How Do Digital Advertising Auctions Impact Product Prices?*

Dirk Bergemann† Alessandro Bonatti‡ Nicholas Wu§

June 22, 2024

Abstract

We present a model of digital advertising with three key features: (i) advertisers can reach consumers on and off a platform, (ii) additional data enhances the value of advertiser-consumer matches, and (iii) the allocation of advertisements follows an auction-like mechanism. We contrast data-augmented auctions, which leverage the platform's data advantage to improve match quality, with managed-campaign mechanisms that automate match formation and price-setting.

The platform-optimal mechanism is a managed campaign that conditions the on-platform prices for sponsored products on the off-platform prices set by all advertisers. This mechanism yields the efficient on-platform allocation but inefficiently high off-platform product prices. It attains the vertical integration profit for the platform and the advertisers, and it increases off-platform product prices while decreasing consumer surplus, relative to data-augmented auctions.

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." An extended abstract appeared in the Proceedings of the 24th ACM Conference on Economics and Computation (EC '23). We thank the editor, Andrea Galeotti, and four referees for productive suggestions. We also thank Santiago Balseiro, Gianluca Binelli, Michael Ostrovsky, and Balu Sivan for helpful comments, and Roi Orzach for excellent research assistance.

†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

Digital advertising facilitates the matching of consumers and advertisers online. Large platforms utilize their extensive consumer data to connect online shoppers with their preferred firms and products. In turn, advertisers join these platforms in order to target a wide range of potential consumers beyond their existing customer base. As a result, sponsored content is ubiquitous on the Internet: advertising makes up nearly all the revenue of search engines and social media platforms, a growing fraction of the revenue of retail platforms such as Amazon and Instacart, and a large fraction of other retail platforms' revenue, such as Alibaba's Taobao marketplace.

The role of these platforms' proprietary datasets becomes apparent when we decompose the value of digital paid traffic across the web. Google, Meta, Amazon, and other platforms place advertising *directly* on their own sites (through sponsored search results, stories, and products) and also serve as *intermediaries* that place advertising on third-party sites. The most recent data from public filings show that nearly 60% of the worldwide digital advertising revenue (which exceeds $600 billion) accrued on the platforms' own websites (Lebow, 2023). In particular, Google received revenues of $191 billion from digital advertising on its own sites, (e.g., google.com, youtube.com, etc.) and merely $33 billion from ad placement on third-party websites.[1] Thus, the majority of revenue accrues precisely where the deployment of proprietary data is completely unrestricted and not accessible to competing marketers. Indeed, the evolution of this marketplace suggests a reversal in the traditional assumptions on asymmetric information in digital advertising.

As the market for digital advertising has grown and become more complex, the prevailing mechanisms by which platforms sell ads have also shifted. Digital platforms increasingly act as intermediaries that run managed campaigns for advertisers, who set a fixed budget, specify high-level objectives for their campaigns, and leave the task of bidding to "auto-bidding" algorithms offered by the platform. Recent estimates suggest that over 80% of digital advertising is now generated by managed campaigns (Deng et al., 2022a,b). For example, over 80% of Google advertisers were using automated bidding in 2023.[2]

[1] See Google's 2023 financial report to the Securities and Exchange Commission, available at `https://www.statista.com/study/163755/alphabet-google-annual-report-2023`.

[2] See `https://ads.google.com/home/measurement/bidding/`. Furthermore, Competition & Markets Authority (2020, §5.201, §5.76) reported that 40-50% of Google's 2019 search advertising revenue in the UK came from advertisers using automated bidding. Finally, 90-100% of UK advertisers on Facebook were using the default auto-bidding feature, which does not allow advertisers to specify a maximum bid.

In this paper, we provide an equilibrium treatment of how data-intensive mechanisms for selling advertising impact product prices and welfare both on and off the platform. Our approach takes two fundamental aspects of digital advertising into account. First, consistent with the revenue breakdown described above, platforms possess valuable data that can enhance the matching efficiency. Second, advertisers have parallel sales channels, i.e., they can reach their customers on and off the platform.

We consider a monopolist digital platform that sells access to its users. Advertisers determine their pricing strategy on and off the platform and their advertising strategy on the platform. On-platform consumers act as *shoppers* and choose the product that offers the highest net value. Because these consumers compare the advertised offers to all firms' off-platform prices, advertisers endogenously behave as if under a "showrooming" constraint: they wish to 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 the option of charging higher personalized prices to on-platform shoppers.[3]

A key innovation in our model is that the platform actively influences the firms' advertising campaigns. With access to the platform's data, advertisers can offer prices that reflect the consumers' willingness to pay. This form of price discrimination broadens the market and enhances the efficiency of matching on the platform. Off the platform, advertisers lack additional data and offer a uniform price. We contrast two main mechanisms for allocating advertising space on the platform: *data-augmented auctions* and *managed campaigns*. Both these mechanisms, as well as our simple advertising model, are of course simplifications. Throughout the paper, we discuss the relevance of each of our modeling choices to real-world advertising markets: our model of advertised prices in Section 2; the auction format in Section 3; and the managed-campaign mechanism in Section 4.

Our model demonstrates how any analysis of the 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.

[3] In 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 faces the showrooming constraint directly. The advertiser wants to pay for the listing only if it leads to a sale, as the offline transaction would have occurred without the advertising. In 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.

1.2 Results

We begin our analysis with a second-price auction for a single advertising slot where the platform augments the bidders' information by soliciting bids based on the match values with each consumer. We refer to this as *data-augmented bidding*: each advertiser submits a bid for the slot and a price at which to offer its product if it wins the slot.

We derive the optimal bidding and pricing strategy of the advertisers. On 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 (Theorem 1). Additionally, each advertiser must set the price at which to offer its product to *loyal customers* off the platform. In equilibrium, the advertisers raise their off-platform prices, relative to the prices 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. Consequently, the off-platform prices increase with the number of on-platform shoppers.

Next, we introduce the concept of a *managed campaign*. In this more centralized mechanism, the platform proposes a steering policy and a pricing policy for each advertiser's product on the platform. Contextually, the platform requests a fixed fee from each advertiser, which we can interpret as a required advertising budget. Each advertiser simultaneously decides whether to enter into the managed campaign or not, and how to price its product off the platform. We show that the platform optimizes its revenue by matching firms and consumers efficiently and by offering a *best-value pricing* policy. This policy ensures the efficient firm always makes the offer with the best value to the consumer, even if its competitors deviate in their posted prices (Theorem 2). In doing so, the platform weakens competition and leads the firms to raise their posted prices off the platform in order to extract more surplus from online consumers.

Best-value pricing is not only revenue-optimal for the platform; the joint producer surplus attains the vertical integration benchmark where one firm controls all the advertisers and the platform (Theorem 3). In consequence, the posted prices off-platform are higher than under the data-augmented auction (Theorem 4). By comparing the prices charged to consumers and the advertising costs across these two mechanisms, we can then quantify a notion of pass-through (Proposition 3).

1.3 Policy Relevance

The digital platforms that offer managed campaigns also enjoy significant market power, which has raised regulatory concerns. In a recent report, the UK regulator 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." (Competition & Markets Authority, 2020, §6.19, §6.20.)

The Competition & Markets Authority (2020, Chapter 5 and Appendix Q) sets forth the principle that platforms should act in customers' best interests when making choices on their behalf. Our baseline model raises the concern that automated bidding options in Google and Facebook could be used to increase platform revenues to the consumers' detriment instead.

We therefore deploy our model to examine two competition- and privacy-policy interventions. The first policy we consider restricts the platform's auto-bidding algorithms by requiring the pricing and steering policies to be *independent* of all off-platform posted prices. We show that *any* independent managed campaign that steers consumers efficiently leads to lower on-platform prices than the fully optimal managed campaign (Theorem 5). In particular, limiting the signals that the pricing policy can use restores the possibility for on-platform consumers to be poached by other firms through off-platform price cuts. This force fosters competition and benefits consumers on both sales channels.

The second policy we consider is a privacy restriction that prevents the platform from steering consumers and setting prices on the basis of the consumers' detailed data. Instead, we allow the platform to condition its steering 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 perfectly 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 (Proposition 4). The privacy restriction reduces off-platform prices compared to the benchmark of a managed campaign, but may reduce total surplus on the platform, because low-value consumers no longer receive personalized discounts.

1.4 Related Literature

Our paper contributes to the literature on online advertising auctions. Recent work in this field studies learning in repeated auctions (Balseiro and Gur, 2019; Kanoria and Nazerzadeh, 2020; Nedelec et al., 2022), collusion (Decarolis et al., 2020, 2022), discriminatory effects (Celis et al., 2019; Ali et al., 2019; Nasr and Tschantz, 2020), and the role of stochastic quality scores (Ostrovsky and Skrzypacz, 2022). Our goal, instead, is to compare auctions with other allocation mechanisms under a given information structure and with parallel sales channels. As such, our approach is distinct from, Bar-Isaac and Shelegia (2022), who compare auctions and auto-bidding mechanisms in a single market under exogenous limits to the ability to steer and to price discriminate. Motta and Penta (2022) study a model of targeted bidding where the number of organic search results is fixed. In their setting, sponsored content may crowd out organic information when the same firm wins both types of links. This limits competition, facilitates market segmentation, and reduces welfare.

Several papers (Golrezaei et al., 2021; Liaw et al., 2022; Mehta, 2022; Deng et al., 2022b) study online auction design in the presence of autobidders and return-on-investment constraints.[4] Our setting adds a dimension related to advertised prices: firms submit both bids for a sponsored link and tailored prices to offer consumers. While Li and Lei (2023) also investigate mechanisms that allow for advertised prices, we further explore the interaction of these mechanisms with off-platform activity.

Our paper also relates to the literature on information design in auctions and markets. In particular, Bergemann et al. (2015), Haghpanah and Siegel (2022), and Elliott et al. (2024) 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.

As in Varian (1980), the advertisers in our model face two segments of consumers, *shoppers* on the platform and *loyals* off the platform. The design of the auction is therefore subject to the showrooming constraint, i.e., to competition from a separate and distinct market. Earlier papers on "partial mechanism design" or "mechanism design with a competitive fringe" studied mechanism design in settings where the agents' outside option consists of participating in alternative markets, e.g., Philippon and Skreta (2012), Tirole (2012), Calzolari and Denicolò (2015), and Fuchs and Skrzypacz (2015).

[4] A recent literature on auto-bidding algorithms allows for objective functions by the bidders outside of the class of quasilinear utility models common 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.

The showrooming constraint in our model is related to a growing literature on digital platforms with competing advertisers or multiple sales channels. Recent contributions include de Cornière and de Nijs (2016), Ganuza and Llobet (2018), Bar-Isaac and Shelegia (2020), Miklós-Thal and Shaffer (2021), and Wang and Wright (2020). In our setting, advertisers deter showrooming in order to capture the added value of making data-augmented offers on the platform. In parallel work, Bergemann and Bonatti (2023) study on- and off-platform competition with multi-product firms and nonlinear pricing. They focus on the implications of managed campaigns for equilibrium product quality, relative to our paper's exploration of showrooming and its impact on pricing strategies in the off-platform markets.

Finally, a recent contribution by 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 his separable model, an exogenous increase in advertising costs does not necessarily lead to an increase in product prices.

2 Model

Payoffs and Information There are J advertisers (or firms) indexed by $j = 1, 2, ..., J$, each selling unique indivisible products and a single digital platform. Each firm's production cost is normalized to zero. There is a unit mass of consumers, each demanding a single product. The willingness to pay v_j for each firm's product is drawn independently across consumers and firms according to a distribution function F that admits a strictly positive density f on its support $V = [\underline{v}, \bar{v}] \subset \mathbb{R}_+$. The consumer's *value* is given by the vector of willingness to pay

$$v = (v_1, ..., v_J) \in V^J \subset \mathbb{R}^J_+.$$

The utility for a consumer with value v of purchasing product j at price p_j is given by

$$v_j - p_j.$$

Initially, values are observed by the consumers and by the platform, but not by the firms.[5]

[5] 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 information only.

Firms and Platform The platform presents consumers with a single "sponsored" result followed by a list of non-sponsored products. The platform allocates the sponsored position using either a *data-augmented auction* or a *managed campaign*. We describe these two mechanisms in Sections 3 and 4, respectively, and we connect them to current practices in digital advertising markets. Under either mechanism, an on-platform consumer with value v receives a personalized offer to buy some firm j's product at a price $p_j(v)$. Thus, the firm in the sponsored slot can condition its price on the J-dimensional consumer value. In addition to the on-platform prices $p_j(v)$, each firm j posts a price $\bar{p}_j$ for its product off the platform.

On-platform Consumers A measure $\lambda \in [0, 1]$ of consumers are on-platform "shoppers." These consumers observe $J+1$ prices: the advertised price $p_j(v)$ by the firm j that is awarded the sponsored slot, as well as the prices $\bar{p}_k$ posted by all firms $k = 1, ..., J$. We can view these prices as organic results shown by the platform, or equivalently interpret the model as allowing for free search: only a "sponsored" firm can target a price offer to an on-platform consumer, but the consumer can search and find the prices posted by any firm.[6]

Under either interpretation, each firm j is subject to a *showrooming constraint* when setting its on-platform prices: for all v, the prices it advertises on the platform must satisfy $p_j(v) \leq \bar{p}_j$. Thus, in this model, firms offer the lowest prices on the platform, regardless of whether the platform imposes price-parity or most-favored-nation clauses.[7]

Off-platform Consumers The remaining $1 - \lambda$ measure of consumers are "loyals" who visit only a single firm off the platform (e.g., its physical store or website). The off-platform consumer population is divided into J captive segments of size $(1-\lambda)/J$. Segment j considers firm j only: these consumers buy if and only if the off-platform price $\bar{p}_j$ is lower than their willingness to pay v_j. Figure 1 summarizes our model.

