Marketing Science Institute Working Paper Series 2025

Report No. 25-143

Experimentation in Online Advertising

W. Jason Choi and Amin Sayedi

Experimentation in Online Advertising

W. Jason Choi Amin Sayedi*

November 2025

*Amin Sayedi (aminsa@uw.edu) is Professor of Marketing at the Foster School of Business, University of Washington, Seattle, WA 98195. W. Jason Choi (wjchoi@umd.edu) is Assistant Professor of Marketing at the Smith School of Business, University of Maryland, College Park, MD 20742. The authors would like to thank Ron Berman, Yuxin Chen, Preyas Desai, Elea Feit, Kinshuk Jerath, Jessie Liu, Fei Long, Wendy Moe, Woochoel Shin, Upender Subramanian, Kannan Srinivasan, Raphael Thomadsen, Michael Trusov, Robert Zeithammer, Dennis Zhang, Z. John Zhang, Bobby Zhou, participants at the 2025 Marketing Theory Seminar, 2025 Washington University Junior Faculty Forum in Marketing, 2025 Smith School Research Seminar Series for their valuable comments.

1 Introduction

Digital advertising expenditure in the United States reached $310 billion in 2024 and is projected to surpass $450 billion by 2028.[1] A central challenge in this space is measuring incrementality, the share of conversions *caused* by ad exposure (Blake et al., 2015). Randomized controlled trials (RCTs) implemented as field experiments remain the gold standard for isolating this effect: by comparing outcomes between exposed and holdout groups, RCTs disentangle ad-driven conversions from those that would have occurred organically. RCTs enable advertisers to identify high-return customer segments and calibrate bids and budgets with precision (Gordon et al., 2021). Without such evidence, firms risk overpaying for non-incremental conversions, misallocating spend, and eroding campaign ROI.[2]

However, incrementality measurement is fraught with challenges (Johnson, 2023). Small effect sizes require large samples to achieve statistical significance (Lewis and Rao, 2015), and ad exposure is confounded by external factors—such as algorithmic ad delivery (Braun and Schwartz, 2025) and competition among advertisers (Waisman et al., 2025b)—that are beyond the experimenter's control. Moreover, advertisers incur significant costs in setting up and running experiments. According to industry surveys, managers "often shy away" from experiments because they are seen as "costly" and "difficult to get right" (Campbell et al., 2022, Page 4). Firms pay implementation costs associated with coordinating across various stakeholders in the advertising ecosystem (Johnson et al., 2017a). Gluck (2011) estimates rigorous experiments to cost between $50,000 and $80,000.[3] Even after the experimentation has been set up, "performing lift tests can be expensive ... [as advertisers pay] for the ads that are delivered to the test group" (Ervasti, 2022) and advertisers pay "the opportunity cost

[1] www.statista.com/statistics/242552/digital-advertising-spending-in-the-us/

[2] Practitioners also use observational approaches such as Marketing Mix Models (MMMs), which offer aggregate, channel-level insights but often obscure the contributions of individual campaigns, keywords, or audience segments; see https://funnel.io/blog/mta-vs-mmm.

[3] In part due to exorbitant costs, "the industry has resigned itself to accepting a sub-par research methodology in exchange for lower cost and efficiency" (Gluck, 2011, Page 9).

of withholding ads" (Mackinnon, 2024). Overall, these anecdotal evidence suggest that (a) operational costs involved in setting up experiments, (b) price of ads shown to the treatment group, and (c) opportunity cost from holdout are important consideration factors for ad experimentation.

Recognizing these challenges, several advertising platforms have introduced tools to lower the frictions of experimental ad measurement (Gordon et al., 2021). For instance, Meta Ads Manager provides user-friendly infrastructure and detailed documentation to help advertisers split holdout and test groups.[4] Such tools, however, are not offered universally. While Google facilitates experimentation for its display advertising network, comparable tools are not available for search advertising.[5] Platforms have at times restricted, and at other times expanded, the experimentation capabilities they offer, suggesting that the ad platforms' decisions to offer experimentation tools are made strategically.[6]

This research examines the strategic interactions between an advertising platform and advertisers in a setting where learning ad incrementality requires costly experimentation. Complementing the empirical literature on ad measurement that study optimal experimental designs conditional on firms experimenting, we study the economics of experimentation conditional on the availability of optimal experimental designs. Specifically, we seek to answer the following questions:

1. When should advertisers experiment? How do ad competition and ad prices influence advertisers' incentives to experiment?

[4] For example, see https://developers.facebook.com/docs/marketing-api/guides/lift-studies and www.facebook.com/business/news/suite-of-truth.

[5] Google states "You can use Conversion Lift based on users for Video, Discovery, and Demand Gen campaigns. If you're interested in using Conversion Lift based on users for Display, Search, Shopping, or Performance Max campaigns, reach out to your Google account representative for more information" (https://support.google.com/google-ads/answer/12005564). Similarly, Facebook limits experimentation with campaign requirements: "to run a Conversion Lift test, your ad account must have a campaign that started in the past year with a spend of $5,000 USD or more. In addition, this campaign needs to have a minimum of 500 conversions with a 1-day click, 7-day click or 1-day view attribution setting" (www.facebook.com/business/help/221353413010930).

[6] For example, Meta disabled holdout test in 2021 after introducing them in 2017 (www.facebook.com/business/help/169589021376481).

2. How do advertisers' experimentation decisions impact the platform's ad pricing strategy and profitability? Under what conditions should the platform facilitate or deter experimentation?

To address these questions, we develop a game-theoretic model with two advertisers competing for impressions on an advertising platform. Advertisers can conduct costly experiments to estimate incrementality. The platform sets the reserve price and decides whether to facilitate or deter experimentation. A key model feature is that experimentation involves dual costs: (a) the fixed cost of setting up experiments, and (b) the opportunity cost from experimental holdout (Feit and Berman, 2019; Gordon et al., 2021; Johnson et al., 2017a). The fixed cost can be reduced by the platform (e.g., by offering free experimentation tools), whereas the opportunity cost (i.e., foregone revenue from holdout) cannot be reduced. Importantly, the opportunity cost from holdout is borne *jointly* by the platform: holdout softens bidding competition, which lowers the platform's ad revenue.

Our analysis yields several key insights. First, we show that the advertisers' incentive to experiment is non-monotonic in the platform's reserve price. Experimentation incentive first increases and then decreases in the reserve price. As the reserve price increases, experimentation becomes more valuable because it helps advertisers avoid overpaying for non-incremental impressions. As the reserve price increases further, however, the advertisers' efficiency gains from learning incremental impressions are extracted away by the platform, thereby reducing the value of experimentation. This highlights a novel strategic role of ad platform's reserve prices in the context of endogenous experimentation beyond the standard surplus extraction role (Choi and Mela, 2023; Coey et al., 2021; Ostrovsky and Schwarz, 2023).

Second, some advertisers do not conduct experiments, even when the informational value of experimentation is large and there are no experimental fixed costs. The intuition revolves around the relationship between ad price and opportunity cost. The ad platform's revenue decreases from experimental holdout which softens bidding competition. Therefore, the platform may deter advertiser experimentation, and it achieves this by distorting the reserve price. A low reserve price inflates advertisers' opportunity cost of holdout, as advertisers forego potentially valuable ad impressions that otherwise could have been bought at low prices. Consequently, two ex ante symmetric advertisers may adopt asymmetric strategies wherein one advertiser experiments and the other does not. This strategic perspective of experimentation, combined with the opportunity costs of conducting experiments, provides a rationalization of the low experimentation rates observed in practice (Runge and Skokan, 2025; Wernerfelt et al., 2025).[7]

Third, a platform's stance on experimentation is determined by (a) the expected incrementality and (b) the fixed cost of experimentation incurred by advertisers. When the expected incrementality of ad impressions is moderate, the platform's decision of whether to facilitate or deter experimentation depends crucially on the fixed cost of experimentation. If fixed costs are high, the platform deters experimentation. High fixed costs diminish the appeal of experimentation for advertisers; therefore, the platform can deter experimentation and avoid opportunity costs from holdout without signigicantly lowering the reserve price. This allows the platform to extract large surplus from non-incremental impressions and save on opportunity costs.

The main contribution of our paper is to highlight the understudied strategic interactions between advertisers and the platform in an environment where advertisers must actively learn incrementality through costly experimentation (see Figure 1). Previous theoretical studies on information disclosure and targeting in online advertising assume that if the platform enables fine-grained targeting, advertisers automatically and costlessly learn ad effectiveness for every consumer segment (e.g., Bergemann and Bonatti, 2011; Choi and Sayedi, 2024;

[7]According to Runge et al. (2020), 1.4%—29.4% of firms in different industries have run at least one ad experiment on Meta in the one-year observation period. Moreover, Rao and Simonov (2019) show that the publication of "negative information on brand search ad effectiveness" by Blake et al. (2015) did not increase firms' propensity to conduct ad experiments, even though "[T]he ease of experimentation was aptly summarized in the popular press and held up as a contrast to the staggering value it offered eBay" (Rao and Simonov, 2019, Page 107). See Runge (2020) and Feger (2025) for more discussion on the under-usage of experimentation.

Hummel and McAfee, 2016; Levin and Milgrom, 2010; Rafieian and Yoganarasimhan, 2021). In practice, however, advertisers must conduct experiments to learn this information—a process which entails costs for both the advertisers and the platform. By explicitly modeling advertisers as learning agents who decide whether to experiment, our paper reveals novel economic forces—distinct from market-thinning effect—that influence the platform's ad pricing strategies and the advertisers' experimentation incentives. Notably, in contrast to previous literature (Hummel and McAfee, 2016), we show that the platform is more inclined to promote advertiser knowledge about ad effectiveness when the expected incrementality (equivalent to valuation in the information disclosure literature) is low, as advertiser experimentation in these cases imposes lower costs on the platform.

Related literature

Our paper contributes to two research streams. First, it complements the rich empirical literature on ad measurement (e.g., Choi et al., 2020; Johnson, 2013; Johnson et al., 2017b; Lewis et al., 2015; Sahni, 2015, 2016; Sahni and Zhang, 2024; Sahni et al., 2017; Waisman and Gordon, 2025; Waisman et al., 2025a; Yang et al., 2024; Zantedeschi et al., 2017). A central insight from this literature is that observational measures of ad effectiveness are imperfect substitutes for causal estimates obtained through experiments (Blake et al., 2015). Despite the abundance of granular data, observational methods fail to replicate causal estimates that are derived from exogenous variation (Gordon et al., 2023a, 2019). Additionally, prior research highlights the challenges of conducting randomized experiments in online advertising (Johnson, 2023), including algorithmic ad delivery (Boegershausen et al., 2025; Johnson et al., 2017a) to privacy regulations (Grosso and Runge, 2025) to competitive interference (Waisman et al., 2025b).[8] To address these concerns, innovative solutions such as ghost ads (Johnson et al., 2017a), causal machine learning methodologies (Chernozhukov et al.,

[8]A related issue is one-sided non-compliance: the lack of control advertisers have in ensuring that the consumer assigned to the treatment group actually receives the treatment (Lemmens et al., 2025).

2025; Wager and Athey, 2018), and predictive modeling leveraging prior experimental data (Gordon et al., 2023b) have been proposed. Our theoretical model builds on these empirical findings, generating novel insights and empirically testable hypotheses.

Second, our work contributes to the theoretical literature on online advertising that explores information disclosure policies. Choi and Sayedi (2019) examine a setting where ad valuation is common knowledge but advertisers learn their *quality scores* by winning ad auctions; they show that the endogenous learning process distorts bidding strategies. Our study differs by focusing on how advertisers learn their *incrementality* through endogenous experimentation, which entails both fixed costs and opportunity costs. Berman and Heller (2025) analyzes a "Naive Analytics Equilibrium," in which firms strategically avoid unbiased estimation of ad effectiveness to induce competitors to overspend on advertising, thereby expanding category demand. While our model yields a similar outcome (i.e., that not all advertisers experiment to improve ad measurement accuracy), the driving mechanism is fundamentally different. In our setting, advertisers refrain from experimentation not to influence competitors' spending, but in response to the ad platform's pricing strategy and the experimentation behavior of competing advertisers.

Prior research has explored how information disclosure affects auctions bids and revenues (e.g., Arefeva and Meng, 2021; Board, 2009; Milgrom and Weber, 1982). An important finding in this literature is the market-thinning effect, where more granular targeting reduces the number of participants in ad auctions, and ultimately lowers publisher revenue (Bergemann and Bonatti, 2011; Levin and Milgrom, 2010; Rafieian and Yoganarasimhan, 2021). Hummel and McAfee (2016) show that efficiency gains from reserve price optimization can counteract this market-thinning effect such that the ad seller's profit may be higher when information is disclosed than when it is withheld. Choi and Sayedi (2024) and D'Annunzio and Russo (2024) analyze information disclosure strategies in the presence of strategic intermediaries. A common assumption across these studies is that advertisers costlessly learn their ad valuations given the publisher's information disclosure. In contrast, our paper explicitly models the costly learning process through experimentation and examines the platform's strategic role in facilitating or deterring the endogenous learning process.

The remainder of the paper is structured as follows. Section 2 introduces the main model. Section 3 first analyzes a benchmark model comparing scenarios where ad incrementality are exogenously known versus unknown. We then investigate the full model, demonstrating how endogenizing the learning process through costly experimentation creates novel economic forces. Section 4 establishes robustness of core insights by studying various extensions of the main model. Finally, Section 5 discusses limitations and future directions, and then concludes.

2 Model

The game consists of an ad platform and two advertisers. The platform, denoted by P, sells ad impressions via second-price auction across two periods $t \in \{1, 2\}$. The platform sets a reserve price, denoted by R, and decides whether to reduce advertisers' fixed experimentation cost, denoted by $k > 0$. The reserve price is the minimum amount advertisers need to bid to participate in the auction. We assume that the same reserve price R is applied across both periods.[9] The platform's second decision variable of reducing advertisers' experimentation cost will be discussed in more detail below. The advertisers, indexed by $j \in \{A, B\}$, decide whether to experiment, incurring a fixed cost and opportunity cost due to holdouts. Advertisers then bid for ad impressions on the platform. We assume that each ad impression is associated with a consumer action such as clicks, visits, and purchases. We assume that both the advertisers' bids and their payments are *per action*; e.g., the advertiser submits a per-click bid, and pays only if the ad is clicked. Each action yields to the advertiser a payoff $v > 0$, which we normalize to 1. Only a fraction $\alpha_j \in [0, 1]$ of actions are *incremental*

[9] This reflects the notion that platforms have limited capacity to condition reserve prices on advertisers' experimentation decisions.

for Advertiser j. In other words, $1-\alpha_j$ actions would have occurred even without the ad, whereas α_j are incremental actions caused by the ad. For example, if the action is clicking on the ad to visit the advertiser's product page, $1-\alpha_j$ fraction of the customers who click on the ad would have visited the advertiser's product page even without the ad exposure (see Figure 2). In the main model, we focus on the simplest binary case where α_j is distributed

$$\mathbb{P}\{\alpha_j = 1\} = 1 - \mathbb{P}\{\alpha_j = 0\} = \mu \in [0, 1]. \tag{1}$$

In Section 4.1, we consider a continuous distribution of α and show that the insights continue to hold.

