State-Dependent Predictive Value and Information Asymmetry in Online Advertising

Wee Chaimanowong, Francisco Cisternas, Stylianos Despotakis*

February 13, 2026

Abstract

In online advertising auctions, the predictive value of consumer data can depend on an unobservable consumer state, such as whether a traveler is on business or vacation. We call this *state-dependent predictive value* and study its interaction with asymmetric data access across advertisers. We develop a game-theoretic model that features latent consumer intent, heterogeneous advertiser information, and a data component whose relevance varies with the unobservable state. We analyze three information settings (full information, partial information, and information asymmetry) under both correlated and independent valuation structures. When detailed data shifts advertisers' valuations in a correlated direction, information asymmetry distorts equilibrium bidding: constrained advertisers systematically overbid or underbid, reducing both allocative efficiency and publisher revenue. Restricting data access for all advertisers eliminates this distortion, simultaneously improving conversion rates and revenue, overturning the presumed tradeoff between privacy and market efficiency. When valuations are independent, more information improves conversion rates but can reduce revenue, and information asymmetry can outperform both symmetric settings. Finally, advertisers' preferences over data regimes do not always align with their information status: informed advertisers may prefer that their competitors also gain data access, eliminating their own advantage, while uninformed advertisers may prefer that informed competitors retain access.

Keywords: Online advertising; Targeting; Information asymmetry; Conversion rate; Auction efficiency; Auction theory.

*The authors are listed in alphabetical order. Wee Chaimanowong: chaimanowongw@gmail.com (CUHK Business School, The Chinese University of Hong Kong), Francisco Cisternas: fcisternas@cuhk.edu.hk (CUHK Business School, The Chinese University of Hong Kong), Stylianos Despotakis (corresponding author): sdespota@cityu.edu.hk (City University of Hong Kong). The authors thank Amin Sayedi, Tony Ke, Chenxi Liao, Jesse Yao, and various seminar participants for their helpful feedback and comments. Despotakis acknowledges financial support from the Research Grants Council of Hong Kong (Project No. CityU 21509121).

1 Introduction

In online advertising markets, the value of observable user data often depends on unobservable consumer intent. Consider a user visiting a travel website to book a flight. Several hotel advertisers compete to display an ad to this user. The user may be planning a business trip or a vacation, an unobservable distinction that can affect their probability of booking any given hotel. Now consider a specific user characteristic, such as price sensitivity. Some advertisers can infer this from proprietary behavioral data or third-party cookies, while others cannot. Importantly, the predictive value of this information depends on the user's unobservable intent: for a business traveler whose expenses are reimbursed, personal price sensitivity may be largely irrelevant to hotel choice; for a vacation planner, it may be a key determinant of conversion. We call this phenomenon *state-dependent predictive value*: the relevance of observable data varies with an unobservable consumer state. This interaction between latent consumer heterogeneity and asymmetric advertiser information access characterizes many online advertising markets, from restaurant searches to retail shopping, yet has received limited theoretical attention.

Understanding this interaction is particularly important given ongoing shifts in how consumer data is accessed and used. Major browsers are phasing out third-party cookies, privacy regulations increasingly limit data collection and sharing, and platforms face growing pressure to restrict advertiser access to user information. These changes are commonly framed as a tradeoff between privacy and economic efficiency: protecting consumer data at the cost of less relevant ads and lower revenues. Yet reported empirical evidence complicates this narrative. When the Dutch public broadcaster NPO stopped using tracking cookies, it experienced a substantial increase in advertising revenue (Edelman, 2020), and subsequent experiments revealed higher conversion rates with contextual rather than microtargeted ads in certain cases (Snelders et al., 2020). The New York Times reported similar outcomes after discontinuing behavioral targeting (Davies, 2019). Reduced intermediary costs may partially explain the revenue gains (Hsiao, 2020), but they do not account for the observed improvements in conversion rates. While these observations are context-specific and potentially confounded by concurrent changes in supply-side dynamics, they motivate a theoretical puzzle: Under what conditions can reducing the information available to bidders improve not only seller revenue but also the quality of the match between ads and consumers?

Existing auction theory provides partial but incomplete answers. The classical linkage principle (Milgrom and Weber, 1982) establishes that a seller can increase revenue by publicly revealing information it holds, because doing so reduces bidder uncertainty and strengthens competition. On the other hand, in the advertising context, Levin and Milgrom (2010) shows that better-informed advertisers can thin out markets, reducing competition and prices, a finding that highlights how information can decrease revenues through market structure effects. However, these frameworks generally treat the relevance of available information as uniform across consumers, rather than allowing it to vary with an unobservable consumer state. Furthermore, while revenue effects have been studied extensively, the impact of information structures on the quality of the ad-to-consumer match (i.e., conversion rates) has received comparatively little attention. The prevailing assumption, that more information should yield better ad-to-consumer matches even when it reduces revenues for competitive reasons, has not been subjected to formal theoretical scrutiny in settings where information relevance itself is state-dependent.

This paper develops a game-theoretic model of advertising auctions designed to address this gap. Our model incorporates latent consumer intent that is unobservable to advertisers, heterogeneous data access across advertisers, and, critically, the feature that additional data's predictive value varies with the consumer's unobservable state. We further allow detailed data to shift advertisers' valuations in either correlated or independent ways, capturing different market structures. We analyze three information settings: *Full Information*, where all advertisers access detailed consumer data; *Partial Information*, where no advertiser does; and *Information Asymmetry*, where only a subset of advertisers has access. For each setting, we derive equilibrium bidding strategies and compare outcomes along three dimensions: publisher revenue, conversion rates, and advertiser payoffs.

Our analysis produces three main results. First, when detailed data generates correlated shifts in advertiser valuations (e.g., a user's price sensitivity affects all hotel advertisers' conversion probabilities similarly), restricting data access uniformly can improve both publisher revenue and conversion rates relative to the asymmetric setting. The mechanism is that information asymmetry distorts equilibrium bidding: uninformed advertisers face adverse selection and systematically overbid or underbid, reducing allocative efficiency. Removing the asymmetry eliminates this distortion and improves average match quality between ads and users. Second, when valuations are independent across advertisers, Information Asymmetry can generate higher revenues than either Full or Partial Information, suggesting that selective data access policies may outperform uniform ones. Third, advertisers' preferences over information settings need not align with their information status: informed advertisers may prefer that their competitors also gain data access, eliminating their own advantage, while uninformed advertisers may prefer that informed competitors retain access rather than face a uniformly uninformed market, because the competitive effects of information can dominate direct informational benefits.

These results contribute to several ongoing discussions. For platforms evaluating cookie deprecation or data-sharing restrictions, our findings identify specific market conditions (namely, correlated valuation shifts and pre-existing information asymmetry) under which privacy-enhancing policies improve both revenue and ad relevance simultaneously. A platform can assess the applicability of these conditions by examining whether detailed consumer data (e.g., browsing history, behavioral signals) tends to shift advertisers' valuations in the same direction across competitors in a given auction, or whether its effects are idiosyncratic. For advertisers, the analysis reveals that information advantages do not reliably translate into competitive benefits, particularly in markets where detailed data shifts valuations in independent ways. For regulators, we show that the welfare implications of data restrictions depend on the correlation structure of advertiser valuations and the existing degree of information asymmetry, factors that must be understood before the net effects of privacy regulation can be assessed.

2 Related Literature

This research contributes to the growing literature on targeted advertising and online advertising auctions. We organize our discussion of the most closely related theory papers around the following themes: the effects of targeting on firm competition, the role of information structures in advertising auctions, and the implications of privacy and data restrictions.

Targeting and Firm Competition

A substantial body of work examines how improved targeting ability affects competition among advertisers. A general finding in this literature is that targeting tends to increase firms' profits and the number of consumer-product matches, but the effects can be non-monotonic or even negative under certain conditions. Iyer et al. (2005) describe a model of competing firms who can target different segments of consumers with advertising and show that targeted advertising will improve the firms' profits and, moreover, it can sometimes be more valuable than targeted pricing. Bergemann and Bonatti (2011) show that better targeting causes an increase in the number of consumer-product matches, but prices of ads change non-monotonically in the targeting capacity.

However, several papers identify conditions under which targeting can harm firms. Chen et al. (2001) study the effects of imperfect targetability on prices for different segments of consumers. Interestingly, they find that improving the targetability of a firm can sometimes benefit both the firm and its competitor. Brahim et al. (2011) study a model with two firms competing in prices and targeted advertising. They show that firms' profits can be lower with targeted relative to random advertising. Despotakis and Yu (2022) study a multidimensional targeting model and show that sometimes the use of multiple dimensions of data to target consumers can have negative effects for a firm. Johnson (2013) considers targeted advertising in combination with advertising avoidance technology. He shows that targeting will increase firms' profits, but it can make consumers worse off.

Our model departs from this stream in that we do not study the effects of improved targeting *per se*. Instead, we examine what happens when the same consumer data has different predictive value depending on an unobservable consumer state, and when advertisers differ in their ability to access that data. This combination of state-dependent data value and information asymmetry generates effects that are absent from models with symmetric advertisers or state-independent data.

Information Structures in Advertising Auctions

A second strand of literature studies how information is, or should be, structured in advertising auction environments. The foundational work of Milgrom and Weber (1982) on affiliated-value auctions established the linkage principle: that a seller benefits from publicly revealing information it holds, because doing so reduces the winner's curse and intensifies competition.

In the specific context of advertising, De Corniere and De Nijs (2016) show that when a platform chooses to reveal the information it has about a consumer to advertisers, the advertisers will set higher prices in anticipation of a better matching. This will benefit the advertisers and the platform.

Our setting differs in that revealing information can actually worsen the matching between the advertisers and the consumer, resulting in a lower social welfare. This is because in addition to the full disclosure or non-disclosure of information, we also consider the case where not all advertisers have access to the same information about the consumer, and this asymmetry plays a significant role in our model. Ada et al. (2022) study the impact of providing ad context information in ad exchanges and find that in most cases ad exchanges can boost publishers' revenues by sharing context information with ad buyers. Bobkova (2024) discusses the role of the advertisers' (bidder) endogenous information choice in auction design and focuses on the symmetric Bayesian Nash equilibrium outcomes. This is unlike our work where we assume the information to be controlled by the publisher which allows us to focus on the implications of information asymmetry.

Our model differs from this literature in two important respects. First, whereas much of the information design literature considers a principal who controls the precision of information provided symmetrically to agents, we study settings where information access is inherently asymmetric across advertisers, and the question is whether to eliminate this asymmetry by restricting or expanding access. Second, in existing models, revealing information generally improves the match between the auction winner and the object; in our setting, revealing information to an asymmetric set of advertisers can *worsen* the match between ads and consumers, because the resulting bidding distortions outweigh the informational gains.

Privacy, Data Restrictions, and Market Outcomes

A third group of papers examines how restrictions on data access affect advertising market outcomes. Levin and Milgrom (2010) provides an influential analysis showing that superior information can thin out competition in advertising auctions. Building on this insight, Rafieian and Yoganarasimhan (2021) show that the revenues of ad-networks can increase when users preserve their privacy, because more precise targeting can reduce competition. However, when this happens, the targeting becomes less efficient. Our model replicates this effect in the case of independent valuations among the advertisers, but in the case of dependent valuations, we show that revenue and targeting efficiency can move in the same direction, both improving when data access is restricted.

Several papers study related aspects of privacy and data in advertising markets. Zhang and Katona (2012) study how contextual advertising affects product market competition. Esteves and

Resende (2016) study how targeted advertising can be used by competing firms to price discriminate different segments of consumers. Shen and Miguel Villas-Boas (2018) study advertising based on the past purchase behavior of consumers and examine how it affects product prices for a monopolist. Hummel and McAfee (2016) study how the number of bidders in an auction affects a seller's revenue under two different settings (bundling vs. targeting), though unlike our model, they assume independent valuations.

Advertising Market Structure

Finally, a growing literature examines how the organizational structure of advertising markets interacts with information. Choi and Sayedi (2023b) investigate the effects of private exchanges on the display advertising market, finding that while private exchanges offer higher quality impressions compared to open exchanges, they can also create information asymmetry among advertisers, which can hurt publisher's revenue. Choi and Sayedi (2023a) examine the effects of ad agencies on the online advertising market, revealing that publishers face a trade-off when deciding whether to withhold targeting information from agencies, which can either mitigate "bid rotation" and attract direct advertisers or reduce the efficiency of allocation for agency-using advertisers. Shin and Shin (2022) demonstrate that irrelevant advertising can stem from strategic decisions within the ad agency-advertiser relationship, rather than simply technological imperfections. The study also explores how contractual restrictions can lead to inefficiencies in ad delivery, and suggests that the prevalence of irrelevant ads may decrease, but not disappear, as the number of impressions available in the market increases.

Contribution

This paper contributes to the targeted advertising literature by introducing and analyzing two features that have not been studied jointly: *state-dependent data value*, whereby the predictive value of observable consumer data depends on an unobservable consumer state, and *information asymmetry* across advertisers in access to that data. We show that the interaction of these two features generates some unique predictions: in the common-value case, both publisher revenue and conversion rates can increase simultaneously when data access is restricted, a result that does not arise with symmetric advertisers. The paper also contributes to the online advertising auction literature by providing a unified analysis across different valuation structures (correlated and independent) and different information settings (full, partial, and asymmetric), offering a comprehensive account of how information asymmetries shape equilibrium market outcomes.

3 Model

3.1 General Setup

We consider n advertisers who compete in an ad auction for a user impression provided by a publisher.

Users enter the market for one of two possible reasons on a given occasion: Reason 1 or Reason 2. The visiting reason is idiosyncratic and cannot be directly observed by advertisers, but they can assign a predictive probability based on available individual characteristics. For a user j with characteristic data $\mathcal{D}_j$, the probability the visit is due to Reason 1 is given by $\kappa(\mathcal{D}_j) := \Pr[\text{Reason 1} \mid \mathcal{D}_j] = 1 - \Pr[\text{Reason 2} \mid \mathcal{D}_j]$.

For example, a user j may visit a travel website either for a vacation or a business trip. Based on prior market knowledge, a hotel advertiser can predict that the user's reason to visit on this particular occasion is a vacation with probability $\kappa(\mathcal{D}_j)$ and a business trip with probability $1 - \kappa(\mathcal{D}_j)$. Similarly, a user searching for restaurants might be planning a romantic date (Reason 1) or catching up with a friend (Reason 2), and a user shopping online might be buying a gift (Reason 1) or purchasing something for themselves (Reason 2). In each case, advertisers cannot observe the actual reason on a particular occasion, so they assign probabilities based on their market knowledge.[1]

Not every advertiser i will be an equally good match for a user j with a particular visiting reason. Given a user j with characteristic data $\mathcal{D}_j$, if they see an ad from advertiser i, the probability of conversion conditional on a Reason 1 visit is $b_i(\mathcal{D}_j) := \Pr[\text{Conversion} \mid \text{Advertiser } i,\ \mathcal{D}_j,\ \text{Reason 1}]$ and the probability of conversion conditional on a Reason 2 visit is $c_i(\mathcal{D}_j) := \Pr[\text{Conversion} \mid \text{Advertiser } i,\ \mathcal{D}_j,\ \text{Reason 2}]$, where $b_i(\cdot)$, $c_i(\cdot)$ are advertiser-specific functions from user data to a value in $[0, 1]$.

[1] The two-state structure is a deliberate modeling choice. While real-world consumer intent may be richer (a traveler could be browsing, comparing prices, planning a future trip, or booking imminently), the essential feature we aim to capture is that some observable data is relevant for predicting conversion under one intent but not another. The binary structure is the simplest framework that generates this state-dependent data value.

Given the above, for an advertiser i, the expected probability of conversion after showing an ad to a user j of unknown Reason is

$$v_{ij} := \Pr[\text{Conversion} \mid \text{Advertiser } i,\ \mathcal{D}_j] = \kappa(\mathcal{D}_j) \cdot b_i(\mathcal{D}_j) + (1 - \kappa(\mathcal{D}_j)) \cdot c_i(\mathcal{D}_j).$$

For simplicity, we assume that each advertiser i is willing to pay up to a normalized amount of 1 for a conversion, so that their valuation (maximum willingness to pay) to show an ad to a user j of unknown Reason equals v_{ij}.[2]

We define the *conversion rate* of the auction as the expected probability that the winning advertiser's ad leads to a conversion, i.e., the expected value of v_{ij} for the winning advertiser (where the expectation is taken over the randomness in advertiser valuations and bidding). This measures the allocative efficiency of the auction: a higher conversion rate means the impression is allocated to an advertiser who is, on average, a better match for the user.

3.2 Information Constraints and Asymmetry

Not all of the data $\mathcal{D}_j$ may be necessary to determine each model primitive. Let $\mathcal{A}_j \subseteq \mathcal{D}_j$ denote the subset on which the visit-reason probability κ depends, and let $\mathcal{B}_j, \mathcal{C}_j \subseteq \mathcal{D}_j$ denote the subsets on which b_i and c_i depend, respectively (these subsets may overlap). We can then write $\kappa(\mathcal{A}_j)$, $b_i(\mathcal{B}_j)$, and $c_i(\mathcal{C}_j)$ instead of $\kappa(\mathcal{D}_j)$, $b_i(\mathcal{D}_j)$, and $c_i(\mathcal{D}_j)$.

If all advertisers could compute their v_{ij}'s, the auction would reduce to a standard complete-information setting. However, not all advertisers may know their v_{ij}'s, either because they lack access to some necessary data or because they lack the ability to translate that data into accurate conversion probability estimates.

The key asymmetry we study is that $\mathcal{A}_j$ and $\mathcal{C}_j$ are accessible to all advertisers, whereas $\mathcal{B}_j$ is not. The assumption that the information asymmetry falls on $\mathcal{B}_j$ (the data needed to estimate the Reason 1 conversion probability) rather than on $\mathcal{A}_j$ or $\mathcal{C}_j$ captures a common feature of online advertising markets: baseline conversion probabilities under the more common or "default" consumer

[2] This normalization assumes that advertisers differ in their conversion probabilities for a given user but not in their willingness to pay per conversion. In practice, heterogeneity in per-conversion value (e.g., due to different profit margins) is an important dimension of advertiser competition that could interact with the information structure we study. For instance, if high-value advertisers systematically invest more in data infrastructure, the correlation between information status and per-conversion value would be relevant. We abstract from this interaction to isolate the effects of information asymmetry on conversion-probability-based valuations. The model can be extended by replacing the range $[0, 1]$ of the functions $b_i(\cdot)$, $c_i(\cdot)$ with advertiser-specific ranges $[0, a_i]$, representing expected profits rather than conversion probabilities; we leave this generalization for future work.

state can often be estimated from widely available data (e.g., the destination of a flight booking, which is visible to all advertisers on the page), whereas predicting conversion under a less common or more nuanced state requires proprietary data (e.g., the user's price sensitivity from browsing history) or more sophisticated modeling capabilities. Continuing with the vacation-versus-business-trip example, a business traveler may be price sensitive as a person, but in the context of a business trip their price sensitivity may be irrelevant due to employer reimbursement. Therefore, $b_i(\mathcal{B}_j)$ can depend on price sensitivity data, while $\kappa(\mathcal{A}_j)$ and $c_i(\mathcal{C}_j)$ do not.[3]

For a specific user j, we call an advertiser who can utilize $\mathcal{A}_j$, $\mathcal{B}_j$, and $\mathcal{C}_j$ an *informed* advertiser, and an advertiser who can utilize only $\mathcal{A}_j$ and $\mathcal{C}_j$ a *constrained* advertiser. From now on, we focus on a specific user of unknown Reason and simplify notation by dropping the index j: we write κ, b_i, and c_i in place of $\kappa(\mathcal{A}_j)$, $b_i(\mathcal{B}_j)$, and $c_i(\mathcal{C}_j)$, and v_i in place of v_{ij} (keeping in mind that these values can vary across users). An informed advertiser knows their $v_i = \kappa \cdot b_i + (1 - \kappa) \cdot c_i$, while a constrained advertiser can only compute their expected valuation $\mathbb{E}_{\mathcal{B}}[v_i] = \kappa \cdot \mathbb{E}_{\mathcal{B}}[b_i] + (1 - \kappa) \cdot c_i$. Therefore, an informed advertiser bids based on v_i and a constrained advertiser bids based on $\mathbb{E}_{\mathcal{B}}[v_i]$.

3.3 Information Settings, Distributional Assumptions, and Auction Format

To study the effects of information asymmetry on the auction outcomes, we consider the following three information settings:

- **Full Information (FI)**: All n advertisers are informed.
- **Information Asymmetry (IA)**: n_1 advertisers are informed and n_2 advertisers are constrained (where $n_1, n_2 > 0$ and $n_1 + n_2 = n$).
- **Partial Information (PI)**: All n advertisers are constrained.

An Information-Asymmetry market can become a Partial-Information market if the publisher restricts data access for informed advertisers (e.g., by disabling third-party cookies so that $\mathcal{B}_j$ becomes unavailable). Conversely, it can become a Full-Information market if constrained advertisers gain the ability to estimate b_i. We compare outcomes across these three settings to determine when such transitions are beneficial.

[3] Alternatively, $\mathcal{B}_j$ may represent data in a format that requires specialized processing (e.g., unstructured browsing logs), while $\mathcal{A}_j$ and $\mathcal{C}_j$ correspond to structured, publicly available signals.

As additional terminology, when an advertiser bids based on their actual v_i (using both b_i and c_i), we call this *full targeting*. When an advertiser bids based on $\mathbb{E}_{\mathcal{B}}[v_i]$ (using only c_i), we call this *partial targeting*. Disabling full targeting (e.g., by removing third-party cookie access) corresponds to moving from Information Asymmetry to Partial Information.

Correlation structure of b-values. The values $b_1, b_2, \ldots, b_n$ of the different advertisers can be correlated or independent. For instance, a user's low price sensitivity might indicate a high b_i for all hotel advertisers, whereas certain past brand preferences might indicate a high b_i for some advertisers and a low b_i for others. We consider two extreme cases:

- **Common-value case (CV)**: $b_1 = b_2 = \ldots = b_n =: b$, where b is drawn from a distribution with CDF F.
- **Independent-values case (IV)**: $b_1, b_2, \ldots, b_n$ are i.i.d. draws from a distribution with CDF F.

For parsimony and analytical tractability, we model F as a Bernoulli distribution on $\{0, 1\}$, with $\Pr[b_i = 1] = p$ for some known probability $p \in [0, 1]$.

