

HOW DO DIGITAL ADVERTISING AUCTIONS IMPACT
PRODUCT PRICES?

By

Dirk Bergemann, Alessandro Bonatti, and Nicholas Wu

July 2023

COWLES FOUNDATION DISCUSSION PAPER NO. 2366



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

How Do Digital Advertising Auctions Impact Product Prices?*

Dirk Bergemann[†] Alessandro Bonatti[‡] Nicholas Wu[§]

July 26, 2023

Abstract

We ask how the advertising mechanisms of digital platforms impact product prices. We present a model that integrates three fundamental features of digital advertising markets: (i) advertisers can reach customers on and off-platform, (ii) additional data enhances the value of matching advertisers and consumers, and (iii) bidding follows auction-like mechanisms. We compare data-augmented auctions, which leverage the platform’s data advantage to improve match quality, with managed campaign mechanisms, where advertisers’ budgets are transformed into personalized matches and prices through auto-bidding algorithms.

In data-augmented second-price auctions, advertisers increase off-platform product prices to boost their competitiveness on-platform. This leads to socially efficient allocations on-platform, but inefficient allocations off-platform due to high product prices. The platform-optimal mechanism is a sophisticated managed campaign that conditions on-platform prices for sponsored products on off-platform prices set by all advertisers. Relative to auctions, the optimal managed campaign raises off-platform product prices and further reduces consumer surplus.

Keywords: Data, Advertising, Competition, Digital Platforms, Auctions, Automated Bidding, Managed Advertising Campaigns, Matching, Price Discrimination.

JEL Codes: D44, D82, D83.

*We acknowledge financial support through NSF Grant SES-1948692, the Omidyar Network, and the Sloan Foundation. An earlier version of this paper circulated under the title “Managed Campaigns and Data-Augmented Auctions for Digital Advertising,” and an extended abstract appears in the Proceedings of the 24th ACM Conference on Economics and Computation (EC ’23).

[†]Department of Economics, Yale University, New Haven, CT 06511, dirk.bergemann@yale.edu

[‡]Sloan School of Management, MIT, Cambridge, MA 02142, bonatti@mit.edu

[§]Department of Economics, Yale University, New Haven, CT 06511, nick.wu@yale.edu

1 Introduction

1.1 Motivation and Results

Digital advertising facilitates the matching of consumers and advertisers online. Large platforms utilize their extensive consumer data to connect online shoppers with their preferred products and brands. In turn, advertisers join these platforms in order to target a wider range of potential consumers beyond their existing customer base.

Large digital platforms also enjoy significant market power, which has raised regulatory concerns. In a recent report, the UK Competition & Markets Authority argues:

“Where an advertising platform has market power [...] advertiser bids in its auctions are higher, resulting in higher prices. In addition, the platforms may be able to use levers including the use of reserve prices or mechanisms such as automated bidding to extract more rent from advertisers. [...]

Higher advertising prices matter because they represent increased costs to the firms producing goods and services which are purchased by consumers. We would expect these costs to be passed through to consumers in terms of higher prices for goods and services, even if the downstream market is highly competitive.”

(See Competition & Markets Authority (2020) 6.19, 6.20, p 314.)

Indeed, digital platforms typically organize the competition for attention among advertisers through an auction-based allocation mechanism. As the market for digital advertising has grown and become more complex, digital platforms and other advertising intermediaries often implement bidding on behalf of advertisers. These intermediaries run managed campaigns for advertisers, choosing how to bid across numerous opportunities to create matches. These managed campaigns are implemented through auto-bidding algorithms that bid on behalf of the advertisers with certain objectives and relevant constraints explicitly stated.¹

In this paper, we explore how the mechanisms by which large digital platforms sell access to consumers affect product prices both on and off the platform. We take into account three fundamental aspects of digital advertising. First, advertisers can reach their customers on and off the platforms. Second, platforms possess valuable data that can enhance matching efficiency. Third, on-platform matching between viewers and advertisers is governed by

¹A recent literature has developed around autobidding algorithms when the bidders formulate their objective outside of the class of quasilinear utility models in mechanism design. For example, the bidder may seek to maximize return on investments and have budget or spending constraints. Aggarwal et al. (2019), Balseiro et al. (2021), and Deng et al. (2021) offer excellent introductions this rapidly growing research area.

bidding mechanisms. Our model considers a monopolistic digital platform that sells access to consumers. Advertisers determine their pricing strategy on and off the platform and their advertising strategies on the platform.

With access to the platform’s data, advertisers can offer prices that reflect consumers’ willingness to pay. This type of third-degree price discrimination broadens the market and enhances the efficiency of matching on the platform. Thus, the advertiser is using the additional information to reach more shoppers and improve the matches formed on the platform. Off the platform, advertisers lack additional data and offer a uniform price. At this price, those customers with values above the product price will receive an information rent, while the others will be priced out of the market.

On the platform, consumers act as *shoppers* and choose the product that offers the highest net value. As these consumers compare the advertised offers to all firms’ off-platform prices, advertisers face a “showrooming” constraint: they must ensure their on-platform offers are at least as attractive as their off-platform offers. Conversely, consumers off the platform are *loyal* and buy from a single brand. Consequently, advertisers face a trade-off between setting optimal prices for their loyal customers off-platform and charging higher personalized prices to on-platform shoppers.²

Our focus is on how data-intensive mechanisms for selling advertising space affect welfare both on and off the platform. Our results suggest that any analysis of pass-through of online advertising costs must account for cross-channel distortions. Indeed, we show that advertisers raise prices off the platform to gain a competitive edge on the platform. In particular, under the platform-optimal mechanism, the higher costs of online advertising are passed on to consumers by means of higher product prices *off* rather than *on* the platform.

We begin our analysis with a second-price auction for a single advertising slot where the platform augments the bidders’ information and solicits bids based on estimated match values between consumers and advertisers. We refer to this as *data-augmented bidding*. Each advertiser can offer a bid for the slot and a price at which to offer the product associated with the advertising slot. Additionally, each advertiser must decide on the price at which to offer their product to *loyal customers* off the platform.

We derive the optimal bidding and pricing strategy of the advertisers (Theorem 1). On

²In pure advertising platforms, where the matching fee is typically incurred before the transaction (e.g., through pay-per-impression or pay-per-click fees), the advertiser directly faces the showrooming constraint. The advertiser wants to pay for the listing only if it leads to a sale, as the offline transaction could have occurred without the advertising. In other platforms where the fee is based on transactions, such as referral fees on shopping services like Amazon, the platform often imposes the showrooming constraint through a most favored nation clause. This clause requires the advertiser to offer the most favorable price online.

the platform, the second-price auction implements an efficient allocation, and the additional data allows the advertisers to sell successfully to consumers with lower values without the need to price them out of the market (Proposition 1). Off the platform, the advertisers raise their price to their loyal customers relative to the price they would have charged in a stand-alone market (Proposition 2). By offering their product only at a higher price, each advertiser can weaken the showrooming constraint and extract more surplus on the platform. The off-platform prices increase with the number of on-platform shoppers (Proposition 3) and decrease with the number of bidders (Proposition 4). Finally, as bidders are homogeneous ex-ante, the platform can impose a participation fee that extracts all their surplus without affecting the subsequent prices and bids (Proposition 5).

Next, we introduce the concept of a *managed campaign*. In this more centralized mechanism, the platform proposes to each advertiser an advertising budget, an autobidding algorithm, and a pricing function for the product on the platform. The autobidding algorithm governs how the managed campaign for every advertiser bids for potential matches as a function of the value of the match and the off-platform prices. Each advertiser simultaneously decides whether to enter into the managed campaign or not, and how to price its product off the platform.³ We show that the platform optimizes its revenue by offering best-value pricing; that is, the platform implements the efficient allocation but ensures that the efficient firm always makes the offer with the best value to the consumer (Theorem 2). However, in doing so, the platform weakens competition, and so the firms raise their posted prices off the platform in order to extract more surplus from online consumers. We show that best-value pricing is revenue-optimal for the platform (Theorem 3), and in fact, joint producer surplus attains the vertical integration benchmark where one firm controls all the advertisers and the platform (Corollary 1). In consequence, the posted prices off-platform are higher than under the data-augmented auction (Theorem 4).

Finally, we deploy our model to assess how policy interventions affect the platform, firms, and consumers. We examine two interventions. First, we consider restricting the platform’s pricing policy to be *independent* of the posted prices of competing firms. Here, we restrict the platform to price based on the consumer’s value for the advertiser’s own product only, and not on the consumer’s value for other advertisers’ products. The optimal independent managed campaign mechanism implements an efficient allocation of advertising

³The autobidding algorithm that allocates budgets can be interpreted as maximizing profit subject to a return on investment constraint. Alternatively, we can decompose the advertising budget into a payment per winning bid for each consumer value. In this case, one can show that the bidding algorithm boosts the bids of the advertisers, but never beyond the value of the match. Thus, the autobidding mechanism satisfies an ex-post participation constraint for every (winning and losing) bid.

slots (Proposition 6), but relative to either a second-price auction or even a revenue optimal auction design, it differs along a number of dimensions that have a substantial impact on the outcomes (Theorem 5). First, by charging the bidders up front for expected matches, the digital platform can capture a larger share of the surplus yet do so without hurting the efficiency of the allocation. Under a sufficient condition on the value distribution, the off-platform posted prices are equal to the stand-alone monopoly prices (Proposition 7). An independent managed campaign can lead to a more efficient outcome both on- and off-platform. We show that under a relatively mild condition, posted prices are lowest with an independent managed campaign, followed by data-augmented bidding (Proposition 8).

The second policy we examine is a privacy restriction, which prevents the platform from steering consumers and setting prices on the basis of consumers’ detailed data. Instead, we allow the platform to condition its steer and pricing decision on the basis of coarse information only, i.e., on the identity of each consumer’s favorite firm. This restriction is equivalent to removing the ability to price discriminate using the platform’s data. In this scenario, the firms sell to both on- and off-platform consumers via the same posted price, which we characterize in Proposition 10. The privacy restriction reduces off-platform prices compared to the benchmark of a sophisticated managed campaign, but not necessarily relative to data-augmented bidding. However, the welfare implications of privacy restrictions depend on the size of the platform. In summary, privacy restrictions lead to lower off-platform prices but may affect efficiency on the platform, particularly in the case of large platforms.

1.2 Related Literature

In our digital advertising model, each advertiser has a parallel sales channel available off the platform and faces two segments of consumers, *shoppers* on the platform and *loyal customers* off the platform, as in Varian (1980). The design of the auction is therefore subject to competition from a separate and distinct market. Earlier papers referred to mechanism design subject to alternative markets as “partial mechanism design,” or “mechanism design with a competitive fringe,” e.g., Philippon and Skreta (2012), Tirole (2012), Calzolari and Denicolò (2015), and Fuchs and Skrzypacz (2015). In these papers, the platform is limited in its ability to monopolize the market since the firms have access to an outside option. We focus on digital advertising through auctions rather than competition for the consumer between on and off-platform firms. Varian (2022) analyzes the relationship between advertising costs and product prices through the lens of a single (representative) online merchant. The size of the advertising audience increases sales proportionally at every price level, with a convex cost

of increasing the audience size. In this simple and separable model, an exogenous increase in advertising costs does not necessarily lead to an increase in product prices.

Our paper also contributes to the literature on online ad auctions. Recent works have studied learning in repeated auctions (Balseiro and Gur, 2019; Kanoria and Nazerzadeh, 2020; Nedelec et al., 2022), discriminatory effects (Celis et al., 2019; Ali et al., 2019; Nasr and Tschantz, 2020), and collusion (Decarolis et al., 2020, 2022). Our focus, instead, lies in comparing the effects of an auction to other allocation mechanisms in the presence of off-platform markets in a static setting with a fixed information structure. As such, our approach is closely related to, yet distinct from, Bar-Isaac and Shelegia (2022), who compare auctions and auto-bidding mechanisms in a single marketplace under exogenous limits to the ability to steer and to price discriminate.

