Open and Private Exchanges in Display Advertising

W. Jason Choi   Amin Sayedi*

August 2022

*W. Jason Choi (wjchoi@umd.edu) is Assistant Professor of Marketing at the Smith School of Business, University of Maryland. Amin Sayedi (aminsa@uw.edu) is Associate Professor of Marketing at the Foster School of Business, University of Washington. The authors would like to thank Elea Feit, Kinshuk Jerath, Miklos Sarvary, Jiwoong Shin, seminar participants at the Chinese University of Hong Kong, the NYC Media Seminar, University of Florida, University of Maryland, and University of Science and Technology of China for their valuable comments.

Open and Private Exchanges in Display Advertising

Abstract

We study the impact of the emergence of private exchanges (PX) on the display advertising market. Unlike open exchanges (OX), the original exchange types that are open to all publishers and advertisers, the newly emerged PX is only available to a smaller set of pre-screened advertisers and publishers through an invite-only process. The OX exposes advertisers to ad fraud and brand safety risks, whereas the PX ensures that advertisers purchase high-quality impressions from reputable publishers. While the assurance of higher quality increases advertisers' valuation for the PX impressions, we find that selling through both the OX and PX can hurt publishers by creating an information asymmetry among advertisers. In equilibrium, the publisher may sell through the OX, PX, or both, depending on the baseline fraud intensity and the advertisers' average valuations. Finally, our model sheds light on the OX's incentive to fight fraud. In the absence of the PX, the OX has low incentive to combat fraud because it earns commission from fraudulent transactions. However, the introduction of the PX may create competitive pressure such that the OX screens fake impressions; i.e., the PX may induce the market to self-regulate.

Keywords: display advertising, real-time bidding, first-price auction, private exchange, open exchange, advertising fraud

1 Introduction

Display ad spending in the US is projected to reach \$108 billion in 2021, accounting for 57% of total digital ad spending. Approximately one fourth of the display ad spending, around \$27 billion in 2021, is allocated to real-time bidding (RTB).[1] RTB was initially created as an efficient means to clear inventory that was left unsold through the traditional sales method, whereby brands and publishers connect one-to-one and negotiate the media sales contract. However, advances in programmatic ad technology, combined with the proliferation of impressions on the web, have drastically increased the demand for RTB, which offered scalable, individual-level ad targeting technology.

In RTB, advertisers submit their bids in real time to online marketplaces, known as *exchanges*, where publishers sell their inventory. Exchanges act as intermediary auction houses that connect publishers to advertisers. There are two types of exchanges in the RTB market: *open*

[1] https://forecasts-na1.emarketer.com/584b26021403070290f93a56/5851918a0626310a2c1869c4

exchanges and *private exchanges* (also known as private marketplaces). An open exchange, as the name suggests, is open to all publishers and advertisers. Examples of open exchanges include DoubleClick, Xandr, and OpenX.

While open exchanges (mainly Google's DoubleClick) dominated the RTB market since their inception, the opacity and complexity of the multi-tiered supply chain rendered open exchanges vulnerable to ad fraud. eMarketer projects that in 2023, advertisers will lose $100 billion of their ad spend to fraud (He, 2019). Common forms of ad fraud include domain spoofing, non-human traffic, and click spamming (Davies, 2018; Fou, 2020). For example, in domain spoofing, a fraudster presents itself as a reputable publisher and deceives advertisers into buying fake inventory. In a recent study, to assess the degree of ad fraud, Financial Times tried to buy impressions in open exchanges allegedly originating from FT.com, Financial Times' own website. The company found that over 300 fake accounts were selling, under the guise of FT.com, the equivalent of one month's supply of bona fide FT.com video inventory in a single day (Davies, 2017).[2]

In response to the growing fraud risks in open exchanges, publishers set up their own private exchanges, where they have more control over their inventory sales. A private exchange is an exclusive exchange where a publisher, or a small group of publishers, sells their inventory only to select advertisers through an invite-only process.[3] Ad spending in private marketplaces has grown rapidly in recent years, and in 2020, it surpassed that in open exchanges for the first time.

The advantages of private marketplaces over open exchanges are manifold. First, private marketplaces can mitigate ad fraud because only trusted publishers and advertisers have access to the exchange. Second, since advertisers are pre-screened in private marketplaces, publishers share more information about contexts (e.g., webpage content) and consumers (e.g., browsing history) in private marketplaces than in open exchanges (Vrountas, 2020). Third, advertisers

[2]For a comprehensive report on advertising fraud, see Cheq (2020).

[3]Note that this is different from programmatic direct advertising where advertisers and publishers connect one-to-one and negotiate terms of the advertising campaign akin to the traditional media-buying process. For more information, see Zawadzinski (2021).

trust publishers in private marketplaces more than in open exchanges; therefore, advertisers are less concerned about brand safety issues; e.g., having their ad shown next to objectionable content (Hsu and Lutz, 2020). Moreover, publishers can benefit from private marketplaces because milder fraud and higher-quality information allow them to sell inventory at higher prices.

Private exchanges do not come without any downsides. Advertisers can access a private exchange only if they are invited. The invitation process allows the publisher to control the operational costs of running the private exchange, and gives publishers more control over who sees what type of information about customers. However, since private exchanges are available only to a smaller set of invited advertisers, the average number of bids per impression (also known as bid density) is lower than in open exchanges. As such, the impressions may sell at lower prices than in open exchanges. Publishers have sought to address this problem by sending their request-for-bids to private and open exchanges simultaneously, in a process known as header bidding. Header bidding allows a publisher to send a request-for-bid to multiple (open and private) exchanges at the same time, and allocate the impression to the exchange with the highest clearing price. While header bidding mitigates the negative impact of softened competition on publishers' revenues, it cannot necessarily eliminate it.

This paper studies how the introduction of private exchanges affects advertisers', publishers' and open exchanges' revenues, as well as their strategies. We compare the benchmark where private exchanges do *not* exist to the situation where they co-exist with open exchanges and are accessible by a subset of advertisers. We address the following research questions.

1. How does the existence of a private exchange affect the strategies and the expected utilities of the advertisers?
2. How does the existence of a private exchange affect the expected utility of the publisher that offers the private exchange? How should the publisher set reserve prices in the private and open exchanges?
3. How does the existence of a private exchange affect the expected utility of an open exchange? How does it influence the open exchange's incentive to fight fraud?

To answer these questions, we use a game-theoretic model with two advertisers, a publisher, an open exchange, and a private exchange. In answering the first question, we show that the existence of a private exchange distorts the information structure of the game by giving an advantage to the *connected* advertiser, who has access to the private exchange, compared to the *unconnected* advertiser, who does not have access to the private exchange. The private exchange enables the connected advertiser to better identify legitimate impressions; therefore, conditioned on winning, the impression bought by an unconnected advertiser is more likely to be a fake impression when the private exchange exists than when it does not. This informational disadvantage lowers the unconnected advertiser's willingness-to-pay for impressions in the open exchange, which in turn softens bidding competition. Therefore, advertisers who have access to the private exchange benefit not only from pruning off fake impressions, but also from softened competition.

As for the publishers, selling through private exchanges mitigates ad fraud and allows publishers to set discriminatory reserve prices in open and private exchanges. However, the introduction of the private exchange can also hurt the publishers. First, the existence of the private exchange can *disperse competition*. When there are two exchanges, the bids in one exchange cannot be used as a clearing price in the other exchange. Therefore, the publisher's revenue may decrease in the presence of a private exchange as the advertisers are thinned out across multiple auctions. Interestingly, we show that if the publisher uses first-price auctions instead of second-price auctions in its exchanges, the competition dispersion effect is completely eliminated. Intuitively, when there are two exchanges with first-price auctions, advertisers in each exchange take into account competitors in other exchanges when submitting their bids.

The addition of the private exchange has a second negative effect on the publisher's revenue: the competition softening effect induced by the information asymmetry among advertisers. The existence of a private exchange informationally disadvantages the unconnected advertiser, thereby lowering its willingness-to-pay for impressions in the open exchange. We call this the devaluation effect. This in turn allows the connected advertiser to win impressions with lower bids. As a result, selling through a private exchange may reduce the publisher's revenue. Interestingly, the equilibrium market structure may not include a private exchange, even if publishers can adopt it to reduce fraud. If the baseline fraud intensity is mild and the advertisers' average willingness-to-pay for a legitimate impression high, then the devaluation effect outweighs the gains from mitigating fraud such that the publisher does not set up a private exchange. More generally, we characterize the conditions under which a publisher sells through only an open exchange, only a private exchange, or both exchanges at the same time.

Finally, we analyze how the addition of a private exchange impacts the open exchange's revenue and its incentive to fight ad fraud. We find that the existence of the private exchange lowers the open exchange's revenue because impression sales in the open exchange are lost to the private exchange. Open exchanges have been criticized for their inadequate anti-fraud efforts as they take a cut from those fraudulent transactions (Rowntree, 2019). We show that this is indeed the case in the absence of the private exchange. With the introduction of the private exchange, however, competitive pressure may incentivize the open exchange to fight fraud. While filtering fraudulent impressions reduces the transaction volume, it increases the advertisers' valuation for impressions in the open exchange.

Overall, our work sheds light on how the emergence of private exchanges in the RTB market affects advertisers and publishers. We highlight the information asymmetry induced by the introduction of a private exchange as an important economic force in this market. We provide managerially relevant insights for advertisers and publishers regarding bidding strategies and reserve prices. We also elucidate the nuanced implications for advertisers who have access to the private exchange and those who do not. For publishers, we characterize the optimal exchange configurations (i.e., sell through an open exchange only, a private exchange only, or open and private exchanges simultaneously) as well as the optimal reserve prices under different market conditions.

The rest of this paper is structured as follows. First, we discuss related papers to our work. In Section 2, we describe the model. In Section 3, we present the analysis and discuss the results for publishers and advertisers. In Section 4, we study the open exchange's incentive to fight fraud in response to the introduction of a private exchange. In Section 5, we test the robustness of the key insights by considering numerous extensions. In Section 6, we suggest avenues for future research and conclude. All proofs are relegated to the appendix.

Related Literature

Our work is related to the growing literature on online advertising auctions. Katona and Sarvary (2010) and Jerath et al. (2011) study advertisers' incentives in obtaining lower vs. higher positions in search advertising auctions. Sayedi et al. (2014) investigate advertisers' poaching behavior on trademarked keywords, and their budget allocations across traditional media and search advertising. Desai et al. (2014) analyze the competition between brand owners and their competitors on brand keywords. Lu et al. (2015) and Shin (2015) study budget constraints, and budget allocation across keywords. Zia and Rao (2019) look at the budget allocation problem across search engines. Wilbur and Zhu (2009) find the conditions under which it is in a search engine's interest to allow some click fraud. Cao and Ke (2019) and Jerath et al. (2018) study manufacturer and retailers' cooperation in search advertising and show how it affects intra- and inter-brand competition. Amaldoss et al. (2015) show how a search engine can increase its profits and also improve advertisers' welfare by providing first-page bid estimates. Berman and Katona (2013) study the impact of search engine optimization, and Amaldoss et al. (2016) analyze the effect of keyword management costs on advertisers' strategies. Katona and Zhu (2017) show how quality scores can incentivize advertisers to invest in their landing pages and to improve their conversion rates. Long et al. (2021) study the informational role of search advertising on the organic rankings of an online retail platform. Our work is different from these papers as we study display advertising auctions in real-time bidding. In the RTB market, the publisher can sell an impression in multiple auctions (open and private exchanges) in parallel, whereas in the search advertising market, impressions are only sold in single auctions that are owned and operated by search engines. As such, the competition between multiple exchanges, and the information asymmetry that emerges by the introduction of private exchanges do not exist in search advertising markets.

Our work contributes to the vast literature on display advertising. Empirical works in this area have assessed the effectiveness of display advertising in various contexts (e.g., Bruce et al., 2017; Hoban and Bucklin, 2015; Lambrecht and Tucker, 2013; Rafieian and Yoganarasimhan, 2021). Ada et al. (2021) exploit a change in information disclosure policy and find that context information disclosure to advertisers increases the publisher's average per-impression revenue. On the theoretical front, Sayedi et al. (2018) study advertisers' bidding strategies when publishers allow advertisers to bid for exclusive placement on the website. Selling through a private exchange is similar to selling exclusively as both mechanisms can thin the market by selling to a small subset of advertisers and create asymmetry in advertisers' valuations. However, a unique aspect of our model that does not exist in Sayedi et al. (2018) is the information asymmetry that selling through private exchanges creates. Since connected advertisers who have access to the private exchange can cherry-pick legitimate impressions, the expected value of an unconnected advertiser for an impression in the open exchange conditional on winning decreases when the private exchange is introduced. In contrast, in Sayedi et al. (2018), while the introduction of exclusivity can lower the probability of winning of the advertisers who cannot get exclusivity, it does *not* decrease their valuation per impression conditional on winning. This distinction leads to the devaluation effect that arises in our model but not in Sayedi et al. (2018).

Zhu and Wilbur (2011) and Hu et al. (2015) study the trade-offs involved in choosing between "cost-per-click" and "cost-per-action" contracts. Berman (2018) explores the effects of advertisers' attribution models on their bidding behavior and their profits. Despotakis et al. (2021b) and Gritckevich et al. (2021) look at how ad blockers affect the online advertising ecosystem, and Dukes et al. (2020) show how skippable ads affect publishers' and advertisers' strategies and profits. Choi et al. (2022) analyze consumers' privacy choices in a setting where their choices affect the advertisers' ability to track and target consumers along the purchase journey. Kuksov et al. (2017) study firms' incentives in hosting the display ads of their competitors on their websites. Choi and Sayedi (2019) study the optimal selling mechanism when a publisher does not know, but benefits from learning, the performance of advertisers' ads. In contrast to these papers, which do not study the roles of intermediaries (i.e., exchange platforms) in the market, we investigate the emergence of private exchanges in the RTB market and its impact on the advertisers', publishers' and exchanges' utilities and their strategies.

In the context of real-time bidding auctions, Johnson (2013) estimates the financial impact of privacy policies on publishers' revenue and advertisers' surplus. Rafieian (2020) characterizes the optimal mechanism when the publisher uses dynamic ad sequencing. Zeithammer (2019) shows that introducing a soft reserve price, a bid level below which a winning bidder pays his own bid instead of the second-highest bid, cannot increase publishers' revenue in RTB auctions when advertisers are symmetric; however, it can increase the revenue when advertisers are asymmetric. Sayedi (2018) analyzes the interaction between selling impressions through real-time bidding and selling through reservation contracts; it shows that, in order to optimize their revenue, publishers should use a combination of RTB and reservation contracts. Both in our paper and Sayedi (2018), the cherry-picking of impressions in one market — the RTB market in Sayedi (2018) and the private exchange in ours — negatively affects the prices in the competing market. However, a novel effect that exists in our model, and not in Sayedi (2018), is that since some advertisers have access to both exchanges, the lower prices in the open exchange may in turn lower the prices in the private exchange as well, and reduce the publisher's overall profit. As such, unlike Sayedi (2018), in our model the publisher sometimes benefits from selling in only one market. Choi and Mela (2019) study the problem of optimal reserve prices in the context of RTB, and, using a series of experiments, estimate the demand curve of advertisers as a function of the reserve price. Since the dataset in Choi and Mela (2019) is from 2016, the publishers primarily rely on open exchanges. The most closely related paper is Despotakis et al. (2021a), where the authors study a market with multiple exchanges. Despotakis et al. (2021a) examine how the transition from waterfalling to header bidding alters the competition between exchanges, and how this change motivates the exchanges to move from second- to first-price auctions. The exchanges in Despotakis et al. (2021a) are symmetric, and the authors do not look at the issue of ad fraud. In contrast, we model different types of exchanges, the asymmetries that arise from that, and how those relate to ad fraud. Despotakis et al. (2017) also study the strategic implications of information asymmetry among bidders, but in a dynamic setting with exogenous asymmetry. The observability of competitor's bids introduces signaling, which may motivate non-experts to bid above their valuation. In our paper, the information structure is endogenously determined by the publisher's exchange choices. Moreover, we show that in the absence of signaling, information asymmetry induced by the co-existence of two types of exchanges *lowers* the uninformed advertiser's bid as it increases the risk of winning fake impressions in the open exchange. This novel mechanism stems from endogenous information distortions and is orthogonal to the competition dispersion effect, also known as the market thinning effect, documented in the literature (e.g., Bergemann and Bonatti, 2011; Levin and Milgrom, 2010; Rafieian and Yoganarasimhan, 2021).[4] While the marketing thinning effect is derived from the horizontal differentiation of advertisers' impression valuations (e.g., Bergemann and Bonatti, 2011; Hummel and McAfee, 2016), our insights stem from information asymmetries under vertical differentiation where all advertisers want legitimate impressions and no advertiser wants fake ones. Choi et al. (2020) present a summary of the literature and key trends in the area of display advertising; they highlight the emergence of private marketplaces, and how it affects advertisers' and publishers' strategies, as an area for future research.

Finally, the existence of fraud and the fact that bidders only value legitimate impressions make display advertising auctions a special case of *common-value* auctions. In this context, advertisers in the open exchange who do not have access to the private exchange face the *winner's curse* problem: when they win, they know that they probably overpaid for the impression (e.g., Kagel and Levin, 2009). Due to the winner's curse, advertisers shade their bids below their ex ante expected value for the impression, which in turn negatively impacts the publisher's revenue. There are two main differences between our paper and the previous literature on common-value auctions. First, in our model the winner's curse is caused by the information

[4] See Section A2.4 of the Appendix for details.

asymmetry that the publisher creates when it introduces a private exchange; in the previous literature, the winner's curse is caused by the bidders' (noisy) *private signals* of the (common) value of the item. Second, in our model, the negative impact of the winner's curse on the publisher's revenue happens *indirectly* through the lower optimal reserve prices across multiple exchanges (instead of lower bids in a single auction as in the previous literature).

2 Model

The game consists of one publisher and two advertisers. One advertiser is *connected* (denoted by C-advertiser) and the other is *unconnected* (denoted by U-advertiser). The publisher and the advertisers can transact through two platforms, a private exchange and an open exchange (hereafter, PX and OX, respectively). The PX, owned and operated by the publisher, sells ad inventory exclusively to the C-advertiser. In contrast, the OX is open to all ad buyers and sellers (including fraudsters).

In practice, while publishers can provide any advertiser access to their PX, this process involves considerable costs. There are operational costs such as signing contracts and non-disclosure agreements, as well as costs of renting ad tech solutions (akin to renting cloud computing services) from companies that run ad exchanges.[5] These costs deter publishers from indiscriminately providing PX access (Graham, 2020a). Publishers may also incur proprietary and reputational costs. Publishers must trust advertisers to share proprietary information with them as this information could potentially be disclosed to the publishers' competitors (O'Reilly, 2015).[6] Furthermore, the prevalence of malvertising — the use of online advertising to spread malware — presents additional risks of inviting advertisers to private exchanges.[7]

[5]For example, see Admeld (`www.admeld.com`), which was acquired by Google in 2011.

[6]Moreover, data privacy regulations (e.g., the General Data Protection Regulation) increase the risk of non-compliance when publishers share information with third-party advertisers (Benes, 2018).

[7]For instance, The New York Times was hit by malvertising (Hern, 2016) and tweeted its readers to avoid clicking on an "unauthorized ad." The New York Times' Tweet on September 13, 2009 reads "Attn: NYTimes.com readers: Do not click pop-up box warning about a virus – it's an unauthorized ad we are working to eliminate" (`https://twitter.com/nytimes/status/3958547840`).

In sum, publishers incur significant costs in inviting advertisers to their private exchanges. Consequently, the set of advertisers that have access to the PX is a strict subset of those who have access to the OX, such that fewer advertisers compete in the PX than in the OX. To parsimoniously capture this feature, we assume that only one of the two advertisers (i.e., the C-advertiser) has access to the PX. Apart from accessibility to the PX, the C- and U-advertisers are ex ante symmetric.[8]

When the publisher sells through the OX, advertisers bidding in the OX face the fraudster's fake impression with probability β, and the publisher's legitimate impression with probability $1-\beta$. Thus, β measures fraud intensity in the OX.[9] On the other hand, when the publisher sells through the PX, the C-advertiser, who has access to it, faces the publisher's legitimate impression with probability 1, conditional on receiving a request-for-bid.