Correlation structure of c-values. The values $c_1, c_2, \ldots, c_n$ could also be identical or independent. Because the c_i's are known to all advertisers, the case where they are all identical is less interesting (there is no private information in this component). We defer the common c-value case to Appendix C and focus on independent c-values in the main analysis:

- The random variables c_i are i.i.d. draws from a distribution with CDF G.

We model G as a uniform distribution on $[0, 1]$ in the main exposition.[4]

Auction format. The impression is sold using a second-price auction run by the publisher. Advertisers bid based on their expectations about v_i, and the highest bidder wins and pays the second-highest bid.[5]

[4] This is for expositional clarity, but our results hold for arbitrary distributions G, as we demonstrate in Sections 5 and 6.

[5] In practice, first-price auctions have become the dominant format in programmatic advertising (see e.g., Despotakis et al., 2021). We use a second-price auction for analytical tractability, but our main results are robust to alternative selling mechanisms. In particular, Lemma 4 shows that the conversion rate result holds for a wide class of mechanisms, including first-price auctions and multiple parallel auctions. Proposition 12 in Appendix B.3 further demonstrates that the revenue result also holds under first-price auctions.

Notation. For each dependence setting $\sigma \in \{\mathrm{CV}, \mathrm{IV}\}$ (common-value, independent-values) and each information setting $\tau \in \{\mathrm{FI}, \mathrm{IA}, \mathrm{PI}\}$, we denote by W_τ^σ and V_τ^σ the publisher's expected revenue and the expected conversion rate, respectively. Similarly, E_τ^σ and D_τ^σ denote the informed advertiser's and the constrained advertiser's expected payoffs, respectively. Table 1 summarizes all notation.

The remainder of the paper is structured as follows. In Section 4, we present the main results and insights with $n = 2$ advertisers. Sections 5 and 6 then establish robustness in the general setting with any number of advertisers and arbitrary distributions G. Section 5 provides analytical results where closed-form proofs are feasible despite the lack of a closed-form bidding function.[6] Section 6 establishes the existence of a pure symmetric equilibrium for the general model[7] and numerically approximates the bidding function for general examples. All proofs are in Appendices A.1 and B.1, and a summary of key formulas appears in Appendix B.2. Appendix B.3 analyzes the first-price auction variant, and Appendix C treats the common c-value case.

4 Analysis and Main Insights

In this section, we start by presenting the results and intuitions for two advertisers: one informed advertiser and one constrained advertiser. In subsection 4.1 we consider the common-value case (CV) and in subsection 4.2 we consider the independent-values case (IV). In subsection 4.3 we compare and discuss the differences between the common-value and independent-values cases in terms of publisher's revenue, conversion rates, and advertisers' payoffs.

4.1 Common-value case

Under the common-value case in the PI and FI settings, both advertisers will have the same information, therefore, in the second-price auction they will truthfully bid their valuation (in FI) or their expected valuation (in PI, where they do not know the actual valuation) (see e.g., Krishna, 2009). However, in the IA setting the informed advertiser is more informed than the constrained advertiser. As a consequence of this asymmetry, the informed advertiser will still bid their true

[6] The bidding function of constrained advertisers is the solution to a differential equation that does not always admit a closed-form solution (see Eq. (4) and the proof of Lemma 3).

[7] Specifically, we show that a pure symmetric equilibrium bidding strategy exists for constrained advertisers under the common-value IA setting.

Information Settings	
FI	Full Information. All n advertisers are informed.
IA	Information Asymmetry. n_1 advertisers are informed and n_2 are constrained ($n_1, n_2 > 0$, $n = n_1 + n_2$).
PI	Partial Information. All n advertisers are constrained.
b-Value Dependence Settings	
CV	Common Value. The b-value is the same for all advertisers.
IV	Independent Values. The advertisers' b-values are independent.
Market Metrics (for a b-value dependence setting $\sigma \in \{\mathrm{CV}, \mathrm{IV}\}$ and an information setting $\tau \in \{\mathrm{FI}, \mathrm{IA}, \mathrm{PI}\}$)	
V_τ^σ	Expected conversion rate (the expected v_i of the winning advertiser).
W_τ^σ	Publisher's expected revenue.
E_τ^σ	Informed advertiser's expected payoff.
D_τ^σ	Constrained advertiser's expected payoff.
Parameters	
κ	The probability a user visits for Reason 1.
p	b-probability. The probability that the b-value b_i is high for an advertiser.
G	c-distribution. The CDF of the distribution of the c-value c_i.
Others	
$v_i = \kappa b_i + (1-\kappa) c_i$	Advertiser i's valuation (equivalent to the advertiser's expected conversion probability).
$\beta(c_i)$	Equilibrium bidding function of a constrained advertiser under the common-value IA setting, where the constrained advertiser does not know b_i but knows c_i.

Table 1: Summary of Notation

valuation, but the constrained advertiser might not always do that. We start off with Lemma 1 on the bidding function of the constrained advertiser.

Lemma 1 (Advertisers' bidding behavior)**.** *Under the common-value* IA *setting, the informed advertiser bids their true valuation, while the constrained advertiser with a c-value* $c \in [0,1]$ *bids:*

$$\beta(c) := \begin{cases} (1-\kappa)c, & \text{if } 0 \le c < \min\left\{\frac{\sqrt{1-p}}{\sqrt{1-p}+\sqrt{p}}, \frac{\sqrt{1-p}\kappa}{1-\kappa}\right\}, \\ \kappa p + (1-\kappa)c, & \text{if } \min\left\{\frac{\sqrt{1-p}}{\sqrt{1-p}+\sqrt{p}}, \frac{\sqrt{1-p}\kappa}{1-\kappa}\right\} \le c < \max\left\{\frac{\sqrt{1-p}}{\sqrt{1-p}+\sqrt{p}}, 1-\frac{\sqrt{p}\kappa}{1-\kappa}\right\}, \\ \kappa + (1-\kappa)c, & \text{if } \max\left\{\frac{\sqrt{1-p}}{\sqrt{1-p}+\sqrt{p}}, 1-\frac{\sqrt{p}\kappa}{1-\kappa}\right\} \le c \le 1. \end{cases} \tag{1}$$

The intuition behind Lemma 1 is the following. If the c-value c of the constrained advertiser is relatively low, they bid as if the common b-value b is 0. This is because if they assume some other value $b = x > 0$, they risk overpaying for the impression in the case where $b = 0$ and $(1-\kappa)\cdot c < (1-\kappa)\cdot c' < \kappa \cdot x + (1-\kappa)\cdot c$ (where c' is the c-value of the informed advertiser), where they end up with a negative payoff of $(1-\kappa)\cdot(c-c')$. When $c < \min\left\{\frac{\sqrt{1-p}}{\sqrt{1-p}+\sqrt{p}}, \frac{\sqrt{1-p}\kappa}{1-\kappa}\right\}$, this risk is too high to take. However, when the c-value c is high $\left(c \ge \max\left\{\frac{\sqrt{1-p}}{\sqrt{1-p}+\sqrt{p}}, 1-\frac{\sqrt{p}\kappa}{1-\kappa}\right\}\right)$, it is very likely that $c > c'$, therefore they are not afraid to bid as if $b = 1$, because they have a higher incentive to win and avoid losing impressions with good b-values. For medium values of c, both the risks of overpaying for a bad impression and losing a good impression are too high to make any assumption about b, therefore the advertiser simply bids their expected valuation (note that the expected value of b is p).

Note that as κ increases, i.e. as the b-value becomes more important, the middle interval of c where the constrained advertiser bids their expected valuation shrinks, and for $\kappa \ge 1/2$ it disappears, i.e. the advertiser either underbids or overbids depending on c (see also Figure 1 where the bidding function is shown for different values of κ). Since the role of information asymmetry is more important for larger values of κ and it is where the more interesting results occur, for some of the results we will focus on the case where $\kappa \ge 1/2$.

Note also that when the value of p is low, the region of c where the underbidding occurs is wider compared to the overbidding region, but the amount of underbidding (κp) is smaller compared to the amount of overbidding ($\kappa(1-p)$). On the other hand, when p is high, overbidding is more common but the amount of overbidding is lower. This is illustrated in Figure 2.

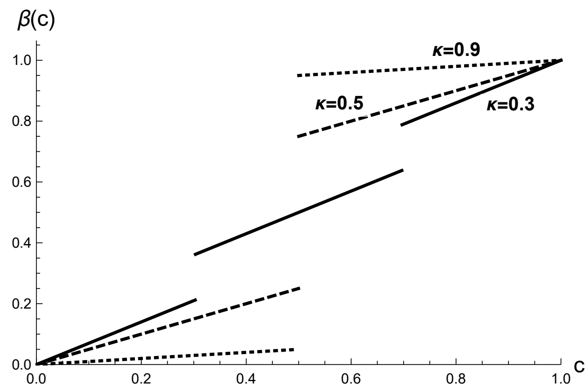

Figure 1: Bidding function of the constrained advertiser for different values of κ (solid line for $\kappa = 0.3$, dashed line for $\kappa = 0.5$, and dotted line for $\kappa = 0.9$), $n_1 = n_2 = 1$, $p = 1/2$, and $G(x) = x$. Notice that for large c-values c, as κ increases there is more overbidding, while for small c-values c, as κ increases there is more underbidding.

In Lemma 2 of Section 5 we show a generalization of Lemma 1 for any $n_1 \geq 1$, any distribution G, $p \in [0, 1]$, and $\kappa \geq 1/2$. Lemma 3 in Section 6 is a further generalization for the more general case with $n_2 \geq 1$ (where the bidding function does not always have a closed-form expression). The same intuition as for Lemma 1 applies to Lemma 2 and Lemma 3 as well.

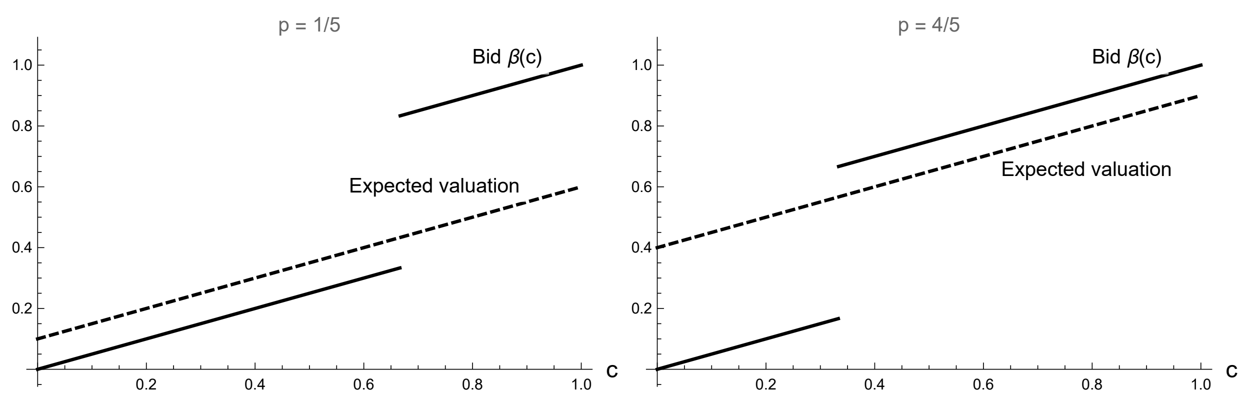

(a) $p = 1/5$. (b) $p = 4/5$.

Figure 2: Bidding function of the constrained advertiser (solid line) compared to their expected valuation (dashed line), for different values of p, $n_1 = n_2 = 1$, $\kappa = 1/2$, and $G(x) = x$. Notice that for small values of p (left) the region of overbidding is smaller than the region of underbidding, but the amount of overbidding ($\kappa(1-p)$) is larger than the amount of underbidding (κp). For large values of p (right) the opposite happens.

The bidding function of Lemma 1 sometimes results in an inefficient market under the IA information setting. More specifically, both the underbidding and the overbidding can result in lower conversion rate compared to the PI setting (where every advertiser bids their expected valuation).

This is illustrated in Example 1.

Example 1 (Inefficiency of non-truthful bidding)**.** Let $\kappa = p = 1/2$. Then the constrained advertiser bids $c/2$ if $c < 1/2$ and $(1 + c)/2$ if $c \geq 1/2$, where c is their c-value. The following two examples illustrate the inefficiency caused by the non-truthful bidding of the constrained advertiser. They show that both underbidding and overbidding can result in lower conversion rates.

- **Inefficiency of underbidding**

  (IA setting) Suppose that the common b-value is high, i.e. $b = 1$, the informed advertiser has c-value $c_1 = 1/6$, and the constrained advertiser has c-value $c_2 = 1/3$. The actual valuations of the two advertisers are $v_1 = (1 + c_1)/2 = 7/12$ for the informed advertiser and $v_2 = (1 + c_2)/2 = 8/12$ for the constrained advertiser. The informed advertiser bids their actual valuation $\beta_1 = 7/12$, but the constrained advertiser underbids, i.e. $\beta_2 = c_2/2 = 2/12$. As a result, the constrained advertiser loses the auction even though they have a higher valuation. Thus, the conversion rate ends up being lower than what it could be in a more efficient auction.

  (PI setting) If none of the advertisers knew the b-value b, then both advertisers would bid their expected valuations, i.e. $\beta_1 = 1/4 + c_1/2 = 4/12$ and $\beta_2 = 1/4 + c_2/2 = 5/12$. Thus, the advertiser with the highest valuation would win, leading to a higher conversion rate.

- **Inefficiency of overbidding**

  (IA setting) Suppose that the common b-value is low, i.e. $b = 0$, the informed advertiser has c-value $c_1 = 5/6$, and the constrained advertiser has c-value $c_2 = 2/3$. The actual valuations of the two advertisers are $v_1 = c_1/2 = 5/12$ for the informed advertiser and $v_2 = c_2/2 = 4/12$ for the constrained advertiser. The informed advertiser bids their actual valuation $\beta_1 = 5/12$, but the constrained advertiser overbids, i.e. $\beta_2 = (1 + c_2)/2 = 10/12$. As a result, the constrained advertiser wins the auction even though they have a lower valuation. Thus, the conversion rate ends up being lower than what it could be in a more efficient auction.

  (PI setting) If none of the advertisers knew the b-value b, then both advertisers would bid their expected valuations, i.e. $\beta_1 = 1/4 + c_1/2 = 8/12$ and $\beta_2 = 1/4 + c_2/2 = 7/12$. Thus, the advertiser with the highest valuation would win, leading to a higher conversion rate.

As we can see in Example 1, there are cases where under the IA setting the advertiser with the highest valuation does not win, either due to the underbidding or due to the overbidding of the constrained advertiser. In contrast, under the PI setting, the highest-valuation advertiser always wins, because every bidder bids their expected valuation, and the winner is determined based on the c-values. This results in a higher conversion rate for the PI setting, as shown in Proposition 1.

For the publisher's revenue, things are less clear. On the one hand, the underbidding that occurs under IA can hurt the publisher, but on the other hand, the overbidding can benefit the publisher because it can increase the prices. Surprisingly, the opposite can happen too; underbidding can sometimes increase publisher's revenue, and overbidding can decrease it, as illustrated in Example 2.

Example 2 (The effects of non-truthful bidding on revenue)**.** Let $\kappa = 1/2$ and $p = 1/3$. Then the constrained advertiser bids $c/2$ if $c < 2 - \sqrt{2}$ and $(1+c)/2$ if $c \geq 2 - \sqrt{2}$, where c is their c-value. The following two examples illustrate that, counter-intuitively, underbidding can sometimes increase publisher's revenue, and overbidding can sometimes decrease it.

- **Underbidding can increase publisher's revenue**

  (IA setting) Suppose that the common b-value is high, i.e. $b = 1$, the informed advertiser has c-value $c_1 = 1/12$, and the constrained advertiser has c-value $c_2 = 1/2$. The actual valuations of the two advertisers are $v_1 = (1 + c_1)/2 = 13/24$ for the informed advertiser and $v_2 = (1 + c_2)/2 = 18/24$ for the constrained advertiser. The informed advertiser bids their actual valuation $\beta_1 = 13/24$, but the constrained advertiser underbids, i.e. $\beta_2 = c_2/2 = 6/24$. The informed advertiser wins and pays β_2, therefore, the publisher's revenue is $6/24$.

  (PI setting) If none of the advertisers knew the b-value b, then both advertisers would bid their expected valuations, i.e. $\beta_1 = 1/6 + c_1/2 = 5/24$ and $\beta_2 = 1/6 + c_2/2 = 10/24$. Then the constrained advertiser would win and pay β_1. Therefore, publisher's revenue would be $5/24$, which is lower than the revenue under the IA setting.

- **Overbidding can decrease publisher's revenue**

  (IA setting) Suppose that the common b-value is low, i.e. $b = 0$, the informed advertiser has c-value $c_1 = 5/6$, and the constrained advertiser has c-value $c_2 = 3/4$. The actual valuations of the two advertisers are $v_1 = c_1/2 = 10/24$ for the informed advertiser and $v_2 = c_2/2 = 9/24$ for the constrained advertiser. The informed advertiser bids their actual valuation $\beta_1 = 10/24$, but the constrained advertiser overbids, i.e. $\beta_2 = (1 + c_2)/2 = 21/24$. The constrained advertiser wins and pays β_1, therefore, publisher's revenue is $10/24$.

(PI setting) If none of the advertisers knew the b-value b, then both advertisers would bid their expected valuations, i.e. $\beta_1 = 1/6 + c_1/2 = 14/24$ and $\beta_2 = 1/6 + c_2/2 = 13/24$. Then the informed advertiser would win and pay β_2. Therefore, publisher's revenue would be $13/24$, which is higher than the revenue under the IA setting.

Despite valuation instances like those in Example 2, in Proposition 1 we show that the overall publisher's expected revenue is higher under the PI setting.

Proposition 1 (Common-value and information asymmetry)**.** *Under the common-b-value setting, the publisher can improve both the conversion rate and the expected revenue if it hides the b-data from all the advertisers. In other words, we have* $V_{\text{IA}}^{\text{CV}} \leq V_{\text{PI}}^{\text{CV}}$ *and* $W_{\text{IA}}^{\text{CV}} \leq W_{\text{PI}}^{\text{CV}}$*.*

Proposition 1 shows that if the publisher has some useful information about a consumer but cannot provide this information to all advertisers, it can achieve a higher conversion rate by hiding the information from everyone rather than giving it only to some advertisers. As an added benefit, the publisher can also simultaneously increase its revenue by hiding this information for all advertisers. The main reason this happens is the inefficiency of the non-truthful bidding of the constrained advertiser under the IA setting, as illustrated in Example 1.

Given the result of Proposition 1, one may wonder if the same can happen when there is no information asymmetry between the advertisers. In other words, if all advertisers have access to the same information, is it still possible that less information can simultaneously increase the conversion rate and the publisher's revenue? In Proposition 2 we show that this cannot happen under the common-b-value setting.

Proposition 2 (Common-value and full information)**.** *Under the common-b-value setting, both the conversion rate and the expected revenue remain unchanged when all advertisers have access to the same information (i.e., when all advertisers are informed or all advertisers are constrained). In other words, it holds that* $V_{\text{FI}}^{\text{CV}} = V_{\text{PI}}^{\text{CV}}$ *and* $W_{\text{FI}}^{\text{CV}} = W_{\text{PI}}^{\text{CV}}$*.*

The equality $V_{\text{FI}}^{\text{CV}} = V_{\text{PI}}^{\text{CV}}$ is relatively easy to see, whereas the equality $W_{\text{FI}}^{\text{CV}} = W_{\text{PI}}^{\text{CV}}$ is less straightforward. Under the common-value setting, since all advertisers have the same b-value, when all have the same information, the b-part of their bids is the same for everyone; therefore, the winner of the auction is purely determined by their c-values in both the FI and the PI settings. As a result, the conversion rate remains unchanged.

For the revenue, when the common b-value is high, publisher's revenue is higher under the PI setting because every advertiser bids above their actual valuation. In contrast, when the common b-value is low, the publisher's revenue is lower under the PI setting because every advertiser bids below their actual valuation. Due to the linearity of the expectation, the average revenue remains the same in the two information settings.

Note that as we move from the FI to the IA and then to the PI information setting, the overall information to the advertisers is reduced. As a result, the inequality $W_{\mathrm{FI}}^{\mathrm{CV}} \geq W_{\mathrm{IA}}^{\mathrm{CV}}$ agrees with the linkage principle (Milgrom and Weber, 1982) which would suggest that revealing information is better for the revenue, but the inequality $W_{\mathrm{IA}}^{\mathrm{CV}} \leq W_{\mathrm{PI}}^{\mathrm{CV}}$ violates the principle which happens due to the information asymmetry.[8]

In this section, we have seen that under the common-value setting there is a non-monotonic relationship between the amount of information available to the advertisers and the efficiency of the auction; as we reduce the information, efficiency (i.e. the conversion rate) first goes down and then goes up again. We have also seen that a similar effect occurs for the publisher's revenue. In Section 4.2 we show that this is no longer true when the b-values are independent.

4.2 Independent-values case

In contrast to the common-value case, when the b-values of advertisers are independent, all advertisers will bid truthfully according to their (expected) valuation.

In the independent-values case, the intuitive result that less information to the advertisers decreases the conversion rate now holds. This is still not true for every valuation instance, as illustrated in Example 3, but it is true for the expected conversion rates, as shown in Proposition 3.

Example 3. Let $\kappa = p = 1/2$.

(IA setting) Suppose that the informed advertiser has a b-value $b_1 = 1$ and a c-value $c_1 = 3/8$, and the constrained advertiser has a b-value $b_2 = 1$ and a c-value $c_2 = 5/8$. The actual valuations

[8] For some other cases where the principle is violated for different reasons, see e.g. Perry and Reny (1999); Fang and Parreiras (2003); Krishna (2009); Despotakis et al. (2017).

of the two advertisers are $v_1 = (1+c_1)/2 = 11/16$ for the informed advertiser and $v_2 = (1+c_2)/2 = 13/16$ for the constrained advertiser. The informed advertiser bids their actual valuation $\beta_1 = 11/16$, but the constrained advertiser bids their expected valuation, i.e. $\beta_2 = 1/4 + c_2/2 = 9/16$. As a result, the constrained advertiser loses the auction even though they have higher valuation.

(PI setting) If none of the advertisers knew their b-values b_i, then both advertisers would bid their expected valuations, i.e. $\beta_1 = 1/4 + c_1/2 = 7/16$ and $\beta_2 = 1/4 + c_2/2 = 9/16$. Thus, the advertiser with the highest valuation would win, leading to a higher conversion rate than the IA setting.