A key innovation in our model is that the platform actively manages the firms' advertising campaigns. Managed campaigns have become the predominant mode of selling advertisements in real-world digital markets, where advertisers set a fixed budget, specify high-level objectives for their campaigns, and leave the task of bidding to "autobidders" offered by the platform. Several recent papers have focused on auction design in the presence of autobidders (Liaw et al. (2022); Mehta (2022); Deng et al. (2022)) and return-on-investment constraints (Golrezaei et al. (2021)). Our setting adds an additional dimension related to display prices: advertisers submit both bids for the sponsored link and tailored prices to offer consumers. While Li and Lei (2023) also investigate mechanisms that allow for these display prices, we further explore the impact of activity off the platform on allocations as well as pricing.

Our paper also relates to the information design literature. In particular, Bergemann et al. (2015), Haghpanah and Siegel (2022), and Elliott et al. (2022) study the effect of market segmentations and the achievable combinations of consumer and producer surplus, i.e., how to use data to make markets more or less competitive.

Finally, the showrooming constraint is related to a growing literature on digital platforms with competing advertisers or multiple sales channels. Recent contributions on these topics include de Cornière and de Nijs (2016), Bar-Isaac and Shelegia (2020), Miklós-Thal and Shaffer (2021), and Wang and Wright (2020). However, our model differs in that advertisers in our setting are concerned about showrooming because selling on the platform can be more profitable, thanks to the added value of making data-augmented offers. In parallel work, Bergemann and Bonatti (2022) study on- and off-platform competition with multi-product firms and associated nonlinear pricing. Their focus is on the implications of

managed campaigns for the equilibrium product quality, relative to our paper’s exploration of showrooming and its impact on pricing strategies in the presence of off-platform markets.

2 Model

Payoffs and Information There are J firms indexed $1, 2, \dots, J$, each selling unique indivisible products, and a single digital platform. Each firm has zero production costs. There is a unit mass of consumers, each demanding a single product. Willingness to pay v_j for each firm j ’s product is drawn independently across consumers and firms according to a distribution function F with support on $V = [0, 1]$. We assume F admits a log-concave density f on its support.⁴ The vector of the willingness-to-pay is the consumer’s *value*

$$v = (v_1, \dots, v_J) \in V^J = [0, 1]^J.$$

The utility for consumer v of purchasing product j at price p_j is

$$v_j - p_j.$$

Initially, values are observed by the consumers and by the platform, but not by the firms. Because the consumers and the platform share the same information, we are implicitly assuming that the platform has already learned everything about consumer preferences.

The symmetry in the information is helpful for the welfare comparison but is clearly a stark assumption. The equilibrium implications are robust to a more general formulation in which the platform is endowed with partial and potentially endogenous information.

Platform A measure $\lambda \in [0, 1]$ of consumers uses the platform. The platform presents on-platform consumers with a single “sponsored” result first, followed by organic search results, i.e., a list of non-sponsored products. The platform sells the sponsored position using either a second-price auction or a managed campaign. Under either mechanism, the firm in the sponsored slot can condition its price on the consumer’s value.

Let v denote the full J -dimensional consumer value. An on-platform consumer with value v will see a sponsored offer, which offers some firm j ’s product at some price. In the remaining sections of the paper, we discuss mechanisms for the platform to determine which

⁴This is a technical assumption that ensures that first-order conditions for maximization problems we consider later are well-defined.

firm’s offer gets shown to on-platform consumers, and at what price. Note the platform does not ex-ante commit to steer consumers efficiently; that is, value v does not have to see a sponsored offer from j such that $j = \arg \max_i v_i$.

Firms and Showrooming In addition to the on-platform prices $p_j(v)$ displayed in the sponsored slot, each firm j posts a price \bar{p}_j for its product off the platform. We use the upper-bar notation here since the showrooming constraint implies that the posted price \bar{p}_j is an upper bound on the amount that any consumer will pay for j ’s good. Indeed, each firm j is subject to a *showrooming* constraint: for all v , the prices it advertises on the platform must satisfy $p_j(v) \leq \bar{p}_j$. One interpretation of this constraint is that on-platform consumers can search for free at any off-platform website or store. Alternatively, the platform may impose most-favored-nation clauses requiring firms to offer their lowest prices on the platform.

On-platform Consumers The on-platform consumers observe their willingness to pay v , the “sponsored” offer $p_j(v)$ for the firm j that wins the sponsored slot auction, and the posted prices \bar{p}_k for all firms k . Equivalently, we can interpret the model as allowing for free online search; that is, only a “sponsored” firm can target a price offer to an online consumer, but the online consumer can search and find the posted prices of all firms, including those that did not make the sponsored offer.

Off-platform Consumers We assume that the remaining $1 - \lambda$ mass of consumers are loyal, and visit only a single firm’s non-platform store (e.g., physical store, store website). Thus the off-platform consumer population is divided into J segments of size $(1 - \lambda)/J$, where the j th segment shops directly from firm j . Firm j is the only firm in the consideration set of the j th segment of off-platform consumers. The off-platform consumers view the off-platform price of the single firm in their consideration set, and choose to buy if and only if the off-platform price is lower than their willingness to pay.

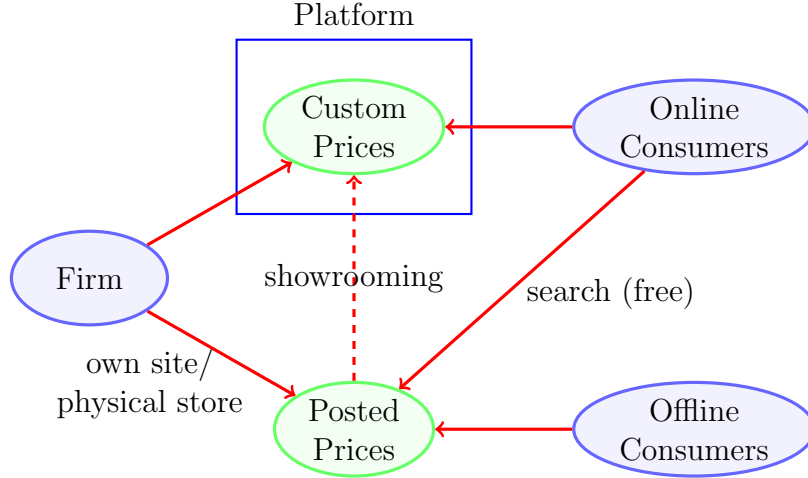


FIGURE 1: Depiction of the model.

3 Data-Augmented Bidding

In this section, we characterize the symmetric Bayesian Nash equilibrium of the bidding and pricing game among the advertisers. Each firm j submits a bid function $b_j : V^J \rightarrow \mathbb{R}_+$ and a sponsored price function $p_j : V^J \rightarrow \mathbb{R}_+$, in addition to (simultaneously) posting a price \bar{p}_j . We refer to this as *data-augmented bidding*, because the platform’s proprietary data enables the advertisers to condition bids and sponsored prices on the consumer’s full value vector v .

Let us first discuss some of the economic intuition for how the presence of the platform impacts the prices posted by the firms before presenting the formal analysis. Recall that the off-platform consumers are loyal, and so in the absence of a platform, all firms post the monopoly price for their market segment. Adding the platform and the on-platform consumers has two contrasting effects on the prices posted by the firms. The first effect is upward pressure due to the increased ability to price discriminate; i.e. since the posted price sets an upper bound on the prices that a firm can offer to on-platform consumers, the potential to price discriminate more effectively on-platform pushes firms to raise their posted prices. However, there is an opposite effect, where competition for the on-platform consumers introduces an incentive to lower the posted prices—the ability to undercut its competitors by advertising a lower off-platform price, in order to win more on-platform consumers.

3.1 Bidding Equilibrium

The following result helps characterize the equilibrium strategies of the firms for this setting. Effectively, the proposition shows that *regardless* of the profile of posted prices set by the

firms, the bidding equilibrium on the platform results in a *symmetric* assignment, where each on-platform consumer sees a sponsored offer from the firm they like best. This implies that the sponsored slot allocation resulting from data-augmented bidding is efficient.

Proposition 1 (Efficient Bidding Outcome)

Fix any vector of posted prices \bar{p} off-platform. Consider an on-platform consumer with value v . If $v_j > v_k$, firm j bids at least as much as firm k for consumer v in any bidding equilibrium.

The proof mostly proceeds by casework, but we provide an intuition here. Suppose a consumer arrives, and the consumer’s favorite firm is firm 1. Consider a competitor, say firm 2. If firm 1 has set a higher posted price than firm 2, then firm 1 has a larger ability to price discriminate than firm 2, and hence the consumer is intuitively worth more to firm 1, thus allowing it to bid more. However, if firm 1 has a lower posted price than firm 2, then firm 2 must concede rent to the consumer because even if 2 won the sponsored slot, the consumer could still search and find the posted price for firm 1; hence, this disciplines firm 2’s bid, and we show that this actually constrains 2’s bid to be lower than 1’s.

Proposition 1 is useful because it allows us to separate the bidding stage from the posted prices; that is, the matches (though not the bids) in the bidding game are invariant with respect to the posted prices. As a consequence of this Proposition, the set of online consumers who purchase from firm j is exactly those for whom $j = \arg \max_i v_i$; that is, the consumers with the highest value for firm j ’s product.

Since we are looking for symmetric equilibria, we suppose all the other firms post price p' and consider the best response problem of a single firm:

$$\max_p \left\{ \frac{1-\lambda}{J} p(1-F(p)) + \lambda \Omega(p; p') \right\}, \quad (1)$$

where

$$\Omega(p; p') = \int \int^v (\min(v - \max(v' - p', 0), p) - (\min(v' - \max(v - p, 0), p'))_+) dF^{J-1}(v') dF(v).$$

This term denotes the expected profit from on-platform consumers that a firm would expect to make by setting a posted price at p when all other firms set a posted price p' . The term integrates over $v' = \max_{j \neq i} v_j$, which is the highest value the consumer has for any other firm besides i . Since the firm must concede utility $\max(v' - p', 0)$ to the threat of the on-platform consumer going to the competitor, the firm setting price p will bid $\min(v - \max(v' - p', 0), p)$. The highest competitor bids $(\min(v' - \max(v - p, 0), p'))_+$, where $(\cdot)_+$ denotes the nonnegative

part. It turns out that with some casework, we can simplify this expression for the on-platform sales further:

Lemma 1 (On-platform Bidding Profit)

The expected on-platform profit satisfy

$$\Omega(p; p') = \int \int^v \min(v - v', p) dF^{J-1}(v') dF(v).$$

Note that v' is only integrated on values less than v . The proof is algebraic and left to the Appendix. The result, however, is quite intuitive. In a standard second-price auction, the expected surplus of a bidder is the expected gap between the bidder value v and the value of the second highest bid v' ; this form shows that with the showrooming constraints and strategic bidding behavior in the presence of the off-platform interaction, the firm profit is $v - v'$, capped by the posted price.

To solve for the symmetric equilibria, we compute the derivative of Ω with respect to p . Long but straightforward algebra yields the following expression:

$$\frac{\partial \Omega(p; p')}{\partial p} = \int_p F^{J-1}(v - p) dF(v). \quad (2)$$

Finally, we can write out the first-order condition for profit maximization using (2):

$$\frac{1 - \lambda}{J} (1 - F(p) - pf(p)) + \lambda \left(\int_p F^{J-1}(v - p) dF(v) \right) = 0.$$

Rearranging this condition, we summarize the equilibrium characterization as follows.

Theorem 1 (Bidding Equilibrium)

In the unique symmetric equilibrium, the posted prices of the firms satisfy

$$p_B = \frac{1 - F(p_B) + \frac{\lambda J}{1 - \lambda} \left(\int_{p_B} F^{J-1}(v - p_B) dF(v) \right)}{f(p_B)}. \quad (3)$$

Firms bid their true value $\max(v_j, p_B)$ for each consumer on-platform. On-platform consumers buy the sponsored offer, and off-platform consumers buy from the firm they are loyal to if and only if the posted price p_B is below their value v_j .

We denote the symmetric equilibrium price for the product off the platform in the presence of the bidding mechanism on the platform by p_B , where we use subscript B as this is the bidding equilibrium.

