
Causal Inference from Competing Treatments

Ana-Andreea Stoica¹ Vivian Yvonne Nastl^{1,2} Moritz Hardt¹

Abstract

Many applications of RCTs involve the presence of multiple treatment administrators—from field experiments to online advertising—that compete for the subjects’ attention. In the face of competition, estimating a causal effect becomes difficult, as the position at which a subject sees a treatment influences their response, and thus the treatment effect. In this paper, we build a game-theoretic model of agents who wish to estimate causal effects in the presence of competition, through a bidding system and a utility function that minimizes estimation error. Our main technical result establishes an approximation with a tractable objective that maximizes the sample value obtained through strategically allocating budget on subjects. This allows us to find an equilibrium in our model: we show that the tractable objective has a pure Nash equilibrium, and that any Nash equilibrium is an approximate equilibrium for our general objective that minimizes estimation error under broad conditions. Conceptually, our work successfully combines elements from causal inference and game theory to shed light on the equilibrium behavior of experimentation under competition.

1. Introduction

Randomized controlled trials (RCTs) have become ubiquitous in a variety of fields, to the point that they are “not so much the ‘gold standard’ as just a standard tool in the toolbox” (Banerjee et al., 2016). Application areas range from A/B testing in industry to informing policy choices at a large scale (Levy, 2007; Alatas et al., 2012; de Souza Leão & Eyal, 2019).

¹Social Foundations of Computation, Max Planck Institute for Intelligent Systems, Tübingen, Germany and Tübingen AI Center, Germany ²Department of Mathematics, ETH Zürich, Zürich, Switzerland. Correspondence to: Ana-Andreea Stoica <astoica@tuebingen.mpg.edu>, Vivian Yvonne Nastl <vivian.nastl@tuebingen.mpg.de>.

Given the abundance of experimentation, an agent who wishes to measure the effect of their treatment through an RCT may not be alone in doing so. Competing treatments arise when multiple agents share the population on which they wish to experiment. Subjects may then receive multiple treatments sequentially determined by some form of competition between the treatment administrators. Treatments that came before may modify the effectiveness of subsequent treatments. Precisely estimating effects therefore becomes challenging under competition: sequential treatments may pose the issue of external validity (Duflo, 2006) or mislead companies in their assessment of products (e.g. as the click-through-rate plummets for low-position content).

Our concrete running example is online advertising. Multiple advertisers wish to estimate the effectiveness of their campaign. However, due to competition between advertisers, their ad will be displayed at a certain position on screen. The lower the rank of the ad, the lesser its effectiveness (Agarwal et al., 2011). Running the campaign, an advertiser collects data of impressions and clicks on ads displayed at different ranks. From this data, the advertiser wishes to answer the question:

What would have been the effect of the treatment had it been the first treatment applied?

In other words, how effective is the ad if it is displayed in the top rank? Equipped with knowledge of this effect, an advertiser can distinguish between two scenarios: (1) an ad was not effective because it was mostly displayed at lower ranks, or (2) an ad was not effective because of its design. An advertiser can then take appropriate action to increase its effectiveness in subsequent campaigns. The advertiser faces two problems: one game-theoretic, the other statistical. First, how should the advertiser bid on subjects, anticipating that other advertisers will bid strategically, too? Second, how should data from different ranks be weighted optimally to obtain the best possible estimate of the causal effect?

Specifically, we show what the optimal estimation error is that an advertiser can achieve at equilibrium in a game that models the competition between advertisers. In other words, our work shows how to spend an experimental budget rationally on subjects when the goal is to estimate a causal effect optimally.

1.1. Our Contribution

We introduce a novel model of causal inference under competing treatments. Our model combines elements from game theory and statistics to find an answer that neither toolkit can provide on its own. Specifically, how should treatment administrators allocate their budget so as to minimize statistical estimation error, while facing competition?

We assume that the causal effect of interest is a bounded quantity τ that we interpret to be the effect of treatment had it come first. A number k of treatment administrators compete over a pool of n subjects. Administrators have budgets that they can spend by making bids on subjects. A probabilistic allocation rule determines the order in which each subject receives the k treatments. We assume that the treatment effectiveness decays with the order in which the treatment is applied. A treatment applied at position r has a causal effect $\tau_r < \tau$. We assume that $\tau_r = \alpha_r \cdot \tau$, where α_r is a known structural parameter of our model.

In our model, administrators bid strategically so as to achieve the smallest possible estimation error. The optimal estimation error depends on what we call the *sample profile*, that is, the number of samples n_r at each rank r . One of our key contributions is to approximate the optimal error objective by a more tractable objective that we call *sample value*. The sample value corresponds to the weighted sum $S = \sum_r n_r \cdot \alpha_r^2$. The sample value quantifies the effective utility of the sample where each rank is discounted appropriately.

We prove that, in the typical regime where budgets are larger than k^2 , any Nash equilibrium in the sample value objective is also an approximate Nash equilibrium with respect to the optimal error objective. Using this approximation result, we can understand the competition over optimal error by instead analyzing the competition over sample value. This objective has many nice properties. We show that it has a pure Nash equilibrium, by finding a class of equilibria in closed form in which rational administrators spend their budget exhaustively in a way that maximally avoids competition.

Conceptually, we find a fruitful bridge between causal inference and game theory. On the one hand, causal inference literature tells us how to account for known position effects in an optimal way. On the other hand, game theory provides solution concepts for dealing with competition. A novelty is that we not only apply game theory to budget allocation strategies but rather to estimate causal effects. Our work opens the door for numerous interesting questions, which we discuss in Section 5, alongside current limitations of our framework.

1.2. Related Work

Several fields provide evidence that the position at which content is seen affects engagement, from early theories in experimental psychology (Ebbinghaus, 1885) to survey design (McFarland, 1981). Because of position effects randomized experiments in development economics run the risk of reducing the external validity of sequential studies run on the same population (Duflo, 2006). Perhaps the most extensive work studies the digital space, where search results displayed lower on a webpage receive less engagement (Ansari & Mela, 2003; Teevan, 2008) and low-positioned ads are recalled less (Varian, 2007; Agarwal et al., 2011). In particular, Narayanan & Kalyanam (2015) measure the effect of positioning ads under selection biases such as competition using regression-discontinuity, without modeling the competition for ranks explicitly. Moreover, Hardt et al. (2022) define a framework for measuring the effect of position on market power acquisition.