An important model feature is that while μ is common knowledge, advertisers do not *a priori* know the value of α_j for a given *target* (e.g., a customer segment or context).[10] To learn α_j, advertisers can conduct randomized experiments. In practice, estimating the causal effect of advertising involves comparing conversion rates between a holdout group and an ad-exposed group across various targets. These estimates then allow advertisers to predict the incremental impact of future ad impressions for each target.

Running experiments involves two types of costs: fixed cost and opportunity cost. The fixed cost k encompasses various operational costs ranging from computing costs to labor costs of setting up and running experiments.[11] The opportunity cost arises endogenously from holding out a certain fraction of impressions as control; the opportunity cost, unlike the fixed cost, cannot be reduced by the platform. To sharpen insights, we model opportunity cost as the experimenting advertiser foregoing its entire Period 1 payoff. In Section 4.2, we consider an extension with an interior holdout-treatment split such that an experimenting advertiser may earn positive Period 1 payoff from the treatment group, and show that the

[10]The prior μ can be interpreted as the information advertisers have obtained through their own analytics (e.g., marketing mix models) prior to experimentation.

[11]For example, Johnson et al. (2017a) discuss the high cost of running experiments with public service announcements for the control group as they "require coordination among advertisers, publishers, and third-party charities" (Page 867). Furthermore, Gluck (2011) estimates that 8–15 hours of staff training are needed to conduct rigorous experiments.

qualitative insights carry over.[12]

We denote Advertiser j's experimentation decision in Period 1 by the indicator variable

$$e_j = \mathbb{I}_{\{\text{Advertiser } j \text{ experiments}\}}. \tag{2}$$

If $e_j = 0$, then Advertiser j does not learn (before Period 2) the true incrementality α_j and simply strategizes based on the prior $\mathbb{E}[\alpha_j] = \mu$. If $e_j = 1$, then Advertiser j learns whether α_j is 0 or 1, and strategizes accordingly in Period 2. In Section 4.3, we consider continuous experimentation decisions where the precision with which advertisers learn incrementality increases proportionately with the holdout size.

We describe the players' payoffs. Suppose Advertiser j decides e_j in Period 1, and then bids b_{j1} and b_{j2} in Periods 1 and 2, respectively. Advertiser j's net expected profit is a weighted sum of payoffs from Periods 1 and 2; i.e.,

$$\pi_j = (1 - \delta)\pi_{j1} + \delta\pi_{j2}, \tag{3}$$

where $\delta \in [0, 1]$ is a weight parameter, and π_{jt} is Advertiser j's expected payoff in Period t.

Period 1 payoff is

$$\pi_{j1} = \begin{cases} (\mu - p)\,\mathbb{I}_{\{b_{j1} \geq P1\}} & \text{if } e_j = 0, \\ -k & \text{if } e_j = 1, \end{cases} \tag{4}$$

where

$$p_1 \equiv \max\left\{R, \mathbb{I}_{\{e_{-j}=0\}} b_{-j1}\right\}, \tag{5}$$

which represents the winning advertiser's payment in a second-price auction. The first branch

[12] The holdout setup in our model is consistent with ghost ads technology (Johnson et al., 2017a): experimenting advertisers do not need to pay to display neutral messages (e.g., public service announcements) to the holdout group. Their ads are suppressed from the audience who otherwise would have been eligible to see their ads.

of (4) represents Advertiser j's opportunity cost of experimentation; i.e., the Period 1 payoff that would be foregone should Advertiser j experiment. Observe that this opportunity cost, through the Period 1 payment p_1 in (5), depends on the platform's reserve price R and the competing advertiser's experimentation decision e_{-j}. The second branch of (4) represents the fixed cost of experimentation incurred in Period 1. As discussed above, we assume that an experimenting advertiser foregoes all Period 1 payoffs; we later relax this assumption in Section 4.2.

Note that both forms of experimentation cost appear similar on the surface: they are both incurred in Period 1 and they are incurred only if the advertiser experiments. However, there are two important distinctions. First, only the opportunity cost depends endogenously on the platform's reserve price; second, only the fixed cost can be reduced by the platform's actions (e.g., offering experimentation tools). These subtle distinctions will play a key role in shaping the equilibrium of the experimentation game.

In Period 2, Advertiser j conditions its bid b_{j2} on its information set, which would be more refined had Advertiser j experimented in Period 1 than if it had not. Specifically, if $e_j = 1$, then Advertiser j can selectively bid for impressions that generate incremental conversions (valued at 1) and avoid impressions that do not (valued at 0). In contrast, if $e_j = 0$, then it cannot bid selectively as such; it bids some fixed amount across all impressions, whether they generate incremental conversions or not. Taken together, Advertiser j's expected payoff in Period 2 is

$$\pi_{j2} = \mathbb{I}_{\{b_{j2}(e_j) \geq P2\}} \cdot \begin{cases} (\mu - p_2) & \text{if } e_j = 0, \\ \mu\,(1 - p_2) & \text{if } e_j = 1, \end{cases} \tag{6}$$

where

$$p_2 \equiv \max\left\{R, b_{-j2}(e_{-j})\right\}. \tag{7}$$

We illustrate the key trade-off that advertisers face in the experimentation context. Suppose Advertiser A experiments whereas Advertiser B does not (i.e., $e_A = 1$ and $e_B = 0$). Advertiser A's net payoff in Period 1 is negative due to the fixed cost payment (i.e., $-k$); however, it may earn a high expected payoff in Period 2 from improved bidding efficiency due to incrementality learning. Advertiser B's net payoff in Period 1 may be positive as it does not pay the experimentation cost nor withhold any Period 1 advertising. By not experimenting, however, Advertiser B's Period 2 payoff may be low because Advertiser B optimizes based on a coarser information set compared to its experimenting counterpart's.

Finally, the platform's payoff is

$$\pi_P = (1-\delta)\pi_{P1} + \delta\pi_{P2}, \tag{8}$$

where

$$\pi_{P1} = \mathbb{I}_{\{\max\{b_{j1},b_{-j1}\}\geq R,\, e_j=e_{-j}=0\}} \max\{R, \min\{b_{j1}, b_{-j1}\}\} + \sum_{j\in\{1,2\}} \mathbb{I}_{\{b_{-j1}\geq R,\, e_j>e_{-j}\}} R, \tag{9}$$

and

$$\pi_{P2} = \mathbb{I}_{\{\max\{b_{j2},b_{-j2}\}\geq R\}} \max\{R, \min\{b_{j2}, b_{-j2}\}\}, \tag{10}$$

per the standard second-price auction payment rules.

The platform's Period 1 payoff in (9) reveals that the platform *also* incurs opportunity cost of advertisers' experimentation. For example, if $e_j = e_{-j} = 1$, then the platform's Period 1 ad revenue is low because advertisers' holding out softens bidding competition. As we later demonstrate, this joint bearing of opportunity cost will constitute another key force that shape the platform's equilibrium strategy.

The game timing is as follows:

0. The platform sets the reserve price R.

1. Advertisers decide whether to experiment $e_j \in \{0, 1\}$. If Advertiser j experiments (i.e., $e_j = 1$), then it incurs fixed cost of experimentation k. Advertisers submit Period 1 bids and Period 1 payoffs are realized.

2. Experimenting advertisers learn incrementality, after which advertisers submit Period 2 bids; Period 2 payoffs are realized.

3 Analysis

We solve for the subgame perfect Nash equilibrium via backward induction. In Period 2, Advertiser j's expected ad incrementality depends on whether it experiments in Period 1:

$$\mathbb{E}[\alpha_j | e_j] = \begin{cases} \mu & \text{if } e_j = 0, \\ \alpha_j & \text{if } e_j = 1. \end{cases} \tag{11}$$

Because Period 1 auction outcome does not impact Period 2 auction, advertisers' weakly dominant strategy in each period is to bid truthfully. Thus, the advertisers' bids in Periods 1 and 2, respectively, are μ and (11). Their expected payoffs in Periods 1 and 2, respectively, are

$$\pi_{j1} = \begin{cases} (\mu - R)^+ & \text{if } e_j = 0 \text{ and } e_{-j} = 1, \\ 0 & \text{otherwise,} \end{cases} \tag{12}$$

and

$$\pi_{j2} = \begin{cases} 0 & \text{if } e_j = e_{-j} = 0, \\ (1-\mu)(\mu - R)^+ & \text{if } e_j = 0, e_{-j} = 1, \\ \mu\left(1 - \max\{R, \mu\}\right) & \text{if } e_j = 1, e_{-j} = 0, \\ \mu(1-\mu)(1-R) & \text{if } e_j = e_{-j} = 1. \end{cases} \tag{13}$$

We compare the advertisers' payoffs across different experimentation conditions to separately examine the opportunity cost of experimentation (realized in Period 1) and the returns from experimentation (realized in Period 2).

Advertiser j's Period 1 opportunity cost from experimenting is

$$\pi_{j1}(e_j = 0) - \pi_{j1}(e_j = 1) = e_{-j}(\mu - R)^+. \tag{14}$$

The expression in (14) reveals an important preliminary insight. The competitor's experimenting decreases the opportunity cost of experimentation, and the opportunity cost weakly decreases in the reserve price. If the competitor experiments, then it withdraws from ad auctions for its holdout group, softening Period 1 bidding competition. This effectively reduces the ad price and makes missing out on impressions that potentially generate incremental conversions more costly (compare dashed lines across Figures 3a and 3b). Moreover, the higher the reserve price, the more expensive are the ad impressions, and thus the lower the opportunity cost (see weakly decreasing dashed lines in Figures 3a and 3b). Therefore, all else equal, the advertisers' incentive to experiment decreases if the competing advertiser experiments and increases in R.

Advertiser j's Period 2 returns from experimenting can be derived from (13):

$$\pi_{j2}(e_j = 1) - \pi_{j2}(e_j = 0) = \begin{cases} \mu(1 - \max\{R, \mu\}) & \text{if } e_{-j} = 0, \\ (1 - \mu)\min\{(1 - \mu)R, \mu(1 - R)\} & \text{if } e_{-j} = 1. \end{cases} \tag{15}$$

This Period 2 gain is positive for any competitor experimentation decision $e_{-j} \in \{0, 1\}$ (see solid lines above 0 in Figures 3a and 3b), and is greater when the competitor does not experiment than when it does (compare solid lines across Figures 3a and 3b).

A key takeaway from this preliminary stage of the analysis is that the returns and opportunity cost of running experiments depend crucially on the platform's reserve price as well as on the competitor's experimentation decision. Interestingly, this endogenous relationship will create situations where some advertisers do not experiment even if the fixed cost of running experiments is negligible. We summarize these findings in the following lemma.

Lemma 1. *All else equal, the opportunity cost of experimentation*

- *is lower if the competitor experiments than if it does not; and*
- *weakly decreases in R.*

The returns from experimentation

- *is lower if the competitor experiments than if it does not; and*
- *weakly increases in R if $R \leq \mu$ and then decreases in R if $\mu < R$.*

Lemma 1 reveals two important insights. First, advertisers should be acutely mindful of competitor strategies when deciding their experimentation policy. An advertiser's incentive to experiment decreases when competitors conduct their own experiments. The reason is two-fold. First, the return on experimentation decreases because of the asymmetric *ex post* surplus outcomes: if the competitor realizes a high incrementality for the same consumer segment as the focal advertiser, the advertising value is competed away, whereas if the competitor realizes a low incrementality, the surplus gain is capped by the reserve price. Second, the opportunity cost increases if the competitor experiments because as the competitor holds out, bidding competition softens in Period 1. This in turn makes the focal advertiser's holding out during experimentation more costly.[13]

The second insight from Lemma 1 concerns the novel strategic role of reserve prices beyond the traditional surplus extraction role. The ad platform's reserve price materially

[13]This economic argument based on opportunity costs complements the methodological argument put forth by Waisman et al. (2025b), who show that parallel experimentation by competing advertisers creates interference effects that complicate the focal advertiser's measurement of treatment effects.

impacts the advertisers' experimentation incentives. For example, as the reserve price increases, the returns from experimentation *increases* while the associated opportunity cost *decreases* (see Figure 3b for $R \leq \mu$). Return from experimentation increases because experimentation enables the advertiser to avoid overpaying for impressions associated with low incrementality. Further boosting advertiser's experimentation motive, the opportunity cost of experimentation decreases as reserve price increases. The reason is that even if the advertiser had foregone experimenting and not held out in Period 1, an increase in the reserve price implies lower Period 1 payoff.

In total, Lemma 1 highlights a critical link between ad price and advertisers' experimentation incentives. As we later demonstrate, this link creates important distortions in the equilibrium reserve price set by the ad platform.

3.1 Benchmark model

To elucidate the impact of endogenous learning through experimentation that we analyze in the main model, we first solve a benchmark model in which we compare two exogenous information schemes. In the first scheme (denoted by Ω), advertisers know incrementality, and in the second scheme (denoted by $\emptyset$), advertisers do not know incrementality.

Under the knowledge regime (Ω), advertisers truthfully bid their incremental value, which is normalized to 1. The platform's reserve price-optimized payoff is

$$\max_{R \in [0,1]} \pi_P^{\Omega}(R) = \max_{R \in [0,1]} \mu^2 1 + 2\mu(1-\mu)R = \mu(2-\mu). \tag{16}$$

Under the ignorance regime ($\emptyset$), advertisers always bid μ, the expected incrementality value, such that the platform's reserve price-optimized payoff is

$$\max_{R \in [0,1]} \pi_P^{\emptyset}(R) = \max_{R \in [0,1]} \mu \, \mathbb{I}_{\{R \leq \mu\}} = \mu. \tag{17}$$

This profit is never greater than $\pi_P^{\Omega *}$ in (16).[14] We thus obtain the following benchmark result.