3.2 Welfare and Comparative Statics

We discuss the efficiency implications of the outcome of data-augmented bidding. Proposition 1 implies that the allocation on-platform is socially efficient; since the sponsored offer is always made by the consumer’s most preferred firm, each on-platform consumer purchases the product they like best. Off-platform consumers face two sources of inefficiency; first, they might be unaware of the existence of a firm that they would prefer, and second, since the firms only sell to off-platform consumers via posted prices, consumers with value for their firm’s product below the posted price will not buy.

To characterize the efficiency implications for the off-platform consumers, we first define the posted price a firm would set if it only had its loyal off-platform population:

$$p_M = \frac{1 - F(p_M)}{f(p_M)}. \quad (4)$$

We term this p_M , as this is analogous to the monopoly price.

Examining the price equations, one can see that since the expression on the right-hand side of (3) is a larger function of the price than in (4), the price p_B is larger than p_M . Note that higher posted prices entail greater welfare loss off-platform than if the platform did not exist. Since higher prices mean fewer sales, lower consumer surplus, and less efficiency, the presence of the on-platform consumers induces firms to price out some off-platform consumers in order to gain sales on the platform. Recall the two effects discussed at the beginning of the section—the incentive to raise posted prices to price discriminate and the incentive to compete for on-platform consumers through posted prices. We show that the former effect dominates firm competition on-platform.

Proposition 2 (Posted Prices)

Data-augmented bidding results in higher posted prices than would occur without the platform, $p_B \geq p_M$. The presence of the platform induces lower consumer surplus, higher posted prices, and lower total welfare off the platform.

We can generalize the insight to show comparative statics with respect to the share λ of consumers that are on the platform, fixing the total measure of consumers to 1. We will first define several welfare objects of interest, as functions of the posted price p . The expected consumer surplus of an off-platform consumer is:

$$CS_{\text{off}}(p) = \int_p^1 (v - p) dF(v).$$

The expected consumer surplus of an on-platform consumer is:

$$CS_{\text{on}}(p) = \int_p^1 (v - p) dF^J(v).$$

Total consumer surplus is then:

$$CS(p) = (1 - \lambda)CS_{\text{off}}(p) + \lambda CS_{\text{on}}(p).$$

Because F^J describes the distribution of the highest-order statistic, the expected welfare of an on-platform consumer is always larger than an off-platform consumer's.

The off-platform profit of a firm, per unit measure of loyal consumers, is given by

$$\Pi_{\text{off}}(p) = p(1 - F(p)).$$

The on-platform firm profit per sponsored offer is given by

$$\Pi_{\text{on}}(p) = J\Omega(p; p) = J \int \left[\int^v \min(v - v', p) dF^{J-1}(v') \right] dF(v),$$

by Lemma 1, where the J term comes from the fact that the firm only makes a sponsored offer to a $1/J$ fraction of the on-platform consumers. Note that total firm profit is

$$\Pi(p) = (1 - \lambda)\Pi_{\text{off}}(p) + \lambda\Pi_{\text{on}}(p).$$

Moving on to platform revenue, we note that the revenue generated by a sale to a consumer on the platform by firm j is $\min(v_j, p)$. The total platform revenue is given by the expected value of this minus the value conceded to firms, or

$$\begin{aligned} R(p) &= \lambda \left(\int \min(v, p) dF^J(v) - J \int \left[\int^v \min(v - v', p) dF^{J-1}(v') \right] dF(v) \right) \\ &= \lambda J \left(\int \left[\min(v, p) F^{J-1}(v) - \int^v \min(v - v', p) dF^{J-1}(v') \right] dF(v) \right) \\ &= \lambda J \left(\int^p \int^v v' dF^{J-1}(v') dF(v) + \int_p^1 \int_{v-p}^v v' dF^{J-1}(v') dF(v) \right). \end{aligned}$$

Lastly, the total welfare per consumer off-platform is given by

$$W_{\text{off}}(p) = \int_p^1 v dF(v),$$

since only consumers with $v_j \geq p$ buy. On-platform, total welfare per consumer is

$$W_{\text{on}} = \int v dF^J(v),$$

since there is allocative efficiency on-platform regardless of the posted prices. The total welfare is

$$W(p) = (1 - \lambda)W_{\text{off}}(p) + \lambda W_{\text{on}}.$$

We then have the following comparative statics in λ , the market share of the platform.

Proposition 3 (Platform Size λ)

The following comparative statics hold:

1. *The posted price with data-augmented bidding p_B is increasing in λ .*
2. *The expected surplus of on-platform and off-platform consumers is decreasing in λ .*
3. *The expected off-platform firm profit per consumer Π_{off} is decreasing in λ , and the expected on-platform firm profit per consumer Π_{on} is increasing in λ .*
4. *Platform revenue is increasing in λ .*
5. *Off-platform welfare per consumer W_{off} is decreasing in λ .*

We also have the following comparative statics with respect to the number of firms J .

Proposition 4 (Number of Bidders J)

If $J > -1/(\ln F(1 - p_M))$, then the following hold:

1. *The equilibrium posted price with data-augmented bidding p_B is decreasing in J .*
2. *Expected consumer surplus both off- and on- platform are increasing in J , and so total consumer surplus increases in J .*
3. *Welfare per consumer off- and on- platform are both increasing in J , and so total welfare also increases in J .*

The proofs are left to the appendix, but we will illustrate many of these comparative statics with a simple example.

Example Consider the setting where the distribution F is uniform on $[0, 1]$. Note that in this setting, since the distribution is uniform, the monopoly price is $p_M = 0.5$. We plot the equilibrium posted prices, total firm profit, and consumer surplus resulting from data-augmented bidding for $J = 3, 5, 7$ in Figure 2. As shown in Proposition 3, for any J , the prices are increasing in λ .




FIGURE 2: Posted prices as a function of λ . Results are plotted for $J = 3, 5, 7$.

Figure 3a depicts the consumer surplus as a function of λ . We note that total consumer surplus is increasing in λ . Initially, the welfare gains from moving consumers from being loyal to shopping over all firms dominate (moving consumers from welfare level CS_{off} to CS_{on}) but as the platform becomes too large, the increasing ability to price discriminate on the platform dominates and consumers lose welfare. Hence, total consumer surplus is nonmonotone in λ .

Figure 3b depicts firm profit as a function of λ . Here, firm profits for $J = 3$ are nonmonotone. As mentioned in Proposition 3, the profit per consumer off-platform is decreasing in λ and the profit per consumer on-platform is increasing in λ , and so the overall effect on total profit depends on which force dominates.

Figure 3c depicts the platform revenues as a function of λ . As expected, platform revenues are increasing in λ . However, the interesting feature of this example in platform revenue is that for very large platforms, λ close to 1, the platform revenue can be nonmonotone in J , the number of firms. The two contrasting forces here are that with more firms, the expected value of second-highest bids will be higher, which would suggest that platform revenue should be increasing in J . However, with more firms, as shown in Figure 2, posted prices can be pushed down, thus reducing the price-discriminating ability of the firms on-platform and




FIGURE 3: Consumer surplus, firm profit, platform revenue, and welfare with data-augmented bidding as a function of the share of consumers on the platform, λ . Results are plotted for $J = 3, 5, 7$.

pushing down the platform revenues.

Figure 3d shows that total welfare is increasing in both λ and J , as would be expected.

3.3 Participation Fees

In the bidding model discussed so far, the platform received revenues only from the bids of the advertisers. We now ask whether tools from optimal auction design such as participation fees or reserve prices may increase the revenue of the platform.⁵ In particular, as advertisers have no prior information about the consumers, we investigate how a participation fee for the second-price auction would affect the division of surplus between the platform and advertisers. Thus, we consider the following game:

1. The platform sets a participation fee T .

⁵The importance of such tools in online ad auctions has been widely documented, e.g., by Ostrovsky and Schwarz (2023) for the case of reserve prices.

2. The firms choose whether to pay the participation fee and set their posted prices.
3. If all firms accept, the platform runs a second-price auction for the on-platform consumers. If any firm rejects, the platform can assign the sponsored offers however it would like.

The platform maximizes revenue, and we will assume a firm that is indifferent about accepting chooses to accept. As such, the platform extracts all the producer surplus, up to an outside option the firm could obtain by refusing to participate.

Proposition 5 (Equilibrium with Participation Fees)

In equilibrium, all firms join and the off-platform posted prices are given by (3). Firms bid their true value $\max(v_j, p_B)$. Firm profits are held to their outside option:

$$\Pi_O = \max_p \left\{ \frac{1-\lambda}{J} p(1-F(p)) + \lambda \int_p p F^{J-1}(v-p) dF(v) \right\}. \quad (5)$$

The transfer charged by the platform holds firms to this outside option.

Intuitively, the pricing and bidding behavior follow as in Theorem 1 due to subgame perfection. The transfer charged is as large as possible to make firms indifferent between accepting and rejecting. The proof is in the Appendix.

To gain some intuition for the outside option profit expression, the first part of the expression is the profit from selling to loyal consumers; the second integral expression denotes the profit the firm makes due to the ability of on-platform consumers to search; upon rejecting the platform's service, the firm could still be found by consumers with a sufficiently high value for its product, provided the consumer's value v satisfies $v - p > v'$, where v' is the consumer's value for the best competitor.

4 Managed Advertising Campaigns

In a managed advertising campaign, the platform determines which firm wins the sponsored slot, and makes an offer to that consumer on behalf of that firm. The platform collects an ex-ante fee for this service from each participating firm. Thus, the platform (rather than the advertising firms) selects the bidding functions and the product prices. The key difference is therefore that the firms relinquish agency over the on-platform allocation process to the platform, though they still collect the revenue from on-platform sales. The firms giving up

this agency is why we refer to this as a *managed* campaign. However, the firms still make decisions on participation and on posted prices. Here, we consider *sophisticated* managed campaigns where the platform conditions the on-platform sponsored pricing on all the posted prices of firms. We will be explicit about the extensive form of the game, and we consider subgame perfect equilibrium.

The platform offers automated pricing on-platform, firms pay a participation fee and set posted prices. Let $a_j \in \{0, 1\}$ denote firm j 's acceptance decision. The game has the following extensive form:

1. The platform proposes to all firms a mechanism (s, p, T) , where $s : V^J \times \{0, 1\}^J \times \mathbb{R}_+^J \rightarrow J$ is a *steering policy*, $p : V^J \times \{0, 1\}^J \times \mathbb{R}_+^J \rightarrow \mathbb{R}_+$ is a *pricing policy*, and $T \in \mathbb{R}_+^J$ is a profile of lump-sum transfers.
2. The firms simultaneously decide whether to accept ($a_j = 1$) or reject ($a_j = 0$) the platform's offer and what off-platform price \bar{p}_j to post.
3. If a firm accepts the offer, that firm pays the transfer T , and its product will be offered to a subset of on-platform consumers according to policies s and p .

Intuitively, a steering policy maps consumer value and a profile of acceptance to a choice of firm to steer the consumer towards. The pricing policy maps the consumer value for the steered firm's product and the acceptance profile into a price. The dependence on the acceptance profile allows the platform to react to the firms' participation decisions. In particular, we use $\vec{1}$ to denote the vector $(1, 1, \dots, 1)$ of all firms participating.

In this mechanism, the platform collects an ex-ante fee for its on-platform consumers and its data that allows for price discrimination. Thus, it bundles both access and price-discriminating ability and charges a fee for the bundle. Note that the bundling of these two services implies that if firms set posted prices off the equilibrium path in the third stage, the platform still makes price-discriminating offers and some consumers could potentially be poached by other firms via the search ability of online consumers. That is, the steering policy guarantees the firms the opportunity to price discriminate on the segment of on-platform consumers, but the firm could still lose consumers to search. However, in the equilibrium characterization, we show that this does not happen on-path.

The pricing policy is a function $V^J \times \{0, 1\}^J \times \mathbb{R}_+^J \rightarrow \mathbb{R}_+$; in other words, the platform can condition its pricing policy on the posted prices set by the firms and the full value vector of the consumer.