Our work draws on this literature and models position effects that emerge from competition: agents who wish to run randomized controlled trials on a population encounter a position effect when competing with each other over ‘winning’ the attention of subjects. Such competition has been studied in game-theoretic settings. In particular in applications like online advertising, pricing mechanisms have been developed in order to achieve stability in the market (Aggarwal et al., 2006; Edelman et al., 2007). A related line of work is that of position-based auctions. Early works in this research line studied equilibria and truthfulness through designs that encode positions in the agents’ valuations for individual users (Varian, 2007; Athey & Ellison, 2011; Börgers et al., 2013). Another recent line of work studies optimal platform design choices under proxy bidding (or ‘autobidding’), where agents’ strategy space is limited to budget choices and targeting mechanisms (Aggarwal et al., 2019; Balseiro et al., 2021; Conitzer et al., 2022). Ghosh & Sayedi (2010) study alternative designs for auctions with negative externalities generated by competing advertisers in online conversion-based auctions, without an explicit position-based model. Our work differs from this literature in the following ways: in our set-up, agents do not have an intrinsic value for particular individuals—a common auction-modeling choice in early position-based auctions—but rather a sample value that determines the estimation error of their. This brings our problem closer to the ‘autobidding’ literature, yet, differs from it in that our optimization problem includes both bidding strategies and the choice of an optimal estimator for precise estimation of causal effects.

Our game-theoretic design is inspired by and directly generalizes the work of Maehara et al. (2015). Their design optimizes the sample size of customers obtained by competing advertisers, in a line of work that bridges

influence maximization with budget allocation (Alon et al., 2012). Whereas in Maehara et al. (2015)’s design there is no explicit rank (an advertiser either wins a customer or not, maximizing the size of their sample), our model generalizes theirs by including multiple positions in which advertisers may be shown to users. This generalized game may be of independent research interest, as we can characterize particular Nash equilibria.

Estimating treatment effects under equilibrium conditions has been studied in the case of unit interference (Wager & Xu, 2021) and two-sided markets (Johari et al., 2022). They consider dynamic models in which the treatment effect is dependent on the available market supply and unit’s interactions, assuming a single treatment coming from a decision-maker (often a platform designer).

Adaptive experimentation is a rapidly growing area in estimating causal effects under sequential treatments. Hadad et al. (2021) and Dwivedi et al. (2022) provide statistical inference guarantees in the problem of estimating counterfactual means, without explicitly modeling multiple ranks or strategic interactions. Methods such as multi-armed bandits have been employed to provide guarantees on the efficacy of treatments through optimal data collection in the presence of position effects (Lagrée et al., 2016; Zhou et al., 2023). They primarily assume a *single* learner. While bearing similarities in objectives (as our objective also directly optimizes for estimation error), our work differs from these by modeling data acquisition through competition that emerges from *multiple* learners.

2. Games and Objectives

We formally introduce a game that captures the setting of k treatment administrators (also called admins) who compete over a set of n subject slots. Linking back to our online advertising example, the different admins are represented by different advertisers and the subject slots are positions at which users will see the displayed ads on an online platform. Conceptually, subject slots are placeholders for i.i.d. samples drawn from a data-generating distribution. Administrators have varying budgets that they can spend exhaustively by bidding on the subject slots. The outcome of the game is an allocation of the subject slots to the k administrators. Thus, each admin receives up to n slots for displaying their ad. For each slot, they can run a treatment-control experiment and measure the effect of an ad. The formal definition of the data-generating process is given in Section 3. We assume that the causal effect of the treatment at rank 1 is given by a bounded quantity τ with $|\tau| \leq 1$.¹ We also assume that

¹The effect just needs to be bounded quantity. We assume it is bounded by 1 without loss of generality.

treatment at a higher rank $r > 1$ is less effective so that

$$\tau_r = \tau \cdot \alpha_r \quad \text{with} \quad 0 < \alpha_r < 1.$$

The discount factors $\{\alpha_r\}$ are known structural parameters of our model. We discuss the case where the discount factors are not *a priori* known in Section 3. Samples at rank 1 are most valuable, but samples from lower ranks can still be useful. A treatment administrator bids strategically so as to minimize their mean squared error in estimating the causal effect τ . To fully specify the game, we need the following ingredients:

- We can represent the bids as a nonnegative integer matrix $\mathbf{x} \in \mathbb{N}_{\geq 0}^{n \times k}$, where \mathbf{x}_{ai} is the bid of admin a on subject slot i . We denote the budget of admin a by $B^{(a)} = \sum_i \mathbf{x}_{ai}$. An admin’s strategy is a potential bid allocation over user slots. We denote the set of all possible strategies of an admin a by \mathcal{D}_a .
- An allocation rule \mathcal{A} maps bids to an allocation of subject slots. That is, given a set of bids, the allocation rule determines an assignment of ranks for each admin. A rank assignment of a particular admin is a tuple $\bar{r} = (r_1, r_2, \dots, r_n)$, where r_i is the rank assigned for subject slot i . Under a probabilistic rule, we draw the rank assignment from the distribution given by the randomized allocation. The *sample profile* is just the count of the number of slots at each rank: $\bar{n} = (n_1, n_2, \dots, n_k)$, where $n_r = \sum_i \mathbb{1}\{r_i = r\}$. We denote the total amount of slots obtained by admin a by $n^{(a)} = \sum_k n_k$. The specific allocation rule we define later is randomized.
- The utility function $f_a: \mathcal{D} \rightarrow \mathbb{R}$ of admin a maps a tuple of strategies from $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2 \cdots \mathcal{D}_k$ to the utility of admin a given the strategies of all admins.

Our goal is to understand the Nash equilibria of the game. These are the strategies in which each admin is unilaterally maximizing her utility. In other words, admins are simultaneously best responding to each other. Specifically, we want to understand what estimation error admins can hope to achieve at equilibrium.

2.1. Objective Design

Given an assignment of ranks \bar{r} drawn from the distribution induced by the allocation function \mathcal{A} , an administrator has to decide how to use the samples so as to optimally estimate the causal effect τ . This corresponds to solving the optimization problem, from the perspective of a single admin:

$$\inf_{\hat{\tau}} \mathbb{E}_{Z \sim \mathcal{X}(\bar{r})} \left[(\hat{\tau}(Z) - \tau)^2 \right] \quad (1)$$

over all possible estimators $\hat{\tau}$. Here, $\mathcal{X}(\bar{r})$ denotes the data-generating distribution under the realized rank assignment \bar{r} ,

defined as: given a population \mathcal{X} and a ranking assignment \bar{r} , a subject is drawn independently at each subject slot i from the interventional distribution $\mathcal{X}^{\text{do}(\text{rank}=r_i)}$, where r_i is the rank that the admin obtained at subject slot i (coordinate i in the assignment rank vector \bar{r}). The data Z drawn from $\mathcal{X}(\bar{r})$ is the collection $Z_i = (r_i, T_i, Y_i), i \in [n^{(a)}]$, where T_i is the treatment assignment and Y_i denotes the outcome of a subject sampled for slot i at rank r_i . We formalize the data-generating process, together with the modeling choices for the outcome and treatment assignment in Section 3. Recall that treatment effects vary with the rank at which an admin’s campaign is run, and therefore the data-generating distribution depends on the rank assignment.