Note that the volume of impressions available for sale in the PX is smaller than that in the OX. Specifically, for every impression in the OX, there are $1-\beta$ impressions available for sale in the PX. Following the previous literature on online advertising, in the main model we assume that the advertisers consider bidding on every impression opportunity (i.e., the market is *supply-constrained*); as such, we multiply the advertisers' utilities by $1-\beta$ in the PX to account for the smaller volume of impressions available for sale in the PX. In Section 5.1 we consider a *demand-constrained* market where (e.g., due to the large volume of impressions in ad exchanges) advertisers cannot process every impression opportunity. In this case, the fact that there are more impressions available for sale in the OX does not impact the advertisers' utilities, and we show that our qualitative insights continue to hold.

For $j \in \{C, U\}$, the j-advertiser's value for an impression i, denoted by v_{ij}, consists of impression-specific and advertiser-specific factors; i.e.,

$$v_{ij} = \lambda_i \nu_j,$$

[8] In Section A2.3, we analyze a model where the advertisers are ex ante asymmetric and show that our results continue to hold.

[9] While we keep β exogenous in the main model, in Section 4, we consider a scenario in which the OX may endogenously reduce β through anti-fraud efforts.

where λ_i equals 1 if the impression is legitimate and 0 if it is fraudulent. ν_j is the j-advertiser's value for displaying its ad on the publisher's website; it is independently and identically distributed across advertisers according to

$$\nu_j = \begin{cases} \overline{\nu} & \text{with probability } \mu, \\ \underline{\nu} & \text{with probability } 1-\mu, \end{cases} \tag{1}$$

where $0 \leq \underline{\nu} < \overline{\nu}$. We normalize $\underline{\nu}$ to 0 and $\overline{\nu}$ to 1. Advertisers privately know their own realized value of ν before bidding for an impression; the publisher and other advertisers only know the distribution (1). The value of a fraudulent impression is zero for all advertisers. In (1), μ is the probability that an advertiser has a high valuation for an impression (e.g., there is a targeting match) conditional on the impression being legitimate. Given the normalizations of $\underline{\nu}$ and $\overline{\nu}$, μ can also be interpreted as the expected value of an advertiser for a legitimate impression. Depending on their accessibility to the PX, advertisers may or may not know whether an impression is fraudulent before bidding for the impression. We assume that μ and β are common knowledge.[10]

The publisher sells its ad inventory via first-price auctions with reserve prices R^{PX} and R^{OX} in the PX and OX, respectively.[11] When there are two exchanges, the publisher sends request-for-bids for the impression generated on its website to both exchanges simultaneously. Each exchange auctions off the publisher's impression independently and sends the clearing price to the publisher. In a first-price auction, the clearing price equals the highest bid if the bid is greater than or equal to the reserve price, and zero otherwise. After receiving the clearing prices, the publisher allocates the impression to the exchange with the highest clearing price. The publisher's payoff for the impression (from its website) is thus the maximum of the two exchanges' clearing prices.[12] When there is only one exchange, the impression is sold via a

[10]In practice, advertisers can rely on historic data (e.g., previous viewability rates, click-through rates, and conversion rates) to infer μ and β (Fou, 2019).

[11]For more information on the emergence and prevalence of first-price auctions in the RTB market, see Despotakis et al. (2021a).

[12]The process of sending the impression to multiple exchanges simultaneously, and allocating it to the exchange

standard first-price auction and the publisher's payoff is the clearing price. The publisher sets reserve prices R^{PX} and R^{OX} to maximize its expected payoff.

A central feature of our model is the information structure. By virtue of its exclusive connection to the PX, the C-advertiser can identify an impression coming through the PX as originating from the publisher. On the other hand, if the same impression is sent to the OX, then neither the C-advertiser nor the U-advertiser can discern whether it is legitimate or fake. This is because the fraudster sends its request-for-bid in the OX disguised as the publisher, mimicking all aspects of the publisher's request-for-bid, including the reserve price set by the publisher.[13]

Conditional on its bid for impression i exceeding the reserve price, the j-advertiser's expected payoff, when advertisers do not know the impression's legitimacy (i.e., whether $\lambda_i = 0$ or 1), is

$$\begin{aligned}\pi_j(b_j) &= F_{-j}(b_j)\left((\nu_j - b_j)\,\mathbb{P}\{\lambda_i = 1\} + (0 - b_j)\,\mathbb{P}\{\lambda_i = 0\}\right) \\ &= F_{-j}(b_j)\left(\nu_j(1-\beta) - b_j\right),\end{aligned}$$

where F_{-j} is the cumulative distribution function of the competitor's bid b_{-j}. Similarly, its expected payoff when it knows λ_i is

$$\pi_j(b_j|\lambda_i) = F_{-j}(b_j)\cdot\begin{cases}\nu_j - b_j & \text{if } \lambda_i = 1,\\ -b_j & \text{if } \lambda_i = 0.\end{cases}$$

with the highest price, is known as *header bidding* (Sluis, 2016).

[13]Even though the C-advertiser observes legitimate impressions in the PX, it is difficult for the advertiser to identify the same impressions in the OX. First, different exchanges use different identifiers and cookies such that if a publisher's impression is distributed to multiple exchanges, matching them is often not possible. Second, stringent privacy regulations (e.g., the General Data Protection Regulation and the California Privacy Rights Act) coupled with companies' anti-tracking efforts (e.g., Apple's Intelligent Tracking Prevention) have exacerbated the cross-exchange matching problem for advertisers that bid on multiple exchanges (Zawadzinski and Sweeney, 2020). According to IAB Tech Lab (`https://bit.ly/3tb6ArP`), Apple's anti-tracking move "eliminates a major method of establishing a consistent, cookie-based identity for use by third-party advertising systems." Moreover, the authors' conversations with industry experts reveal that practitioners have attempted but failed cross-exchange matching due to technical challenges. Nonetheless, we conduct a robustness check in Section A2.1 of the Appendix wherein the C-advertiser can leverage its information from the PX to distinguish legitimate impressions from fraudulent ones in the OX. We show that the C-advertiser's added informational advantage reduces the PX's "premium" while preserving the devaluation effect. Therefore, the publisher never sells through both exchanges simultaneously in this case.

The above payoff expressions imply that if $\nu = 0$, then regardless of whether the impression is legitimate or fraudulent, the advertiser is better off withdrawing from the auction. Put differently, advertisers submit positive bids only if $\nu = 1$. For ease of exposition, whenever we discuss advertisers with positive bids, we hereafter refer to the *high-valuation advertisers* with $\nu = 1$ simply as "advertisers" without the "high-valuation" qualifier.

The timing of the game is as follows.

1. The publisher sets reserve prices R^{PX} and R^{OX}.
2. The j-advertiser realizes its value ν_j and submits its bids. The C-advertiser bids b_C^{PX} and b_C^{OX} in the PX and OX, respectively. The U-advertiser bids b_U^{OX} in the OX.
3. • For a legitimate impression, each exchange runs a first-price auction and sends its clearing price to the publisher. The publisher allocates the impression to the exchange with the highest clearing price, provided it is greater than 0;[14] otherwise the impression is left unsold.
   • For a fraudulent impression, only the OX runs first-price auctions. If the highest bid is greater than or equal to the reserve price, the fraudulent impression is allocated to the highest bidder; otherwise, the impression is left unsold.

   Finally, payments are made and players' utilities are realized.

3 Analysis

We begin the analysis with the benchmark case in which only the OX exists (see Figure 1a). We then analyze the publisher's ad exchange choices with the option to sell through both the PX and OX (see Figure 1b). The OX-only benchmark corresponds to the earlier days of RTB when the vast majority of RTB inventory was sold through open exchanges. The benchmark analysis serves to elucidate the impact of the introduction of the PX on the RTB market.

[14] A clearing price of 0 implies none of the bids submitted in the exchange exceeded the reserve price.

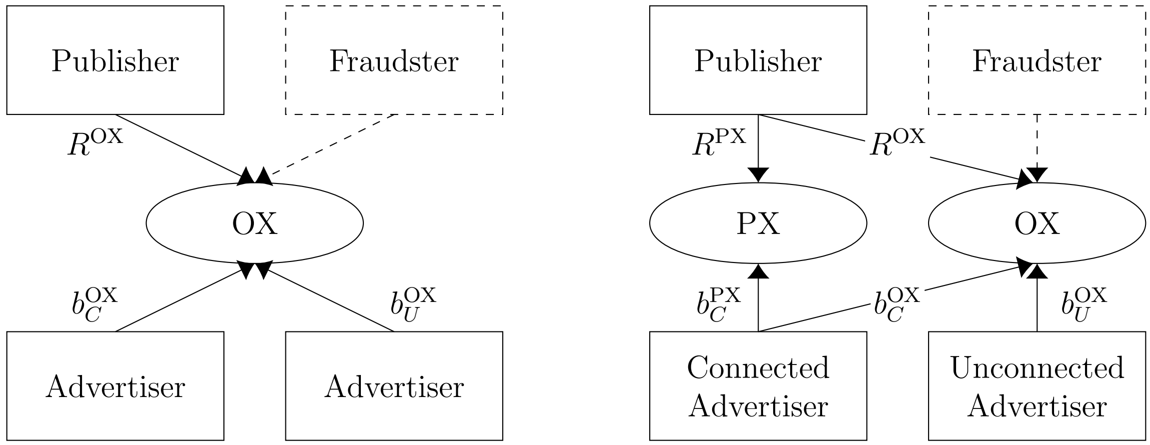

(a) Only OX *(b) PX and OX*

Figure 1: Exchange Configurations

3.1 OX-Only Benchmark

Suppose the publisher can only sell ad inventory through the OX. Due to the open nature of the exchange, advertisers buying in the OX are prone to fraud. Specifically, advertisers cannot distinguish the publisher's legitimate impressions from the fraudster's fake impressions because the fraudster presents itself as the publisher. Note that the fraudster always sets the same reserve price as the publisher; since the advertisers' valuation for the fraudster's impressions is always zero, the game cannot have a separating equilibrium.

Upon seeing a request-for-bid for impression i in the OX, high-valuation advertisers (i.e., advertisers with $\nu_j = 1$) value the impression at

$$\underbrace{1}_{v_{ij}|\text{legitimate}} \cdot \underbrace{(1-\beta)}_{\mathbb{P}\{\text{legitimate}\}} + \underbrace{0}_{v_{ij}|\text{fraudulent}} \cdot \underbrace{\beta}_{\mathbb{P}\{\text{fraudulent}\}} = 1 - \beta.$$

Low-valuation advertisers (i.e., advertisers with $\nu_j = 0$) value it at 0. Based on these valuations, we derive the equilibrium reserve price and bidding strategies, which we summarize in the following lemma.

Lemma 1. *In the OX-only benchmark, the equilibrium reserve price and the advertisers' (sym-*

metric) bids are

$$R^{OX\text{-}only} = b^{OX\text{-}only} = 1 - \beta. \tag{2}$$

The advertisers' expected profits are 0*, and the publisher's expected profit is*

$$\pi_P^{OX\text{-}only} = (2 - \mu)\mu(1 - \beta)^2, \tag{3}$$

where the profit expectations are taken prior to the realizations of advertisers' valuations ν *and of ad type* λ*.*

Lemma 1 shows that the publisher sets the reserve price to the expected value of the high-valuation advertisers ($R^{\text{OX-only}} = 1 - \beta$), and high-valuation advertisers bid the reserve price ($b^{\text{OX-only}} = 1 - \beta$). If only one of the two advertisers is high-valuation, the high-valuation advertiser wins the impression at price $1 - \beta$. If both advertisers are high-valuation, both bid the same amount for the impression and the winner is chosen randomly.[15] If the impression is legitimate, the winning advertiser obtains a positive payoff $1 - (1 - \beta) = \beta$, whereas if the impression is fraudulent, it obtains a negative payoff of $-(1 - \beta)$.

Lemma 1 also shows that the publisher's profit under the OX-only regime is decreasing in β. This reflects the direct, negative effect of fraud: the larger the β, the lower the advertisers' valuations for impressions sold through the OX. Therefore, as β increases, the advertisers bid less and the publisher's profit declines.

3.2 Introduction of PX

We turn to the main analysis where the publisher has the option to sell through a private exchange. The publisher adopts one of the following three regimes: (i) sell only through the OX, (ii) sell only through the PX, and (iii) sell through the PX and OX simultaneously.[16]

[15] The qualitative insights are robust to other tie-breaking rules.

[16] In Section 5.2, we analyze a scenario in which the publisher sells its impressions sequentially, and show that sequential selling does not affect the results.

Under the third regime, the publisher distributes its request-for-bid to both exchanges, and selects the winner based on the exchanges' clearing prices. We compute the publisher's subgame equilibrium profits, and then characterize the publisher's equilibrium exchange choices.

Since the analysis for the OX-only regime is provided in Section 3.1, we omit it here. We analyze in turn the latter two regimes in which the publisher sells only through the PX, and sells through the PX and OX simultaneously.

If the publisher sells exclusively through the PX, then the separation of exchanges reveals fraudster's impressions in the OX. Therefore, no transactions occur in the OX. On the other hand, in the PX, the C-advertiser has valuation 1 with probability μ and valuation 0 with probability $1-\mu$ for the publisher's impression. Therefore, the publisher sets reserve price $R^{\text{PX-only}} = 1$ and the high-valuation C-advertiser bids $b_C^{\text{PX}} = 1$. The following lemma summarizes the advertisers' and the publisher's strategies and their profits under the PX-only regime.

Lemma 2. *In the PX-only regime, the equilibrium reserve price and the C-advertiser's bid in the PX are*

$$R^{PX\text{-}only} = b_C^{PX\text{-}only} = 1. \tag{4}$$

The advertisers' expected profits are 0*, and the publisher's expected profit is*

$$\pi_P^{PX\text{-}only} = \mu(1-\beta), \tag{5}$$

where the profit expectations are taken prior to the realizations of advertisers' valuations ν and of ad type λ.

Comparing the publisher's OX-only profit (3) and PX-only profit (5) reveals the publisher's margin-volume trade-off. If the publisher sells exclusively through the PX, then compared to selling exclusively through the OX, demand for ad slots is lower (i.e., $\mu \leq \mu(2-\mu)$) since only the C-advertiser can bid in the PX. On the other hand, the margin per transaction is higher if it sells exclusively through the PX (i.e., $1 \geq 1-\beta$) because the C-advertiser's knowledge that the PX impressions are legitimate increases its bid in the PX. Finally, note that even though the advertiser pays 1 with probability μ to the publisher, the publisher's utility is $(1-\beta)\mu$ instead of μ. This accounts for the fact that the volume of legitimate impressions in the OX and PX should be the same. Since there are only $(1-\beta)$ legitimate impressions in the OX, we assume that there are only $(1-\beta)$ legitimate impressions available for sale in the PX as well.[17]

Next, consider the third regime in which the publisher sells through the PX and OX simultaneously. Both the PX and OX have positive volume of transactions in equilibrium *only if* the following conditions hold:

1. (individual rationality) the reserve prices are no greater than the advertisers' expected valuations;
2. (incentive compatibility) the C-advertiser's expected profit from bidding in the OX is no greater than that from bidding in the PX; and
3. the C-advertiser's bid in the PX exceeds the U-advertiser's bid in the OX.

The third condition is necessary to sustain the PX-OX regime, because if $b_C^{\mathrm{PX}} < b_U^{\mathrm{OX}}$, the publisher would always sell its legitimate impression to the U-advertiser at a higher price in the OX. In this case, no transactions would occur through the PX.

Before solving for the publisher's optimal reserve prices under these constraints, we characterize the C-advertiser's bidding behavior under the PX-OX regime when it receives two requests-for-bid, one from each exchange.

Lemma 3. *If the C-advertiser receives two requests-for-bids, one from the PX and another from the OX, bidding in both exchanges is weakly dominated by bidding in only the PX.*

Lemma 3 implies that the C-advertiser "single-homes" under the PX-OX regime: in equilibrium, it only bids in the PX. Intuitively, if instead of submitting two bids (one to the PX and another to the OX), the C-advertiser only bids in the PX the *maximum* of those two bids, it maintains the probability of and payoff from winning legitimate impressions, while eliminating the risk of buying fake impressions in the OX.

[17] In Section 5.1, we analyze the case where the volume of impressions is not adjusted (e.g., due to demand constraints), and show that our main results continue to hold.

To illustrate, suppose that the C-advertiser bids $b_C^{\mathrm{OX}} \geq \max\left[b_U^{\mathrm{OX}}, R^{\mathrm{OX}}\right]$ and $b_C^{\mathrm{PX}} \geq R^{\mathrm{PX}}$ in the OX and PX, respectively. In addition, suppose that $\max\left[b_C^{\mathrm{OX}}, b_C^{\mathrm{PX}}\right]$ exceeds the U-advertiser's bid b_U^{OX} in the OX. If the impression is fraudulent (which occurs with probability β), the C-advertiser wins it in the OX for a payoff of $0 - b_C^{\mathrm{OX}}$; if the impression is legitimate (which occurs with probability $1-\beta$), the C-advertiser wins it either in the OX (if $b_C^{\mathrm{PX}} < b_C^{\mathrm{OX}}$) for a payoff of $1 - b_C^{\mathrm{OX}}$, or in the PX (if $b_C^{\mathrm{PX}} \geq b_C^{\mathrm{OX}}$) for a payoff of $1 - b_C^{\mathrm{PX}}$. Thus, the high-valuation C-advertiser's expected payoff (post realization of its ad valuation ν_C and prior to realization of ad type λ) from bidding the tuple $\left(b_C^{\mathrm{OX}}, b_C^{\mathrm{PX}}\right)$ is

$$\beta\left(-b_C^{\mathrm{OX}}\right) + (1-\beta)\left(1 - \max\left[b_C^{\mathrm{OX}}, b_C^{\mathrm{PX}}\right]\right). \tag{6}$$

Now, suppose instead that the C-advertiser bids $\max\left[b_C^{\mathrm{OX}}, b_C^{\mathrm{PX}}\right]$ in the PX and does not bid in the OX. If the impression is fraudulent, its payoff would be 0 since it does not participate in the OX auction; if the impression is legitimate, the C-advertiser wins it in the PX for a payoff of $1 - \max\left[b_C^{\mathrm{OX}}, b_C^{\mathrm{PX}}\right]$. Thus, the C-advertiser's expected payoff from "single-homing" is

$$(1-\beta)\left(1 - \max\left[b_C^{\mathrm{OX}}, b_C^{\mathrm{PX}}\right]\right). \tag{7}$$

Since (7) is greater than (6), "single-homing" in the PX dominates "multi-homing" in both exchanges. Other cases (e.g., $\max\left[b_C^{\mathrm{OX}}, b_C^{\mathrm{PX}}\right] < b_U^{\mathrm{OX}}$) yield the same result (see proof of Lemma 3).

To derive the high-valuation U-advertiser's individual rationality constraint, note that the U-advertiser can win the legitimate impression in the OX only if the C-advertiser has low valuation; otherwise, the C-advertiser buys the impression (at a higher price) in the PX. Therefore, the U-advertiser's expected profit (post realization of its ad valuation ν_U and prior to ad type realization λ) from bidding the reserve price in the OX is

$$(1-\mu)(1-\beta)(1-R^{\mathrm{OX}}) + \beta(0 - R^{\mathrm{OX}}), \tag{8}$$

where the first summand is the expected payoff from a legitimate impression, and the second the expected payoff from a fake impression. For the U-advertiser to bid in the OX, the expected utility (8) has to be non-negative. This individual rationality constraint simplifies to $R^{\mathrm{OX}} \leq \frac{(1-\mu)(1-\beta)}{1-(1-\beta)\mu}$, which binds in equilibrium.