Despite valuation instances like those in Example 3, in Proposition 3 we show that the overall expected conversion rate increases with more information, under the independent-values setting.

Proposition 3 (Independent-values, conversion rates)**.** *Under the independent-b-values setting, the less information advertisers have overall, the lower the conversion rate is. More specifically,* $V_{\text{FI}}^{\text{IV}} \geq V_{\text{IA}}^{\text{IV}} \geq V_{\text{PI}}^{\text{IV}}$.

Proposition 3 shows that the dependence between the b-values of different advertisers is an essential element for the result of Proposition 1, since for independent values it no longer holds.

With respect to the publisher's revenue, the result is less intuitive. Proposition 4 shows that as we provide more information in general to advertisers, publisher revenue decreases.

Proposition 4 (Independent-values, publisher's revenues)**.** *Under the independent-b-values setting, the less information advertisers have overall, the higher publisher's revenue is. More specifically, we have* $W_{\text{FI}}^{\text{IV}} \leq W_{\text{IA}}^{\text{IV}} \leq W_{\text{PI}}^{\text{IV}}$.[9]

The result of Proposition 4 is sensitive to the number of advertisers (in contrast to the previous results that hold for arbitrary number of advertisers; see Section 5). What happens in general is that, for a small number of advertisers, less information is better, but for a large number of advertisers, more information is better. This is due to a version of the market-thinning effect (Levin and Milgrom, 2010). When there are few advertisers in the market, as they become more informed their values spread out, and there is less competition in the high valuations. But as

[9]We want to highlight that this result holds for a low number of advertisers (e.g. two, like in the main model), but unlike the other results it does not always generalize for more advertisers. In Section 5.2 we consider the general case and discuss the details on this.

the number of advertisers becomes larger and the competition increases, more information should improve publisher's revenue. More specifically, when $n_1 = n_2 = 1$ (i.e. there is one advertiser of each type) it holds that more information decreases revenue, but as n_1 and n_2 increase, at some point this stops being true. The exact threshold for the number of advertisers where monotonicity changes depends on the value of p, with a lower p increasing the threshold, the weight κ, with higher κ increasing the threshold, and the c-distribution G (see Proposition 9 and Figure 7 for more details).[10]

4.3 Comparison of the b-Value Settings and Advertisers' Payoffs

Now that we have the results for the simple model with two advertisers, we can compare the two b-value settings (the common-value and the independent-values cases) to each other in terms of their consequences for the publisher's revenue, conversion rates, and advertisers' payoffs.

We start with the publisher's revenue in Figure 3. In the two plots of Figure 3, we see the revenue under the three different information settings as the b-probability p changes in $[0, 1]$. In the common-value case in Figure 3(a), we can see that starting from the IA setting and eliminating the information asymmetry by going towards FI or PI, the publisher's revenue increases. This is due to the underbidding and overbidding behavior that occurs in IA, as discussed in Section 4.1. In contrast to Figure 3(a), in the independent-values case in Figure 3(b) we observe a monotonic change in revenue. As we add information to the market (moving from PI to IA and then to FI), publisher's revenue goes down, as described in Proposition 4.

Next, we move to the conversion rates in Figure 4. What we observe in Figure 4(a) is one of our main findings. What happens here is that in the IA setting, according to Lemma 1, a constrained advertiser with high valuation often bids conservatively and loses to an informed advertiser with a lower valuation. In addition, a constrained advertiser with low valuation often bids aggressively and wins against an informed advertiser with a higher valuation. Both of these bidding behaviors create an inefficient auction because an advertiser with lower valuation wins the consumer's impression, resulting in a lower conversion rate compared to the settings without information asymmetry.

Often in the literature, we see that in markets with thin competition when the publisher's rev-

[10] It is interesting to note that there are also cases where the expected revenue is non-monotonic with respect to the total amount of information that is available to the advertisers. In other words, all six different orderings of $W^{\mathrm{IV}}_{\mathrm{FI}}$, $W^{\mathrm{IV}}_{\mathrm{IA}}$, and $W^{\mathrm{IV}}_{\mathrm{PI}}$ are possible under different conditions (see Figure 8).

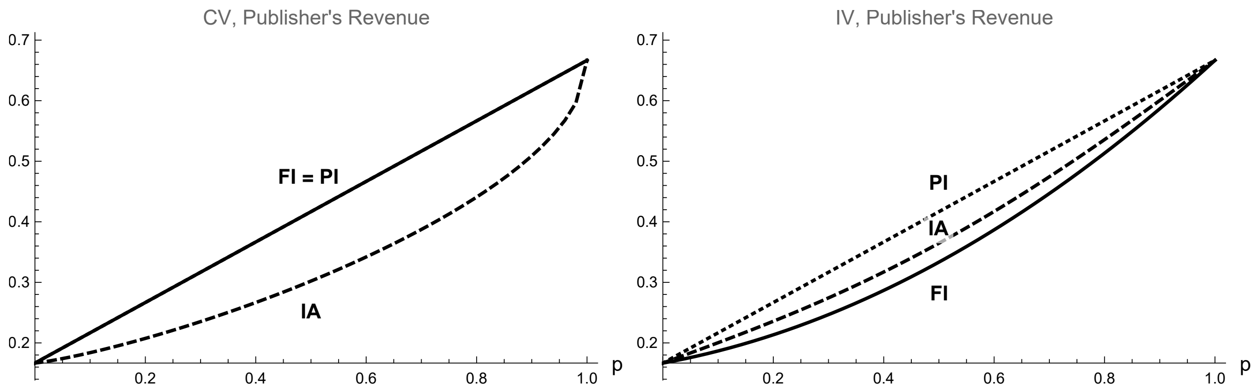

(a) Common-value case. (b) Independent-values case.

Figure 3: Publisher's revenue under the different information settings for different values of $p \in [0,1]$, $n_1 = n_2 = 1$, $\kappa = 1/2$, and $G(x) = x$.

enue goes down (Figure 3(b)), conversion rate (Figure 4(b)) and the advertisers' payoffs (Figure 6) go up as we add information to the market (moving from PI to IA and then to FI). Here, we verify this for our model. However, Proposition 1 states that this is not the case when the b-values are correlated. In fact, both publisher revenue and conversion rate can move in the same direction, as we observe in Figures 3(a) and 4(a) in contrast to Figures 3(b) and 4(b).

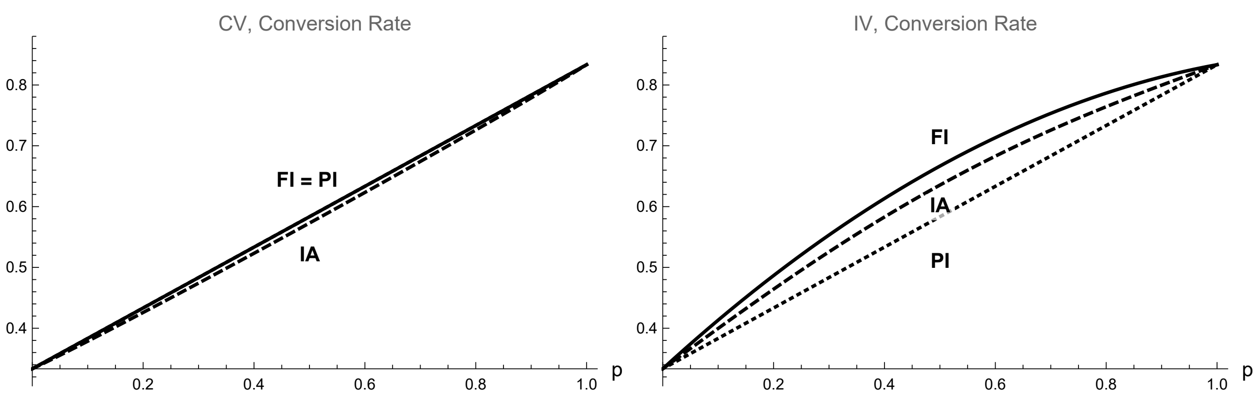

(a) Common-value case. (b) Independent-values case.

Figure 4: Conversion rate under the different information settings as a function of $p \in [0,1]$, for $n_1 = n_2 = 1$, $\kappa = 1/2$, and $G(x) = x$.

Regarding advertisers' payoffs, in Figure 5 we can see that in the common-value case they change non-monotonically both in terms of the information that is available to the advertisers and in terms of p. First, in Figure 5(a), we see that the constrained advertiser's payoff decreases slightly in the IA setting compared to the FI and PI settings, while in Figure 5(b) we see that the informed advertiser's payoff increases significantly. This is expected because the informed advertiser has a strong competitive advantage in the IA setting, while in the FI and PI settings both advertisers are similar. Second, in terms of p, we see that in the IA setting, the constrained advertiser's payoff is minimum for $p = 1/2$ where the uncertainty about the common b-value is maximized. However, the informed advertiser's payoff is maximized for a value $p > 1/2$, which gives the informed advertiser a higher probability of a high valuation in addition to the advantageous uncertainty of the constrained advertiser.

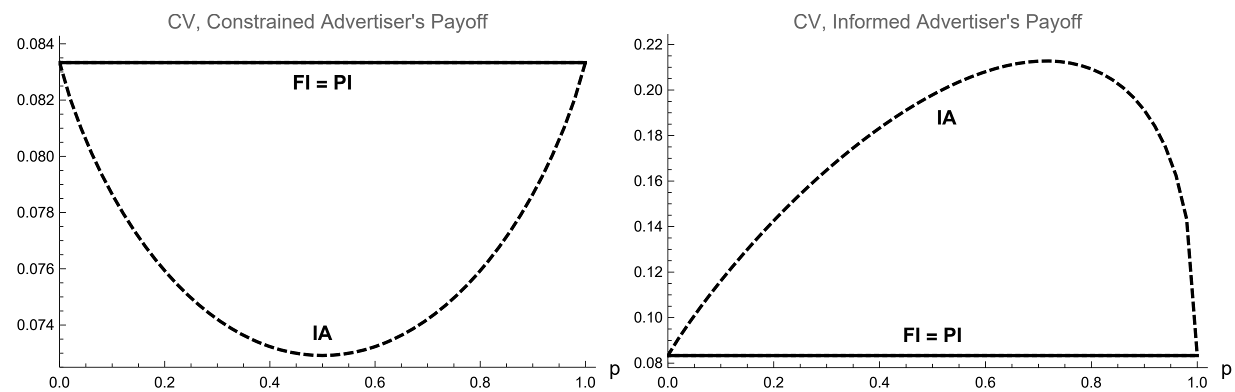

(a) Constrained advertiser. (b) Informed advertiser.

Figure 5: Advertisers' payoffs in the common-value case under the different information settings for different values of $p \in [0, 1]$, $n_1 = n_2 = 1$, $\kappa = 1/2$, and $G(x) = x$.

In contrast to Figure 5, in the independent-values case in Figure 6 we see that both payoffs go down monotonically as we remove information from the market. Furthermore, we see that both types of advertisers have identical payoffs in all settings under IV, including the IA setting where the informed advertiser would normally be expected to have an advantage. Although the constrained advertiser has more fluctuations in their payoff in IA under different realizations of the valuations, their average payoff is the same as the informed advertiser's one, because the advantage of extra information is not that big when the valuations are independent. This perhaps surprising result is independent of any distributional assumptions, but it is a consequence of the fact that there are only two advertisers in the simple version of the model. In Section 6 we discuss the more general case, which is more intuitive in the sense that the informed advertiser has a higher payoff under the IA setting, but still interesting in terms of how the payoff changes as a function of p (see Figure 10).

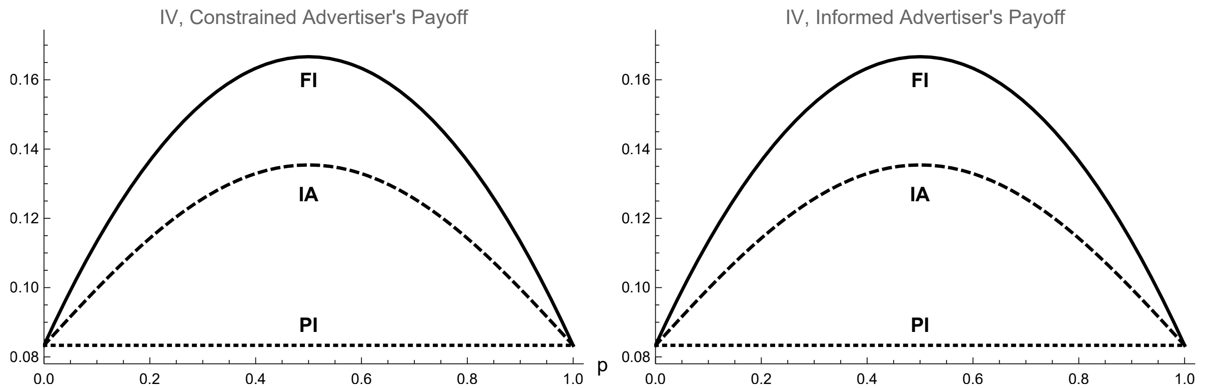

(a) Constrained advertiser.

(b) Informed advertiser.

Figure 6: Advertisers' payoffs in the independent-values case under the different information settings for different values of $p \in [0, 1]$, $n_1 = n_2 = 1$, $\kappa = 1/2$, and $G(x) = x$.

5 Generalizations

In this section, we consider the more general version of the main model, for an arbitrary c-distribution G (instead of uniform) and more than two advertisers. More specifically, there are $n \geq 2$ advertisers competing for the impression, a subset of $n_1 \leq n$ of them are informed advertisers, and the remaining $n_2 = n - n_1$ are constrained advertisers. For the common-value case, the results shown in Section 4.1 extend to the more general setting; this is discussed in subsection 5.1. In the independent-values case there are some interesting differences when we increase the number of advertisers, which we discuss in subsection 5.2.

5.1 Common-value case

Lemma 2 is an analog of Lemma 1 for arbitrary distributions G and more than one informed advertisers.

Lemma 2 (Advertisers' bidding behavior)**.** *For any distribution G, any $n_1 \geq 1$, $n_2 = 1$, and $\kappa \geq 1/2$, under the common-value* IA *setting, all the informed advertisers bid their true valuations while there exists $\underline{c}(p) \in [0, 1]$ such that the constrained advertiser's bidding function is*

$$\beta(c) := \begin{cases} (1-\kappa)c, & \text{if } 0 \leq c < \underline{c}(p), \\ \kappa + (1-\kappa)c, & \text{if } \underline{c}(p) \leq c < 1. \end{cases}$$

Moreover, $\underline{c}$ is independent of κ, and it is a continuously differentiable decreasing function in p, with $\underline{c}(0) = 1$, $\underline{c}(1/2) = n_1 \mathbb{E}[c \cdot G(c)^{n_1 - 1}]$, and $\underline{c}(1) = 0$.

Like in Lemma 1, we see that the constrained advertiser sometimes underbids, for low values of c, and sometimes overbids, for high values of c. Also, as p increases, they overbid more than they underbid. Due to this non-truthful bidding, similar results to those in Section 4.1 continue to hold for $n_1 > 1$.[11] Propositions 5 and 6 generalize the results of Propositions 1 and 2 for arbitrary distributions G. Proposition 5 is shown here for any number of advertisers $n_1, n_2 \geq 1$. We further check the robustness of Proposition 6 for $n_1, n_2 > 1$ in Section 6.

Proposition 5. *For any distribution G, any $n_1, n_2 \geq 1$, and $\kappa \in [0, 1]$, under the common-value case, we have that $V^{\mathrm{CV}}_{\mathrm{IA}} \leq V^{\mathrm{CV}}_{\mathrm{FI}} = V^{\mathrm{CV}}_{\mathrm{PI}}$.*

Proposition 6. *For any distribution G, $n_1 = n_2 = 1$, and $\kappa \geq 1/2$, under the common-value case, we have that $W^{\mathrm{CV}}_{\mathrm{IA}} \leq W^{\mathrm{CV}}_{\mathrm{FI}} = W^{\mathrm{CV}}_{\mathrm{PI}}$.*

5.2 Independent-values case

Propositions 7 and 8 generalize the results of Propositions 3 and 4. The intuition for Proposition 7 is similar to that in the simple model version (as the amount of information available to advertisers decreases, the efficiency of the auction decreases).

Proposition 7. *For any distribution G, any $n_1, n_2 \geq 1$, and $\kappa \in [0, 1]$, under the independent-values case, we have that $V^{\mathrm{IV}}_{\mathrm{FI}} \geq V^{\mathrm{IV}}_{\mathrm{IA}} \geq V^{\mathrm{IV}}_{\mathrm{PI}}$.*

Proposition 8. *For any distribution G, $n_1 = n_2 = 1$, and $\kappa \in [0, 1]$, under the independent-values case, we have that $W^{\mathrm{IV}}_{\mathrm{FI}} \leq W^{\mathrm{IV}}_{\mathrm{IA}} \leq W^{\mathrm{IV}}_{\mathrm{PI}}$.*

In contrast to the common-value setting, under independent b-values, publisher revenue behaves somewhat differently in general (for $n \geq 2$) than what we showed in Propositions 4 and 8. Proposition 9 describes the general phenomenon.

Proposition 9. *For any distribution G and $\kappa \geq 1/2$, under the independent-values case, we have $W^{\mathrm{IV}}_{\mathrm{FI}} \leq W^{\mathrm{IV}}_{\mathrm{PI}}$ for sufficiently small n, and $W^{\mathrm{IV}}_{\mathrm{FI}} > W^{\mathrm{IV}}_{\mathrm{PI}}$ for sufficiently large n. The threshold for n where the inequality is reversed depends on p, κ, and G.*

[11] Lemma 3 in Section 6 generalizes this result for $n_2 > 1$ as well.

As described in Section 4.2, the market-thinning effect that occurs under the IV setting makes hiding information from advertisers beneficial for the publisher's revenue when the number of advertisers is low. However, when there is a sufficiently large number of advertisers, revealing more information increases revenue.

The threshold for the number of advertisers n where the inequality in Proposition 9 reverses depends on the parameters p and κ, and the distribution G. Figure 7 illustrates this. In Figure 7(a) we see that as p decreases and as κ increases, we need more and more advertisers to make the full-information setting give higher revenue than the partial-information setting (i.e. $W_{\mathrm{FI}}^{\mathrm{IV}} \geq W_{\mathrm{PI}}^{\mathrm{IV}}$). In Figure 7(b) we see the thresholds for some examples of different Beta distributions for various parameters α and β. Generally, we observe that c-distributions G with higher average have higher threshold. Also, if the average is low, a lower variance gives a higher threshold, but when the variance is high, a higher variance gives a higher threshold.

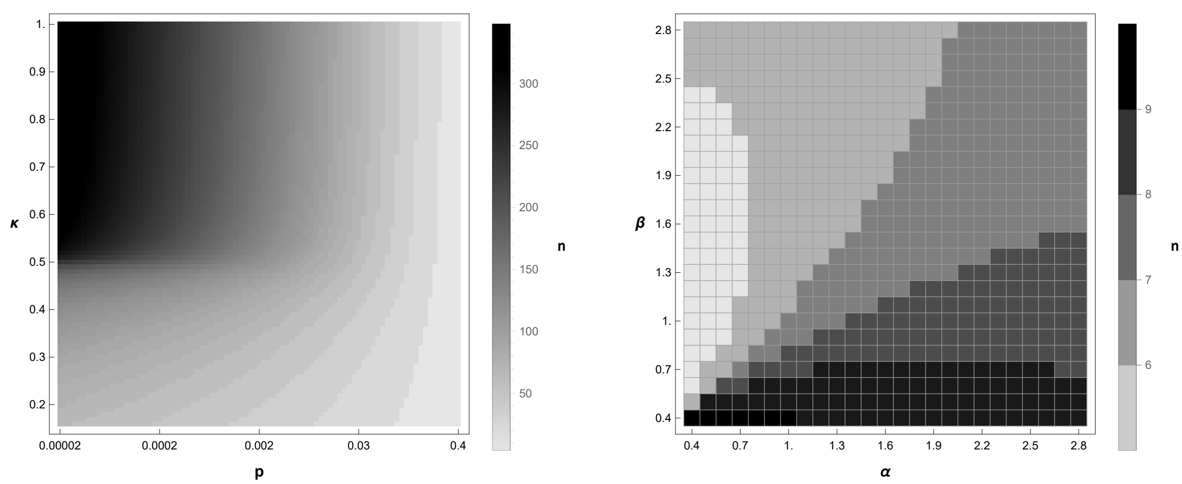

(a) $G(x) = x$, various p and κ.

(b) $G \sim \textsc{Beta}(\alpha, \beta)$, $p = 1/20$, $\kappa = 1/2$.

Figure 7: Minimum number of advertisers n such that $W_{\mathrm{FI}}^{\mathrm{IV}} \geq W_{\mathrm{PI}}^{\mathrm{IV}}$ for various values of p, κ, and different c-distributions G.

Given Proposition 9, a natural question to ask is how $W_{\mathrm{IA}}^{\mathrm{IV}}$ (the revenue under independent values with information asymmetry) compares to $W_{\mathrm{FI}}^{\mathrm{IV}}$ and $W_{\mathrm{PI}}^{\mathrm{IV}}$ under different conditions. The intuition behind Proposition 9 potentially suggests that $W_{\mathrm{FI}}^{\mathrm{IV}} \leq W_{\mathrm{IA}}^{\mathrm{IV}} \leq W_{\mathrm{PI}}^{\mathrm{IV}}$ for low n and $W_{\mathrm{FI}}^{\mathrm{IV}} \geq W_{\mathrm{IA}}^{\mathrm{IV}} \geq W_{\mathrm{PI}}^{\mathrm{IV}}$ for high n. Surprisingly, this is not always the case. In fact, as illustrated in Figure 8, all six different orderings between the revenues $W_{\mathrm{FI}}^{\mathrm{IV}}$, $W_{\mathrm{IA}}^{\mathrm{IV}}$, and $W_{\mathrm{PI}}^{\mathrm{IV}}$ are possible under different conditions. The information asymmetry between advertisers adds an additional element of complexity that the

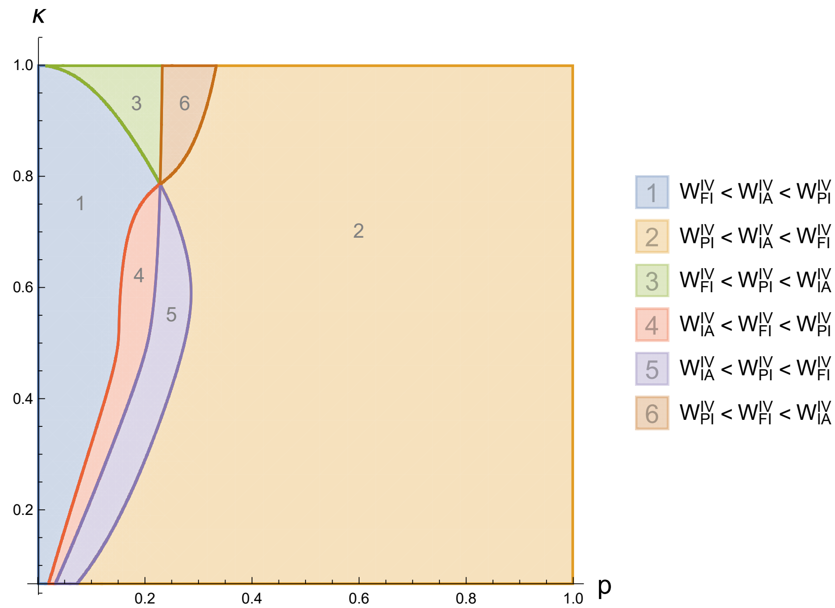

Figure 8: Publisher's revenues comparisons between different information settings under independent b-values, for various values of p and κ, $G(x) = x$, and $n_1 = n_2 = 2$.

market-thinning effect alone is not sufficient to explain.