Digital Advertising through the Lens of the Model Digital platforms offer a variety of advertisement formats, such as sponsored links, images, or videos. The content, often a product-price pair selected from the advertiser's portfolio, can differ across media channels

[6] The dual presence of sponsored and organic search describes most closely the practice of platforms such as Google or Amazon. However, in our model, the key feature is having at least some organic search results, not necessarily having both types under one platform. Thus, our model applies even to social networks like Meta, which have less search functionality compared to Google or Amazon.

[7] We use the upper bar notation for off-platform prices because the posted price $\bar{p}_j$ is an upper bound on the amount that any on-platform consumer will pay for firm j's product. As we will see in Sections 3 and 4, the platform would prefer to stop shoppers from accessing off-platform prices. In practice, however, this is not always possible or profitable.

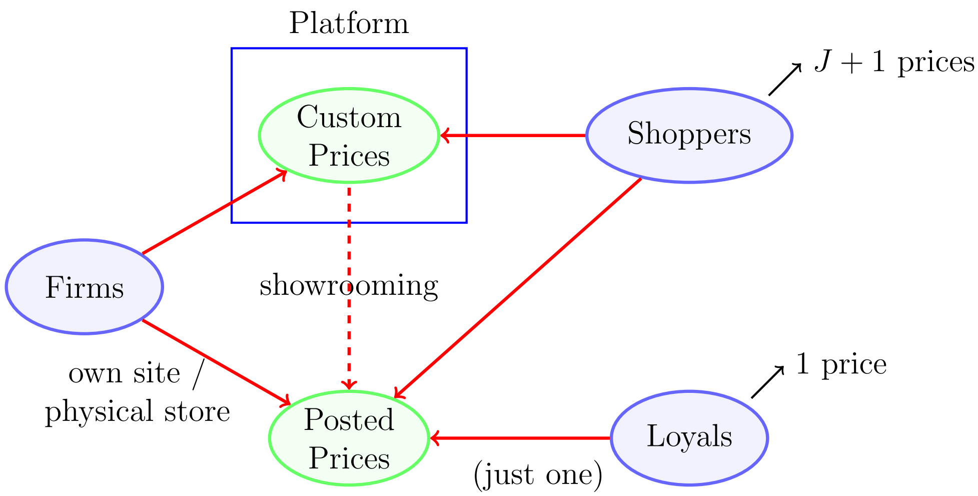

Figure 1: Model Depiction

and is tailored to individual consumers. Advertisers face three key decisions: identifying the target users, selecting the appropriate advertisement for each user, and determining the bid for each user's attention. The best strategy depends on the platform's nature and the advertiser's product line. For instance, a brand with multiple product lines would adopt a different approach than a single-product firm, adjusting its campaign according to the type of platform, e.g., search engines, social networks, or third-party publishers.

In our model, each firm offers a single product with fixed characteristics. Thus, the content of an advertisement is limited to a specific brand and to a personalized price. As such, our model is certainly an abstraction from the rich practice of digital advertising, where brands have multiple product lines and products have many features.

At the same time, our model allows us to capture two crucial aspects features of real-world digital advertising: the platform's ability to match consumers with their preferred firms and the value created through personalized pricing, which offers discounts to lower-value consumers who would not purchase at the monopoly price.

The model also accommodates a broader interpretation where each firm offers a range of products varying in quality and price. The platform's information enables firms to guide each consumer to a different quality-price pair within their product line, a process known as *product steering*. This process combines value creation and extraction, similar to our single-product model. As the variation in product quality diminishes (i.e., the products of each firm become more alike), product steering becomes akin to personalized pricing.[8]

[8] See Bar-Isaac and Shelegia (2022), Bergemann and Bonatti (2023), and Teh and Wright (2022) for recent models of product steering on digital platforms.

Finally, in our model, personalized pricing is exclusive to the platform and only applies to the firm that wins the sponsored slot. This is due to the consumers' unit demand and the significant difference in the information each firm possesses on- and off-platform. However, in real-world scenarios, multiple forms of price discrimination, such as market segmentation and nonlinear pricing, can occur both on and off the platform. In that sense, our model accentuates the differences between these two sales channels.

3 Data-Augmented Auctions

In this section, the platform runs an auction to determine which firm makes a personalized offer to each consumer. The platform provides the advertisers with information about the characteristics of the consumer, summarized by the vector of values v. Because the advertisers can make their bidding and pricing decisions contingent on the information disclosed by the platform, we refer to this mechanism as *data-augmented bidding.*

3.1 Data-Augmented Bidding: Mechanism

The platform runs a second-price auction for each realized consumer value v separately, breaking ties in favor of the consumer's most preferred firm. In each auction, the platform enables the advertisers to condition bids and sponsored prices on the consumer's value. Formally, each advertiser j adopts a bidding strategy

$$b_j : V^J \to \mathbb{R}_+,$$

and a sponsored pricing strategy

$$p_j : V^J \to \mathbb{R}_+.$$

This game proceeds as follows. First, all firms (simultaneously) post off-platform prices $\bar{p}_j$. Second, each firm j submits a bid function $b_j : V^J \to \mathbb{R}_+$ and a sponsored price function $p_j : V^J \to \mathbb{R}_+$ to the platform. Third, a consumer value v is realized, and a second-price auction (with no reserve) determines which firm j and which price $p_j(v)$ are advertised to the consumer. We characterize the symmetric Bayesian Nash equilibria of the bidding and pricing game among the advertisers.

3.2 Data-Augmented Bidding: Practice

Manual bidding is the original mechanism for selling advertising online and is still in use, though far less common.[9] When bidding manually, an advertiser typically specifies their willingness to pay for a click on a search result, display ad, or sponsored product listing. The advertisers can also modify their bids and their messages according to the platform's information on each consumer. Thus, the platform monetizes its data through an indirect sale of information (Admati and Pfleiderer, 1990; Bergemann and Bonatti, 2019), whereby advertisers can act *as if* they had direct access to the consumers' characteristics. In our model, where the platform has complete information about the consumer's preferences, the entire value profile v acts as a *targeting category*.

The rules for the allocation of sponsored placements vary across digital platforms and publishers. Broadly speaking, second-price mechanisms are used by the digital platforms on their own websites, e.g., by Google for sponsored search on google.com (Edelman et al., 2007), by Meta on its social networks, and by Amazon for its sponsored product listings. By contrast, the pricing of display advertising by third party publishers such as nytimes.com or wsj.com, which is often mediated by the digital platforms, has recently seen a transition from second price auctions to first-price auctions.[10] In what follows, we focus on the second-price data-augmented auction for its prevalence in digital advertising and its simplicity.[11]

3.3 Data-Augmented Bidding: Equilibrium

To help characterize the firms' equilibrium bidding and pricing strategies for this setting, we first establish a useful property. Proposition 1 below shows that, regardless of the prices posted by the firms off the platform, a bidding equilibrium in undominated strategies results in a *symmetric and efficient* assignment of on-platform consumers. In other words, each on-platform consumer sees a sponsored offer from the firm they like best.

---

[9]See, for example, `https://support.google.com/google-ads/answer/2390250`. At the same time, Google also suggests to advertisers that setting bids manually may result in lower performance, as reported in `https://growthmindedmarketing.com/blog/google-ads-mistakes-new-campaigns`.

[10]This shift is largely due to the organization of this market through ad exchanges (Goke et al., 2022). In recent work, Ostrovsky and Skrzypacz (2022) establish several properties of the generalized first-price auction when stochastic quality scores are used to weigh bids.

[11]Yet the details of the auction do not matter for our characterization of equilibrium bids and prices. Indeed, as the platform enables the firms to bid in a complete-information auction for each consumer type, both the winner and the price paid for each v are identical in a first- and second-price auction. Thus, the platform revenue and the equilibrium posted price are also equivalent for our data-augmented auctions.

Proposition 1 (Efficient Bidding Outcome)
Fix a profile of posted prices $\overline{p}$ *and 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 in undominated strategies.*

Proof of Proposition 1 Fix the vector of posted prices $\bar{p}$ and consider a second-price auction for a consumer with value v. Note first that, because bids can condition on v, this auction is effectively run under complete information; hence, it is weakly dominant for each firm to bid its valuation for the consumer. Furthermore, a firm's valuation for a consumer is equal to the price it would charge if it won the sponsored slot, $b_j(v) = p_j(v)$. Having observed all posted prices $\bar{p}$, each firm knows that consumer v has the option to buy from the most attractive off-platform offer,

$$\underline{u}(v, \bar{p}) \triangleq \max_{k=1,\ldots,J} (v_k - \bar{p}_k)_+ \, , \tag{1}$$

where $(\cdot)_+$ denotes the nonnegative part throughout the paper. Given this outside option, firm j can offer consumer v (and therefore bid)

$$p_j(v) = b_j(v) = (v_j - \underline{u}(v, \bar{p}))_+ \, ,$$

which in particular implies that the showrooming constraint $p_j(v) \le \bar{p}_j$ is satisfied. Because the outside option $\underline{u}$ in (1) is common to all firms, the highest-value firm $j = \arg\max_k v_k$ can offer consumer v the highest price p_j and still make a sale on the platform. Consequently, firm j also makes the highest bid b_j. $\square$

Proposition 1 allows us to separate the outcome of the bidding stage from the posted prices: the equilibrium matches in the bidding game are invariant with respect to the posted prices. Our main result in this section (Theorem 1), uses this property to characterize the unique symmetric equilibrium (in undominated strategies) of the data-augmented auctions and the associated posted price p_B.

Theorem 1 (Symmetric Equilibrium)
There exists a symmetric equilibrium in undominated strategies. In any such equilibrium:

1. *Consumer v receives and buys a sponsored offer from $j = \arg\max_k v_k$.*

2. *Each firm k posts price $\bar{p}_k = p_B$ satisfying*

$$(1-\lambda)(1 - F(p_B) - p_B f(p_B)) + \lambda J \int_{p_B}^{\bar{v}} F^{J-1}(v - p_B) dF(v) = 0. \tag{2}$$

   Further, when there is a unique solution to (2), the symmetric equilibrium in undominated strategies is unique.

3. *Firm $j = \arg\max_k v_k$ bids $b_j(v) = p_j(v) = \min(v_j, p_B)$.*

4. *Firm $j \neq \arg\max_k v_k$ bids $b_j(v) = p_j(v) = (v_j - \max_k(v_k - p_B)_+)_+$.*

The proofs of all results, unless noted otherwise, are collected in the Appendix. To gain intuition for the equilibrium posted prices, consider the case of a single bidder ($J = 1$). Because the losing bid is zero, the profit of the firm is given by

$$(1-\lambda)(1 - F(p))p + \lambda \int_{\underline{v}}^{\bar{v}} \min\{v, p\} dF(v). \tag{3}$$

In other words, the firm offers personalized discounts to λ on-platform consumers and posts a single price for the remaining $1-\lambda$ consumers. This formulation of the personalized pricing problem first appears in Ganuza and Llobet (2018).

In our setting with J competing firms, equation (2) is the first-order condition for the competitive analog to the monopoly profit (3). The equilibrium posted price p_B balances the winning firm's profit on the two sales channels. By showrooming, the posted price sets an upper bound on the prices that can be advertised to the on-platform consumers. Therefore, the potential to price discriminate more effectively on-platform pushes firms to raise their posted prices. This effect is captured by the second term in the first-order condition (2), which is positive.

While it is intuitive that higher posted prices enable higher advertised prices, equation (2) illustrates a more nuanced, important property of data-augmented auctions. A marginal increase in $\bar{p}_j$ above p_B benefits firm j only if (i) a consumer values firm j's product at least p_B, and (ii) a consumer values all other brands k less than $v_j - p_B$, so that the second highest bid $b_k(v)$ is nil. If the second condition is not met, i.e., if the auction is sufficiently competitive, the second highest bid is given by $b_k(v) = \bar{p}_j + v_k - v_j$ as in part (4.) of Theorem 1. This bid equals the price that firm k would advertise if it won the auction. Critically, firm k's bid increases one-to-one with firm j's posted price $\bar{p}_j$: firm k bids more aggressively for consumer v when winning the auction would enable it to charge a higher price and still make a sale. Thus, a higher posted price relaxes showrooming but makes a firm a softer competitor in the more competitive auctions, which dampens the benefits of raising prices.

3.4 Welfare Implications

We now discuss the welfare implications of data-augmented bidding. Theorem 1 shows that the on-platform allocation is socially efficient: every consumer participates and buys their favorite product. Relative to the on-platform channel, the off-platform market suffers from two sources of inefficiency: first, consumers are loyal to a random firm, i.e., they might be unaware of the existence of a firm they prefer; and second, since firms optimally post a single off-platform price, those consumers with values below the posted price do not buy at all.

Turning to the welfare implications for consumers, part (3.) of Theorem 1 shows that the winning firm extracts all consumer surplus on the platform, up to the equilibrium posted price. Thus, the expected surplus of off- and on-platform consumers are given by

$$CS_{\text{off}}(p_B) = \int_{p_B}^{\bar{v}} (v - p)\, dF(v), \quad \text{and} \quad CS_{\text{on}}(p_B) = \int_{p_B}^{\bar{v}} (v - p)\, dF^J(v). \tag{4}$$

On both channels, only consumers with values above p_B obtain a positive surplus.

To capture the effect of the platform on consumer surplus, we then consider the posted prices. We first define p_M as the monopoly price for distribution F,

$$p_M \triangleq \arg\max_p p\,(1 - F(p), \tag{5}$$

and we assume $p_M > \underline{v}$ throughout.

All firms would post price p_M if they had a loyal off-platform population only, as can be seen by setting $\lambda = 0$ in (2). For any $\lambda > 0$, the second term on the right-hand side of (2) pushes the equilibrium price p_B above p_M. Formally, part (1.) of Proposition 2 uses a monotone comparative statics argument (which we shall invoke repeatedly) to show that the equilibrium price is increasing in λ. Therefore, posted prices are larger than the monopoly price p_M. We trace out the welfare implications of this result in part (2.) of Proposition 2.

Proposition 2 (Posted Prices and Welfare Effects)

1. *The symmetric equilibrium posted price p_B is increasing in λ.*

2. *Off-platform per capita total surplus and consumer surplus are both decreasing in λ.*

In traditional models of search stemming from brick and mortar stores with non-posted prices, the increased presence of consumers who obtain more quotes has a positive externality on the other consumers. Our model generates the opposite prediction because the growth of the platform is unambiguously harmful for off-platform consumers.