Estimation error objective. Solving objective (1) directly is a difficult task. We therefore replace the optimal error with a minimax lower bound on the error. We prove (in Section 3) that for any estimator $\hat{\tau}$, there exists a model instance M of the data-generating distribution $\mathcal{X}(\bar{r})$ such that the estimation error is lower bounded by

$$\min \left(c \cdot \left(\sum_{r=1}^k n_r \cdot \alpha_r^2 \right)^{-1}, 1 \right) \quad (2)$$

for a constant $c > 0$. In our subsequent analysis, we set w.l.o.g. $c = 1$.² When the number of samples obtained by an admin approaches 0, the first term in the bound expressed in equation (2) tends to infinity. However, in this case, since the causal effect is bounded, $|\tau| \leq 1$, the admin’s mean squared error is also bounded by 1 (this bound is given by choosing the constant estimator, for example). We provide a detailed discussion of this property in the Appendix. We show (in Section 3) that the minimax lower bound in (2) is attainable up to constant factors. For example, we can use an optimally weighted average of Horvitz-Thompson treatment effect estimators at each rank, which render an estimation error of $\mathcal{O} \left(\left(\sum_{r=1}^k n_r \cdot \alpha_r^2 \right)^{-1} \right)$. This motivates defining the **estimation error objective** as

$$f_a(\mathbf{x}) = -\mathbb{E}_{\bar{r} \sim \mathcal{A}(\mathbf{x})} \min \left(\mathbb{E}_{\mathcal{X}(\bar{r})} \left(\sum_{r=1}^k n_r \cdot \alpha_r^2 \right)^{-1}, 1 \right). \quad (3)$$

The negative sign serves the purpose of turning an estimation error minimization objective into a maximization objective. In essence, the estimation error objective captures an admin’s goal of minimizing their estimation error (up to constant factors) in expectation over the sample profile distribution given by the allocation rule.

²The value of the constant c does not change the results: as all admins minimize the error by selecting the best sample distribution, the constant c can be ignored in this objective. We obtain identical results using the expression $\min \left(\left(\sum_{r=1}^k n_r \cdot \alpha_r^2 \right)^{-1}, 1/c \right)$ in our objectives. Our results only rely on this expression being bounded by a constant.

Sample value objective. Even the estimation error objective is not necessarily tractable. In particular, it appears difficult to reason about Nash equilibria with respect to the estimation error objective. One of our main contributions is to relate the estimation error objective to a tractable objective. Given a sample profile, we call the weighted sum $\sum_r n_r \cdot \alpha_r^2$ the *sample value* given by a rank allocation \bar{r} . The **sample value objective** is defined as

$$f_a(\mathbf{x}) = \mathbb{E}_{\bar{r} \sim \mathcal{A}(\mathbf{x})} \mathbb{E}_{\mathcal{X}(\bar{r})} \left[\sum_{r=1}^k n_r \cdot \alpha_r^2 \right]. \quad (4)$$

Intuitively, the sample value objective and the estimation error objective are linked, as they are both a function of the sample value $\sum_{r=1}^k n_r \cdot \alpha_r^2$, obtained by an admin through a strategy choice. Formally, we can show that for typical parameter settings, any Nash equilibrium in the sample value objective is an approximate Nash equilibrium in the estimation error objective. Moreover, we can characterize a set of Nash equilibria in the sample value objective.

2.2. Main Results

We denote by \mathcal{G}_{MSE} and \mathcal{G}_{SV} the games associated to the estimation error objective and the sample value objective, respectively. Figure 1 illustrates our setup and results: through a causal inference analysis, we find an explicit expression for the minimax lower bound of an estimator as a function of the sample value (Section 3); the sample value objective’s associated game can be solved through finding a Nash equilibrium (Theorem 2.1), described in detail in Section 4, and any such solution approximates the error minimization objective under mild assumptions (Theorem 2.2), which is the main result of our paper.

Theorem 2.1. *A budget allocation game \mathcal{G}_{SV} with multiple treatment administrators has a pure Nash equilibrium for all sequences $(\alpha_r)_r$ such that $\alpha_r \in (0, 1)$ and $\alpha_1 = 1$. Moreover, for $B^{(a)} \leq n$, an allocation in which all treatment administrators split their entire budget uniformly on subjects and maximally avoid competition on subjects is always a Nash equilibrium.*

Theorem 2.2. *Any Nash equilibrium for the \mathcal{G}_{SV} game with k treatment administrators and n subjects is an (ϵ, η) -approximate Nash equilibrium for the \mathcal{G}_{MSE} game with*

$$\epsilon = O \left(\frac{k}{\sqrt{B}} \right) \quad \text{and} \quad \eta = O \left(\exp(-B/k^2) \right),$$

where $B = \min_a B^{(a)}$ and the maximum bid per subject slot is bounded.

In particular, if all administrators have budget $B^{(a)} = \omega(k^2)$, we get a $(o(1), o(1))$ -approximate Nash equilibrium. The proofs for all results are detailed in the Appendix.

Nash equilibrium. We assume that players act rationally and selfish in maximizing their utilities. We assume this is a




Figure 1. Illustration of the objective set-up and results.

complete information game (i.e. players know each other’s utility functions). This set-up is common in game designs applicable in online advertising settings, as motivated by prior works (Maehara et al., 2015; Varian, 2007). We employ the well-known notions of *pure* and *approximate* Nash equilibrium. In short, a strategy set \mathbf{x} is a pure Nash equilibrium (Nash Jr, 1950) if no deviation from it can increase utility for any admin (i.e. every admin’s strategy is a best response). The concept of (ϵ, η) -approximate equilibria has been widely used in finding approximate solutions (Alon et al., 2012; Roughgarden, 2016). In its essence, it assumes that an admin will not change his strategy if his current one is an approximate best response.

