

Open and Private Exchanges in Display Advertising

W. Jason Choi

Amin Sayedi*

September 2021

*W. Jason Choi (jason.choi@rutgers.edu) is Assistant Professor of Marketing at the Rutgers Business School, Rutgers University. Amin Sayedi (aminsa@uw.edu) is Associate Professor of Marketing at the Foster School of Business, University of Washington. The authors contributed equally.

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 private exchanges are only available to a smaller set of pre-screened advertisers and publishers through an invite-only process. OX potentially exposes advertisers to ad fraud and brand safety issues, whereas PX ensures that advertisers purchase high-quality impressions from reputable publishers. While the assurance of higher quality increases advertisers' valuation for PX impressions, we find that selling through both OX and PX can hurt publishers by creating an information asymmetry among advertisers. Intuitively, the presence of PX informationally advantages the “connected” advertisers, who have access to PX, while simultaneously disadvantaging the “unconnected” advertisers, who only have access to OX. Therefore, compared to the OX-only benchmark where advertisers are equally uninformed about impression quality, the unconnected advertisers anticipate a higher chance of winning low-quality impressions and thus lower their valuations. This devaluation effect softens bidding competition and reduces the publisher's expected profit. We show that the publisher's optimal exchange choice—sell through OX, PX, or both—depends crucially on the baseline fraud intensity and the advertisers' average valuations. Finally, our model sheds light on OX's incentive to fight fraud. In the absence of PX, OX has low incentive to combat fraud because it earns commission from fraudulent transactions. However, the introduction of PX may create competitive pressure such that OX screens fake impressions; i.e., 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 terms of 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 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 Google’s 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 them 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

¹<https://forecasts-na1.emarketer.com/584b26021403070290f93a56/5851918a0626310a2c1869c4>

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 decided to set up their own private exchanges, where publishers 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 been growing rapidly in recent years, and in 2020, it surpassed that of open exchanges for the first time. Ad spending growth in private marketplaces is projected to outpace that in open exchanges by approximately 3 to 1 in 2021 (Fisher, 2020).

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 (Vrontas, 2020). Third, advertisers 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. One of the main drawbacks for publishers selling through private marketplaces is softer bidding competition. Since private exchanges are available only to a small 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

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

³Note 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).

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. Since the bids in one exchange cannot be used for pricing in another exchange, receiving bids from multiple exchanges may still reduce the publishers' revenues.

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? How do the effects of private exchange depend on the advertisers' accessibility to the private exchange?
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, 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.

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. 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. We find that if ad fraud in the system is sufficiently mild and the advertisers’ average willingness-to-pay for a legitimate impression is sufficiently high, then the publisher should not set up a private exchange. More generally, we characterize the conditions under which a publisher should sell 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, simply because 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 out fraudulent impressions reduces the supply-side ad inventory, it increases the advertisers’ demand for 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. Our results provide guidance on how advertisers should adjust their bidding strategies in the presence of a private exchange. We highlight 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. To the best of our knowledge, our work is the first in the marketing and economics literature to study the impact of private exchanges on advertisers’ and publishers’ strategies in the RTB market.

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 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 allocation 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. \(2015a\)](#) 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. \(2015b\)](#) 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](#)). On the theoretical front, [Sayedi et al. \(2018\)](#) study advertisers' bidding strategies when publishers allow advertisers to bid for exclusive placement on the website. [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 [Gritkevich et al. \(2018\)](#) look at how ad blockers affect the online advertising ecosystem, and [Dukes et al. \(2020\)](#) show how skippable ads affect publishers' and advertisers' strategies as well as their profits. [Choi et al. \(2021\)](#) 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. The models in [Zeithammer \(2019\)](#) and [Sayedi \(2018\)](#) have only one exchange, and, therefore, cannot distinguish between open and private exchanges. In contrast, our model focuses on the differences between the two types of exchanges in RTB and how they impact advertisers’ and publishers’ strategies. [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. [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. To the best of our knowledge, our paper is the first in the marketing and economics literature that studies the impact of private exchanges on advertisers’ and publishers’ strategies.

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 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 a publisher can provide PX access to any

advertiser, this involves considerable costs. First, the publisher must trust the advertiser to share proprietary information with them as this information could potentially be revealed to the publisher’s competitors.⁴ Second, there are fixed costs such as signing contracts and non-disclosure agreements that deter the publisher from providing PX access to all advertisers. In sum, the set of advertisers that has access to PX is a strict subset of those who have access to the OX; therefore, fewer advertisers compete in PX than in OX. To parsimoniously capture this institutional feature, we assume that only one of the two advertisers (i.e., the C -advertiser) has access to PX. Apart from accessibility to PX, the C - and U -advertisers are ex ante symmetric.