In contrast, the effect of λ on on-platform consumer surplus is more nuanced. Because every on-platform consumer is matched with their favorite firm (as captured by the distribution F^J), the expected welfare of an on-platform consumer in (4) is always larger than an off-platform consumer's. Moreover, every consumer gains weakly ex post (after learning their v) by joining the platform.

This creates an important participation externality, however, because more consumers joining the platform increases λ, which raises all off- and on-platform prices.[12] As $\lambda \to 1$, the equilibrium posted price $p_B \to \bar{v}$. This means total surplus is at the first-best level, but the firms extract all consumer surplus on and off the platform.[13]

4 Managed Advertising Campaigns

We now contrast data-augmented auctions with the more novel auto-bidding and managed-campaign mechanisms. In a *managed* advertising campaign, the platform determines which firm wins the sponsored slot for each consumer value and makes an offer to the consumer on behalf of that firm. The platform collects a fixed upfront fee for this service from each participating firm. In turn, the firms relinquish agency over the on-platform allocation process, but they still collect the resulting revenue and post the off-platform prices.

[12] In recent work, Kirpalani and Philippon (2021) and Bergemann et al. (2022) document the externalities that consumers impose on each other through their decisions to share data with a two-sided platform. Yu (2024) empirically explores the welfare effects of sponsored-product advertising on retail platforms, documenting a tradeoff between better information and higher equilibrium prices.

[13] Away from the limit as $\lambda \to 1$, the disparate effects of the participation externality and the efficient firm-consumer matching render a definite answer as to whether consumer surplus increases as the platform becomes larger impossible. Intuitively, when there are many firms, the matching benefit of joining the platform is larger, and consumer surplus may be initially increasing in λ; with few firms, it is possible that consumer surplus is always decreasing in the size of the platform.

4.1 Managed Campaigns: Mechanism

A managed campaign is a mechanism where the platform conditions the advertised products and prices on all available information: the consumer's value $v \in V^J$, the firms' participation decisions in the mechanism $a \in \{0,1\}^J$, and the posted prices $\bar{p} \in \mathbb{R}_+^J$. We thus consider the following extensive form:

1. The platform proposes a mechanism (s, p, T) to all firms, where $s : V^J \times \{0,1\}^J \times \mathbb{R}_+^J \to J$ is a *steering policy*, $p : V^J \times \{0,1\}^J \times \mathbb{R}_+^J \to \mathbb{R}_+$ is a *pricing policy*, and $T \in \mathbb{R}_+^J$ is a profile of fixed fees (advertising budgets).

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 firm j accepts the platform's offer, it pays a fee T_j. Its product is offered to a subset of on-platform consumers according to the steering policy s and priced according to the policy p.

In other words, the steering policy steers each consumer to a firm, depending on all firms' participation decisions, their posted prices, and the consumer's own value. The pricing policy maps those same variables into an advertised price. In the remainder of this section, we focus on a specific instance of a managed campaign and then show that these pricing and steering policies are revenue optimal for the platform.

Definition 1 (Best-Value Pricing)
The best-value pricing policy sets

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

where $j = s(v, a, \bar{p})$ is the firm selected by the platform's steering policy.

Correspondingly, the efficient steering policy selects the consumer's favorite firm among those that participate in the mechanism.

Definition 2 (Efficient Steering)
The efficient steering policy sets $s(v, a, \bar{p}) = \arg\max_j a_j v_j$.

When combined with efficient steering, the best-value pricing policy ensures that the consumer's favorite firm always offers the best deal, so that no other firm can poach the consumer by posting a lower price $\bar{p}_k$. In other words, the platform commits to competing with any firm that cuts its price. Indeed, the advertised prices $p_j(v)$ in (6) are exactly the ones that firm j would choose if it retained control over the on-platform prices.

4.2 Managed Campaigns: Practice

All three key elements of our model of managed campaigns connect to the current practices of large digital platforms. First, relative to data-augmented auctions, ex-ante fixed fees replace individual, per-auction payments. In the model, the fixed fees represent advertising budgets that firms submit to the platform. This is the predominant mechanism on pure advertising platforms, such as Google, Facebook, or Tiktok that match advertisers and consumers, but do not charge any transaction fees. In all these markets, the firms delegate the spending decisions to the platform, subject to constraints on the returns to their investment.[14]

By contrast, retail platform such as Amazon or Instacart typically receive revenue from a mixture of advertising and sales commissions. Significantly, Amazon's advertising revenue is catching up to those of Google and Meta (Konstantinovic, 2023), which suggests that the relevance of the advertising mechanisms is extending to retail platforms. Other retail platforms, most notably Alibaba, have very low sales commissions and generate most of their revenue from sponsored listings. In particular, Alibaba's Taobao operates as a fee-free consumer-to-consumer marketplace where users can pay to rank higher in the search results, thus generating all its revenue from advertising (Stapleton, 2021).

Second, the platform controls both the allocation of sponsored slots and the prices of the firms' products. Many advertisers run advertising campaigns that target different consumers with promotional offers, which can involve personalized prices, varieties, or product versions. For example, Meta's Advantage+ Catalog Ads automatically delivers relevant product recommendations to people based on their revealed intent. Meta describes this service as follows.

> "You can create a catalog with all your products and create one campaign that drives sales on your website or app. When someone expresses interest in an item from your catalog [or in the types of products or services you are offering], Meta can dynamically generate an ad for that person and deliver it automatically on mobile, tablet and desktop."[15]

[14] Our model is also consistent with this type of arrangement. See the discussion after Theorem 3.

[15] Meta offers two targeting options that reach customers across different stages of the buying journey, depending on whether they have manifested interest in an advertiser's own products or in their industry's products. See `https://www.facebook.com/business/help/397103717129942`.

Likewise, Amazon and Google offer portfolio-bidding strategies consisting of "AI-powered, goal-driven bid strategies that help you optimize bids across multiple campaigns," i.e., that choose which offers to target to which user.[16]

In our model (see the discussion at the end of Section 2), each firm sells a single product, and therefore the platform's choice of personalized advertising content reduces to a targeted promotional discount.

Third, the platform conditions the advertised prices on all off-platform posted prices. Real-world managed-campaign algorithms such as Google's Performance Max and Meta's Advantage+ can be viewed as implementing our static mechanism by adapting behavior over time. The connection is as follows. Google's algorithm "uses Google AI across bidding, budget optimization, audiences, creatives, attribution, and more."[17] Thus, the algorithm adjusts not only the automated bids, but also the creative content shown to each consumer in order to achieve the campaign goals. For example, if the algorithm detects a drop in clicks on a given advertisement by a certain consumer segment, it can advertise a cheaper product to those consumers, so to improve sales.[18] In practice, the adjustment process occurs gradually. In our model, this process is instantaneous: the platform modifies the advertised prices as soon as a firm deviates from the equilibrium posted price.

Table 1 offers two interpretations of our model that summarize the above discussion. The narrow interpretation is the focus of our model, the broader interpretation links our model to a more extensive set of practices and tools.

	Narrow	**Broad**
Product Line	Single product	Multiple products
Targeting	Personalized pricing	Product steering
Algorithm	Advertised price reacts to all posted prices	"Portfolio bids" and "catalog ads" select most profitable product
Timing	One-time pricing	Learning over time

Table 1: Mapping the model to real-world managed campaigns

[16]See https://support.google.com/google-ads/answer/6263058 for Google's portfolio bidding and https://advertising.amazon.com/blog/introducing-portfolios-for-sponsored-ads for Amazon's.

[17]For a detailed description, see https://support.google.com/google-ads/answer/10724817.

[18]In particular, Google's Smart Bidding "... continues to update your bidding algorithms to align with any corresponding shifts in performance. Fluctuations in performance [...] can also be influenced by external factors like seasonality or competition." See https://support.google.com/google-ads/answer/10970825.

4.3 Managed Campaigns: Equilibrium with Best-Value Pricing

We shall focus on symmetric equilibria of the managed campaign with efficient steering and best-value pricing where all firms participate. We now define and then characterize such equilibria (henceforth "full participation symmetric equilibria").

Definition 3 (Full Participation Symmetric Equilibrium)
A full participation symmetric equilibrium is a pair $(\bar{p}^*, t^*)$ such that:

1. *Each firm j posts price $\bar{p}^* = \arg\max_p \Pi_j(p, \bar{p}^*)$, where the firm's profits are given by*

$$\Pi_j(p, \bar{p}) := \frac{1-\lambda}{J} p(1-F(p)) + \lambda \int_{\underline{v}}^{\bar{v}} \int_{\underline{v}}^{v_j} \min(v_j, p, v_j - v_k + \bar{p})_+ dF^{J-1}(v_k) dF(v_j). \quad (7)$$

2. *The platform charges fee $t^* \leq \Pi_j(\bar{p}^*, \bar{p}^*) - \Pi_j^O$, where each firm's outside option is*

$$\Pi_j^O = \max_p \left\{ \frac{1-\lambda}{J} p(1-F(p)) + \lambda \int_p^{\bar{v}} pF^{J-1}(v_j - p) \, dF(v_j) \right\}. \quad (8)$$

A key consequence of best-value pricing and efficient steering is that firms are insulated from competition. As best-value pricing eliminates the incentives to compete in posted prices, each firm maximizes profits on their captive consumers and on the on-platform consumers that like its product the best; note that (7) is the firm's best-response profit when participating and when all other firms participate and set posted price $\bar{p}$.

The firms' outside option consists of not participating in the mechanism and posting a price that maximizes profit on the captive consumers and competes with the advertised prices set by the best-value pricing policy. This outside option yields the profit level in (8). The first term is the profit from selling to loyal consumers. The second term denotes the profit the deviating firm can obtain on-platform: upon rejecting the platform's offer, this firm can make a sale to any on-platform consumer with a sufficiently high value for its product. In this case, the best-value pricing policy, which attempts to sell the second-highest valued product to the consumer, charges a price of zero. Therefore, the deviating firm j that posted price p makes a sale only when the consumer's value v satisfies $v_j - p > v_k$, where v_k is the value for the best competitor.

Theorem 2 characterizes all full participation symmetric equilibria of the managed campaign with efficient steering and best-value pricing.

Theorem 2 (Best-Value Pricing Managed Campaign Equilibrium)
In any full-participation symmetric equilibrium:

1. *On-platform consumers v with $v_j = \max_k v_k$ buy from firm j at $p_j(v) = \min\{v_j, p_V\}$.*

2. *The posted price p_V is characterized by the following equation:*

$$(1-\lambda)(1 - F(p_V) - p_V f(p_V)) + \lambda J \int_{p_V}^{\bar{v}} F^{J-1}(v) dF(v) = 0. \tag{9}$$

When (9) *has a unique solution, the full-participation symmetric equilibrium is unique.*

The best-value pricing policy (6) ensures that each firm makes a sale to its favorite customers regardless of the posted prices. Therefore, part (1.) shows that each firm sets its posted price like a monopolist with exogenous market segments, subject to showrooming.

The characterization of the equilibrium p_V in (9) follows from the first-order condition for each firm's profit. Indeed, each firm posts a price that balances the profit off-platform with the relaxation of the showrooming constraint. In particular, the second term in (9) shows that a marginal posted price increase yields a one-for-one benefit to firm j when facing any consumer that values product j the most and values it more than p_V.

As it turns out, the best-value pricing campaign is revenue-optimal for the platform. Indeed, it admits an equilibrium in which the platform is able to charge a fee that attains an exogenous upper bound across all managed campaigns.[19]

Theorem 3 (Optimal Managed Campaign)
The highest-price, full-participation symmetric equilibrium:

1. *maximizes revenue for the platform among all steering and pricing policies;*

2. *attains the integrated (collusive) gross profit for the firms.*

The argument proceeds by considering the problem of a vertically integrated platform that jointly maximizes the profit of the 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 optimal joint solution is then decentralized

[19] Depending on the level of the platform's fee, the best-value pricing campaign may also admit symmetric equilibria with no seller participation. However, the platform could uniquely implement full participation by designing transfers in the divide-and-conquer style of Segal (2003) or using other techniques from the literature in contracting with externalities.

by charging a fixed fee that extracts all of the firms' surplus, net of an exogenous outside option for the firms.[20]

If we allowed on-platform prices to be negative, the platform would be able to worsen the firms' outside option further by ensuring that a non-participating firm never makes a sale on the platform. Even in this case, however, the optimal campaign cannot charge firms their entire revenues on the platform, because firms distort off-platform prices away from p_M. In this sense, the platform could frame the managed-campaign mechanism as delivering a positive return on investment.

4.4 Comparing Advertising Mechanisms

We now compare the equilibrium posted prices and the welfare properties of the data-augmented second-price auction and the optimal managed campaign. We refer to p_B and p_V as the highest symmetric equilibrium prices under data-augmented bidding and best-value pricing, respectively.

Theorem 4 (Welfare and Posted Price Comparison)
The posted price p_V in the optimal managed campaign is higher than the posted price p_B under data-augmented bidding:

$$p_V \geq p_B \geq p_M.$$

Total consumer surplus and total welfare are lower in the optimal managed campaign than under data-augmented bidding.

In our model, the impact of digital advertising auctions on product prices is entirely due to the different competitive responses under the data-augmented auctions and the managed-campaign mechanism. Theorems 1 and 2 showed that both data-augmented auctions and the optimal managed campaign yield an efficient matching of all on-platform consumers to firms. Moreover, under both mechanisms, the on-platform consumers buy from their favorite firm (say, j) at a price $p_j(v) = \min\{v_j, \bar{p}_j\}$. Therefore, both mechanisms create a common benefit of raising the posted price $\bar{p}_j$, namely to increase revenue on all consumers that like firm j best *and* value product j more than the posted price.