Definition 2.3. A strategy set $\mathbf{x} \in \mathcal{D}$ is an (ϵ, η) -approximate Nash equilibrium for a utility-maximization game \mathcal{G} with non-positive utility function $f_a : \mathcal{D} \rightarrow \mathbb{R}_-$ if for all players a ,

$$f_a(x_a, x_{-a}) \geq (1 + \epsilon) \cdot f_a(x'_a, x_{-a}) - \eta, \forall x'_a \in \mathcal{D}_a \quad (5)$$

where \mathcal{D}_a denotes the strategy set of player a and $\epsilon, \eta > 0$.

An equivalent definition for a utility minimization game with non-negative utilities occurs for $\eta < 0$.

Allocation rule: We use a probabilistic allocation rule that generalizes a model proposed by Maehara et al. (2015). We defer a formal definition to Section 4. Intuitively, each administrator has a chance to win the first rank of any given subject slot with a probability that is an increasing function of the administrator’s bid for the slot. After determining a winner for rank r , the allocation rule proceeds to the next rank until all ranks have been allocated. The main properties of the allocation rule are that the probability for an admin of winning a rank is component-wise concave and increasing in the bid, which facilitates characterizing a Nash equilibrium. Furthermore, the allocation rule benefits from the property that the sample value does not grow in expectation slower than $B^{(a)}/k$: $\mathbb{E} \left[\sum_{r=1}^k n_r \cdot \alpha_r^2 \right] = \Omega(B^{(a)}/k)$ under a Nash equilibrium for \mathcal{G}_{SV} . This allows the sample value variable to concentrate around its mean in the regime where budgets are larger than k^2 .

Proof sketch for Theorem 2.2: The intuition behind the proof relies on a few properties of the model and the alloca-

tion rule. Mainly, the allocation allows the sample objective to concentrate around its expectation with high probability in the bounded parameter regime. This allows us to show that if an allocation is indeed an optimal solution for the sample value objective, it will be approximately optimal for the estimation error objective. The ‘failure’ to concentrate probability determines a distance η from the utility of any other allocation; through a market-efficiency interpretation, it determines the necessary overhead above which an admin would be incentivized to change its strategy. We discuss alternative allocation rules and connections to common auction designs in Section 5, noting that our analysis is not contingent on the specific choice of an allocation rule, but rather on the properties mentioned in the proof sketch.

3. A Causal Inference Analysis of the Error Minimization Objective

We treat the estimation problem from the perspective of a particular admin a . For simplicity, we drop in the notation the index a when there is no confusion. We define a population distribution \mathcal{X} from which subjects will be independently drawn and used in the estimation problem for the admin. The admin runs an RCT on each subject and uses the data they gathered to estimate a causal effect. We employ classic modeling choices from the causal inference literature to model the outcomes and treatment assignments of subjects under an RCT.

Data generating process. We detail the data-generating process for an admin a , briefly described in Section 2. The process consists of two parts, sketched in Figure 2 in the Appendix:

1. Admin a competes over the n subject *slots* with the other admins by placing bids. A set of bids \mathbf{x} placed by the admins induce a distribution of rank allocations $\mathcal{A}(\mathbf{x})$, where \mathcal{A} is the rank allocation function. We sample a rank assignment \bar{r} from $\mathcal{A}(\mathbf{x})$ for admin a . The rank assignment \bar{r} is a vector containing the assigned rank of admin a for each subject slot. The sample profile (n_1, n_2, \dots, n_k) gives the number of subjects drawn at each rank, on whom the admin will run their campaign. In total, the admin obtained

$n^{(a)} = \sum_k n_k$ slots out of the total n slots.³

- The subjects are drawn independently at each slot from the interventional distribution $\mathcal{X}^{\text{do}(\text{rank}=r_i)}$, where r_i is the rank that admin a obtained at subject slot i and treatment or control are allocated in a randomized manner (through an RCT). The data the admin collects consists of the tuples (r_i, T_i, Y_i) : rank r_i , treatment assignment T_i , and outcome Y_i of each subject $i \in [n^{(a)}]$. As the RCT eliminates the influence from any possibly confounding factors, the treatment effect at a specific rank r is identifiable (Pearl, 2009).

We point out that the budget allocation happens *prior* to the sampling process. In particular, the individuals for each slot are not yet realized when the ranking assignment occurs. This process ensures that the individuals are independently drawn. We use potential outcomes to model the outcome of each individual sampled after the second step of the data-generating process occurs, described as follows.

Potential outcomes model. We model the outcome of a subject $i \in n^{(a)}$ sampled from $\mathcal{X}^{\text{do}(\text{rank}=r_i)}$ as

$$Y_i = c_{i,0} + c_{i,1}^{(r_i)} \cdot T_i + \varepsilon_i, \quad r \in [k] \quad (6)$$

where T_i denotes the treatment assignment of subject i by the admin. We interpret $T_i = 1$ as assigning subject i to the treatment group, and $T_i = 0$ to the control group. We model $T_i \sim \text{Ber}(q_i)$, for $q_i \in [q, 1 - q]$ for a constant $q \in (0, 0.5]$, $\forall i \in [n^{(a)}]$. Essentially, we assume the admin implements a non-uniform Bernoulli randomized design. The observational noise ε_i is drawn from a normal distribution and does not depend on the subject's rank,

$$\varepsilon_i \sim N(0, \sigma^2).$$

We denote the potential outcome of subject i for being in the treatment and control groups by $Y_i(1)$ and $Y_i(0)$, respectively.

Estimand: ATE but only for the first rank. We are interested in the treatment effect an administrator would have obtained if they had experienced no competition from other administrators. In reality, competition causes an admin to show up at lower positions for some slots. The lower the rank for a slot, the smaller the treatment effect. The causal effect of interest is thus the *average treatment effect had it been the first treatment applied*. We formally define the estimand as

$$\tau = \mathbb{E}[Y_i(1) - Y_i(0) | \text{do}(\text{rank} = 1)], \quad (7)$$

³Each admin may obtain a different number of data slots. For example, if an admin bids 0 on a slot, they will not obtain any rank for that slot; an admin bidding all 0s will not obtain any slots, and therefore $n^{(a)} = 0$ for him.

where the expectation is taken over the intervened population $\mathcal{X}^{\text{do}(\text{rank}=1)}$ that always encounter the content of admin a first.

Assumption 3.1 (Bounded effect). We assume that the treatment effect at the first rank τ is bounded with $|\tau| \leq 1$.

Given that a treatment administrator may acquire data from subjects at different ranks, we also define the average treatment effect of an administrator at rank $r \in [k]$ as

$$\tau_r = \mathbb{E}[Y_i(1) - Y_i(0) | \text{do}(\text{rank} = r)]. \quad (8)$$

Note that we identify the treatment effect at rank r by $\tau_r = \frac{1}{n_r} \sum_{i:r_i=r} c_{i,1}^{(r)}$ if $n_r \geq 1$.