An ad impression is generated from either the publisher or the fraudster. With probability (w.p.) β , the impression is drawn from the fraudster, and with probability $1 - \beta$, it is drawn from the publisher. Thus, β measures fraud intensity in the system. β could also be interpreted as the inverse measure of the extent and sophistication of the industry’s anti-fraud enforcement. In the main model, we keep β exogenous; however, in Section 4, we allow players’ actions to influence β . For $j \in \{C, U\}$, the j -advertiser’s value for an impression i consists of both impression-specific and advertiser-specific factors; i.e.,

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

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 i.i.d. across advertisers according to

$$\nu_j = \begin{cases} \bar{\nu} & \text{w.p. } \mu, \\ \underline{\nu} & \text{w.p. } 1 - \mu, \end{cases} \quad (1)$$

where $0 \leq \underline{\nu} < \bar{\nu}$. We normalize $\underline{\nu}$ to 0 and $\bar{\nu}$ to 1. Advertisers privately know their own

⁴For example, see O’Reilly (2015). 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).

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), $\mu \in [0, 1]$ 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 $\bar{\nu}$, μ can also be interpreted as the expected value of an advertiser for a legitimate impression. Depending on their accessibility to 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.⁵

The publisher sells its ad inventory via first-price auctions with reserve prices R^{PX} and R^{OX} in PX and OX, respectively.⁶ When there are two exchanges, the publisher sends request-for-bids for the impression to both exchanges simultaneously. Each exchange auctions off the 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 is thus the maximum of the two exchanges’ clearing prices.⁷ When there is only one exchange, the impression is sold via a standard first-price auction where 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 ad impression coming through PX as originating from the legitimate publisher. On the other hand, if the same impression is sent to OX, then neither the C -advertiser nor the U -advertiser can discern whether the

⁵In 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).

⁶For more information on the emergence and prevalence of first-price auctions in the RTB market, see Despotakis et al. (2021a).

⁷The process of sending the impression to multiple exchanges simultaneously, and allocating it to the exchange with the highest price, is known as *header bidding*; it is common practice in the industry (Sluis, 2016).

impression is legitimate or fraudulent. This is because the fraudster sends its request-for-bid in OX disguised as the legitimate publisher, mimicking all aspects of the publisher’s request-for-bid, including the reserve price set by the publisher.⁸

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

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

where F_{-j} is a cumulative distribution function denoting the j -advertiser’s belief about its competitor’s bid b_{-j} . Similarly, its expected payoff when it knows λ_i is

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

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. An impression is drawn either from a publisher w.p. $1 - \beta$, or from a fraudster w.p. β .
3. The j -advertiser realizes its value ν_j and submits its bids. The C -advertiser submits bids b_C^{PX} and b_C^{OX} to the private exchange and the open exchange, respectively. The

⁸Even though the C -advertiser observes legitimate impressions in PX, it cannot identify the same impressions coming through OX because different exchanges use different identifiers and cookies.

U -advertiser submits its bid b_U^{OX} to the open exchange.

4.
 - 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; otherwise the impression is left unsold.
 - For a fraudulent impression, only the OX runs a first-price auction. If the highest bid is greater than or equal to the reserve price, the fraudulent impression is allocated to the highest bidder; otherwise, the fraudulent impression is left unsold.

Finally, the payments are made and the 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 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 will help elucidate the impact of the introduction of private exchanges on the RTB market.

3.1 OX-Only Benchmark

Suppose the publisher can only sell ad inventory through OX. Due to the open nature of the exchange, advertisers buying in OX are prone to fraud. Specifically, the 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.

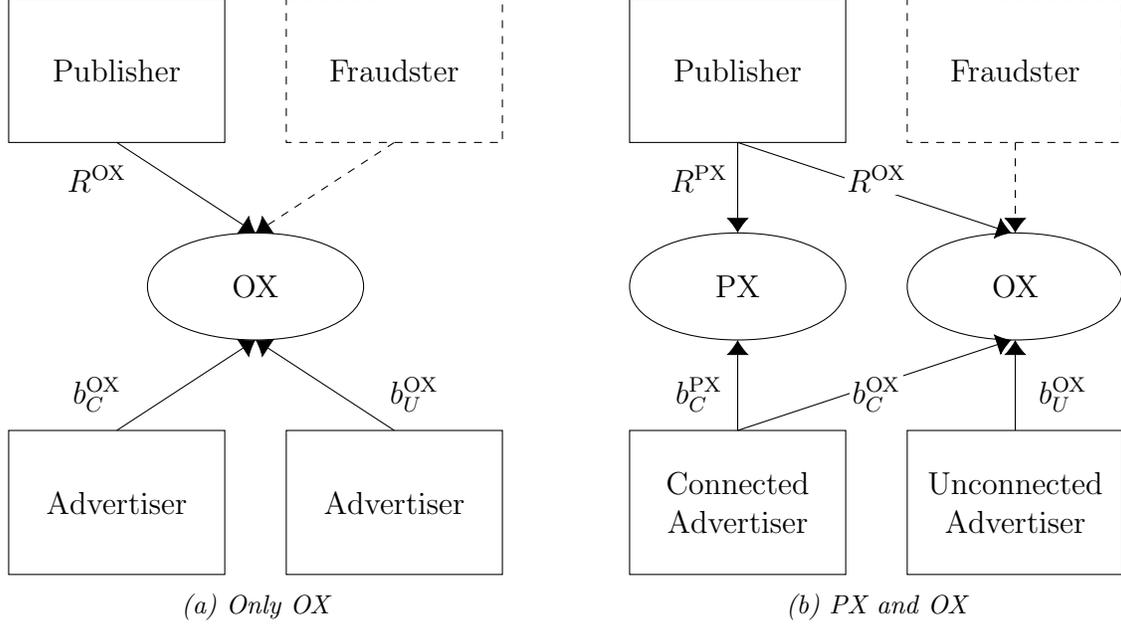


Figure 1: Exchange Configurations

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

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

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' (symmetric) bids are*

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

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

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

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.⁹ 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 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. In particular, the publisher adopts one of the following three regimes: (i) sell only through OX, (ii) sell only through PX, and (iii) sell through PX and OX simultaneously. 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 under each regime, 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 PX, and sells through PX and OX simultaneously.