The intuition behind Figure 8 is as follows. In Proposition 9 we saw that a low n makes hiding information from advertisers beneficial for the publisher, due to a thinner market. For a similar reason, a low p also makes hiding information beneficial. This is because when p is low, there is a low probability that the second highest bidder at the auction will have a b-value $b_i = 1$, which means that the clearing price will most likely be of the form $(1-\kappa)c_i$ if all advertisers know their b-values. Thus, when p is low, the publisher prefers to hide information from as many advertisers as possible to make them bid their expected valuation $\kappa p + (1-\kappa)c_i$ instead of their actual valuation. On the contrary, when p is high, the publisher prefers to reveal the b-data to as many advertisers as possible so that they can bid their (likely high) actual valuation. In other words, when p is low we have that $W_{\text{FI}}^{\text{IV}} < W_{\text{IA}}^{\text{IV}} < W_{\text{PI}}^{\text{IV}}$ (Region 1) and when p is high we have that $W_{\text{FI}}^{\text{IV}} > W_{\text{IA}}^{\text{IV}} > W_{\text{PI}}^{\text{IV}}$ (Region 2). This also explains why in Regions 1, 3, and 4 it is $W_{\text{FI}}^{\text{IV}} < W_{\text{PI}}^{\text{IV}}$, while in Regions 2, 5, and 6 it is $W_{\text{FI}}^{\text{IV}} > W_{\text{PI}}^{\text{IV}}$.

To understand why the IA setting generates higher revenue than the other two information settings in Regions 3 and 6 where p is medium and κ is high, let us consider the extreme case where $\kappa = 1$. In this extreme case, the valuations of the advertisers are just b_i, without a c-part.

Under the IA setting, the informed advertisers will bid their actual values b_i, while the constrained advertisers will bid their expected valuation which is just p. When p is high, there is a high chance that there will be at least two advertisers with high b_i's, so the publisher wants the advertisers to learn their b_i's to have a high chance of getting a clearing price of 1 (Region 2). When p is low, it is less likely that there will be at least two b_i's that are high, so the publisher prefers if the advertisers bid p instead of their b_i which is more likely 0. However, having many advertisers bidding p has no additional benefit compared to just two advertisers bidding p, since the clearing price will be p in both cases. Therefore, the optimal revenue for the publisher when p is low is achieved when there is information asymmetry, where the publisher guarantees a clearing price of at least p from the constrained advertisers and there is also a (small) chance of something higher from the informed advertisers (Regions 3 and 6).

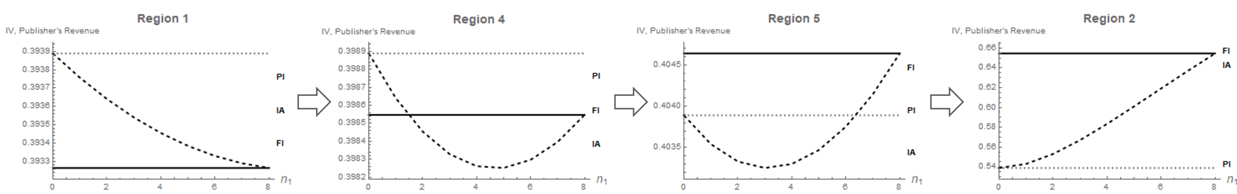

Figure 9: Publisher's revenues for the different information settings under independent b-values, for $n = 8$ advertisers, $n_1 \in [0, n]$, $n_2 = n - n_1$, $\kappa = 1/2$, $G(x) = x$, and $p \in \{0.01, 0.02, 0.03, 0.3\}$ (from left to right).

Finally, when κ is low, the importance of the b-value on the advertisers' valuations is low. The c-part of the valuations dominates in determining the winner. As a result, the benefit of information asymmetry described above, where it is good for the publisher to have both informed and constrained advertisers, is not essential anymore since the c-values are known by both. In Regions 4 and 5, the IA setting has worse revenue than the other two settings because of a third effect.

Under the IA setting, there are two groups of advertisers, n_1 informed advertisers and n_2 constrained advertisers. If we fix the total number of advertisers $n = n_1 + n_2$, then we can think of the FI and the PI information settings as extreme versions of the IA setting. More specifically, FI is like IA with $(n_1, n_2) = (n, 0)$ and PI is like IA with $(n_1, n_2) = (0, n)$. With that view in mind, to understand how $W^{\text{IV}}_{\text{FI}}$, $W^{\text{IV}}_{\text{IA}}$, and $W^{\text{IV}}_{\text{PI}}$ compare to each other, it is useful to look at the function $W^{\text{IV}}_{\text{IA}}(n_1, n_2) = W^{\text{IV}}_{\text{IA}}(n_1, n - n_1)$ as n_1 goes from 0 to n, while everything else is fixed.

In Figure 9 we can see some examples of this function (represented by the dashed line) for four different values of p, starting from a low p in the first plot on the left and increasing it towards the right (we also consider $n = 8$ advertisers to make the effect clearer). For $n_1 = 0$ the function gives the revenue under the PI setting (dotted line) and for $n_1 = n$ it gives the revenue under the FI setting (solid line). We see that for low p this function is decreasing and it gradually becomes increasing as p increases. While it transitions from decreasing to increasing, at some point, for medium values of p it becomes non-monotone (first decreasing and then increasing). This is the point where the IA setting can give lower revenue for the publisher than both the FI and the PI settings (Regions 4 and 5 in Figure 8).

The explanation for this is as follows. As n_1 increases from 0 to n, what we do is we move advertisers one by one from the group of constrained advertisers to the group of informed advertisers. The average of the bids in both groups is the same; therefore, the average bid is not affected as we move advertisers. However, what changes is the variance of the distribution of the bids. More specifically, the bids of the constrained advertisers are more concentrated around the mean, while the bids of the informed advertisers are more spread out. When we move the first few advertisers from the constrained group to the informed group, we make the bid distribution of the constrained group slightly worse. However, the advertiser who determines the clearing price of the overall auction is still more likely in the constrained group, as it has significantly more advertisers. Therefore, what happens is that as we start moving advertisers, we make the clearing price lower. However, after we reach a critical mass of advertisers in the informed group, suddenly the clearing price will more likely be determined by the informed group (i.e. there is a high chance that there will be at least two informed advertisers with high b_i's). From that point onwards, as we make the informed group larger, we increase the expected clearing price. This is the reason for the non-monotonicity of the function in the second and third plots of Figure 9. This transition phase is also what explains the existence of Regions 4 and 5 in Figure 8.

All three effects described above combined generate the six different regions we see in Figure 8.

Advertisers' payoffs in the independent-values case. In Figure 10 we can see the payoffs of each type of advertiser for the three information settings and different values of p in $[0, 1]$ when the b-values of the advertisers are independent. The two plots of Figure 10 describe the more general

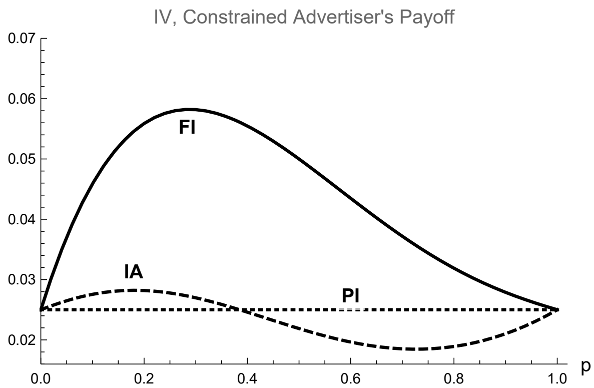

(a) Constrained advertiser.

(b) Informed advertiser.

Figure 10: Advertisers' payoffs in the independent-values case under the different information settings for different values of $p \in [0,1]$, $n_1 = n_2 = 2$, $\kappa = 1/2$, and $G(x) = x$.

behavior of the payoffs when there are more than two advertisers in total (in contrast to Figure 6 which was for one advertiser of each type). There are a few interesting things to note regarding the payoffs. First, in Figure 10(a), we see that for low values of p it is $D^{\mathrm{IV}}_{\mathrm{PI}} \leq D^{\mathrm{IV}}_{\mathrm{IA}} \leq D^{\mathrm{IV}}_{\mathrm{FI}}$, while for high values of p it is $D^{\mathrm{IV}}_{\mathrm{IA}} \leq D^{\mathrm{IV}}_{\mathrm{PI}} \leq D^{\mathrm{IV}}_{\mathrm{FI}}$. In other words, when p is low, a constrained advertiser prefers the asymmetric setting where informed advertisers have more information than them, over the partial-information setting where all advertisers have similar information. Second, in Figure 10(b), we see that for low values of p it is $E^{\mathrm{IV}}_{\mathrm{PI}} \leq E^{\mathrm{IV}}_{\mathrm{IA}} \leq E^{\mathrm{IV}}_{\mathrm{FI}}$, while for high values of p it is $E^{\mathrm{IV}}_{\mathrm{PI}} \leq E^{\mathrm{IV}}_{\mathrm{FI}} \leq E^{\mathrm{IV}}_{\mathrm{IA}}$. In other words, when p is low, an informed advertiser prefers the full-information setting where constrained advertisers have as much information as them, over the asymmetric setting where the informed advertiser has more information than the constrained advertisers. The following result shows that these observations hold more generally.

Proposition 10. *For a uniform distribution G, any $n_1, n_2 \geq 1$, and $\kappa \geq 1/2$, when p is sufficiently low, it holds that $D^{\mathrm{IV}}_{\mathrm{PI}} \leq D^{\mathrm{IV}}_{\mathrm{IA}}$ and $E^{\mathrm{IV}}_{\mathrm{IA}} \leq E^{\mathrm{IV}}_{\mathrm{FI}}$.*

The intuition behind Proposition 10 is the following. As an advertiser, it is often advantageous for you if other advertisers gain more information than they currently have. This is because when an advertiser does not know their actual valuation they bid their expected valuation, but when p is low it is more likely than not that their actual valuation is lower than their expected valuation. In other words, when p is sufficiently low, you want the other advertisers to learn their actual valuations because then it is very likely that they will lower their bids.

6 More Robustness Checks

In this section, we check the robustness of Proposition 6 (which is the remaining result not proven analytically for the case where $n_1, n_2 > 1$, due to the lack of a closed-form general bidding function for the constrained advertisers under the common-value IA setting). We first start by establishing the existence of a pure strategy symmetric equilibrium bidding function for the general case.

Lemma 3 (Advertisers' bidding behavior)**.** *For any strictly increasing and smooth distribution G, any $n_1, n_2 \geq 1$, and $\kappa \geq 1/2$, under the common-value* IA *setting, all informed advertisers bid their true valuations and there exists a pure strategy symmetric equilibrium bidding function β for the constrained advertisers satisfying $\beta(c) \in \{(1-\kappa)c\} \cup [\kappa, \kappa + (1-\kappa)c]$ for $c \in [0, 1]$.*

Lemma 3 is a generalization of Lemmas 1 and 2. Based on Lemma 3, we can numerically approximate the function β for any $n_1, n_2 \geq 1$ by solving the differential equation $\frac{\partial u(\tilde{\beta};\beta,c)}{\partial \tilde{\beta}}\Big|_{\tilde{\beta}=\beta(c)} = 0$, where u is defined in Eq. (4). In Figure 11 we can see one example of the equilibrium bidding function when there are two informed advertisers and two constrained advertisers. Like in Lemma 1, for small c-values c, constrained advertisers underbid, while for large values of c they overbid.

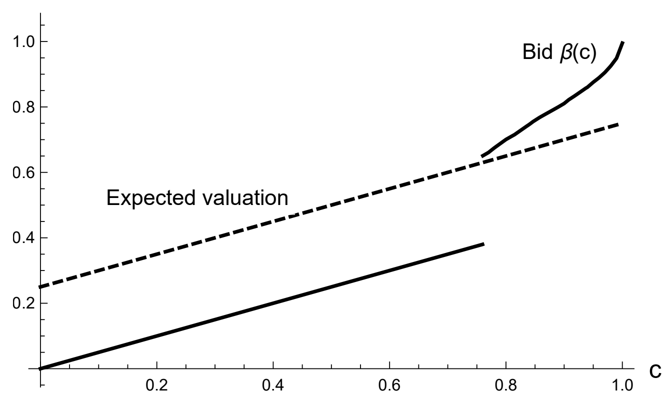

Figure 11: Bidding function of the constrained advertisers (solid line) compared to their expected valuation (dashed line), for $n_1 = n_2 = 2$, $p = 1/2$, $\kappa = 1/2$, and $G(x) = x$.

Despite the lack of a closed-form bidding function, the intuition for the bidding behavior is the same as the one discussed in Section 4.1. As a result, Proposition 6 continues to hold for a large number of advertisers. In Figure 12 we can see a demonstration of this. In Figure 12(a) we consider different values of n, i.e. the total number of advertisers, and assuming that there is an equal number of informed and constrained advertisers, we estimate the bidding function of

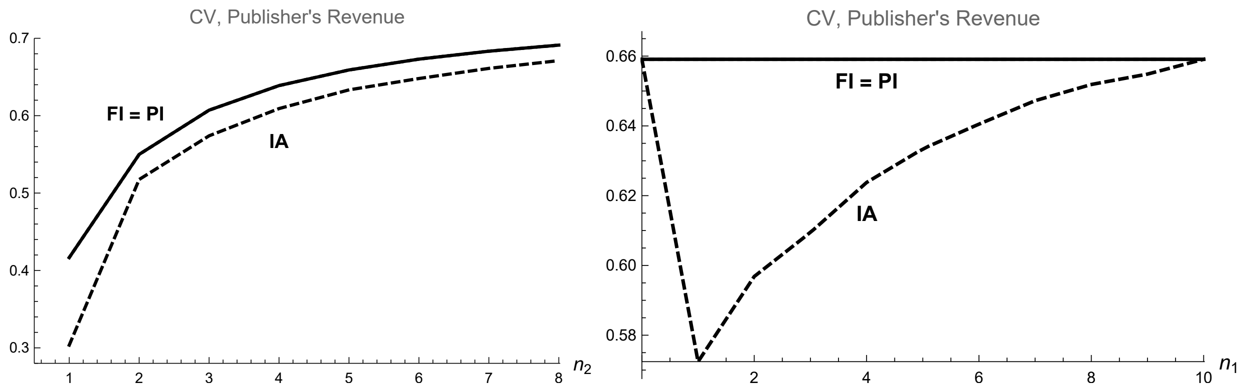

(a) $n \in \{2, 4, 6, \ldots, 16\}$, $n_1 = n_2 = n/2$. (b) $n = 10$, $n_1 \in \{0, 1, \ldots, n\}$, $n_2 = n - n_1$.

Figure 12: Publisher's revenue under the different information settings in the common-value case for different combinations of $n_1, n_2 \geq 1$, $p = 1/2$, $\kappa = 1/2$, and $G(x) = x$.

the constrained advertisers and calculate the publisher's revenue. We see that for all cases it is $W_{\text{IA}}^{\text{CV}} \leq W_{\text{FI}}^{\text{CV}} = W_{\text{PI}}^{\text{CV}}$. In Figure 12(b) we fix the total number of advertisers n and consider all different combinations of n_1 and n_2. As before, we establish that $W_{\text{IA}}^{\text{CV}} \leq W_{\text{FI}}^{\text{CV}} = W_{\text{PI}}^{\text{CV}}$ for all cases. Different choices for the number of advertisers and the other parameters generate similar plots (see also Appendix B.2 for all the key formulas used to generate the plots).

7 Conclusion

This paper introduces and analyzes the concept of state-dependent predictive value in advertising auctions: the idea that the usefulness of observable consumer data for predicting conversion depends on an unobservable consumer state. We develop a game-theoretic model that combines this feature with asymmetric data access across advertisers and study how different information structures affect publisher revenue, conversion rates, and advertiser payoffs under both correlated and independent valuation shifts. Figure 13 provides a concise summary of the main findings.

The central finding is that the interaction between state-dependent data value and information asymmetry produces a distinctive bidding distortion. When the detailed data component shifts all advertisers' valuations in the same direction (the common-value case), constrained advertisers who lack access to this data face a dilemma: they cannot condition their bids on the realization of the common component, leading them to underbid when their baseline value is low and overbid when it is high. This distortion reduces allocative efficiency, causing impressions to be won by

Should the publisher **restrict data access** to maximize **conversion rate** and **auction efficiency**?

	Asymmetric Information	Symmetric Information
Common Value	**Yes** $V_{IA}^{CV} \leq V_{PI}^{CV}$ (Propositions 1, 5)	**No** $V_{FI}^{CV} = V_{PI}^{CV}$ (Propositions 2, 5)
Independent Values	**No** $V_{IA}^{IV} \geq V_{PI}^{IV}$ (Propositions 3, 7)	**No** $V_{FI}^{IV} \geq V_{PI}^{IV}$ (Propositions 3, 7)

Should the publisher **restrict data access** to maximize **revenue**?

	Asymmetric Information	Symmetric Information
Common Value	**Yes** $W_{IA}^{CV} \leq W_{PI}^{CV}$ (Propositions 1, 6)	**No** $W_{FI}^{CV} = W_{PI}^{CV}$ (Propositions 2, 6)
Independent Values	**Yes**, for low p and intermediate κ $W_{IA}^{IV} \leq W_{PI}^{IV}$ **No**, otherwise $W_{IA}^{IV} \geq W_{PI}^{IV}$ (Propositions 4, 8, Figure 8)	**Yes**, for low p and high κ $W_{FI}^{IV} \leq W_{PI}^{IV}$ **No**, otherwise $W_{FI}^{IV} \geq W_{PI}^{IV}$ (Propositions 4, 8, 9, Figures 7, 8)

Figure 13: Summary of main findings.

advertisers who are not the best match for the consumer. Eliminating the asymmetry by restricting data access for all advertisers removes this distortion, simultaneously improving both publisher revenue and conversion rates. This result overturns the common intuition that privacy-enhancing data restrictions necessarily involve a tradeoff between revenue and ad relevance. Importantly, when all advertisers have symmetric access to the same information, whether full or partial, these gains disappear: it is the asymmetry itself, rather than the level of information, that drives the inefficiency.

The picture changes when detailed data affects advertisers' valuations independently. In this case, more information does improve allocative efficiency, consistent with standard intuitions. However, the revenue implications are nuanced and depend on market thickness. In thin markets with few advertisers, restricting information increases revenue through a market-thinning mechanism: informed advertisers' valuations spread out, weakening competition at the top. As the number of advertisers grows, this effect is overwhelmed by the competitive benefits of better-informed bidding. Furthermore, information asymmetry can, under certain parameter configurations, generate higher publisher revenue than either full or partial information, suggesting that selective rather than uniform data-access policies may sometimes be optimal. The fact that all six possible orderings of revenues across the three information settings can arise under different conditions underscores the complexity of the relationship between information structure and market performance in this case.

Our analysis of advertiser payoffs reveals further counterintuitive patterns. In the independent-values case, informed advertisers can prefer the symmetric settings where their informational advantage is eliminated, because the constrained competitors' bidding strategy under asymmetry creates an inefficient competitive environment that harms all participants. Additionally, when the probability of a high data-dependent valuation component is low, constrained advertisers may prefer the asymmetric setting over one where all advertisers are equally uninformed, because giving competitors access to data makes them more likely to discover that their valuations are low and bid accordingly. These findings suggest that advertisers' preferences over data-access regimes cannot be straightforwardly inferred from their information status.

These results carry implications for several stakeholders. For publishers and platforms evaluating whether to restrict data access, such as through cookie deprecation, our findings identify a specific and empirically assessable condition under which restrictions are beneficial: when the data in question shifts advertiser valuations in a correlated manner across competitors. A platform can evaluate this by examining whether signals like browsing history or behavioral data tend to make a given user more or less attractive to most competing advertisers simultaneously, or whether the effects are idiosyncratic to each advertiser. When the correlated-shift condition holds and data access is currently asymmetric, restricting access improves both revenue and the quality of ad-to-consumer matches. For regulators and policymakers, our analysis cautions against evaluating data-restriction policies without accounting for the correlation structure of advertiser valuations and the pre-existing degree of information asymmetry. The welfare consequences of the same policy can be qualitatively different depending on these market fundamentals: a restriction that improves outcomes in a market with correlated valuations may reduce efficiency in one with independent valuations. For advertisers, the analysis highlights that restricting data access to competitors does not always yield a net competitive advantage, particularly in markets where the additional information leads to independent valuation shifts.

The model involves several simplifying assumptions that suggest directions for future research. First, the binary structure of the unobservable consumer state and the Bernoulli distribution for the data-dependent valuation component, while sufficient to generate the key mechanisms, could be extended to richer state spaces and continuous distributions to examine whether the effects we identify are amplified or attenuated. Second, we normalize the per-conversion value to be identical across advertisers, abstracting from heterogeneity in profit margins. In practice, advertisers with higher per-conversion values may invest more heavily in data infrastructure, creating a correlation between information status and willingness to pay that could interact with the mechanisms we study. Third, our model treats the information structure as exogenous to the auction, controlled by the publisher's data-access policy. Endogenizing advertisers' decisions to invest in data acquisition, or modeling a platform's dynamic data-sharing strategy, would enrich the analysis. Fourth, we focus on a single auction for a single impression. Extending the framework to a repeated setting where advertisers learn over time and update their beliefs about the distribution of consumer states would be a natural next step, particularly for understanding how the transition from one information regime to another unfolds in practice. Finally, incorporating consumer welfare more explicitly, for instance by modeling how ad relevance affects user experience and long-run platform engagement, would provide a more complete picture of the welfare implications of data-access policies.