Next, we derive the high-valuation C-advertiser's incentive compatibility constraint. If it participates in the PX, its expected profit (post realization of its ad valuation ν_C and prior to the realization of ad type λ) is $(1-\beta)\left(1-R^{\mathrm{PX}}\right)$. If the C-advertiser deviates to the OX and bids infinitesimally higher than R^{OX}, it would win in the OX for expected profit of $(1-\beta)(1-R^{\mathrm{OX}})+\beta(0-R^{\mathrm{OX}})$. Thus, the incentive compatibility constraint simplifies to

$$(1-\beta)\left(1-R^{\mathrm{PX}}\right) \geq 1-\beta-R^{\mathrm{OX}}, \tag{9}$$

from which it follows that $R^{\mathrm{PX}} \leq \frac{1}{1-\beta} R^{\mathrm{OX}}$; this constraint binds in equilibrium. Note that, akin to standard product line design settings (e.g., Moorthy, 1984), the C-advertiser's incentive compatibility constraint is stricter than the C-advertiser's individual rationality constraint, which is $R^{\mathrm{PX}} \leq 1$; thus, the C-advertiser receives positive surplus in equilibrium. The following lemma summarizes the players' strategies and payoffs under the dual exchange regime.

Lemma 4. *Let b_C^{PX} and b_U^{OX} denote the C-advertiser's bid in the PX and the U-advertiser's bid in the OX, respectively. In the PX-OX regime where the publisher sells through the PX and OX simultaneously, the equilibrium reserve prices and the advertisers' bids are*

$$R^{PX} = b_C^{PX} = \frac{1-\mu}{1-(1-\beta)\mu} \text{ and } R^{OX} = b_U^{OX} = \frac{(1-\mu)(1-\beta)}{1-(1-\beta)\mu}. \tag{10}$$

The C-advertiser's expected profit is

$$\pi_C^{PX\text{-}OX} = \frac{\mu^2(1-\beta)\beta}{1-(1-\beta)\mu}, \tag{11}$$

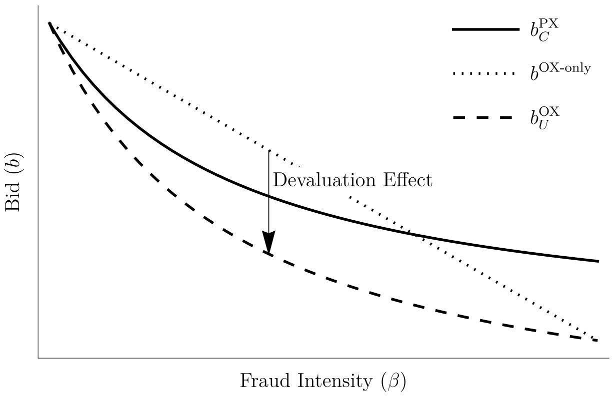

(a) With Respect to β

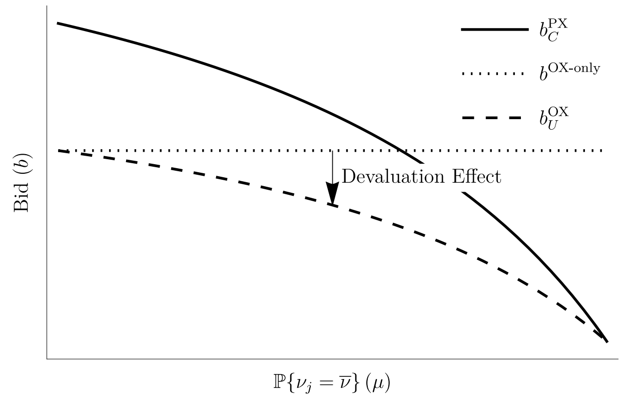

(b) With Respect to μ

Figure 2: Advertisers' Bids and Devaluation Effect

the U-advertiser's expected profit is 0, and the publisher's expected profit is

$$\pi_P^{PX\text{-}OX} = \frac{(1-\mu)\mu(2-\mu-\beta(1-\mu))(1-\beta)}{1-(1-\beta)\mu}, \tag{12}$$

where the profit expectations are taken prior to realizations of advertisers' valuations ν and of ad type λ.

Lemma 4 reveals important insights regarding the U-advertiser's bidding strategy under the PX-OX regime. First, the U-advertiser bids lower under the PX-OX regime than under the OX-only regime; i.e., b_U^{OX} in (10) is less than $b_U^{\text{OX-only}}$ in (2) (see Figure 2). The intuition is as follows. In the presence of the PX, the U-advertiser knows that it competes against the informationally advantaged C-advertiser, who bids high in the PX (for the publisher's legitimate impression) and bids nothing in the OX. Thus, conditioned on winning, the U-advertiser's probability of having won a fraudulent impression is higher, compared to the OX-only benchmark where both advertisers are equally uninformed. In other words, the introduction of the PX creates an information asymmetry between the advertisers that dampens the U-advertiser's valuation for impressions in the OX. We call this the *devaluation effect.*

Second, the reserve price in the PX is set lower, and the C-advertiser bids lower, under the PX-OX regime than under the PX-only regime; i.e., $\frac{1-\mu}{1-(1-\beta)\mu} \leq 1$. In contrast to the PX-only regime, under the PX-OX regime, the publisher cannot raise the reserve price in the PX to 1, even though the high-valuation C-advertiser in the PX knows that the impression is legitimate and values it at 1. The reason is that under the PX-only regime, the C-advertiser has no outside option: if it does not win the impression in the PX, its expected payoff is zero. This allows the publisher to maximally raise the reserve price to 1, thereby extracting all of the C-advertiser's surplus. Under the PX-OX regime, however, if the reserve of the PX is set too high, the C-advertiser switches to buying (potentially fraudulent) impressions in the OX. That is, if the reserve price in the PX is too high, the C-advertiser's expected payoff will be higher buying in the OX with the risk of ad fraud than buying a guaranteed legitimate impression in the PX at a high price. In sum, the OX cannibalizes the PX and reduces the publisher's revenue from the PX.

The devaluation and cannibalization effects jointly lower the publisher's revenue from the C-advertiser. The devaluation effect lowers the U-advertiser's bids in the OX, which in turn softens bidding competition for the C-advertiser, leading to lower bids in the PX. Due to the cannibalization effect, the publisher cannot set a high reserve price in the PX to offset the devaluation effect in the OX. That is, if R^{PX} is too high, the C-advertiser will switch to bidding in the OX. In the following proposition, we summarize the central force generated by the introduction of the PX.

Proposition 1 (Devaluation Effect)**.** *The introduction of the PX may soften bidding competition. Specifically, the U-advertiser bids lower under the PX-OX regime than under the OX-only regime. Moreover, the C-advertiser bids lower under the PX-OX regime than under the OX-only regime if and only if $\mu > \frac{1}{2}$ and $\beta \leq 2 - \frac{1}{\mu}$.*

Figures 2a and 2b reveal interesting relationships between the devaluation effect and the parameters β and μ. First, the devaluation effect (the difference between $b^{\text{OX-only}}$ and b_U^{OX}, represented by dotted and dashed lines in Figure 2, respectively) first amplifies then diminishes in β. It amplifies in β because the U-advertiser's probability of winning fraudulent impressions increases in β, which lowers the U-advertiser's valuation. The devaluation effect then diminishes in β because regardless of the presence of the PX, the U-advertiser's valuation of ad impressions in

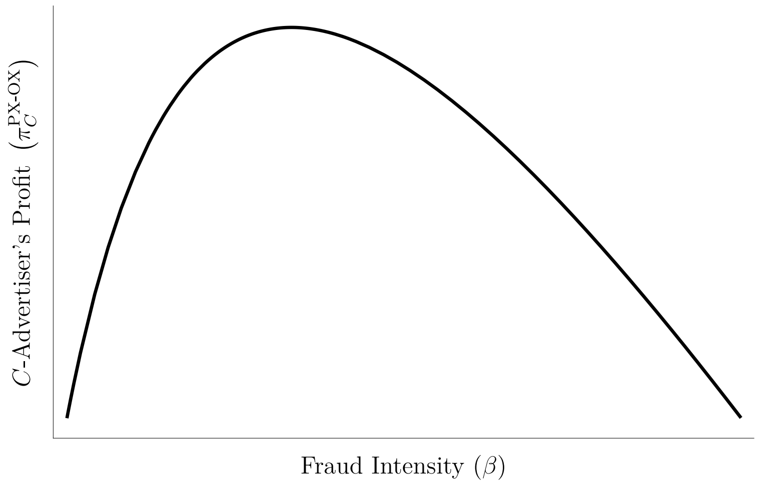

Figure 3: C-Advertiser's Profit Under PX-OX Regime

the OX decrease to 0 as β approaches 1.

Second, Figure 2b illustrates the devaluation effect amplifying in μ, the probability that advertisers realize high valuations. The reason is that as μ increases, the U-advertiser anticipates a higher probability of facing a high-valuation C-advertiser, who bids higher for the legitimate impression in the PX than the U-advertiser does in the OX. This implies that conditioned on winning, the U-advertiser has a higher probability of having won a fraudulent impression. Therefore, the U-advertiser discounts its bid more steeply as μ increases.

Interestingly, under the PX-OX regime, higher fraud intensity has non-monotonic effects on the C-advertiser's expected profit (see Figure 3). On the one hand, fraud amplifies the U-advertiser's devaluation effect: beyond the direct, negative effect of fraud on the U-advertiser's valuation, the U-advertiser's informational disadvantage under the PX-OX regime hurts the U-advertiser more acutely as β increases. In response, the publisher lowers the reserve price in the OX, which in turn exacerbates the OX's cannibalization of the PX. Overall, the publisher leaves more surplus for the C-advertiser in the PX as β increases.

On the other hand, higher fraud depresses the C-advertiser's profit as larger β implies that at any point in time, an impression generated is more likely to be from a fraudster than from the publisher. Thus, on average, the C-advertiser faces fewer opportunities to buy legitimate impressions through the PX as β increases. These countervailing forces make the impact of β

on the C-advertiser's profit non-monotonic. The following proposition summarizes this finding.

Proposition 2 (Connected Advertiser Profit)**.** *Under the PX-OX regime, the C-advertiser's expected profit (prior to realizations of advertiser valuation ν and of ad type λ) increases in β if $\beta \leq \frac{\sqrt{1-\mu}-1+\mu}{\mu}$, and decreases in β otherwise.*

In summary, the comparison of the OX-only benchmark with the regimes with the PX sheds light on important insights regarding the impact of the introduction of the PX on the RTB market. The introduction of the PX distorts the advertisers' information structure such that the U-advertiser values impressions less than it does without the PX. The publisher proportionately lowers the reserve price in the OX, which makes bidding in the OX more attractive for the C-advertiser. Therefore, due to the cannibalization effect, the publisher lowers the reserve price in the PX as well. This allows the C-advertiser to win PX impressions at a low price. As such, the C-advertiser's profit under the PX-OX regime may increase in fraud intensity.

In the following section, we discuss the impact of the various forces related to fraud (e.g., devaluation effect and cannibalization effect) on the publisher's exchange choices.

3.3 Equilibrium

In this section, we characterize the publisher's equilibrium exchange choices. The following proposition shows that all three regimes — OX-only, PX-only, and PX-OX — can emerge in equilibrium.

Proposition 3 (Exchange Configuration)**.** *The publisher's equilibrium ad exchange choices are as follows:*

1. *if* $\frac{3-2\beta-\sqrt{4\beta^2-8\beta+5}}{2(1-\beta)} < \mu$ *and* $\beta \leq \frac{1-\mu}{2-\mu}$*, the publisher sells only through the OX;*
2. *if* $\max\left[\frac{1-\mu}{2-\mu}, \frac{(1-\mu)^2}{\mu^2-\mu+1}\right] < \beta$*, the publisher sells only through the PX;*
3. *otherwise, the publisher sells through both the PX and OX.*

Proposition 3 shows that even if the publisher has the option to sell through the PX, which

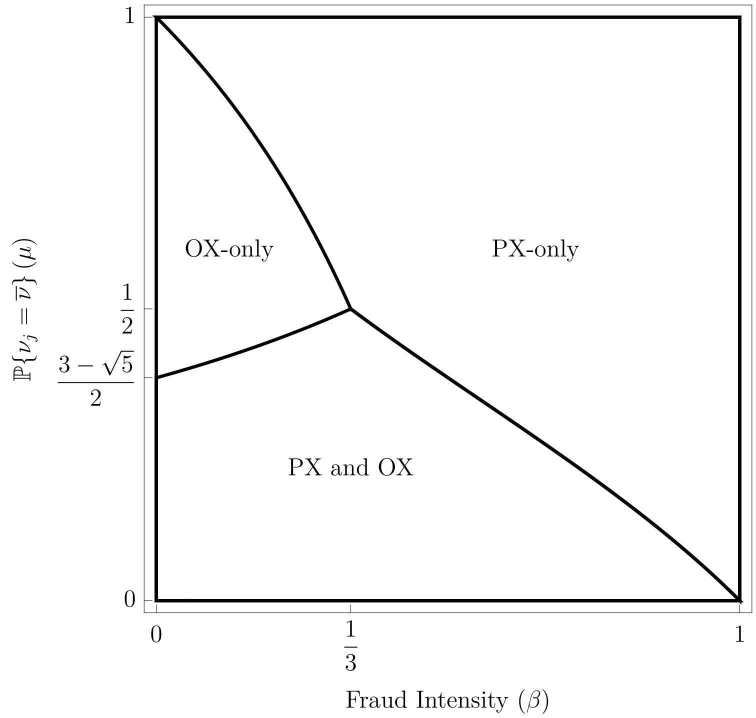

Figure 4: Publisher's Equilibrium Exchange Choices

helps the C-advertiser distinguish legitimate impressions from fake ones, it does not always choose to do so (see Figure 4). To understand the underlying forces at play, we partition the parametric region into two sub-regions: small β (i.e., $\beta \leq 1/3$) and large β (i.e., $\beta > 1/3$).

Suppose β is small. As μ increases, the exchange configuration changes from the PX-OX, to OX-only, to PX-only. If μ is small, the publisher sells through PX and OX, thereby capitalizing on the OX's thick market and on the C-advertiser's high valuation in the PX. However, the devaluation effect amplifies with μ as the U-advertiser anticipates a higher probability of the legitimate impression being "poached" by the high-valuation C-advertiser in the PX. Therefore, as μ increases from low to intermediate range, the publisher switches from PX-OX to OX-only. Interestingly, if μ increases further, the publisher switches from OX-only to PX-only. Intuitively, if μ is large, the C-advertiser will likely draw a high valuation and buy the publisher's legitimate impression in the PX at a high price. This attenuates the OX's relative benefit of a thick market such that the publisher switches from OX-only to PX-only for large μ.[18]

[18] If $\beta = 0$, the publisher is indifferent between selling through OX-only and selling through PX-OX, since the respective profits equal $(2-\mu)\mu$ (see Equations (3) and (12)).

Finally, suppose β is large. Then selling through OX-only is unprofitable because the advertisers' valuations for the OX impressions decrease with fraud intensity. Therefore, for small μ where the devaluation effect is mild, the publisher sells through the PX and OX. On the other hand, if μ is large, the cannibalization effect of the OX outweighs the benefit of selling to the U-advertiser. Thus, by selling through PX-only, the publisher capitalizes on the C-advertiser's high valuation without significantly increasing the risk of the ad slot going unfilled.[19]

The publisher's exchange choice is similar to the product line design problem (e.g., Desai, 2001; Moorthy, 1984; Villas-Boas, 2004). It involves determining the optimal type and number of exchanges to offer to advertisers in the presence of cannibalization effects. However, the exchange choice is qualitatively different from the standard product line design setting due to its effect on the advertisers' information structure. The introduction of the PX not only ensures a "higher quality" for advertisers who buy in the PX, but also informationally disadvantages advertisers who do not have access to the PX, which in turn lowers their valuations. In total, while the cannibalization effect deters the publisher from selling through the OX, the low-quality analogue, the devaluation effect induced by the information asymmetry deters it from selling through the PX, the high-quality analogue. Under the first condition outlined in Proposition 3, the devaluation effect is so severe that the publisher forgoes selling through the PX altogether. That is, a product line-optimizing monopolist forgoes offering the high quality option due to its information distortion effect that softens bidding competition.

4 OX and Anti-Fraud Efforts

The main model assumed the OX to be passive. While this assumption allowed us to obtain sharp insights about the effect of introducing the PX on the RTB market, the OX may play a more active role in gatekeeping the types of ad impressions it sells (Graham, 2020b). In this

[19]The equilibrium exchange configurations in Proposition 3 carry over even if the OX can fight fraud (see Section 4), except if $\frac{3-\sqrt{5}}{2} < \mu \leq \min\left[3 - 2\beta - \sqrt{4\beta^2 - 8\beta + 5}, 2 - \beta - \sqrt{(4-3\beta)\beta}\right]/2(1-\beta)$. In this parametric region, OX's anti-fraud efforts induce a switch from PX-OX to OX-only.

section, we explore the OX's incentives (or lack thereof) to fight fraud and analyze how the OX's strategy may affect the qualitative insights from the main model. To that end, we augment the main model such that the OX filters out $\gamma \in [0, 1]$ fraction of the fraudulent impressions in the OX.[20] In practice, exchanges filter out fraudulent impressions in real time; as such, in our model we assume that γ is set in Period 1 of the game, at the same time as the publisher sets the reserve price(s).

To facilitate exposition, we assume that it is costless for the OX to identify and remove fake impressions. Furthermore, consistent with industry practice, we assume that the OX's profit is based on a fixed commission rate per transaction occurring through the OX (Hsiao, 2020). We fix the commission rate arbitrarily small to mute the effects associated with the magnitude of the commission rate.[21]

We begin the analysis for the OX-only benchmark, and then analyze the OX's equilibrium filter level with the PX. In the benchmark scenario without the PX, fighting fraud has two effects on the OX's profit. First, it reduces the OX's profit because filtering out fraudulent impressions decreases the OX's transaction volume. Second, fighting fraud increases the OX's margin per transaction because advertisers' valuations increase as fraud decreases. The following lemma shows that the former negative effect associated with volume-reduction always dominates. This result is consistent with reports of publishers complaining to open exchanges about their lack of anti-fraud efforts (Rowntree, 2019).

Lemma 5. *In the OX-only benchmark, fighting fraud reduces the OX's expected profit.*

The benchmark analysis reveals that the OX has no incentive to fight fraud if the publisher sells exclusively through the OX. The intuition is as follows. Even though fighting fraud increases the advertisers' valuations in the OX and leads to higher equilibrium bids by the advertisers, the incremental surplus per impression will be shared between the OX, the publisher, and the advertisers. In other words, the OX does not fully reap the benefits of increasing the advertisers'

[20] It is worth mentioning that if the publisher were able to costlessly fight fraud in the OX, then it would always filter out all fake impressions (i.e., $\gamma^* = 1$); see Claim 4 in Section WA7 of the Web Appendix for details.

[21] In Section 5.3, we analyze a scenario in which the OX's commission rate is fully endogenized.

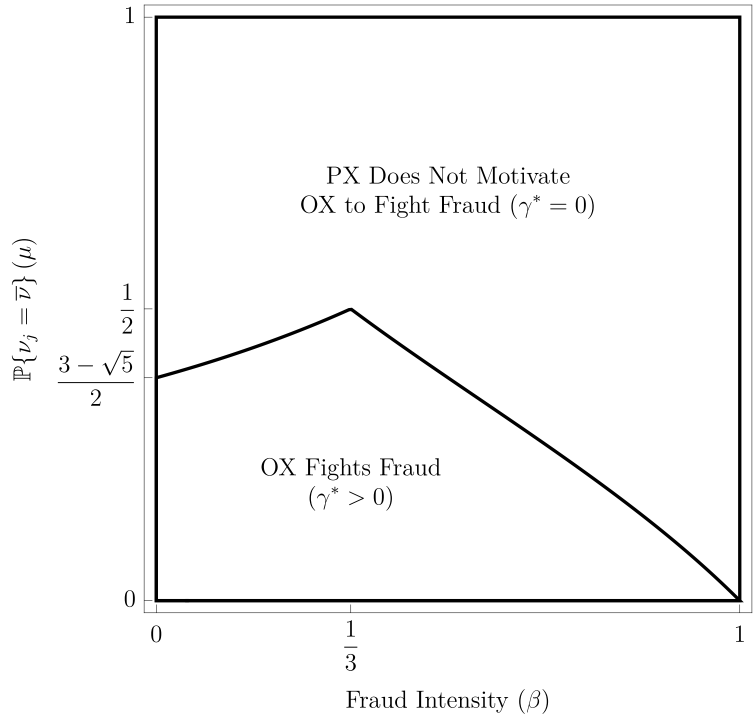

Figure 5: OX's Equilibrium Anti-Fraud Efforts

valuations while it fully internalizes the cost of reduced transaction volume.[22]

Interestingly, the OX's anti-fraud incentive changes qualitatively if the publisher has the option to sell through the PX. In particular, the introduction of the PX creates competitive pressure that induces the OX to combat fraud. By reducing fraudulent request-for-bids coming through the OX, the OX induces the C-advertiser to switch from bidding in the PX to bidding in the OX. The following proposition characterizes the conditions under which the OX combats fraud.