If the publisher sells exclusively through PX, then the separation of exchanges reveals fraudster’s impressions in OX. Therefore, no transactions occur in OX. On the other hand, in

⁹The qualitative insights are robust to other tie-breaking rules.

PX, the C -advertiser has valuation 1 w.p. μ and valuation 0 w.p. $1 - \mu$ 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*

$$R^{\text{PX-only}} = b_C^{\text{PX-only}} = 1. \quad (4)$$

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

$$\pi_P^{\text{PX-only}} = \mu \cdot 1 = \mu. \quad (5)$$

Comparison of the publisher's OX-only profit (3) and its PX-only profit (5) reveals the publisher's margin-volume trade-off. If the publisher sells exclusively through PX, then compared to selling exclusively through OX, demand for ad slots is lower (i.e., $\mu \leq 1 - (1 - \mu)^2$) since only the C -advertiser can bid in PX. On the other hand, the margin per transaction is higher if it sells exclusively through PX (i.e., $1 \geq 1 - \beta$) because the C -advertiser's knowledge that PX impressions are legitimate increases its bid in PX.

Finally, consider the third regime in which the publisher sells through PX and OX simultaneously. The publisher decides R^{PX} and R^{OX} , the reserve prices in PX and OX, respectively. The following proposition summarizes the players' strategies and payoffs under the dual exchange regime.

Lemma 3. *Let b_C^{PX} and b_U^{OX} denote the C -advertiser's bid in PX and the U -advertiser's bid in OX, respectively. In the PX-OX regime where the publisher sells through PX and OX simultaneously, the equilibrium reserve prices and the advertisers' bids are*

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

and

$$R^{OX} = b_U^{OX} = \frac{(1 - \beta)(1 - \mu)}{1 - (1 - \beta)\mu}. \quad (6)$$

The C -advertiser's expected profit is

$$\pi_C^{PX-OX} = \frac{(1 - \beta)\beta\mu}{1 - (1 - \beta)\mu}, \quad (7)$$

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

$$\pi_P^{PX-OX} = \frac{(2 - \mu - \beta(1 - \mu))(1 - \mu)\mu}{1 - (1 - \beta)\mu}.$$

Lemma 3 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., $b_U^{OX\text{-only}}$ in (6) is less than b_U^{OX} in (2) (see Figure 2). The intuition is as follows. In the presence of PX, the U -advertiser knows that it competes against the informationally advantaged C -advertiser, who bids high in PX (for the legitimate publisher's impression) and bids nothing in OX.¹⁰ 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 total, the introduction of PX creates an information asymmetry between the advertisers that dampens the U -advertiser's valuation for impressions in OX. We call this the *devaluation effect*.

Second, the reserve price in 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 PX to 1, even though the high-valuation C -advertiser in PX knows that the impression is legitimate (and thus values the impression at 1). The reason is that under the PX-only regime, the C -

¹⁰Note that it is never optimal for the C -advertiser to bid in both exchanges simultaneously. See Claim 1 in the appendix for details.