Despite these limitations, the paper establishes a core theoretical insight: in advertising markets characterized by state-dependent data value and asymmetric information access, the conventional wisdom that more data yields better targeting outcomes does not hold in general. An important determinant, other than how much information is available, is how it is distributed across competitors and how its predictive value interacts with unobservable consumer heterogeneity. As the online advertising industry continues to navigate the tension between data-driven targeting and privacy protection, understanding these structural features of the market is essential for designing information policies that serve the interests of publishers, advertisers, and consumers.

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.

Athey, S. and Gans, J. S. (2010). The impact of targeting technology on advertising markets and media competition. *American Economic Review*, 100(2):608–13.

Bergemann, D. and Bonatti, A. (2011). Targeting in advertising markets: implications for offline versus online media. *The RAND Journal of Economics*, 42(3):417–443.

Bobkova, N. (2024). Information choice in auctions. *American Economic Review*, 114(7):1883–1915.

Brahim, N. B. E.-B., Lahmandi-Ayed, R., and Laussel, D. (2011). Is targeted advertising always beneficial? *International Journal of Industrial Organization*, 29(6):678–689.

Chandra, A. and Kaiser, U. (2014). Targeted advertising in magazine markets and the advent of the internet. *Management Science*, 60(7):1829–1843.

Chen, Y., Narasimhan, C., and Zhang, Z. J. (2001). Individual marketing with imperfect targetability. *Marketing Science*, 20(1):23–41.

Choi, H., Mela, C. F., Balseiro, S. R., and Leary, A. (2020). Online display advertising markets: A literature review and future directions. *Information Systems Research*, 31(2):556–575.

Choi, W. J., Jerath, K., and Sarvary, M. (2023). Consumer privacy choices and (un) targeted advertising along the purchase journey. *Journal of Marketing Research*, page 00222437221140052.

Choi, W. J. and Sayedi, A. (2019). Learning in online advertising. *Marketing Science*, 38(4):584–608.

Choi, W. J. and Sayedi, A. (2023a). Agency bidding in online advertising. *Available at SSRN 4324955*.

Choi, W. J. and Sayedi, A. (2023b). 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*.

Davies, J. (2019). After GDPR, The New York Times cut off ad exchanges in Europe and kept growing ad revenue. https://digiday.com/media/gumgumtest-new-york-times-gdpr-cut-off-ad-exchanges-europe-ad-revenue/.

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

Deng, Y. and Mela, C. F. (2018). Tv viewing and advertising targeting. *Journal of Marketing Research*, 55(1):99–118.

Despotakis, S., Hafalir, I., Ravi, R., and Sayedi, A. (2017). Expertise in online markets. *Management Science*, 63(11):3895–3910.

Despotakis, S., Ravi, R., and Sayedi, A. (2021). First-price auctions in online display advertising. *Journal of Marketing Research*, 58(5):888–907.

Despotakis, S. and Yu, J. (2022). Multidimensional targeting and consumer response. *Management Science*, forthcoming.

Dong, X., Manchanda, P., and Chintagunta, P. K. (2009). Quantifying the benefits of individual-level targeting in the presence of firm strategic behavior. *Journal of Marketing Research*, 46(2):207–221.

D'Annunzio, A. and Russo, A. (2024). Intermediaries in the online advertising market. *Marketing Science*, 43(1):33–53.

Edelman, G. (2020). Can killing cookies save journalism? https://www.wired.com/story/can-killing-cookies-save-journalism/.

Esteban, L., Gil, A., and Hernandez, J. M. (2001). Informative advertising and optimal targeting in a monopoly. *The Journal of Industrial Economics*, 49(2):161–180.

Esteves, R.-B. and Resende, J. (2016). Competitive targeted advertising with price discrimination. *Marketing Science*, 35(4):576–587.

Fang, H. and Parreiras, S. O. (2003). On the failure of the linkage principle with financially constrained bidders. *Journal of Economic Theory*, 110(2):374–392.

Gal-Or, E., Gal-Or, M., May, J. H., and Spangler, W. E. (2006). Targeted advertising strategies on television. *Management Science*, 52(5):713–725.

Goldfarb, A. and Tucker, C. (2011). Online display advertising: Targeting and obtrusiveness. *Marketing Science*, 30(3):389–404.

Gordon, B. R., Jerath, K., Katona, Z., Narayanan, S., Shin, J., and Wilbur, K. C. (2021). Inefficiencies in digital advertising markets. *Journal of Marketing*, 85(1):7–25.

Hsiao, S. (2020). How our display buying platforms share revenue with publishers. https://blog.google/products/admanager/display-buying-share-revenue-publishers.

Hummel, P. and McAfee, R. P. (2016). When does improved targeting increase revenue? *ACM Transactions on Economics and Computation (TEAC)*, 5(1):1–29.

Iyer, G., Soberman, D., and Villas-Boas, J. M. (2005). The targeting of advertising. *Marketing Science*, 24(3):461–476.

Johnson, J. P. (2013). Targeted advertising and advertising avoidance. *The RAND Journal of Economics*, 44(1):128–144.

Krishna, V. (2009). *Auction theory*. Academic press.

Lee, D., Hosanagar, K., and Nair, H. S. (2018). Advertising content and consumer engagement on social media: Evidence from facebook. *Management Science*, 64(11):5105–5131.

Levin, J. and Milgrom, P. (2010). Online advertising: Heterogeneity and conflation in market design. *American Economic Review*, 100(2):603–07.

Long, F., Jerath, K., and Sarvary, M. (2022). Designing an online retail marketplace: Leveraging information from sponsored advertising. *Marketing Science*, 41(1):115–138.

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

Perry, M. and Reny, P. J. (1999). On the failure of the linkage principle in multi-unit auctions. *Econometrica*, 67(4):895–900.

Rafieian, O. and Yoganarasimhan, H. (2021). Targeting and privacy in mobile advertising. *Marketing Science*, 40(2):193–218.

Roy, S. (2000). Strategic segmentation of a market. *International Journal of Industrial Organization*, 18(8):1279–1290.

Sayedi, A. (2018). Real-time bidding in online display advertising. *Marketing Science*, 37(4):553–568.

Shen, Q. and Miguel Villas-Boas, J. (2018). Behavior-based advertising. *Management Science*, 64(5):2047–2064.

Shin, J. and Shin, W. (2022). A theory of irrelevant advertising: An agency-induced targeting inefficiency. *Management Science*.

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

Snelders, E., Worp, L., and Sheng, S. (2020). A future without advertising cookies? https://www.ster.nl/media/h5ehvtx3/ster_a-future-without-advertising-cookies.pdf.

Summers, C. A., Smith, R. W., and Reczek, R. W. (2016). An audience of one: Behaviorally targeted ads as implied social labels. *Journal of Consumer Research*, 43(1):156–178.

Trusov, M., Ma, L., and Jamal, Z. (2016). Crumbs of the cookie: User profiling in customer-base analysis and behavioral targeting. *Marketing Science*, 35(3):405–426.

Tucker, C. E. (2014). Social networks, personalized advertising, and privacy controls. *Journal of marketing research*, 51(5):546–562.

Zhang, K. and Katona, Z. (2012). Contextual advertising. *Marketing Science*, 31(6):980–994.

Zia, M. and Rao, R. C. (2019). Search advertising: Budget allocation across search engines. *Marketing Science*, 38(6):1023–1037.

A Appendix

A.1 Proofs of Lemmas 1 and 2 and Propositions 1 to 4

Proof of Lemma 1

The constrained advertiser's expected utility when their c-value is c and they bid β is:

$$u(\beta, c) := p(1-\kappa)\int_0^{\max\left\{\frac{\beta-\kappa}{1-\kappa},0\right\}} (c-c')d(G(c')^{n_1}) + (1-p)(1-\kappa)\int_0^{\min\left\{\frac{\beta}{1-\kappa},1\right\}} (c-c')d(G(c')^{n_1}).$$

Suppose that $\kappa < 1-\kappa$. First, let us consider the case when $0 \le \beta < \kappa$, we have $u(\beta, c) = (1-p)(1-\kappa)\int_0^{\frac{\beta}{1-\kappa}} (c-c')d(G(c')^{n_1})$, which means, $\frac{\partial u}{\partial \beta} = n_1(1-p)\left(c - \frac{\beta}{1-\kappa}\right) G\left(\frac{\beta}{1-\kappa}\right)^{n_1-1} G'\left(\frac{\beta}{1-\kappa}\right) = 0 \implies \beta = (1-\kappa)c$. For, the case when $\kappa < \beta \le 1-\kappa$, we have $u(\beta, c) = p(1-\kappa)\int_0^{\frac{\beta-\kappa}{1-\kappa}} (c-c')d(G(c')^{n_1}) + (1-p)(1-\kappa)\int_0^{\frac{\beta}{1-\kappa}} (c-c')d(G(c')^{n_1})$, which means, $\frac{\partial u}{\partial \beta} = n_1 p\left(c - \frac{\beta-\kappa}{1-\kappa}\right) G\left(\frac{\beta-\kappa}{1-\kappa}\right)^{n_1-1} G'\left(\frac{\beta-\kappa}{1-\kappa}\right) + n_1(1-p)\left(c - \frac{\beta}{1-\kappa}\right) G\left(\frac{\beta}{1-\kappa}\right)^{n_1-1} G'\left(\frac{\beta}{1-\kappa}\right)$. When $n_1 = 1$ and $G(x) = x$, $\frac{\partial u}{\partial \beta} = 0$ implies that $\beta = \kappa p + (1-\kappa)c$. For the case when $1-\kappa < \beta \le 1$, we have $u(\beta, c) = p(1-\kappa)\int_0^{\frac{\beta-\kappa}{1-\kappa}} (c-c')d(G(c')^{n_1}) + (1-p)(1-\kappa)\int_0^1 (c-c')d(G(c')^{n_1})$, which means, $\frac{\partial u}{\partial \beta} = n_1 p\left(c - \frac{\beta-\kappa}{1-\kappa}\right) G\left(\frac{\beta-\kappa}{1-\kappa}\right)^{n_1-1} G'\left(\frac{\beta-\kappa}{1-\kappa}\right) = 0 \implies$

$\beta = \kappa + (1-\kappa)c$. The global maximum of u occurs either at $\beta = (1-\kappa)c, \kappa p + (1-\kappa)c, \kappa + (1-\kappa)c$, or at one of the singular points $\beta = \kappa, 1-\kappa$. Let

$$u_1(c) := u(\beta = (1-\kappa)c, c) = \begin{cases} \frac{1}{2}(1-p)(1-\kappa)c^2, & c \le \frac{\kappa}{1-\kappa}, \\ \frac{1}{2}(1-\kappa)c^2 - \frac{p}{2}\left(\frac{\kappa^2}{1-\kappa}\right), & \frac{\kappa}{1-\kappa} < c \le 1, \end{cases}$$

$$u_2(c) := u(\beta = \kappa p + (1-\kappa)c, c) = \begin{cases} \frac{1}{2}(1-p)(1-\kappa)c^2 - \frac{1}{2}(1-p)p^2\left(\frac{\kappa^2}{1-\kappa}\right), & c \le \frac{(1-p)\kappa}{1-\kappa}, \\ \frac{1}{2}(1-\kappa)c^2 - \frac{1}{2}(1-p)p\left(\frac{\kappa^2}{1-\kappa}\right), & \frac{(1-p)\kappa}{1-\kappa} < c \le 1 - \frac{p\kappa}{1-\kappa}, \\ \frac{1}{2}p(1-\kappa)c^2 - \frac{1}{2}p(1-p)^2\left(\frac{\kappa^2}{1-\kappa}\right) & \\ \qquad +(1-p)(1-\kappa)\left(c - \frac{1}{2}\right), & 1 - \frac{p\kappa}{1-\kappa} < c \le 1, \end{cases}$$

$$u_3(c) := u(\beta = \kappa + (1-\kappa)c, c) = \begin{cases} \frac{1}{2}(1-\kappa)c^2 - \frac{1-p}{2}\left(\frac{\kappa^2}{1-\kappa}\right), & c \le 1 - \frac{\kappa}{1-\kappa}, \\ \frac{1}{2}p(1-\kappa)c^2 + (1-p)(1-\kappa)\left(c - \frac{1}{2}\right), & 1 - \frac{\kappa}{1-\kappa} < c \le 1, \end{cases}$$

$$u_4(c) := u(\beta = \kappa, c) = (1-p)(1-\kappa)\int_0^{\frac{\kappa}{1-\kappa}} (c-c')d(G(c')^{n_1}) = (1-p)\kappa c - \frac{1}{2}(1-p)\left(\frac{\kappa^2}{1-\kappa}\right),$$

$$\begin{aligned} u_5(c) := u(\beta = 1-\kappa, c) &= p(1-\kappa)\int_0^{\frac{1-2\kappa}{1-\kappa}} (c-c')d(G(c')^{n_1}) + (1-p)(1-\kappa)\int_0^1 (c-c')d(G(c')^{n_1}) \\ &= p(1-2\kappa)c - \frac{1}{2}p\left(\frac{(1-2\kappa)^2}{1-\kappa}\right) + (1-p)(1-\kappa)\left(c - \frac{1}{2}\right). \end{aligned}$$

Clearly, $u_4(c) \le u_1(c)$ (in fact u_4 is tangent to $\frac{1}{2}(1-p)(1-\kappa)c^2$ at $c = \frac{\kappa}{1-\kappa}$) and $u_5(c) \le u_3(c)$ (in fact u_5 is tangent to $\frac{1}{2}p(1-\kappa)c^2 + (1-p)(1-\kappa)\left(c - \frac{1}{2}\right)$ at $c = 1 - \frac{\kappa}{1-\kappa}$), so we can ignore u_4 and u_5. Then $\beta = (1-\kappa)c$ when $u_1(c) > u_2(c), u_3(c)$, $\beta = \kappa p + (1-\kappa)c$ when $u_2(c) > u_1(c), u_3(c)$, and $\beta = \kappa + (1-\kappa)c$ when $u_3(c) > u_1(c), u_2(c)$, and we break ties arbitrarily. We note that u_1, u_2, u_3 are all continuous in c and that $\frac{du_1}{dc} \le \frac{du_2}{dc} \le \frac{du_3}{dc}$, therefore u_1 can only be overtaken by u_2, u_3 and u_2 can only be overtaken by u_3, and u_3 cannot be overtaken. So, for $\kappa < 1-\kappa$, there must exist $\underline{c}$ and $\bar{c}$ such that $\beta(c) = (1-\kappa)c$ if $c < \underline{c}$, $\beta(c) = \kappa p + (1-\kappa)c$ if $\underline{c} < c < \bar{c}$, and $\beta(c) = \kappa + (1-\kappa)c$. Let us now find c_{12}, the point where u_2 overtakes u_1, suppose that $\frac{(1-p)\kappa}{1-\kappa} \le c_{12} \le \frac{\kappa}{1-\kappa}$: $\frac{1}{2}(1-p)(1-\kappa)c_{12}^2 = \frac{1}{2}(1-\kappa)c_{12}^2 - \frac{1}{2}(1-p)p\left(\frac{\kappa^2}{1-\kappa}\right) \implies c_{12} = \frac{\sqrt{1-p}\kappa}{1-\kappa} \in \left[\frac{(1-p)\kappa}{1-\kappa}, \frac{\kappa}{1-\kappa}\right]$. We do not need to further check other intervals due to the uniqueness of the intersection point. Similarly, we find the location of the point c_{23} where u_3 overtakes u_2: $1 - \frac{\kappa}{1-\kappa} \le c_{23} \le 1 - \frac{p\kappa}{1-\kappa}$: $\frac{1}{2}p(1-\kappa)c_{23}^2 + (1-p)(1-\kappa)\left(c_{23} - \frac{1}{2}\right) = \frac{1}{2}(1-\kappa)c_{23}^2 - \frac{1}{2}(1-p)p\left(\frac{\kappa^2}{1-\kappa}\right)$. This is a quadratic equation in c_{23} that has roots: $1 \pm \frac{\sqrt{p}\kappa}{1-\kappa}$. We take the negative root $c_{23} = 1 - \frac{\sqrt{p}\kappa}{1-\kappa} \in \left[1 - \frac{\kappa}{1-\kappa}, 1 - \frac{p\kappa}{1-\kappa}\right]$. Finally, we consider the point c_{13} where u_3 overtakes u_1, suppose that $1 - \frac{\kappa}{1-\kappa} < c_{13} < \frac{\kappa}{1-\kappa}$:

$\frac{1}{2}(1-p)(1-\kappa)c_{13}^2 = \frac{1}{2}p(1-\kappa)c_{13}^2 + (1-p)(1-\kappa)\left(c_{13}-\frac{1}{2}\right)$. This is a quadratic equation in c_{13} with two roots: $\frac{\sqrt{1-p}}{\sqrt{1-p}\pm\sqrt{p}}$. Since $c_{13}\in[0,1]$, we take the positive root: $c_{13}=\frac{\sqrt{1-p}}{\sqrt{1-p}+\sqrt{p}}$. Finally, we take $\underline{c} := \min\{c_{12}, c_{13}\}$, and $\bar{c} := \max\{c_{13}, c_{23}\}$, this also ensures that c_{13} is relevant only if $1-\frac{\kappa}{1-\kappa} < c_{23} < c_{13} < c_{12} < \frac{\kappa}{1-\kappa}$.

Now, suppose that $\kappa \geq 1-\kappa$. Let us consider the case when $0\leq\beta\leq 1-\kappa$, we have $u(\beta,c) = (1-p)(1-\kappa)\int_0^{\frac{\beta}{1-\kappa}}(c-c')d(G(c')^{n_1})$, as before, $\frac{\partial u}{\partial c}=0$ implies $\beta=(1-\kappa)c$. For $1-\kappa\leq\beta\leq\kappa$, we find that $u(\beta,c)=(1-p)(1-\kappa)\int_0^1(c-c')d(G(c')^{n_1})$, which is a constant in β. For $\kappa<\beta\leq 1$, we have $u(\beta,c)=p(1-\kappa)\int_0^{\frac{\beta-\kappa}{1-\kappa}}(c-c')d(G(c')^{n_1})+(1-p)(1-\kappa)\int_0^1(c-c')d(G(c')^{n_1})$, which means $\frac{\partial u}{\partial\beta}=0$ implies $\beta=\kappa+(1-\kappa)c$. This time we let

$$u_1(c) := u(\beta=(1-\kappa)c, c) = \frac{1}{2}(1-p)(1-\kappa)c^2,$$
$$u_2(c) := u(\beta=\kappa+(1-\kappa)c, c) = \frac{1}{2}p(1-\kappa)c^2 + (1-p)(1-\kappa)\left(c-\frac{1}{2}\right).$$

And, as before, $u_3(c) := u(\beta=\kappa, c)$, $u_4(c) := u(\beta=1-\kappa, c)$, which we can check that they satisfy $u_3(c)\leq u_1(c)$ and $u_4(c)\leq u_2(c)$, so we can ignore them. Since $\frac{du_1}{dc}=(1-p)(1-\kappa)\leq\frac{du_2}{dc}=(1-\kappa)c$, we conclude that u_2 can only overtake u_1 and cannot be overtaken. Hence, there exists $\underline{c}=\bar{c}$ such that $\beta(c)=(1-\kappa)c$ if $c<\underline{c}$ and $\beta(c)=\kappa+(1-\kappa)c$ if $c>\bar{c}$. Further inspection reveals that $\underline{c}=\bar{c}=c_{13}=\frac{\sqrt{1-p}}{\sqrt{1-p}+\sqrt{p}}$ as previously found. This completes the proof. ■

Proofs of Propositions 1 to 4

Propositions 1 to 4 are special cases of Propositions 5 to 8. We present the proofs of the more general statements in Appendix B. ■

Proof of Lemma 2

When $\kappa\geq 1/2$, we have $\kappa\geq 1-\kappa$, and we only need to consider two cases: $0\leq\beta\leq 1-\kappa$ where the constrained advertiser expected utility is $u(\beta,c)=(1-p)(1-\kappa)\int_0^{\frac{\beta}{1-\kappa}}(c-c')d(G(c')^{n_1})$ and $\kappa\leq\beta\leq 1$ where the constrained advertiser expected utility is $u(\beta,c)=p(1-\kappa)\int_0^{\frac{\beta-\kappa}{1-\kappa}}(c-c')d(G(c')^{n_1})+(1-p)(1-\kappa)\int_0^1(c-c')d(G(c')^{n_1})$. We do not need to consider the $1-\kappa<\beta<\kappa$ case since $u(\beta,c)=(1-p)(1-\kappa)\int_0^1(c-c')d(G(c')^{n_1})$ is constant in β over that domain. It follows that if $0\leq\beta\leq 1-\kappa$, then $\frac{\partial u}{\partial\beta}=n_1(1-p)\left(c-\frac{\beta}{1-\kappa}\right)G\left(\frac{\beta}{1-\kappa}\right)^{n_1-1}G'\left(\frac{\beta}{1-\kappa}\right)=0 \implies \beta=(1-\kappa)c$. Similarly, if $\kappa\leq\beta\leq 1$, then $\frac{\partial u}{\partial\beta}=n_1p\left(c-\frac{\beta-\kappa}{1-\kappa}\right)G\left(\frac{\beta-\kappa}{1-\kappa}\right)^{n_1-1}G'\left(\frac{\beta-\kappa}{1-\kappa}\right)=0 \implies \beta=$

$\kappa + (1-\kappa)c$. For any fixed c, the global maximum of u occurs either at $\beta = (1-\kappa)c, \kappa + (1-\kappa)c$ or at one of the singular points $\beta = \kappa, 1-\kappa$. Let

$$
\begin{aligned}
u_1(c) &:= u(\beta = (1-\kappa)c, c) = (1-p)(1-\kappa)\int_0^c (c-c')d(G(c')^{n_1}),\\
u_2(c) &:= u(\beta = \kappa + (1-\kappa)c, c) = p(1-\kappa)\int_0^c (c-c')d(G(c')^{n_1}) + (1-p)(1-\kappa)\int_0^1 (c-c')d(G(c')^{n_1}),\\
u_3(c) &:= u(\beta = \kappa, c) = (1-p)(1-\kappa)\int_0^{\frac{\kappa}{1-\kappa}} (c-c')d(G(c')^{n_1}),\\
u_4(c) &:= u(\beta = 1-\kappa, c) = p(1-\kappa)\int_0^{1-\frac{\kappa}{1-\kappa}} (c-c')d(G(c')^{n_1}) + (1-p)(1-\kappa)\int_0^1 (c-c')d(G(c')^{n_1}).
\end{aligned}
$$

Since $u_3(c)$ is the tangent line to $u_1(c)$ at $c = \frac{\kappa}{1-\kappa}$ and $u_4(c)$ is the tangent line to $u_2(c)$ at $c = 1-\frac{\kappa}{1-\kappa}$, and both u_1, u_2 are convex, we have $u_3 \le u_1$ and $u_4 \le u_2$, so we can ignore u_3, u_4. Next, we note that $\frac{du_1}{dc} = (1-p)(1-\kappa)G(c)^{n_1} < p(1-\kappa)G(c)^{n_1} + (1-p)(1-\kappa) = \frac{du_2}{dc}$ for all $c \in [0,1]$. We conclude that u_1 can only be overtaken by u_2. Note also that $u_1(0) = 0 > u_2(0) = -(1-p)(1-\kappa)\int_0^1 c' d(G(c')^{n_1})$ and $u_2(1) = p(1-\kappa)\int_0^c (1-c')d(G(c')^{n_1}) + u_1(1) > u_1(1)$, so the intersection point $\underline{c}(p) \in [0,1]$ exists and is unique. For a given distribution G, we can find $\underline{c}$ from the relation $u_1(\underline{c}) = u_2(\underline{c})$. Equivalently, the relation for $\underline{c}$ may be written as

$$
\int_0^1 (\underline{c} - c')d(G(c')^{n_1}) = \frac{1-2p}{1-p}\int_0^{\underline{c}} (\underline{c} - c')d(G(c')^{n_1}). \tag{2}
$$

Clearly, $1-\kappa$ cancels out and $\underline{c}$ is independent of κ. Furthermore, $u_1 - u_2$ is continuously differentiable in p and in $\underline{c}$ with nonvanishing derivative, and hence $\underline{c}$ is continuously differentiable in p by the Implicit Function Theorem. Differentiating $u_1 - u_2 = 0$ with respect to p, we get $-\frac{1}{1-p}\int_0^{\underline{c}} (\underline{c} - c')d(G(c')^{n_1}) = [(1-p)(1-G(\underline{c})^{n_1}) + pG(\underline{c})^{n_1}]\,\frac{d\underline{c}}{dp}$. The factor in the square bracket is positive and also $\int_0^{\underline{c}} (\underline{c} - c')d(G(c')^{n_1}) > 0$, hence $\frac{d\underline{c}}{dp} < 0$.