Lemma 2. *The platform's expected profit is higher if advertisers know incrementality than if they do not know incrementality.*

Lemma 2 highlights the efficiency effect of advertisers' incrementality knowledge.[15] If advertisers know incrementality, impressions are more likely to be allocated to the advertiser who values them most, creating additional surplus that the platform can extract through reserve price optimization.[16] Lemma 2 then raises the question: under endogenous learning, will the platform always facilitate experimentation to capitalize on the efficiency effect? Surprisingly, we find that despite the efficiency gains the platform does not always facilitate experimentation, even if it is costless for it to do so.

3.2 Main: Endogenous experimentation

Consider the main scenario in which advertisers must conduct costly experimentation to learn ad incrementality. Recall from Lemma 1 that the two critical determinants of the returns and opportunity costs of experimentation are the platform's reserve price and the competing advertiser's experimentation decision. For instance, if the reserve price is low, then the returns may be low because non-incremental clicks are cheap anyway; simultaneously, the opportunity cost may be high because the missed ad opportunities (potentially generating incremental conversions) from holdout are cheap. Given these effects on the returns and costs of experimentation, the following proposition characterizes the advertisers' subgame

[14] Similarly, it can be shown that the platform's optimal profit under Ω is higher than that under partial knowledge, wherein only advertiser knows incrementality.

[15] Mathematically, the efficiency effect boils down to the Jensen's inequality: $\mathbb{E}[\max_j \alpha_j] \geq \max_j \mathbb{E}[\alpha_j]$.

[16] In advertising auction literature, the market-thinning effect refers to how increased targeting precision can reduce the number of competing advertisers for any given impression, potentially reducing platform revenue. In contrast to Hummel and McAfee (2016), we show that with optimal reserve prices, the market-thinning effect may be outweighed by efficiency gains. In Section 4.1, we analyze a continuous distribution of incrementality and induce the market-thinning effect.

equilibrium experimentation policies conditional on the reserve price.

Proposition 1 (Subgame Experimentation)**.** *Let the tuple* $(e_A, e_B) \in \{0,1\}^2$ *denote the advertisers' experimentation decisions. The subgame equilibrium is as follows:*

- *if* $k > \frac{\delta}{1-\delta}\mu(1-\max\{R,\mu\})$, *then* $(e_A, e_B) = (0,0)$*; i.e., neither advertisers experiment;*
- *if* $k \leq \frac{\delta}{1-\delta}\mu(1-\mu)^2$ *and* $\frac{(1-\delta)(\mu+k)}{1-\delta(2-\mu)\mu} < R \leq 1 - \frac{(1-\delta)k}{\delta(1-\mu)\mu}$, *then* $(e_A, e_B) = (1,1)$*; i.e., both advertisers experiment; and*
- *otherwise,* $(e_A, e_B) = (0,1)$*;*[17] *i.e., only one advertiser experiments.*

A high-level takeaway from Lemma 1 and Proposition 1 is that the reserve price is an important strategic lever the platform wields to influence experimentation dynamics. Figure 4 illustrates how reserve prices and experimentation costs jointly determine advertisers' experimentation strategies. For instance, if the fixed cost of experimentation is sufficiently low, then the platform can incentivize both advertisers to experiment by setting an intermediate value of R, and it can deter experimentation by setting more extreme values of R.

The advertisers' reaction to the platform's reserve price introduces important constraints on the platform's reserve price problem. In particular, the platform's optimum profit in the benchmark model, $\pi_P^{\Omega}(R^* = 1) = \mu(2-\mu)$ in (16), cannot be attained by simply setting $R = 1$ under endogenous advertiser experimentation. The reason is that a high reserve price discourages advertisers from experimenting as the returns from experimentation are low under high reserve price. This contrast against the benchmark result highlights the significance of considering the advertisers' experimentation motives in the platform's reserve optimization problem, to which we turn next.

An important implication of Proposition 1 is that when optimizing the reserve price, the platform should judiciously weigh the surplus extraction and the strategic implications for advertiser experimentation. Recall from Lemma 2 that the platform's profit is higher when

[17] For ease of exposition, we hereafter omit the "mirror" equilibrium $(e_A, e_B) = (1,0)$.

advertisers know incrementality than when they do not. Thus, one may predict that the platform obtains higher profit by facilitating advertiser experimentation than deterring it. As the following proposition shows, this is not always the case. Under certain conditions, the platform deters experimentation by lowering the reserve price, even if it could have raised the reserve price and promoted experimentation.

The following proposition summarizes the advertisers' experimentation policies and the platform's optimal reserve price, as a function of the model primitives; i.e., expected incrementality μ, discount factor δ, and the advertisers' fixed cost of experimentation k.

Proposition 2 (Equilibrium Experimentation Strategy). *Let $\tilde{k}$ and $\tilde{\delta}$ be as defined in the proof. The equilibrium experimentation strategies (e_A^*, e_B^*) and reserve price are as follows:*

- *if $k > \frac{\delta\mu(1-\mu)}{1-\delta}$, then $(e_A^*, e_B^*) = (0, 0)$ (i.e., neither advertiser experiments) and $R^* = \mu$;*
- *if $k \leq \min\left\{\frac{\delta\mu(1-\mu)^2}{1-\delta}, \tilde{k}\right\}$ and $\delta > \tilde{\delta}$, then $(e_A^*, e_B^*) = (1, 1)$ (i.e., both advertisers experiment) and $R^* = 1 - \frac{(1-\delta)k}{\delta\mu(1-\mu)}$;*
- *otherwise, $(e_A^*, e_B^*) = (0, 1)$ (i.e., only one advertiser experiments) and $R^* = \min\left\{\mu, \frac{(1-\delta)(\mu+k)}{1-\delta\mu(2-\mu)}\right\}$.*

If k is large, then the advertisers do not experiment for any reserve price. Therefore, advertisers hold same ad valuations equal to μ across both periods, and the platform sets the reserve price to this level.

Interestingly, if k is not large, then the platform uses the reserve price as an instrument to promote or deter experimentation (see gray and black regions in Figure 5). As the second part of the proposition shows, if k is small and δ large, then the platform induces both advertisers to experiment by setting a high reserve price (i.e., $R^* > \mu$). Recall from Lemma 2 that the efficiency effect incentivizes the platform to induce advertiser experimentation. However, because experimentation is costly for advertisers, the platform needs to subsidize experimentation by discounting reserve prices.[18] Therefore, the smaller is k, the less experimentation cost the platform internalizes. In this case, inducing experimentation becomes profitable for the platform: it can reap the efficiency gains with little subsidization. A large δ (i.e., post-experimentation payoff looms sufficiently large to the players relative to experimentation-stage payoff) is necessary for the platform to promote experimentation. The reason is that the platform jointly bears the advertisers' opportunity cost from holdout: as experimenting advertisers hold out, bidding competition softens and the ad revenue falls. This reduction in Period 1 ad revenue is mitigated by a large δ. Taken together, if k is small and δ large, the platform induces experimentation by raising the reserve price.

Finally, if either k is intermediate or δ small, the platform sets a low reserve price (i.e., $R^* \leq \mu$) and induces one advertiser to experiment while the other does not. This differentiation in advertiser experimentation (described in the third bullet of Proposition 2) reflects the platform's strategy shift from reaping efficiency gains via full-experimentation (i.e., $(e_A, e_B) = (1, 1)$) to saving on the first period opportunity cost via partial-experimentation (i.e., $(e_A, e_B) = (0, 1)$). This is because k is large enough that the fixed cost passed on from the advertisers to the platform caps the reserve price. This constrains the platform's capacity to monetize the efficiency effect. Simultaneously, k is small enough that one advertiser finds experimentation worthwhile. Neither advertiser has incentive to deviate from the asymmetric strategy profile because experimentation are strategic substitutes: returns from experimentation are high if the competitor is not experimenting, and low otherwise (see Figure 3). The platform's experimentation deterrence strategy also emerges if δ is small, because the opportunity cost of experimentation in Period 1 looms large for small δ.

Next, we consider the platform's ability to costlessly reduce advertisers' experimentation cost k by offering experimentation tools that facilitate experimentation. Would the platform always reduce fixed cost k to facilitate experimentation and thereby mitigate the subsidy

[18] In a frictionless world where learning occurred costlessly, the platform would simply set $R^* = 1$ and maximally capitalize on the efficiency gains. In the presence of experimentation costs, however, advertisers would not experiment under $R^* = 1$.

effect? Would the platform ever obtain higher profit from *not* reducing k than from facilitating experimentation? As the following proposition shows, even though the platform can costlessly reduce k to minimize the experimentation subsidy in the form of discounted reserve price, the platform does not always do so.

Proposition 3 (Equilibrium Experimentation). *Let $\underline{\mu}$, $\overline{\mu}$, and $\overline{k}$ be as defined in the proof.*

- *If (i) $\mu \leq \underline{\mu}$ or (ii) $\underline{\mu} < \mu \leq \overline{\mu}$ and $k \leq \overline{k}$, then the platform reduces $k \downarrow 0$ and induces $(e_A^*, e_B^*) = (1, 1)$; i.e., both advertisers experiment.*
- *Otherwise, the platform does not reduce k and induces either $(e_A^*, e_B^*) = (0, 1)$ or $(0, 0)$; i.e., at least one advertiser does not experiment.*

The opportunity cost of experimentation is a key driver of the platform's strategy of reducing experimentation cost. If μ is small (i.e., $\mu < \underline{\mu}$ in Figure 6), then the opportunity cost is low, and therefore, the platform reduces k to promote experimentation and reap the efficiency gains (see Lemma 2). In contrast, if μ is large (i.e., $\mu \geq \overline{\mu}$ in Figure 6), then the opportunity cost in the form of missed ad revenue during experimentation outweighs the efficiency effect. In this case, the platform does not facilitate experimentation and keeps advertisers' fixed cost as is. This insight underscores the key conceptual distinction between fixed cost and opportunity cost in our model. Even though the platform can completely eliminate the fixed cost of experimentation, it cannot reduce the opportunity cost arising from holdout. Therefore, the platform deters experimentation if μ is large.

The platform's strategy is more nuanced when the expected incrementality is intermediate (i.e., $\underline{\mu} < \mu \leq \overline{\mu}$ in Figure 6). As described in part (ii) of the first bullet in Proposition 3, the platform reduces k if and only if k is sufficiently small. The platform chooses between deterrence and promotion strategies: set a low reserve price to deter experimentation and earn high Period 1 payoff at the expense of Period 2 efficiency gains (which is large for intermediate μ), or set a high reserve price to promote experimentation to reap Period 2 efficiency gains at the expense of Period 1 payoff.

The profitability of these two options depends on the two types of cost: fixed cost k and opportunity cost (proxied by μ). If k is small, then experimentation is appealing to advertisers such that the deterrence strategy requires setting an unprofitably low reserve price. Therefore, the platform opts for the promotion strategy: maximize the efficiency gains by reducing $k \downarrow 0$ and raise the reserve price to extract surplus. On the other hand, if k is large, then experimentation is less appealing for advertisers such that the platform can deter experimentation without significantly lowering the reserve price. This allows the platform to obtain a relatively high Period 1 payoff from the non-experimenting advertiser. In this case, the platform opts for the deterrence strategy.

4 Extensions

We extend the baseline model in two directions to assess the robustness of our insights. The first extension allows for a continuous distribution of ad incrementality, the second explores an interior-split experimental design that is more closely aligned with practice, and the third relaxes the binary experimentation choice to a continuous experimentation strategy. We find that the core mechanisms identified in the main model carry over in all extensions. Furthermore, we highlight novel insights obtained from these extension analyses.

4.1 Continuous incrementality distribution

In the main model, ad incrementality was assumed to follow a binary distribution. This assumption simplified analysis and highlighted the efficiency-opportunity cost trade-off associated with experimentation. In this section, we analyze the case where incrementality is continuously distributed and demonstrate the robustness of the key insights.

To that end, assume ad incrementality α is distributed uniformly on $[0, 1]$. For example, if $\alpha = .01$, then 1% of the observed conversions associated with ad exposure are incremental, whereas 99% of them would have occurred even without exposure. Interestingly, we find that the preliminary efficiency result in Lemma 2 does not carry over under the uniform distribution assumption. That is, the platform's expected profit is higher when advertisers know incrementality than when they do not. This reversal arises due to the amplification of the market-thinning effect under a continuous support.

Lemma 3. *The platform's expected profit is higher when advertisers know incrementality than when they do not know.*

Overall, Lemma 3 implies that the platform has greater incentive to deter experimentation relative to the binary case. The reason is that the efficiency gains from advertiser learning are counterbalanced by thinner competition. However, we find that even when the marketing-thinning effect is operative, the qualitative insights from the main model continue to hold. We summarize the subgame experimentation results, holding reserve price fixed, in the following proposition (see Figure 7).

Proposition 4 (Subgame Experimentation)**.** *Let $\underline{R}$ and $\overline{R}$ be as defined in the proof. Let the tuple $(e_A, e_B) \in \{0, 1\}^2$ denote the advertisers' experimentation decisions. The subgame equilibrium is as follows:*

- *if $k > \frac{\delta}{2(1-\delta)} \left(1 - \max\left\{R, \frac{1}{2}\right\}\right)^2$, then $(e_A, e_B) = (0, 0)$; i.e., neither advertiser experiments;*
- *if $k \leq \frac{\delta}{12(1-\delta)}$ and $\underline{R} < R \leq \overline{R}$, then $(e_A, e_B) = (1, 1)$; i.e., both advertisers experiment; and*
- *otherwise, $(e_A, e_B) = (0, 1)$; i.e., only one advertiser experiments.*

The subgame equilibrium characterization closely mirrors that from the main model (compare Figures 4 and 7). The parallels in the subgame equilibrium carry over to the total equilibrium. Thus, we find that the core insights from the main model are robust to the incrementality distributional assumptions. The following proposition characterizes the equilibrium under the continuous incrementality distribution.

Proposition 5. *Let $\hat{k}$ and R_1 be as defined in the proof. The equilibrium experimentation strategies (e_A^*, e_B^*) and reserve price are as follows:*

- *if $k > \frac{\delta}{8(1-\delta)}$, then $(e_A^*, e_B^*) = (0, 0)$ (i.e., neither advertiser experiments) and $R^* = \frac{1}{2}$;*
- *if $k \leq \hat{k}$ and $\delta > \frac{12}{13}$, then $(e_A^*, e_B^*) = (1, 1)$ (i.e., both advertisers experiment) and $R^* = \frac{1}{2}$;*
- *otherwise, $(e_A^*, e_B^*) = (0, 1)$ (i.e., only one advertiser experiments) and $R^* = \min\left\{\frac{1}{2}, R_1\right\}$.*

When incrementality follows a continuous distribution, the model continues to exhibit the same strategic trade-offs. However, the market-thinning effect becomes more salient: advertiser learning thins competition, occasionally offsetting efficiency gains. Thus, the platform's incentive to deter experimentation strengthens relative to the binary case. Overall, the equilibrium mirrors the baseline model in that the non-experimentation, asymmetric experimentation, and full experimentation regions persist in qualitatively similar conditions.