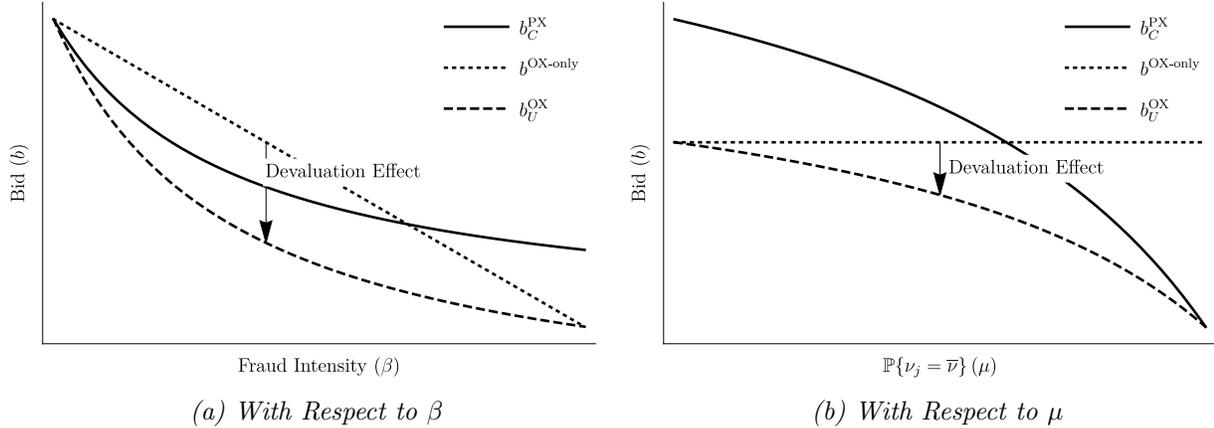


Figure 2: Advertisers' Bids and Devaluation Effect

advertiser has no outside option: if it does not win the impression in 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 PX is set too high, the C -advertiser switches to buying (potentially fraudulent) impressions in OX. In other words, if the reserve price in PX is too high, the C -advertiser's expected payoff will be higher buying in OX with the risk of ad fraud than buying a guaranteed legitimate impression in PX at a high price. In sum, OX cannibalizes PX and reduces the publisher's revenue from 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 OX; this, in turn softens bidding competition for the C -advertiser, leading to lower bids in PX. Due to the cannibalization effect, the publisher cannot set a high reserve price in PX to offset the devaluation effect in OX. That is, if R^{PX} is too high, the C -advertiser will switch to bidding in OX. In the following proposition, we summarize the central force generated by the introduction of PX.

Proposition 1 (Devaluation Effect). *The introduction of 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*

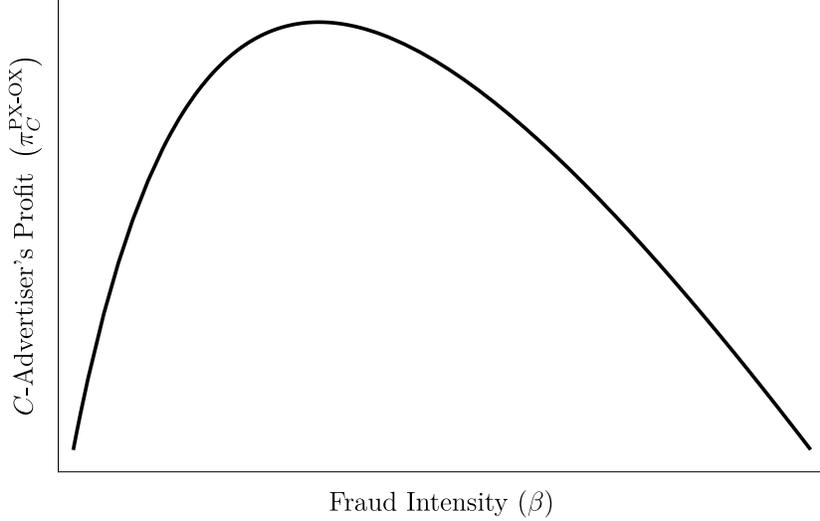


Figure 3: *C-Advertiser's Profit Under PX-OX Regime*

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 $b^{OX\text{-only}}$ 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 PX, the U -advertiser's valuation of ad impressions in 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, informationally-advantaged C -advertiser, who bids higher for the legitimate impression than the U -advertiser does in OX. Again, 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 deeply as μ increases.

Interestingly, under the PX-OX regime, higher fraud intensity has non-monotonic effects on

the C -advertiser’s profit (see Figure 3). For large β , higher fraud depresses the C -advertiser’s profit as larger β implies fewer opportunities to buy legitimate impressions through PX. In contrast, for small β , the U -advertiser’s devaluation effect amplifies with β . This softens bidding competition, allowing the C -advertiser to win legitimate impressions in PX with lower bids. Note that due to the cannibalization effect, the publisher cannot set a high reserve price in PX to offset the devaluation effect in OX. The following proposition summarizes this finding.

Proposition 2. *Under the PX-OX regime, the C -advertiser’s profit increases in β if $\beta \leq (\sqrt{1-\mu} - (1-\mu)) / \mu$, and decreases in β otherwise.*

In summary, the comparison of the OX-only benchmark with the regimes with PX sheds light on important insights regarding the impact of the introduction of PX on the RTB market. First, the introduction of PX distorts the advertisers’ information structure such that the U -advertiser values impressions less than it does without PX. This lowers the U -advertiser’s bid in the OX, and softens competition for the C -advertiser. The U -advertiser’s lower bid in OX also makes bidding in OX more attractive for the C -advertiser. Therefore, due to the cannibalization effect, the publisher lowers the reserve price in PX. This allows the C -advertiser to win impressions in PX at a lower 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 implications of the various forces related to fraud (i.e., direct effect of fraud, 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. *The publisher’s equilibrium ad exchange choices are as follows:*