However, with data-augmented auctions, raising $\bar{p}_j$ raises all rivals' bids b_k by the same amount, because any firm $k \neq j$ that wins the auction can now charge a higher price and still

[20] 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 auto-bidding mechanism satisfies an ex-post participation constraint for every (winning and losing) bid.

induce the consumer to buy its product rather than shop for firm j's. (Recall the discussion after Theorem 1.) Therefore, raising the posted price $\bar{p}_j$ helps firm j only if all other bids are nil ($b_{k\neq j} = 0$), which occurs when the consumer v is willing to pay a large enough premium for firm j's product, i.e., $v_j > \bar{p}_j + \max_{k\neq j} v_k$. The latter effect is absent under the optimal managed campaign, where fixed fees replace variable, endogenous payments for each consumer. Therefore, raising posted prices is more profitable under managed campaigns.[21]

Finally, our model admits a formal notion of "advertising cost pass-through" driven by the mechanisms for selling ads, namely as platforms move from data-augmented bidding to managed campaigns. To compute this measure of pass-through, we fix the bargaining power of the platform to that of the sophisticated campaign; that is, the platform charges fees in order to hold firm profit to their outside option. We take these fees as a proxy for advertising cost, and compare the firm net transfer under data-augmented bidding to the transfer under the optimal sophisticated campaign. More precisely, let T_B be the total transfer paid by an individual firm under revenue-maximizing data-augmented bidding,[22] and let T_V be the total transfer paid by an individual firm under the sophisticated managed campaign. Then we define the "pass-through" of the change in mechanisms as:

$$\eta = \frac{p_V - p_B}{T_V - T_B}, \tag{10}$$

as how the change in advertising impacts the off-platform posted prices. We can thus characterize the pass-through more formally.

Proposition 3 (Advertising Mechanism Pass Through)
The pass-through rate satisfies $\eta > J$. The increase in advertising costs induced by a managed campaign relative to bidding are reflected by an amplified increase in off-platform prices.

Proposition 3 and numerical examples suggest that the cost pass-through is greater the more firms there are. Intuitively, prices differ more dramatically across mechanisms when more firms are present because the sophisticated managed campaign softens competition.

[21] Consistent with this intuition, the proof of Theorem 4 uses a monotone comparative statics argument that does not require assumptions on the distribution of consumer values.

[22] To fix the bargaining power of the platform, we consider the data-augmented bidding where the platform can charge a participation fee before bidding occurs. This results in no changes to the pricing outcome and makes the platform's bargaining power comparable.

5 Policy Interventions

In this section, we investigate the impact of potential interventions that a policymaker might impose on the platform. We specifically consider restricting the platform's ability to use sophisticated algorithms that respond to posted prices, and its ability to condition advertising and prices on the consumer' full value profile.

5.1 Competition Management

A first regulatory question is whether fully automated systems should be kept in check. The Competition & Markets Authority (2020, §6.15) expresses concerns that "Although both Google's and Facebook's core services can be accessed by consumers at no direct cost, consumers therefore nevertheless suffer financially from the exercise of market power." The alleged concern is that the platform's market power raises the cost of advertising, which is then passed on to consumers.

To address these concerns, we analyze a policy that limits the platform's active role in managing firm competition. In particular, we assume the platform's pricing and steering policies can condition on the consumer's full value profile, but not on the posted prices.

Current Practice Limiting auto-bidding algorithms to enable *rule-based bidding* only is an example of such a policy intervention. According to Amazon, rule-based bidding is an existing automated bidding strategy that "take[s] the guesswork out of adjusting bids," but lets advertisers introduce fixed rules for showing creatives to (and bidding on) specific consumer segments.[23] Likewise, Google's "Demand Gen" campaigns allow advertisers to manually select specific channels for ad display, "offering more control over where and how ads appear."[24] Letting advertisers retain partial control over the bidding rules necessarily slows down the algorithm's adjustment process. In our static model, we capture these algorithms by means of managed campaigns that do not react to evolving market conditions (as proxied by deviations in posted prices).

Independent Managed Campaign We now restrict the platform's pricing and steering policy space by removing the platform's ability to condition on off-platform prices. The

[23]For a complete description of the bidding strategies available on Amazon, see `https://advertising.amazon.com/help/GCU2BUWJH2W3A8Z7`.

[24]See `https://ads.google.com/home/campaigns/demand-gen`.

platform can now only propose a pricing policy $p : V^J \times \{0,1\}^J \to \mathbb{R}_+$ and a steering policy $s : V^J \times \{0,1\}^J \to J$ that depend on the value v and the participation decision of the firms.

In the previous section, we showed that the optimal managed campaign dampens competition between firms, resulting in higher posted prices off-platform than in data-augmented bidding. Theorem 5 shows that forcing the platform to price *independently* of the posted price decisions curtails the ability of the platform to soften competition. We denote the posted price off the platform induced by the independent campaign by p_I.

Theorem 5 (Independent Managed Campaigns)
Any independent managed campaign with efficient steering results in lower prices relative to the optimal campaign: $p_I < p_V$, higher social welfare, and higher consumer surplus.

The critical economic feature that an independent campaign introduces is the potential for consumers to be poached by other firms. In particular, since the on-platform prices cannot condition on the off-platform prices, a deviating price downwards by firm i can induce some consumers whose favorite firm is $j \neq i$ to buy from i instead. This downward pressure helps mitigate the platform's ability to soften competition.

To illustrate how an independent campaign restores a downward competitive pressure on prices, we present two examples where the platform steers consumers efficiently and prices as in the optimal managed campaign (i.e. $p_i(v) = \min(v_i, p_V)$). We show that the posted price p_I in the independent campaign can even fall below the monopoly price p_M. In such an independent managed campaign, the best-response problem of the firm simplifies to

$$\max_p \left[\frac{1-\lambda}{J} p(1-F(p)) + \lambda \int_{\underline{v}}^{\bar{v}} \min(v,p,\hat{p}) F^{J-1}(v - \min(v,p,\hat{p}) + p') dF(v) \right].$$

The first-order condition for p at or below $\hat{p}$ then implies that:

$$\frac{1-\lambda}{J}(1 - F(p) - pf(p)) + \lambda \int_p^{\bar{v}} (F^{J-1}(v) - p dF^{J-1}(v)) dF(v) = 0. \tag{11}$$

Note the presence of the term $p\, dF^{J-1}(v)$, which captures the poaching gain/loss from consumers who at the margin are nearly indifferent between firms.

Example $J = 2 : p_I > p_M$. 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. Computing the best-response price using the first-order condition (11), we obtain

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

Example: $J = 3 : p_I < p_M$. 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 consider three firms. In this case, we obtain

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

By adding one firm to the previous example, the competitive effect becomes stronger, and the posted price falls to a level below the monopoly price p_M. Note that the pricing in sophisticated managed campaign—which allows for perfect price discrimination up to the showrooming constraint—may not be the optimal independent pricing policy for the platform. In particular, as the example shows, this pricing policy induces stronger competition by firms and sometimes lower posted prices than even the monopoly price. Note that a pricing policy that offers a product for free to consumers whose values are all below p_M, and price p_M for a consumer's favorite firm otherwise, can induce an equilibrium posted price of p_M; however, such a pricing policy necessarily concedes rent to the consumer or reduces aggregate welfare. This illustrates the platform's trade-off under the independent pricing restriction: the more aggressive the price discrimination offered by the platform, the stronger the incentives for firms to undercut each other, and the lower the posted prices. However, to raise the posted prices, the platform must concede utility to the consumers; an optimal independent pricing policy must therefore balance these two forces.

5.2 Privacy and Data

We now assess the impact of privacy regulation by considering policies that limit the firms' access to the consumers' information. Specifically, we consider cohort-based privacy, which is a restriction in line with the recent Google Privacy Sandbox proposals to replace third-party cookies. Under this policy, the platform in our model informs the firms about the consumer's ranking of their products, without disclosing the consumer's exact value for any specific product.[25]

Formally, the platform's steering policy selects a firm to advertise to each *cohort* of consumers, and each consumer within a cohort has the same preference *ranking* over the J firms. In what follows, we maintain the efficient steering policy (i.e., the platform shows the consumer her favorite product), which yields exactly J distinct consumer cohorts. We then

[25]See the complete Google proposal at `https://privacysandbox.com/`. In this Section, we focus on exogenous restrictions on information disclosure. Voluntary information disclosure by the consumer is another important, though different dimension. See Ali et al. (2023) for a treatment of this question.

restrict the platform's pricing policy space to

$$p : J \times \{0,1\}^J \times \mathbb{R}_+^J \to \mathbb{R}_+.$$

Thus, the platform cannot price based on the consumer's individual value vector v, but it can condition the advertised price on the consumer's cohort, the firms' participation 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. We denote the resulting equilibrium off-platform price with *privacy protection* by p_P.

Proposition 4 (Cohort Privacy)
In the platform-optimal managed campaign with cohort privacy, the posted price is p_P with:

$$p_P = \frac{1-(1-\lambda)F(p_P) - \lambda F^J(p_P)}{(1-\lambda)f(p_P) + \lambda J F^{J-1}(p_P) f(p_P)}. \tag{12}$$

This managed campaign can be implemented by the platform pricing each segment at the lowest off-platform price: $p(i,\cdot,\bar{p}) = \min_i \bar{p}_i$. On path, the on-platform price is also p_P, and the equilibrium posted price p_P satisfies $p_M \le p_P \le p_V$.

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 on-platform shoppers who are matched to the efficient firm, i.e., their values are 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, but showrooming limits its ability to do so.

Proposition 4 also shows that the off-platform posted prices are lower than under the optimal managed campaign, which implies greater consumer surplus and total welfare off the platform. However, the inability to price discriminate on-platform means the privacy restriction reduces total welfare on the platform too, because low-value consumers are priced out. Hence, as the platform size λ grows, total welfare can be worse under privacy restrictions than under data-augmented bidding or managed campaigns. However, consumer surplus grows because low-value consumers' surplus is nil in both settings without privacy protection; that is, the loss in welfare comes entirely from reduced producer surplus.

6 Extensions and Robustness

We discuss two extensions of the basic model that speak to the robustness of the analysis. The first variation concerns the nature of the off-platform market, the second the nature of the platform, in particular the revenue model of the platform.

6.1 Off-Platform Competition

Previously, we modeled each firm as operating as a monopolist of a market segment off the platform. We now show that the analysis and consequent results extend to a more competitive structure in the off-platform markets. Suppose then that the off-platform consumers are divided into K markets, and each firm operates in one market only, so each off-platform market has $N = J/K$ firms, where N is assumed to be an integer.

Now, the posted price impacts the off-platform market slightly differently. In particular, by setting a posted price p when the competitors in the off-platform market set price p', the firm wins an off-platform consumer if and only if $v - p \geq v' - p'$, where v' is the value of the best competitor. Hence, the firm's profit off the platform is given by

$$\int_{p}^{\bar{v}} pF^{N-1}(v - p + p')dF(v). \tag{13}$$

In the absence of the platform, the symmetric equilibria of the game with payoffs (13) yield the oligopoly prices, which we denote by p_O^K. These are the prices the firms would charge if there were K segmented markets. The basic model considered $K = J$ segmented markets, so that each firm was acting as a monopolist in its market; thus for $K = J$, $p_O^J = p_M$. We now investigate how a more competitive off-platform market affects the behavior in the bidding and managed-campaign mechanism.

To this end, note that the on-platform profit terms in our previous analysis are not affected by changes in the off-platform market as suggested above. We now consider the symmetric equilibrium off-platform prices with competition in K market segments, both under data-augmented bidding (p_B^K) and under the fully optimal managed campaign (p_V^K). We obtain the following comparison.

Proposition 5 (Off-Platform Competition)
In the highest-price symmetric equilibria with off-platform competition in K markets, the off-platform posted prices satisfy

$$p_O^K \leq p_B^K \leq p_V^K.$$

Thus, the equilibrium ordering of the off-platform prices, and the corresponding welfare results are invariant to the structure of competition in the off-platform markets. In particular, the above ordering is same as in Theorem 4 where we observed: $p_M \leq p_B \leq p_V$.

6.2 Platform Revenue Models

In our model of managed campaigns, the platform requests an up-front participation fee or advertising budget. This aligns with current practices on pure advertising platforms such as Google, Facebook, and Tiktok. These platforms match advertisers and consumers but do not charge any transaction fees. By contrast, shopping or retail platforms such as Amazon and Instacart generate revenue from a mixture of advertising and sales commissions. One might thus wonder whether there are multiple, payoff equivalent mechanisms that attain the same platform revenue. Here, we focus on one such alternative in which the platform charges a constant transaction fee t_j to each firm and does not impose a fixed payment T_j. We maintain the features of the optimal managed campaign in Section 4, i.e., efficient steering and best-value pricing. We then show that for modest transaction fees, the firms' incentives for setting prices off the platform are exactly as in in Theorem 2.

Proposition 6 (Transaction Fees)
Suppose the platform proposes a transaction fee t_j for each sale of firm j on the platform. There exists a $\bar{t} > 0$ such that if the fee satisfies $t_j \leq \bar{t}$, then the off-platform prices $\bar{p}$ are the same as in the best-value managed campaign of Theorem 2.

In the current revenue model, we replaced the ex-ante advertising budget with a constant sales (or referral) fee. The key insight is that the referral fee does not influence the marginal incentives of the firm's off-platform pricing decision; under efficient steering and best-value pricing, each firm's volume of sales is constant regardless of the off-platform prices. Hence, provided the referral fee is not too large and the firm is still willing to participate, the pricing decision off the platform is the same as in Theorem 2. Thus, there are revenue models that are revenue equivalent to the platform-optimal managed-campaign mechanism.

Unlike the mechanism of Theorem 2, this result relies on the platform retaining control over the firms' advertised prices. If the platform delegated the personalized pricing decisions to the firms (which was without loss in the baseline model), the firms would induce all consumers who pay the posted price to showroom instead. Moreover, they would forego selling to some low-value consumers altogether. Attaining the optimal revenue would require the platform to adopt tools like proportional sales fees and most-favored-nation clauses.

7 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 combination of manual and automated bidding mechanisms to select valuable advertisements to display to each user and to set prices for these ads. 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 and supports clicks and other forms of user engagement with the platform.