Assumption 3.2 (Position effect). We formalize the assumption that the treatment effect at a lower rank is lower than at the first rank as

$$\tau_r = \alpha_r \cdot \tau, \forall r \in [k] \quad (9)$$

with $(\alpha_r)_r$ satisfying $\alpha_1 = 1$ and $\alpha_r \in (0, 1)$ for all $r \in [k]$.

An admin wishing to estimate τ may of course just use the data from samples at rank 1 (n_1 samples). However, the position effect assumption allows the admin to identify the treatment effect at the first rank *from* the treatment effect at rank $r \in [k]$ and thus, to use the entire sample profile,

$$\tau = \frac{\tau_r}{\alpha_r}.$$

We refer to

$$\tau^{(r)} := \frac{\tau_r}{\alpha_r} \quad (10)$$

as the r -estimand.

Remark 3.3. Estimating the treatment effect at rank 1 is a powerful primitive, as it allows one to extrapolate the effect at any *distribution* over ranks, given the discount factors $(\alpha_r)_r$. As mentioned in the introduction, this knowledge is important in order for an admin to gain insight into the quality of their treatment: was a lower effect observed because of a lower rank, or because of the treatment itself? Equipped with this knowledge, an admin can choose the best course of action: increase their bid or improve on the treatment design (e.g. in the case of online advertising, change the design of an ad). We emphasize that the admin has prior knowledge of the parameters $(\alpha_r)_r$ in this analysis. We discuss the case of jointly estimating τ and $(\alpha_r)_r$ later in this section.

3.1. An Optimality Argument

The problem an admin is solving is now that of optimally using their collected data in order to achieve a minimum estimation error. Exactly *how* an admin should combine

their samples from different ranks towards minimizing the estimation error is non-trivial. We show however that the estimation error of any possible estimator is bounded below by a function of the sample value $\sum_{r=1}^k n_r \cdot \alpha_r^2$.

To formalize this result, we denote the set of our potential outcomes model instances as \mathcal{M} . In particular,

$$\mathcal{M} = \left\{ \left(c_{i,0}, c_{i,1}^{(r_i)} \right)_{i \in [n^{(a)}]} : |\tau| \leq 1, \tau_r = \alpha_r \cdot \tau \right\}. \quad (11)$$

Theorem 3.4 (Minimax Lower Bound). *For any estimator $\hat{\tau}$, there exists an instance $M \in \mathcal{M}$ such that the minimax squared error is bounded below by*

$$\mathbb{E}_M[(\hat{\tau} - \tau)^2] \geq \min \left(\frac{\sigma^2}{16 \cdot (1-q) \cdot \sum_{r=1}^k n_r \cdot \alpha_r^2}, 1 \right).$$

We provide the detailed proof in the Appendix. The proof relies on Le Cam's method (Le Cam, 1973) for hypothesis testing and on leveraging the assumption of independence between individuals. The particular choice of an optimal estimator for τ is beyond the scope of this paper. Instead, we showcase an estimator that achieves the minimax lower bound (up to a constant).

Estimator examples. We find examples of estimators that aggregate the effect at each rank and obtain an error upper bound of $\mathcal{O} \left(\left(\sum_{r=1}^k n_r \cdot \alpha_r^2 \right)^{-1} \right)$. Estimators that aggregate the effect at each rank are often used in meta-analysis studies or stratified experiments (Hartung et al., 2011; de Chaisemartin, 2022). We refer to these as *decomposable* estimators. Formally, a decomposable estimator can be written as $\hat{\tau} = \sum_{r=1}^k \omega_r \cdot \hat{\tau}_r^{(r)}$ with $\hat{\tau}_1^{(r)}$ an estimator for the r -estimand for $r \in [k]$. For unbiased estimators $\hat{\tau}_1^{(r)}$, the well-known inverse variance estimator (Markowitz, 1952; 1959; Cochran, 1954; Shahar, 2017) provides a closed form solution for optimally weighing the r -estimators while preserving unbiasedness. The optimal weights ω^* are obtained by solving a variance minimization problem through Lagrange multipliers,

$$\omega_r^* = \frac{1}{\text{Var}(\hat{\tau}_1^{(r)})} \cdot \frac{1}{\sum_{r=1}^k 1/\text{Var}(\hat{\tau}_1^{(r)})}. \quad (12)$$

An immediate corollary is that the minimum MSE achieved for weights ω^* is

$$\mathbb{E}[(\tau_1 - \hat{\tau}_{\omega^*})^2] = \frac{1}{\sum_{r=1}^k 1/\text{Var}(\hat{\tau}_1^{(r)})}. \quad (13)$$

Within our potential outcomes model, there are unbiased estimators for the r -estimand that achieve an

error $\mathbb{E}[(\hat{\tau}_r - \tau_r)^2] = \mathcal{O}(n_r^{-1})$. Any such estimators, weighted optimally, would attain a total error bounded by $\mathcal{O} \left(\left(\sum_{r=1}^k n_r \cdot \alpha_r^2 \right)^{-1} \right)$, as we will further argue. Given the parametric assumption on the sequence $(\alpha_r)_r$, we know that α_r amplifies the variance of the r -estimator by a squared term,

$$\text{Var}(\hat{\tau}_1^{(r)}) = \frac{\text{Var}(\hat{\tau}_r)}{\alpha_r^2} \quad (14)$$

and thus, $\mathbb{E}[(\hat{\tau}_1^{(r)} - \tau_1^{(r)})^2] = \mathcal{O}((n_r \cdot \alpha_r^2)^{-1})$. It then immediately follows from equation (13) that the estimation error of the inverse variance estimator $\hat{\tau}_{\omega^*}$ is

$$\mathbb{E}[(\tau_1 - \hat{\tau}_{\omega^*})^2] = \mathcal{O} \left(\left(\sum_{r=1}^k n_r \cdot \alpha_r^2 \right)^{-1} \right) \quad (15)$$

As examples of unbiased estimators for the r -estimand, the OLS estimator and the Horvitz-Thompson estimator (Horvitz & Thompson, 1952) have variance of order $\mathcal{O}(n_r^{-1})$. The OLS estimator attains the lowest actual MSE, including constants, for unbiased linear estimators (see Gauss-Markov Theorem in Johnson & Wichern (2007)), while the Horvitz-Thompson estimator is preferred in more general models that include unequal treatment probability (unequal q_i across subjects). We do not reproduce the bound computation for the OLS estimator as it is a known result, but we detail, for the interested reader, properties of the Horvitz-Thompson estimator in our potential outcomes model.