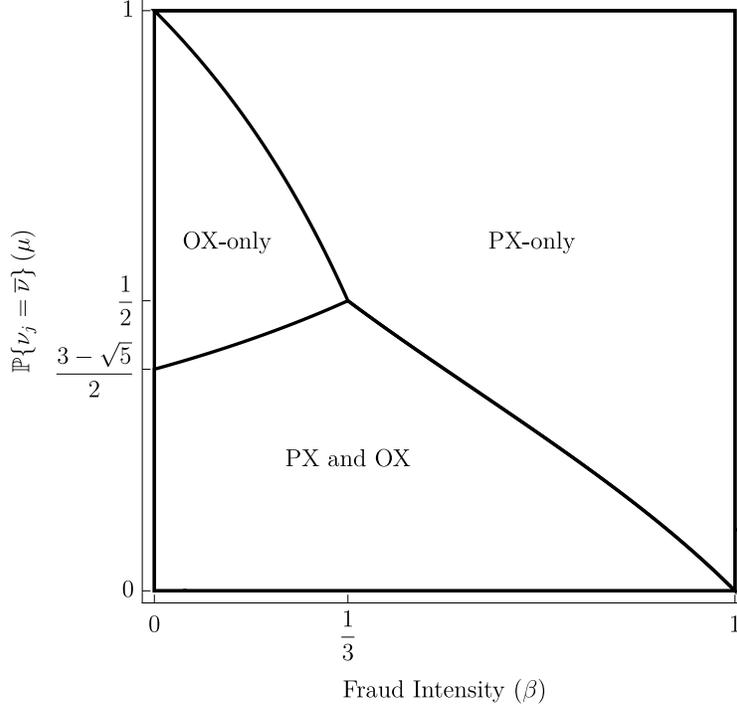


Figure 4: Publisher's Equilibrium Exchange Choices

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 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 PX;
3. otherwise, the publisher sells through both PX and OX.

Proposition 3 shows that even if the publisher has the option to sell through PX, which helps connected advertisers distinguish legitimate impressions from fake ones, it does not always choose to do so. Specifically, if μ is large and β is small, then the publisher sells exclusively through OX (see Figure 4). The intuition is that the devaluation effect is severe under large μ as the U -advertiser anticipates a higher probability of facing an informationally advantaged, high-valuation C -advertiser. Moreover, small β ensures that bids under the OX-only regime are sufficiently high. Taken together, if μ is large and β small, foregoing PX is more profitable for the publisher than selling through PX.

On the other hand, if μ and β are small, the publisher sells through both exchanges. This is because the devaluation effect is mitigated for small μ , and the direct, negative effect of β in

OX is mild for small β . In this case, the publisher sells through both PX and OX, thereby capitalizing on both the C -advertiser's high valuation in PX, and the market expansion effect in OX.

Finally, if μ or β is large, the publisher's optimal strategy is to sell only through PX. If μ is large, the devaluation effect is strong, and if β is large, most impressions sold in the OX are fraudulent. Both of these conditions dampen the U -advertiser's willingness-to-pay in OX. Consequently, the U -advertiser's bid is sufficiently low that selling through OX has limited upside for the publisher. In this case, the cannibalization effect of OX dominates the positive impact of selling to the U -advertiser. As such, if μ or β is large, the publisher sells exclusively through PX.

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 also qualitatively different from the standard product line design setting due to its effect on the advertisers' information structure. Specifically, the introduction of the PX not only ensures a "higher quality" for advertisers who buy in PX, but also informationally disadvantages advertisers who do not have access to PX, which in turn lowers their valuations. In total, while the cannibalization effect deters the publisher from selling through OX, the low-quality analogue, the devaluation effect induced by the information asymmetry deters it from selling through 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 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.

Competition Dispersion Effect

Before concluding 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 i.i.d. valuations $U[0, 1]$ compete in a second-price auction with (the optimal) reserve price $1/2$. The publisher’s expected revenue from this auction is $5/12$.¹¹ On the other hand, if the two advertisers bid 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 $5/12$.¹²

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 advertiser 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

¹¹With 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 $1/2 \cdot 1/2 + 1/4 \cdot 2/3 = 5/12$.

¹²It can be shown that even if the publisher optimizes the reserve prices under two separate second-price auctions, the qualitative insight holds.

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 4. *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.*

Note that the result of Lemma 4 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.

4 OX and Anti-Fraud Efforts

We have hitherto assumed OX to be passive. While this assumption allowed us to obtain sharp insights about the effect of introducing PX on the RTB market, OX may play a more active role in gatekeeping the types of ad impressions it sells (Graham, 2020). In this section, we explore the OX’s incentives (or lack thereof) to fight fraud and analyze how OX’s strategy may affect the qualitative insights from the main model. To that end, we augment the main model such that the OX decides $\gamma \in [0, 1]$ fraction of fraudulent request-for-bids to filter out, simultaneously as the publisher sets the reserve price(s).

To sharpen insights, we assume that it is costless for the OX to identify and remove fake impressions. Consistent with industry practice, we assume that the OX’s profit is based on a fixed α -commission rate per transaction occurring through OX, for some $\alpha \in (0, 1)$ (Hsiao, 2020).

We begin the analysis for the OX-only benchmark, and then analyze the OX’s equilibrium filter level with PX. In the benchmark scenario without PX, fighting fraud has two effects on the OX’s profit. First, it reduces the OX’s profit because filtering out fraudulent impressions decreases 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 OX. This is because filtering out fraudulent request-for-bids reduces the volume of transactions that occur through OX, and since OX's revenue is on a commission-per-trade basis, its expected profit decreases.

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

Proposition 4. *If the publisher has the option to sell through 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, the OX's anti-fraud efforts may 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 Figure 5). These are the conditions under which the publisher has incentive to sell through both PX and OX (see Figure 4). In other words, if it is optimal for the publisher to sell only through OX even without any anti-fraud efforts from 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 OX at higher margins.