As a note to break ties, if the platform can propose two revenue-equivalent mechanisms but one mechanism results in more on-platform consumers purchasing their sponsored offers, the platform prefers the mechanism where more on-platform consumers purchase their sponsored offers. We first argue that the platform finds it optimal to steer efficiently:

Proposition 6 (Efficient Platform Steering)

An optimal strategy for the platform is to steer the consumer efficiently among the participating firms:

$$s(v, a) = \arg \max_j v_j \text{ s.t. } a_j = 1.$$

4.1 Best-Value Pricing

We now focus on a specific instance of a sophisticated managed campaign and then show that this specific pricing policy is revenue-optimal for the platform. Formally, define *best-value pricing* as the pricing policy dictated by:

$$p(v, a, \bar{p}) = \min_{k \neq j} (v_j, \bar{p}_j, v_j - v_k + \bar{p}_k) \tag{6}$$

where $j = s(v, a)$ is the firm the platform steers the consumer towards. Note that the arguments in Proposition 6 still hold in this setting, and so the platform steers efficiently.

Intuitively, best-value pricing ensures that there will never be poaching even off the equilibrium path, or equivalently that the sponsored offer always guarantees the best value to the consumer. In this sense, the best-value pricing guarantee is stronger than a most-favored-nation clause that ensures firms offer their goods at a lower price on-platform than off-platform. In addition to doing so, the guarantee in (6) makes sure no competing firm offers a lower price than the sponsored firm.

We then obtain the following equilibrium characterization, where we subscript the off-platform price by V to denote that this results from (best-) value pricing.

Theorem 2 (Best-Value Managed Campaign Equilibrium)

The symmetric managed campaign equilibrium with best-value pricing has the platform offer efficient steering and the posted prices are characterized by the following implicit equation:

$$p_V = \frac{1 - F(p_V) + \frac{\lambda J}{1-\lambda} \left(\int_{p_V} F^{J-1}(v) dF(v) \right)}{f(p_V)}. \tag{7}$$

The proof is algebraic and involves writing out the profit expressions of the firms and

deriving the implicit price characterization in (7) from the first order condition, so it is left to the Appendix.

As it turns out, best-value pricing is revenue-optimal for the platform. That is, best-value pricing attains the maximum revenue a platform can achieve in the sophisticated managed campaign setting.

Theorem 3 (Optimal Managed Campaign)

The best-value pricing managed campaign is platform revenue-maximizing among all sophisticated managed campaigns.

Proof. To show this, consider the problem of a vertically integrated platform that jointly maximizes the profit of firms and the platform. The vertically integrated platform can jointly coordinate on-platform and off-platform pricing but still faces the showrooming constraint due to consumer search. The vertically integrated firm's problem is then to maximize

$$\max_p \left\{ (1 - \lambda)p(1 - F(p)) + \lambda \int \min(v, p) dF^J(v) \right\}.$$

The first-order condition of the planner problem is

$$(1 - \lambda)(1 - F(p) - pf(p)) + \lambda \int_p dF^J(v) = 0.$$

Expanding dF^J , and dividing through by J , we get

$$\frac{1 - \lambda}{J}(1 - F(p) - pf(p)) + \lambda \int_p F^{J-1}(v) dF(v) = 0.$$

But by definition, p_V exactly satisfies this first order condition, and by the characterization in Theorem 4, p_V are exactly the off-platform prices in the sophisticated managed campaign. Thus, this implies that the sophisticated managed campaign necessarily maximizes the joint surplus of the platform and firms.

Now, note that the firms are guaranteed their outside option value (defined in (5)) since in any managed campaign, the firms could refuse to participate. Additionally, note that in the sophisticated managed campaign described, the firms make exactly their outside option, since the transfer the platform charges to each firm makes them exactly indifferent between joining the platform and not. Since the sophisticated campaign maximizes the joint surplus of the platform and firm and concedes the smallest possible surplus to the firms, it follows that the platform earns the most revenue in the sophisticated managed campaign over any

managed campaign. □

In fact, in the proof, we actually showed that the joint surplus obtained by the firms and the platform is maximized for best-value pricing; that is,

Corollary 1 (Producer Surplus)

Producer surplus (sum of firm profit and platform revenue) is maximized for best-value pricing and equals the profit of a vertically integrated platform that owns the firms.

4.2 Comparing Advertising Mechanisms

We now compare the equilibrium posted prices and the welfare implications under these two distinct mechanisms, the data-augmented second price auction and the optimal sophisticated managed campaign mechanism. We start with the comparison of the prices off the platform. Recall the pricing equations (3) and (7).

Theorem 4 (Welfare and Price Comparison)

The posted prices in the optimal sophisticated managed campaign are higher than the posted prices under data-augmented bidding

$$p_V \geq p_B \geq p_M$$

Total consumer surplus and total welfare are lower in the managed campaign than under the bidding equilibrium.

Proof. Consider the derivative of the best-response profit maximization problem with respect to the posted price for each of the three models. In the bidding model is, we have

$$\frac{1-\lambda}{J} (1 - F(p) - pf(p)) + \lambda \left(\int_p F^{J-1}(v-p) dF(v) \right)$$

In the sophisticated campaign, we have

$$\frac{1-\lambda}{J} (1 - F(p) - pf(p)) + \lambda \left(\int_p F^{J-1}(v) dF(v) \right).$$

Note that the second expression is larger than the first, since $F^{J-1}(v) \geq F^{J-1}(v-p)$. Hence, we must have $p_V \geq p_B$. Note that this also implies $p_V \geq p_B \geq p_M$ by Proposition 2. Since welfare and total consumer surplus are both decreasing in posted price, the welfare comparative statics follow. □

Theorem 4 shows that the platform offering a best-value pricing policy eliminates the threat of poaching and weakens competition between firms; this reduced competition thus results in higher posted prices than bidding. Further, an implication of Theorem 3 is that platform revenue is higher for the sophisticated managed campaign:

Corollary 2 (Platform Revenue Comparison)

Platform revenue in the optimal sophisticated managed campaign is higher than in the data-augmented bidding equilibrium.

Intuitively, the sophisticated managed campaign gives the platform the ability to offer pricing and steering that replicate the bidding equilibrium strategies, and by doing so the platform could at least attain its revenue from the bidding equilibrium.

5 Policy Interventions

In this section, we investigate the impact of potential interventions that a policymaker might consider imposing on the platform’s ability to use information on values and posted prices.

5.1 Competition Management

We first analyze a regulatory policy that targets the platform’s role in managing competition between firms. Suppose that the platform’s pricing policy can only condition on the value the consumer has for the steered firm and on the participation decisions, and not the posted prices of other firms. In particular, this forces the platform to price *independently* of the posted price decisions and the consumer’s value for alternatives, which curtails the ability of the platform to soften competition.

Formally, recall that in the sophisticated managed campaign, the platform’s available pricing policy space was $p : V^J \times \{0, 1\}^J \times \mathbb{R}_+^J \rightarrow \mathbb{R}_+$. In this section, we restrict the pricing policy space to $p : V \times \{0, 1\}^J \rightarrow \mathbb{R}_+$; that is, the platform cannot make its pricing policy contingent on the posted off-platform prices or the consumer’s value for other firms anymore.

Theorem 5 (Independent Managed Campaign Equilibrium)

The equilibrium posted price in the independent managed campaign is

$$p_I = \max(p_C, p_M). \quad (8)$$

where p_C is defined by the solution to

$$p_C = \frac{1 - F(p_C) + \frac{\lambda J}{1-\lambda} \left(\int_{p_C} (F^{J-1}(v) - p_C dF^{J-1}(v)) dF(v) \right)}{f(p_C)}. \quad (9)$$

The proof of this result is in the Appendix. The subscript I is used here since p_I denotes the off-platform posted price of the firms for the independent managed campaign; it is equal to the candidate p_C identified by the first-order condition (16) if it is larger than the monopoly price, but is the monopoly price otherwise. We present two examples that show that the independent campaign price can indeed be equal to the monopoly price.

Example ($p_I > p_M$): Take a uniform distribution of values ($F(x) = x$), and suppose there is an equal share of on-platform and off-platform consumers ($\lambda = 1/2$), and consider two firms. From the pricing equation (9), we get

$$p_C = 1 - p_C + 2 \int_{p_C} (v - p_C) dv \approx 0.59 > 0.5 = p_M.$$

Example ($p_I = p_M$): Consider almost the exact same environment as the previous example (uniform distribution of values, equal share of consumers on- and off-platform) but now, we introduce a third firm. From the pricing equation (9), we get

$$p_C = 1 - p_C + 3 \int_{p_C} (v^2 - p_C(2v)) dv \approx 0.43 < 0.5 = p_M.$$

By adding one firm to the previous example, the competitive effect becomes stronger, and the posted price falls to the level of the monopoly price p_M . This insight regarding competition generalizes; with enough firms and a regularity condition on F , the independent campaign does not distort off-platform posted prices.

Proposition 7 (Price Comparison)

Suppose $\exists B > 0$ such that $F(v)/f(v) < B \forall v$. Then for all sufficiently high J , $p_I = p_M$.

In Proposition 8, we discuss the posted price and welfare implications of independent managed campaigns, relative to the sophisticated managed campaign (1.) and to data-

augmented bidding (2.).

Proposition 8 (Price and Welfare Comparisons)

1. *The off-platform prices satisfy $p_V > p_I$. Both consumer surplus and welfare are higher in the independent managed campaign than in the sophisticated managed campaign.*
2. *Suppose that F^{J-1} is convex. Then $p_I \leq p_B$, and total welfare and total consumer surplus are higher in the independent managed campaign than in data-augmented bidding. Further, if F^{J-1} is concave, all the inequalities are reversed.*

In words, posted prices are higher in the bidding equilibrium than in the independent managed campaign equilibrium. Since total welfare and total consumer surplus are both decreasing, total welfare and consumer surplus are both higher under the managed campaign. Proposition 8 shows that allowing the platform to run an independent managed campaign can create a more competitive environment relative to data-augmented bidding; the threat of poaching is larger, and the competition for on-platform consumers dominates.


To interpret the condition that F^{J-1} is convex, note that F^{J-1} represents the cumulative distribution function of the maximum of $J - 1$ values drawn from F . For large enough J , this cumulative distribution function is convex under relatively weak conditions. Indeed, if the density f is such that f'/f is bounded below, then there always exists a J large enough such that F^{J-1} is convex.

We now discuss the implications of independent managed campaigns for platform revenue. Intuitively, since the joint profit of the platform and firms increases with posted prices up to p_V , the platform revenue ordering between bidding and the independent campaign should follow the price ranking. More precisely,

Proposition 9 (Revenue Comparison)

If $p_B \geq p_I$, platform revenue in a bidding model with participation fees is weakly higher than in the independent managed campaign. Otherwise, the platform earns less in the bidding model with participation fees relative to the independent managed campaign.

Proof. Note that in both models, the firms are held to their outside options. Hence, whether the platform earns more depends exactly on the producer surplus extracted. By Theorem 3, the off-platform price p_V induced by the sophisticated managed campaign maximizes producer surplus. By Theorem 4, $p_V \geq p_B, p_I$. Since producer surplus is concave in the off-platform price, p_V maximizes producer surplus, and $p_V \geq p_B, p_I$, the producer surplus is larger in the bidding model iff $p_B \geq p_I$. □



(A) Platform revenue as a function of λ , $J = 2$

(B) Platform revenue as a function of λ , $J = 3$

FIGURE 4: Platform revenue with consumer values distributed as $F(v) = v^{3/4}$. With $J = 2$, F^{J-1} is concave, and with $J = 3$, F^{J-1} is convex. In the first figure for $J = 2$, the platform revenue from the independent managed campaign lies in between the sophisticated managed campaign revenue and bidding with participation fees. In the second figure, the relative ordering of bidding and independent campaigns are switched.

In Figure 4, we plot the revenue generated by the platform in the bidding model and the independent managed campaign as functions of λ when consumer values are drawn from value distribution $F(v) = v^{3/4}$. Figure 4a shows the revenue when there are $J = 2$ firms, and Figure 4b plots the revenue for $J = 3$ firms. Figure 4a demonstrates a scenario where the independent managed campaign yields more revenue, and Figure 4b demonstrates a case where data-augmented bidding yields more revenue. However, if we allow the platform to charge participation fees, it is clear the platform earns more revenue than in the standard bidding model without a participation fee. It is also true that in a bidding model with participation fees, the platform earns more than in an independent managed campaign. Finally, the sophisticated managed campaign results in higher platform revenue than the independent managed campaign and bidding, as would be expected by Theorem 3.