4.2 Interior treatment-holdout split

In the main model, we make the stylized assumption that experimentation involves holding out the entire target audience. This assumption helped sharpen our insights by accentuating the opportunity cost of experimentation. However, one may wonder to what extent our insights generalize to a setting where, consistent with practice, experimentation involves some interior split between treatment and holdout fraction. For instance, according to Runge et al. (2020), the average holdout size of ad experiments on Meta is 16% with an average experimentation duration of 32 days. In this section, we examine the impact of an interior split on the equilibrium outcomes.

Suppose advertisers hold-out $h \in (0, 1)$ fraction of their target segment in Period 1 for experimentation. The model features of Period 2 remain unchanged. At a high level, the key difference this interior split creates compared to the main model is a reduction in the opportunity cost of experimentation. Because advertisers are no longer assumed to hold-out the whole target population, they miss out on fewer ad opportunities associated with potentially positive incrementality.

To simplify exposition, we assume $h = 1/2$; however, all insights from this section carry over for any interior level $h \in (0, 1)$.

To illustrate the main departure from the baseline model, we analyze Advertiser A's opportunity cost of experimentation. We first compute Advertiser A's expected Period 1 payoffs under different experimentation policies (see Table 1). If its competitor, Advertiser B, does not experiment (i.e., $e_B = 0$), then Advertiser B bids the expected incrementality μ for all of its target segment such that Advertiser A cannot obtain any positive surplus in Period 1, whether it experiments or not.

If the competing advertiser experiments (i.e., $e_B = 1$), then random half of ad opportunities in Period 1 are bid on by Advertiser B whereas the other half are held out. From the focal Advertiser A's perspective, if it does not experiment in Period 1, then it bids μ for all consumers and earns surplus $(\mu - R)^+$ from the segment held out by Advertiser B. If Advertiser A does experiment, then it earns the same weakly positive surplus only from the segment for which it "treats" and Advertiser B holds out. Taken together, conditional on $e_B = 1$, Advertiser A's expected Period 1 payoff is $(\mu - R)^+/2$ if $e_A = 0$ and $(\mu - R)^+/4$ if $e_A = 1$. Modulo fixed cost k, Advertiser A's opportunity cost of experimentation is thus 0 if $e_B = 0$ and $(\mu - R)^+/4$ if $e_B = 1$.

The latter opportunity cost is smaller than that from the main model, which was $(\mu - R)^+$ (see Figure 3b). Because the Period 2 returns from experimentation are the same as the main model, the reduced opportunity cost compared to the main model implies that advertisers will have greater incentive to experiment than in the main model, all else equal. This intuition is reflected in the enlarged parameter region in which both advertisers experiment compared to the main model (observe that the black parameter region in Figure 8a is larger than that in Figure 4). The following lemma formalizes the interior-split-induced moderation of the qualitative insight from the main model.

Lemma 4 (Subgame Experimentation)**.** *The parameter range in which both advertisers experiment is larger in the interior-split extension model than in the main model.*

For completeness, we also present the total equilibrium outcomes under the platform's optimal reserve prices. As summarized in the following proposition, the qualitative results from Proposition 2 carry over. For example, consistent with the results from Proposition 2, the platform induces neither advertiser to experiment if the fixed cost of experimentation is sufficiently high (compare Figures 5 and 8b). Moreover, it induces both to experiment if the fixed cost is low and the discount factor is sufficiently large (i.e., the advertisers and the platform weigh post-experimentation payoffs sufficiently high compared to experimentation-stage payoffs).

Proposition 6 (Equilibrium Experimentation Strategy)**.** *Let $\tilde{k}'$ and $\tilde{\delta}'$ be as defined in the proof. The equilibrium experimentation strategies (e_A^*, e_B^*) and reserve price are as follows:*

- *if $k > \frac{\delta\mu(1-\mu)}{1-\delta}$, then $(e_A^*, e_B^*) = (0,0)$ (i.e., neither advertiser experiments) and $R^* = \mu$;*
- *if $k \leq \tilde{k}'$ and $\delta > \tilde{\delta}'$, then $(e_A^*, e_B^*) = (1,1)$ (i.e., both advertisers experiment) and*

$$R^* = \begin{cases} \mu & \text{if } k > \frac{\mu(\delta(11-8(2-\mu)\mu)-3)^+}{8(1-\delta)} \text{ and } \delta \leq \frac{3}{11-8\mu(3-(3-\mu)\mu)} \\ 1 - \frac{(1-\delta)k}{\delta\mu(1-\mu)} & \text{otherwise;} \end{cases} \tag{18}$$

- *otherwise,* $(e_A^*, e_B^*) = (0,1)$ *(i.e., only one advertiser experiments) and*

$$R^* = \begin{cases} \frac{(1-\delta)(4k+\mu)}{1+\delta(3-4(2-\mu)\mu)} & \textit{if } k \leq \frac{\delta\mu(1-\mu)^2}{1-\delta} \\ \mu & \textit{otherwise.} \end{cases} \tag{19}$$

Allowing an interior holdout fraction reduces the opportunity cost of experimentation as advertisers no longer forgo all impressions during testing. Consequently, the region in which both advertisers experiment expands relative to the main model, while preserving the core comparative statics. The qualitative mechanisms (i.e., reserve price distortion and platform trade-offs between efficiency and opportunity cost) remain robust.

4.3 Continuous experimentation strategy

In this section, we extend the experimentation decision from binary (experiment or not) to a continuous choice: advertisers select a holdout fraction $h_j \in [0, 1]$, which determines the precision of their learning signal. Larger holdout sizes yield more accurate inference about α but entail higher opportunity costs in the first period.

We model the signal accuracy as

$$\rho_j = \mathbb{P}\{s_j = \alpha_j | \alpha_j\} = \frac{1}{2}(1 + h_j), \tag{20}$$

where $s_j \in \{0, 1\}$ is a signal of α_j that Advertiser j obtains from running the experiment. For instance, if Advertiser j does not hold out any sample (i.e., $h_j = 0$), then it is equivalent to not running any experiment and Advertiser j receives random noise. On the other hand, $h_j = 1$ corresponds to a perfect experiment that reveals the true incrementality.[19]

In Period 2, Advertiser j's expected ad incrementality depends on the level of experi-

[19]That experimental precision is maximized at $h_j = 1$ is a stylization. The functional form can be adjusted, without altering the qualitative insights, so that measurement precision is maximized at, e.g., $h_j = 1/2$.

mentation h_j set in Period 1. Specifically,

$$\mathbb{E}[\alpha_j|s_j] = 1 \cdot \mathbb{P}\{\alpha_j = 1|s_j\} + 0 \cdot \mathbb{P}\{\alpha_j = 0|s_j\} = \begin{cases} \frac{\rho_j\mu}{\rho_j\mu+(1-\rho_j)(1-\mu)} & \text{if } s_j = 1, \\ \frac{(1-\rho_j)\mu}{(1-\rho_j)\mu+\rho_j(1-\mu)} & \text{if } s_j = 0. \end{cases} \tag{21}$$

Consistent with the main model, advertisers' weakly dominant strategy in each period is to bid truthfully. Accordingly, their expected payoffs in Periods 1 and 2 are

$$\pi_{j1} = (1-h_j)h_{-j}\left(\mathbb{E}[\alpha_j] - R\right)^+, \tag{22}$$

$$\pi_{j2} = \begin{cases} \mathbb{P}\{s_j = 1\}\left(\mathbb{E}[\alpha_j|s_j = 1] - \mathbb{E}_{s_{-j}}\left[\max\{R, \mathbb{E}[\alpha_{-j}|s_{-j}]\}\right]\right)^+ & \text{if } h_j \geq h_{-j}, \\ \mathbb{P}\{s_{-j} = 0\}\mathbb{E}_{s_j}\left[\left(\mathbb{E}[\alpha_j|s_j] - \max\{R, \mathbb{E}[\alpha_{-j}|s_{-j} = 0]\}\right)^+\right] & \text{if } h_j < h_{-j}, \end{cases} \tag{23}$$

respectively.

We find that the equilibrium experimentation strategies collapse to corner solutions; i.e., $(h_1^*, h_2^*) \in \{0,1\}^2$. Therefore, the equilibrium results coincide with those from the main model with binary experimentation decisions. We state this result in the following lemma.

Lemma 5. *In the extension model with continuous experimentation strategy space, advertisers choose either no experimentation ($h = 0$) or full experimentation ($h = 1$) in equilibrium. In other words, the extension model is isomorphic to the binary baseline model.*

When experimentation becomes continuous, advertisers optimally choose corner solutions (i.e., either full or no experimentation). Partial experimentation does not survive in equilibrium because the marginal learning benefits are outweighed by opportunity costs at intermediate holdout levels. This result confirms that modeling experimentation as a binary choice in the main model entails little loss of generality.

5 Conclusion

This study examines the role of costly experimentation in online advertising markets, where advertisers seek to determine the incrementality of their campaigns through controlled experiments. We develop a game-theoretic framework modeling interactions between competing advertisers and an ad platform that sets reserve prices and decides whether to facilitate or discourage experimentation.

Our analysis yields several important insights about the strategic dynamics of experimentation in advertising markets. First, we find that advertisers do not always conduct experiments, even when the potential learning value is substantial. This is because opportunity costs—ad impressions foregone in experimental holdout groups—can outweigh the efficiency gains from learning true incrementality. This finding helps explain the puzzlingly low experimentation rates observed in practice, despite the widely acknowledged benefits of causal measurement.

Second, our analysis reveals that the platform's reserve price has non-monotonic effects on advertisers' experimentation incentives: experimentation value initially increases with reserve price as experimentation helps advertisers avoid overpaying for impressions that do not generate incremental conversions, but then declines as higher reserve prices erode advertisers' surplus from incremental impressions. This highlights a novel strategic function of reserve prices beyond their traditional surplus extraction role.

Third, we demonstrate that ad platforms may strategically deter experimentation by depressing reserve prices, despite the potential efficiency gains from advertiser learning. When expected incrementality is high, the opportunity costs of experimentation—reduced revenue from holdout groups—may outweigh efficiency gains, leading platforms to discourage experimentation. Conversely, when expected incrementality is low, platforms are more likely to facilitate experimentation as the opportunity costs are minimal.

These findings contribute to both the experimental measurement literature in digital advertising and the theoretical literature on information disclosure in online advertising markets. Unlike previous work that assumes advertisers automatically learn incrementality when platforms enable targeting, our model explicitly considers the costly experimentation process required for learning. This reveals novel economic forces—distinct from the market-thinning effect documented in prior research—that influence platform pricing strategies and advertiser experimentation incentives.

We acknowledge several limitations and suggest avenues for future research. Our model assumes a simplified two-symmetric-advertiser setting to capture competition parsimoniously; future research could extend this framework to investigate the impact of thicker markets on experimentation incentives. Additionally, while we have explored continuous incrementality distributions and continuous experimentation strategies separately, it would be interesting to examine the interaction between the two model features. Exploring these extensions would further enrich our understanding of the complex strategic interactions that shape experimentation in digital advertising markets.

Declarations

Funding and competing interests: All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no funding to report.

References

Arefeva, A. and Meng, D. (2021). Revealing information in auctions: The optimal auction versus the second-price auction. *Economics Letters*, 204:109895.

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. and Heller, Y. (2025). Naive analytics: The strategic advantage of algorithmic heuristics. *Games and Economic Behavior (Forthcoming)*.

Blake, T., Nosko, C., and Tadelis, S. (2015). Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. *Econometrica*, 83(1):155–174.

Board, S. (2009). Revealing information in auctions: the allocation effect. *Economic Theory*, 38(1):125–135.

Boegershausen, J., Cornil, Y., Yi, S., and Hardisty, D. J. (2025). On the persistent mischaracterization of google and facebook a/b tests: How to conduct and report online platform studies. *International Journal of Research in Marketing*.

Braun, M. and Schwartz, E. M. (2025). Where a/b testing goes wrong: How divergent delivery affects what online experiments cannot (and can) tell you about how customers respond to advertising. *Journal of Marketing*, 89(2):71–95.

Campbell, C., Runge, J., Bates, K., Haefele, S., and Jayaraman, N. (2022). It's time to close the experimentation gap in advertising: Confronting myths surrounding ad testing. *Business Horizons*, 65(4):437–446.

Chernozhukov, V., Demirer, M., Duflo, E., and Fernandez-Val, I. (2025). Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in india. *Econometrica (Forthcoming)*.

Choi, H. and Mela, C. F. (2023). Optimizing reserve prices in display advertising auctions. *Working Paper, University of Rochester*.

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

Choi, W. J. and Sayedi, A. (2024). Agency market power and information disclosure in online advertising. *Marketing Science*, 43(6):1279–1298.

Coey, D., Larsen, B. J., Sweeney, K., and Waisman, C. (2021). Scalable optimal online auctions. *Marketing Science*, 40(4):593–618.

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

Ervasti, M. (2022). "Why Calibrate Marketing Mix Modeling with Facebook Lift Tests?" (accessed October 8, 2025). https://sellforte.com/blog/why-should-you-calibrate-marketing-mix-modeling-results-with-facebook-lift-tests.

Feger, A. (2025). "The incrementality equation: Solving retail media's ROI challenge" (accessed October 9, 2025). www.emarketer.com/content/incrementality-equation-solving-retail-media-roi-challenge.

Feit, E. M. and Berman, R. (2019). Test & roll: Profit-maximizing a/b tests. *Marketing Science*, 38(6):1038–1058.

Gluck, M. (2011). Best practices for conducting online ad effectiveness research. *Technical Report, Internet Advertising Bureau.*

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

Gordon, B. R., Moakler, R., and Zettelmeyer, F. (2023a). Close enough? A large-scale exploration of non-experimental approaches to advertising measurement. *Marketing Science*, 42(4):768–793.

Gordon, B. R., Moakler, R., and Zettelmeyer, F. (2023b). Predictive Incrementality by Experimentation (PIE) for ad measurement. *Working paper, Northwestern University.*

Gordon, B. R., Zettelmeyer, F., Bhargava, N., and Chapsky, D. (2019). A comparison of approaches to advertising measurement: Evidence from big field experiments at facebook. *Marketing Science*, 38(2):193–225.

Grosso, B. and Runge, J. (2025). "The Uncomfortable Truth About Advertising Effectiveness: Why Marketers Avoid True Experimentation" (accessed April 2, 2025). www.adexchanger.com/data-driven-thinking/the-uncomfortable-truth-about-advertising-effectiveness-why-marketers-avoid-true-experimentation.