Proposition 4 highlights another interesting aspect of the OX's anti-fraud efforts. Since the

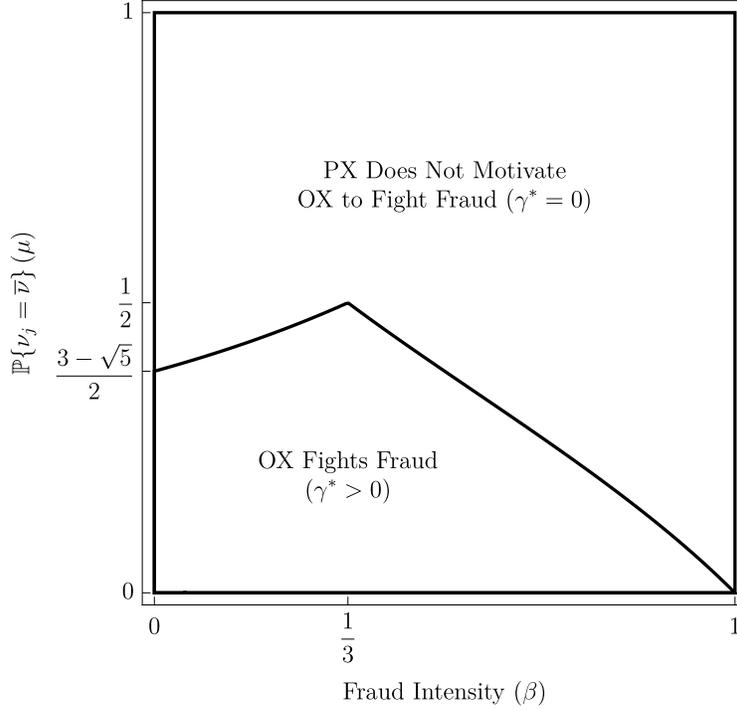


Figure 5: OX's Equilibrium Anti-Fraud Efforts

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 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 valuation for impressions; this in turn shrinks the total revenue the publisher can extract from the unconnected advertisers. Second, as the unconnected advertisers lower their bids, the bidding competition softens, and in response, even the connected advertisers lower 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 depending on the market conditions. In particular, we show that if the baseline fraud in the system is sufficiently mild and the advertisers' average ad valuations are sufficiently high, then the negative 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 combat ad fraud in the form of filtering

out fraudulent 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 PX and OX, then the OX strategically responds by fighting fraud to lure the connected advertisers, who have access to PX, to transact through the open exchange.

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 select 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

- Amaldoss, W., Desai, P. S., and Shin, W. (2015a). Keyword Search Advertising and First-Page Bid Estimates: A Strategic Analysis. *Management Science*, 61(3):507–519.
- Amaldoss, W., Jerath, K., and Sayedi, A. (2015b). Keyword Management Costs and “Broad Match” in Sponsored Search Advertising. *Marketing Science*, 35(2):259–274.
- 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/3151WcK>.
- 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. (2021). Consumer purchase journey, targeted advertising, and privacy choices.
- 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., 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.
- Dukes, A. J., Liu, Q., and Shuai, J. (2020). Skippable ads: Interactive advertising on digital media platforms. *Available at SSRN 3169629*.
- Fisher, L. (2020). “Private Marketplace Ad Spending to Surpass Open Exchange in 2020,” (accessed September 7, 2021). <https://bit.ly/3BVtd5G>.
- 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, M. (2020). “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>.
- Gritkevich, A., Katona, Z., and Sarvary, M. (2018). Ad blocking. *Columbia Business School Research Paper*, (18-11).
- He, A. (2019). “The Size of the Ad Fraud Problem in Digital Marketing Is Varying,” (accessed September 7, 2021). <https://bit.ly/3jXJqkJ>.
- 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.

- 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.
- 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.
- 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.
- 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*.
- 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/3hm5bZ0>.
- Rafieian, O. (2020). Revenue-optimal dynamic auctions for adaptive ad sequencing. Technical report, Working paper: Cornell University.

- Rafeian, 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>.
- Sayed, A. (2018). Real-time bidding in online display advertising. *Marketing Science*, 37(4):553–568.
- Sayed, A., Jerath, K., and Baghaie, M. (2018). Exclusive placement in online advertising. *Marketing Science*, 37(6):970–986.
- Sayed, 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>.
- Villas-Boas, J. M. (2004). Communication strategies and product line design. *Marketing Science*, 23(3):304–316.
- Vrontas, 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>.
- 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

A Proofs

A.1 Proof of Lemma 1

Proof. If the publisher sells exclusively through 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 $1 - \beta$ 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 is

$$\pi_P^{\text{OX-only}} = (1 - \beta) (1 - (1 - \mu)^2) R^{\text{OX-only}} = (2 - \mu)\mu(1 - \beta)^2 \quad (8)$$

■

A.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 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 is thus

$$\pi_P^{\text{PX-only}} = \mathbb{P}\{\text{legitimate}\} \mathbb{P}\{\nu_C = 1\} \cdot 1 = (1 - \beta)\mu \quad (9)$$

■

A.3 Proof of Lemma 3

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 last condition ensures that the C -advertiser does not deviate from bidding in PX, and the publisher does not deviate from selling through both PX and OX simultaneously.