We have proposed an integrated model that considers how auction mechanisms and data availability jointly determine match formation and surplus extraction both on and off large digital platforms. The auction mechanisms employed by 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, advertisers raise prices to gain a competitive edge on the platform.

The cross-channel distortions become more pronounced the more tools the platform has at its disposal, relative to the traditional (generalized) second price auctions. Indeed, 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. These results suggest the need for further analysis of how algorithmic bidding impacts competition and welfare in all markets, particularly off the large digital platforms.

A Appendix

The proof of Proposition 1 is given in the text.

Proof of Theorem 1 (1.) This part follows directly from Proposition 1.

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

$$\max_p \left\{ \frac{1-\lambda}{J} p(1-F(p)) + \lambda\Omega(p;p') \right\}, \tag{14}$$

where

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

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 $(\min(v' - \max(v-p,0), p'))_+$.

With some casework, we can show that the expected on-platform profit satisfies

$$\Omega(p;p') = \int_{\underline{v}}^{\bar{v}} \int_{\underline{v}}^{v} \min(v - v', p)\ dF^{J-1}(v')\ dF(v), \text{ for all } p'. \tag{15}$$

Expression (15) shows that each firm's profit in the second-price auction cannot exceed the difference between its own value and the next highest value. Furthermore, this difference is capped by the firm's own posted price p due to the showrooming constraint. Additionally, this expression is independent of p', so we suppress the dependence on p' in the notation and write $\Omega(p)$ instead.

To characterize the symmetric equilibria, we compute the derivative of Ω with respect to p. Straightforward algebra yields the following expression:

$$\Omega'(p) = \int_p^{\bar{v}} F^{J-1}(v-p)\ dF(v). \tag{16}$$

Finally, we can write out the first-order condition for profit maximization using (16):

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

(3.)-(4.) These results follow from setting $\bar{p}_j = p_B$ in the buyer's outside option (1) and recalling from the proof of Proposition 1 that firm j bids $b_j(v,\bar{p}) = (v_j - \underline{u}(v,\bar{p}))_+$. □

Proof of Proposition 2 (1.) The equilibrium price maximizes the profit function

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

where $\Omega(p)$ is given in (15). Equivalently, p_B maximizes the rescaled profit function

$$p(1-F(p))+\frac{\lambda J}{1-\lambda}\int_{\underline{v}}^{\bar{v}}\int_{\underline{v}}^{v}\min(v-v',p)\ dF^{J-1}(v')\ dF(v).$$

Because the second term is strictly increasing in p and λ, and it is multiplicatively separable, this function is supermodular in (p,λ); hence, by Topkis's theorem (Topkis, 1978), the profit-maximizing posted price is nondecreasing in λ.

(2.) The expected consumer surplus of an off-platform consumer and the expected welfare per consumer off-platform are given by

$$CS_{\text{off}}(p)=\int_p^{\bar{v}}(v-p)\ dF(v) \quad \text{and} \quad W_{\text{off}}(p)=\int_p^{\bar{v}} v\ dF(v),$$

respectively. Both quantities are strictly decreasing in p, and hence also in λ. □

Proof of Theorem 2 (1.) This follows from the definition of the best-value pricing and efficient steering policies when all firms post an identical price $\bar{p}_j \equiv p_V$.

(2.) To characterize the symmetric equilibrium prices, recall that, under the best-value pricing policy, no firm can poach a on-platform consumer for which it is not the high-value firm. Suppose first the firm posts a price $p < p'$. In this case, the firm collects $\min(v,p)$ on all such consumers, regardless of p'. The firm's profit is then given by

$$\Pi(p,p')=\frac{1-\lambda}{J}p(1-F(p))+\lambda\left(\int_{\underline{v}}^{p} vF^{J-1}(v)\ dF(v)+\int_p^{\bar{v}} pF^{J-1}(v)\ dF(v)\right). \qquad (17)$$

The derivative with respect to p is given by

$$\begin{aligned}
&\frac{1-\lambda}{J}(1-F(p)-pf(p))+\lambda\left(pF^{J-1}(p)f(p)-pF^{J-1}(p)f(p)+\int_p^{\bar{v}}F^{J-1}(v)\;dF(v)\right)\\
=&\frac{1-\lambda}{J}(1-F(p)-pf(p))+\lambda\left(\int_p^{\bar{v}}F^{J-1}(v)\;dF(v)\right). \qquad (18)
\end{aligned}$$

Now, suppose the firm posts a price $p > p'$. The firm's profit function is given by

$$\Pi(p,p') = \frac{1-\lambda}{J}p(1-F(p))+\lambda\begin{pmatrix} \int_{\underline{v}}^{p'} vF^{J-1}(v)\;dF(v)+\int_{p'}^{p}\int_{\underline{v}}^{p'} v\;dF^{J-1}(v')\;dF(v) \\ +\int_{p'}^{p}\int_{p'}^{v}(v-(v'-p'))\;dF^{J-1}(v')\;dF(v) \\ +\int_{p}^{\bar{v}}\int_{\underline{v}}^{p'+(v-p)} p\;dF^{J-1}(v')\;dF(v) \\ +\int_{p}^{\bar{v}}\int_{p'+(v-p)}^{v}(v-(v'-p'))\;dF^{J-1}(v')\;dF(v)\end{pmatrix}. \qquad (19)$$

With some algebra, one can show that the derivative of this expression with respect to p is

$$\frac{1-\lambda}{J}(1-F(p)-pf(p))+\lambda\left(\int_p^{\bar{v}}F^{J-1}(p'+v-p)\;dF(v)\right). \qquad (20)$$

Comparing (18) and (20), the derivative matches from the left and right at $p = p'$, and so the best-response function is continuously differentiable at p' with derivative

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

This expression is strictly positive at $p = \underline{v}$ and strictly negative at $p = \bar{v}$. Therefore, a necessary condition for a symmetric equilibrium is given by the first-order condition (9) that sets this derivative to zero. Furthermore, if a single price satisfies (9), then the symmetric equilibrium is unique. □

Proof of Theorem 3 The vertically integrated platform, which controls all the firms' prices, sets identical prices $\bar{p}_j = p$ to maximize

$$\Pi_C(p) \triangleq \left\{(1-\lambda)p(1-F(p))+\lambda\int_{\underline{v}}^{\bar{v}}\min(v,p)\;dF^J(v)\right\}.$$

Now compare this problem to a firm's best reply to a common price p' posted by its competitors: the firm's profit $\Pi(p,p')$ in (17) coincides with $\Pi_C(p)/J$ on $p \in [\underline{v}, p']$; and the firm's profit $\Pi(p,p')$ in (19) satisfies $\Pi(p,p') < \Pi_C(p)/J$ on $[p', \bar{v}]$. Now let $p^* \in \arg\max_p \Pi_C(p)$

denote a solution to the vertically integrated platform's problem. By construction, we have that $p^* \in \arg\max_p \Pi(p, p^*)$. Therefore, any p^* is supported in a symmetric equilibrium of the best-value pricing managed campaign with sufficiently low transfers.

Consider the largest such p^*; by the previous paragraph, there exists an equilibrium with posted price p^*. To show that the largest p^* is also the highest symmetric equilibrium price of the managed campaign, it suffices to argue that there cannot be some $p' > p^*$ which is an equilibrium. Suppose, for sake of contradiction, that such an equilibrium existed at p'. By our observation in the previous paragraph, $\Pi(p, p') = \Pi_C(p)/J$ for $p \leq p'$; since $p^* < p'$, this implies that $\Pi(p^*, p') = \Pi_C(p^*)/J > \Pi_C(p')/J = \Pi(p', p')$, where the inequality follows since p^* is the largest maximizer of Π_C. However, this is a contradiction of equilibrium at p', since a firm benefits by deviating down to p^*. Therefore, there cannot be equilibria with posted prices above the largest p^*, so the equilibrium with the highest price must maximize Π_C and hence must attain the integrated producer surplus level.

To show the optimality of best-value pricing and efficient steering, consider the platform revenue, which equals the joint surplus generated by this mechanism, net of the firms' outside option value defined in (8). This level of the outside option is a lower bound on the profit of a firm that refuses to participate in any mechanism. Therefore, the managed campaign we are considering maximizes the joint surplus of the platform and firms, and it concedes the smallest possible surplus to the firms. It follows that the best-value pricing campaign maximizes the platform's revenue. □

Proof of Theorem 4 We can nest the optimal pricing problem across the three problems (monopoly, auction, campaign) with a parameter γ. Consider the choice of posted price in each of the three models. Define the auxiliary profit function

$$\Pi(p, \gamma) \triangleq \frac{1-\lambda}{J}\left(1 - F(p)\right) p + \lambda \left(\int_{\underline{v}}^{\bar{v}} \int_{\underline{v}}^{v} \min(v - \gamma v', p) dF^{J-1}(v') dF(v) \right).$$

In the data-augmented auctions, each firm's profit function is given by $\Pi(p, 1)$. The profit function of the vertically integrated firm (which yields the equilibrium price in the best-value pricing managed campaign by Theorem 3) is given by $\Pi(p, 0)$. It is straightforward to verify that Π is submodular in (p, γ):

$$\frac{\partial^2 \Pi(p, \gamma)}{\partial \gamma \partial p} = -\int_{p}^{\bar{v}} (v - p)(J-1)F^{J-2}(v - p) f(v - p) f(v) \; dv < 0.$$

Thus, by Topkis's theorem, the largest maximizer of $\max_p \Pi(p, \gamma)$ is nonincreasing in γ. Since $p_V = \max\{\arg\max_p \Pi(p, 0)\}$ and $p_B = \max\{\arg\max_p \Pi(p, 1)\}$, it follows that $p_V \geq p_B$. Note that this also implies $p_V \geq p_B \geq p_M$ by Proposition 2. Finally, note that the matching of consumers to firms is identical across the two mechanism. Thus, the comparison of total surplus and consumer surplus is entirely driven by the posted price. Because both surplus levels are decreasing in p, the welfare comparative statics follow. $\square$

Proof of Proposition 3 Note that the joint profit outcome of the platform and firms in both bidding and the sophisticated managed campaign takes the form:

$$\Pi_J(p) = (1-\lambda)p(1-F(p)) + \lambda \int_{\underline{v}}^{\bar{v}} \min(v, p) dF^J(v),$$

and the bidding profit outcome is $\Pi_J(p_B)$ while the sophisticated managed campaign outcome is $\Pi_J(p_V)$. Fixing the bargaining power of the platform, the total transfer charged to all firms is $\Pi_J(p)/J - \Pi_O$, and so we can compute η as

$$\frac{p_V - p_B}{(1-\lambda)/J\left[p_V(1-F(p_V)) - p_B(1-F(p_B))\right] + \lambda/J \int_{p_B}^{\bar{v}} (\min(v, p_V) - p_B) F^{J-1}(v) dF(v)},$$

$$= \frac{J}{(1-\lambda)\left[1 - \frac{p_V F(p_V) - p_B F(p_B)}{p_V - p_B}\right] + \lambda \int_{p_B}^{p_V} \frac{v - p_B}{p_V - p_B} dF^J(v) + \frac{\lambda}{J} \int_{p_V}^{\bar{v}} dF^J(v)}.$$

We claim that the denominator of the above expression is less than 1. To see this, note the bracketed term is less than 1 because $p_V > p_B$, so the first term is at most $1 - \lambda$. The second term is integrated for $v \in [p_B, p_V]$, so $0 \leq v - p_B \leq p_V - p_B$. Hence the sum of the last two terms is bounded above by the integral $\lambda \int_{p_B}^{\bar{v}} dF^J(V) = \lambda(1 - F^J(p_B)) < \lambda$. Thus, the denominator is at most 1, with strict inequality for $\lambda > 0$, and so $\eta > J$. $\square$

Proof of Theorem 5 Suppose, for sake of contradiction, that in the optimal independent managed campaign, $p_I > p_V$, and the platform chose some pricing policy $p^*(v)$. Consider the best-response problem of a firm in the independent managed campaign. The best-response profit function for all $p < p_I$ is

$$\Pi_I(p) = \frac{1-\lambda}{J} p(1-F(p)) + \lambda \begin{pmatrix} \int_{\underline{v}}^{\bar{v}} \min(p^*(v), p) F^{J-1}(v)\ dF(v) \\ + \int_p p\left[\int_{v}^{\bar{v}} \mathbb{1}[v_j - p^*(v) < v - p] dF^{J-1}(v_j)\right]\ dF(v) \end{pmatrix}.$$

We can split this into two components. Denote the first component

$$\Pi_A(p) \triangleq \frac{1-\lambda}{J}p(1-F(p)) + \lambda \int_{\underline{v}}^{\bar{v}} \min(p^*(v), p) F^{J-1}(v)\ dF(v),$$

which is the profit from off-platform sales and on-platform sales in the segment that of on-platform consumers that prefer firm i, and the second component

$$\Pi_B(p) \triangleq \lambda \int_p^{\bar{v}} p \left[\int_v^{\bar{v}} \mathbb{1}[v_j - p^*(v) < v - p] dF^{J-1}(v_j) \right]\ dF(v),$$

which is the segment of consumers who prefer some other $j \neq i$ but are poached by i. Note that by construction, $\Pi_I(p) = \Pi_A(p) + \Pi_B(p)$. By the contradiction supposition, $p_I > p_V$ was the outcome of the optimal independent managed campaign. Note that $\Pi_B(p_I) = 0$, since no consumers can be poached when i sets the same posted price as all other firms. Since $\Pi_B \geq 0$ for all $p < p_I$, it follows that $\Pi_B(p_V) - \Pi_B(p_I) \geq 0$.