We define the Horvitz-Thompson estimator formally for estimating the effect at rank r ,

$$\hat{\tau}_r = \frac{1}{n_r} \cdot \sum_{i:r_i=r} \left(\frac{Y_i \mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{Y_i \mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right) \quad (16)$$

Similar to the OLS estimator, it is unbiased and has a variance of order $\mathcal{O}(n_r^{-1})$ (albeit with a larger constant factor). We show these properties in Propositions 3.5 and 3.6.

Proposition 3.5. *The estimator $\hat{\tau}_r$ is an unbiased estimator for the average treatment effect τ_r at rank r .*

Proposition 3.6. *The error of the rank r estimator $\hat{\tau}_r$, $\mathbb{E}[(\tau_r - \hat{\tau}_r)^2]$, can be upper bounded by*

$$\frac{1}{n_r} \cdot \left[\frac{\sigma^2}{q(1-q)} + Y_{\max}^2 \cdot \left(\frac{(1-q)^2 + q^2}{q(1-q)} + 2 \right) \right] \quad (17)$$

with $Y_{\max} := \max_{i \in [n^{(a)}]} |c_{i,0} + c_{i,1}^{(r_i)}|$.

These results show that we can find an optimal estimator up to constants. While we do not require decomposable

estimators in order to obtain the minimax lower bound on the error in Theorem 3.4, they provide a good intuition for the form of the sample value expression. The intuition is that whereas in a classical non-competitive causal inference task, the sample value is just a function of the sample size, the α_r downgrade in treatment effect from rank 1 to rank r scales down the sample value of rank r by α_r^2 . Thus, we interpret the sum $\sum_{r=1}^k n_r \cdot \alpha_r^2$ as the sample value that an admin has in order to estimate a treatment effect. The choice of different strategies offers an admin the possibility of *optimizing* the sample value.

3.2. Estimating the Discount Factors

In our analysis so far, we assumed that the discount factors $(\alpha_r)_r$ are known by the admins. The parameters are perhaps given by a central platform or estimated in previous studies.⁴ In settings where the discount factors $(\alpha_r)_r$ are not *a priori* known, they may be estimated from historical data. Related set-ups have utilized the EM algorithm in estimating the position bias in online advertisement problems (Chuklin et al., 2016; Dempster et al., 1977). In many applications of multiple treatments, such as online advertisement, the estimated discount factors do not depend on the particular admin but rather on the display position (Lagrée et al., 2016). In these cases, a pilot experiment is usually run by randomizing results or performing pairwise comparisons (Joachims et al., 2017; Wang et al., 2018), in order to get good estimates of the discount factors. The estimated discount factors are then used in subsequent campaigns, under the assumption that they do not need to be re-estimated.

For completeness, we show that estimating the discount factors jointly with the treatment effect through data splitting does not reduce the estimation error, even in the most simple case: take a pilot experiment in which an online platform has collected data on an admin’s treatments at different ranks, with n_r samples at each rank. Without prior knowledge of the position effect parameters, the admin has two choices in estimating τ : (1) use just the data where their treatment was shown on the first rank, or (2) split the data at all positions, using part of it for estimating the position effect parameters $(\alpha_r)_r$, and then using these estimates together with the rest of the data for estimating the treatment effect τ . Even if the admin has to split the data available in (2), he might still have more samples for the treatment estimation task than in scenario (1). We show in a simple argument that the treatment effect estimation error is not reduced in (2) as compared to (1), despite the fact that the admin is using more data in (2) and no matter *how* the split is being done (Lemma A.6 in the Appendix).

⁴For example, Google reports the decay rate in click-through-rates with lower positions in online advertising (Bailyn, 2024).

4. Nash Equilibria for the Sample Value Game

We define a probabilistic allocation rule $\mathcal{A}_{\text{prob}}$ and prove Theorem 2.1 by characterizing a set of pure Nash equilibria for the game \mathcal{G}_{SV} .

Allocation rule $\mathcal{A}_{\text{prob}}$: We use a probabilistic allocation rule that generalizes the model proposed in Maehara et al. (2015). We define a probabilistic allocation rule called $\mathcal{A}_{\text{prob}}$ that allocates ranks to the k admins as follows.

- *Activation probability:* When an admin i allocates budget to a subject slot, we define his probability of winning that subject slot at a particular rank (determined by the order in which admins play) as

$$\text{PA}_a(\mathbf{x}_a, t) := 1 - (1 - p)^{\mathbf{x}_a(t)}, \quad (18)$$

where p is a parameter governing how relevant admin a is for a subject slot t .⁵ We use interchangeably the notation $\mathbf{x}_a(t)$ or \mathbf{x}_{at} for the budget that treatment administrator a allocates to slot t . We note that the relevance parameter p cannot depend on the specific subject drawn at subject slot t , since the admin does not know *a priori* which subject is realized. Thus, we can think of p as an expected value over the population distribution. If a treatment administrator does not bid on a slot, he does not get a sample from that slot (at any rank).

- *Random ordering:* Then, we must define a way for advertisers to compete over the subject slots. We propose the following procedure for administrators to get allocated ranks to show their content to each subject slot: for each subject slot $t \in [n]$ and for each rank $r \leq k$, set \mathcal{S}_{k-r} the set of administrators who have not been allocated a rank $r' < r$ yet for this subject slot t . Then, draw a permutation of the elements of \mathcal{S}_{k-r} , $\sigma \in \mathcal{P}(\mathcal{S}_{k-r})$, uniformly at random (where $\mathcal{P}(S)$ denotes the set of all permutations of elements of a set S). Now, according to the order that the permutation σ defined, each administrator will try to win rank r for subject slot t , based on the budget they allocated to this subject slot (note that an administrator allocates budget to a subject, not to each rank). See Algorithm 1 for a formal description.

We characterize a set of Nash equilibria for the case $B^{(a)} \leq n$ in the following lemma, thus proving Theorem 2.1. The results easily generalize to $B^{(a)} > n$, by considering allocations in which treatment administrators split their budget uniformly over subjects.

⁵In this first analysis, we take p to be the same for all admins and subjects; future work can model a different p_a for each admin a , noting that the equilibrium becomes more complicated to characterize in closed form.