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

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

Johnson, G. A. (2023). Inferno: A guide to field experiments in online display advertising. *Journal of Economics and Management Strategy*, 32(3):469–490.

Johnson, G. A., Lewis, R. A., and Nubbemeyer, E. I. (2017a). Ghost ads: Improving the economics of measuring online ad effectiveness. *Journal of Marketing Research*, 54(6):867–884.

Johnson, G. A., Lewis, R. A., and Reiley, D. H. (2017b). When less is more: Data and power in advertising experiments. *Marketing Science*, 36(1):43–53.

Lemmens, A., Roos, J. M., Gabel, S., Ascarza, E., Bruno, H., Gordon, B. R., Israeli, A., Feit, E. M., Mela, C. F., and Netzer, O. (2025). Advancing personalization: How to experiment, learn and optimize. *Working Paper, Erasmus University*.

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

Lewis, R., Rao, J. M., and Reiley, D. H. (2015). *Measuring the Effects of Advertising*. University of Chicago Press.

Lewis, R. A. and Rao, J. M. (2015). The unfavorable economics of measuring the returns to advertising. *Quarterly Journal of Economics*, 130(4):1941–1973.

Mackinnon, S. (2024). "The 4 Reasons Why Marketers Are Hesitant About Incrementality" (accessed October 8, 2025). www.crealytics.com/blog/the-4-reasons-why-marketers-are-hesitant-about-incrementality.

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

Ostrovsky, M. and Schwarz, M. (2023). Reserve prices in internet advertising auctions: A field experiment. *Journal of Political Economy*, 131(12):3352–3376.

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

Rao, J. M. and Simonov, A. (2019). Firms' reactions to public information on business practices: The case of search advertising. *Quantitative Marketing and Economics*, 17(2):105–134.

Runge, J. (2020). "Marketers Underuse Ad Experiments. That's a Big Mistake." (accessed March 15, 2025). https://hbr.org/2020/10/marketers-underuse-ad-experiments-thats-a-big-mistake.

Runge, J., Geinitz, S., and Ejdemyr, S. (2020). Experimentation and performance in advertising: An observational survey of firm practices on facebook. *Expert Systems with Applications*, 158:113554.

Runge, J. and Skokan, I. (2025). "From Theory To Practice: How Organizations Can Embrace Experimentation In Marketing Measurement" (accessed April 2, 2025). www.adexchanger.com/data-driven-thinking/from-theory-to-practice-how-organizations-can-embrace-experimentation-in-marketing-measurement/.

Sahni, N. S. (2015). Effect of temporal spacing between advertising exposures: Evidence from online field experiments. *Quantitative Marketing and Economics*, 13:203–247.

Sahni, N. S. (2016). Advertising spillovers: Evidence from online field experiments and implications for returns on advertising. *Journal of Marketing Research*, 53(4):459–478.

Sahni, N. S. and Zhang, C. (2024). Are consumers averse to sponsored messages? The role of search advertising in information discovery. *Quantitative Marketing and Economics*, 22(1):63–114.

Sahni, N. S., Zou, D., and Chintagunta, P. K. (2017). Do targeted discount offers serve as advertising? Evidence from 70 field experiments. *Management Science*, 63(8):2688–2705.

Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

Waisman, C. and Gordon, B. R. (2025). Multicell experiments for marginal treatment effect estimation of digital ads. *Management Science (Forthcoming)*.

Waisman, C., Nair, H. S., and Carrion, C. (2025a). Online causal inference for advertising in real-time bidding auctions. *Marketing Science*, 44(1):176–195.

Waisman, C., Sahni, N. S., Nair, H. S., and Lin, X. (2025b). Parallel experimentation and competitive interference on online advertising platforms. *Marketing Science*, 44(2):437–456.

Wernerfelt, N., Tuchman, A., Shapiro, B. T., and Moakler, R. (2025). Estimating the value of offsite tracking data to advertisers: Evidence from meta. *Marketing Science*, 44(2):268–286.

Yang, J., Sahni, N. S., Nair, H. S., and Xiong, X. (2024). Advertising as information for ranking e-commerce search listings. *Marketing Science*, 43(2):360–377.

Zantedeschi, D., Feit, E. M., and Bradlow, E. T. (2017). Measuring multichannel advertising response. *Management Science*, 63(8):2706–2728.

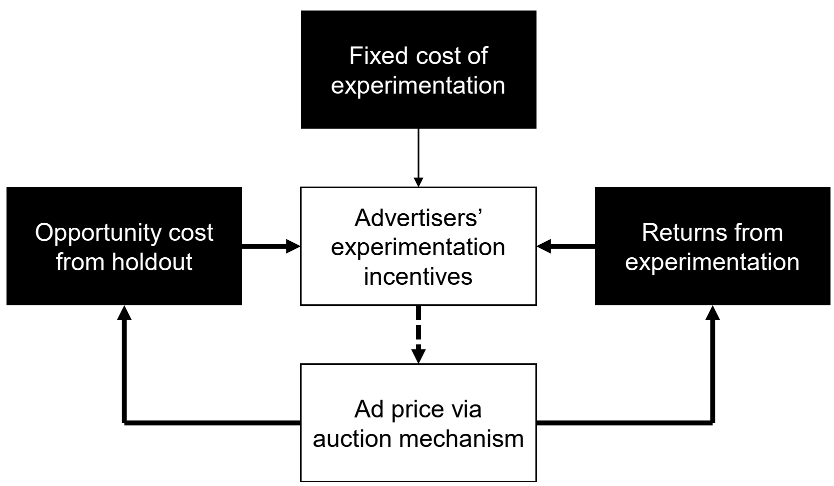

Figure 1: Overview of key forces. Thick, solid lines indicate indirect channels through which ad platform's reserve price impacts experimentation incentives. Black boxes represent factors that influence advertisers' experimentation incentives.

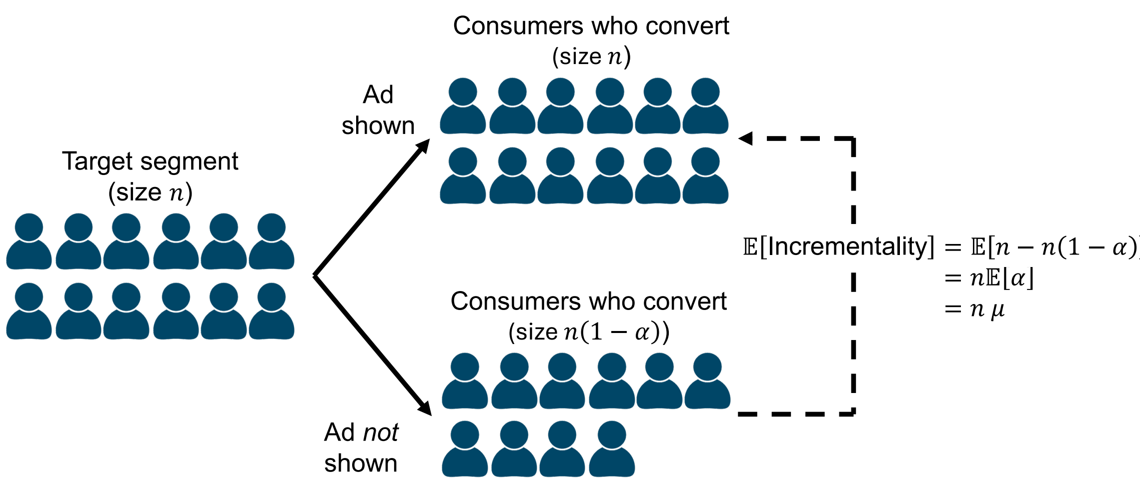

Figure 2: Advertising and conversion incrementality

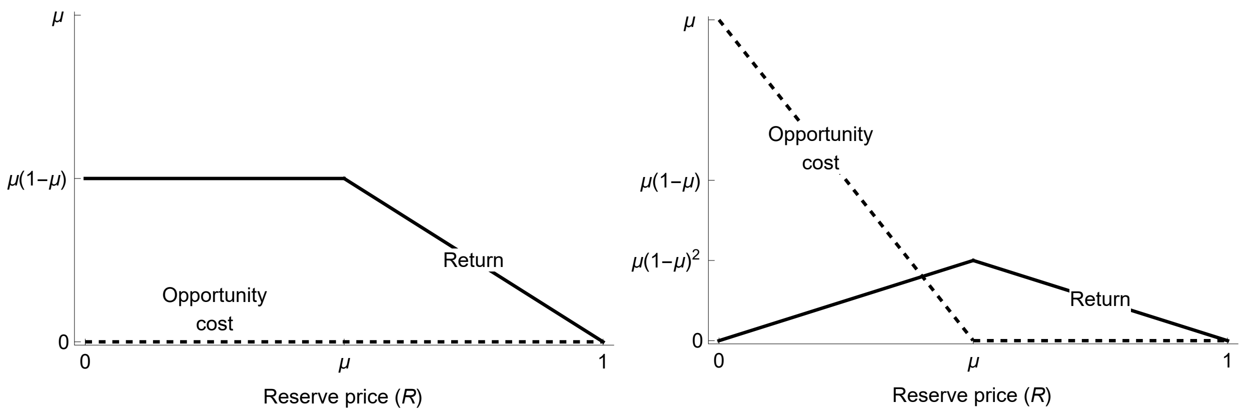

(a) Competitor does not experiment ($e_{-j} = 0$) *(b) Competitor experiments ($e_{-j} = 1$)*

Figure 3: Return and opportunity cost of experimentation

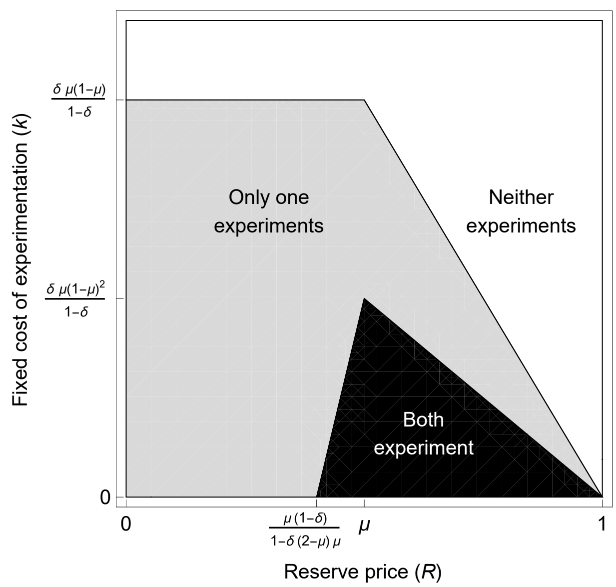

Figure 4: Subgame equilibrium.

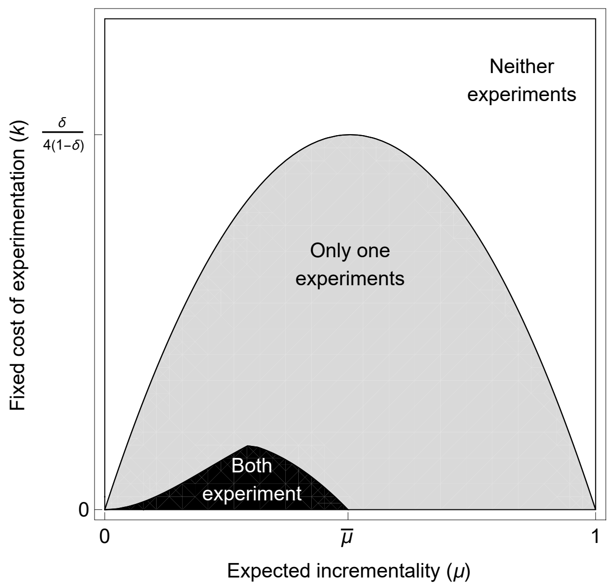

Figure 5: Reserve-optimized equilibrium.

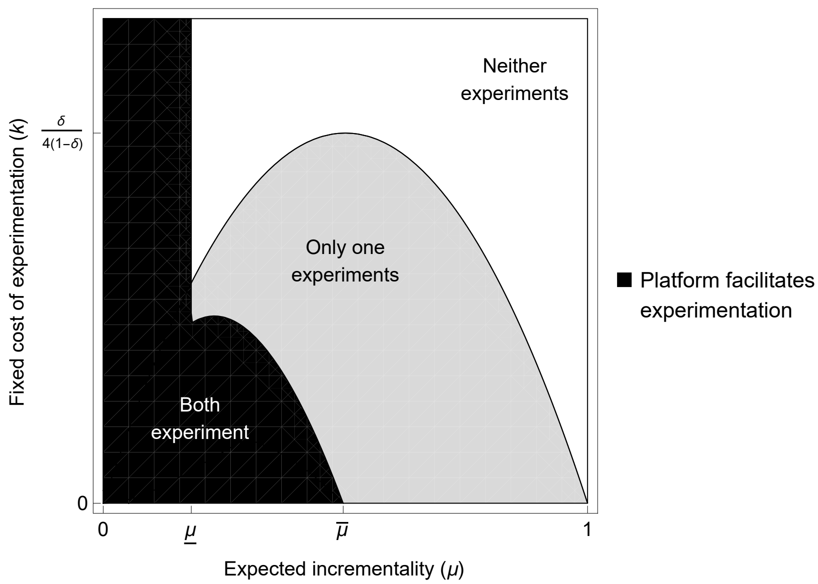

Figure 6: Total equilibrium.

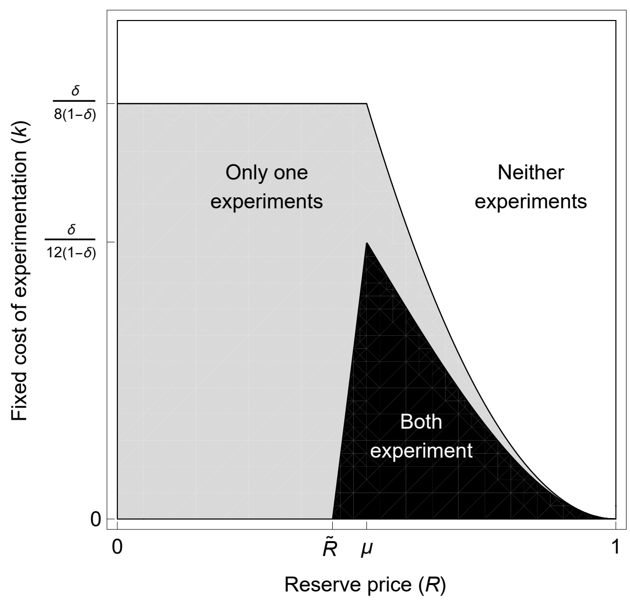

Figure 7: Subgame equilibrium ($\tilde{R}$ defined in proof of Proposition 4)

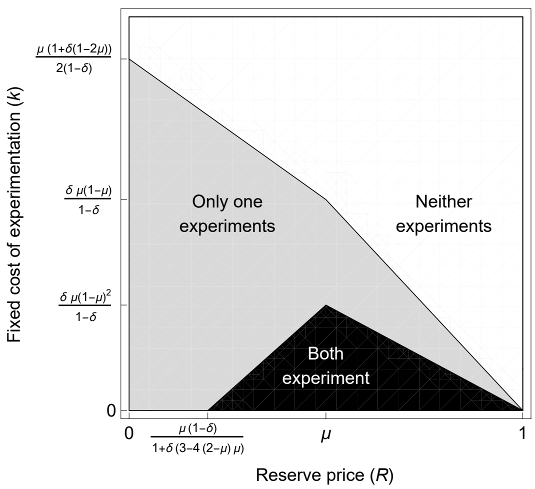

(a) Subgame equilibrium.