5.2 Privacy and Data

We now consider assessing the impact of privacy regulation. Suppose the platform cannot observe the willingness-to-pay of on-platform consumers for each firm, but only which firm the consumer prefers most. That is, the platform can still steer cohorts of consumers towards firms, but cannot price discriminate within the cohort. This eliminates the firms' ability to price contingent on the consumer data vector. Formally, recall that in the sophisticated managed campaign, the platform's available pricing policy space was $p : V^J \times \{0, 1\}^J \times \mathbb{R}_+^J \rightarrow$

\mathbb{R}_+ . Here, the pricing policy space is restricted to $p : \{0, 1\}^J \times \mathbb{R}_+^J \rightarrow \mathbb{R}_+$; that is, the platform cannot price based on the consumer's individual type vector, but can condition pricing on other observables (i.e. the accept/reject decisions and the posted prices). This is in contrast to the independent managed campaign, which conditions advertised prices on the consumer's value but not on posted prices.

Proposition 10 (Cohort Privacy)

In equilibrium, the platform steers efficiently. The equilibrium posted price under the privacy restriction is p_P with:

$$p_P = \frac{1 - (1 - \lambda)F(p_P) - \lambda F^J(p_P)}{(1 - \lambda)f(p_P) + \lambda J F^{J-1}(p_P)f(p_P)}. \quad (10)$$

The product price on the platform is identical to the off-platform price p_P .

The proof is in the Appendix. Intuitively, firms face a distributional mixture of consumers; a measure $(1 - \lambda)/J$ of consumers are loyal with values distributed according to F , and a measure of λ/J consumers are considering them on-platform with values distributed as F^J . Hence, the firm would like to be able to set higher prices to take advantage of a more favorable distribution of consumer values.

Proposition 11 (Price Impact of Cohort Privacy)

The equilibrium posted price under the privacy restriction satisfies $p_M < p_P < p_V$.


Proof. The first-order condition for p_P is

$$\frac{1 - \lambda}{J}(1 - F(p) - pf(p)) + \frac{\lambda}{J}(1 - F^J(p) - JpF^{J-1}(p)f(p)) = 0. \quad (11)$$


The first-order condition for p_V is given by

$$\frac{1 - \lambda}{J}(1 - F(p) - pf(p)) + \frac{\lambda}{J}(1 - F^J(p)) = 0. \quad (12)$$

Since the left-hand side of (12) is larger than the left-hand side of the (11) and both are decreasing in p , we have $p_P < p_V$. To see that $p_P > p_M$, note that p_P is equal to the inverse hazard rate of the distribution $(1 - \lambda)F + \lambda F^J$ at p_P , and p_M is equal to the inverse hazard rate of F at p_M . The likelihood ratio of $(1 - \lambda)F + \lambda F^J$ relative to F is $(1 - \lambda) + \lambda J F^{J-1}(v)$ which is monotonically increasing, so the inverse hazard rate of $(1 - \lambda)F + \lambda F^J$ is greater than the inverse hazard rate of F . Hence, we have $p_P > p_M$. \square




(A) Off-platform prices as a function of λ , $J = 2$




(B) Off-platform prices as a function of λ , $J = 3$

FIGURE 5: Off-platform prices under the privacy restriction, with and without sponsored offers, compared to the benchmarks (sophisticated managed campaign and data-augmented bidding). Prices are plotted as a function of λ for a uniform distribution of values and $J = 2, 3$. In the large λ limit for $J = 2$, the no-offer price is larger than the with-offer price.

Figure 5 depicts the posted prices for a uniform distribution of values, for $J = 2, 3$ firms. The plots vary the share of on-platform consumers λ . Note that for a uniform distribution, the monopoly price $p_M = 0.5$. As would be expected from Proposition 11, the sophisticated managed campaign price is highest, and the privacy price with offers p_P is above $p_M = 0.5$ but below the managed campaign price p_V . However, the relative ordering of the bidding price and the privacy price is ambiguous: for $J = 3$, on smaller platforms the bidding price p_B can be lower than p_P . Intuitively, for sufficiently many firms, the competitive effect of bidders on each other profits sufficiently outweighs the incentives to raise prices. The welfare implications of privacy are more ambiguous, as we plot in Figure 6 for a uniform distribution of consumer values.



(A) Consumer surplus as a function of λ , $J = 2$



(B) Total welfare as a function of λ , $J = 2$

FIGURE 6: Welfare implications of privacy, with a uniform distribution of consumer values.

Proposition 11 shows that the off-platform posted prices are lower than under the managed campaign, which implies greater consumer surplus and total welfare off the platform. However, the privacy restriction reduces total welfare on the platform too: the inability to price discriminate means that low-value on-platform consumers are priced out by the privacy restrictions. Hence, in the limit of large platforms, the total welfare can be worse under privacy restrictions than under data-augmented bidding or managed campaigns. The consumer surplus is larger, however, since the low-value consumers get their surplus extracted away in the non-privacy benchmarks; that is, the loss in welfare comes primarily from reduced producer surplus.

6 Conclusion

Many digital platforms such as Google, Meta, Amazon, and TikTok generate revenue through advertising by placing ads or sponsored slots on their own and partner websites. These platforms use a bidding and auction mechanism to determine advertisers' willingness to pay and a ranking and recommendation mechanism to select the most suitable ad to display to the viewer. The platform's knowledge about the match value between consumers and products is critical to the success of both mechanisms. This knowledge helps generate the most competitive bids from advertisers in the auction and supports more clicks and other engagement in the ranking mechanism. We proposed an integrated model that considers how auction and data jointly determine match formation on digital platforms. We also highlighted the value of information and data for the platform in the joint deployment of these services on both sides of the market.

The auction mechanisms on the platform have substantial implications for product prices. On the platform, the data made available to the advertisers allows for efficient matching, yet most of the surplus accrues to the platform. Off the platform, the implications of the bidding algorithms are more dramatic. Advertisers raise prices off the platform to gain a competitive edge on the platform. The cross-channel distortions become more pronounced under the sophisticated managed campaign than in the traditional (generalized) second price auction. This suggests the need for further analysis of how algorithmic bidding on platforms impacts competition and welfare in all markets, particularly off the platform. Indeed, we have shown that the higher costs of online advertising under a more extractive mechanism are passed on to consumers by means of higher product prices *off* rather than *on* the platform.

A Appendix

Proof of Proposition 1

Proof. Suppose $v_j > v_k$. Note that since the platform mechanism is a second-price auction, it is weakly dominant for each firm to bid exactly what the online consumer is worth to that firm. We proceed using casework. As a useful reference, denote by

$$\underline{u} = \max(\max_{i \neq j, k} (v_i - \bar{p}_i), 0),$$

the utility the consumer would get from all firms except j and k . \underline{u} thus is a lower bound on the utility that is conceded to any consumer that does purchase from j or k .

First, consider the cases where $\bar{p}_k \geq \bar{p}_j$. Note that this implies $v_j - \bar{p}_j > v_k - \bar{p}_k$. There are two subcases to consider: either $v_j > \bar{p}_j$ or $v_j \leq \bar{p}_j$. In the first subcase, the highest price that firm j can charge is restricted by showrooming and the nonnegativity constraint, so $b_j(v) = \max(\min(\bar{p}_j, v_j - \underline{u}), 0)$. If firm k were to win the auction, then firm k 's offer must guarantee at least $v_j - \bar{p}_j$ utility to the consumer to dissuade the consumer from going off-platform, and hence the most that firm k can offer is

$$\begin{aligned} b_k(v) &\leq \max(\min(v_k - (v_j - \bar{p}_j), v_k - \underline{u}), 0) \\ &= \max(\min(\bar{p}_j - (v_j - v_k), v_k - \underline{u}), 0) \\ &< \max(\min(\bar{p}_j, v_j - \underline{u}), 0) = b_j(v). \end{aligned}$$

For the second subcase, since $v_j \leq \bar{p}_j$, the consumer is worth v_j to firm j , so $b_j(v) = \max(v_j - \underline{u}, 0)$. Then $v_k < v_j \leq \bar{p}_j < \bar{p}_k$, so the consumer is worth v_k to firm k , and the bids satisfy the following condition

$$b_k(v) = \max(v_k - \underline{u}, 0) \leq \max(v_j - \underline{u}, 0) = b_j(v).$$

Now, consider the cases where $\bar{p}_j > \bar{p}_k$. We have four subcases here.

1. $v_k < \bar{p}_k$ and (a) $v_j < \bar{p}_j$ or (b) $v_j \geq \bar{p}_j$,
2. $v_k \geq \bar{p}_k$ and (a) $v_j \leq v_k + \bar{p}_j - \bar{p}_k$ or (b) $v_j > v_k + \bar{p}_j - \bar{p}_k$.

In subcase (1)(a), $v_j < \bar{p}_j$ and $v_k < \bar{p}_k$, so the highest price j can charge is v_j and the highest price k can charge is v_k . Then $b_k(v) = \max(v_k - \underline{u}, 0)$ and $b_j(v) = \max(v_j - \underline{u}, 0)$,

and so

$$b_k(v) = \max(v_k - \underline{u}, 0) \leq \max(v_j - \underline{u}, 0) = b_j(v).$$

In subcase (1)(b), $v_j \geq \bar{p}_j > \bar{p}_k > v_k$, so $v_j - \underline{u} \geq v_k - \underline{u}$ and $\bar{p}_j \geq v_k - \underline{u}$. Hence

$$b_k(v) = \max(v_k - \underline{u}, 0) \leq \max(\min(\bar{p}_j, v_j - \underline{u}), 0) = b_j(v).$$

In subcase (2)(a), $v_k \geq \bar{p}_k$ and $v_j \leq v_k + \bar{p}_j - \bar{p}_k$. Firm j must concede at least $v_k - \bar{p}_k$ utility to the consumer (else the consumer would buy k 's product), and hence the bid

$$\begin{aligned} b_k(v) &= \max(\min(\bar{p}_k, v_k - \underline{u}), 0) \\ &\leq \max(\min(\bar{p}_k + (v_j - v_k), v_j - \underline{u}), 0) \\ &= \max(\min(v_j - (v_k - \bar{p}_k), v_j - \underline{u}), 0) \\ &= b_j(v). \end{aligned}$$

Finally, in the last subcase, note that $v_j - \bar{p}_j \geq v_k - \bar{p}_k$. The bids are

$$b_k(v) = \max(\min(\bar{p}_k, v_k - \underline{u}), 0) \leq \max(\min(\bar{p}_j, v_k - \underline{u}), 0) \leq b_j(v).$$

In all cases, $b_k(v) \leq b_j(v)$. □

Proof of Lemma 1

First consider the regime $p < p'$. Since the firm only wins the consumers for which $v > v'$, it follows that $v - p > v' - p'$, so the firm never concedes rents to the threat of the second-best firm. There are three distinct regions to consider here: if $v < p$, the firm earns v and pays v' . If $v \geq p$, then the firm earns p , but the firm pays 0 if $v' < v - p$ the firm pays 0, else the firm pays $v' - (v - p)$. The Ω term in this region is thus given by:

$$\begin{aligned} \Omega(p; p') &= \int^p \left(\int^v (v - v') dF^{J-1}(v') \right) dF(v) \\ &\quad + \int_p \left(\int^{v-p} p dF^{J-1}(v') + \int_{v-p}^v (p - (v' - (v - p))) dF^{J-1}(v') \right) dF(v) \\ &= \int^p \int^v (v - v') dF^{J-1}(v') dF(v) \\ &\quad + \int_p \left(\int^{v-p} p dF^{J-1}(v') + \int_{v-p}^v (v - v') dF^{J-1}(v') \right) dF(v) \\ &= \int \int^v \min(v - v', p) dF^{J-1}(v') dF(v). \end{aligned}$$

Now, suppose $p \geq p'$. We again proceed with casework. If $v < p'$, then no constraints bind and the firm earns v and pays the bid v' . If $v \in [p', p]$ and $v' < p'$, the firm once again earns v and pays v' . If $v \in [p', p]$ and $v' \in (p', v]$, then the firm earns $v - (v' - p')$ and pays p' . If $v > p$, then there are 3 subcases for v' . If $v' < v - p$, the firm earns p and pays 0. If $v' \in [v - p, p' + v - p]$, the firm earns p and pays $v' - (v - p)$. If $v' \geq p' + v - p$, then the firm earns $v - (v' - p')$ and pays p' .