Algorithm 1 Winning ranks for a subject slot t through the probabilistic allocation rule $\mathcal{A}_{\text{prob}}$.

```

Denote  $\mathcal{S}_{k-1} = [k]$  the set of all administrators.
for rank  $r \in [k]$ : do
    Draw a random permutation  $\sigma \in \mathcal{P}(\mathcal{S}_{k-r})$ 
    for  $i \in [k]$ : do
        Admin  $\sigma(i)$  tries to win subject  $t$  at rank  $r$  with
        probability  $\text{PA}_{\sigma(i)}(x_{\sigma(i)}, t)$ 
        if success then
            Define  $\mathcal{S}_{k-r-1} = \mathcal{S}_{k-r} \setminus \sigma(i)$ 
            break and move on to the next rank
        end if
    end for
end for
    
```

Lemma 4.1. An allocation \mathbf{x} with the following properties is always a Nash equilibrium for \mathcal{G}_{SV} :

- Admins prefer to split their budget uniformly over subjects: $\mathbf{x}_a(t) \in \{0, 1\}, \forall a \in [k], t \in [n]$;
- Admins prefer to spend all their budget:

$$\sum_{t=1}^n \mathbf{x}_a(t) = B^{(a)}, \forall a \in [k]$$

- Admins prefer to minimize competition with each other:

$$\left| \sum_{a=1}^k \mathbf{x}_a(t) - \sum_{a=1}^k \mathbf{x}_a(t') \right| \leq 1, \forall t, t' \in [n]$$

This means that an admin will prioritize spending budget on the subjects with the least bids.

The proof for Lemma 4.1 is detailed in the Appendix and relies on the following properties for an allocation \mathbf{x} as described: an admin a (1) cannot increase their utility through permuting his bids (re-arranging his allocation \mathbf{x}_a of budget units over subjects); and (2) prefers to split his budget in units of 1 over subjects rather than aggregating it over fewer subjects. The intuition behind proving these properties relies on the form of the utility function $f_a(\cdot)$, as $\text{PA}_a(\cdot)$ is component-wise concave and increasing: an admin gains more from uniform bidding since it has a chance of winning a first rank over multiple subjects, rather than increasing his bid on a single subject. Although increasing the bid on a single subject increases his chance to win the first rank, it does not increase it by much in comparison. Moreover, we note that an allocation in which an admin has unused budget cannot be an equilibrium, as the admin can increase their sample value by bidding on new subjects, given the probabilistic nature of the allocation rule. The nature of this equilibrium is quite subtle: it's not clear from the start whether

the equilibrium is always the equal bids strategy, or to split the subject slots (bidding high on a few subject slots and zero on the rest, thus splitting the subjects between different admins). We conjecture that the allocations \mathbf{x} as described in Lemma 4.1 are the only Nash equilibria that can occur when the relevance probability p is the same for all admins:

Conjecture 4.2. The only pure Nash equilibria for the game \mathcal{G}_{SV} are allocations \mathbf{x} that satisfy the properties stated in Lemma 4.1, when the activation probability $p_a := p$ is the same for all admins $a \in [k]$.

We can formally show this property for two admins, $k = 2$, and any number of subjects n :

Proposition 4.3. The only pure Nash equilibria for the game \mathcal{G}_{SV} for two admins ($k = 2$) are allocations \mathbf{x} that satisfy the properties stated in Lemma 4.1, when the activation probability $p_a := p$ is the same for all admins $a \in [k]$.

5. Discussion and Future Directions

While our analysis uses the probabilistic allocation rule $\mathcal{A}_{\text{prob}}$, it is not limited to it. As the approximation proof sketch suggests, a main ingredient is that the inherent randomness in the allocation rule renders the sample value objective S close to its mean, allowing for concentration inequalities to apply and extend to the estimation error objective. Albeit the random ordering of the agents, the allocation rule has the property that higher bids increase the probability of winning higher ranks, due to the sequential rank allocation. The Nash equilibria in the sample value optimization game bear a noteworthy resemblance to Feldman et al. (2007)'s randomized solution of the associated optimization problem under a generalized second-price auction. In their setting, the authors show that the optimization problem induced by the utility function of agents with position-based value over subjects is intractable, yet randomized strategies find nearly optimal solutions. Thus, future directions could analyze our objectives with other allocation rules similar to those used in generalized second-price or first-price auctions (Vickrey, 1961).

Another promising future avenue of research includes studying the quality of Nash equilibria within the class of estimation error games, in terms of bounding the Price of Anarchy (PoA) (Koutsoupas & Papadimitriou, 2009). In the \mathcal{G}_{SV} game, this becomes more challenging to study as it is not necessarily a potential game like its simpler version in Maehara et al. (2015), although with the added novelty that we can fully characterize a Nash equilibrium. Furthermore, it would be interesting to study the distribution of utility across admins through the lens of inequality, perhaps through related solution concepts such as Nash social welfare, often employed to study fairness issues in resource allocation games (Brânzei et al., 2017).

Acknowledgements

We thank Florian Dorner for his insightful feedback on the manuscript. We are also grateful to Claire Vernade, Aaron Schein, Kevin Jamieson, Lalit Jain, Robert Nowak, and Scott Sievert for helpful discussions on earlier versions of the project. Finally, we thank the anonymous reviewers for their helpful feedback on the paper. This project has received funding from the Max Planck ETH Center for Learning Systems (CLS).