Proposition 4 (Open Exchange's Anti-Fraud Effort)**.** *If the publisher has the option to sell through the PX, then the OX fights fraud (i.e., $\gamma^* > 0$ in equilibrium) if and only if $\mu \leq \frac{3-2\beta-\sqrt{4\beta^2-8\beta+5}}{2(1-\beta)}$ and $\beta \leq \frac{(1-\mu)^2}{1-\mu+\mu^2}$. Furthermore, under these conditions, the OX's anti-fraud efforts decrease the C-advertiser's profit.*

Proposition 4 shows that the OX fights fraud if and only if μ and β are sufficiently small (see

[22]Lemma 5 is robust to settings where (a) advertisers have a positive outside option, or (b) the OX can publicly and credibly commit to γ for the publisher to adjust its reserve price (see Claim 5 in Section WA8 of the Web Appendix). However, the OX fights fraud if it can publicly and credibly commit to γ before the reserve prices are set *and* it incurs reputational costs from selling fraudulent impressions (see Claim 6 in Section WA9 of the Web Appendix).

Figure 5). These are the conditions under which the publisher has incentive to sell through both the PX and OX (see Figure 4).If it is optimal for the publisher to sell only through the OX even without any anti-fraud efforts from the OX (i.e., $\gamma = 0$), the OX has no incentive to fight fraud. On the other hand, if the market conditions are such that the publisher has incentive to sell through both exchanges, the OX benefits from fighting fraud. The intuition is that fighting fraud mitigates the devaluation effect, which in turn increases the U-advertiser's bid; and higher bids implies more transactions through the OX at higher margins.[23]

Proposition 4 highlights another interesting aspect of the OX's anti-fraud efforts. Since the anti-fraud efforts of the OX mitigates the informational disadvantage of the U-advertiser, these efforts may hurt the C-advertiser. In other words, OX's anti-fraud efforts induce the U-advertiser to bid higher, which in turn intensifies bidding competition and, ultimately, lowers the C-advertiser's profit.

Finally, Proposition 4 reveals a hidden blessing of PX from a regulatory perspective. If fraud in the system is sufficiently mild, then competition will induce the market to self-regulate fraud, albeit not completely. On the other hand, if fraud is severe, then exchanges will have little incentive to combat fraud. In such cases, regulatory intervention may be required to protect the RTB industry from fraud-based welfare losses, which industry experts estimate to be substantial (He, 2019).

5 Extensions

In this section, we test the robustness of the key insights — i.e., the devaluation effect and its impact on the exchange configurations — by analyzing alternative model specifications. In Section 5.1, we analyze a scenario in which advertisers are constrained from buying all impressions available in the exchanges. In Section 5.2, we investigate the impact of the publisher's sequen-

[23]Proposition 4 implies that if the OX can fight fraud costlessly, then the PX-OX equilibrium under the third condition of Proposition 3 changes to OX-only regime. However, we demonstrate in Section A2.2 of the Appendix that when the OX incurs a positive marginal cost for filtering fraudulent impressions, all three regimes can emerge in equilibrium despite the OX's anti-fraud efforts.

tial (vs. simultaneous) bid solicitation on the advertisers' bidding strategies and the publisher's exchange configuration. In Section 5.3, we extend the main model by allowing the OX to endogenously determine the commission rate for its platform. We discuss the robustness of the main results as well as shed light on novel insights stemming from the model variations.

5.1 Demand-Constrained Market

In the main model, we assumed a supply-constrained market where advertisers analyze (and consider buying) all impressions available for sale. This assumption is consistent with how online advertising markets largely operate, and is standard in the academic literature. In this section, we investigate a slightly different model wherein advertisers can only look at one impression (out of many) in each exchange. In other words, due to the operational costs and technological limitations, advertisers cannot process (and calculate their valuations for) every impression.[24] In a demand-constrained model, the volume of impressions in each exchange does not impact the advertisers' utilities; as such, we no longer multiply their utilities by $(1-\beta)$ to account for the lower volume of impressions in the PX.

The publisher's problem in the PX-only regime is the same as the main model. Since legitimate and fake impressions are separated by the exchanges, the OX collapses and the publisher sets the reserve price in the PX to 1, the C-advertiser's maximum valuation. With probability μ, the C-advertiser has high-valuation and buys the impression; therefore, the publisher's expected profit is μ. Note that the only difference between the publisher's expected profit here and that in the main model, $(1-\beta)\mu$, is the volume adjustment multiplier $1-\beta$. In the OX-only regime, the publisher sets the same reserve price of $1-\beta$ as in the main model. The probability that at least one advertiser has high valuation is $(2-\mu)\mu$, and the probability that the sold impression is indeed the publisher's impression (as opposed to a fake impression) is $1-\beta$. Thus, the publisher's expected profit in the OX-only regime is $(2-\mu)\mu(1-\beta)^2$, which is the same as (3)

[24] According to DriveScale (2017), AppNexus (now Xandr) executed "more than 10 billion transactions per day," which translates to approximately 116,000 transactions per second, "each a real-time auction conducted in a fraction of a second."

in the main model.[25]

The impact of the C-advertiser's demand constraint on the publisher's and advertisers' strategies is more nuanced under the PX-OX regime as it affects the C-advertiser's incentive compatibility constraint. The publisher's problem under the PX-OX regime is

$$\begin{aligned}
\max_{R^{\mathrm{PX}},\, R^{\mathrm{OX}}} \quad & \mu R^{\mathrm{PX}} + (1-\mu)\mu(1-\beta)R^{\mathrm{OX}} \\
\text{subject to} \quad & 1 - R^{\mathrm{PX}} \geq 1 - \beta - R^{\mathrm{OX}}, && (13) \\
& R^{\mathrm{PX}} \leq 1, \\
& (1-\beta)(1-\mu)(1-R^{\mathrm{OX}}) - \beta R^{\mathrm{OX}} \geq 0. && (14)
\end{aligned}$$

The three constraints represent, respectively, the C-advertiser's incentive compatibility (IC) constraint, the C-advertiser's individual rationality (IR) constraint, and the U-advertiser's IR constraint. The U-advertiser's IR constraint (14) is identical to that from the main model in that the U-advertiser's valuation of an OX impression is depressed by its knowledge that the C-advertiser takes advantage of its PX connection to cherry-pick the legitimate impression, raising the probability of winning fraudulent impressions in the OX.

The C-advertiser's IC constraint (13), however, is different from the main model. While the C-advertiser still compares its profit from buying a legitimate impression in the PX with that from buying an uncertain impression in the OX, it does so without accounting for the thinner volume in the PX (i.e., without the $1-\beta$ factor as in the main model). This makes the PX "premium" even greater for the C-advertiser under the demand-constraint assumption, and allows the publisher to charge higher reserve prices in the PX than in the main model.

[25] Since advertisers buy a single impression in the demand-constrained market, the publisher's unconditional profit should account for the measure of a single impression, which is $\epsilon \to 0$. Therefore, to be more precise, the publisher's expected profit in each regime should equal the above expressions multiplied by ϵ (e.g., $\epsilon\mu$ instead of μ in the PX-only regime, and $\epsilon(2-\mu)\mu(1-\beta)^2$ instead of $(2-\mu)\mu(1-\beta)^2$ in the OX-only regime). To facilitate exposition, however, we suppress ϵ as it only serves as a scaling factor.

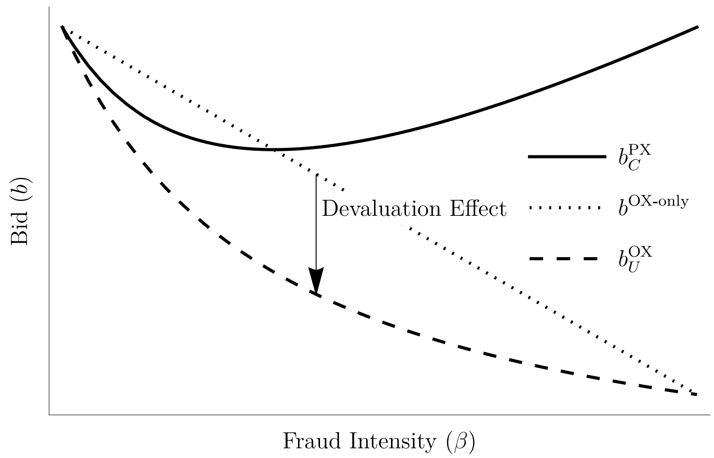

(a) Advertisers' Bids

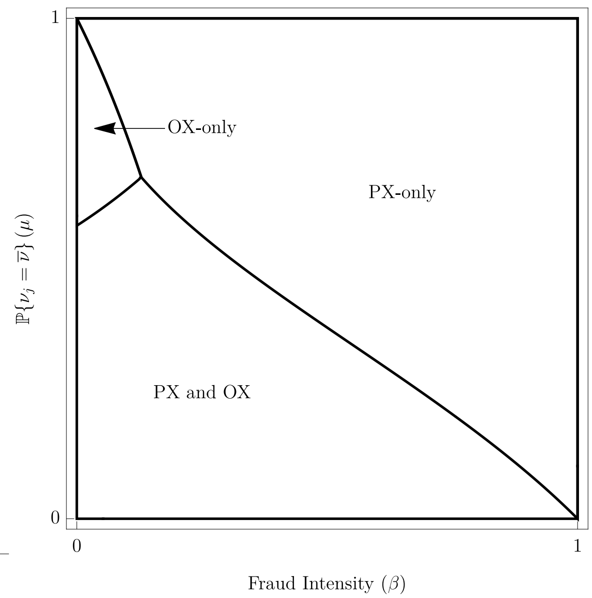

(b) Publisher's Equilibrium Exchange Choices

Figure 6: Bidding Strategies and Exchange Choices Under Demand-Constrained Model

The publisher's equilibrium reserve prices are

$$R^{\text{PX}} = R^{\text{OX}} + \beta, R^{\text{OX}} = \frac{(1-\mu)(1-\beta)}{1-(1-\beta)\mu},$$

and its expected profit is

$$\pi_P^{\text{PX-OX}} = \frac{\mu\left(\beta^2\left(\mu^2-\mu+1\right) - 2\beta(1-\mu)^2 + \mu^2 - 3\mu + 2\right)}{1-(1-\beta)\mu}.$$

The following proposition summarizes the publisher's equilibrium exchange configuration when it sells under advertisers' demand constraints.

Proposition 5 (Exchange Configuration Under Demand Constraint)**.** *The publisher's equilibrium ad exchange choices are as follows:*

1. *if* $\frac{\beta^2-3\beta-\sqrt{(2-\beta)(1-\beta)^3}+2}{(1-\beta)^2} < \mu$ *and* $\beta \leq 1 - \frac{\sqrt{2-\mu}}{2-\mu}$*, the publisher sells only through OX;*
2. *if* $\max\left[1 - \frac{\sqrt{2-\mu}}{2-\mu}, \frac{(1-\mu)^2}{\mu^2-\mu+1}\right] < \beta$*, the publisher sells only through the PX;*
3. *otherwise, the publisher sells through both the PX and OX.*

Proposition 5 shows that the qualitative insights from the main model largely carry over. In particular, the devaluation effect that critically shapes the equilibrium outcome in the main model remains operative under the demand-constrained assumption.

Interestingly, however, the C-advertiser's bid in the PX-OX regime no longer monotonically decreases in β; instead, it is U-shaped in β (see Figure 6a). For small β, the C-advertiser's bid decreases in the fraud rate for the same reason as in the main model: the C-advertiser anticipates softer competition from the OX due to the devaluation effect. However, if β is large, the probability of buying a legitimate impression in the OX is low under the demand constraint, such that the advertisers' total revenue from the OX decreases. Therefore, for large β, the relative benefit of the PX increases in β, which in turn increases the C-advertiser's valuation of PX impressions. It follows that the C-advertiser's bid increases in β for large β.

Figure 6b illustrates the impact of fraud and average ad valuation on the publisher's exchange choices under the demand-constraint assumption. The demand-constrained model mutes the PX's thin-volume effect, thereby increasing the PX's appeal for advertisers. As such, the publisher's incentive to sell through the OX decreases. Comparing Figure 6b to Figure 4 (from the main model) shows that the demand constraint shrinks the parametric region under which the publisher sells exclusively through the OX. While the equilibrium boundaries are modified quantitatively, the qualitative insights from the main model are preserved.

5.2 Sequential Bid Request (Waterfalling)

The main model assumed that, consistent with the shift in industry practice, the publisher adopts header bidding; i.e., it broadcasts its request-for-bids simultaneously to all exchanges (Despotakis et al., 2021a). In this section, we analyze the case of waterfalling, in which the publisher prioritizes its ad inventory sales to the C-advertiser in the PX, and then sells remnant inventory to all advertisers in the OX.

We find that the advertisers' bidding strategies, and hence the publisher's profit under sequential selling are identical to those under simultaneous selling. First, note that for the sequential PX-

OX configuration to be an equilibrium, the publisher has to set reserve prices such that the C-advertiser buys only from the PX in Period 1 and the U-advertiser bids alone in the OX in Period 2. If the reserve price were set such that the C-advertiser buys from the OX in Period 2, this would be equivalent to the regime in which the publisher sells exclusively through the OX.

Second, consider the C-advertiser's bidding strategy. If it bids in the PX, then its expected profit (post realization of its ad valuation ν_C and prior to ad type realization λ) is $(1-\beta)(1-R^{\text{PX}})$. On the other hand, if it foregoes buying the PX impression in Period 1, then the C-advertiser knows that the publisher's unsold impression will be sold through the OX in Period 2. Thus, if it bids in the OX infinitesimally higher than the U-advertiser's bid of R^{OX}, it would obtain expected profit (post ν_C realization and prior to λ realization) of $1-\beta-R^{\text{OX}}$. This yields the same incentive compatibility constraint as that from the simultaneous selling case in (9). Therefore, under PX-OX, the C-advertiser's bidding strategy, the publisher's profit, and its equilibrium exchange choices in the simultaneous case carry over to the sequential setting. We state this result in the following lemma.

Lemma 6. *The equilibrium outcomes under simultaneous selling is equivalent to those under sequential selling.*

5.3 Endogenous OX Commission

In this extension, we test the robustness of our insights in a setting where the OX plays a more active role. While the main model assumes a passive OX with an exogenously fixed commission rate α, which is arbitrarily close to 0, this extension allows the OX to set the commission rate $\alpha \in [0,1]$ to maximize its expected profit.[26] The OX first sets α, and then the publisher decides its exchange configuration (i.e., OX-only, PX-only, or PX-OX). If an impression (be it legitimate or fraudulent) is sold through the OX, then the OX takes α share of the transaction amount; if the impression is legitimate, the remaining $1-\alpha$ share goes to the publisher.

[26] In practice, commission rates vary across platforms (`https://adalytics.io/blog/adtech-supply-fees`). According to Sluis (2020), Google charged approximately 16% for programmatic ads.

We find that the qualitative insights from the main model, where α is arbitrarily small, carry over. While the equilibrium commission rate is generally higher under endogenous commission than in the main model, the equilibrium exchange configurations from the main model are largely preserved. Specifically, endogenizing the OX's commission rate changes the equilibrium exchange configurations only if OX-only is the equilibrium outcome in the main model.

To understand the impact of endogenous commission, consider the three equilibrium regimes from the main model. If PX-only is the equilibrium in the main model, then it must be so under endogenous commission as well. The reason is that α can only be larger under endogenous commission than in the main model; this makes the PX-OX and OX-only configurations even less profitable for the publisher than in the main model. Similarly, if PX-OX is equilibrium in the main model, it must also be so under endogenous commission. As α increases, selling through OX-only becomes less profitable for the publisher than PX-OX because more transactions occur through the OX in the former; therefore, if the PX-OX configuration is the equilibrium outcome in the main model, the equilibrium outcome cannot be OX-only under endogenous commission. Similarly, while the OX may set a sufficiently large α such that the publisher switches from PX-OX to PX-only, doing so would reduce the OX's profit to zero; therefore, if the PX-OX is the equilibrium outcome in the main model, the equilibrium outcome under endogenous commission cannot be PX-only. Taken together, endogenizing the commission rate does not change the equilibrium exchange configuration if PX-only or PX-OX is the equilibrium in the main model. The following lemma states this preliminary insight.

Lemma 7. *If the OX sets endogenous commission rate α, then the equilibrium exchange configurations change relative to the main model only if $\frac{3-2\beta-\sqrt{4\beta^2-8\beta+5}}{2(1-\beta)} < \mu$ and $\beta \leq \frac{1-\mu}{2-\mu}$.*

Lemma 7 implies that, to assess the robustness of the exchange configurations from the main model, it suffices to examine how OX's endogenous commission decision affects the equilibrium outcome in the OX-only regime. Within this parametric region (i.e., $\frac{3-2\beta-\sqrt{4\beta^2-8\beta+5}}{2(1-\beta)} < \mu$ and $\beta \leq \frac{1-\mu}{2-\mu}$), we find that the equilibrium outcome under endogenous α diverges from the main model if and only if μ is sufficiently small. Specifically, the OX induces the publisher to sell

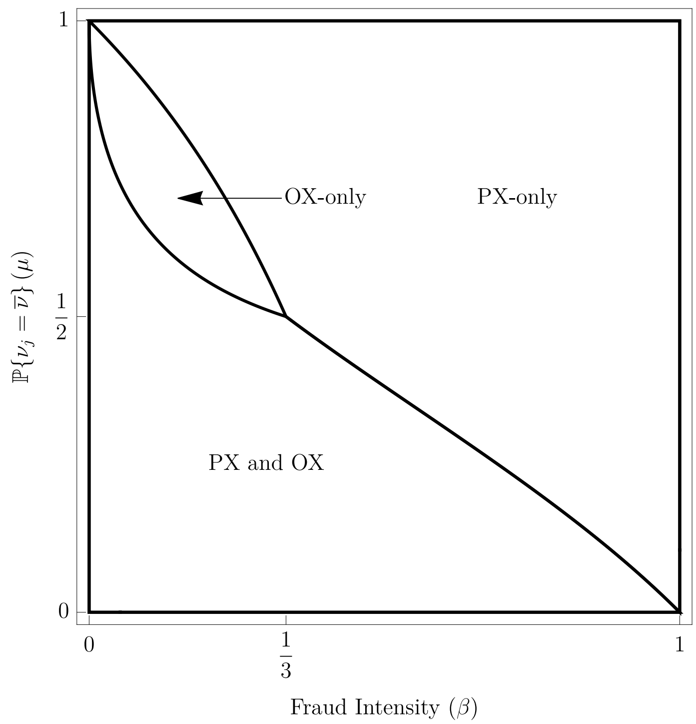

Figure 7: Publisher's Equilibrium Exchange Choices Under Endogenous OX Commission

through PX-OX (instead of OX-only as in the main model) if μ is small (see Figure 7).

Intuitively, small μ implies that the devaluation effect under the PX-OX configuration is mild (see the discussion of Figure 2b). In this case, the publisher has high valuation of selling through OX, which the OX extracts by raising the commission rate. On the other hand, if μ is large, the devaluation effect is severe such that the publisher has low valuation of selling through OX. In this case, the OX lowers its commission and induces the publisher to sell through OX-only. The following proposition summarizes these findings.

