbf{c}) \in [0,1)$, the fractional part of $r_m(\mathbf{c}) \cdot 10^m$ equals $\phi(\mathbf{c})$. Because the rule $r_m(\cdot)$ is publicly committed, each invited advertiser can recover $\phi(\mathbf{c})$ from the observed r_m and, by injectivity of ϕ, recover the full vector $\mathbf{c}$. In particular, advertiser i learns its own c_i and bids $c_i \theta_i$, exactly as in model $P_C A_{\text{ALL}}$.

The reserve-price error satisfies $|r_m(\mathbf{c}) - r^*_{P_C A_{\text{ALL}}}(\mathbf{c})| \leq 1/10^m$ for all $\mathbf{c}$. Since the platform's expected revenue is continuous in the reserve price, the revenue under this strategy converges to $\Pi^*_{P_C A_{\text{ALL}}}$ as $m \to \infty$.

More precisely, for any $\delta > 0$, there exists an $m \geq 1$ such that the commitment strategy $(\mathcal{S}, r_m)$ yields expected revenue at least $\Pi^*_{P_C A_{\text{ALL}}} - \delta$. Since the platform optimizes over all commitment strategies, its optimal committed revenue satisfies $\Pi^*_{P_C^{\text{COM}} A_\Theta} \geq \Pi^*_{P_C A_{\text{ALL}}} - \delta$ for all $\delta > 0$, which implies $\Pi^*_{P_C^{\text{COM}} A_\Theta} \geq \Pi^*_{P_C A_{\text{ALL}}}$.

The inequality is strict for small K (e.g., $K \leq 9$), as can be verified from Tables 3 and 5: the threshold commitment strategy alone already outperforms model $P_C A_{\text{ALL}}$ for these values. For large K ($K \geq 10$), the comparison between the threshold strategy and $P_C A_{\text{ALL}}$ favors the latter (Table 3), but the platform's optimal commitment strategy (which is not restricted to the threshold class) achieves at least $\Pi^*_{P_C A_{\text{ALL}}}$ by the argument above. ■

A.2 Numerical Examples

K	$\Pi^*_{P_C A_\Theta}$	$SW^*_{P_C A_\Theta}$	$A^*_{P_C A_\Theta}$
2	0.208	0.292	0.083
3	0.266	0.352	0.086
4	0.319	0.422	0.103
5	0.357	0.452	0.095
6	0.394	0.498	0.104
7	0.4211	0.533	0.112
8	0.4479	0.547	0.099
9	0.470	0.574	0.104
10	0.489	0.597	0.108

Table 1: Model $P_C A_\Theta$ equilibrium, for different values of K.

K	$\Pi^*_{P_{\text{ALL}} A_\Theta}$	$SW^*_{P_{\text{ALL}} A_\Theta}$	$A^*_{P_{\text{ALL}} A_\Theta}$
2	0.333	0.333	0
3	0.375	0.375	0
4	0.400	0.400	0
5	0.417	0.417	0
6	0.429	0.429	0
7	0.438	0.438	0
8	0.444	0.444	0
9	0.450	0.450	0
10	0.455	0.455	0

Table 2: Model $P_{\text{ALL}} A_\Theta$ equilibrium, for different values of K.

K	$\Pi^*_{P_C A_{\text{ALL}}}$	$SW^*_{P_C A_{\text{ALL}}}$	$A^*_{P_C A_{\text{ALL}}}$
2	0.202	0.312	0.110
3	0.262	0.398	0.136
4	0.312	0.461	0.150
5	0.354	0.510	0.156
6	0.390	0.549	0.159
7	0.4210	0.580	0.159
8	0.4480	0.605	0.157
9	0.472	0.627	0.155
10	0.493	0.645	0.152
11	0.512	0.660	0.148
12	0.529	0.673	0.145
13	0.544	0.685	0.141
14	0.558	0.696	0.138

Table 3: Model $P_C A_{\text{ALL}}$ equilibrium, for different values of K.

K	$\Pi^*_{P_{\text{ALL}} A_{\text{ALL}}}$	$SW^*_{P_{\text{ALL}} A_{\text{ALL}}}$	$A^*_{P_{\text{ALL}} A_{\text{ALL}}}$
2	0.370	0.370	0
3	0.445	0.445	0
4	0.498	0.498	0
5	0.538	0.538	0
6	0.569	0.569	0
7	0.594	0.594	0
8	0.616	0.616	0
9	0.634	0.634	0
10	0.650	0.650	0

Table 4: Model $P_{\text{ALL}} A_{\text{ALL}}$ equilibrium, for different values of K.

K	$\Pi^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau^*)$	$SW^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau^*)$	$A^{\text{TH}}_{P_C^{\text{COM}} A_\Theta}(\tau^*)$
2	0.210	0.297	0.087
3	0.274	0.371	0.097
4	0.324	0.425	0.102
5	0.364	0.467	0.103
6	0.398	0.501	0.103
7	0.427	0.529	0.103
8	0.451	0.553	0.101
9	0.473	0.573	0.100
10	0.493	0.591	0.098
11	0.510	0.607	0.097
12	0.526	0.621	0.095
13	0.540	0.633	0.094
14	0.553	0.645	0.092

Table 5: Model $P_C^{\text{COM}} A_\Theta$ under the optimal threshold strategy, for different values of K.