We now claim that $\Pi_A(p_V) - \Pi_A(p_I) > 0$. Consider Π_A'. Let ϕ^* denote the preimage of p^*: $\phi^*(p) = \{v \in V | p^*(v) \geq p,\ i = \arg\max v_i\}$. Then

$$\Pi_A'(p) = \frac{1-\lambda}{J}(1 - F(p) - pf(p)) + \lambda\mu(\phi^*(p)),$$

where μ is the probability measure over the type space. Note that since it is without loss for the platform pricing policy to never set a price larger than the consumer's value, $\phi^*(p) \subseteq \{v \in V | v \geq p, i = \arg\max v_i\} =: \bar{v}(p)$. Thus, $\mu(\phi^*(p)) \leq \mu(\bar{v}(p)) = \int_p^{\bar{v}} F^{J-1}(v)dF(v)$. Substituting this in, we get

$$\begin{aligned}
\Pi_A'(p) &= \frac{1-\lambda}{J}(1 - F(p) - pf(p)) + \lambda\mu(\phi^*(p)) \\
&\leq \frac{1-\lambda}{J}(1 - F(p) - pf(p)) + \lambda \int_p^{\bar{v}} F^{J-1}(v)dF(v) = \Pi_V'(p),
\end{aligned}$$

where Π_V was the joint vertical integration profit from the sophisticated managed campaign. Therefore, we have $\Pi_A'(p) \leq \Pi_V'(p)$, so

$$\Pi_A(p_V) - \Pi_A(p_I) = \int_{p_V}^{p_I} -\Pi_A'(p)dp \geq \int_{p_V}^{p_I} -\Pi_V'(p)dp = \Pi_V(p_V) - \Pi_V(p_I) > 0,$$

where the last inequality follows because p_V by definition is the largest maximizer of Π_V. But this implies that $\Pi_A(p_V) - \Pi_A(p_I) > 0$ and $\Pi_B(p_V) - \Pi_B(p_I) \geq 0$, and hence $\Pi_I(p_V) - \Pi_I(p_I) >$

0, which contradicts our supposition that the individual best response was to set posted price p_I. Thus, it cannot be that $p_I > p_V$.

To show that we cannot have equality, suppose for sake of contradiction that $p_I = p_V$, and consider $\Pi_I'(p_V)$. By the same argument above, we have that $\Pi_A'(p_V) \leq \Pi_V'(p_V)$. If $\Pi_A'(p_V) > \Pi_V'(p_V)$ strictly, then since $\Pi_B(p_V) = 0$ and $\Pi_B(p_V) \geq 0$, there exists an ϵ such that $\Pi_I'(p_V - \epsilon) > \Pi_I'(p_V)$, a contradiction. If $\Pi_A'(p_V) = \Pi_V'(p_V)$, then it must be that $\phi^*(p) = \bar{v}(p)$, so every consumer with value at least p sees a price at least p. But this implies that poaching can happen; more precisely, the left-derivative $(\Pi_B)'_-(p_v) = -\int_p^{\bar{v}} p dF^{J-1}(v)dF(v)$. Since the left derivative of Π_B is strictly negative, and $\Pi_A'(p_V) = \Pi_V'(p_V)$, for small enough ϵ then we must have $\Pi_I(p_V - \epsilon) > \Pi_I(p_V)$, a contradiction again. Hence we cannot have $p_I = p_V$. Thus $p_I < p_V$. Since the welfare and off-platform surplus are decreasing in posted price for $p \geq p_M$, the welfare comparative statics follow. □

Proof of Proposition 4 First, consider the problem of a vertically integrated platform facing the cohort-privacy constraint:

$$\Pi_P(p) := \max_{p' \leq p} \frac{1-\lambda}{J} p(1 - F(p)) + \frac{\lambda}{J} p'(1 - F^J(p')).$$

Let p_P denote the largest maximizer of Π_P. We can write the Lagrangian:

$$L(p, \mu) := \frac{1-\lambda}{J} p(1 - F(p)) + \frac{\lambda}{J} p'(1 - F^J(p')) + \mu(p' - p),$$

where μ is the multiplier associated with the showrooming constraint. Because F^J satisfies the monotone-likelihood ratio with respect to F, the unconstrained maximum must have $p' > p$; hence, the showrooming constraint binds, $p' = p$, and the optimal p_P satisfies:

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

Note that $\Pi_P(p_P)$ can be rewritten as

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

The necessary first-order condition for optimality thus requires p_P to be equal to the inverse hazard rate of the distribution $(1-\lambda)F + \lambda F^J$; since $(1-\lambda)F + \lambda F^J$ satisfies the monotone likelihood ratio property with respect to F, it follows that $p_P \geq p_M$. Similarly, since F^J satisfies the monotone likelihood ratio property with respect to $(1-\lambda)F + \lambda F^J$, it follows that the maximizer $p_J := \arg\max_p p(1 - F^J(p))$ is larger than p_P.

We now show that the platform can implement an equilibrium where all firms post p_P. Consider the platform pricing policy, which sets the price for cohort i as the minimum posted price of firms $j \neq i$. The best-response profit of firm i when all other firms set price p_P is

$$\begin{cases} \frac{1-\lambda}{J} p(1 - F(p)) + \frac{\lambda}{J} p(1 - F^J(p)) & p < p_P, \\ \frac{1-\lambda}{J} p(1 - F(p)) + \frac{\lambda}{J} p_P(1 - F^J(p_P)) & p \geq p_P. \end{cases}$$

Suppose, towards a contradiction, that p_P was not a maximizer, and some other p^* was the maximizer. If $p^* < p$, that implies that $\Pi_P(p^*) > \Pi_P(p_P)$, a contradiction of the definition of p_P. If $p^* > p$, then there exists $p' = p_P$, $p = p^*$ such that

$$\frac{1-\lambda}{J} p(1 - F(p)) + \frac{\lambda}{J} p'(1 - F^J(p')) > \Pi_P(p_P),$$

again a contradiction of the optimality of p_P. Hence, the best-response of each firm is also to set price p_P. Since this attains the vertical integration profit, it is optimal for the platform to set such a pricing policy.

It remains to show that $p_V \geq p_P$. Note that p_P maximizes

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

and p_V maximizes

$$\Pi_V(p) = \frac{1-\lambda}{J} p(1 - F(p)) + \frac{\lambda}{J} \left(\int_{\underline{v}}^{p} v dF^J(v) + \int_{p}^{\bar{v}} p dF^J(v) \right).$$

To complete the argument, define the auxiliary profit function again:

$$\Pi(p, \gamma) = \frac{1-\lambda}{J} p(1 - F(p)) + \frac{\lambda}{J} \left(\gamma \int_{\underline{v}}^{p} v dF^J(v) + \int_{p}^{\bar{v}} p dF^J(v) \right),$$

and note that $\Pi(\cdot, 0) = \Pi_P$ and $\Pi(\cdot, 1) = \Pi_V$. Finally, note that

$$\frac{\partial^2 \Pi(p, \gamma)}{\partial \gamma \partial p} = p \, dF^J(p) > 0.$$

By Topkis's theorem, the maximizer of $\Pi(\cdot, \gamma)$ is nondecreasing in γ; hence $p_P \leq p_V$. $\square$

Proof of Proposition 5 We consider the benchmark without a platform (denoted by O for oligopoly), the data-augmented bidding B, and the best-value pricing campaign V. To compare these three cases, consider an auxiliary game with payoffs

$$\begin{aligned}\Pi_{\text{off}}(p,\bar{p};\lambda,\gamma) =& \frac{1}{M}p\int_{p}^{\bar{v}} F^{N-1}(v-(p-\bar{p}))f(v)\ dv \\ &+ \frac{\lambda}{1-\lambda}\int_{\underline{v}}^{\bar{v}}\int_{\underline{v}}^{v} \min(v-\gamma v', p)dF^{J-1}(v')dF(v). \end{aligned} \tag{21}$$

The game with $\lambda = 0$ describes the case without a platform; the game with $\gamma = 1$ describes the data-augmented auctions; and a similar argument as in the proof of Theorem 3 establishes that every symmetric equilibrium in the game with $\gamma = 0$ is also an equilibrium under best-value pricing. We can then define the following prices:

$$\begin{aligned} p_O^K &\in \arg\max_p \Pi_{\text{off}}(p, p_O^K; 0, 1), \\ p_B^K &\in \arg\max_p \Pi_{\text{off}}(p, p_B^K; \lambda, 1), \\ p_V^K &\in \arg\max_p \Pi_{\text{off}}(p, p_V^K; \lambda, 0). \end{aligned}$$

To compare p_O^K and p_B^K, consider the best reply functions $p^*(\bar{p})$ in the game with $\lambda = 0$ and in the game with $\lambda > 0$ and $\gamma = 1$. The payoff function Π_{off} in (21) has increasing differences in (p, λ) for any $\bar{p}$ when $\gamma = 1$. Indeed, we have

$$\frac{\partial^2 \Pi_{\text{off}}(p,\bar{p};\lambda,1)}{\partial p \partial \lambda} = \frac{1}{(1-\lambda)^2}\int_p^{\bar{v}} F^{J-1}(v-p)dF(v) > 0.$$

By Topkis's theorem, a higher λ increases the best reply p^* to any price $\bar{p}$. Furthermore, the best response function satisfies $p^*(\bar{v}) < \bar{v}$ in all three cases. Therefore, in the highest-price symmetric equilibrium, the best-reply function crosses the line $p^* = \bar{p}$ from above, and an upward shift in the best replies raises the highest symmetric equilibrium price, i.e., $p_B^K > p_O^K$.

Similarly, the payoff function (21) has decreasing differences in (p, γ) for all $\bar{p}$ and $\lambda > 0$:

$$\frac{\partial^2 \Pi_{\text{off}}(p,\bar{p};\lambda,\gamma)}{\partial p \partial \gamma} = -\frac{\lambda}{1-\lambda}\int_p^{\bar{v}} (v-p)(J-1)F^{J-2}(v-p)f(v-p)f(v)\ dv < 0.$$

By Topkis's theorem, the best replies in the game with $\gamma = 0$ (i.e., the managed campaign) are pointwise larger than in the game with $\gamma = 1$ (i.e., the data-augmented auctions). Furthermore, the best replies when $\gamma = 1$ satisfy $p(\bar{v}) < \bar{v}$, and hence the largest symmetric equilibrium price increases as best replies shift up, i.e., $p_V^K > p_B^K$. Finally, because every equilibrium in the auxiliary game is also an equilibrium under the best-value pricing campaign, the largest symmetric equilibrium in the original game is strictly larger than p_B^K. $\square$

Proof of Proposition 6 Under best-value pricing, each firm j sells to the $1/J$ fraction of consumers who like its product best, since by construction $v_j - p(v, a, \bar{p}) \geq v_k - \bar{p}_k$. Thus, the pricing decisions of the firms are the same as in Theorem 2. The platform could then charge transaction fees instead of an upfront budget with the same outcome, provided the transaction fee satisfies the participation constraint of the firms: equivalently $\lambda t_j / J \leq T_j$, where T_j is what the platform would have charged in the original managed campaign. $\square$

References

Admati, Anat and Paul Pfleiderer (1990): "Direct and Indirect Sale of Information," *Econometrica*, 58, 901–928.

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.

Ali, S. Nageeb, Gregory Lewis, and Shoshana Vasserman (2023): "Voluntary Disclosure and Personalized Pricing," *Review of Economic Studies*, 90, 538–571.

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 (2019): "Markets for Information: An Introduction," *Annual Review of Economics*, 11, 85–107.

——— (2023): "Data, Competition, and Digital Platforms," Tech. rep., Yale University.

Bergemann, Dirk, Alessandro Bonatti, and Tan Gan (2022): "The Economics of Social Data," *RAND Journal of Economics*, 53, 263–296.

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, Negin Golrezaei, Patrick Jaillet, Jason Cheuk Nam Liang, and Vahab Mirrokni (2022a): "Individual Welfare Guarantees in the Autobidding World with Machine-learned Advice," *arXiv preprint*.

Deng, Yuan, Jieming Mao, Vahab Mirrokni, Hanrui Zhang, and Song Zuo (2022b): "Efficiency of the First-Price Auction in the Autobidding World," `https://arxiv.org/pdf/2208.10650.pdf`.

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.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz (2007): "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97, 242–259.

Elliott, Matthew, Andrea Galeotti, Andrew Koh, and Wenhao Li (2024): "Market Segmentation through Information," Tech. rep., Cambridge University.

Fuchs, William and Andrzej Skrzypacz (2015): "Government Interventions in a dynamic Market with Adverse Selection," *Journal of Economic Theory*, 158, 371–406.

Ganuza, Juan-José and Gerard Llobet (2018): "Personalized Prices in the Digital Economy," in *Economic Analysis of the Digital Revolution*, FUNCAS Social and Economic Studies.

Goke, Shumpei, Gabriel Y. Weintraub, Ralph Mastromonaco, and Sam Seljan (2022): "Bidders' Responses to Auction Format Change in Internet Display Advertising Auctions," in *Proceedings of the 23rd ACM Conference on Economics and Computation*.

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:2002.07331*.

Kirpalani, R. and T. Philippon (2021): "Data Sharing and Market Power with Two-Sided Platforms," Tech. Rep. 28023, NBER.

Konstantinovic, D. (2023): "Ad duopoly looks over its shoulder as Amazon ad revenues jump 26%," `https://www.insiderintelligence.com/content/ad-duopoly-looks-over-its-shoulder-amazon-ad-revenues-jump-26`.

Lebow, S. (2023): "Digital Ad Spend Worldwide to Pass $600 Billion This Year," `https://www.insiderintelligence.com/content/digital-ad-spend-worldwide-pass-600-billion-this-year`.

Li, Bin and Yahui Lei (2023): "Mechanism Design for Ad Auctions with Display Prices," *arXiv:2303.13192*.

Liaw, Christopher, Aranyak Mehta, and Andres Perlroth (2022): "Efficiency of non-truthful auctions under auto-bidding," *arXiv:2207.03630*.

Mehta, Aranyak (2022): "Auction Design in an Auto-bidding Setting: Randomization Improves Efficiency: Beyond VCG," in *WWW '22: 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.

Motta, Massimo and Antonio Penta (2022): "Market effects of sponsored search auctions," Tech. Rep. DP17401, CEPR.

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, Michael and Michael Schwarz (2023): "Reserve Prices in Internet Advertising Auctions: A Field Experiment," *Journal of Political Economy*, forthcoming.