Proposition 6 (Exchange Configuration Under Endogenous Commission)**.** *Let* $\hat{\mu} = \{\mu \in \left[\frac{1}{2}, 1\right] : 1+\beta+(2\beta^2-5\beta-3)\mu+(-5\beta^2+5\beta+3)\mu^2+(2\beta^2-\beta-1)\mu^3 = 0\}$*. The publisher's equilibrium ad exchange choices are as follows:*

1. *if* $\hat{\mu} < \mu$ *and* $\beta \leq \frac{1-\mu}{2-\mu}$*, the publisher sells only though OX;*
2. *if* $\max\left[\frac{(1-\mu)^2}{1-\mu+\mu^2}, \frac{1-\mu}{2-\mu}\right] < \beta$*, the publisher sells only through PX;*
3. *otherwise, the publisher sells through the PX and OX.*

6 Conclusion

This paper studies how the emergence of private exchanges affects advertisers and publishers in the RTB market. We show that, while publishers can mitigate ad fraud by setting up private exchanges, doing so is not without any downsides. In particular, the presence of a private exchange can soften competition among advertisers by creating an information asymmetry between them. Our results provide important managerial implications for advertisers and publishers in the RTB industry.

When a publisher introduces a private exchange, advertisers who have access to the private exchange (i.e., connected advertisers) will be, at least partially, protected from buying fraudulent impressions. This implies that the impressions bought by advertisers who do not have access to the private exchange (i.e., unconnected advertisers) are now more likely to be fraudulent impressions. As such, the expected value of unconnected advertisers for the impressions in the open exchange declines with the introduction of a private exchange.

This information asymmetry hurts the publisher in two distinct ways. First, the unconnected advertisers' informational disadvantage lowers their impression valuation; this shrinks the total surplus the publisher can extract from the unconnected advertisers. Second, the unconnected advertisers' lower bids soften bidding competition such that even the connected advertisers reduce their bids. For the publisher, the positive impact of reduced ad fraud may or may not be sufficient to compensate for the negative, competition-softening effect induced by the information asymmetry. In particular, if the baseline fraud in the system is mild and the advertisers' average ad valuations are high, then the negative devaluation effect dominates such that the publisher is better off not introducing a private exchange, even if it is costless for the publisher to do so.

Finally, we study the open exchange's incentive to fight fraud in the form of filtering fake impressions. The open exchange faces a trade-off between lower transaction volume from forgoing sales of fraudulent impressions and higher transaction margin from alleviating advertisers' fraud concerns. If the publisher has strong incentive to sell through both the private and open exchanges, then the open exchange fights fraud to lure the connected advertisers, who have access to the private exchange, to transact through the open exchange.

While the focus of our paper is on ad fraud, the results can be applied more broadly to settings where advertisers' valuations are positively correlated through an unknown common-value component. For example, advertisers who sell digital cameras have higher valuations for impressions shown to customers who are in the market for digital cameras. Private exchanges can provide an informational advantage, similar to what we discuss in this paper, to connected advertisers by sharing first-party data such as patterns in the customers' browsing history. This informational advantage creates similar trade-offs to those discussed in this paper.

We acknowledge limitations of our model and suggest avenues for future research. First, we assume exogenous connections between advertisers and the publishers that set up private exchanges. In practice, the process of publishers inviting advertisers to join the private exchange, and whether advertisers accept or decline, may involve nuanced strategic decisions. It would be interesting to extend our current framework to analyze the endogenous private exchange formation process. Second, our paper restricts attention to the case where the open exchange combats ad fraud by identifying and filtering out fraudulent impressions. Another fruitful avenue for future research would be to consider imperfect identification of fraudulent impressions and alternative approaches to combating fraud, such as working with third-party ad verification providers or offering refunds to advertisers for fraudulent transactions (O'Reilly, 2017). Analyzing different forms of anti-fraud efforts and comparing their efficacy with respect to various welfare metrics could provide meaningful insights for regulators and policymakers.

References

Ada, S., Abou Nabout, N., and Feit, E. M. (2021). Context information can increase revenue in ad auctions: Evidence from a policy change. *Journal of Marketing Research (Forthcoming)*.

Amaldoss, W., Desai, P. S., and Shin, W. (2015). Keyword search advertising and first-page bid estimates: A strategic analysis. *Management Science*, 61(3):507–519.

Amaldoss, W., Jerath, K., and Sayedi, A. (2016). Keyword management costs and "broad match" in sponsored search advertising. *Marketing Science*, 35(2):259–274.

Benes, R. (2018). "How GDPR Is Moving Ad Budgets Toward Private Marketplaces," (accessed September 7, 2021). https://bit.ly/3l5lWcK.

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

Berman, R. (2018). Beyond the last touch: Attribution in online advertising. *Marketing Science*, 37(5):771–792.

Berman, R. and Katona, Z. (2013). The role of search engine optimization in search marketing. *Marketing Science*, 32(4):644–651.

Bruce, N. I., Murthi, B., and Rao, R. C. (2017). A dynamic model for digital advertising: The effects of creative format, message content, and targeting on engagement. *Journal of Marketing Research*, 54(2):202–218.

Cao, X. and Ke, T. T. (2019). Cooperative search advertising. *Marketing Science*, 38(1):44–67.

Cheq (2020). "The Economic Cost of Bad Actors on the Internet: Ad Fraud 2020," (accessed September 21, 2021). https://bit.ly/3u9Njq3.

Choi, H. and Mela, C. F. (2019). Monetizing online marketplaces. *Marketing Science*, 38(6):948–972.

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

Choi, W. J., Jerath, K., and Sarvary, M. (2022). Consumer privacy choices and (un)targeted advertising along the purchase journey. *Working paper, University of Maryland*.

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

Davies, J. (2017). "The FT warns advertisers after discovering high levels of domain spoofing," (accessed September 7, 2021). https://bit.ly/3njwsj8.

Davies, J. (2018). "Why publishers don't name and shame vendors over ad fraud," (accessed September 7, 2021). https://bit.ly/3nayK4a.

Desai, P. S. (2001). Quality segmentation in spatial markets: When does cannibalization affect product line design? *Marketing Science*, 20(3):265–283.

Desai, P. S., Shin, W., and Staelin, R. (2014). The company that you keep: When to buy a competitor's keyword. *Marketing Science*, 33(4):485–508.

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

Despotakis, S., Ravi, R., and Sayedi, A. (2021a). First-price auctions in online display advertising. *Journal of Marketing Research*.

Despotakis, S., Ravi, R., and Srinivasan, K. (2021b). The beneficial effects of ad blockers. *Management Science*, 67(4):2096–2125.

DriveScale (2017). "AppNexus Scales with DriveScale" (accessed June 13, 2022). https://bit.ly/3NRhjjo.

Dukes, A. J., Liu, Q., and Shuai, J. (2020). Skippable ads: Interactive advertising on digital media platforms. *Working paper, University of Southern California*.

Fou, A. (2019). "Hidden Costs in Digital Media Supply Path," (accessed September 7, 2021). https://bit.ly/3zcabGi.

Fou, A. (2020). "How Publishers Commit Ad Fraud," (accessed September 7, 2021). https://bit.ly/3hhBgBP.

Graham, K. (2020a). "What is a Private Marketplace for Publishers?" (accessed September 21, 2021). https://bit.ly/3wascor.

Graham, M. (2020b). "To show how easy it is for plagiarized news sites to get ad revenue, I made my own," (accessed September 7, 2021). https://cnb.cx/39copN3.

Gritckevich, A., Katona, Z., and Sarvary, M. (2021). Ad blocking. *Management Science (Forthcoming)*.

He, A. (2019). "The Size of the Ad Fraud Problem in Digital Marketing Is Varying," (accessed September 7, 2021). https://bit.ly/3jXJqkJ.

Hern, A. (2016). "Major sites including New York Times and BBC hit by 'ransomware' malvertising" (accessed September 21, 2021). https://bit.ly/3uA6V6Q.

Hoban, P. R. and Bucklin, R. E. (2015). Effects of internet display advertising in the purchase funnel: Model-based insights from a randomized field experiment. *Journal of Marketing Research*, 52(3):375–393.

Hsiao, S. (2020). "How our display buying platforms share revenue with publishers," (accessed September 7, 2021). https://bit.ly/2Vw8g1w.

Hsu, T. and Lutz, E. (2020). "More Than 1,000 Companies Boycotted Facebook. Did It Work?" (accessed September 21, 2021). https://nyti.ms/3CAKIIC.

Hu, Y., Shin, J., and Tang, Z. (2015). Incentive problems in performance-based online advertising pricing: cost per click vs. cost per action. *Management Science*, 62(7):2022–2038.

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

Jatain, V. (2021). “The Pros and Cons of Private Marketplace (PMP) Deals,” (accessed September 7, 2021). https://bit.ly/3la5B6o.

Jeffery, E. (2020). “What Is a Private Marketplace (PMP)?” (accessed September 7, 2021). https://bit.ly/3zRLxvP.

Jerath, K., Ke, T. T., and Long, F. (2018). The logic and management of ‘digital co-op’ in search advertising. *Working paper, Columbia University*.

Jerath, K., Ma, L., Park, Y.-H., and Srinivasan, K. (2011). A “position paradox” in sponsored search auctions. *Marketing Science*, 30(4):612–627.

Johnson, G. (2013). The impact of privacy policy on the auction market for online display advertising. *Working paper, Boston University*.

Kagel, J. H. and Levin, D. (2009). *Common value auctions and the winner’s curse*. Princeton University Press.

Katona, Z. and Sarvary, M. (2010). The race for sponsored links: Bidding patterns for search advertising. *Marketing Science*, 29(2):199–215.

Katona, Z. and Zhu, Y. (2017). Quality score that makes you invest. *Working paper, University of California, Berkeley*.

Kuksov, D., Prasad, A., and Zia, M. (2017). In-store advertising by competitors. *Marketing Science*, 36(3):402–425.

Lambrecht, A. and Tucker, C. (2013). When does retargeting work? information specificity in online advertising. *Journal of Marketing Research*, 50(5):561–576.

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

Long, F., Jerath, K., and Sarvary, M. (2021). Designing an online retail marketplace: Leveraging information from sponsored advertising. *Marketing Science (Forthcoming)*.

Lu, S., Zhu, Y., and Dukes, A. (2015). Position auctions with budget constraints: Implications for advertisers and publishers. *Marketing Science*, 34(6):897–905.

Moorthy, K. S. (1984). Market segmentation, self-selection, and product line design. *Marketing Science*, 3(4):288–307.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

O'Reilly, L. (2015). "It looks like Amazon shut down a popular ad product to stop companies like Google getting access to its data," (accessed September 7, 2021). https://bit.ly/2VrAryB.

O'Reilly, L. (2017). "Google Issuing Refunds to Advertisers Over Fake Traffic, Plans New Safeguard," (accessed September 7, 2021). https://on.wsj.com/3hm5bZO.

Rafieian, O. (2020). Revenue-optimal dynamic auctions for adaptive ad sequencing. *Working paper, Cornell University*.

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

Rowntree, L. (2019). "Ad Fraud: We Haven't Solved It Because We Don't Want To," (accessed September 7, 2021). https://bit.ly/3yVtMKP.

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

Sayedi, A., Jerath, K., and Baghaie, M. (2018). Exclusive placement in online advertising. *Marketing Science*, 37(6):970–986.

Sayedi, A., Jerath, K., and Srinivasan, K. (2014). Competitive poaching in sponsored search advertising and its strategic impact on traditional advertising. *Marketing Science*, 33(4):586–608.

Shin, W. (2015). Keyword search advertising and limited budgets. *Marketing Science*, 34(6):882–896.

Sluis, S. (2016). "The Year Header Bidding Went Mainstream," (accessed September 7, 2021). https://bit.ly/2VpPEAb.

Sluis, S. (2020). "Google Opens Its Black Box And Shares Fees Across DV360, Google Ads And Google Ad Manager" (accessed March 14, 2022). https://bit.ly/3CKA979.

Villas-Boas, J. M. (2004). Communication strategies and product line design. *Marketing Science*, 23(3):304–316.

Vrountas, T. (2020). "What are Private Marketplaces (PMP) and Why Should Advertisers Care?" (accessed September 7, 2021). https://bit.ly/3A0iwhI.

Wilbur, K. C. and Zhu, Y. (2009). Click fraud. *Marketing Science*, 28(2):293–308.

Zawadzinski, M. (2021). "Understanding RTB, Programmatic Direct and Private Marketplace," (accessed September 7, 2021). https://bit.ly/3kuJjwZ.

Zawadzinski, M. and Sweeney, M. (2020). "Identity in AdTech: Unravelling The ID Problem" (accessed June 1, 2022). https://bit.ly/3aeSIX8.

Zeithammer, R. (2019). Soft floors in auctions. *Management Science*, 65(9):4204–4221.

Zhu, Y. and Wilbur, K. C. (2011). Hybrid advertising auctions. *Marketing Science*, 30(2):249–273.

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

Appendix

A1 Proofs

A1.1 Proof of Lemma 1

Proof. If the publisher sells exclusively through the OX, then advertisers cannot distinguish between legitimate and fake ad impressions. Therefore, advertisers' expected impression valuation is $(1-\beta)\cdot 1+\beta\cdot 0=1-\beta$ if $\nu=1$ and 0 if $\nu=0$. It follows that the publisher's optimal reserve price is $R^{\text{OX-only}}=1-\beta$. This reserve price completely extracts the high-valuation advertisers' surplus, so their profits are 0. On the other hand, the publisher's profit is $R^{\text{OX-only}}$ if at least one of the advertisers draws high valuation, an event which occurs with probability $1-(1-\mu)^2=(2-\mu)\mu$. Therefore, the publisher's expected profit (before the advertisers' valuations and impression type are realized) is

$$\pi_P^{\text{OX-only}}=\left(1-(1-\mu)^2\right)R^{\text{OX-only}}=(2-\mu)\mu(1-\beta)^2 \tag{A1}$$

∎

A1.2 Proof of Lemma 2

Proof. If the publisher sells exclusively through PX, then the C-advertiser with $\nu_C=1$ values the ad impressions coming through PX at 1. Since the publisher does not sell its impression through OX, advertisers know in equilibrium that ad impressions coming through the OX are fraudulent. Therefore, no transactions occur in OX. The publisher's optimal reserve price for impressions sent exclusively to PX is raised as high as the high-valuation C-advertiser's impression valuation, which is 1. The publisher's expected profit (before the advertiser's valuation and impression type are realized) is thus

$$\pi_P^{\text{PX-only}} = \mathbb{P}\{\text{legitimate impression}\}\mathbb{P}\{\nu_C = 1\} \cdot 1 = (1-\beta)\mu \tag{A2}$$

■

A1.3 Proof of Lemma 3

Proof. Suppose the C-advertiser bids $b^{\text{PX}} \geq R^{\text{PX}}$ and $b^{\text{OX}} \geq R^{\text{OX}}$. Let b_U and π_C denote the U-advertiser's bid in OX and the C-advertiser's expected profit, respectively. We show that bidding b^{PX} in PX and b^{OX} in OX yields a weakly lower profit for the C-advertiser than bidding $\max\left[b^{\text{OX}}, b^{\text{PX}}\right]$ in PX only. Consider the C-advertiser's profit if it bids in both exchanges.

- If $\max\left[b^{\text{OX}}, b^{\text{PX}}\right] < b_U$, then the U-advertiser always wins the auction (both legitimate and fake impressions); therefore, $\pi_C = 0$.
- If $\max\left[b^{\text{OX}}, b^{\text{PX}}\right] > b_U$, then the publisher always chooses the C-advertiser's highest bid and allocates the impression to it, and the fraudster also allocates the impression to the highest bidder; therefore, $\pi_C = (1-\beta)\left(1 - \max\left[b^{\text{OX}}, b^{\text{PX}}\right]\right) - \beta \mathbb{I}_{\left\{b^{\text{OX}} > b_U\right\}} b^{\text{OX}}$, where $\mathbb{I}_{\{x\}}$ is an indicator function which equals 1 if x is true and 0 otherwise.
- If $\max\left[b^{\text{OX}}, b^{\text{PX}}\right] = b_U$, then $\pi_C = \alpha_1(1-\beta)\left(1 - \max\left[b^{\text{OX}}, b^{\text{PX}}\right]\right) - \alpha_2 \beta \mathbb{I}_{\left\{b^{\text{OX}} = b_U\right\}} b^{\text{OX}}$, where α_1 and α_2 are probabilities that the C-advertiser wins in the respective auctions under general tie-breaking rules.

In sum,

$$\pi_C = \begin{cases} 0 & \text{if } \max\left[b^{\text{OX}}, b^{\text{PX}}\right] < b_U, \\ (1-\beta)\left(1 - \max\left[b^{\text{OX}}, b^{\text{PX}}\right]\right) - \beta \mathbb{I}_{\left\{b^{\text{OX}} > b_U\right\}} b^{\text{OX}} & \text{if } \max\left[b^{\text{OX}}, b^{\text{PX}}\right] > b_U, \\ \alpha_1(1-\beta)\left(1 - \max\left[b^{\text{OX}}, b^{\text{PX}}\right]\right) - \alpha_2 \beta \mathbb{I}_{\left\{b^{\text{OX}} = b_U\right\}} b^{\text{OX}} & \text{if } \max\left[b^{\text{OX}}, b^{\text{PX}}\right] = b_U. \end{cases}$$

On the other hand, if the C-advertiser deviates to bidding $\max\left[b^{\mathrm{OX}}, b^{\mathrm{PX}}\right]$ in PX only, then its profit would be

$$\pi_C = (1-\beta)\cdot\begin{cases} 0 & \text{if } \max\left[b^{\mathrm{OX}}, b^{\mathrm{PX}}\right] < b_U, \\ 1-\max\left[b^{\mathrm{OX}}, b^{\mathrm{PX}}\right] & \text{if } \max\left[b^{\mathrm{OX}}, b^{\mathrm{PX}}\right] > b_U, \\ \alpha_1\left(1-\max\left[b^{\mathrm{OX}}, b^{\mathrm{PX}}\right]\right) & \text{if } \max\left[b^{\mathrm{OX}}, b^{\mathrm{PX}}\right] = b_U. \end{cases}$$

Therefore, the deviation strategy weakly dominates the original bidding strategy. ■

A1.4 Proof of Lemma 4

Proof. Advertisers who draw low valuations (i.e., $\nu_j = 0$) do not participate in the market; therefore, for ease of exposition, the advertisers discussed in the proof refer to those who draw high valuations (i.e., $\nu_j = 1$), unless specified otherwise.

For the regime in which the publisher sells through both PX and OX simultaneously to be an equilibrium, we need the following conditions:

1. (individual rationality) the reserve prices are no greater than the advertisers' valuations;
2. (incentive compatibility) the C-advertiser's profit from bidding in OX is no greater than that from bidding in PX; and
3. the C-advertiser's bid in PX is greater than the U-advertiser's bid in OX such that the C-advertiser wins.

The last two conditions are required to sustain the market for ad impressions in the PX.

The publisher sets the reserve prices as high as possible under the above constraints. We first determine the U-advertiser's valuation. To that end, note that the U-advertiser's expected profit from bidding the reserve price in OX equals

$$\pi_U(R) = (1-\mu)(1-\beta)(1-\beta)\left(1-R\right) + \beta\left(-R\right), \tag{A3}$$

where $1-\mu$ the probability that the C-advertiser's valuation is low (and therefore, the U-advertiser wins), $1-\beta$ the probability that the impression is legitimate, $1-R$ the U-advertiser's payoff if it wins the legitimate impression, β the probability that the impression is fake, and $-R$ the U-advertiser's payoff if it wins a fake impression (note that the U-advertiser always wins fake impressions because the C-advertiser only bids for legitimate impressions in the PX).