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

$$\pi_U(R) = (1 - \beta)(1 - \mu)(1 - R) + \beta(-R), \quad (10)$$

where $1 - \beta$ is the probability that the impression is legitimate, $1 - \mu$ the probability that the C -advertiser's valuation is low (and therefore, the U -advertiser wins), $1 - R$ is the U -advertiser's payoff if it wins the legitimate impression, β is the probability that the impression

is fake, and $-R$ is the U -advertiser's payoff of winning a fake impression (note that the U -advertiser always wins the 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

$$R^{\text{OX}} = \max\{R : \pi_U(R) \geq 0\} = \frac{(1-\beta)(1-\mu)}{1-(1-\beta)\mu}. \quad (11)$$

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 above, which simplifies to

$$\pi_C^{\text{PX}} \geq \max_{b_C^{\text{OX}} \geq R^{\text{OX}}} \pi_C^{\text{OX}} \Leftrightarrow (1-\beta)(1-R^{\text{PX}}) \geq 1-\beta-R^{\text{OX}} \Leftrightarrow R^{\text{PX}} \leq \frac{R^{\text{OX}}}{1-\beta}. \quad (12)$$

Using the optimal OX reserve price (11), we obtain the optimal reserve price in PX:

$$R^{\text{PX}} = \frac{R^{\text{OX}}}{1-\beta} = \frac{1-\mu}{1-(1-\beta)\mu}. \quad (13)$$

We check the three necessary conditions above. Individual rationality is satisfied because $\pi_U(R^{\text{OX}}) \geq 0$ due to (11), and $\pi_C(R^{\text{PX}}) \geq 0$ due to (12); incentive compatibility holds by construction of R^{PX} ; and finally, for the publisher's impression, the C -advertiser's bid, which equals (13) is higher than the U -advertiser's, which equals (11).

The C -advertiser's expected profit is

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

and the publisher's expected profit is

$$\pi_P^{\text{PX-OX}} = (1 - \beta) (\mu R^{\text{PX}} + \mu(1 - \mu)R^{\text{OX}}) = \frac{(1 - \beta)(2 - \mu - \beta(1 - \mu))(1 - \mu)\mu}{1 - (1 - \beta)\mu}. \quad (14)$$

■

A.4 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 $\mu > \frac{1}{2}$.

■

A.5 Proof of Proposition 2

Proof. The result follows from $\frac{d^2}{d\beta^2} \pi_C^{\text{PX-OX}} = -\frac{2(1-\mu)\mu}{(1-(1-\beta)\mu)^3} < 0$ and $\frac{d}{d\beta} \pi_C^{\text{PX-OX}} = 0 \Leftrightarrow \beta = \frac{\sqrt{1-\mu-(1-\mu)}}{\mu}$.

■

A.6 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 (8), the OX-only regime yields $(2 - \mu)\mu(1 - \beta)^2$; from (9), the PX-only regime yields $(1 - \beta)\mu$; and from (14), 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)

$$\pi_P^{\text{OX-only}} \geq \pi_P^{\text{PX-only}} \Leftrightarrow 1 - \mu - \beta(2 - \mu) \geq 0 \Leftrightarrow \beta \leq \frac{1 - \mu}{2 - \mu}, \quad (15)$$

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 , $\text{LHS}|_{\mu=0} = -1$ and $\text{LHS}|_{\mu=1} = 1 - \beta$; therefore, the inequality simplifies to μ being greater than the root of LHS; therefore,

$$\mu \geq \frac{3 - 2\beta - \sqrt{4\beta^2 - 8\beta + 5}}{2(1 - \beta)}. \quad (16)$$

2. (PX-only)

$$\pi_P^{\text{PX-only}} \geq \pi_P^{\text{PX-OX}} \Leftrightarrow -(1 - \mu)^2 + \beta(1 - \mu + \mu^2) \geq 0 \Leftrightarrow \beta \geq \frac{(1 - \mu)^2}{1 - \mu + \mu^2}, \quad (17)$$

and

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

from the complement of (15).

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 (16), 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 (17). ■

A.7 Proof of Lemma 4

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

$$v = \begin{cases} \bar{v} & \text{w.p. } \mu, \\ 0 & \text{w.p. } 1 - \mu, \end{cases} \quad (18)$$

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

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 2 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 (18) satisfies regularity, as defined by Myerson (1981), because

$$0 - \frac{1 - F(0)}{f(0)} = 0 - \frac{\mu}{1 - \mu} < \bar{v} - \frac{1 - F(\bar{v})}{f(\bar{v})} = \bar{v}.$$

Therefore, it follows from Myerson (1981) that the publisher's optimal profit is achieved under a second-price auction with reserve price $\inf\{z \in \{0, \bar{v}\} : z - (1 - F(z))/f(z) \geq 0\} = \bar{v}$. 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.