$$\begin{aligned}
\Omega(p; p') &= \int^{p'} \left(\int^v (v - v') dF^{J-1}(v') \right) dF(v) \\
&\quad + \int_{p'}^p \left(\int^{p'} (v - v') dF^{J-1}(v') + \int_{p'}^v (v - (v' - p') - p') dF^{J-1}(v') \right) dF(v) \\
&\quad + \int_p^p \left(\int^{v-p} p dF^{J-1}(v') + \int_{v-p}^{p'+(v-p)} (p - (v' - (v - p))) dF^{J-1}(v') \right) dF(v) \\
&\quad + \int_p^p \left(\int_{p'+(v-p)}^v (v - (v' - p') - p') dF^{J-1}(v') \right) dF(v) \\
&= \int^{p'} \left(\int^v (v - v') dF^{J-1}(v') \right) dF(v) \\
&\quad + \int_{p'}^p \left(\int^{p'} (v - v') dF^{J-1}(v') + \int_{p'}^v (v - v') dF^{J-1}(v') \right) dF(v) \\
&\quad + \int_p^p \left(\int^{v-p} p dF^{J-1}(v') + \int_{v-p}^v (v - v') dF^{J-1}(v') \right) dF(v) \\
&= \int \int^v \min(v - v', p) dF^{J-1}(v') dF(v).
\end{aligned}$$

Finally, to make sure the first-order condition is valid, we take the second derivative to check that the objective is concave:

$$\frac{\partial^2 \Omega(p)}{\partial p^2} = \int_p (-dF^{J-1}(v - p)) dF(v) = - \int_p (J - 1) F^{J-2}(v - p) f(v - p) f(v),$$

which is always negative; so the on-platform profit term is concave. \square

Proof of Proposition 2

This result is implied by Proposition 3.

Proof of Proposition 3

Let the right hand sides of the pricing equations (3) and (4) be

$$\Phi_B(x) = \frac{1 - F(x) + \frac{\lambda J}{1-\lambda} \left(\int_x F^{J-1}(v-x) dF(v) \right)}{f(x)}. \quad (13)$$

and

$$\underline{\Phi}(x) = \frac{1 - F(x)}{f(x)}.$$

respectively. Clearly $\Phi_B(x) \geq \underline{\Phi}(x)$, with inequality holding strictly if $x < 1$ and $\lambda > 0$. Since, by regularity $\underline{\Phi}(x)$ is decreasing, and p_M is a fixed point of $\underline{\Phi}$ and p_B is a fixed point of Φ_B , we must have $p_B \geq p_M$.

To see the first statement, note that $\lambda/(1-\lambda)$ is increasing in λ , and so the right hand side of the implicit price equation (13), is increasing in λ . It follows that p_B is also increasing in λ . Then the second statement follows from the first and the fact that CS_{off} and CS_{on} are both decreasing in p . Examining Π_{off} , note that

$$\frac{d\Pi_{\text{off}}}{dp} = 1 - F(p) - pf(p) = f(p) \left(\frac{1 - F(p)}{f(p)} - p \right).$$

Since by Proposition 2, $p_B \geq p_M$,

$$\frac{1 - F(p_B)}{f(p_B)} - p_B \leq \frac{1 - F(p_M)}{f(p_M)} - p_M = 0.$$

So since p_B is increasing in λ by the first statement, and Π_{off} is decreasing in p , it follows that Π_{off} is decreasing in λ . For on-platform consumers, Π_{on} is clearly increasing in p , so Π_{on} is also increasing in λ . For the platform revenue, recall that

$$R(p) = \lambda J \left(\int^p \int^v v' dF^{J-1}(v') dF(v) + \int_p \int_{v-p}^v v' dF^{J-1}(v') dF(v) \right).$$

It suffices to show that the parenthesized part is increasing in p , since p is increasing in λ .

Taking the derivative of the parenthesized part, we get

$$\begin{aligned} & f(p) \int^p v' dF^{J-1}(v') - f(p) \int^p v' dF^{J-1}(v') + \int_p (v-p) dF^{J-1}(v-p) dF(v) \\ &= \int_p (v-p) dF^{J-1}(v-p) dF(v) > 0. \end{aligned}$$

Hence, the platform revenue is increasing in λ . Finally, the total off-platform welfare per consumer is decreasing in p , and has no other λ dependence, so W_{off} is decreasing in λ . \square

Proof of Proposition 4

Recall the right hand side of the pricing equation (3) is

$$\Phi_B(x) = \frac{1 - F(x) + \frac{\lambda J}{1-\lambda} \left(\int_x F^{J-1}(v-x) dF(v) \right)}{f(x)}.$$

The partial derivative of this expression with respect to J gives

$$\begin{aligned} \frac{\partial \Phi_B}{\partial J} &= \frac{\lambda}{(1-\lambda)f(x)} \left(\int_x (F^{J-1}(v-x) + JF^{J-1}(v-x) \ln F(v-x)) dF(v) \right) \\ &= \frac{\lambda}{(1-\lambda)f(x)} \left(\int_x (1 + J \ln F(v-x)) F^{J-1}(v-x) dF(v) \right). \end{aligned}$$

Since by assumption $1 < -J \ln F(1 - p_M)$, $1 + J \ln F(v-x) \leq 1 + J \ln F(1 - p_M) < 0$ for $x \geq p_M$. Hence this derivative is negative with respect to J . Since p_B is the fixed point of Φ_B and from Proposition 2 $p_B \geq p_M$, it thus follows that p_B must be decreasing in J .

Since p_B must be decreasing in J , it follows then $CS_{\text{off}}(p_B)$ and $W_{\text{off}}(p_B)$ are increasing in J , since they are decreasing in p and has no other J dependence. Note that W_{on} is equivalently the expected value of the max of J i.i.d random variables distributed as F , and hence is increasing in J . C_{on} is the expected value of an increasing function of v $\max(0, v-p)$, where v is distributed as a max of J i.i.d random variables. Hence the partial derivative of C_{on} with respect to J is positive. Because C_{on} is also decreasing in p and p_B decreases in J , it follows that $C_{\text{on}}(p_B)$ is increasing in J . Since CS is a fixed (not J -dependent) linear combination of CS_{off} and CS_{on} and likewise for W , CS and W are both increasing in J . \square

Proof of Proposition 5

By Lemma 1, the firm willing to pay the most for any consumer regardless of off-platform prices is the firm which the consumer has the highest value for; hence, it is not revenue optimal for the platform to exclude any firm from participating. Consider the subgame after all firms have paid the participation fee. By Theorem 1, the pricing condition for off-platform prices is given by (3), and firms bid their true value $\max(v_j, p_B)$. It is then straightforward to see that the maximum participation fee must hold the firm's profit to what they could get from being excluded; hence, it is optimal for the platform to charge transfer fees that make

the firm indifferent between joining and not. Since the exclusion profit is given by (5), the result follows. \square

Proof of Proposition 6

If the platform's steering policy were inefficient (i.e., $s(v, a) = j$ for a positive measure of consumer values v whose highest value is not for firm j), the platform could instead steer those consumers to their most preferred firms. By doing so, the platform could charge a higher transfer from each consumer's most preferred firm, since the consumer is worth weakly more to the most preferred firm. Hence it is weakly dominant for the platform to offer each consumer their favorite product. \square

Proof of Theorem 2

Once again, we consider the best response problem given other firms setting price p' .

First, consider firms setting price $p < p'$. Here, since the firm will not poach anyone, the firm collects p on all values above p and the value from all values below p . The firm's profit is

$$\frac{1-\lambda}{J}p(1-F(p)) + \lambda \left(\int^p v F^{J-1}(v) dF(v) + \int_p p F^{J-1}(v) dF(v) \right).$$

The derivative with respect to p is

$$\begin{aligned} & \frac{1-\lambda}{J}(1-F(p) - pf(p)) + \lambda \left(pF^{J-1}(p)f(p) - pF^{J-1}(p)f(p) + \int_p F^{J-1}(v) dF(v) \right) \\ &= \frac{1-\lambda}{J}(1-F(p) - pf(p)) + \lambda \left(\int_p F^{J-1}(v) dF(v) \right). \end{aligned} \quad (14)$$

Now, consider firms setting price $p > p'$. The firm profit function is

$$\frac{1-\lambda}{J}p(1-F(p)) + \lambda \left(\begin{aligned} & \int^{p'} v F^{J-1}(v) dF(v) + \int_{p'}^p \int^{p'} v dF^{J-1}(v') dF(v) \\ & + \int_{p'}^p \int_{p'}^v (v - (v' - p')) dF^{J-1}(v') dF(v) + \int_p \int^{p'+(v-p)} p dF^{J-1}(v') dF(v) \\ & + \int_p \int_{p'+(v-p)}^v (v - (v' - p')) dF^{J-1}(v') dF(v) \end{aligned} \right).$$

The derivative of this expression with respect to p is

$$\frac{1-\lambda}{J}(1-F(p)-pf(p)) + \lambda \left(\begin{array}{l} \int_{p'}^{p'} p dF^{J-1}(v') dF(p) + \int_{p'}^p (p - (v' - p')) dF^{J-1}(v') dF(p) \\ - \int_{p'}^{p'} p dF^{J-1}(v') dF(p) - \int_p^p p dF^{J-1}(p' + (v - p)) dF(v) \\ + \int_p \int_{p'+(v-p)}^{p'+(v-p)} dF^{J-1}(v') dF(v) - \int_{p'}^p (p - (v' - p')) dF^{J-1}(v') dF(p) \\ + \int_p p dF^{J-1}(p' + (v - p)) dF(v) \end{array} \right).$$

Everything cancels except the first term in the third line, so with some algebra, we obtain

$$= \frac{1-\lambda}{J}(1-F(p)-pf(p)) + \lambda \left(\int_p F^{J-1}(p' + v - p) dF(v) \right). \quad (15)$$

Comparing (14) and (15), the derivative matches from the left and right at $p = p'$, and so the best-response function is smooth. One can also easily see that $\int_p F^{J-1}(v) dF(v)$ is decreasing in p , so the objective is concave and we can take the first order condition:

$$0 = \frac{1-\lambda}{J}(1-F(p)-pf(p)) + \lambda \int_p F^{J-1}(v) dF(v).$$

Rearranging, we get the implicit characterization of posted prices in (7). □

Proof of Theorem 5

We first prove two lemmas.

Lemma 2 (Platform Pricing Policy)

It is weakly optimal for the platform to offer first-degree price discrimination up to some cap; that is, $p(v_j, \vec{1}) = \min(v_j, \hat{p})$, where \hat{p} is the cap.

Proof. Suppose the platform chooses a price policy, and the subgame posted price equilibrium resulting from this policy results in posted prices at \bar{p} . If the price policy offered prices larger than \bar{p} , then since the platform weakly prefers sales to occur on-platform, the platform would instead prefer to cap its prices at \bar{p} . Hence, the platform anticipates the posted prices set by the firms and never offers a price larger than the subsequent posted price equilibrium.

Now suppose the platform offers some pricing policy \tilde{p} , and the largest price offered to any value, \hat{p} , is at most the resulting subgame posted price \bar{p} . If the platform is not first-degree price discriminating up to \hat{p} , then

$$\int \tilde{p}(v) dF^J(v) < \int \min(v, \hat{p}) dF^J(v).$$

Since setting a price cap of 0 earns no profit, the intermediate value theorem implies that there exists a $\hat{p}' < \hat{p}$ such that a first-degree price discriminating policy with a cap at \hat{p}' is revenue-equivalent to the original pricing policy. That is, \hat{p}' exists such that

$$\int \tilde{p}(v) dF^J(v) = \int \min(v, \hat{p}') dF^J(v).$$

Since the cap of this alternative pricing policy is lower than \hat{p} , and each firm gets at least as much profit in sales on-platform under this alternative pricing policy if it set a posted price \bar{p} . If $\bar{p} > p_M$, then the firms would also lower off-platform posted prices and gain more surplus, and so the platform could charge a weakly higher transfer for this policy. It could not be the case that $\bar{p} < p_M$, since this would imply that since prices online were capped at $\hat{p} \leq \bar{p}$, some firm would have had an incentive to raise its price to p_M . So the last case to consider is $\bar{p} = p_M$. In this case, to check that this policy is without loss, it suffices to check that firms setting posted prices at p_M is still a subgame equilibrium. Since the marginal incentives for firms to raise or lower posted prices around \bar{p} are unchanged by switching to the price discriminating policy capped at \hat{p}' , it remains to argue that there is no profitable undercutting incentive introduced. But this follows since the first-degree price discriminating policy minimizes the maximum price charged to an on-platform consumer fixing the value of on-platform sales, and so it also minimizes the profit from undercutting deviations. \square

Lemma 3 (Outside Option)

An optimal strategy for the platform sets $p(\cdot, a) = 0$ if $a \neq \vec{1}$.