The U-advertiser's maximum willingness-to-pay (WTP) for an OX ad impression, and hence the U-advertiser's bid and the publisher's optimal OX reserve price, is

$$R^{\text{OX}} = \max\{R : \pi_U(R) \geq 0\} = \frac{(1-\beta)(1-\mu)}{1-(1-\beta)\mu}. \tag{A4}$$

The C-advertiser's maximum WTP for a PX ad impression is 1. However, the PX reserve price cannot be set as high as 1 due to the incentive compatibility constraint (9), which implies

$$R^{\text{PX}} \leq \frac{R^{\text{OX}}}{1-\beta}. \tag{A5}$$

Using the optimal OX reserve price (A4), we obtain the optimal reserve price in PX:

$$R^{\text{PX}} = \frac{R^{\text{OX}}}{1-\beta} = \frac{1-\mu}{1-(1-\beta)\mu}. \tag{A6}$$

We check the three necessary conditions above. Individual rationality is satisfied because $\pi_U\left(R^{\text{OX}}\right) \geq 0$ due to (A4), and $\pi_C\left(R^{\text{PX}}\right) \geq 0$ due to (A5); incentive compatibility holds by construction of R^{PX}; and finally, for the publisher's impression, the C-advertiser's bid, which equals (A6) is higher than the U-advertiser's, which equals (A4).

The C-advertiser's expected profit (before its valuation and impression type are realized) is

$$\pi_C^{\text{PX-OX}} = \mu(1-\beta)\left(1-R^{\text{PX}}\right) = \frac{(1-\beta)\beta\mu^2}{1-(1-\beta)\mu},$$

and the publisher's expected profit is

$$\pi_P^{\text{PX-OX}} = (1-\beta)\left(\mu R^{\text{PX}} + \mu(1-\mu)R^{\text{OX}}\right) = \frac{(2-\mu-\beta(1-\mu))(1-\beta)(1-\mu)\mu}{1-(1-\beta)\mu}. \tag{A7}$$

■

A1.5 Proof of Proposition 1

Proof. We show that the U-advertiser's bid in the PX-OX regime, b_U^{OX}, is lower than that under the OX-only regime, $1-\beta$:

$$b_U^{\text{OX}} \leq 1-\beta \Leftrightarrow \frac{(1-\beta)(1-\mu)}{1-(1-\beta)\mu} \leq 1-\beta \Leftrightarrow 1-\mu \leq 1-(1-\beta)\mu,$$

which is true for all $\beta \in [0,1]$ and $\mu \in [0,1]$.

The C-advertiser's bid in the PX-OX regime is lower than that under the OX-only regime if and only if

$$\frac{1-\mu}{1-(1-\beta)\mu} \leq 1-\beta \Leftrightarrow \beta-(2-\beta)\beta\mu \leq 0 \Leftrightarrow \beta \leq 2-\frac{1}{\mu},$$

which is possible only if $\mu > \frac{1}{2}$. ■

A1.6 Proof of Proposition 2

Proof. From (11), we obtain

$$\frac{d}{d\beta}\pi_C^{\text{PX-OX}} = \frac{d}{d\beta}\frac{(1-\beta)\beta\mu^2}{1-(1-\beta)\mu} = -\frac{\mu^2\left((\beta-1)^2\mu+2\beta-1\right)}{((\beta-1)\mu+1)^2},$$

from which it follows that

$$\frac{d}{d\beta}\pi_C^{\text{PX-OX}} > 0 \Leftrightarrow \beta < \frac{\sqrt{1-\mu}-(1-\mu)}{\mu},$$
$$\frac{d}{d\beta}\pi_C^{\text{PX-OX}} < 0 \Leftrightarrow \beta > \frac{\sqrt{1-\mu}-(1-\mu)}{\mu}.$$

■

A1.7 Proof of Proposition 3

Proof. The publisher compares the subgame optimal profits in the OX-only, PX-only, and PX-OX regime, and chooses the regime that yields the highest profit. From (A1), the OX-only regime yields $(2-\mu)\mu(1-\beta)^2$; from (A2), the PX-only regime yields $(1-\beta)\mu$; and from (A7), the PX-OX regime yields $\frac{(1-\beta)(2-\mu-\beta(1-\mu))(1-\mu)\mu}{1-(1-\beta)\mu}$.

Based on these profit expressions, we derive the conditions under which each of the subgame optimal profits is the maximum of the three:

1. (OX-only)

$$\pi_P^{\text{OX-only}} \geq \pi_P^{\text{PX-only}} \Leftrightarrow 1-\mu-\beta(2-\mu) \geq 0 \Leftrightarrow \beta \leq \frac{1-\mu}{2-\mu}, \tag{A8}$$

and

$$\pi_P^{\text{OX-only}} \geq \pi_P^{\text{PX-OX}} \Leftrightarrow -(1-\beta)\mu^2+(3-2\beta)\mu-1 \geq 0,$$

but LHS (i.e., $-(1-\beta)\mu^2+(3-2\beta)\mu-1$) is concave in μ, LHS$|_{\mu=0}= -1$ and LHS$|_{\mu=1}= 1-\beta$; therefore, the inequality simplifies to μ being greater than the root of LHS; therefore,

$$\mu \geq \frac{3-2\beta-\sqrt{4\beta^2-8\beta+5}}{2(1-\beta)}. \tag{A9}$$

2. (PX-only)

$$\pi_P^{\text{PX-only}} \geq \pi_P^{\text{PX-OX}} \Leftrightarrow -(1-\mu)^2 + \beta(1-\mu+\mu^2) \geq 0 \Leftrightarrow \beta \geq \frac{(1-\mu)^2}{1-\mu+\mu^2}, \tag{A10}$$

and

$$\pi_P^{\text{PX-only}} \geq \pi_P^{\text{OX-only}} \Leftrightarrow \beta \geq \frac{1-\mu}{2-\mu},$$

from the complement of (A8).

3. (PX-OX)

$$\pi_P^{\text{PX-OX}} \geq \pi_P^{\text{OX-only}} \Leftrightarrow \mu \leq \frac{3-\sqrt{4\beta^2-8\beta+5}-2\beta}{2(1-\beta)},$$

from the complement of (A9), and

$$\pi_P^{\text{PX-OX}} \geq \pi_P^{\text{PX-only}} \Leftrightarrow \beta \leq \frac{(1-\mu)^2}{1-\mu+\mu^2},$$

from the complement of (A10).

■

A1.8 Proof of Lemma 5

Proof. Consider OX's unilateral deviation to filtering γ proportion of fraudulent impressions, given the publisher's reserve price. Then, there are

$$\frac{1-\beta}{1-\beta+\beta(1-\gamma)} \tag{A11}$$

share of legitimate ad impressions coming through OX. Therefore, the advertisers' valuation is (A11). Under this unilateral deviation, the reserve price is fixed at $1-\beta$, which is less than the post-filter valuation (A11). Moreover, the auction format is first-price, so there is no pure strategy equilibrium in the advertisers' bids. Since the publisher's reserve price in OX is fixed at $1-\beta$ (recall that we are considering the OX's unilateral deviation), the high-valuation advertisers mix on the interval $[1-\beta, \overline{b}]$ according to distribution H where $\overline{b}$ and H satisfy the following indifference conditions:

$$\begin{aligned}
\pi(1-\beta) &= (1-\beta+\beta(1-\gamma))\,(1-\mu)\left(\frac{1-\beta}{1-\beta+\beta(1-\gamma)}-(1-\beta)\right)\\
&= (1-\beta+\beta(1-\gamma))(1-\mu)(1-\beta)\frac{\beta\gamma}{1-\beta\gamma}\\
= \pi(b) &= (1-\beta+\beta(1-\gamma))\,(1-\mu+\mu H(b))\left(\frac{1-\beta}{1-\beta+\beta(1-\gamma)}-b\right) \text{ for all } b\in(1-\beta,\overline{b})\\
= \pi(\overline{b}) &= (1-\beta+\beta(1-\gamma))\left(\frac{1-\beta}{1-\beta+\beta(1-\gamma)}-\overline{b}\right),
\end{aligned}$$

where the arguments of $\pi(\cdot)$ denote the advertiser's bid. Thus, we obtain

$$\overline{b} = \frac{1-\beta}{1-\beta\gamma}\,(1-(1-\mu)\beta\gamma) \text{ and } H(b) = \begin{cases} 0 & \text{if } b < 1-\beta, \\ \frac{(1-\mu)(b+\beta-1)(1-\beta\gamma)}{b\beta\mu\gamma+\mu(-b-\beta+1)} & \text{if } 1-\beta \le b < \overline{b}, \\ 1 & \text{if } \overline{b} \le b. \end{cases}$$

the OX's expected profit in this mixed strategy equilibrium under arbitrarily small commission α is

$$\alpha(1-\beta+\beta(1-\gamma))\left((1-\mu)^2\cdot 0 + 2\mu(1-\mu)\int_{1-\beta}^{\overline{b}} b_j\, dH(b_j) + \mu^2 \iint_{1-\beta}^{\overline{b}} \max\left[b_C, b_U\right]\, dH\left(b_C\right) dH\left(b_U\right)\right),$$

which simplifies to

$$\pi_{\text{OX}} = \alpha(1-\beta)\mu(2-\mu-2\beta\gamma(1-\mu)). \tag{A12}$$

Since $\frac{d}{d\gamma}\pi_{\text{OX}} = -2(1-\beta)\beta(1-\mu)\mu < 0$, we obtain $\gamma^* = 0$. ■

A1.9 Proof of Proposition 4

Proof. Consider the OX's anti-fraud incentive for each of the different regimes. First, in the PX-only regime, OX's profit is always zero, so OX is indifferent between any γ. Second, OX does not fight fraud in the OX-only regime (see Lemma 5). Finally, consider the PX-OX regime. Claim 2 from Section WA5 in the Web Appendix proves that if the pre-filter regime is PX-OX regime, then the OX always sets $\gamma^* > 0$.

For the second part of the proposition, observe that the C-advertiser's profit is positive only under the PX-OX equilibrium. Following Proposition 3 and Claim 2, if $\mu < \frac{3-2\beta-\sqrt{4\beta^2-8\beta+5}}{2(1-\beta)}$ and $\beta \leq \frac{(1-\mu)^2}{1-\mu+\mu^2}$, then the OX filters fake impressions such that the resultant equilibrium is the OX-only regime. In this case, OX's anti-fraud efforts reduce the C-advertiser's profit from $\pi_C^{\text{PX-OX}} = \frac{\beta\mu}{1-(1-\beta)\mu}$ to $\pi_C^{\text{OX-only}} = 0$. ■

A1.10 Proof of Lemma 6

Proof. The result follows immediately from the fact that, as described in the text, the C-advertiser's incentive compatibility constraint under the simultaneous PX-OX regime coincides with that under the sequential setting. ■

We should note that the result in Lemma 6 is different from the findings of Despotakis et al. (2021a) where the authors show that simultaneous selling (under header bidding) leads to a higher equilibrium revenue for the publisher than sequential selling (under waterfalling). This is because in our model, all advertisers have access to the second exchange (OX) in the sequence, whereas in Despotakis et al. (2021a) the set of advertisers who bid through the first exchange and the second exchange are mutually exclusive. As such, in Despotakis et al. (2021a), if the publisher sets a high reserve price in the first exchange, it would lose bids from advertisers that bid only through the first exchange. On the other hand, in our model, even if the publisher sets a high reserve price in the first exchange (PX), it can still sell the impression to all advertisers because they would all bid in the second exchange (OX).

A1.11 Proof of Proposition 5

Proof. From the main text, we have

$$
\begin{aligned}
\pi_P^{\text{PX-only}} &= \mu, \\
\pi_P^{\text{OX-only}} &= (2-\mu)\mu(1-\beta)^2, \\
\text{and } \pi_P^{\text{PX-OX}} &= \frac{\mu\left(\beta^2\left(\mu^2-\mu+1\right)-2\beta(1-\mu)^2+\mu^2-3\mu+2\right)}{1-(1-\beta)\mu}.
\end{aligned}
$$

Comparing the three profits, we derive the conditions under which each regime is optimal for the publisher:

1. (OX-only)

$$\pi_P^{\text{OX-only}} \geq \pi_P^{\text{PX-only}} \Leftrightarrow \beta \leq 1-\frac{\sqrt{2-\mu}}{2-\mu}$$

   from (A8) and

$$\pi_P^{\text{OX-only}} \geq \pi_P^{\text{PX-OX}} \Leftrightarrow \mu > \frac{\beta^2-3\beta-\sqrt{(2-\beta)(1-\beta)^3}+2}{(1-\beta)^2}. \tag{A13}$$

2. (PX-only)

$$\pi_P^{\text{PX-only}} \geq \pi_P^{\text{OX-only}} \Leftrightarrow 1-\frac{\sqrt{2-\mu}}{2-\mu} < \beta \tag{A14}$$

   from the complement of (A8) and

$$\pi_P^{\text{PX-only}} \geq \pi_P^{\text{PX-OX}} \Leftrightarrow \frac{(1-\mu)^2}{\mu^2-\mu+1} < \beta. \tag{A15}$$

3. (PX and OX)

$$\pi_P^{\text{PX-OX}} \geq \pi_P^{\text{OX-only}} \Leftrightarrow \mu \leq \frac{\beta^2-3\beta-\sqrt{(2-\beta)(1-\beta)^3}+2}{(1-\beta)^2} \tag{A16}$$

from the complement of (A13), and

$$\pi_P^{\text{PX-OX}} \leq \pi_P^{\text{PX-only}} \Leftrightarrow \beta \leq \frac{(1-\mu)^2}{\mu^2-\mu+1} \tag{A17}$$

from the complement of (A15).

Finally, algebraic manipulations yield the simplified expressions in the proposition. ■

A1.12 Proof of Lemma 7

Proof. Conditional on an exchange configuration (i.e., either OX-only, PX-only, or PX-OX), the OX's profit monotonically increases in α. Therefore, under endogenous α, the OX's optimal α must bind the publisher's incentive compatibility constraint. To derive these constraints, note that the publisher's profits under the three possible exchange configurations are

$$\begin{aligned}
\pi_P^{\text{OX-only}} &= (1-\alpha)(1-(1-\mu)^2)(1-\beta) \\
&= (1-\alpha)(2-\mu)\mu(1-\beta)^2 \\
\pi_P^{\text{PX-only}} &= \mu, \\
\pi_P^{\text{PX-OX}} &= \mu\frac{1-\mu}{1-(1-\beta)\mu} + (1-\alpha)\mu(1-\mu)\frac{(1-\beta)(1-\mu)}{1-(1-\beta)\mu} \\
&= \frac{(1-\mu)\mu(2-\alpha(1-\beta)(1-\mu)-(1-\beta)\mu-\beta)}{1-(1-\beta)\mu}
\end{aligned}$$

It follows immediately that $\frac{d}{d\alpha}\pi_P^{\text{OX-only}} < 0$, $\frac{d}{d\alpha}\pi_P^{\text{PX-OX}} < 0$, and $\frac{d}{d\alpha}\pi_P^{\text{PX-only}} = 0$. We also have $\frac{d}{d\alpha}\pi_P^{\text{OX-only}} < \frac{d}{d\alpha}\pi_P^{\text{PX-OX}}$ because:

$$\frac{d}{d\alpha}\pi_P^{\text{OX-only}} < \frac{d}{d\alpha}\pi_P^{\text{PX-OX}} \Leftrightarrow \frac{(1-\beta)\mu(\beta(2-\mu)\mu+1-\mu)}{1-(1-\beta)\mu} > 0,$$

which is true for all $\mu, \beta \in (0,1)$.

Therefore, if the equilibrium configuration is PX-only with α arbitrarily close to 0 (i.e., in the main model), then raising α would only reduce the publisher's profits in the other two regimes such that PX-only sustains in the endogenous commission equilibrium. Similarly, if the equilibrium configuration is PX-OX with α arbitrarily close to 0, then raising α will only make the PX-OX configuration more profitable relative to OX-only (due to $\frac{d}{d\alpha}\pi_P^{\text{OX-only}} < \frac{d}{d\alpha}\pi_P^{\text{PX-OX}}$). This implies that by raising α, the OX may induce the publisher to switch to PX-only, but this would reduce the OX's profit to 0. Therefore, under endogenous commission, PX-OX regime will sustain. ■

A1.13 Proof of Proposition 6

Proof. Due to Lemma 7, it suffices to consider the parametric region for which OX-only is the equilibrium outcome in the main model. Since the OX's profit is 0 under the PX-only regime, OX will only consider inducing either OX-only or PX-OX. Next, consider the following two cases:

- Suppose $\mu \leq 1 - \sqrt{\frac{\beta}{1+\beta}}$.
  - $\pi_P^{\text{OX}} \geq \max\left[\pi_P^{\text{PX-OX}}, \pi_P^{\text{PX-only}}\right] \Leftrightarrow 0 < \alpha \leq 1 - \frac{1-\mu}{(1-\beta)(\beta(2-\mu)\mu-\mu+1)}$
  - $\pi_P^{\text{PX-OX}} \geq \max\left[\pi_P^{\text{OX-only}}, \pi_P^{\text{PX-only}}\right] \Leftrightarrow 1 - \frac{1-\mu}{(1-\beta)(\beta(2-\mu)\mu-\mu+1)} < \alpha \leq \frac{\beta\mu(1-\mu)-\beta+(1-\mu)^2}{(1-\beta)(1-\mu)^2}$.

  Therefore, OX compares its optimal profit under each of these regimes:

$$\pi_{\text{OX}}^{\text{OX-only}}\left(\alpha = 1 - \frac{1-\mu}{(1-\beta)(\beta(2-\mu)\mu-\mu+1)}\right) = \frac{\beta(2-\mu)\mu(\mu(-\beta(2-\mu)-\mu+3)-1)}{\beta(2-\mu)\mu-\mu+1} \tag{A18}$$

$$\pi_{\text{OX}}^{\text{PX-OX}}\left(\alpha = \frac{\beta\mu(1-\mu)-\beta+(1-\mu)^2}{(1-\beta)(1-\mu)^2}\right) = \mu\left(\beta\left(\mu - \frac{1}{1-\mu}\right) - \mu + 1\right) \tag{A19}$$

  Since (A18) $\geq$ (A19) $\Leftrightarrow \mu \geq \hat{\mu}$, where $\hat{\mu}$ is the root of $1 + \beta + (2\beta^2 - 5\beta - 3)\mu + (-5\beta^2 + 5\beta + 3)\mu^2 + (2\beta^2 - \beta - 1)\mu^3$ in the interval $[1/2, 1]$, we obtain that the OX sets $\alpha = 1 - \frac{1-\mu}{(1-\beta)(\beta(2-\mu)\mu-\mu+1)}$ and induces OX-only if $\mu \geq \hat{\mu}$ and sets $\alpha = \frac{\beta\mu(1-\mu)-\beta+(1-\mu)^2}{(1-\beta)(1-\mu)^2}$ and induces PX-OX otherwise.

- Suppose $1 - \sqrt{\frac{\beta}{1+\beta}} < \mu$. the OX cannot induce PX-OX for any α; therefore, the OX's optimal commission is to set α such that the publisher is indifferent between OX-only and PX-only, and induces OX-only.

Taken together, within the OX-only regime from the main model, if $\mu \leq \hat{\mu}$, the OX sets commission such that the equilibrium regime changes to PX-OX regime. ■

A2 Additional Analyses

A2.1 Selective Bidding in OX

In the main model, we assumed that even though the C-advertiser has access to PX, it was equally informationally disadvantaged as the U-advertiser when bidding in the OX. In this section, we relax this assumption and examine the opposite case where the C-advertiser can leverage information from the PX to perfectly identify and bid selectively for legitimate impressions in the OX.[27]

Suppose the C-advertiser can gain informational advantage from its access to PX, allowing the C-advertiser to bid selectively for legitimate impressions in the OX. An important implication of this alternative information structure is that it decreases the publisher's capacity to price-discriminate among the advertisers by selling through both exchanges. The intuition is that since the C-advertiser can "free-ride" on PX's information to buy legitimate impressions in the OX, the C-advertiser's incentive to buy in the PX decreases relative to the main model. Thus, the devaluation effect that emerges under the PX-OX regime dominates the potential gains from price-discrimination. Overall, if the C-advertiser gains additional informational advantage in the form of identifying and bidding selectively for legitimate impressions in the OX, the publisher does not sell through both exchanges simultaneously. The following proposition summarizes this result.