Ostrovsky, Michael and Andrzej Skrzypacz (2022): "Pure-Strategy Equilibrium in the Generalized First-Price Auction," Tech. rep., Stanford GSB.

Philippon, Thomas and Vasiliki Skreta (2012): "Optimal interventions in markets with adverse selection," *American Economic Review*, 102, 1–28.

Segal, Ilya (2003): "Coordination and discrimination in contracting with externalities: Divide and conquer?" *Journal of Economic Theory*, 113, 147–181.

Stapleton, C. (2021): "Amazon's vs. Alibaba's Business Models: What's the Difference?" `https://www.investopedia.com/articles/investing/061215/difference-between-amazon-and-alibabas-business-models.asp`.

Teh, Tat-How and Julian Wright (2022): "Intermediation and steering: Competition in prices and commissions," *American Economic Journal: Microeconomics*, 14, 281–321.

Tirole, Jean (2012): "Overcoming adverse selection: How public intervention can restore market functioning," *American Economic Review*, 102, 29–59.

Topkis, D.M. (1978): "Minimizing a Submodular Function on a Lattice," *Mathematics of Operations Research*, 27, 305–323.

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.

Yu, Chuan (2024): "The Welfare Effects of Sponsored-Product Advertising," Tech. rep., Stanford University.

B Supplementary Appendix (Not for Publication)

B.1 Comparative Statics

B.1.1 Data-Augmented Bidding

We present additional welfare comparative statics for the data-augmented bidding model. 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 the difference between their willingness-to-pay and price, which is

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

Note that only off-platform consumers with value above p receive a positive surplus, since consumers with low values do not buy. The expected consumer surplus of an on-platform consumer is:

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

Again, only consumers with sufficiently high values receive positive surplus, since low-value consumers are price discriminated against. 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.

Similarly, the expected welfare per consumer off-platform is given by

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

Note that we focus on off-platform welfare because under our model, the platform has full information and sales are always made, there is full welfare on-platform.

Proposition 7 (Comparative Statics)
Suppose $J > -1/(\ln F(\bar{v} - p_M))$. The following comparative statics hold:

1. *The equilibrium posted price with data-augmented bidding p_B is decreasing in J.*

2. *The expected consumer surplus of on-platform and off-platform consumers is increasing in J.*

3. *Off-platform welfare per consumer W_{off} is increasing in J.*

Proof. By rescaling the best-response profit condition solved by the bidding problem, the bidding price is a maximizer of

$$\Pi(p) = (1-\lambda)p(1-F(p)) + \lambda\left(\int_{\underline{v}}^{\bar{v}}\int_{\underline{v}}^{v} \min(v-v',p) J dF^{J-1}(v') dF(v)\right).$$

Note that the cross-partial derivative with respect to J, p is

$$\frac{\partial^2 \Pi}{\partial J \partial p} \;\; = \;\; \int_{p}^{\bar{v}} (1 + J \ln F(v-p)) F^{J-1}(v-p) dF(v).$$

Since by assumption $1 < -J \ln F(\bar{v} - p_M)$, $1 + J \ln F(v-p) \leq 1 + J \ln F(\bar{v} - p_M) < 0$ for $p \geq p_M$. Hence this derivative is negative, so by Topkis's theorem, it thus follows that p_B must be decreasing in J.

The second statement follows from the first and the fact that CS_{off} and CS_{on} are both decreasing in p. Since CS_{on} is increasing in J, it follows that both quantities are increasing in J. Finally, the total off-platform welfare per consumer is decreasing in p, and has no other λ or J dependence, so W_{off} is increasing in J. ☐

We 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 7, for any J, the prices are increasing in λ.

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 profit for $J = 3$ are nonmonotone. As mentioned in Proposition 7, 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

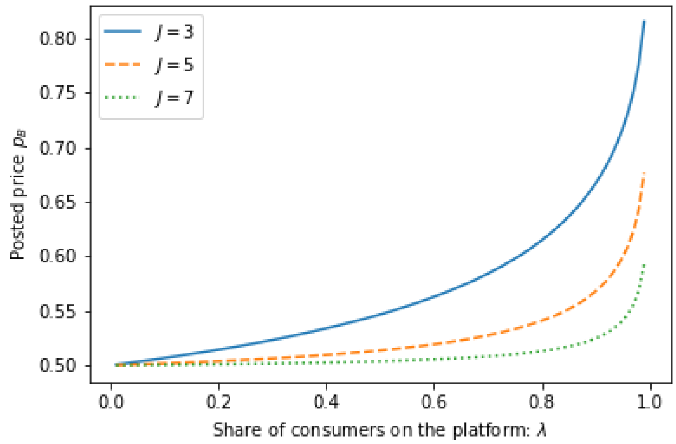

FIGURE 2: Posted prices as a function of λ. Results are plotted for $J = 3, 5, 7$.

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 pushing down the platform revenues.

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

B.1.2 Managed Campaign

Importantly, the optimal sophisticated managed campaign influences the effect of competition on consumer welfare. Recall that in the data-augmented bidding model, Proposition 7 showed that the expected welfare of both on-platform and off-platform consumers is increasing in the number of firms J. However, under the best-value managed campaign, this effect is reversed for the off-platform consumers, and may be reversed for on-platform consumers as well.

Proposition 8 (Managed Campaign Comparative Statics)
The best-value managed campaign price is increasing in J. The expected surplus of an off-platform consumer is decreasing in J.

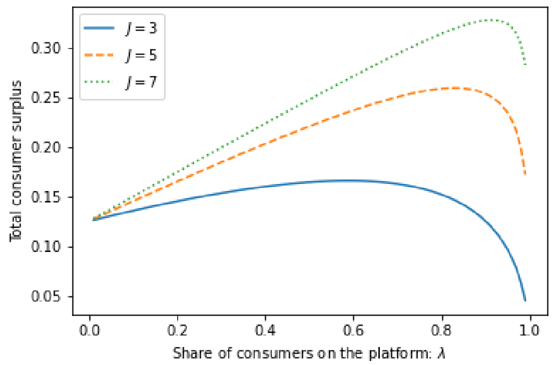

(A) Total Consumer Surplus

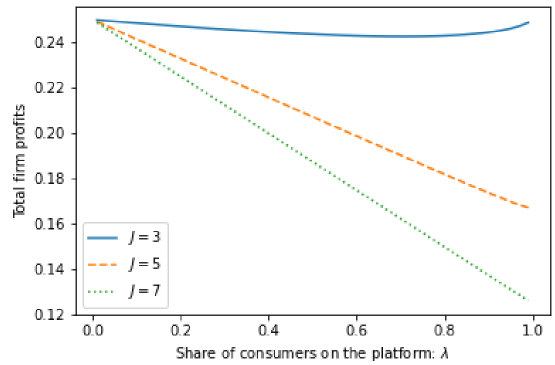

(B) Total Firm profit

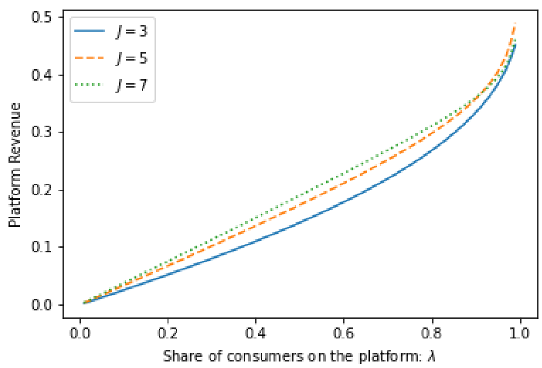

(C) Platform Revenue

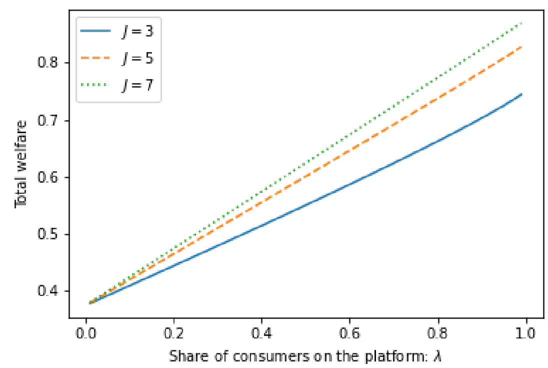

(D) Total Welfare

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 $J = 3, 5, 7$.

Proof. The cross-partial derivative of the vertical integration profit (multiplied by J) is

$$\frac{\partial^2 \Pi_C}{\partial p \partial J} = -\lambda F^J(p) \ln F(p) > 0.$$

By Topkis's theorem, the largest maximizer p_V is therefore increasing in J. The consumer surplus comparative static follows since the off-platform consumer surplus is weakly decreasing in p_V. □

Proposition 8 shows that the expected surplus of off-platform consumers decreases with the number of firms, since the off-platform prices *increase* with the number of firms on the platform.

On the other hand, for on-platform consumers, more firms implies that an on-platform consumer's favorite firm is likely to be more valuable (the first-order statistic is larger), which counteracts the higher off-platform prices and greater price discrimination. We present an

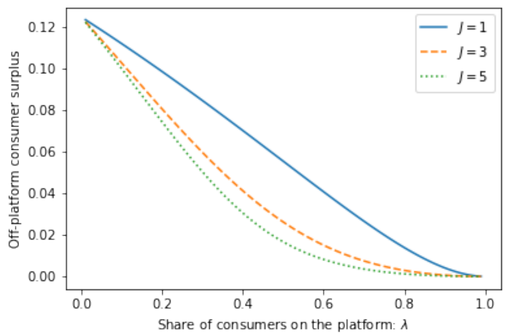

(A) Expected consumer surplus off-platform

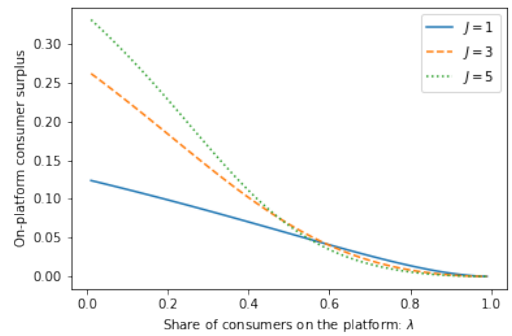

(B) Expected consumer surplus on-platform

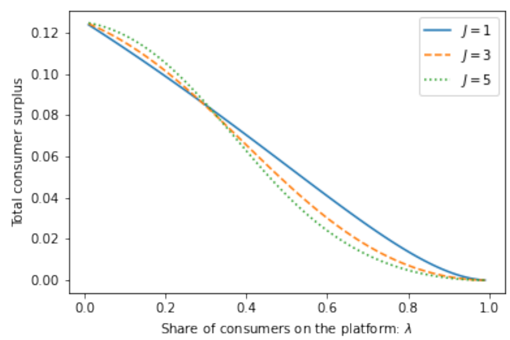

(C) Total consumer surplus

FIGURE 4: Managed campaign consumer surplus and competition. Distribution of consumer values F is uniform on $[0, 1]$, plotted for varying λ with $J = 1, 3, 5$.

example illustrating the tension between these two forces in Figure 4. In the example, we take the consumer values to be distributed uniformly on the unit interval, and plot the expected surplus of an off-platform consumer, of an on-platform consumer, and total consumer surplus varying the number of firms for $J = 1, 3, 5$. As indicated by Proposition 8, the off-platform consumer surplus is decreasing for larger J, regardless of platform size. For on-platform consumers, when the platform is relatively small (lower λ), the larger off-platform market dampens on-platform price discrimination, and the surplus of on-platform consumers increases with more firms, as the on-platform consumers gain from having more firms to choose between. However, when the platform is relatively large, the increase in price discrimination with more firms results in reduced surplus for on-platform consumers despite having more firms to choose from. As such, the total consumer surplus is increasing in J for small platforms (low λ) but decreasing in J for large platforms (high λ).

B.1.3 Independent Campaigns vs. Bidding

We now discuss the implications of independent managed campaigns relative to data-augmented bidding for both posted price and welfare. We compare the independent managed campaign induced by perfect price discrimination (pricing policy $p_i(v) = \min(v, p_i)$) to data-augmented bidding.

Proposition 9 (Price and Welfare Comparisons)
Suppose the monopoly profit function $p(1-F(p))$ is concave. If 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. If F^{J-1} is concave, all the inequalities are reversed.

Proof. If the monopoly profit function is concave, then p_I and p_B are the unique solutions to their respective first-order conditions. Recall that p_B satisfies

$$(1-\lambda)(1-F(p_B)-p_Bf(p_B)) + \lambda J \int_{p_B}^{\bar{v}} F^{J-1}(v-p_B)dF(v) = 0,$$

and p_I satisfies

$$(1-\lambda)(1-F(p_I)-p_If(p_I)) + \lambda J \int_{p_I}^{\bar{v}} (F^{J-1}(v) - p_I dF^{J-1}(v))\mathbb{1}[p_I \in P(v;p_I)]dF(v) = 0.$$

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

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

since the right-hand side is a first-order expansion of F^{J-1} around v. Additionally, $F^{J-1}(v-p) \geq 0$. Therefore, the left-hand sides of the first-order condition for p_I is lower than the first-order condition for p_B, and hence $p_B \geq p_I$.

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

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

since the right-hand side is a first-order expansion of F^{J-1} around v. As $F^{J-1}(v) - pdF^{J-1}(v)$ is positive, it follows that $P(v;p_I) = \{p_I\}$, since the maximizing function is always increasing; hence, p_I solves the first-order condition

$$(1-\lambda)(1-F(p_I)-p_If(p_I)) + \lambda J \int_{p_I}^{\bar{v}} (F^{J-1}(v) - p_I dF^{J-1}(v))dF(v) = 0.$$

Using the same argument as before, $p_B \leq p_I$. As total welfare and total consumer surplus decrease in p, the welfare comparative statics follow. $\square$

Proposition 9 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 10 (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 $p_B \geq p_I$. $\square$

In Figure 5, 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 5a shows the revenue when there are $J = 2$ firms, and Figure 5b plots the revenue for $J = 3$ firms. Figure 5a demonstrates a scenario where the independent managed campaign yields more revenue, and Figure 5b 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.