Proof. Since, by Proposition 6, the platform finds it weakly optimal to make sponsored offers efficiently among participating firms, the platform makes the most revenue from such a steering policy only when all firms accept since the steering policy is most efficient only when all firms accept (as otherwise, there is loss due to some on-platform consumers being shown a sponsored offer for a firm that is not their favorite). Hence, it is optimal for the platform to set transfers such that all firms are willing to accept. Therefore, the platform must offer each firm the difference between rejecting and best responding to the resulting steering policy and accepting. So the optimal strategy of the platform must be to reduce the firm's value from rejection as much as possible. Consider the profit firm j could earn by rejecting. Since Proposition 6 implies that the consumers who would have seen j 's product in the sponsored slot now a sponsored offer from the next-best, and the rejecting firm j can only sell via posted price now, firm j 's profit from rejecting is at least Π_O , the outside option defined in (5). Since offering $p(\cdot, a) = 0$ for $a \neq \vec{1}$ exactly attains this lower bound because

firm j only can sell to consumers who value j 's product p more than any other firm, this is an optimal strategy for the platform. \square

We now move on to the proof of the theorem.

Proof: Note that Proposition 6 and Lemma 3 hold, so we consider the pricing policy of the platform when all firms accept. Because it is optimal for the platform to steer efficiently, and firms are ex-ante symmetric, we look for symmetric equilibria in the pricing subgame on the equilibrium path where all firms join the platform. By rearranging (9), p_C satisfies

$$0 = \frac{1-\lambda}{J} (1 - F(p) - pf(p)) + \lambda \left(\int_p (F^{J-1}(v) - p dF^{J-1}(v)) dF(v) \right). \quad (16)$$

We consider two subcases; when $p_C \geq p_M$ and when $p_C < p_M$.

First, suppose $p_C \geq p_M$. Suppose firms $k \neq j$ set off-platform prices p' and consider the best-response problem of firm j . By Lemma 2, the pricing policy of the platform has a cap \hat{p} . When $p' < \hat{p}$, the profit function takes two forms, depending on whether $p \geq p'$ or $p < p'$. If firm j deviates by raising its posted price to $p \geq p'$, the firm gets consumers poached away if $v_k - p' > v_j - p$, or $v_k > v_j + p' - p$. So the profit function is

$$\frac{1-\lambda}{J} p(1-F(p)) + \lambda \left(\int^{p'} v F^{J-1}(v) dF(v) + \int_{p'}^p v F^{J-1}(p') dF(v) + \int_p p F^{J-1}(v + p' - p) dF(v) \right).$$

The derivative with respect to p in this regime is

$$\begin{aligned} & \frac{1-\lambda}{J} (1 - F(p) - pf(p)) + \lambda \left(\begin{aligned} & p F^{J-1}(p') f(p) - p F^{J-1}(p') f(p) \\ & + \int_p (F^{J-1}(v + p' - p) - p dF^{J-1}(v + p' - p)) dF(v) \end{aligned} \right) \\ & = \frac{1-\lambda}{J} (1 - F(p) - pf(p)) + \lambda \left(\int_p (F^{J-1}(v + p' - p) - p dF^{J-1}(v + p' - p)) dF(v) \right). \end{aligned}$$

In the other regime, firm j , by deviating to a price $p < p'$, can poach some consumers whose maximum other value is for some other firm k 's product, but $v_k - p' < v_j - p$. Note that since firm j is undercutting firm k , j can potentially poach consumers whose value $v_k > v_j$. So the profit of the firm from such a deviation is given by

$$\frac{1-\lambda}{J} p(1 - F(p)) + \lambda \left(\int_p v F^{J-1}(v) dF(v) + \int_p p F^{J-1}(\min(v + p' - p, 1)) dF(v) \right).$$

The derivative with respect to p is

$$\begin{aligned} & \frac{1-\lambda}{J} (1 - F(p) - pf(p)) + \lambda (pF^{J-1}(p)f(p) - pF^{J-1}(p')f(p)) \\ & + \lambda \int_p (F^{J-1}(\min(v + p' - p, 1)) - p dF^{J-1}(\min(v + p' - p, 1))) dF(v). \end{aligned}$$

As we are interested in symmetric equilibria, we take $p \rightarrow p'$. A quick check confirms that the profit of the firm is both continuous at p' and the left- and right- derivatives match at $p = p'$. Hence, we get the first order condition

$$0 = \frac{1-\lambda}{J} (1 - F(p) - pf(p)) + \lambda \left(\int_p (F^{J-1}(v) - p dF^{J-1}(v)) dF(v). \right) \quad (17)$$

Then it is clear that if $\hat{p} \geq p_M$, and $p_C < \hat{p}$, the resulting posted price subgame equilibrium has firms setting prices p_C . Now, if $p' \geq \hat{p}$, the left derivative of the best-response profit function is the same as before, but the right derivative changes; specifically, since the platform is already capping the price offers at \hat{p} , any price increase only affects the offline population: that is, the right derivative at $p \rightarrow p'$ is $\frac{1-\lambda}{J} (1 - F(p) - pf(p))$, which is nonpositive since $p' \geq \hat{p} \geq p_M$. Hence the best response value function of the firms kinks at p' , but the right-derivative is always negative at the kink. So if $p' \geq p_C$, then the optimal best response is p_C , but if $p' < p_C$, then p' is the best response. In particular, this implies that for a particular set of subgames $p_M \leq \hat{p} < p_C$ (which turn out to be off-path), there are multiple equilibria in the subgame. That is, any price $p' \in [\hat{p}, p_C]$ is a subgame equilibrium if $p_C \geq \hat{p}$. However, since the platform is profit maximizing and make a higher transfer profit for more extractive pricing policies (higher $\hat{p} \geq p_M$), if $p_C \geq p_M$, the platform's subgame optimal strategy then is to choose $\hat{p} = p_C$, after which p_C is the unique equilibrium in the posted price subgame.

Now, we turn to the case where $p_C < p_M$. Again, by Lemma 2 the platform pricing policies cap prices at \hat{p} . Consider the best response for a single firm, supposing all other firms are pricing at p' . If $p' < \hat{p}$, then the derivative of the best response function (which we analyzed taking $p \rightarrow p'$ above) is

$$\frac{1-\lambda}{J} (1 - F(p) - pf(p)) + \lambda \left(\int_p (F^{J-1}(v) - p dF^{J-1}(v)) dF(v) \right).$$

Note that this is decreasing, and zero at p_C . So p_C is a locally optimal response. However, if $p_M > \hat{p}$, then the best-response profit function has a right-derivative of $\frac{1-\lambda}{J} (1 - F(p) - pf(p))$ at \hat{p} , which is positive since $\hat{p} < p_M$. This expression also governs the derivative of the best

response on $[\hat{p}, p_M]$; so the best response has a kink at \hat{p} . Hence, the only two potential symmetric equilibria are either p_C or the monopoly price, p_M . Then clearly, if $\hat{p} \geq p_M$, the pricing subgame equilibrium is p_C , and if $\hat{p} \leq p_C$, the pricing subgame equilibrium is p_M .

So it remains to characterize the subgame equilibria if $\hat{p} \in (p_C, p_M)$. Recall that the platform seeks to obtain the maximum transfer from the firms for its service; hence, the platform sets pricing policies to maximize the joint surplus of the platform and firms given the resulting pricing equilibrium. Since the resulting pricing equilibrium is either $p < \hat{p}$ or p_M , if the pricing equilibrium is p , the joint surplus is

$$(1 - \lambda)p(1 - F(p)) + \int \min(v, p) dF^J(v).$$

However, if the platform sets some \hat{p} such that the resulting pricing equilibrium in the subgame is p_M , the joint surplus is

$$(1 - \lambda)p_M(1 - F(p_M)) + \int \min(v, \hat{p}) dF^J(v).$$

Note that by definition of p_M , $(1 - \lambda)p(1 - F(p)) < (1 - \lambda)p_M(1 - F(p_M))$, and $\hat{p} > p$; hence, it follows that the platform's optimal strategy is to choose the largest possible \hat{p} such that p_M is a pricing subgame equilibrium. (That is, raising \hat{p} increases surplus extraction from on-platform consumers, but the platform cannot raise prices so high that firms want to undercut each other). Define the deviating profit of a single firm from undercutting (setting a price below \hat{p}) as

$$\Pi_U(\hat{p}) = \max_{p \leq \hat{p}} \left[\frac{1 - \lambda}{J} p(1 - F(p)) + \lambda \left(\int^p v F^{J-1}(v) dF(v) + \int_p p F^{J-1}(\min(v + \hat{p} - p, 1)) dF(v) \right) \right].$$

We subscript this profit expression with a U to indicate that this is the largest profit that the firm could get by undercutting other firms. In order for the firms to set p_M as the posted price, the value of Π_U must be less than the firm profit from setting p_M , which is

$$\Pi_M(\hat{p}) = \frac{1 - \lambda}{J} p_M(1 - F(p_M)) + \lambda \left(\int \min(v, \hat{p}) F^{J-1}(v) dF(v) \right).$$

Intuitively, Π_M is the profit a firm receives from setting p_M when the platform's pricing policy caps at \hat{p} . Note that $\Pi_U(p_C) < \Pi_M(p_C)$, since the firm's best response price to $\hat{p} = p_C$ is p_C . However, $\Pi_U(p_M) > \Pi_M(p_M)$, since the derivative of the maximand of Π_U is negative at $p = \hat{p} = p_M$. So there exists a largest price $p^* \in [p_C, p_M]$ such that $\Pi_M(p^*) \geq \Pi_U(p^*)$, and it is optimal for the platform to cap prices at p^* , and firms to price at p_M . \square

Proof of Proposition 7

The right-hand side of the pricing condition characterizing p_C is

$$\Phi_M(p) = \frac{1 - F(p) + \frac{\lambda J}{1-\lambda} \left(\int_p (F^{J-1}(v) - p dF^{J-1}(v)) dF(v) \right)}{f(p)}.$$

Note that p_C is the fixed point of Φ_M , and $\Phi_M(p)$ is decreasing in p . Thus, it suffices to show that for some large enough J , $\Phi_M(p_M) < p_M$, as that implies the fixed point of Φ_M is less than p_M . Recall that p_M satisfies the FOC $1 - F(p_M) - p_M f(p_M) = 0$. So

$$\begin{aligned} \Phi_M(p_M) &= p_M + \frac{\lambda}{(1-\lambda)f(p_M)} \left(\int_{p_M} J (F^{J-1}(v) - p_M dF^{J-1}(v)) dF(v) \right) \\ &= p_M + \frac{\lambda}{(1-\lambda)f(p_M)} \left(\int_{p_M} J \left(\frac{F(v)}{f(v)} - p_M(J-1) \right) F^{J-2}(v) f^2(v) dv \right) \\ &< p_M + \frac{\lambda}{(1-\lambda)f(p_M)} \left(\int_{p_M} J (B - (J-1)p_M) F^{J-2}(v) f(v) dv \right). \end{aligned}$$

If $J > 1 + B/p_M$, the term in parentheses is negative, so $\Phi_M(p_M) < p_M$, and $p_C < p_M$.