Lemma 8. *If the C-advertiser can identify legitimate impressions in the OX and bid selec-*

[27] As discussed in Footnote 13, technical capability for cross-exchange impression matching remains largely limited due to (i) the use of different cookies and identifiers across exchanges, and (ii) privacy regulations that empower consumers to limit data collection and data sharing between third-parties. For instance, successful cookie-match rate between different AdTech firms lies in the range of 40%–60%, implying that roughly half of the impressions from one domain are incorrectly identified in another (Zawadzinski and Sweeney, 2020).

tively for them, then in equilibrium, the publisher does not sell through both the PX and OX simultaneously.

Proof. See Section WA1 of the Web Appendix. ■

As the C-advertiser's informational advantage from the PX can be transferred to the OX, the information asymmetry between the advertisers becomes more pronounced than in the main model. This exacerbates the devaluation effect which arises under the PX-OX regime such that the publisher does not sell through both exchanges. In total, the publisher sells only through the OX if fraud is sufficiently mild, such that the benefit of a large transaction volume in the OX outweighs the benefit of selling impressions exclusively to the C-advertiser in the PX at a slightly higher margin. If fraud is severe such that advertisers have low ad valuations in the OX, the publisher sells exclusively through the PX. The following proposition summarizes the publisher's exchange choices, depicted in Figure A1, when the C-advertiser can bid selectively in the OX.

Proposition 7 (Exchange Configuration Under Selective Bidding)**.** *The publisher's equilibrium ad exchange choices are as follows:*

1. *if $\beta \leq \frac{1-\mu}{2-\mu}$, the publisher sells only through the OX;*
2. *otherwise, the publisher sells only through the PX.*

Proof. See Section WA2 of the Web Appendix. ■

A2.2 Equilibrium Fraud Filter

In the main model, we find that the OX fights fraud if and only if the publisher induces the PX-OX regime in the absence of the OX's anti-fraud efforts (see Proposition 4). In this section, we numerically analyze the OX's equilibrium fraud filter level and discuss how it varies with

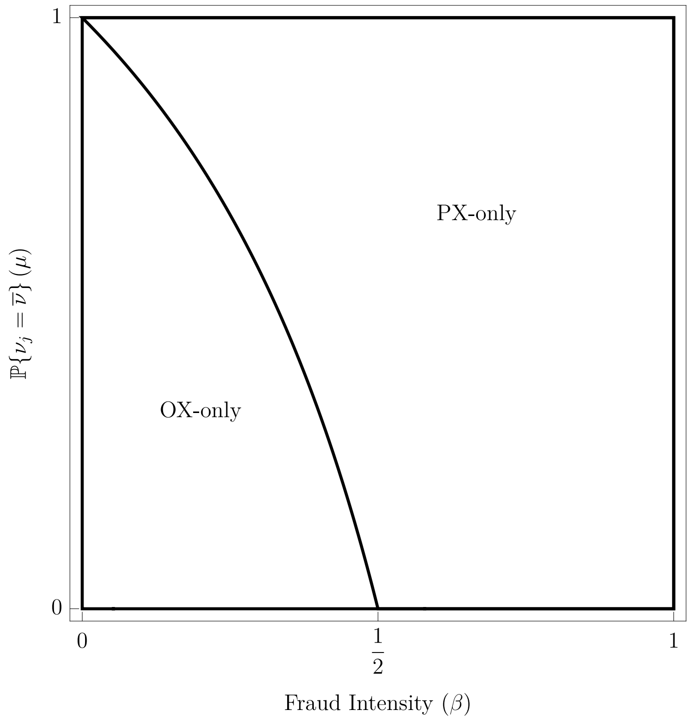

Figure A1: Publisher's Equilibrium Exchange Choices Under Selective OX Bidding

the baseline fraud level β. To that end, we extend the main model by assuming the OX can filter a fraudulent impression at marginal cost $k > 0$.[28]

One may intuit that as baseline fraud increases, the OX fights fraud proportionately more to offset the negative effect of fraud; i.e., the equilibrium filter level γ^* increases in β. Interestingly, we find that γ^* is non-monotonic in β. For small μ, γ^* first increases then decreases in β (see Figure A2).

If baseline fraud is mild, the OX fights fraud more aggressively as fraud intensity increases. However, if baseline fraud is severe, the OX fights fraud less aggressively as β increases. This is because the advertisers' valuations depend not on the *absolute* intensity of fraud in the system, but on the *ratio* of legitimate vs. fraudulent impressions. The marginal decrease in the ratio of fraudulent impressions when the OX filters a unit of fraudulent impression decreases in β; i.e., $\frac{\partial^2}{\partial\beta\partial\gamma}\frac{\beta-\gamma}{\beta-\gamma+1-\beta} = \frac{1}{(1-\gamma)^2} > 0$. In other words, the OX's marginal return from reducing fraud *decreases* in β such that if β is large, the OX filters less fraudulent impressions as β increases.

[28] Moreover, given the cost of filtering fraud, we re-define γ as the absolute number of fraudulent impressions the OX filters, instead of the fraction impressions filtered in the main model.

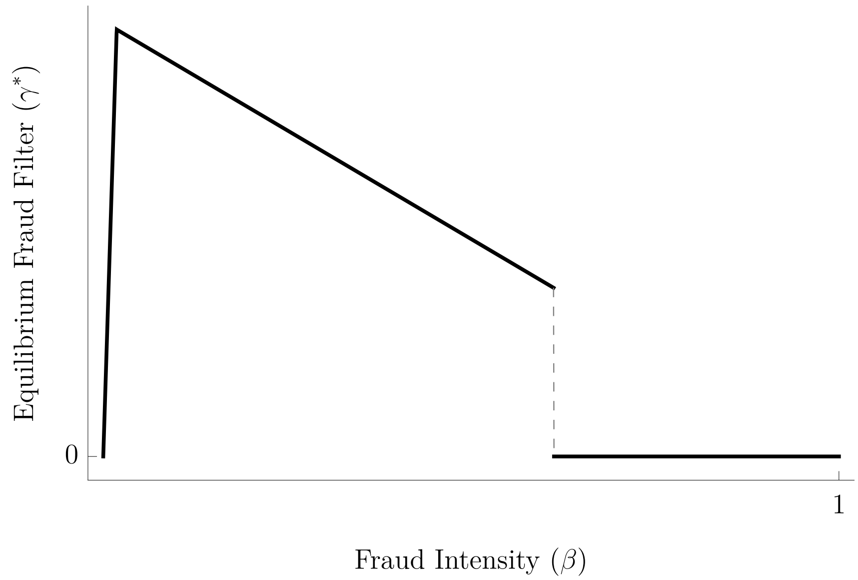

Figure A2: Equilibrium Fraud Filter; $\mu = 0.25, k = 8$

Moreover, if β is sufficiently large such that the publisher sells exclusively through the PX (see Proposition 3), the separation of legitimate and fraudulent impressions causes the OX market to collapse such that reducing fraud creates zero returns. In this case, $\gamma^* = 0$.

Building on the policy implications from the main model (see discussion below Proposition 4), this finding provides additional insights for regulators. For example, it suggests that while the market will likely self-regulate at intermediate ranges of fraud due to competition between the exchanges, regulatory intervention would be crucial at extreme levels of fraud (i.e., either very small or very large β). In particular, if fraud is sufficiently severe, the publisher may sell exclusively through the PX such that the OX market collapses; this in turn would undermine market efficiency.

A2.3 Advertiser Heterogeneity

A simplifying assumption from the main model was advertiser homogeneity. In practice, however, publishers that run the PX may invite advertisers selectively, e.g., based on their total ad spend.[29] Moreover, advertisers who are connected to the PX may benefit from greater opportu-

[29]For example, see https://bit.ly/3QvCdXk (accessed January 24, 2022).

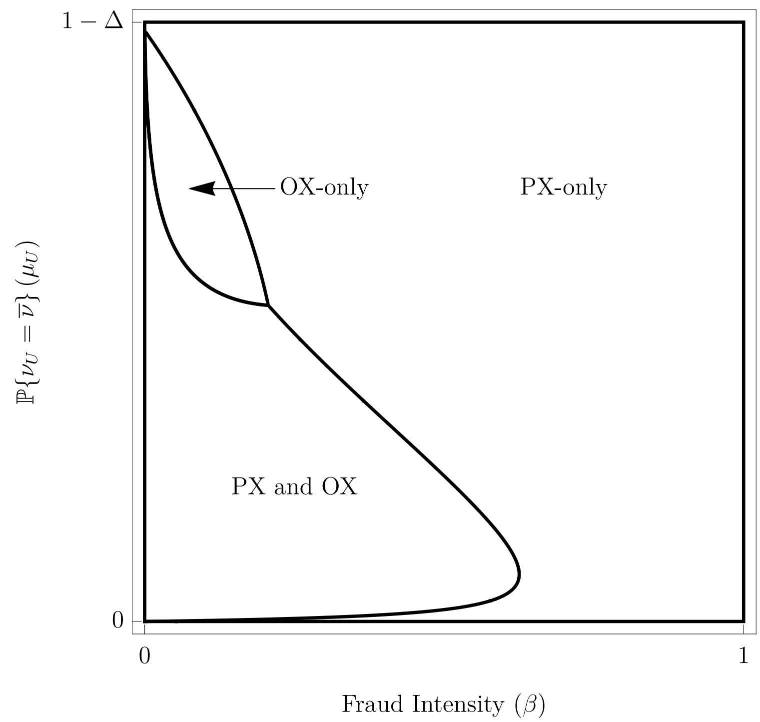

Figure A3: Publisher's Equilibrium Exchange Choices Under Advertiser Heterogeneity; $\Delta = 0.1$

nities for premium placement as well as from higher quality information for targeting purposes (e.g., ad context and consumer profiles). Such factors may create asymmetries in advertisers' valuations depending on their connection to the PX. In this section, we examine the extent to which the qualitative insights from the main model carry over under advertiser heterogeneity. We assume that the C-advertiser's valuation distribution first-order stochastically dominates the U-advertiser's. Specifically, let $\mu_C = \mu_U + \Delta$, where $\mu_j = \mathbb{P}\{\nu_j = \overline{\nu}\} = 1 - \mathbb{P}\{\nu_j = \underline{\nu}\}$ denotes the j-advertiser's probability of realizing a high valuation, and $\Delta \in [0, 1 - \mu_U]$ proxies the degree of asymmetry in the advertisers' valuation distributions. Note that the extension model reduces to the main model if $\Delta = 0$.

We conduct numerical analyses over a wide range of parameter values and demonstrate that while the boundaries of publisher's equilibrium exchange configurations shift with Δ, the qualitative insights from the main model are largely preserved. As illustrated in Figure A3, the parametric regions associated with the three regimes mirror that from the main model (compare with Figure 4).

An interesting departure from the main model occurs for small μ_U. If $\mu_U \approx 0$, then advertiser heterogeneity induces the publisher to switch its exchange configuration from selling through both exchanges to selling through PX-only (see Figure A3). The intuition is as follows. Without advertiser heterogeneity, small μ alleviates the devaluation effect such that the publisher sold through the PX and OX, thereby capitalizing on the C-advertiser's high bid in the PX and additional demand from the U-advertiser in the OX. In contrast, with advertiser heterogeneity, small μ_U exacerbates cannibalization (similar to the case of large β in the main model) such that if the publisher sells through both exchanges, the OX erodes the extractable surplus from the PX. Therefore, if μ_U is small, the publisher sells through PX-only and effectively skims the C-advertiser's high valuation.

A2.4 Competition Dispersion Effect

In this section, we discuss another potential downside of selling through both PX and OX that has been widely documented in the online advertising literature: the *competition dispersion* effect, also known as the market thinning effect (e.g., Amaldoss et al., 2016; Bergemann and Bonatti, 2011; Levin and Milgrom, 2010; Rafieian and Yoganarasimhan, 2021; Sayedi, 2018). The intuition for the competition dispersion effect is as follows. If a publisher offers an impression through multiple channels, advertisers will be divided into multiple groups, each bidding for the impression through one channel. Advertisers within each group compete with one another for the impression; however, competition among advertisers across different groups may be weakened. Overall, competition dispersion may lower the publisher's profit, as the following example demonstrates.

Example. Suppose two advertisers with independent and identical valuation distribution $U[0, 1]$ compete in a second-price auction with (the optimal) reserve price $1/2$. The publisher's expected revenue from this auction is $5/12$.[30] On the other hand, if the two advertisers bid

[30]With probability $1/2$, only one advertiser beats the reserve price, in which case, the revenue would be $1/2$. With probability $1/4$, both advertisers beat the reserve price, in which case, the revenue would be $2/3$. Therefore, the total expected revenue is $1/2 \cdot 1/2 + 1/4 \cdot 2/3 = 5/12$.

in two separate second-price auctions with the same reserve price $1/2$, the publisher's revenue would be $3/8$, which is less than $5/12$.[31]

The reason the publisher's revenue under separate auctions is lower than that under a single auction is the following. When advertisers compete in the same auction, in situations where more than one advertisers outbid the reserve price, the bid of one advertiser can be used as the price for the other advertiser. In contrast, when the advertisers are separated into two auctions, the bid of one auction cannot be used as the price for the other auction.

In this subsection, we highlight that the exchanges' recent transition from second- to first-price auctions has eliminated the competition dispersion effect, which industry experts have documented as a potential drawback of introducing private exchanges (e.g., Jatain, 2021; Jeffery, 2020). Under second-price auctions, the introduction of PX would have hurt publishers due to the competition dispersion effect (see example above); however, under first-price auctions, the negative impact of competition dispersion disappears. Intuitively, this is because when there are two exchanges with first-price auctions, advertisers in each exchange take competitors in other exchanges into account when submitting their bids. Put differently, the publisher does not forego PX for fear of competition dispersion. Instead, the publisher's exchange choice is driven by its effect on the advertisers' *information structure*. Had the publisher's exchange choices preserved the advertisers' *ex ante* information symmetry, then the publisher's exchange choice would have no material impact. That is, under information symmetry, selling through two separate auctions with one advertiser participating in each and selling through a single, integrated auction with both advertisers yield the same optimal revenue. We state this finding in the following lemma.

Lemma 9. *If the advertisers have symmetric information, then the publisher's optimal revenue under two separate first-price auctions with two reserve prices (one in each auction) is the same as its optimal revenue under a single, integrated first-price auction with one reserve price.*

[31]It can be shown that even if the publisher optimizes the reserve prices under two separate second-price auctions, the qualitative insight holds.

Proof. See Section WA3 of the Web Appendix. ■

Note that the result of Lemma 9 is specific to first-price auctions; in particular, as the example above demonstrates, it does not apply to second-price auctions. The intuition is as follows. In first-price auctions, bidders shade their bids according to the intensity of bidding competition; the higher the competitors' bids, the less the bidder shades in equilibrium. Now, even if advertisers are divided into multiple groups, they know that to win the impression, they must outbid not only the advertisers within their own exchange, but also those in other exchanges. As such, when shading their bids, they behave as if they are directly competing with every other advertiser in every other exchange.

In sum, our results shed light on a novel effect of introducing PX on the display ad market. While the publisher's selling through PX helps mitigate fraud for some advertisers, it also creates information asymmetry between advertisers that softens bidding competition and lowers the publisher's profit. An important managerial implication is that publishers should be cognizant of the distortions in information structures created by the PX. In particular, publishers considering selling through PX should exercise caution when the devaluation effect is most pronounced; i.e., the advertisers' average valuation is high and baseline fraud is mild.

Web Appendix

WA1 Proof of Lemma 8

Proof. It suffices to show that the publisher's maximum profit from inducing the PX-OX regime never exceeds that from selling through OX-only. First, note that since the U-advertiser cannot distinguish between legitimate and fake impressions in the OX, its valuation of an OX impression does not exceed $1 - \beta$. Therefore, the U-advertiser does not bid more than $1 - \beta$.

Second, to induce the U-advertiser's participation in the OX (a necessary condition for the PX-OX regime), the publisher must set the OX reserve price no greater than the U-advertiser's maximum valuation $1 - \beta$.

Third, to induce the C-advertiser's participation in the PX (a necessary condition for the PX-OX regime), the publisher cannot set the reserve price in PX greater than the reserve price in OX. Otherwise, the C-advertiser would identify and bid selectively for the legitimate impression in the OX where the reserve price is lower, such that the PX market collapses.

Therefore, we obtain $R^{\text{PX}} \leq R^{\text{OX}} \leq 1 - \beta$. Combined with the fact that the U-advertiser does not bid more than $1 - \beta$, the C-advertiser, too, does not bid more than $1 - \beta$.

Overall, in any subgame equilibrium in which the publisher induces the PX-OX regime, the publisher's expected profit is bounded from above by

$$(2 - \mu)\mu(1 - \beta)^2, \tag{WA1}$$

where $(2 - \mu)\mu$ is the probability that at least one of the two advertisers have a high valuation, $1 - \beta$ is the upper bound of the two advertisers' bids, and the second $1 - \beta$ factor accounts for the probability that the publisher's impression is generated, and not the fraudster's. Since the upper-bound in (WA1) is equal to the publisher's OX-only profit in (3), this completes the proof. ■

WA2 Proof of Proposition 7

Proof. Due to Lemma 8, the publisher considers either PX-only or OX-only. The profits under each of these regimes are unaffected by the C-advertiser's ability to bid selectively in the OX; therefore, the publisher's profits from the main model carry over. It follows that the publisher sells through OX-only if and only if

$$\pi_P^{\text{PX-only}} \leq \pi_P^{\text{OX-only}} \Leftrightarrow (1-\beta)\mu \leq (2-\mu)\mu(1-\beta)^2 \Leftrightarrow \beta \leq \frac{1-\mu}{2-\mu}.$$

∎

WA3 Proof of Lemma 9

Proof. In a discrete valuation setting, the advertiser's valuations under information symmetry can be generalized as

$$v = \begin{cases} \overline{v} & \text{with probability } \mu, \\ 0 & \text{with probability } 1-\mu, \end{cases} \tag{WA2}$$

for some $\overline{v} \in (0, 1]$. For example, under full information, both advertisers know the publisher is legitimate, so $\overline{v} = 1$. Under no information, neither advertiser knows the ad impression's legitimacy, so $\overline{v} = (1-\beta) \cdot 1 + \beta \cdot 0 = 1 - \beta$.

We denote the publisher's profit in the separated auction as $\pi_S(R_1, R_2)$, where R_1 and R_2 are (possibly different) reserve prices in the respective parallel auctions, and the publisher's profit in the integrated auction as $\pi_I(R)$, where R is the reserve price in the integrated auction.

We first show that $\pi_S^* \equiv \max_{R_1,R_2} \pi_S(R_1, R_2) \geq \max_R \pi_I(R) \equiv \pi_I^*$. It suffices to show that the publisher can replicate any profit under the integrated auction using the separated auction. It follows from Claim 1 that for any R, there exist R_1 and R_2 such that $\pi_S(R_1, R_2) \geq \pi_I(R)$. Therefore, $\pi_S^* \geq \pi_I^*$.