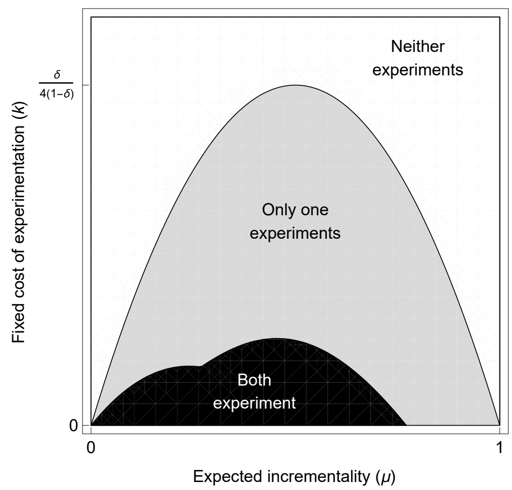

(b) Equilibrium.

Figure 8: Interior-split outcomes

π_{A1}	$e_B = 0$	$e_B = 1$
$e_A = 0$	0	$(\mu - R)^+/2$
$e_A = 1$	$-k$	$(\mu - R)^+/4 - k$

Table 1: Advertiser A's Period 1 payoff

Appendix

A1 Proofs

Proof of Lemma 2

The result follows immediately from (17) $\geq$ (16) for all $\mu \in [0,1]$.

Proof of Lemma 1

If the competitor does not experiment (i.e., $e_{-j} = 0$), then Advertiser j competes with a competitor that always bids in Period 2 the average incrementality μ. If Advertiser j experiments, then for incremental impressions, it gets a positive surplus of $1 - \max\{R, \mu\}$, where $\max\{R, \mu\}$ is Advertiser j's payment to the platform, whereas it gets zero surplus from non-incremental impressions. If Advertiser j does not experiment, then both advertisers compete away their surplus, resulting in a Period 2 payoff of zero. Thus, we have

$$\pi_{j2}(e_j = 1 | e_{-j} = 0) = \mu(1 - \max\{R, \mu\}) \text{ and } \pi_{j2}(e_j = 0 | e_{-j} = 0) = 0. \tag{A1}$$

It follows that

$$V_j(e_{-j} = 0) = \pi_{j2}(e_j = 1 | e_{-j} = 0) - \pi_{j2}(e_j = 0 | e_{-j} = 0) = \mu(1 - \max\{R, \mu\}). \tag{A2}$$

If the competitor experiments (i.e., $e_{-j} = 1$), then following similar reasoning above, we have

$$\pi_{j2}(e_j = 1 | e_{-j} = 1) = \mu(1-\mu)(1-R) \text{ and } \pi_{j2}(e_j = 0 | e_{-j} = 1) = (1-\mu)(\mu - R)^+. \tag{A3}$$

Therefore,

$$V_j(e_{-j}=1)=\pi_{j2}(e_j=1|e_{-j}=1)-\pi_{j2}(e_j=0|e_{-j}=1)=(1-\mu)\left(\mu(1-R)-(\mu-R)^+\right). \tag{A4}$$

We obtain

$$V_j(e_{-j}=0)-V_j(e_{-j}=1)=\begin{cases}(1-\mu)(\mu-R(1-\mu)^2) & \text{if } 0\le R\le\mu,\\ \mu^2(1-R) & \text{if } \mu<R\le 1,\end{cases} \tag{A5}$$

which is always positive. This proves the first part of the lemma.

The second part of the lemma follows from

$$\frac{dV_j(e_{-j}=0)}{dR}=\begin{cases}0 & \text{if } 0\le R\le\mu,\\ -\mu & \text{if } \mu<R\le 1,\end{cases} \text{ and } \frac{dV_j(e_{-j}=1)}{dR}=\begin{cases}(1-\mu)^2 & \text{if } 0\le R\le\mu,\\ -\mu(1-\mu) & \text{if } \mu<R\le 1.\end{cases} \tag{A6}$$

Proof of Proposition 1

We derive the advertisers' expected payoffs in the subgame equilibria.

- $(h_1,h_2)=(0,0)$. Without experimentation, both advertisers bid expected incrementality μ in both periods. Therefore, the advertisers' total expected payoffs in the $(0,0)$ subgame are

$$\pi_j(h_1=h_2=0)=0, \text{ for all } j\in\{1,2\}. \tag{A7}$$

- $(h_1,h_2)=(0,1)$. Advertiser B, who does not experiment, bids μ in both periods. Advertiser B bids μ in Period 1 (but earns 0 payoff due to holdout), and bids 1 for incremental impression and 0 for non-incremental impression in Period 2. Therefore,

$$\pi_1(h_1 = 0, h_2 = 1) = (1-\delta)(\mu - R)^+ + \delta(1-\mu)(\mu - R)^+ = (1-\delta\mu)(\mu - R)^+ \quad \text{(A8)}$$

and

$$\pi_2(h_1 = 0, h_2 = 1) = (1-\delta)(0-k) + \delta\mu(1 - \max\{\mu, R\}). \quad \text{(A9)}$$

- $(h_1, h_2) = (1, 1)$. If both advertisers experiment, then they both incur the fixed cost k in Period 1, and their Period 2 payoff depends on the incrementality outcomes of both advertisers. Specifically, if the impression is incremental for both advertisers, they compete away their profits; if the impression is incremental for only Advertiser j, then it earns incremental value 1 and pays the reserve price (because the competitor does not bid for its non-incremental impression); and if the impression is not incremental for either advertisers, neither bids and the payoff is zero. Overall, we have

$$\pi_j(h_1 = h_2 = 1) = (1-\delta)(0-k) + \delta\mu(1-\mu)(1-R), \text{ for all } j \in \{1, 2\}. \quad \text{(A10)}$$

Based on these advertisers' payoffs, we derive the equilibrium conditions for the different experimentation subgames.

- $(h_1, h_2) = (0, 0)$ is equilibrium if and only if

$$\pi_2(h_1 = h_2 = 0) > \pi_2(h_1 = 0, h_2 = 1) \Leftrightarrow 0 > \delta\mu(1 - \max\{\mu, R\}) - (1-\delta)k \quad \text{(A11)}$$
$$\Leftrightarrow k > \frac{\delta}{1-\delta}\mu(1 - \max\{\mu, R\}) \quad \text{(A12)}$$

- $(h_1, h_2) = (0, 1)$ is equilibrium if and only if

$$\pi_1(h_1 = 0, h_2 = 1) > \pi_1(h_1 = 1, h_2 = 1) \text{ and } \pi_2(h_1 = 0, h_2 = 1) > \pi_2(h_1 = 0, h_2 = 0). \quad \text{(A13)}$$

The second condition is the complement of above; i.e.,

$$k \leq \frac{\delta}{1-\delta}\mu(1-\max\{\mu, R\}), \tag{A14}$$

and the first condition simplifies to

$$(1-\delta\mu)(\mu-R)^{+} > \delta\mu(1-\mu)(1-R)-(1-\delta)k \tag{A15}$$

$$\Leftrightarrow k > \frac{\delta\mu(1-\mu)(1-R)-(1-\delta\mu)(\mu-R)^{+}}{1-\delta}. \tag{A16}$$

- $(h_1, h_2) = (1, 1)$ is equilibrium if and only if

$$\pi_1(h_1 = h_2 = 1) > \pi_1(h_1 = 0, h_2 = 1), \tag{A17}$$

which is the complement of above; i.e.,

$$k \leq \frac{\delta\mu(1-\mu)(1-R)-(1-\delta\mu)(\mu-R)^{+}}{1-\delta}. \tag{A18}$$

For $R \leq \mu$, this simplifies to

$$\frac{(1-\delta)(\mu+k)}{1-\delta\mu(2-\mu)} < R, \tag{A19}$$

whereas for $\mu < R$, this simplifies to

$$R \leq 1 - \frac{(1-\delta)k}{\delta\mu(1-\mu)}, \tag{A20}$$

combined with the necessary condition that the right-hand side exceed μ, which is equivalent to

$$k \leq \frac{\delta\mu(1-\mu)^2}{1-\delta}. \tag{A21}$$

Proof of Proposition 2

From Proposition 1, we know that if $k > \frac{\delta\mu(1-\mu)}{1-\delta}$, then neither firm experiments. Therefore, the optimal strategy here is any $R \in [0, \mu]$, which yields platform profit $\pi_P = \mu$.

If $\frac{\delta\mu(1-\mu)^2}{1-\delta} < k \leq \frac{\delta\mu(1-\mu)}{1-\delta}$, then the platform can set $R \leq 1 - \frac{(1-\delta)k}{\delta\mu}$ to induce $(0, 1)$ or a higher R (that exceeds μ) to induce $(0, 0)$. The latter strategy yields profit 0, so within the strategy space of inducing $(0, 1)$, the platform considers two boundary candidates $R = \mu$ and $R = 1 - \frac{(1-\delta)k}{\delta\mu}$. Note that any $R < \mu$ is dominated by $R = \mu$ and any $\mu < R < 1 - \frac{(1-\delta)k}{\delta\mu}$ by $R = 1 - \frac{(1-\delta)k}{\delta\mu}$.

If $R = \mu$, then the platform earns Period 1 payoff μ from the non-experimenting advertiser, and the same Period 2 payoff, for a total expected profit of μ. On the other hand, if $R = 1 - \frac{(1-\delta)k}{\delta\mu}$, then the platform foregoes Period 1 payoff because the reserve price exceeds the expected incrementality level that the non-experimenting advertiser bids in Period 1. The platform's Period 2 earns positive payoff if and only if the experimenting advertiser realizes an incremental impression and pays the reserve price; therefore, the platform total expected profit is $\delta\mu\left(1 - \frac{(1-\delta)k}{\delta\mu}\right)$. It can be shown that this payoff is never greater than μ, so that $R = \mu$ is the platform's optimal reserve price in this range.

Finally, if $k \leq \frac{\delta\mu(1-\mu)^2}{1-\delta}$, then the platform can set (i) $R \leq \frac{(1-\delta)(\mu+k)}{1-\delta\mu(2-\mu)}$ to induce $(0, 1)$, (ii) $\frac{(1-\delta)(\mu+k)}{1-\delta\mu(2-\mu)} < R \leq 1 - \frac{(1-\delta)k}{\delta\mu(1-\mu)}$ to induce $(1, 1)$, or (iii) $1 - \frac{(1-\delta)k}{\delta\mu(1-\mu)} < R \leq 1 - \frac{(1-\delta)k}{\delta\mu}$ to induce (0,1), or (iv) $1 - \frac{(1-\delta)k}{\delta\mu} < R$ to induce $(0, 0)$. The fourth strategy is dominated because it yields zero profit. The third strategy is dominated by the second because

$$\pi_P\left(R = 1 - \frac{(1-\delta)k}{\delta\mu}\right) = \delta\mu\left(1 - \frac{(1-\delta)k}{\delta\mu}\right) \tag{A22}$$

$$= \delta\mu - (1-\delta)k \tag{A23}$$

$$< \delta\mu^2 + 2\mu(1-\mu)\delta - (1-\delta)k - (1-\delta)k \tag{A24}$$

$$= \delta\left(\mu^2 + 2\mu(1-\mu)\left(1 - \frac{(1-\delta)k}{\delta\mu(1-\mu)}\right)\right) \tag{A25}$$

$$= \pi_P\left(R = 1 - \frac{(1-\delta)k}{\delta\mu(1-\mu)}\right), \tag{A26}$$

where the inequality in the third line follows from the presupposition $k \leq \frac{\delta\mu(1-\mu)^2}{1-\delta}$. Thus, the platform considers the first strategy, which yields profit

$$\pi_P\left(R = \frac{(1-\delta)(\mu+k)}{1-\delta\mu(2-\mu)}\right) = (1-\delta)\frac{(1-\delta)(\mu+k)}{1-\delta\mu(2-\mu)} + \delta\left(\mu^2 + (1-\mu)\frac{(1-\delta)(\mu+k)}{1-\delta\mu(2-\mu)}\right), \tag{A27}$$

and the second strategy, which yields profit

$$\pi_P\left(R = 1 - \frac{(1-\delta)k}{\delta\mu(1-\mu)}\right) = \delta\left(\mu^2 + 2\mu(1-\mu)\left(1 - \frac{(1-\delta)k}{\delta\mu(1-\mu)}\right)\right). \tag{A28}$$

Consider the profit difference

$$\Delta(k) \equiv \pi_P\left(R = 1 - \frac{(1-\delta)k}{\delta\mu(1-\mu)}\right) - \pi_P\left(R = \frac{(1-\delta)(\mu+k)}{1-\delta\mu(2-\mu)}\right). \tag{A29}$$

We have

$$\frac{d\Delta(k)}{dk} = -\frac{(1-\delta)(3-\delta\mu(5-2\mu))}{1-\delta\mu(2-\mu)} < 0 \text{ for all } \mu, \delta \in (0,1); \tag{A30}$$

therefore, $\Delta(k) > 0$ if and only if $k < \tilde{k}$, where $\Delta(\tilde{k}) = 0$. Algebraic manipulations yield

$$\tilde{k} = \frac{\delta\mu(\delta(2(\mu-3)\mu+5)\mu+\mu-3)+\mu}{(\delta-1)(\delta\mu(2\mu-5)+3)}. \tag{A31}$$

The threshold $\tilde{k}$ needs to be positive, which is equivalent to the condition

$$\delta > \tilde{\delta} \equiv -\frac{\mu + (1-\mu)\sqrt{9-8\mu} - 3}{2\mu(2(\mu-3)\mu+5)}. \tag{A32}$$

In summary,

$$\Delta(k) > 0 \Leftrightarrow k \leq \tilde{k} \text{ and } \delta > \tilde{\delta}. \tag{A33}$$

Proof of Proposition 3

From Proposition 2, we know that if $\tilde{k} < 0$, then the platform cannot induce $(1,1)$. In this case, the platform compares its profit from inducing $(0,0)$ under large k and that from inducing $(0,1)$ under small k. The former yields profit μ, whereas the latter yields profit (A28) (see proof of Proposition 2). Even if this latter profit is maximized by reducing $k \downarrow 0$, the net expected payoff never exceeds μ, so the platform does not reduce k if $\tilde{k} < 0$. Algebraic manipulations yield

$$\tilde{k} < 0 \Leftrightarrow \mu > \overline{\mu} \equiv \{\mu_0 : 1 - 3\delta + \delta(1+5\delta)\mu_0 - 6\delta^2\mu_0^2 + 2\delta^2\mu_0^3 = 0\}^+. \tag{A34}$$

Next, if $\mu \leq \tilde{\mu}$, then for fixed k, the platform induces $(0,0)$ for large k, $(0,1)$ for intermediate k, and $(1,1)$ for small k. Following the reasoning above, we know that the platform earns (weakly) higher profit under $(0,0)$ than under $(0,1)$. Therefore, if the platform's profit under $(0,0)$ is less than that under $(1,1)$ by reducing k down to 0, then the platform will reduce k and induce $(1,1)$. This happens if and only if

$$\mu < \lim_{k\downarrow 0} \delta\left(\mu^2 + 2\mu(1-\mu)\left(1 - \frac{(1-\delta)k}{\delta\mu(1-\mu)}\right)\right) = \delta\mu(2-\mu), \tag{A35}$$

which simplifies to $\mu \leq \underline{\mu} \equiv 2 - 1/\delta$. Therefore, the platform always reduces $k \downarrow 0$ and induces $(1,1)$ if $\mu \leq \underline{\mu}$.