From (2) we can see that $p = 0$ implies $\int_{\underline{c}}^1 (\underline{c} - c')d(G(c')^{n_1}) = 0$, which holds exactly if $\underline{c}(0) = 1$ as the integral is < 0 for all $\underline{c} < 1$. Similarly, we can see that the LHS of (2) is bounded in $[-1,1]$, while the RHS approaches $-\infty$ as $p \to 1^-$, unless $\underline{c} \to 0$, which must be the case. Hence $\underline{c}(1) = 0$. Lastly, the RHS of (2) vanishes when $p = 1/2$, therefore, we are left with $\int_0^1 (\underline{c} - c')d(G(c')^{n_1}) = 0$ or $\underline{c} = \int_0^1 c' d(G(c')^{n_1}) = \mathbb{E}[n_1 c G(c)^{n_1 - 1}]$, as claimed. ■

B Online Appendix

B.1 Proofs of Propositions 5 to 10 and Lemma 3

Proof of Proposition 5

The statement of the proposition holds in a more general setting, which we prove in Lemma 4.

Lemma 4. *Consider any auction mechanism M such that, whenever bidders are symmetric and independent, in equilibrium M allocates the impression to the highest-valuation bidder. Suppose bidder i's valuation is given by some function of random variables corresponding to b-values and c-values: $v_i = v(b_i, c_i)$, where we assume v is increasing in c_i. Then under the common-value case $b_1 = b_2 = \cdots =: b$, for any distribution G, any $n_1, n_2 \geq 1$, and with selling mechanism M, we have $V_{\text{IA}}^{\text{CV}} \leq V_{\text{FI}}^{\text{CV}} = V_{\text{PI}}^{\text{CV}}$.*

Proof. Under both the full-information and the partial-information settings, when the b-value b is common among all the bidders, we have that the bidders are symmetric with their valuations determined by the independently drawn c-values c_i. Therefore, the impression is allocated to the bidder with the highest c_i under M. Under full information, the valuation of any bidder i is $v_i = v(b_i, c_i)$. Under the partial-information setting, the expected valuation of any bidder i is $\mathbb{E}[v_i] = \mathbb{E}[v(b_i, c_i)|c_i]$. It follows that the expected conversion rate is

$$V_{\text{FI}}^{\text{CV}} = \mathbb{E}_b\left[(n_1+n_2)\int_0^1 v(b,c)G(c)^{n_1+n_2-1}G'(c)dc\right] = (n_1+n_2)\int_0^1 \mathbb{E}\left[v(b,c)|c\right]G(c)^{n_1+n_2-1}G'(c)dc = V_{\text{PI}}^{\text{CV}},$$

where in the second equality we applied Fubini's Theorem. Since the mechanism M under full information ensures that the bidder with the highest valuation will win, it must be the case that $V_{\text{FI}}^{\text{CV}}$ is the highest possible conversion rate under any information setting. In particular, $V_{\text{IA}}^{\text{CV}} \leq V_{\text{FI}}^{\text{CV}} = V_{\text{PI}}^{\text{CV}}$. □

Proposition 5 follows by applying Lemma 4 to the case where M is a second-price auction and $v(b_i, c_i) := \kappa b_i + (1-\kappa)c_i$. Since $b_i = 1$ with probability p and $b_i = 0$ with probability $1-p$, it also follows that $\mathbb{E}[v(b_i, c_i)|c_i] = \kappa p + (1-\kappa)c_i$. ■

Proof of Proposition 6

The fact that $W^{\mathrm{CV}}_{\mathrm{PI}} = W^{\mathrm{CV}}_{\mathrm{FI}}$ is general and can be seen by directly comparing their expressions. For $n_1 = n_2 = 1$ and $\kappa \geq 1/2$ we may simplify (5) to:

$$\begin{aligned}
W^{\mathrm{CV}}_{\mathrm{IA}} =& 2p\int_{\underline{c}}^{1}(\kappa+(1-\kappa)c)(1-G(c))G'(c)dc + 2(1-p)\int_0^{\underline{c}}(1-\kappa)c(1-G(c))G'(c)dc \\
&+ p\int_0^{\underline{c}}(1-\kappa)cG'(c)dc + p\int_0^{\underline{c}}(\kappa+(1-\kappa)c)(1-G(\underline{c}))G'(c)dc + (1-p)\int_{\underline{c}}^{1}(1-\kappa)c(1-G(\underline{c}))G'(c)dc \\
=& W^{\mathrm{CV}}_{\mathrm{FI}} + p\left(\int_0^{\underline{c}}(\kappa+(1-\kappa)c)G(c)G'(c)dc - \int_0^{\underline{c}}(\kappa+(1-\kappa)c)(G(\underline{c})-G(c))G'(c)dc - \kappa G(\underline{c})\right) \\
&+ (1-p)\left(\int_{\underline{c}}^{1}(1-\kappa)c(G(c)-G(\underline{c}))G'(c)dc - \int_{\underline{c}}^{1}(1-\kappa)c(1-G(c))G'(c)dc\right) \\
=:& W^{\mathrm{CV}}_{\mathrm{FI}} + W_\Delta,
\end{aligned}$$

where W_Δ is defined to be the sum of the first and the second bracket. Let us show that $W_\Delta \leq 0$ for all $p \in [0,1]$. First, we note that (2) can be written equivalently as

$$pn_1\int_0^{\underline{c}} cG(c)^{n_1-1}G'(c)dc + (1-p)n_1\int_{\underline{c}}^{1} cG(c)^{n_1-1}G'(c)dc = \underline{c}pG(\underline{c})^{n_1} + \underline{c}(1-p)(1-G(\underline{c})^{n_1}). \quad (3)$$

Then, we have

$$\begin{aligned}
W_\Delta =& \frac{1}{2}\kappa pG(\underline{c})^2 + (1-\kappa)p\int_0^{\underline{c}} cG(c)G'(c)dc \\
&- \kappa pG(\underline{c})^2 - (1-\kappa)pG(\underline{c})\int_0^{\underline{c}} cG'(c)dc + \frac{1}{2}\kappa pG(\underline{c})^2 + (1-\kappa)p\int_0^{\underline{c}} cG(c)G'(c)dc - \kappa pG(\underline{c}) \\
&+ (1-\kappa)(1-p)\int_{\underline{c}}^{1} cG(c)G'(c)dc - (1-\kappa)(1-p)G(\underline{c})\int_{\underline{c}}^{1} cG'(c)dc - (1-\kappa)(1-p)\int_{\underline{c}}^{1} c(1-G(c))G'(c)dc \\
=& 2(1-\kappa)p\int_0^{\underline{c}} cG(c)G'(c)dc + (1-\kappa)(1-p)\int_{\underline{c}}^{1} cG(c)G'(c)dc \\
&- (1-\kappa)G(\underline{c})p\int_0^{\underline{c}} cG'(c)dc - (1-\kappa)G(\underline{c})(1-p)\int_{\underline{c}}^{1} cG'(c)dc \\
&- (1-\kappa)(1-p)\int_{\underline{c}}^{1} c(1-G(c))G'(c)dc - \kappa pG(\underline{c}) \\
\leq& (1-\kappa)p\underline{c}G(\underline{c})^2 + (1-\kappa)(1-p)\int_{\underline{c}}^{1} cG(c)G'(c)dc \\
&- (1-\kappa)G(\underline{c})\left(\underline{c}pG(\underline{c}) + \underline{c}(1-p)(1-G(\underline{c}))\right) - (1-\kappa)(1-p)\int_{\underline{c}}^{1} c(1-G(c))G'(c)dc \\
&+ (1-\kappa)(1-p)\underline{c}(1-G(\underline{c})) - (1-\kappa)p\int_0^{\underline{c}} cG'(c)dc - (1-\kappa)(1-p)\int_{\underline{c}}^{1} cG'(c)dc.
\end{aligned}$$

The last inequality can be explained as follows. We rewrite the first line using the inequalities: $2(1-\kappa)p\int_0^{\underline{c}} cG(c)G'(c)dc \leq (1-\kappa)p\underline{c}\int_0^{\underline{c}} d(G(c)^2) = (1-\kappa)p\underline{c}G(\underline{c})^2$. The second line follows from (3) with $n_1 = 1$. Lastly, we rewrite $\kappa pG(\underline{c})$ using (3) and the fact that $1-\kappa \leq \kappa$, $\underline{c} \leq 1$:

$\kappa p G(\underline{c}) \geq (1-\kappa)p\underline{c}G(\underline{c}) = (1-\kappa)p\int_0^{\underline{c}} cG'(c)dc + (1-\kappa)(1-p)\int_{\underline{c}}^1 cG'(c)d - (1-\kappa)\underline{c}(1-p)(1-G(\underline{c}))$.

Back to the main calculation, after some cancellations, the last inequality becomes:

$$\begin{aligned} W_\Delta \leq &(1-\kappa)(1-p)\underline{c}(1-G(\underline{c}))^2 - 2(1-\kappa)(1-p)\int_{\underline{c}}^1 c(1-G(c))G'(c)dc - (1-\kappa)p\int_0^{\underline{c}} cG'(c)dc \\ \leq &(1-\kappa)(1-p)\underline{c}(1-G(\underline{c}))^2 - 2(1-\kappa)(1-p)\underline{c}\int_{\underline{c}}^1 (1-G(c))G'(c)dc - (1-\kappa)p\int_0^{\underline{c}} cG'(c)dc \\ = &-(1-\kappa)p\int_0^{\underline{c}} cG'(c)dc \leq 0. \end{aligned}$$

Therefore, we have that $W_\Delta \leq 0$ as needed. In fact, we can see from $W_\Delta \leq -(1-\kappa)p\int_0^{\underline{c}} cG'(c)dc$, that the equality holds exactly when $p = 0, 1$, i.e. $W_{\text{IA}}^{\text{CV}}|_{p=0,1} = W_{\text{FI}}^{\text{CV}}|_{p=0,1} = W_{\text{PI}}^{\text{CV}}|_{p=0,1}$. ■

Proof of Proposition 7

Let's consider $N := n_1 + n_2$ advertisers which are divided into two disjoint subsets $A = \{a_1^A, \ldots, a_{n_1}^A\}$ and $B = \{a_1^B, \ldots, a_{n_2}^B\}$, $A \coprod B = \{1, 2, \ldots, n_1 + n_2\}$. Set A contains informed advertisers with full information and hence bid their true valuation. Set B contains constrained advertisers with only the c-value, and hence bid the expected value $\kappa p + (1-\kappa)c$. Let's consider an instance of an auction where the c-values in set B are given by $c_1 > \cdots > c_{n_2}$, whereas the highest bid in A is given by the bidder a^* with valuation $\kappa b^* + (1-\kappa)c^*$. Independently, we also draw $b_1, \ldots, b_{n_2}$ b-values for the constrained advertisers in B.

First, we consider the case where a_1^B from B is the winner: $\kappa p + (1-\kappa)c_1 > \kappa b^* + (1-\kappa)c^*$. Suppose that we moved an advertiser $a_i^B \neq a_1^B$ from set B to set A, keeping all the c-values fixed. After the move, either a_i^B becomes the winner or nothing changes. Suppose a_i^B becomes the winner, this means $b_i = 1$, $\kappa + (1-\kappa)c_i > \kappa b^* + (1-\kappa)c^*$ and $\kappa + (1-\kappa)c_i > \kappa p + (1-\kappa)c_1 \implies c_1 - c_i < \frac{\kappa(1-p)}{1-\kappa}$.

With probability p we have $b_1 = 1$, and in this case, the change in the winner's valuation Δv_w is given by $\mathbb{E}[\Delta v_w | b_1 = 1] = (\kappa + (1-\kappa)c_i) - (\kappa + (1-\kappa)c_1) = -(1-\kappa)(c_1 - c_i) > -\kappa(1-p)$.

With probability $1-p$ we have $b_1 = 0$, and in this case, the change in the winner's valuation is given by $\mathbb{E}[\Delta v_w | b_1 = 0] = (\kappa + (1-\kappa)c_i) - (1-\kappa)c_1 = \kappa - (1-\kappa)(c_1 - c_i) > \kappa - \kappa(1-p) = \kappa p$.

Therefore, the expected change of the winner's valuation is $\mathbb{E}[\Delta v_w] = \mathbb{E}[\Delta v_w | b_1 = 1]\mathbb{P}[b_1 = 1] + \mathbb{E}[\Delta v_w | b_1 = 0]\mathbb{P}[b_1 = 0] > -p \cdot \kappa(1-p) + (1-p) \cdot \kappa p = 0$.

Suppose that we moved an advertiser a_1^B from set B to A, keeping all the drawn c-values fixed.

After the move, either a_1^B remains the winner, hence nothing changes, or it is not. If a_1^B is no longer a winner, then either a^* is the winner, in that case, we have an increase in the winner's valuation since $\kappa b^* + (1-\kappa)c^* > \kappa b_1 + (1-\kappa)c_1$. Otherwise, a_2^B is now the winner, so we must have $b_1 = 0$ and $\kappa p + (1-\kappa)c_2 > (1-\kappa)c_1 \implies c_1 - c_2 < \frac{\kappa p}{1-\kappa}$.

With probability p we have $b_2 = 1$, and in this case, the change in the winner's valuation is given by $\mathbb{E}[\Delta v_w | b_2 = 1] = (\kappa + (1-\kappa)c_2) - (1-\kappa)c_1 = \kappa - (1-\kappa)(c_1 - c_2) > \kappa - \kappa p = \kappa(1-p)$.

With probability $1-p$ we have $b_2 = 0$, and in this case, the change in the winner's valuation is given by $\mathbb{E}[\Delta v_w | b_2 = 0] = (1-\kappa)c_2 - (1-\kappa)c_1 = -(1-\kappa)(c_1 - c_2) > -\kappa p$.

Therefore, the expected change of the winner's valuation is $\mathbb{E}[\Delta v_w] = \mathbb{E}[\Delta v_w | b_1 = 1]\mathbb{P}[b_1 = 1] + \mathbb{E}[\Delta v_w | b_1 = 0]\mathbb{P}[b_1 = 0] > p \cdot \kappa(1-p) - (1-p) \cdot \kappa p = 0$.

Now, we consider the case where a^* is the winner: $\kappa b^* + (1-\kappa)c^* > \kappa p + (1-\kappa)c_1$. If a winner changed by moving an a_i^B advertiser from the set B to the set A, keeping all the drawn b-values and c-values fixed, then the moved advertiser must have $\kappa b_i + (1-\kappa)c_i > \kappa b^* + (1-\kappa)c^*$. Therefore, the winner's valuation can only increase in this case.

It follows that the conversion rate increases or remains the same for every advertiser we move from set B to set A. We conclude that $V_{\text{FI}}^{\text{CV}} \geq V_{\text{IA}}^{\text{CV}} \geq V_{\text{PI}}^{\text{CV}}$. ■

Proof of Proposition 8

Let's denote by $w_{\text{FI}}^{\text{IV}}, w_{\text{IA}}^{\text{IV}}, w_{\text{PI}}^{\text{IV}}$ the revenue under each information setting for an instant of auction, so that we have $W_{\text{FI}}^{\text{IV}} := \mathbb{E}[w_{\text{FI}}^{\text{IV}}], W_{\text{IA}}^{\text{IV}} := \mathbb{E}[w_{\text{IA}}^{\text{IV}}], W_{\text{PI}}^{\text{IV}} := \mathbb{E}[w_{\text{PI}}^{\text{IV}}]$. Consider an instance of auction where the c-value of the informed advertiser is c_1 and the c-value of the constrained advertiser is c_2. Both c_1, c_2 are drawn independently from the distribution G. First, we consider $\mathbb{E}\left[w_{\text{FI}}^{\text{IV}} | c_1, c_2\right]$, there are two cases: the case $\max\{c_1, c_2\} > \min\{c_1, c_2\} + \frac{\kappa}{1-\kappa}$, for which we find: $\mathbb{E}\left[w_{\text{FI}}^{\text{IV}} | c_1, c_2\right] = (\kappa + (1-\kappa)\min\{c_1, c_2\}) \cdot p + (1-\kappa)\min\{c_1, c_2\} \cdot (1-p) = \kappa p + (1-\kappa)\min\{c_1, c_2\}$ and the case: $\min\{c_1, c_2\} + \frac{\kappa}{1-\kappa} > \max\{c_1, c_2\}$, where we have

$$\begin{aligned}\mathbb{E}\left[w_{\text{FI}}^{\text{IV}} | c_1, c_2\right] &= (\kappa + (1-\kappa)\min\{c_1, c_2\}) \cdot p^2 + (1-\kappa)\min\{c_1, c_2\} \cdot (1-p) + (1-\kappa)\max\{c_1, c_2\} \cdot (1-p)p \\ &= \kappa p^2 + (1-\kappa)\min\{c_1, c_2\} \cdot (1-p+p^2) + (1-\kappa)\max\{c_1, c_2\} \cdot (1-p)p.\end{aligned}$$

Next, we consider $\mathbb{E}\left[w_{\text{IA}}^{\text{IV}} | c_1, c_2\right]$, there are three cases: the case $c_1 > c_2 + \frac{p\kappa}{1-\kappa}$, for which we find $\mathbb{E}\left[w_{\text{IA}}^{\text{IV}} | c_1, c_2\right] = \kappa p + (1-\kappa)c_2$, the case: $c_2 + \frac{p\kappa}{1-\kappa} > c_1 > c_2 - \frac{(1-p)\kappa}{1-\kappa}$, where we have $\mathbb{E}\left[w_{\text{IA}}^{\text{IV}} | c_1, c_2\right] =$

$(\kappa p + (1-\kappa)c_2)\cdot p + (1-\kappa)c_1\cdot(1-p)$, and the case: $c_2 - \frac{(1-p)\kappa}{1-\kappa} > c_1$, where we have $\mathbb{E}\left[w^{\text{IV}}_{\text{IA}}|c_1,c_2\right] = (\kappa + (1-\kappa)c_1)\cdot p + (1-\kappa)c_1\cdot(1-p) = \kappa p + (1-\kappa)c_1$.

Lastly, in all cases we have that $\mathbb{E}\left[w^{\text{IV}}_{\text{PI}}|c_1,c_2\right] = \kappa p + (1-\kappa)\min\{c_1,c_2\}$.