Next, we show that $\pi_S^* \leq \pi_I^*$, which would complete the proof. Note that (WA2) satisfies regularity, as defined by Myerson (1981), because $0 - \frac{1-F(0)}{f(0)} = 0 - \frac{\mu}{1-\mu} < \overline{v} - \frac{1-F(\overline{v})}{f(\overline{v})} = \overline{v}$. Therefore, it follows from Myerson (1981) that the publisher's optimal profit is achieved under a second-price auction with reserve price $\inf\{z \in \{0, \overline{v}\} : z-(1-F(z))/f(z) \geq 0\} = \overline{v}$. Revenue equivalence then implies that π_I^* (under first-price auction) obtains the same optimum; that is, π_I^* is the optimum publisher profit over all feasible mechanisms. Therefore, $\pi_S^* \leq \pi_I^*$. This completes the proof. ■

WA4 Statement and Proof of Claim 1

Claim 1. *For any $R \in [0, \overline{v}]$, $\pi_S(R, R) = \pi_I(R)$.*

Proof. In the separate auction, each advertiser $j \in \{1, 2\}$ wins if and only if its bid exceeds both R and its competitor's bid, and if it wins, it pays its own bid. That is, advertiser j with valuation v_j solves

$$\max_{b_j(v_j) \geq R} \mathbb{P}\left\{b_j(v_j) > b_{-j}(v_{-j})\right\}(v_j - b_j(v_j)),$$

where b_{-j} is the competitor's bidding strategy, and the probability is with respect to the distribution of the competitor's valuation v_{-j}. But this is equivalent to the problem advertisers solve in the integrated auction. Therefore, the optimal bidding strategies are the same across separate and integrated auctions, under equal reserve prices. Finally, since the allocation and payment rules are also the same, we obtain $\pi_S(R, R) = \pi_I(R)$. ■

WA5 Statement and Proof of Claim 2

Claim 2. *Suppose $\mu \leq \tilde{\mu}(\beta) \equiv \frac{\sqrt{4\beta^2-8\beta+5}+2\beta-3}{2(\beta-1)}$ and $\beta \leq \tilde{\beta} \equiv \frac{(1-\mu)^2}{\mu^2-\mu+1}$ such that the pre-filter equilibrium is the PX-OX regime. A pure strategy equilibrium filter γ^* is in the interval $[\min[1, \tilde{\gamma}], 1]$, where $\tilde{\gamma} = \frac{1-\mu(\beta(\mu-2)-\mu+3)}{\beta(1-\mu)}$.*

Proof. We first show that given μ and β such that the publisher adopts the PX-OX regime, the OX has incentive to fight fraud. It suffices to show that the OX's profit under $\gamma = \left(\frac{\mu}{1-\mu}\right)^2$ is higher than that under $\gamma = 0$. Note that in the parameter region for which PX-OX regime is the pre-filter equilibrium, $\left(\frac{\mu}{1-\mu}\right)^2$ is less than 1 because $\mu < \frac{1}{2}$ (see Claim 3). This is true because if $\gamma > 0$, then the C-advertiser, who was indifferent between bidding in PX and in OX before the filter due to the publisher's best-response reserve prices, switches to bidding in OX. Therefore, the OX's profit is given by (A12). the OX's profit with $\gamma = 0$ is its profit under the PX-OX regime, which is $(\beta\mu+(1-\beta)(1-\mu)\mu)R^{\mathrm{OX}} = (\beta\mu+(1-\beta)(1-\mu)\mu)\frac{(1-\beta)(1-\mu)}{1-(1-\beta)\mu} = (1-\beta)(1-\mu)\mu$. We obtain

$$
\begin{aligned}
(\text{A12}) \geq (1-\beta)(1-\mu)\mu &\Leftrightarrow (1-\beta)(1-2\beta\left(\frac{\mu}{1-\mu}\right)^2(1-\mu))\mu \geq 0 \\
&\Leftrightarrow 1-\mu-2\beta\mu^2 \geq 0 \\
&\Leftarrow 1-\mu-2\left(\frac{(1-\mu)^2}{\mu^2-\mu+1}\right)\mu^2 \geq 0 \because \beta \leq \frac{(1-\mu)^2}{\mu^2-\mu+1} \\
&\Leftrightarrow -2\mu^4+3\mu^3-2\mu+1 \geq 0,
\end{aligned}
$$

which is true because $-2\mu^4+3\mu^3-2\mu+1$ is decreasing in μ for all $\mu \in [0,1]$ and attains its minimum value 0 at $\mu = 1$. Therefore, the OX has incentive to fight fraud. Next, we show that $\gamma \in (0, \tilde{\gamma})$ cannot be equilibrium.

If $\gamma \in (0, \tilde{\gamma})$, then by definition of $\tilde{\gamma}$, we have $\mu \leq \tilde{\mu}(\beta')$ and $\beta' \leq \tilde{\beta}$, where $\beta' = \frac{\beta(1-\gamma)}{\beta(1-\gamma)+1-\beta}$ denotes the post-filter share of fake impressions. Therefore, if $\gamma \in (0, \tilde{\gamma})$, then following Lemma 4, the publisher's best response to γ is to sell through both PX and OX simultaneously at reserve prices $R^{\mathrm{OX}} = \frac{(1-\beta')(1-\mu)}{1-(1-\beta')\mu}$ and $R^{\mathrm{PX}} = \frac{1-\mu}{1-(1-\beta')\mu}$, respectively.

Given the publisher's best response, we show that the OX has incentive to filter *additional* fake impressions, thereby proving that $\gamma \in (0, \tilde{\gamma})$ cannot constitute an equilibrium. To that end, consider the OX's profit given the publisher's best response to $\gamma \in (0, \tilde{\gamma})$:

$$
\pi_{\mathrm{OX}} = (\beta(1-\gamma)\mu + (1-\beta)\mu(1-\mu))\, R^{\mathrm{OX}} = (1-\beta)(1-\mu)\mu. \tag{WA3}
$$

Now, suppose the OX filters *additional* $\gamma' = \left(\frac{\mu}{1-\mu}\right)^2$ fraction of fake impressions. Since the publisher's best-response reserve prices are set such that the C-advertiser is indifferent between bidding in PX and in OX, fighting fraud induces the C-advertiser to switch to OX. Thus, both advertisers bid in OX and they mix on the interval $\left[R^{\mathrm{OX}}, \overline{b}\right]$, where the mixing distribution G and the upper bound of the support $\overline{b}'$ are determined by the following indifference conditions for the j-advertiser, $j \in \{C, U\}$:

$$
\begin{aligned}
\pi_j\left(R^{\mathrm{OX}}\right) &= (1-\beta+\beta(1-\gamma''))(1-\mu)\left(\frac{1-\beta}{1-\beta+\beta(1-\gamma'')} - R^{\mathrm{OX}}\right) \\
= \pi_j\left(\overline{b}'\right) &= (1-\beta+\beta(1-\gamma''))\left(\frac{1-\beta}{1-\beta+\beta(1-\gamma'')} - \overline{b}'\right) \\
= \pi_j(b) &= (1-\beta+\beta(1-\gamma''))(1-\mu+\mu G(b))\left(\frac{1-\beta}{1-\beta+\beta(1-\gamma'')} - b\right)
\end{aligned}
$$

where $1-\gamma'' \equiv (1-\gamma)(1-\gamma')$. It follows that

$$
\overline{b} = (1-\mu)R^{\mathrm{OX}} + \frac{\mu(1-\beta)}{1-\gamma''\beta} \text{ and } G(b) = \begin{cases} 0 & \text{if } b \leq R^{\mathrm{OX}}, \\ \frac{(1-\mu)\left(b-R^{\mathrm{OX}}\right)(1-\beta\gamma'')}{\mu(1-\beta-b(1-\beta\gamma''))} & \text{if } R^{\mathrm{OX}} < b \leq \overline{b}', \\ 1 & \text{if } \overline{b}' < b. \end{cases}
$$

Therefore, the advertisers' expected profits in OX is

$$
\begin{aligned}
\pi_j^{\mathrm{OX}} &= (1-\beta+\beta(1-\gamma''))(1-\mu)\left(\frac{1-\beta}{1-\beta+\beta(1-\gamma'')} - R^{\mathrm{OX}}\right) \\
&= \frac{(1-\beta)\beta(1-\gamma)(1-\mu)(\gamma'(1-\mu)+\mu)}{1-\beta(\gamma-\mu)-\mu}.
\end{aligned}
$$

We show that given the U-advertiser's mixed bid according to G, the C-advertiser has no incentive to deviate to bidding in PX. If the C-advertiser bids R^{PX} in PX, then its profit is $\pi_C^{\mathrm{PX}} = (1-\beta)\left(1-R^{\mathrm{PX}}\right) = \frac{(1-\beta)\beta(1-\gamma)\mu}{1-\beta(\gamma-\mu)-\mu}$. It can be shown that $\pi_j^{\mathrm{OX}} \geq \pi_C^{\mathrm{PX}} \Leftrightarrow \gamma' \geq \left(\frac{\mu}{1-\mu}\right)^2$. Since $\gamma' = \left(\frac{\mu}{1-\mu}\right)^2$, the C-advertiser does not deviate to bidding in PX.

Next, we show that the OX's profit under $\gamma' = \left(\frac{\mu}{1-\mu}\right)^2$ is greater than that if the OX does not filter additional fake impressions (i.e., $\gamma' = 0$).

$$\pi_{\text{OX}}\left(\gamma' = \left(\frac{\mu}{1-\mu}\right)^2\right) = (1-\beta+\beta(1-\gamma''))\left(2\mu(1-\mu)\int_{R^{\text{OX}}}^{\bar{b}} b\, dG(b) + \mu^2 \iint_{R^{\text{OX}}}^{\bar{b}} \max[b_C, b_U]\, dG(b_C)dG(b_U)\right)$$
$$= \frac{(1-\beta)\mu\left((1-\beta)\mu^2 + (2-3\mu)(1-\beta\gamma)\right)}{-((1-\beta)\mu) - \beta\gamma + 1},$$
$$\pi_{\text{OX}}(\gamma' = 0) = (1-\beta)(1-\mu)\mu,$$

where the last equality follows from (WA3). We obtain

$$\pi_{\text{OX}}(\gamma') \geq \pi_{\text{OX}}(0) \Leftrightarrow \frac{(1-\beta)\mu(-\mu(\beta(1-2\gamma)+1) - \beta\gamma + 1)}{-((1-\beta)\mu) - \beta\gamma + 1} \geq 0$$
$$\Leftrightarrow -\mu(\beta(1-2\gamma)+1) - \beta\gamma + 1 \geq 0.$$

We show that the last inequality is true: $-\mu(\beta(1-2\gamma)+1)-\beta\gamma+1$ is decreasing in γ because its derivative with respect to γ is $\beta(2\mu-1) \leq 0$, where the non-positivity follows from the fact that $\mu \leq \tilde{\mu}(\beta)$ and $\beta \leq \tilde{\beta}$ jointly imply $\mu \leq \frac{1}{2}$ (see Claim 3). Therefore, $-\mu(\beta(1-2\gamma)+1) - \beta\gamma + 1$ is decreasing in γ, such that $-\mu(\beta(1-2\gamma)+1) - \beta\gamma + 1 \geq (-\mu(\beta(1-2\gamma)+1) - \beta\gamma + 1)\,|_{\gamma=1} = (1-\beta)(1-\mu) \geq 0$.

Finally, suppose $\gamma \geq \min[1, \tilde{\gamma}]$. If $\tilde{\gamma} > 1$, then at $\gamma = 1$, all fraudulent impressions are filtered out, and following Lemma 4, the publisher best-responds by setting $R^{\text{OX}} = 1$ and $R^{\text{PX}} = 1$. This is an equilibrium because in this PX-OX regime, OX only has incentive to increase γ (see first part of the proof above), but it cannot filter out more than $\gamma = 1$. If $\tilde{\gamma} < 1$, then for all $\gamma > \tilde{\gamma}$, the publisher best-responds by selling exclusively through the OX because $\gamma > \tilde{\gamma}$ implies that the post-filter share of fake impressions in OX is

$$\beta' = \frac{\beta(1-\gamma)}{1-\beta+\beta(1-\gamma)} \leq \frac{(3-\mu)\mu - 1}{(2-\mu)\mu} \Leftrightarrow \mu > \frac{3 - 2\beta' - \sqrt{4(\beta')^2 - 8\beta' + 5}}{2(1-\beta')},$$

and this is the condition under which the publisher's optimal strategy is to sell exclusively through the OX (see Proposition 3). Following Lemma 5, OX has no incentive to filter further fake impressions. This completes the proof.

■

WA6 Statement and Proof of Claim 3

Claim 3. *Suppose* $\mu \leq \tilde{\mu}(\beta) \equiv \frac{\sqrt{4\beta^2-8\beta+5}+2\beta-3}{2(\beta-1)}$ *and* $\beta \leq \tilde{\beta} \equiv \frac{(1-\mu)^2}{\mu^2-\mu+1}$. *If* $\mu \leq \tilde{\mu}(\beta)$ *and* $\beta \leq \tilde{\beta}$, *then* $\mu \leq \frac{1}{2}$.

Proof. First, $\tilde{\beta}$ is decreasing in μ, because $\frac{d\tilde{\beta}}{d\mu} = -\frac{1-\mu^2}{(1-\mu+\mu^2)^2} < 0$. Therefore, $\beta \leq \tilde{\beta}$ is equivalent to $\mu \leq \frac{2-\beta-\sqrt{\beta(4-3\beta)}}{2(1-\beta)}$, where the RHS is the μ-root of $\beta = \tilde{\beta}$.

Second, $\tilde{\mu}(\beta)$ is increasing in β because $\frac{d\tilde{\mu}(\beta)}{d\beta} = \frac{1-\frac{1}{\sqrt{4\beta^2-8\beta+5}}}{2(1-\beta)^2} \propto 1 - \frac{1}{\sqrt{4\beta^2-8\beta+5}} \geq 0$, where the last inequality follows from $4\beta^2 - 8\beta + 5 \geq 1 \Leftrightarrow 4(1-\beta)^2 \geq 0$.

Third, $\frac{2-\beta-\sqrt{\beta(4-3\beta)}}{2(1-\beta)}$ is decreasing in β because its derivative with respect to β is $\frac{\beta+\sqrt{(4-3\beta)\beta}-2}{2(\beta-1)^2\sqrt{(4-3\beta)\beta}}$, which is proportional to $\beta + \sqrt{(4-3\beta)\beta} - 2$, and $\beta + \sqrt{(4-3\beta)\beta} - 2 \leq 0$ because $\beta + \sqrt{(4-3\beta)\beta} - 2 \leq 0 \Leftrightarrow (4-3\beta)\beta \leq (2-\beta)^2 \Leftrightarrow 0 \leq (1-\beta)^2$.

Therefore, the joint condition $\mu \leq \tilde{\mu}$ and $\mu \leq \frac{2-\beta-\sqrt{\beta(4-3\beta)}}{2(1-\beta)}$ implies μ is smaller than the value of μ at which the two bounds meet: $\tilde{\mu} = \frac{2-\beta-\sqrt{\beta(4-3\beta)}}{2(1-\beta)} \Leftrightarrow \beta = \frac{1}{3}, \mu = \frac{1}{2}$. This completes the proof.

■

WA7 Statement and Proof of Claim 4

Claim 4. *The publisher's equilibrium profit weakly decreases in* β.

Proof. The publisher's equilibrium profit is the maximum of its profit under each of the three regimes: OX-only, PX-only, and PX-OX. Since each of these profit is continuous is β, so is the maximum of the three profits. Therefore, to prove the claim, it suffices to show that each of these profits is decreasing in β.

- For OX-only, we obtain from (3) that

$$\frac{\partial \pi_P^{\text{OX-only}}}{d\beta} = -2(2-\mu)\mu < 0.$$

- For PX-only, we obtain from (5) that

$$\frac{\partial \pi_P^{\text{PX-only}}}{d\beta} = 0.$$

- For PX-OX, we obtain from (12) that

$$\frac{\partial \pi_P^{\text{PX-OX}}}{d\beta} = -\frac{(1-\mu)\mu}{(1-(1-\beta)\mu)^2} < 0$$

This completes the proof. ■

WA8 Statement and Proof of Claim 5

Claim 5. *Under the OX-only regime, the OX has no incentive to fight fraud, even if (a) advertisers have outside option $u_0 > 0$, or (b) publishers can adjust the reserve price after observing OX's filter level γ.*

Proof. We prove parts (a) and (b) in turn. First, suppose advertisers have outside option utility $u_0 > 0$. If $u_0 > 1$, then advertisers opt for the outside option such that the OX's profit is 0 regardless of its fraud filter level. Therefore, γ^* is degenerately 0.

Suppose $u_0 \leq 1$. In the OX-only regime, the publisher sets reserve price up to the advertisers' valuations, which is

$$R^{\text{OX}} = \max\{R : \pi(R) \geq u_0\} = \max\{R : (1-R) + \frac{\beta}{1-\beta}(0-R) \geq u_0\} = (1-\beta)(1-u_0).$$

If the OX filters γ fraction of fraudulent impressions, advertisers' valuations are now greater than the reserve price. Therefore, advertisers mix on the interval $[(1-\beta)(1-u_0), \bar{b}]$ according to distribution G which satisfies the following indifference conditions:

$$\begin{aligned}
\pi((1-\beta)(1-u_0)) &= (1-\mu)\left(\frac{1-\beta}{1-\beta+\beta(1-\gamma)} - (1-\beta)(1-u_0)\right) \\
&= \pi(b) = (1-\mu+\mu G(b))\left(\frac{1-\beta}{1-\beta+\beta(1-\gamma)} - b\right) \\
&= \pi(\bar{b}) = \frac{1-\beta}{1-\beta+\beta(1-\gamma)} - \bar{b}.
\end{aligned}$$

This implies

$$\begin{aligned}
\bar{b} &= \frac{(1-\beta)\beta\gamma\mu}{1-\beta\gamma} - \beta - (1-\mu)(1-\beta)u_0 + 1, \\
G(b) &= \frac{(\mu-1)(\beta\gamma-1)(b-(1-\beta)(1-u_0))}{\mu(b(\beta\gamma-1)-\beta+1)}.
\end{aligned}$$

the OX's profit is

$$\begin{aligned}
\pi_{\mathrm{OX}} &= \alpha(1-\beta\gamma)\left(2\mu(1-\mu)\int_{(1-\beta)(1-u_0)}^{\bar{b}} b dG(b) + \mu^2 \iint_{(1-\beta)(1-u_0)}^{\bar{b}} \max[b_1, b_2] dG(b_1) dG(b_2)\right) \\
&= \alpha\mu((1-\beta)(2-\mu-2\beta\gamma(1-\mu)) - 2(1-\mu)(1-\beta)u_0(1-\beta\gamma)).
\end{aligned}$$

Since $\frac{d\pi_{\mathrm{OX}}}{d\gamma} = -2\beta\mu(1-\mu)(1-\beta)(1-u_0) < 0$, we obtain $\gamma^* = 0$.

Second, consider part (b) of the claim: suppose the publisher can observe and adjust reserve prices. For any given β in the OX-only regime, reducing β to $\beta(1-\gamma)$ will not induce any regime change (i.e., to PX-only or PX-OX). This is because the OX-only condition (see Proposition 3) can be re-written as

$$\frac{3-\sqrt{5}}{2} < \mu \text{ and } \beta \leq \min\left[\frac{\mu^2-3\mu+1}{(\mu-2)\mu}, \frac{1-\mu}{2-\mu}\right],$$

such that the OX-only regime sustains for any smaller values of β. Thus, for any $\gamma \in [0,1]$, the publisher's reserve price is

$$R^{\mathrm{OX}} = \frac{1-\beta}{1-\beta+\beta(1-\gamma)},$$

which implies that OX's profit is

$$\pi_{\mathrm{OX}} = \alpha(1-\beta+\beta(1-\gamma))\left(1-(1-\mu)^2\right)R^{\mathrm{OX}} = \alpha\mu(2-\mu)(1-\beta). \tag{WA4}$$

Therefore, $\frac{d\pi_{\mathrm{OX}}}{d\gamma} = 0$. ■

WA9 Statement and Proof of Claim 6

Claim 6. *The OX fights fraud in the OX-only regime under the following conditions: (a) publishers can adjust the reserve price after observing OX's filter level γ, and (b) the OX incurs a reputational cost from sales of fraudulent impressions.*

Proof. Following the derivation of (WA4) in Claim 5, we obtain

$$\begin{aligned}
\pi_{\mathrm{OX}} &= \alpha(1-\beta)\left(1-(1-\mu)^2\right)R^{\mathrm{OX}} + \alpha\delta\beta(1-\gamma)\left(1-(1-\mu)^2\right)R^{\mathrm{OX}} \\
&= \alpha\mu(2-\mu)(1-\beta)\frac{\beta(1-\gamma)\delta-\beta+1}{1-\beta\gamma},
\end{aligned}$$

where $\delta \in (0,1)$ is the discounting associated with the reputational cost from selling fraudulent impressions. Since

$$\frac{d\pi_{\mathrm{OX}}}{d\gamma} = \frac{(1-\beta)^2\beta(1-\delta)(2-\mu)\mu}{(1-\beta\gamma)^2} > 0,$$

we obtain $\gamma^* = 1$. ■