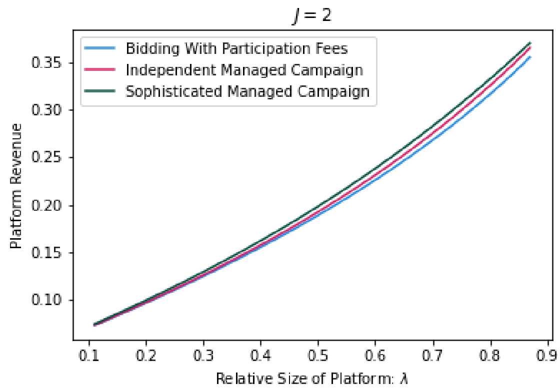

(A) Platform revenue as a function of λ, $J = 2$

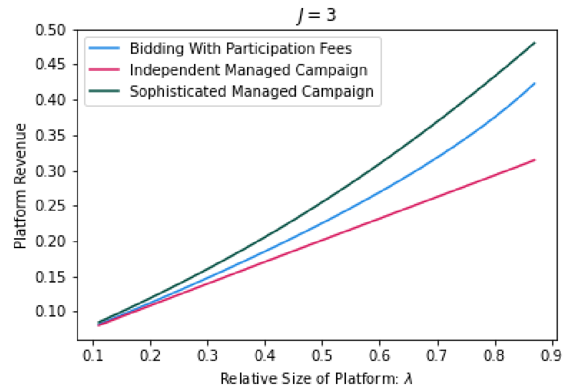

(B) Platform revenue as a function of λ, $J = 3$

FIGURE 5: 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.

Finally, the sophisticated managed campaign results in higher platform revenue than the independent managed campaign and bidding, as would be expected by Theorem 3.

B.1.4 Cohort Privacy

Figure 6 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 is $p_M = 0.5$. As would be expected from Proposition 4, 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 7 for a uniform distribution of consumer values.

B.2 Robustness, RET, and Participation Fees

B.2.1 Bidding Auction Format

In the main text, we chose a second-price auction for convenience. The main insight of Theorem 1 extend to the case of the first-price auction.

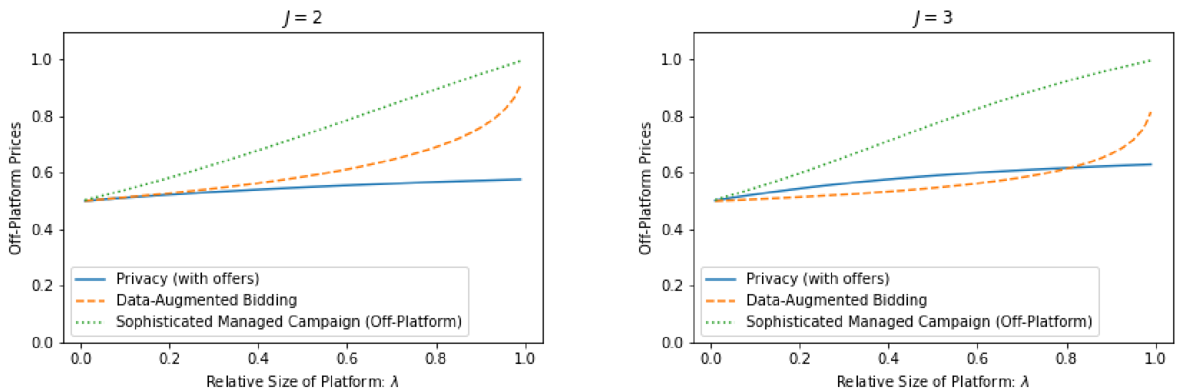

(A) Off-platform prices as a function of λ, $J = 2$ (B) Off-platform prices as a function of λ, $J = 3$

FIGURE 6: 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.

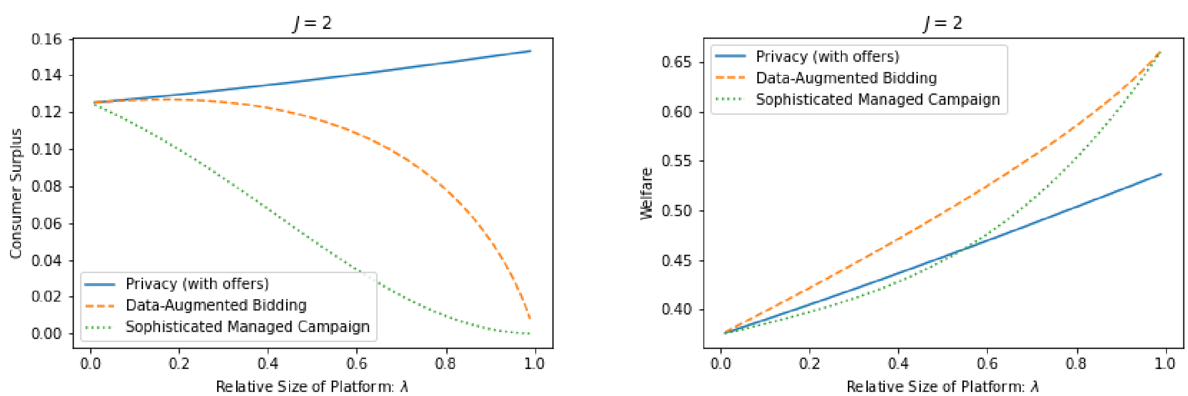

(A) Consumer surplus as a function of λ, $J = 2$ (B) Total welfare as a function of λ, $J = 2$

FIGURE 7: Welfare implications of privacy, with a uniform distribution of consumer values.

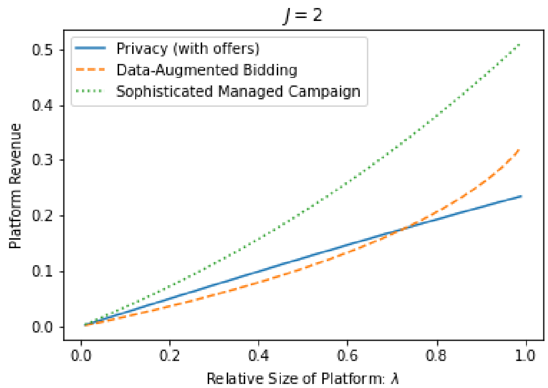

(A) Platform revenue, $J = 2$

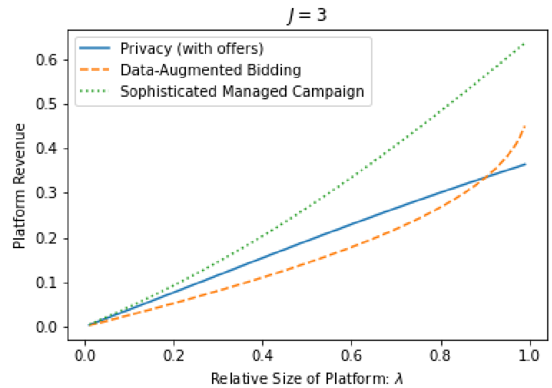

(B) Platform revenue, $J = 3$

FIGURE 8: Privacy implications on platform revenue, $J = 2, 3$.

Proposition 11

Suppose the platform ran a first-price auction instead of a second-price auction (breaking ties in favor of firms with higher value). In the unique symmetric equilibrium in undominated strategies, the firms post price p_B satisfying the same condition (2) *as in Theorem 1.*

Proof. To show that the first-price auction induces the same posted price outcome as in Theorem 1, we argue first that Proposition 1 still holds. Fix any vector of off-platform prices $\bar{p}$. Define $\underline{u}(v, \bar{p})$ as in (1), and once again note that this is the outside option of a consumer on the platform. Hence, each j must concede at least this much utility to the consumer, and can hence charge at most $p_j(v) = (v_j - \underline{u}(v, \bar{p}))_+$. As we argued, since $\underline{u}$ is common to all firms, the highest-valued firm is able to charge the most. In particular, suppose firm j' has the second-highest value, and j the highest value: $v_j > v_{j'}$. Then j always wins in a first-price auction by bidding arbitrarily close to, but above $(v_{j'} - \underline{u}(v, \bar{p}))_+$. Since the platform breaks ties in favor of the firm with the highest value for the consumer, the firm j bids $(v_{j'} - \underline{u}(v, \bar{p}))_+$ and wins.

However, we just showed that firm j bids exactly firm j''s value, and so pays exactly the same that firm j would pay in the second-price auction, implementing the same allocation as the second-price auction. Since the payments are the same for the firms, the rest of the analysis in the proof of Theorem 1 follows, so the equilibrium posted price outcome is the same as in (2). □

B.2.2 Bidding with Participation Fees

In the bidding model introduced in the main text, the platform received revenues only from the bids of the advertisers. We now ask whether tools from optimal auction design (namely, participation fees) may increase the revenue of the platform.[26] 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.
2. The firms choose whether to pay the participation fee and set their posted prices.
3. All firms observe participation decisions and posted prices. The platform runs a second-price auction for each on-platform consumer among the participating firms.

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 12 (Equilibrium with Participation Fees)
In equilibrium, all firms join and the off-platform posted prices are given by (2). *Firms bid the same as in the bidding equilibrium. Firm profits are held to their outside option* (8). *The fee charged by the platform holds firms to this outside option.*

Proof. By Proposition 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. Subgame perfection and Theorem 1 imply that the pricing condition for off-platform prices is given by (2), 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. Consider the profit firm i makes after refusing participation, with all other firms joining the platform. By Proposition 1, the firm can never get a sale from a consumer whose favorite firm is $i \neq j$; further, the highest bidder for a consumer whose favorite firm is i is exactly firm j such that $j = \arg\max_{j \neq i} v_j$. Note that v_j can undercut the posted price firm i sets by pricing down

[26]The 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.

to (possibly) zero; so firm i can only retain on-platform sales from the consumers such that $v_i - p \geq v_j - 0$. That is, the on-platform sales are

$$\int_p^{\bar{v}} pF^{J-1}(v-p)dF(v).$$

Combining, we find the exclusion profit is given by (8), and the result follows. □

Intuitively, the pricing and bidding behavior follow as in Theorem 1 due to subgame perfection. The fee charged is as large as possible to make firms indifferent between accepting and rejecting, and thus holds firms to their outside option profit (8). Qualitatively, the participation fee does not change the posted prices, but redistributes surplus from firms towards the platform.

B.2.3 Managed Campaign: Variable Fees

Alternatively to fixed fees, the platform could charge a commission. That is, instead of taking a fixed fee for each sale, the platform could instead take a fraction α of the revenue of each sale. However, we show that this does result in a difference in the pricing incentives relative to Theorem 2.

Proposition 13 (Commissions)
Suppose the platform charged a commission of $\alpha \in (0,1)$ of the revenue from all sales. Then under best-value pricing, the subgame equilibrium off-platform price set by the firms under commissions satisfies

$$0 = \frac{1-\lambda}{J}(1 - F(p_\alpha) - pf(p_\alpha)) + \lambda(1-\alpha)\left(\int_{p_\alpha} F^{J-1}(v)\ dF(v)\right). \tag{22}$$

Further, $p_\alpha < p_V$, where p_V is the equilibrium price from Theorem 2.

Proof. Note that with a commission α charged on all sales, the firm's best-response pricing profit condition becomes

$$\Pi_\alpha(p,p') = \frac{1-\lambda}{J}p(1-F(p)) + \lambda(1-\alpha)\int_{\underline{v}}^{\bar{v}} \min(v,p)F^{J-1}(v)\ dF(v).$$

for $p < p'$, since the firm only captures the remaining $(1-\alpha)$ of the on-platform sale value. Since the cross-derivative of Π_α in p, α is negative, by Topkis's theorem, it follows that the price solving this first-order condition must be lower than the best-value price p_V. □

Proposition 13 shows that by charging commissions, the platform does induce a distortion relative to p_V; however, this distortion is better for consumers, since the distortion lowers the off-platform prices relative to the managed campaign in Theorem 2. Intuitively, with commissions, the firm does not have as strong an incentive to raise its prices, since the firm's marginal benefit is diminished (multiplied by $1 - \alpha$).

However, this insight is central to showing that a capped commission structure (i.e., the platform charges a commission, but never charges more than a capped amount t per sale) can yield the Theorem 2 outcome; that is, the platform could charge commissions on low-value sales, but implement a cap to induce the firms' to price higher off the platform.

Proposition 14 (Commissions with Caps)
Suppose the platform charged commission of $\alpha \in (0, 1)$ of the revenue from all sales, but capped its maximum commission fee at t. Then if $t/\alpha \leq p_\alpha$, then the equilibrium posted price is the same p_V from (9) as in Theorem 2.

Proof. With commission α capped at some t_j, the firm's on-platform profit changes. If $p \geq t_j/\alpha$, the profit is

$$\int_{\underline{v}}^{t_j/\alpha} (1-\alpha) v F^{J-1}(v) dF(v) + \int_{t_j/\alpha}^{p} (v - t) F^{J-1}(v) dF(v) + \int_{p}^{\bar{v}} (p - t) F^{J-1}(v) dF(v).$$

If $p < t_j/\alpha$,

$$\int_{\underline{v}}^{p} (1-\alpha) v F^{J-1}(v) dF(v) + \int_{p}^{\bar{v}} (1-\alpha) p F^{J-1}(v) dF(v),$$

since the cap never binds in this case. In the first case, note that the derivative of the profit term is exactly

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

which is the same term as in best-value pricing. To argue that the firm never wants to set a price in the second case, it suffices to see that the profit function is strictly increasing in this range. Thus, if $t_j/\alpha < p_\alpha$ (where p_α is the commissions price we characterized in Proposition 13), the profit function of the firm must be strictly increasing on this range. Hence, the firm must set a price at least t_j/α, and so the firm's best response profit is maximized by p_V. □

Intuitively, the cap removes the distortion on the marginal benefit from pricing higher induced by commissions; thus, the implementation of a cap can actually harm consumers, as the cap induces higher firm pricing.