Now we can check that for all possible pairs of c_1, c_2 we have $\mathbb{E}\left[w^{\text{IV}}_{\text{FI}}|c_1,c_2\right] \le \mathbb{E}\left[w^{\text{IV}}_{\text{IA}}|c_1,c_2\right] \le \mathbb{E}\left[w^{\text{PI}}_{\text{IA}}|c_1,c_2\right]$. Taking an expectation over all possible c_1, c_2 we have that $W^{\text{IV}}_{\text{FI}} \le W^{\text{IV}}_{\text{IA}} \le W^{\text{IV}}_{\text{PI}}$ as claimed. ■

Proof of Proposition 9

From Proposition 8 we already know that for $n_1 = n_2 = 1$ we have $W^{\text{IV}}_{\text{FI}} \le W^{\text{IV}}_{\text{PI}}$ for all $p \in [0,1]$. Now, fix $p \in [0,1]$ and consider the $n_1, n_2 > 0$ case. Note that using integration by-parts we can rewrite $W^{\text{IV}}_{\text{FI}}$ as:

$$\begin{aligned}
W^{\text{IV}}_{\text{FI}} =& (n_1+n_2)(1-\kappa)(1-p)^{n_1+n_2-1}p - (n_1+n_2)\kappa(1-p)^{n_1+n_2-1}p \\
&+ (1-\kappa)\int_0^1 \left(c - \frac{1-(1-p)G(c)}{(1-p)G'(c)}\right) d\left((1-p)G(c)\right)^{n_1+n_2} \\
&+ \int_0^1 \left(\kappa + (1-\kappa)c - \frac{(1-\kappa)(1-G(c))}{G'(c)}\right) d\left(pG(c)+(1-p)\right)^{n_1+n_2} \\
=& -(n_1+n_2)(2\kappa-1)(1-p)^{n_1+n_2-1}p + (1-\kappa)(1-p)^{n_1+n_2}\mathbb{E}_{c\sim G(c)^{n_1+n_2}}\left[c - \frac{1-(1-p)G(c)}{(1-p)G'(c)}\right] \\
&+ \mathbb{E}_{c\sim(pG(c)+(1-p))^{n_1+n_2}}\left[\kappa + (1-\kappa)c - \frac{(1-\kappa)(1-G(c))}{G'(c)}\right].
\end{aligned}$$

Where $\mathbb{E}_{c\sim F(c)}[.]$ denotes the expected value with c distributed by $F(c)$. Similarly, we can rewrite $W^{\text{IV}}_{\text{PI}}$ as

$$\begin{aligned}
W^{\text{IV}}_{\text{PI}} &= \int_0^1 \left(\kappa p + (1-\kappa)c - \frac{(1-\kappa)(1-G(c))}{G'(c)}\right) dG(c)^{n_1+n_2} \\
&= \mathbb{E}_{c\sim G(c)^{n_1+n_2}}\left[\kappa p + (1-\kappa)c - \frac{(1-\kappa)(1-G(c))}{G'(c)}\right].
\end{aligned}$$

When n_1 and n_2 are large, the densities of distributions $G(c)^{n_1+n_2}$ and $(pG(c)+(1-p))^{n_1+n_2}$ become concentrated around $c = 1$. Therefore, $W^{\text{IV}}_{\text{PI}}$ tends towards $\kappa p + (1-\kappa)$. On the other hand, the first and second terms in $W^{\text{IV}}_{\text{FI}}$ tend to zero due to $(1-p)^{n_1+n_2}$ but the last term tends to $\kappa + (1-\kappa) = 1$. Hence, we have $W^{\text{IV}}_{\text{FI}} > W^{\text{IV}}_{\text{PI}}$ for all sufficiently large n_1 and n_2. ■

Proof of Proposition 10

Using the formulas in the Online Appendix B.2 for $G(x) = x$ and $\kappa \geq 1/2$, we derive that

$$\frac{\partial}{\partial p}\left(D_{\mathrm{IA}}^{\mathrm{IV}} - D_{\mathrm{PI}}^{\mathrm{IV}}\right)\Big|_{p=0} = \frac{n_1}{n_1+n_2+1}\left(\frac{2\kappa}{n_1+n_2-1} - \frac{1}{n_1+n_2}\right)$$

and

$$\frac{\partial}{\partial p}\left(E_{\mathrm{FI}}^{\mathrm{IV}} - E_{\mathrm{IA}}^{\mathrm{IV}}\right)\Big|_{p=0} = \frac{n_2}{n_1+n_2+1}\left(\frac{2\kappa}{n_1+n_2-1} - \frac{1}{n_1+n_2}\right).$$

Since $\frac{2\kappa}{n_1+n_2-1} - \frac{1}{n_1+n_2} \geq \frac{1}{n_1+n_2-1} - \frac{1}{n_1+n_2} > 0$ when $\kappa \geq 1/2$, both derivatives are strictly positive at $p = 0$. Therefore, there is a neighborhood of $p = 0$, where the result holds. ■

Proof of Lemma 3

The informed advertisers will always bid their true valuation as it is a weakly dominant strategy to do so. Therefore, for the remainder, we will focus on the nontrivial part, which is the constrained advertisers' bidding strategy.

The expected utility for a constrained advertiser with c-value c from bidding $\tilde{\beta}$ when all other $n_2 - 1$ constrained advertisers follow the strategy β is given by:

$$\begin{aligned}
u(\tilde{\beta};\beta,c) :=& p\Bigg[(1-\kappa)\int_0^{\max\left\{\frac{\tilde{\beta}-\kappa}{1-\kappa},0\right\}}(c-c')G(\sup\beta^{-1}[0,\kappa+(1-\kappa)c'))^{n_2-1}\left(n_1G(c')^{n_1-1}G'(c')\right)dc' \\
&+\int_0^{\sup\beta^{-1}[0,\tilde{\beta})}(\kappa+(1-\kappa)c-\beta(c'))G\left(\max\left\{\frac{\beta(c')-\kappa}{1-\kappa},0\right\}\right)^{n_1}\left((n_2-1)G(c')^{n_2-2}G'(c')\right)dc'\Bigg] \\
&+(1-p)\Bigg[(1-\kappa)\int_0^{\min\left\{\frac{\tilde{\beta}}{1-\kappa},1\right\}}(c-c')G(\sup\beta^{-1}[0,(1-\kappa)c'))^{n_2-1}\left(n_1G(c')^{n_1-1}G'(c')\right)dc' \\
&+\int_0^{\sup\beta^{-1}[0,\tilde{\beta})}((1-\kappa)c-\beta(c'))G\left(\min\left\{\frac{\beta(c')}{1-\kappa},1\right\}\right)^{n_1}\left((n_2-1)G(c')^{n_2-2}G'(c')\right)dc'\Bigg].
\end{aligned} \tag{4}$$

Let us restrict our attention to the bidding functions β that belong to the following class of functions: $\mathcal{F} := \{\beta \in L^1[0,1] \mid \beta \text{ is represented by a non-decreasing function } [0,1] \to [0,1]\}$.

Here, $L^1[0,1]$ denotes the usual Banach space of the equivalence classes of Lebesgue-integrable functions on $[0,1]$ equipped with the usual norm $\|f\|_{L^1} := \int_0^1 |f(x)|dx$. It is not hard to verify that $\mathcal{F}$ is a convex and compact subset of $L^1[0,1]$.

We note that the sign of each of the integrals in (4) is determined by the sign of $(c - c')$,

$\kappa + (1-\kappa)c - \beta(c')$, and $(1-\kappa)c - \beta(c')$, respectively, all of which are increasing functions in c and decreasing in c'. Essentially, given β and c, finding the maximum $\tilde{\beta} = \tilde{\beta}_0$ of $u(\tilde{\beta}; \beta, c)$ is to 'integrate until the integrands are negative'. The reality is slightly more subtle, as the upper limit of each integral are different non-linear functions of $\tilde{\beta}$.

Lemma 5. *Given $\beta \in \mathcal{F}$ and $c \in [0,1]$ then $u(\tilde{\beta}; \beta, c)$ as a function of $\tilde{\beta} \in [0,1]$ achieves its global maximum inside $\{(1-\kappa)c\} \cup [\kappa, \kappa + (1-\kappa)c]$.*

Proof. First, let us observe where a maximum of $u(\tilde{\beta}; \beta, c)$ cannot be located. If $\tilde{\beta}(c) \in (\kappa + (1-\kappa)c, 1]$ then the third term of (4) is constant in $\tilde{\beta}$. The first term is strictly decreasing for $\tilde{\beta} > \kappa + (1-\kappa)c$. The second and fourth terms are non-constant if $\beta(c) > \kappa + (1-\kappa)c$ for some c, but then these two terms decrease with $\tilde{\beta}$ because $\beta(c') > \kappa + (1-\kappa)c > (1-\kappa)c$ for $c' = \max \beta^{-1}[0, \tilde{\beta}) \geq \max \beta^{-1}[0, \kappa + (1-\kappa)c)$. Similarly, if $\tilde{\beta} \in [0, (1-\kappa)c)$ then only the third and fourth terms of (4) are non-constant in $\tilde{\beta}$. The third term is strictly increasing for $\tilde{\beta} < (1-\kappa)c$ and for any $c' = \max \beta^{-1}[0, \tilde{\beta}) \leq \max \tilde{\beta}^{-1}[0, (1-\kappa)c)$, which means $\beta(c') < (1-\kappa)c$, hence the fourth term is increasing.

If $\tilde{\beta} \in [(1-\kappa)c, \kappa]$, then every term of (4) is constant except for the fourth term which could be non-constant if $\beta(c) > (1-\kappa)c$ for some c, and in that case, the fourth term is decreasing. In other words, the maximum value of $u(\tilde{\beta}; \beta, c)$ over $[(1-\kappa)c, \kappa]$ is reached at $\tilde{\beta} = (1-\kappa)c$. Since the fourth term of (4) is necessarily strictly decreasing, it is possible that $u(\tilde{\beta}; \beta, c)$ also attains its maximum value at other points in $((1-\kappa)c, \kappa]$, this fact will serve no practical implication for us.

Next, we focus on the case where $\tilde{\beta} \in [\kappa, \kappa + (1-\kappa)c]$, and we shall show that $u(\tilde{\beta}; \beta, c)$ also reaches its maximum over this interval. We note that $u(\tilde{\beta}; \beta, c)$ is left-continuous because $\sup \beta^{-1}[0, \tilde{\beta})$ is left-continuous, and the point where it is not continuous is exactly where $\{c \mid \beta(c) = \tilde{\beta}_0\}$ has non-empty interior. In particular, let $\underline{b} := \inf\{c \mid \beta(c) = \tilde{\beta}_0\}$ and $\bar{b} := \sup\{c \mid \beta(c) = \tilde{\beta}_0\}$ then it follows that $(\underline{b}, \bar{b}) \subset S(\tilde{\beta}_0)$. In that case, we have $\sup \beta^{-1}[0, \tilde{\beta}) \leq \underline{b}$ for $\tilde{\beta} \leq \tilde{\beta}_0$ and $\sup \beta^{-1}[0, \tilde{\beta}) \geq \bar{b}$ for $\tilde{\beta} > \tilde{\beta}_0$. Given that $\tilde{\beta} \in [\kappa, \kappa + (1-\kappa)c]$, the third term of (4) is constant in a neighborhood of $\tilde{\beta}_0$. Let $\delta > 0$ be arbitrarily small, then the first term of (4) will take approximately the same value at $\tilde{\beta}_0$ and at $\tilde{\beta}_0 + \delta$. If $u(\tilde{\beta}_0; \beta, c) < u(\tilde{\beta}_0 + \delta; \beta, c)$ it must be the case that the sum of the second and fourth integrals is positive over $(\underline{b}, \bar{b})$. In particular, $\left(\kappa + (1-\kappa)c - \tilde{\beta}_0\right) G\left(\max\left\{\frac{\tilde{\beta}_0 - \kappa}{1-\kappa}, 0\right\}\right)^{n_1} + \left((1-\kappa)c - \tilde{\beta}_0\right) G\left(\min\left\{\frac{\tilde{\beta}_0}{1-\kappa}, 1\right\}\right)^{n_1} > 0$.

But we also know that for all $c' \in \sup \beta^{-1}[0, \tilde{\beta}_0 + \delta)$ we have $\beta(c') < \tilde{\beta}_0 + \delta$, then from the inequality above we have that the sum of the integrands of the second and fourth integrals in (4) is positive immediately to the right of $\tilde{\beta}_0$ as $\delta > 0$ is arbitrary small. Since $\tilde{\beta}_0 \leq \kappa + (1-\kappa)c$, the integrand of the first integral in (4) is also positive. It follows that $u(\tilde{\beta}; \beta, c)$ continue to increase over some right neighborhood of $\tilde{\beta}_0 + \delta$, hence $\sup_{\tilde{\beta}} u(\tilde{\beta}; \beta, c) > \lim_{\tilde{\beta} \to \tilde{\beta}_0^+} u(\tilde{\beta}; \beta, c)$. □

Lemma 5 allows us to define the best-response set-valued function as follows: $BR(\beta, c) := \arg\max_{\tilde{\beta} \in [0,1]} u(\tilde{\beta}; \beta, c)$. Let us also restrict our attention to β such that $\beta(c) \in \{(1-\kappa)c\} \cup [\kappa, \kappa + (1-\kappa)c]$.

Lemma 6. *The best-response function is closed-valued and* non-decreasing *in the sense that if* $c_1 < c_2$*, then* $\max BR(\beta, c_1) \leq \min BR(\beta, c_2)$.

Proof. The fact that $BR(\beta, c)$ is closed follows since according to Lemma 5, $u(\tilde{\beta}; \beta, c)$ is left-continuous and if $u(\tilde{\beta}; \beta, c)$ is discontinuous at $\tilde{\beta}_0$ then $\limsup_{\tilde{\beta} \to \tilde{\beta}_0^+} u(\tilde{\beta}; \beta, c)$ is always less than the global maximum value of u. In other words, if $\tilde{\beta}_i \in BR(\beta, c), i = 1, 2, \ldots$ and $\tilde{\beta}_i \to \tilde{\beta}_0 \in [0, 1]$ then $u(\tilde{\beta}_0; \beta, c) = u(\tilde{\beta}_i; \beta, c)$ for all i, which means $\tilde{\beta}_0 \in BR(\beta, c)$. Therefore, it makes sense to talk about the maximum and minimum of $BR(\beta, c)$.

Given any $\delta > 0$, we note that it is possible to write $u(\tilde{\beta}; \beta, c+\delta) = u(\tilde{\beta}; \beta, c) + \Delta(\tilde{\beta}; \beta, \delta)$, where $\Delta(\tilde{\beta}; \beta, c)$ is exactly given by (4) but with $(c - c')$, $(\kappa + (1-\kappa)c - \beta(c'))$, and $((1-\kappa)c - \beta(c'))$ factors replaced by $\delta, (1-\kappa)\delta$, and $(1-\kappa)\delta$, respectively. Thus, $\Delta(\tilde{\beta}; \beta, \delta)$ is a non-decreasing function in $\tilde{\beta}$ and strictly increases over $[0, 1-\kappa] \cup [\kappa, 1]$. Then the fact that $BR(\beta, c)$ is non-decreasing follows from the following elementary argument. Let $\tilde{\beta}_0 = \max BR(\beta, c)$ then $u(\tilde{\beta}_0; \beta, c) \geq u(\tilde{\beta}; \beta, c)$ for all $\tilde{\beta} \in [0, \tilde{\beta}_0)$. Therefore, $u(\tilde{\beta}_0; \beta, c+\delta) > u(\tilde{\beta}; \beta, c+\delta)$ for all $\tilde{\beta} \in [0, \tilde{\beta}_0)$ by the strict monotonicity of $\Delta(\tilde{\beta}; \beta, \delta)$, which means any other global maxima of $u(\tilde{\beta}; \beta, c+\delta)$ must be in $[\tilde{\beta}_0, \kappa + (1-\kappa)c]$, proving the lemma. □

Using Lemma 6 it is now possible to define the best-response bidding function to the bidding β of all other $n_2 - 1$ constrained advertisers: $BR : \mathcal{F} \to \mathcal{F}$, $\tilde{\beta} := BR(\beta) : c \mapsto \min BR(\beta, c)$, where we have slightly abused the notation, using both $\tilde{\beta}$ as a particular bidding value and the bidding function, and BR as both the best response bidding set-valued function and the best response bidding function-valued map. However, we hope that any ambiguity can be resolved by context.

Lemma 7. *The best-response function BR is continuous with respect to the L^1 norm.*

We will omit the technical proof, but the intuition is clear. Any two $\beta_1, \beta_2 \in \mathcal{F}$ non-decreasing functions which are 'close' together under L^1 norm must take similar values $\beta_1(c) \approx \beta_2(c)$ at any c they are both continuous. Moreover, the location of any discontinuous points of β_1 and β_2 must be similar. The same is true for their inverses $\sup \beta_1^{-1}[0, \tilde{\beta}) \approx \beta_2^{-1}[0, \tilde{\beta})$. Hence we can expect $u(\tilde{\beta}; \beta_1, c) \approx u(\tilde{\beta}; \beta_2, c)$ for all $\tilde{\beta}$ and c and therefore the maximum point of $u(.; \beta_1, c)$ should be close to the maximum point of $u(.; \beta_2, c)$.

From Lemma 7, the response function BR is continuous with respect to L^1 norm and maps a convex compact subset $\mathcal{F} \subset L^1[0,1]$ into itself. $L^1[0,1]$ is a normed-vector space, hence it is automatically a Hausdorff locally convex topological vector space. From the Kakutani-Fan-Glicksberg Theorem, we know that BR has a fixed point. ■

B.2 Key Formulas

Unless stated otherwise, all formulas in this section are valid for any given $\kappa \in [0,1]$, $p \in [0,1]$, $n_1, n_2 \geq 0$ and an arbitrary c-value distribution G on $[0,1]$.

Common-value case

To deal with any discontinuities of the bidding function β we let $\beta^{-1}[a,b)$ denote the inverse image of β i.e. a set I such that $x \in I \implies \beta(x) \in [a,b)$, and $\sup \beta^{-1}[a,b)$ denotes the supremum of this set.

Advertisers' conversion rate:

The advertisers' conversion rates under each information setting are given by:

$$\begin{aligned} V_{\text{FI}}^{\text{CV}} =& p(n_1+n_2)\int_0^1 (\kappa + (1-\kappa)c) G(c)^{n_1+n_2-1} G'(c) dc \\ &+ (1-p)(n_1+n_2)\int_0^1 (1-\kappa) c G(c)^{n_1+n_2-1} G'(c) dc, \end{aligned}$$

$$\begin{aligned} V_{\text{IA}}^{\text{CV}} =& pn_2 \int_0^1 (\kappa + (1-\kappa)c) G\left(\max\left\{\frac{\beta(c)-\kappa}{1-\kappa}, 0\right\}\right)^{n_1} G(c)^{n_2-1} G'(c) dc \\ &+ pn_1 \int_0^1 (\kappa + (1-\kappa)c) G(c)^{n_1-1} G\left(\sup \beta^{-1}[0, \kappa + (1-\kappa)c)\right)^{n_2} G'(c) dc \end{aligned}$$

$$+ (1-p)n_2 \int_0^1 (1-\kappa)cG\left(\min\left\{\frac{\beta(c)}{1-\kappa}, 1\right\}\right)^{n_1} G(c)^{n_2-1}G'(c)dc$$
$$+ (1-p)n_1 \int_0^1 (1-\kappa)cG(c)^{n_1-1}G\left(\sup \beta^{-1}[0,(1-\kappa)c)\right)^{n_2} G'(c)dc,$$

$$V_{\text{PI}}^{\text{CV}} = (n_1+n_2)\int_0^1 (\kappa p + (1-\kappa)c)G(c)^{n_1+n_2-1}G'(c)dc.$$

We note that $V_{\text{FI}}^{\text{CV}} = V_{\text{PI}}^{\text{CV}}$.

Publisher's expected revenue:

The publisher's expected revenues for each information setting are given by:

$$W_{\text{FI}}^{\text{CV}} = p(n_1+n_2)(n_1+n_2-1)\int_0^1 (\kappa + (1-\kappa)c)(1-G(c))G(c)^{n_1+n_2-2}G'(c)dc$$
$$+ (1-p)(n_1+n_2)(n_1+n_2-1)\int_0^1 (1-\kappa)c(1-G(c))G(c)^{n_1+n_2-2}G'(c)dc,$$

$$\begin{aligned}
W_{\text{IA}}^{\text{CV}} =& pn_1n_2 \int_0^1 (\kappa + (1-\kappa)c)\left(1-G\left(\sup\beta^{-1}[0,\kappa+(1-\kappa)c)\right)\right) G(c)^{n_1-1}G\left(\sup\beta^{-1}[0,\kappa+(1-\kappa)c)\right)^{n_2-1} G'(c)dc \\
&+ pn_1n_2\int_0^1 \beta(c)\left(1-G\left(\max\left\{\frac{\beta(c)-\kappa}{1-\kappa},0\right\}\right)\right) G\left(\max\left\{\frac{\beta(c)-\kappa}{1-\kappa},0\right\}\right)^{n_1-1} G(c)^{n_2-1}G'(c)dc \\
&+ pn_1(n_1-1)\int_0^1 (\kappa+(1-\kappa)c)\left(1-G(c)\right)G(c)^{n_1-2}G\left(\sup\beta^{-1}[0,\kappa+(1-\kappa)c)\right)^{n_2} G'(c)dc \\
&+ pn_2(n_2-1)\int_0^1 \beta(c)\left(1-G(c)\right)G\left(\max\left\{\frac{\beta(c)-\kappa}{1-\kappa},0\right\}\right)^{n_1} G(c)^{n_2-2}G'(c)dc \\
&+ (1-p)n_1n_2\int_0^1 (1-\kappa)c\left(1-G\left(\sup\beta^{-1}[0,(1-\kappa)c)\right)\right)G(c)^{n_1-1}G\left(\sup\beta^{-1}[0,(1-\kappa)c)\right)^{n_2-1} G'(c)dc \\
&+ (1-p)n_1n_2\int_0^1 \beta(c)\left(1-G\left(\min\left\{\frac{\beta(c)}{1-\kappa},1\right\}\right)\right)G\left(\min\left\{\frac{\beta(c)}{1-\kappa},1\right\}\right)^{n_1-1} G(c)^{n_2-1}G'(c)dc \\
&+ (1-p)n_1(n_1-1)\int_0^1 (1-\kappa)c\left(1-G(c)\right)G(c)^{n_1-2}G\left(\sup\beta^{-1}[0,(1-\kappa)c)\right)^{n_2} G'(c)dc \\
&+ (1-p)n_2(n_2-1)\int_0^1 \beta(c)\left(1-G(c)\right)G\left(\min\left\{\frac{\beta(c)}{1-\kappa},1\right\}\right)^{n_1} G\left(c\right)^{n_2-2}G'(c)dc,
\end{aligned} \tag{5}$$

$$W_{\text{PI}}^{\text{CV}} = (n_1+n_2)(n_1+n_2-1)\int_0^1 (\kappa p + (1-\kappa)c)(1-G(c))G(c)^{n_1+n_2-2}G'(c)dc.$$

We note that $W_{\text{FI}}^{\text{CV}} = W_{\text{PI}}^{\text{CV}}$.