Finally, if $\underline{\mu} < \mu \leq \overline{\mu}$, then the platform's profit is higher under $(0,0)$ than under $(1,1)$ with $k \downarrow 0$. Therefore, if $(0,0)$ is the equilibrium under exogenous k, the platform will not reduce k. However, if the equilibrium is $(0,1)$, the platform may benefit from reducing $k \downarrow 0$ and inducing $(1,1)$. This occurs if and only if

$$(1-\delta)\min\left\{\mu, \frac{(1-\delta)(\mu+k)}{1-\delta\mu(2-\mu)}\right\} + \delta\left(\mu^2 + (1-\mu)\min\left\{\mu, \frac{(1-\delta)(\mu+k)}{1-\delta\mu(2-\mu)}\right\}\right) < \delta\mu(2-\mu), \tag{A36}$$

which is equivalent to

$$k \leq \overline{k} \equiv \frac{\mu(\delta(\delta(2(\mu-3)\mu+5)\mu+\mu-3)+1)}{(1-\delta)(\delta\mu-1)}. \tag{A37}$$

Proof of Lemma 3

Because both advertisers' distributions follow the same uniform distribution and the density $f(\cdot) = 1$ is non-increasing, the result is a special case of Theorem 4.1 (Page 8) in Hummel and McAfee (2016).

Proof of Proposition 4

Suppose the reserve price R is fixed.

- $(h_1, h_2) = (0,0)$. If neither advertiser experiments, then each advertiser bids the expected incrementality $1/2$ so that their expected payoffs are zero. The platform's expected payoff is

$$\pi_P(h_1 = h_2 = 0) = \mathbb{I}_{\{R \leq \frac{1}{2}\}} \frac{1}{2}. \tag{A38}$$

- $(h_1, h_2) = (0,1)$. If only one Advertiser B experiments, then Advertiser A's expected payoff is

$$\pi_1(h_1=0,h_2=1) = (1-\delta)\left(\frac{1}{2}-R\right)^+ + \delta\int_0^{1/2}\mathbb{I}_{\left\{R\leq\frac{1}{2}\right\}}\left(\frac{1}{2}-\max\{R,\alpha_2\}\right)d\alpha_2$$
$$= (1-\delta)\left(\frac{1}{2}-R\right)^+ + \frac{\delta}{2}\left(\frac{1}{4}-R^2\right)^+,$$

and Advertiser B's is

$$\pi_2(h_1=0,h_2=1) = (1-\delta)(0-k) + \delta\int_{\max\{R,1/2\}}^1\left(\alpha_2-\max\left\{R,\frac{1}{2}\right\}\right)d\alpha_2$$
$$= \frac{\delta}{2}\left(1-\max\left\{\frac{1}{2},R\right\}\right)^2 - (1-\delta)k.$$

The platform's expected payoff is

$$\pi_P(h_1=0,h_2=1) = (1-\delta)\mathbb{I}_{\left\{R\leq\frac{1}{2}\right\}}R + \delta\cdot\begin{cases}\int_0^{\frac{1}{2}}\max\{R,\alpha_2\}d\alpha_2 + \int_{\frac{1}{2}}^1\frac{1}{2}\,d\alpha_2 & \text{if } R\leq\frac{1}{2},\\ \int_R^1 R\,d\alpha_2 & \text{if } \frac{1}{2}<R\end{cases}$$
$$= (1-\delta)\mathbb{I}_{\left\{R\leq\frac{1}{2}\right\}}R + \delta\cdot\begin{cases}\frac{1}{8}(3+4R^2) & \text{if } R\leq\frac{1}{2},\\ (1-R)R & \text{if } \frac{1}{2}<R.\end{cases}$$

- $(h_1,h_2)=(1,1)$. If both advertisers experiment, then their expected payoff is

$$\pi_j(h_1=h_2=1) = (1-\delta)(0-k) + \delta\int_0^1\int_{\max\{R,\alpha_{-j}\}}^1(\alpha_j-\max\{R,\alpha_{-j}\})\,d\alpha_j d\alpha_{-j}$$
$$= \frac{\delta}{6}(1-R^2)(1+2R) - (1-\delta)k.$$

The platform's expected payoff is

$$\pi_P(h_1=h_2=1) = \delta\iint_0^1\mathbb{I}_{\{\max\{\alpha_1,\alpha_2\}\geq R\}}\max\{\min\{\alpha_1,\alpha_2\},R\}\,d\alpha_1 d\alpha_2$$
$$= \delta\left(\frac{1}{3}+R^2-\frac{4}{3}R^3\right).$$

We express the non-deviation conditions for each of the above three subgames.

- $(h_1, h_2) = (0, 0)$ is equilibrium if and only if

$$0 > \pi_2(h_1 = 0, h_2 = 1) \Leftrightarrow k > \frac{\delta}{2(1-\delta)} \left(1 - \max\left\{\frac{1}{2}, R\right\}\right)^2 \tag{A39}$$

- $(h_1, h_2) = (0, 1)$ is equilibrium if and only if

$$0 < \pi_2(h_1 = 0, h_2 = 1) \Leftrightarrow k < \frac{\delta}{2(1-\delta)} \left(1 - \max\left\{\frac{1}{2}, R\right\}\right)^2 \tag{A40}$$

  and

$$\pi_1(h_1 = 0, h_2 = 1) > \pi_1(h_1 = h_2 = 1) \Leftrightarrow k > \frac{1}{1-\delta} \begin{cases} \frac{13\delta}{24} + \frac{1}{3}\delta\left(R^2 - 3\right)R + R - \frac{1}{2} & \text{if } R \leq \frac{1}{2}, \\ \frac{1}{6}\delta(R-1)^2(2R+1) & \text{if } \frac{1}{2} < R. \end{cases} \tag{A41}$$

- $(h_1, h_2) = (1, 1)$ is equilibrium if and only if

$$\pi_1(h_1 = 0, h_2 = 1) < \pi_1(h_1 = h_2 = 1) \Leftrightarrow k < \frac{1}{1-\delta} \begin{cases} \frac{13\delta}{24} + \frac{1}{3}\delta\left(R^2 - 3\right)R + R - \frac{1}{2} & \text{if } R \leq \frac{1}{2}, \\ \frac{1}{6}\delta(R-1)^2(2R+1) & \text{if } \frac{1}{2} < R. \end{cases} \tag{A42}$$

  This simplifies to

$$k < \frac{\delta}{12(1-\delta)} \text{ and } \underline{R} < R \leq \overline{R}, \tag{A43}$$

  where

$$\underline{R} \equiv \left\{R \in [0, 1/2] : -12 + 13\delta + 24(1-\delta)(R-k) + 8\delta R^3 = 0\right\} \tag{A44}$$

  and

$$\overline{R} \equiv \{R \in [1/2, 1] : -6k(1-\delta) + \delta - 3\delta R^2 + 2\delta R^3 = 0\}. \tag{A45}$$

Because we need $k > 0$, a necessary condition is that $\frac{13\delta}{24} + \frac{1}{3}\delta\left(R^2 - 3\right)R + R - \frac{1}{2} > 0$ for $R \leq 1/2$. This is equivalent to

$$R > \tilde{R} \equiv \left(\frac{\left(-26\delta^3 + 2\sqrt{\delta^3(256 - 3\delta(\delta(29\delta - 152) + 208))} + 24\delta^2\right)^{2/3} + 8\sqrt[3]{2}(\delta - 1)\delta}{4\delta\sqrt[3]{\sqrt{\delta^3(256 - 3\delta(\delta(29\delta - 152) + 208))} + (12 - 13\delta)\delta^2}} \right)^+ . \tag{A46}$$

A1.1 Proof of Proposition 5

If $\frac{\delta}{8(1-\delta)} < k$, then neither advertiser experiments for any R. Therefore, the platform sets the reserve price as high as the expected incrementality 1/2, which induces $(0, 0)$.

If $\frac{\delta}{12(1-\delta)} < k \leq \frac{\delta}{8(1-\delta)}$, then (i) only one advertiser experiments for $R \leq 1 - \sqrt{\frac{2k(1-\delta)}{\delta}}$, and (ii) neither advertiser experiments otherwise. In this range of k, we have $\frac{1}{2} \leq 1 - \sqrt{\frac{2k(1-\delta)}{\delta}}$. Therefore, the platform's profit under $(0, 0)$ is zero. Because $\pi_P(h_1 = 0, h_2 = 1)$ attains its maximum at $R = 1/2$, we obtain that the platform sets $R = 1/2$, which induces $(0, 1)$.

If $\left(\frac{13\delta - 12}{24(1-\delta)}\right)^+ < k \leq \frac{\delta}{12(1-\delta)}$,[20] then (i) only one advertiser experiments for $R \leq R_1$, where

$$R_1 \equiv \{R \in [0, 1/2] : -12 - 24k + 13\delta + 24k\delta + 24(1-\delta)R + 8\delta R^3 = 0\}, \tag{A47}$$

(ii) both advertisers experiment for $R_1 < R \leq R_2$, where

$$R_2 \equiv \{R \in [1/2, 1] : -6k + \delta + 6k\delta - 3\delta R^2 + 2\delta R^3 = 0\}, \tag{A48}$$

(iii) only one advertiser experiments for $R_2 < R \leq 1 - \sqrt{\frac{2k(1-\delta)}{\delta}}$, and (iv) neither advertiser experiments for larger values of R. The locally optimal platform profit under (i)

$$\pi_P(h_1 = 0, h_2 = 1)|_{R=R_1} = (1-\delta)\frac{1}{2} + \frac{\delta}{8}(3 + 4R_1^2), \tag{A49}$$

[20] The lower bound expression derives from $k > 0$ and the condition that $R_1 > 0 \Leftrightarrow k > \frac{13\delta - 12}{24(1-\delta)}$.

that under (ii) is

$$\pi_P(h_1 = h_2 = 1)|_{R=1/2} = \frac{5\delta}{12}, \tag{A50}$$

that under (iii) is

$$\pi_P(h_1 = 0, h_2 = 1)|_{R=R_2^+} = \delta(1 - R_2^+)R_2^+, \tag{A51}$$

and that under (iv) is $R \in \left(1 - \sqrt{\frac{2k(1-\delta)}{\delta}}, 1\right]$ because the platform's profit is zero. Clearly, this last strategy space is dominated. The platform's locally optimal profit under (iii) is strictly less than that under (ii) because $\pi_P(h_1 = 0, h_2 = 1)$ is decreasing in R for $R > \frac{1}{2}$, and

$$\pi_P(h_1 = 0, h_2 = 1)|_{R=R_2^+} < \pi_P(h_1 = 0, h_2 = 1)|_{R=(1/2)^+} = \frac{\delta}{4} < \frac{5\delta}{12} = \pi_P(h_1 = h_2 = 1)|_{R=1/2}. \tag{A52}$$

Therefore, the platform compares locally optimal profits under (i) and (ii). We will show that there exists $\hat{k} \in \left(\left(\frac{13\delta-12}{24(1-\delta)}\right)^+, \frac{\delta}{12(1-\delta)}\right)$ such that (i) dominates if and only if $k > \hat{k}$.

Note that for general R, the platform's profit under (i) is increasing in R because

$$\pi_P(h_1 = 0, h_2 = 1) = (1 - \delta)\frac{1}{2} + \frac{\delta}{8}(3 + 4R^2). \tag{A53}$$

Moreover, we have from the definition of R_1 that R_1 is increasing in k. To see this, we have from the definition: $-12 - 24k + 13\delta + 24k\delta + 24(1 - \delta)R_1 + 8\delta R_1^3 = 0$, which implies

$$\frac{d}{dk}\left(-12 - 24k + 13\delta + 24k\delta + 24(1 - \delta)R_1 + 8\delta R_1^3\right) = 0 \Leftrightarrow \frac{dR_1}{dk}\left(1 - \delta + \delta R_1^2\right) = 1 - \delta,$$

from which it follows that $dR_1/dk > 0$.

Therefore, as a function of k, within the range $\left(\frac{13\delta-12}{24(1-\delta)}\right)^+ < k \leq \frac{\delta}{12(1-\delta)}$, the platform's locally optimal profit under (i) attains its highest value at the upper-bound $k_h \equiv \frac{\delta}{12(1-\delta)}$, and its lowest at the lower-bound $k_l \equiv \left(\frac{13\delta-12}{24(1-\delta)}\right)^+$. The maximum value is

$$\begin{aligned}\pi_P(h_1=0,h_2=1)|_{R=R_1(k_h)} &= (1-\delta)\frac{1}{2}+\frac{\delta}{8}(3+4R_1^2(k_h))\\ &= (1-\delta)\frac{1}{2}+\frac{\delta}{8}(3+4(1/2)^2)\\ &= \frac{1}{2},\end{aligned}$$

which is always greater than $5\delta/12$.