A.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

$$\frac{1 - \beta}{1 - \beta + \beta(1 - \gamma)} \tag{19}$$

share of legitimate ad impressions coming through OX. Therefore, the advertisers' valuation is (19). Under this unilateral deviation, the reserve price is fixed at $1 - \beta$, which is less than the post-filter valuation (19). 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 $[1 - \beta, \bar{b}]$ according to distribution H where \bar{b} and H satisfy

the following indifference conditions:

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

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

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

The OX's expected profit in this mixed strategy equilibrium is

$$\begin{aligned}
\pi_{\text{OX}} &= (1 - \mu)^2 \cdot 0 + 2\mu(1 - \mu) \int_{1 - \beta}^{\bar{b}} b_j dH(b_j) + \mu^2 \int \int_{1 - \beta}^{\bar{b}} \max[b_C, b_U] dH(b_C) dH(b_U) \\
&= (1 - \beta)\mu(2 - \mu - 2\beta\gamma(1 - \mu)), \tag{20}
\end{aligned}$$

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

A.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 3 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 3, 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 positive to zero. \blacksquare

A.10 Statement and Proof of Claim 1

Claim 1. *If the C -advertiser receives two request-for-bids, one from PX and another from OX, bidding in both exchanges is weakly dominated by bidding in only PX.*

Proof of Claim 1. 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 competing 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 than bidding $\max[b^{\text{OX}}, b^{\text{PX}}]$ in PX only. Consider the C -advertiser's profit if it bids in both exchanges.

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

In sum,

$$\pi_C = \begin{cases} 0 & \text{if } \max [b^{\text{OX}}, b^{\text{PX}}] < b_U, \\ (1 - \beta) (1 - \max [b^{\text{OX}}, b^{\text{PX}}]) - \beta \mathbb{I}_{\{b^{\text{OX}} > b_U\}} b^{\text{OX}} & \text{if } \max [b^{\text{OX}}, b^{\text{PX}}] > b_U, \\ \alpha_1 (1 - \beta) (1 - \max [b^{\text{OX}}, b^{\text{PX}}]) - \alpha_2 \beta \mathbb{I}_{\{b^{\text{OX}} = b_U\}} b^{\text{OX}} & \text{if } \max [b^{\text{OX}}, b^{\text{PX}}] = b_U. \end{cases} \quad (21)$$

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

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

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

A.11 Statement and Proof of Claim 2

Claim 2. For any $R \in [0, \bar{v}]$, $\pi_S(R, R) = \pi_I(R)$.

Proof of Claim 2. 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

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

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 $\pi_S(R, R) = \pi_I(R)$. ■

A.12 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 $\beta \leq \tilde{\beta} \equiv \frac{(1 - \mu)^2}{\mu^2 - \mu + 1}$ 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 4). 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 (20). The OX's profit with $\gamma = 0$ is its profit under the PX-OX regime, which is $(\beta\mu + (1 - \beta)(1 - \mu)\mu)R^{\text{OX}} = (\beta\mu + (1 - \beta)(1 - \mu)\mu) \frac{(1 - \beta)(1 - \mu)}{1 - (1 - \beta)\mu} = (1 - \beta)(1 - \mu)\mu$. We obtain

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

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 3, the publisher's best response to γ is to sell through both PX and OX simultaneously at reserve prices

$$R^{\text{OX}} = \frac{(1 - \beta')(1 - \mu)}{1 - (1 - \beta')\mu} \quad (22)$$

and

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

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 $\gamma \in (0, \tilde{\gamma})$:

$$\pi_{\text{OX}} = (\beta(1 - \gamma)\mu + (1 - \beta)\mu(1 - \mu)) R^{\text{OX}} = (1 - \beta)(1 - \mu)\mu. \quad (23)$$

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 $[R^{\text{OX}}, \bar{b}]$, where the mixing distribution G and the upper bound of the support \bar{b}' are determined by the following indifference conditions for the j -advertiser, $j \in \{C, U\}$:

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

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

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

Therefore, the advertisers' expected profits in OX is

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

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

$$\pi_C^{\text{PX}} = (1 - \beta)(1 - R^{\text{PX}}) = \frac{(1 - \beta)\beta(1 - \gamma)\mu}{1 - \beta(\gamma - \mu) - \mu}.$$

It can be shown that $\pi_j^{\text{OX}} \geq \pi_C^{\text{PX}} \Leftrightarrow \gamma' \geq \left(\frac{\mu}{1 - \mu}\right)^2$. 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., $\gamma' = 0$).

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

where the last equality follows from (23). We obtain

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

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 4). 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 3, 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 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 OX (see Proposition 3). Following Lemma 5, OX has no incentive to filter further fake impressions. This completes the proof. ■

A.13 Statement and Proof of Claim 4

Claim 4. 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 $\frac{\beta+\sqrt{(4-3\beta)\beta}-2}{2(\beta-1)^2\sqrt{(4-3\beta)\beta}} \propto \beta + \sqrt{(4-3\beta)\beta} - 2$ and $\beta + \sqrt{(4-3\beta)\beta} - 2 \geq 0 \Leftrightarrow 4(1-\beta)^2 \geq 0$, which is true.

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. ■