Constrained advertisers' payoff:

$$\begin{aligned}
D_{\text{FI}}^{\text{CV}} = \frac{V_{\text{FI}}^{\text{CV}} - W_{\text{FI}}^{\text{CV}}}{n_1+n_2} =& \int_0^1 (p\kappa + (1-\kappa)c)G(c)^{n_1+n_2-1}G'(c)dc \\
&- (n_1+n_2-1)\int_0^1 (p\kappa + (1-\kappa)c)(1-G(c))G(c)^{n_1+n_2-2}G'(c)dc,
\end{aligned} \tag{6}$$

$$
\begin{aligned}
D_{\text{IA}}^{\text{CV}} =&p\int_0^1(\kappa+(1-\kappa)c)G\left(\max\left\{\frac{\beta(c)-\kappa}{1-\kappa},0\right\}\right)^{n_1}G(c)^{n_2-1}G'(c)dc\\
&+(1-p)\int_0^1(1-\kappa)cG\left(\min\left\{\frac{\beta(c)}{1-\kappa},1\right\}\right)^{n_1}G(c)^{n_2-1}G'(c)dc\\
&-pn_1\int_0^1(\kappa+(1-\kappa)c)\left(1-G\left(\sup\beta^{-1}[0,\kappa+(1-\kappa)c)\right)\right)G(c)^{n_1-1}G\left(\sup\beta^{-1}[0,\kappa+(1-\kappa)c)\right)^{n_2-1}G'(c)dc\\
&-p(n_2-1)\int_0^1\beta(c)\left(1-G(c)\right)G\left(\max\left\{\frac{\beta(c)-\kappa}{1-\kappa},0\right\}\right)^{n_1}G(c)^{n_2-2}G'(c)dc\\
&-(1-p)n_1\int_0^1(1-\kappa)c\left(1-G\left(\sup\beta^{-1}[0,(1-\kappa)c)\right)\right)G(c)^{n_1-1}G\left(\sup\beta^{-1}[0,(1-\kappa)c)\right)^{n_2-1}G'(c)dc\\
&-(1-p)(n_2-1)\int_0^1\beta(c)\left(1-G(c)\right)G\left(\min\left\{\frac{\beta(c)}{1-\kappa},1\right\}\right)^{n_1}G\left(c\right)^{n_2-2}G'(c)dc,
\end{aligned}
\tag{7}
$$

$$
D_{\text{PI}}^{\text{CV}}=\frac{V_{\text{PI}}^{\text{CV}}-W_{\text{PI}}^{\text{CV}}}{n_1+n_2}=D_{\text{FI}}^{\text{CV}}. \tag{8}
$$

Informed advertisers' payoff:

$$
E_{\text{FI}}^{\text{CV}}=\frac{V_{\text{FI}}^{\text{CV}}-W_{\text{FI}}^{\text{CV}}}{n_1+n_2}=D_{\text{FI}}^{\text{CV}}=\frac{V_{\text{PI}}^{\text{CV}}-W_{\text{PI}}^{\text{CV}}}{n_1+n_2}=E_{\text{PI}}^{\text{CV}}, \tag{9}
$$

$$
\begin{aligned}
E_{\text{IA}}^{\text{CV}} =&p\int_0^1(\kappa+(1-\kappa)c)G(c)^{n_1-1}G\left(\sup\beta^{-1}[0,\kappa+(1-\kappa)c)\right)^{n_2}G'(c)dc\\
&+(1-p)\int_0^1(1-\kappa)cG(c)^{n_1-1}G\left(\sup\beta^{-1}[0,(1-\kappa)c)\right)^{n_2}G'(c)dc\\
&-pn_2\int_0^1\beta(c)\left(1-G\left(\max\left\{\frac{\beta(c)-\kappa}{1-\kappa},0\right\}\right)\right)G\left(\max\left\{\frac{\beta(c)-\kappa}{1-\kappa},0\right\}\right)^{n_1-1}G(c)^{n_2-1}G'(c)dc\\
&-pn(n_1-1)\int_0^1(\kappa+(1-\kappa)c)\left(1-G(c)\right)G(c)^{n_1-2}G\left(\sup\beta^{-1}[0,\kappa+(1-\kappa)c)\right)^{n_2}G'(c)dc\\
&-(1-p)n_2\int_0^1\beta(c)\left(1-G\left(\min\left\{\frac{\beta(c)}{1-\kappa},1\right\}\right)\right)G\left(\min\left\{\frac{\beta(c)}{1-\kappa},1\right\}\right)^{n_1-1}G(c)^{n_2-1}G'(c)dc\\
&-(1-p)(n_1-1)\int_0^1(1-\kappa)c\left(1-G(c)\right)G(c)^{n_1-2}G\left(\sup\beta^{-1}[0,(1-\kappa)c)\right)^{n_2}G'(c)dc.
\end{aligned}
\tag{10}
$$

Independent-values case

In the independent-values case, all advertisers bid truthfully. An advertiser with a c-value c but without b-data would bid the expected value $\kappa p+(1-\kappa)c$. An advertiser with full information would bid $v=\kappa b+(1-\kappa)c$ where v is distributed by $\tilde{G}(v):=\mathbb{P}[v'\leq v]=pG\left(\max\left\{\frac{v-\kappa}{1-\kappa},0\right\}\right)+(1-p)G\left(\min\left\{\frac{v}{1-\kappa},1\right\}\right)$.

Advertisers' conversion rate:

The advertisers' conversion rates under each information setting are given by:

$$V_{\mathrm{FI}}^{\mathrm{IV}} = (n_1+n_2)\int_0^1 v\tilde{G}(v)^{n_1+n_2-1}\tilde{G}'(v)dv,$$

$$\begin{aligned} V_{\mathrm{IA}}^{\mathrm{IV}} = n_1 &\int_0^1 v\tilde{G}(v)^{n_1-1} G\left(\min\left\{\max\left\{\frac{v-\kappa p}{1-\kappa},0\right\},1\right\}\right)^{n_2}\tilde{G}'(v)dv \\ &+ (1-p)n_2\int_0^1 (1-\kappa)c\tilde{G}(\kappa p+(1-\kappa)c)^{n_1}G(c)^{n_2-1}G'(c)dc \\ &+ pn_2\int_0^1 (\kappa+(1-\kappa)c)\tilde{G}(\kappa p+(1-\kappa)c)^{n_1}G(c)^{n_2-1}G'(c)dc, \end{aligned}$$

$$V_{\mathrm{PI}}^{\mathrm{IV}} = (n_1+n_2)\int_0^1 (\kappa p+(1-\kappa)c)G(c)^{n_1+n_2-1}G'(c)dc.$$

Publisher's expected revenue:

The publisher's expected revenues for each information setting are given by:

$$W_{\mathrm{FI}}^{\mathrm{IV}} = (n_1+n_2)(n_1+n_2-1)\int_0^1 v\left(1-\tilde{G}(v)\right)\tilde{G}(v)^{n_1+n_2-2}\tilde{G}'(v)dv,$$

$$\begin{aligned} W_{\mathrm{IA}}^{\mathrm{IV}} =& n_1(n_1-1)\int_0^1 v\left(1-\tilde{G}(v)\right)\tilde{G}(v)^{n_1-2}G\left(\min\left\{\max\left\{\frac{v-\kappa p}{1-\kappa},0\right\},1\right\}\right)^{n_2}\tilde{G}'(v)dv \\ &+ n_1n_2\int_0^1 v\left(1-G\left(\min\left\{\max\left\{\frac{v-\kappa p}{1-\kappa},0\right\},1\right\}\right)\right)\tilde{G}(v)^{n_1-1}G\left(\min\left\{\max\left\{\frac{v-\kappa p}{1-\kappa},0\right\},1\right\}\right)^{n_2-1}\tilde{G}'(v)dv \\ &+ n_1n_2\int_0^1 (\kappa p+(1-\kappa)c)\left(1-\tilde{G}(\kappa p+(1-\kappa)c)\right)\tilde{G}(\kappa p+(1-\kappa)c)^{n_1-1}G(c)^{n_2-1}G'(c)dc \\ &+ n_2(n_2-1)\int_0^1 (\kappa p+(1-\kappa)c)\,(1-G(c))\,\tilde{G}\,(\kappa p+(1-\kappa)c)^{n_1}\,G(c)^{n_2-2}G'(c)dc, \end{aligned}$$

$$W_{\mathrm{PI}}^{\mathrm{IV}} = (n_1+n_2)(n_1+n_2-1)\int_0^1 (\kappa p+(1-\kappa)c)(1-G(c))G(c)^{n_1+n_2-2}G'(c)dc.$$

Constrained advertisers' payoff:

$$\begin{aligned} D_{\mathrm{FI}}^{\mathrm{IV}} = \frac{V_{\mathrm{FI}}^{\mathrm{IV}}-W_{\mathrm{FI}}^{\mathrm{IV}}}{n_1+n_2} = &\int_0^1 v\tilde{G}(v)^{n_1+n_2-1}\tilde{G}'(v)dv \\ &-(n_1+n_2-1)\int_0^1 v\left(1-\tilde{G}(v)\right)\tilde{G}(v)^{n_1+n_2-2}\tilde{G}'(v)dv, \end{aligned} \tag{11}$$

$$\begin{aligned} D_{\mathrm{IA}}^{\mathrm{IV}} =& (1-p)\int_0^1 (1-\kappa)c\tilde{G}(\kappa p+(1-\kappa)c)^{n_1}G(c)^{n_2-1}G'(c)dc \\ &+ p\int_0^1 (\kappa+(1-\kappa)c)\tilde{G}(\kappa p+(1-\kappa)c)^{n_1}G(c)^{n_2-1}G'(c)dc \end{aligned}$$

$$- n_1 \int_0^1 v \left(1 - G\left(\min\left\{\max\left\{\frac{v-\kappa p}{1-\kappa}, 0\right\}, 1\right\}\right)\right) \tilde{G}(v)^{n_1-1} G\left(\min\left\{\max\left\{\frac{v-\kappa p}{1-\kappa}, 0\right\}, 1\right\}\right)^{n_2-1} \tilde{G}'(v)dv$$
$$- (n_2 - 1) \int_0^1 (\kappa p + (1-\kappa)c)\,(1 - G(c))\, \tilde{G}\,(\kappa p + (1-\kappa)c)^{n_1}\, G(c)^{n_2-2} G'(c)dc, \quad (12)$$

$$D_{\text{PI}}^{\text{IV}} = \frac{V_{\text{PI}}^{\text{IV}} - W_{\text{PI}}^{\text{IV}}}{n_1 + n_2} = \int_0^1 (\kappa p + (1-\kappa)c) G(c)^{n_1+n_2-1} G'(c)dc$$
$$- (n_1 + n_2 - 1) \int_0^1 (\kappa p + (1-\kappa)c)(1 - G(c)) G(c)^{n_1+n_2-2} G'(c)dc.$$

Informed advertisers' payoff:

$$E_{\text{FI}}^{\text{IV}} = \frac{V_{\text{FI}}^{\text{IV}} - W_{\text{FI}}^{\text{IV}}}{n_1 + n_2} = D_{\text{FI}}^{\text{IV}}, \qquad E_{\text{PI}}^{\text{IV}} = \frac{V_{\text{PI}}^{\text{IV}} - W_{\text{PI}}^{\text{IV}}}{n_1 + n_2} = D_{\text{PI}}^{\text{IV}}, \quad (13)$$

$$E_{\text{IA}}^{\text{IV}} = \int_0^1 v \tilde{G}(v)^{n_1-1} G\left(\min\left\{\max\left\{\frac{v-\kappa p}{1-\kappa}, 0\right\}, 1\right\}\right)^{n_2} \tilde{G}'(v)dv$$
$$- (n_1 - 1) \int_0^1 v \left(1 - \tilde{G}(v)\right) \tilde{G}(v)^{n_1-2} G\left(\min\left\{\max\left\{\frac{v-\kappa p}{1-\kappa}, 0\right\}, 1\right\}\right)^{n_2} \tilde{G}'(v)dv$$
$$- n_2 \int_0^1 (\kappa p + (1-\kappa)c) \left(1 - \tilde{G}(\kappa p + (1-\kappa)c)\right) \tilde{G}(\kappa p + (1-\kappa)c)^{n_1-1} G(c)^{n_2-1} G'(c)dc.$$
$$(14)$$

B.3 First-Price Auction

A commonly used auction format in advertising auctions nowadays is the first-price auction format (see e.g., Despotakis et al., 2021). In this section, we modify our model to a first-price auction instead of a second-price auction for selling the impression, to test the robustness of our main findings.

In a first-price auction, since advertisers pay their own bid if they win, constrained advertisers who do not have complete information about their valuations, are even more likely to bid conservatively (underbid) compared to a second-price auction. As a result, information asymmetry can lead to both a lower conversion rate and reduced publisher revenue, compared to the symmetric information settings, for similar reasons this happens in second-price auctions. In other words, our result that restricting data access can simultaneously increase both the conversion rate and publisher revenue remains valid in common-value first-price auctions. Propositions 11 and 12 below confirm this.

Proposition 11. *For any distribution G, any $n_1, n_2 \geq 1$, $\kappa \in [0,1]$, under the common-value case, and when the auction format is first-price, we have that $V_{\text{IA}}^{\text{CV}} \leq V_{\text{FI}}^{\text{CV}} = V_{\text{PI}}^{\text{CV}}$.*

Proof. This result follows from Lemma 4, where the mechanism M is a first-price auction. □

Proposition 12. *For any distribution G, $n_1 > 1$, $n_2 \geq 1$, $\kappa \geq 1/2$, and sufficiently low p, under the common-value case, and when the auction format is first-price, we have that $W_{\text{IA}}^{\text{CV}} \leq W_{\text{FI}}^{\text{CV}} = W_{\text{PI}}^{\text{CV}}$.*

Proof. Suppose that, when there is information asymmetry, at equilibrium the constrained and informed advertisers use the bidding functions $\beta_D : [0,1] \to [0,1]$ and $\beta_E : [0,1]^2 \to [0,1]$, respectively. The expected utility of a constrained advertiser with c-value c who bids $\tilde{\beta}$ is

$$\begin{aligned} u_D(\tilde{\beta}; \beta_D, \beta_E, c) =& p \left(\kappa + (1-\kappa)c - \tilde{\beta}\right) G\left(\sup \beta_E^{-1}([0, \tilde{\beta}), b=1)\right)^{n_1} G(\sup \beta_D^{-1}([0,\tilde{\beta})))^{n_2-1} \\ &+ (1-p)\left((1-\kappa)c - \tilde{\beta}\right) G\left(\sup \beta_E^{-1}([0,\tilde{\beta}), b=0)\right)^{n_1} G(\sup \beta_D^{-1}([0,\tilde{\beta})))^{n_2-1}. \end{aligned}$$

The expected utility of an informed advertiser with c-value c and b-value b who bids $\tilde{\beta}$ is $u_E(\tilde{\beta}; \beta_D, \beta_E, c, b) = \left(\kappa b + (1-\kappa)c - \tilde{\beta}\right) G(\sup \beta_E^{-1}([0,\tilde{\beta}), b))^{n_1-1} G(\sup \beta_D^{-1}([0,\tilde{\beta})))^{n_2}$.

The expected publisher revenue is

$$\begin{aligned} W_{\text{IA}}^{\text{CV}} =& p n_2 \int_0^1 \beta_D(c) G(\sup \beta_E^{-1}([0, \beta_D(c)), 1))^{n_1} G(c)^{n_2-1} G'(c) dc \\ &+ p n_1 \int_0^1 \beta_E(c,1) G(c)^{n_1-1} G(\sup \beta_D^{-1}([0, \beta_E(c,1))))^{n_2} G'(c) dc \\ &+ (1-p) n_2 \int_0^1 \beta_D(c) G(\sup \beta_E^{-1}([0, \beta_D(c)), 0))^{n_1} G(c)^{n_2-1} G'(c) dc \\ &+ (1-p) n_1 \int_0^1 \beta_E(c,0) G(c)^{n_1-1} G(\sup \beta_D^{-1}([0, \beta_E(c,0))))^{n_2} G'(c) dc. \end{aligned}$$

For the other two information settings, $W_{\text{FI}}^{\text{CV}}$ and $W_{\text{PI}}^{\text{CV}}$, the revenue is identical to the second-price auction case by the revenue equivalence principle.

The conversion rate is

$$\begin{aligned} V_{\text{IA}}^{\text{CV}} =& p n_2 \int_0^1 (\kappa + (1-\kappa)c) G\left(\sup \beta_E^{-1}([0, \beta_D(c)), 1)\right)^{n_1} G(c)^{n_2-1} G'(c) dc \\ &+ p n_1 \int_0^1 (\kappa + (1-\kappa)c) G(c)^{n_1-1} G\left(\sup \beta_D^{-1}[0, \beta_E(c,1))\right)^{n_2} G'(c) dc \end{aligned}$$

$$+ (1-p)n_2 \int_0^1 (1-\kappa) c G\left(\sup \beta_E^{-1}([0, \beta_D(c)), 0)\right)^{n_1} G(c)^{n_2-1} G'(c) dc$$
$$+ (1-p)n_1 \int_0^1 (1-\kappa) c G(c)^{n_1-1} G\left(\sup \beta_D^{-1}[0, \beta_E(c, 0))\right)^{n_2} G'(c) dc.$$

If p is sufficiently low,[12] the constrained advertisers will choose to not compete with the informed advertisers if $b = 1$ and will always bid as if $b = 0$ regardless of the actual value of b (which they do not know anyway). In this case, when $b = 0$ the game reduces to a symmetric first-price auction among all $n_1 + n_2$ advertisers, and when $b = 1$ the game reduces to a symmetric first-price auction among n_1 informed advertisers. Consequently, we have a symmetric equilibrium where the bidding functions have simple analytical closed forms as follows:

$$\beta_D(c) = (1-\kappa)\left(c - \int_0^c \left(\frac{G(t)}{G(c)}\right)^{n_1+n_2-1} dt\right),$$
$$\beta_E(c, b) = \kappa b + (1-\kappa)\left(c - \int_0^c \left(\frac{G(t)}{G(c)}\right)^{n_1+(1-b)n_2-1} dt\right). \tag{15}$$

Under the perfect-information and the partial-information settings, the bidders are symmetric and independent, thus it follows from the revenue equivalence principle that $W_{\text{IA}}^{\text{CV}}$ and $W_{\text{FI}}^{\text{CV}}$ are both equal to their second-price auction counterparts, hence they are equal to each other. We also note from (15) that $\beta_D(c) \leq \beta_E(c, b) \leq \kappa b + (1-\kappa)\left(c - \int_0^c \left(\frac{G(t)}{G(c)}\right)^{n_1+n_2-1} dt\right)$, for all c, b, where notice that the RHS is the bidding function under the perfect-information setting. The first inequality holds since $\beta_E(c, 1) - \beta_D(c) \geq \kappa - (1-\kappa) \geq 0$ because $\kappa \geq 1/2$, and the second inequality holds since $\left(\frac{G(t)}{G(c)}\right)^{n_1+(1-b)n_2-1} \geq \left(\frac{G(t)}{G(c)}\right)^{n_1+n_2-1}$ for all $t, c \in [0, 1]$ such that $t \leq c$. In other words, the bids of all advertisers are at least as high in the perfect-information setting as under information asymmetry, hence the revenue $W_{\text{FI}}^{\text{CV}}$ is at least as high as $W_{\text{IA}}^{\text{CV}}$. □

C Common c-Value

In this section, we consider the case where the random variables c_i are not independent, but they are the same for all advertisers, i.e. $c_1 = c_2 = \ldots = c_n =: c$, where c is drawn from distribution G. As in the main model, we consider the following sub-cases. Let $\mathbb{E}[c] := \int_0^1 c dG(c)$ in the following.

[12] E.g. if $\kappa p + (1-\kappa) < \kappa \implies p < 2 - 1/\kappa$, which is a sufficient but not necessary bound.

- Common b-values (CV): All the advertisers have the same valuation $\kappa b + (1-\kappa)c$.
  - Full Information (FI): All advertisers bid $\kappa b + (1-\kappa)c$. For $n_1, n_2 \geq 1$: $V_{\text{FI}}^{\text{CV}} = (\kappa + (1-\kappa)\mathbb{E}[c]) \cdot p + (1-\kappa)\mathbb{E}[c] \cdot (1-p) = \kappa p + (1-\kappa)\mathbb{E}[c]$.
  - Information Asymmetry (IA): Informed advertisers bid $\kappa b + (1-\kappa)c$. If there is at least one informed advertiser, constrained advertisers bid $(1-\kappa)c$. If there is no informed advertiser, constrained advertisers bid $\kappa p + (1-\kappa)c$. For $n_1, n_2 \geq 1$: $V_{\text{IA}}^{\text{CV}} = (\kappa + (1-\kappa)\mathbb{E}[c]) \cdot p + (1-\kappa)\mathbb{E}[c] \cdot (1-p) = \kappa p + (1-\kappa)\mathbb{E}[c]$.
  - Partial Information (PI): All advertisers bid $\kappa p + (1-\kappa)c$. For $n_1, n_2 \geq 1$: $V_{\text{PI}}^{\text{CV}} = \kappa p + (1-\kappa)\mathbb{E}[c]$.

  We can see that $V_{\text{FI}}^{\text{CV}} = V_{\text{IA}}^{\text{CV}} = V_{\text{PI}}^{\text{CV}}$.

- Independent b-values (IV): Advertisers have valuations $\kappa b_i + (1-\kappa)c$.
  - Full Information (FI): All advertisers bid their valuation $\kappa b_i + (1-\kappa)c$. For $n_1, n_2 \geq 1$:

$$\begin{aligned} V_{\text{FI}}^{\text{IV}} &= (\kappa + (1-\kappa)\mathbb{E}[c]) \cdot \left(1 - (1-p)^{n_1+n_2}\right) + (1-\kappa)\mathbb{E}[c] \cdot (1-p)^{n_1+n_2} \\ &= \kappa \left(1 - (1-p)^{n_1+n_2}\right) + (1-\kappa)\mathbb{E}[c]. \end{aligned}$$

  - Information Asymmetry (IA): Informed advertisers bid their valuation $\kappa b_i + (1-\kappa)c$. Constrained advertisers bid $\kappa p + (1-\kappa)c$. For $n_1, n_2 \geq 1$:

$$\begin{aligned} V_{\text{IA}}^{\text{IV}} &= (\kappa + (1-\kappa)\mathbb{E}[c]) \cdot (1 - (1-p)^{n_1}) + (\kappa p + (1-\kappa)\mathbb{E}[c]) \cdot (1-p)^{n_1} \\ &= \kappa \left(1 - (1-p)^{n_1+1}\right) + (1-\kappa)\mathbb{E}[c]. \end{aligned}$$

  - Partial Information (PI): All advertisers bid $\kappa p + (1-\kappa)c$. For $n_1, n_2 \geq 1$: $V_{\text{PI}}^{\text{IV}} = \kappa p + (1-\kappa)\mathbb{E}[c]$.

  Since $1 - (1-p)^n$ is an increasing function in $n \geq 0$, it follows that $V_{\text{FI}}^{\text{CV}} \geq V_{\text{IA}}^{\text{CV}} \geq V_{\text{PI}}^{\text{CV}}$.