To compute the minimum value, consider two separate regions:

- $\delta \le 12/13$: in this case $k_l = 0$, and

$$R_1(k_l) = \{R : -12+13\delta+24(1-\delta)R+8\delta R^3 = 0\}. \tag{A54}$$

But

$$\pi_P(h_1=0,h_2=1)|_{R=R_1(k_l)} = (1-\delta)\frac{1}{2}+\frac{\delta}{8}(3+4R_1^2(k_l)) \tag{A55}$$

decreases in δ so that

$$\begin{aligned}\pi_P(h_1=0,h_2=1)|_{R=R_1(k_l)} &= (1-\delta)\frac{1}{2}+\frac{\delta}{8}(3+4R_1^2(k_l))\\ &\ge \pi_P(h_1=0,h_2=1)|_{R=R_1(k_l),\delta=12/13}\\ &= (1-\delta)\frac{1}{2}+\delta\frac{3}{8}\end{aligned}$$

Finally, this last expression is greater than $5\delta/12$ for all $\delta \le 12/13$; therefore, we conclude that if $\delta \le 12/13$, then inducing $(0,1)$ with $R = R_1$ always yields higher payoff for the platform than inducing $(1,1)$ with $R = 1/2$.

- $12/13 < \delta \le 1$: in this case, $k_l = (13\delta-12)/(24(1-\delta))$, and $R_1(k_l) = 0$. The platform's profit is $(1-\delta)\frac{1}{2}+\delta\frac{3}{8}$, which is less than $5\delta/12$ for all $\delta > 12/13$. In sum, if $\delta > 12/13$, then the platform's locally optimal profit under (i) is less than that under (ii) if $k = k_l$, and it is greater than that under (ii) if $k = k_h$. Therefore, by the Intermediate Value Theorem, there exists $\hat{k} \in (k_l, k_h)$ such that the platform's profit is higher under (i) than under (ii) if $k > \hat{k}$ and higher under (ii) than under (i) if $k \leq \hat{k}$.

Finally, if $k \leq \left(\frac{13\delta-12}{24(1-\delta)}\right)^+$, then the platform cannot induce $(0,1)$ for any $R \leq R_2$. Therefore, the optimal reserve price is $R = 1/2$, which induces $(1,1)$.

A1.2 Proof of Lemma 4

$$\pi_j(h_1 = h_2 = 0) = 0,$$

$$\pi_1(h_1 = 0, h_2 = 1) = \frac{1}{2}(1 - \delta(1 - 2\mu))\,(\mu - R)^+\,,$$

$$\pi_2(h_1 = 0, h_2 = 1) = -(1-\delta)k + \begin{cases} \frac{1}{2}\left(\mu(1 + \delta(1 - 2\mu)) - R(1 - \delta)\right) & \text{if } R \leq \mu, \\ \delta\mu(1 - R) & \text{otherwise,} \end{cases}$$

$$\pi_j(h_1 = h_2 = 1) = -(1-\delta)k + \frac{1}{4}(1-\delta)(\mu - R)^+ + \delta\mu(1-\mu)(1-R)$$

Subgame equilibrium

- (0,0) is eq'm iff

$$\pi_j(h_1 = h_2 = 0) > \pi_2(h_1 = 0, h_2 = 1) \Leftrightarrow k > \frac{1}{2}(\mu - R)^+ + \frac{\delta\mu}{1-\delta}(1 - \max\{\mu, R\})$$

- (1,1) is eq'm iff $\pi_j(h_1 = h_2 = 1) > \pi_1(h_1 = 0, h_2 = 1)$, which simplifies to

$$k \leq \frac{\delta\mu(1-\mu)^2}{1-\delta} \text{ and } \frac{(1-\delta)(4k+\mu)}{\delta(3 - 4(2-\mu)\mu) + 1} < R \leq 1 - \frac{k(1-\delta)}{\delta(1-\mu)\mu}.$$

- In other cases, (0,1) is eq'm.

A1.3 Proof of Proposition 6

Optimal reserve price

- If $k > \frac{\mu(\delta(1-2\mu)+1)}{2(1-\delta)}$, then subgame eq'm if (0,0) for all R, so platform sets $R = \mu$.
- $\frac{\delta\mu(1-\mu)}{1-\delta} < k \leq \frac{\mu(\delta(1-2\mu)+1)}{2(1-\delta)}$: then subgame is (0,1) for $R \leq \frac{\mu(\delta(1-2\mu)+1)}{1-\delta} - 2k$ and (0,0) for $\frac{\mu(\delta(1-2\mu)+1)}{1-\delta} - 2k < R$, where $\frac{\mu(\delta(1-2\mu)+1)}{1-\delta} - 2k < \mu$.
  - if $R = \frac{\mu(\delta(1-2\mu)+1)}{1-\delta} - 2k$, platform's profit is $\pi_P = (1-\delta)\left(\frac{1}{2}\left(\frac{\mu(\delta(1-2\mu)+1)}{1-\delta} - 2k\right) + \frac{1}{2}\mu\right) + \delta\left((1-\mu)\left(\frac{\mu(\delta(1-2\mu)+1)}{1-\delta} - 2k\right) + \mu \cdot \mu\right)$
  - if $R = \mu$, platform's profit is μ
  - latter strategy dominates because the former attains its maximum at the lower bound of k; i.e., $k = \frac{\delta\mu(1-\mu)}{1-\delta}$. at this k value, former profit simplifies to μ.
- $\frac{\delta\mu(1-\mu)^2}{1-\delta} < k \leq \frac{\delta\mu(1-\mu)}{1-\delta}$
  - if $R \leq 1 - \frac{(1-\delta)k}{\delta\mu}$, where $\mu < 1 - \frac{(1-\delta)k}{\delta\mu}$, then (0,1)
  - if $1 - \frac{(1-\delta)k}{\delta\mu} < R$, then (0,0). this strategy yields 0 profit because reserve price exceeds the non-experimentation ex ante expected incrementality.
  - if $R = \mu$, then platform's profit is μ
  - if $R = 1 - \frac{(1-\delta)k}{\delta\mu}$, then platform's profit is $\delta\mu(1 - \frac{(1-\delta)k}{\delta\mu})$ which is less than μ.
  - therefore, $R = \mu$ and (0,1)
- $0 < k \leq \frac{\delta\mu(1-\mu)^2}{1-\delta}$
  - if $R \leq \frac{(1-\delta)(4k+\mu)}{\delta(3-4(2-\mu)\mu)+1}$, then (0,1),
  - if $\frac{(1-\delta)(4k+\mu)}{\delta(3-4(2-\mu)\mu)+1} < R \leq 1 - \frac{k(1-\delta)}{\delta(1-\mu)\mu}$, then (1,1).

    if $R = \mu$, then platform's profit is $\delta\left(\mu^2 \cdot 1 + 2\mu(1-\mu)\mu\right) + (1-\delta)\left(\frac{\mu}{2} + \frac{\mu}{4}\right)$.

    if $R = 1 - \frac{k(1-\delta)}{\delta(1-\mu)\mu}$, then platform's profit is $\delta\left(\mu^2 \cdot 1 + 2\mu(1-\mu)\left(1 - \frac{k(1-\delta)}{\delta(1-\mu)\mu}\right)\right)$

– if $1-\frac{k(1-\delta)}{\delta(1-\mu)\mu} < R \leq 1-\frac{(1-\delta)k}{\delta\mu}$, then (0,1). if $R = 1-\frac{(1-\delta)k}{\delta\mu}$, platform's profit is $\delta\mu(1-\frac{(1-\delta)k}{\delta\mu}) = \delta\mu-(1-\delta)k$, which is dominated by above.

– if $1-\frac{(1-\delta)k}{\delta\mu} < R$, then (0,0), where $\mu \in (\frac{(1-\delta)(4k+\mu)}{\delta(3-4(2-\mu)\mu)+1}, 1-\frac{k(1-\delta)}{\delta(1-\mu)\mu})$. this is dominated because under (0,0) with reserve price exceeding the expected incrementality, platform's profit is 0.

in sum, platform considers three reserve price strategies:

$$\pi_P(R = \frac{(1-\delta)(4k+\mu)}{\delta(3-4(2-\mu)\mu)+1}) = (1-\delta)\left(\frac{(1-\delta)(4k+\mu)}{2(\delta(3-4(2-\mu)\mu)+1)}+\frac{\mu}{2}\right)+\delta\left(\frac{(1-\mu)((1-\delta)(4k+}{\delta(3-4(2-\mu)\mu)+}\right.$$
$$\pi_P(R=\mu) = \delta\left(1\mu^2+2\mu(1-\mu)\mu\right)+(1-\delta)\left(\frac{\mu}{2}+\frac{\mu}{4}\right),$$
$$\pi_P(R = 1-\frac{k(1-\delta)}{\delta(1-\mu)\mu}) = \delta\left(2\mu(1-\mu)\left(1-\frac{k(1-\delta)}{\delta(1-\mu)\mu}\right)+1\mu^2\right),$$

which respectively, induce (0,1), (1,1), and (1,1).

therefore, (1,1) is eq'm iff $0 < k \leq \frac{\delta\mu(1-\mu)^2}{1-\delta}$ and $\max\left\{\pi_P(R=\mu), \pi_P(R=1-\frac{k(1-\delta)}{\delta(1-\mu)\mu})\right\} > \pi_P(R=\frac{(1-\delta)(4k+\mu)}{\delta(3-4(2-\mu)\mu)+1})$.

but $\max\left\{\pi_P(R=\mu), \pi_P(R=1-\frac{k(1-\delta)}{\delta(1-\mu)\mu})\right\}$ decreases in k whereas $\pi_P(R=\frac{(1-\delta)(4k+\mu)}{\delta(3-4(2-\mu)\mu)+1})$ increases in k.

consider the endpoints of k: $k=0$ and $k=\frac{\delta\mu(1-\mu)^2}{1-\delta}$. at $k=\frac{\delta\mu(1-\mu)^2}{1-\delta}$, we have that $\max\left\{\pi_P(R=\mu), \pi_P(R=1-\frac{k(1-\delta)}{\delta(1-\mu)\mu})\right\} > \pi_P(R=\frac{(1-\delta)(4k+\mu)}{\delta(3-4(2-\mu)\mu)+1})$ holds iff $\delta > \delta'$, where

$$\delta' = \begin{cases} \frac{1}{2}\left(\frac{\mu\left(-4\mu^2+6\mu-1\right)+1}{(2\mu-1)(\mu(4\mu((\mu-4)\mu+6)-19)+8)}+\sqrt{\frac{(\mu-1)^2(4\mu(\mu(4(\mu-3)\mu+19)-19)+33)}{(1-2\mu)^2(\mu(4\mu((\mu-4)\mu+6)-19)+8)^2}}\right) & \text{if } 0 \leq \mu \leq \frac{1}{4}, \\ -\frac{2(\mu(5\mu-8)+2)}{6(3-2\mu)^2\mu-3}+2\sqrt{\frac{(\mu-1)^2\mu(\mu(2\mu-3)(2\mu+1)+4)}{(1-2\mu)^2(1-2(3-2\mu)^2\mu)^2}}+\frac{1}{6\mu-3} & \text{if } \frac{1}{4} < \mu \leq \frac{1}{2}, \\ -\frac{2(\mu(5\mu-8)+2)}{6(3-2\mu)^2\mu-3}-2\sqrt{\frac{(\mu-1)^2\mu(\mu(2\mu-3)(2\mu+1)+4)}{(1-2\mu)^2(1-2(3-2\mu)^2\mu)^2}}+\frac{1}{6\mu-3} & \text{if } \frac{1}{2} < \mu \leq 1 \end{cases} \tag{A56}$$

consider the lower endpoint $k=0$. at this value, $\max\left\{\pi_P(R=\mu), \pi_P(R=1-\frac{k(1-\delta)}{\delta(1-\mu)\mu})\right\} >$

$\pi_P(R = \frac{(1-\delta)(4k+\mu)}{\delta(3-4(2-\mu)\mu)+1})$ holds iff $\delta > \tilde{\delta}$, where

$$\tilde{\delta} = \begin{cases} \frac{1}{2}\left(\sqrt{\frac{(\mu-1)^2(4(\mu-9)\mu+33)}{(\mu(-8\mu^2+26\mu-27)+8)^2}} + \frac{\mu(3-2\mu)+1}{(2\mu-1)(\mu(4\mu-11)+8)}\right) & \text{if } \mu \leq \frac{1}{7}, \\ \frac{1-2(\mu-1)\mu}{4(\mu-1)\mu(8(\mu-2)\mu+7)+1} + 2\sqrt{-\frac{(\mu-1)^2\mu(7\mu-8)}{(4(\mu-1)\mu(8(\mu-2)\mu+7)+1)^2}} & \text{if } \frac{1}{7} < \mu \leq \frac{1}{2}, \\ \frac{2\left(-\mu+\sqrt{\mu}\sqrt{8-7\mu}+1\right)\mu-2\sqrt{(8-7\mu)\mu}+1}{4(\mu-1)\mu(8(\mu-2)\mu+7)+1} & \text{if } \frac{1}{2} < \mu. \end{cases} \quad \text{(A57)}$$

therefore, if $\delta < \tilde{\delta}$, then $\max\left\{\pi_P(R = \mu), \pi_P(R = 1 - \frac{k(1-\delta)}{\delta(1-\mu)\mu})\right\} < \pi_P(R = \frac{(1-\delta)(4k+\mu)}{\delta(3-4(2-\mu)\mu)+1})$ for all $k \leq \frac{\delta\mu(1-\mu)^2}{1-\delta}$.

if $\tilde{\delta} < \delta < \delta'$, then $\max\left\{\pi_P(R = \mu), \pi_P(R = 1 - \frac{k(1-\delta)}{\delta(1-\mu)\mu})\right\} > \pi_P(R = \frac{(1-\delta)(4k+\mu)}{\delta(3-4(2-\mu)\mu)+1})$ holds iff $k < \tilde{k}$ for $\tilde{k} \in (0, \frac{\delta\mu(1-\mu)^2}{1-\delta})$ by IVT.

in sum, if $\delta > \tilde{\delta}$, then there exists $\tilde{k}' \in (0, \frac{\delta\mu(1-\mu)^2}{1-\delta}]$ such that $\max\left\{\pi_P(R = \mu), \pi_P(R = 1 - \frac{k(1-\delta)}{\delta(1-\mu)\mu})\right\} > \pi_P(R = \frac{(1-\delta)(4k+\mu)}{\delta(3-4(2-\mu)\mu)+1})$ iff $k < \tilde{k}'$.