Proof of Proposition 8

We first compare p_C and p_V . The FOC defining p_C is:

$$\frac{1-\lambda}{J} (1 - F(p) - pf(p)) + \lambda \left(\int_p (F^{J-1}(v) - p dF^{J-1}(v)) dF(v) \right) = 0.$$

The FOC defining p_V is

$$\frac{1-\lambda}{J} (1 - F(p) - pf(p)) + \lambda \left(\int_p F^{J-1}(v) dF(v) \right) = 0.$$

Note that the term multiplying λ is larger in the sophisticated managed campaign. Thus, the marginal profit from raising posted prices is larger in the sophisticated managed campaign, and hence $p_V \geq p_C$. Note that by Theorem 4, $p_V > p_M$. Hence, $p_V \geq \max(p_M, p_C) = p_I$.

Now, we compare p_C and p_B . The FOC defining p_B is:

$$\frac{1-\lambda}{J} (1 - F(p) - pf(p)) + \lambda \left(\int_p F^{J-1}(v - p) dF(v) \right) = 0.$$

Under the assumption that F^{J-1} is convex, then we have that

$$F^{J-1}(v - p) \geq F^{J-1}(v) - p dF^{J-1}(v).$$

since the right-hand side is a first-order expansion of F^{J-1} around v . Thus, the derivative of profit with respect to posted price is weakly larger in the bidding model, which implies that $p_B \geq p_C$. Since $p_B \geq p_M$ by Proposition 2, $p_B \geq p_I$.

Note that if F^{J-1} is concave, then the inequality order reverses:

$$0 \leq F^{J-1}(v - p) \leq F^{J-1}(v) - p dF^{J-1}(v),$$

since the right-hand side is a first-order expansion of F^{J-1} around v . Using the same argument as before, $p_M \leq p_B \leq p_C$. As total welfare and total consumer surplus decrease in p , the welfare comparative statics follow. \square

Proof of Proposition 10

Note that Proposition 6 and Lemma 3 still hold, so the platform steers efficiently (i.e. $s(j, \vec{1}) = j$). We first consider the posted-price best response to the following pricing policy and then show that this pricing policy is optimal for the platform to offer.

Consider the pricing policy $p(a, \bar{p}) = \min_j \bar{p}_j$. That is, the platform offers the price which is lowest of all off-platform prices. Suppose all other firms set a posted price of \bar{p} . Consider the profit-maximizing price for j to set on the segment of consumers who have the highest willingness to pay for j 's product. If j sets a price $p < \bar{p}$, j wins each consumer in the segment; hence, the profit for setting p would be

$$\frac{1 - \lambda}{J} p(1 - F(p)) + \frac{\lambda}{J} (p(1 - F^J(p)))$$

If $p > \bar{p}$, and p is less than j 's posted price, j loses some on-platform consumers to poaching, and the profit is

$$\frac{1 - \lambda}{J} p(1 - F(p)) + \frac{\lambda}{J} \bar{p}(1 - F^J(\bar{p}))$$

Note that the left derivative of profits at \bar{p} is:

$$\frac{1 - \lambda}{J} (1 - F(p) - p f(p)) + \frac{\lambda}{J} (1 - F^J(p) - J p F^{J-1}(p) f(p))$$

and the right derivative is just $\frac{1 - \lambda}{J} (1 - F(p) - p f(p))$. Setting the left derivative to zero and

rearranging gives the implicit equation for p_P (10). To show that this is an equilibrium for the firms, it suffices to show that $1 - F^J(p_P) - Jp_P F^{J-1}(p_P)f(p_P)$ is positive (and hence the best-response function kinks down at p_P).

Define

$$p^\dagger = \frac{1 - F^J(p^\dagger)}{JF^{J-1}(p^\dagger)f(p^\dagger)}$$

That is, p^\dagger maximizes the on-platform profits facing the distribution of consumer types F^J . Since F^J satisfies the monotone-likelihood ratio property (MLRP) with respect to F , it follows that the inverse hazard rate of F^J is larger than the inverse hazard rate of F , so $p^\dagger > p_M$. Additionally, p_P can be rewritten as equal to the inverse hazard rate of $(1 - \lambda)F + \lambda F^J$. Hence, it suffices to argue that $p_M < p_P < p^\dagger$. The first inequality follows since the likelihood ratio of $(1 - \lambda)F + \lambda F^J$ with respect to F is just $(1 - \lambda) + \lambda JF^{J-1}(v)$ which is increasing in v . The second inequality follows since the likelihood ratio of F^J with respect to $(1 - \lambda)F + \lambda F^J$ is also monotonically increasing in v .

Finally, to show that the pricing policy was optimal, we follow the same argument as Theorem 3. Consider the vertically-integrated firm problem, with the constraint that the firm cannot offer a different price to the segments determined by the efficient steering policy. Since steering is efficient, the vertically integrated firm chooses an off-platform price \bar{p} and an on-platform price p such that:

$$\max_{\bar{p} \geq p} \frac{1 - \lambda}{J} p(1 - F(p)) + \frac{\lambda}{J} (p(1 - F^J(p)))$$

Note that if the optimization were unconstrained, the values would be $\bar{p} = p_M$, $p = p^\dagger$. But we have shown earlier that $p_M < p^\dagger$, which violates the constraint. Hence, we consider the Lagrangian relaxation:

$$\max_{\bar{p}, p} \frac{1 - \lambda}{J} p(1 - F(p)) + \frac{\lambda}{J} (p(1 - F^J(p))) + \xi(\bar{p} - p)$$

Since ξ cannot be zero (as we argued above, the unconstrained optimum is infeasible), we must therefore have $\bar{p} = p$, and hence

$$\frac{1 - \lambda}{J} p(1 - F(p)) + \lambda \int_{\bar{p}} F^{J-1}(v) dF(v) = 0$$

Hence, the price set by the vertically integrated firm on- and off-platform is p_P . Since the transfers hold the firms to their outside option, this policy maximizes platform revenue.

References

- AGGARWAL, GAGAN, ASHWINKUMAR BADANIDIYURU, AND ARANYAK MEHTA (2019): “Autobidding with Constraints,” in *International Conference on Web and Internet Economics*, Springer, 17–30.
- ALI, MUHAMMAD, PIOTR SAPIEZYNSKI, MIRANDA BOGEN, ALEKSANDRA KOROLOVA, ALAN MISLOVE, AND AARON RIEKE (2019): “Discrimination through optimization: How Facebook’s Ad delivery can lead to biased outcomes,” *Proceedings of the ACM on human-computer interaction*, 3, 1–30.
- BALSEIRO, SANTIAGO R, YUAN DENG, JIEMING MAO, VAHAB S MIRROKNI, AND SONG ZUO (2021): “The Landscape of Auto-Bidding Auctions: Value Versus Utility Maximization,” in *Proceedings of the 22nd ACM Conference on Economics and Computation*, 132–133.
- BALSEIRO, SANTIAGO R AND YONATAN GUR (2019): “Learning in repeated auctions with budgets: Regret minimization and equilibrium,” *Management Science*, 65, 3952–3968.
- BAR-ISAAC, HESKI AND SANDRO SHELEGIA (2020): “Search, Showrooming, and Retailer Variety,” Tech. rep., CEPR Discussion Paper No. DP15448.
- (2022): “Monetizing steering,” Tech. rep., Centre for Economic Policy Research.
- BERGEMANN, DIRK AND ALESSANDRO BONATTI (2022): “Data, Competition, and Digital Platforms,” Tech. Rep. 2343, Cowles Foundation for Research in Economics.
- BERGEMANN, DIRK, BENJAMIN BROOKS, AND STEPHEN MORRIS (2015): “The Limits of Price Discrimination,” *American Economic Review*, 105, 921–957.
- CALZOLARI, GIACOMO AND VINCENZO DENICOLÒ (2015): “Exclusive contracts and market dominance,” *American Economic Review*, 105, 3321–51.
- CELIS, ELISA, ANAY MEHROTRA, AND NISHEETH VISHNOI (2019): “Toward controlling discrimination in online ad auctions,” in *International Conference on Machine Learning*, PMLR, 4456–4465.
- COMPETITION & MARKETS AUTHORITY (2020): “Online platforms and digital advertising,” Tech. rep., UK Government.

- DE CORNIÈRE, ALEXANDRE AND ROMAIN DE NIJS (2016): “Online Advertising and Privacy,” *Rand Journal of Economics*, 47, 48–72.
- DECAROLIS, FRANCESCO, MARIS GOLDMANIS, AND ANTONIO PENTA (2020): “Marketing agencies and collusive bidding in online ad auctions,” *Management Science*, 66, 4433–4454.
- DECAROLIS, FRANCESCO, GABRIELE ROVIGATTI, MICHELE ROVIGATTI, AND KSENIA SHAKHGILDYAN (2022): “Artificial Intelligence, Algorithmic Bidding and Collusion in Online Advertising,” Tech. rep., Bocconi University.
- DENG, YUAN, JIEMING MAO, VAHAB MIRROKNI, HANRUI ZHANG, AND SONG ZUO (2022): “Efficiency of the First-Price Auction in the Autobidding World,” .
- DENG, YUAN, JIEMING MAO, VAHAB MIRROKNI, AND SONG ZUO (2021): “Towards Efficient Auctions in an Auto-Bidding World,” in *Proceedings of the Web Conference 2021*, 3965–3973.
- ELLIOTT, MATTHEW, ANDREA GALEOTTI, ANDREW KOH, AND WENHAO LI (2022): “Market Segmentation through Information,” Tech. rep., Cambridge University.
- FUCHS, W. AND A. SKRZYPACZ (2015): “Government Interventions in a dynamic Market with Adverse Selection,” *Journal of Economic Theory*, 158, 371–406.
- GOLREZAEI, NEGIN, ILAN LOBEL, AND RENATO PAES LEME (2021): “Auction design for roi-constrained buyers,” in *Proceedings of the Web Conference 2021*, 3941–3952.
- HAGHPANAH, N. AND R. SIEGEL (2022): “The Limits of Multi-Product Discrimination,” *American Economic Review: Insights*, 4, 443–458.
- KANORIA, YASH AND HAMID NAZERZADEH (2020): “Dynamic reserve prices for repeated auctions: Learning from bids,” *arXiv preprint arXiv:2002.07331*.
- LI, BIN AND YAHUI LEI (2023): “Mechanism Design for Ad Auctions with Display Prices,” *arXiv preprint arXiv:2303.13192*.
- LIAW, CHRISTOPHER, ARANYAK MEHTA, AND ANDRES PERLROTH (2022): “Efficiency of non-truthful auctions under auto-bidding,” *arXiv preprint arXiv:2207.03630*.
- MEHTA, ARANYAK (2022): “Auction design in an auto-bidding setting: Randomization improves efficiency beyond vcg,” in *Proceedings of the ACM Web Conference 2022*, 173–181.

- MIKLÓS-THAL, JEANINE AND GREG SHAFFER (2021): “Input price discrimination by resale market,” *RAND Journal of Economics*, 52, 727–757.
- NASR, MILAD AND MICHAEL CARL TSCHANTZ (2020): “Bidding strategies with gender nondiscrimination constraints for online ad auctions,” in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 337–347.
- NEDELEC, THOMAS, CLÉMENT CALAUZÈNES, NOUREDDINE EL KAROUI, VIANNEY PERCHET, ET AL. (2022): “Learning in repeated auctions,” *Foundations and Trends® in Machine Learning*, 15, 176–334.
- OSTROVSKY, M. AND M. SCHWARZ (2023): “Reserve Prices in Internet Advertising Auctions: A Field Experiment,” *Journal of Political Economy*, forthcoming.
- PHILIPPON, THOMAS AND VASILIKI SKRETA (2012): “Optimal interventions in markets with adverse selection,” *American Economic Review*, 102, 1–28.
- TIOLE, JEAN (2012): “Overcoming adverse selection: How public intervention can restore market functioning,” *American Economic Review*, 102, 29–59.
- VARIAN, HAL (1980): “A Model of Sales,” *American Economic Review*, 70, 651–659.
- (2022): “Advertising Costs and Product Prices,” *The Journal of Law and Economics*, 65, S419–S431.
- WANG, CHENGSI AND JULIAN WRIGHT (2020): “Search platforms: Showrooming and price parity clauses,” *RAND Journal of Economics*, 51, 32–58.