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Agarwal, A., Hosanagar, K., and Smith, M. D. Location, location, location: An analysis of profitability of position in online advertising markets. *Journal of marketing research*, 48(6):1057–1073, 2011.
- Aggarwal, G., Goel, A., and Motwani, R. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pp. 1–7, 2006.
- Aggarwal, G., Badanidiyuru, A., and Mehta, A. Autobidding with constraints. In *Web and Internet Economics: 15th International Conference, WINE 2019, New York, NY, USA, December 10–12, 2019, Proceedings 15*, pp. 17–30. Springer, 2019.
- Alatas, V., Banerjee, A., Hanna, R., Olken, B. A., and Tobias, J. Targeting the poor: evidence from a field experiment in Indonesia. *American Economic Review*, 102(4):1206–1240, 2012.
- Alon, N., Gamzu, I., and Tennenholtz, M. Optimizing budget allocation among channels and influencers. In *Proceedings of the 21st International Conference on World Wide Web*, pp. 381–388, 2012.
- Ansari, A. and Mela, C. E-customization. *Journal of Marketing Research*, 40(2):131–145, 2003.
- Athey, S. and Ellison, G. Position auctions with consumer search. *The Quarterly Journal of Economics*, 126(3): 1213–1270, 2011.
- Bailyn, E. Google click-through rates (CTRs) by ranking position. 2024. URL <https://firstpagesage.com/reports/google-click-through-rates-ctrs-by-ranking-position/>.
- Balseiro, S., Deng, Y., Mao, J., Mirrokni, V., and Zuo, S. Robust auction design in the auto-bidding world. *Advances in Neural Information Processing Systems*, 34: 17777–17788, 2021.
- Banerjee, A. V., Duflo, E., and Kremer, M. The influence of randomized controlled trials on development economics research and on development policy. *The state of Economics, the state of the world*, pp. 482–488, 2016.
- Börgers, T., Cox, I., Pesendorfer, M., and Petricek, V. Equilibrium bids in sponsored search auctions: Theory and evidence. *American economic Journal: microeconomics*, 5(4):163–187, 2013.
- Brânzei, S., Gkatzelis, V., and Mehta, R. Nash social welfare approximation for strategic agents. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 611–628, 2017.
- Charles, D., Chakrabarty, D., Chickering, M., Devanur, N. R., and Wang, L. Budget smoothing for Internet ad auctions: a game theoretic approach. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pp. 163–180, 2013.
- Chuklin, A., Markov, I., and de Rijke, M. Click Models for Web Search and their Applications to IR. In *WSDM '16: Proceedings of the 9th ACM International Conference on Web search and data mining*, pp. 689–690, New York, New York, USA, 2016. ACM Press.
- Cochran, W. G. The combination of estimates from different experiments. *Biometrics*, 10(1):101–129, 1954.
- Conitzer, V., Kroer, C., Panigrahi, D., Schrijvers, O., Stier-Moses, N. E., Sodomka, E., and Wilkens, C. A. Pacing equilibrium in first price auction markets. *Management Science*, 68(12):8515–8535, 2022.
- Cortez-Rodriguez, M., Eichhorn, M., and Yu, C. L. Exploiting neighborhood interference with low-order interactions under unit randomized design. *Journal of Causal Inference*, 11(1):20220051, 2023.
- de Chaisemartin, C. Trading-off bias and variance in stratified experiments and in staggered adoption designs, under a boundedness condition on the magnitude of the treatment effect. Technical report, National Bureau of Economic Research, 2022.
- de Souza Leão, L. and Eyal, G. The rise of randomized controlled trials (RCTs) in international development in historical perspective. *Theory and Society*, 48:383–418, 2019.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. Maximum likelihood from incomplete data via the EM algorithm.

- Journal of the royal statistical society: series B (methodological)*, 39(1):1–22, 1977.
- Duflo, E. Field experiments in Development Economics. *Econometric Society Monographs*, 42:322, 2006.
- Dwivedi, R., Tian, K., Tomkins, S., Klasnja, P., Murphy, S., and Shah, D. Counterfactual inference for sequential experiments. *arXiv preprint arXiv:2202.06891*, 2022.
- Ebbinghaus, H. *Memory: A contribution to experimental psychology*. Dover, 1885.
- Edelman, B., Ostrovsky, M., and Schwarz, M. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1):242–259, 2007.
- Feldman, J., Muthukrishnan, S., Pal, M., and Stein, C. Budget optimization in search-based advertising auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, pp. 40–49, 2007.
- Ghosh, A. and Sayedi, A. Expressive auctions for externalities in online advertising. In *Proceedings of the 19th International Conference on World Wide Web*, pp. 371–380, 2010.
- Hadad, V., Hirshberg, D. A., Zhan, R., Wager, S., and Athey, S. Confidence intervals for policy evaluation in adaptive experiments. *Proceedings of the National Academy of Sciences*, 118(15):e2014602118, 2021.
- Hardt, M., Jagadeesan, M., and Mendler-Dünnner, C. Performative power. *Advances in Neural Information Processing Systems*, 35:22969–22981, 2022.
- Hartung, J., Knapp, G., and Sinha, B. K. *Statistical meta-analysis with applications*. John Wiley & Sons, 2011.
- Horvitz, D. G. and Thompson, D. J. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952.
- Joachims, T., Swaminathan, A., and Schnabel, T. Unbiased learning-to-rank with biased feedback. In *Proceedings of the tenth ACM International Conference on Web search and data mining*, pp. 781–789, 2017.
- Johari, R., Li, H., Liskovich, I., and Weintraub, G. Y. Experimental design in two-sided platforms: An analysis of bias. *Management Science*, 68(10):7069–7089, 2022.
- Johnson, R. A. and Wichern, D. W. *Applied multivariate statistical analysis*. Prentice hall Upper Saddle River, NJ, 2007.
- Koutsoupias, E. and Papadimitriou, C. Worst-case equilibria. *Computer science review*, 3(2):65–69, 2009.
- Lagrée, P., Vernade, C., and Cappe, O. Multiple-play bandits in the position-based model. *Advances in Neural Information Processing Systems*, 29, 2016.
- Le Cam, L. Convergence of estimates under dimensionality restrictions. *The Annals of Statistics*, 1(1):38–53, 1973.
- Levy, S. *Progress against poverty: sustaining Mexico’s Progres-Oportunidades program*. Rowman & Littlefield, 2007.
- Maehara, T., Yabe, A., and Kawarabayashi, K.-i. Budget allocation problem with multiple advertisers: A game theoretic view. In *International Conference on Machine Learning*, pp. 428–437. PMLR, 2015.
- Markowitz, H. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952. ISSN 00221082, 15406261. URL <http://www.jstor.org/stable/2975974>.
- Markowitz, H. M. *Portfolio Selection: Efficient Diversification of Investments*. Yale University Press, 1959. ISBN 9780300013726. URL <http://www.jstor.org/stable/j.ctt1bh4c8h>.
- McDiarmid, C. et al. On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188, 1989.
- McFarland, S. G. Effects of question order on survey responses. *Public Opinion Quarterly*, 45(2):208–215, 1981.
- Narayanan, S. and Kalyanam, K. Position effects in search advertising and their moderators: A regression discontinuity approach. *Marketing Science*, 34(3):388–407, 2015.
- Nash Jr, J. F. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.
- Pearl, J. *Causality*. Cambridge University Press, 2009.
- Roughgarden, T. *Twenty lectures on algorithmic game theory*. Cambridge University Press, 2016.
- Shahar, D. J. Minimizing the variance of a weighted average. *Open Journal of Statistics*, 7(2):216–224, 2017.
- Teevan, J. How people recall, recognize, and reuse search results. *ACM Transactions on Information Systems (TOIS)*, 26(4):Article No. 19., 2008.
- Varian, H. R. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.
- Vickrey, W. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

Wager, S. and Xu, K. Experimenting in equilibrium. *Management Science*, 67(11):6694–6715, 2021.

Wang, X., Golbandi, N., Bendersky, M., Metzler, D., and Najork, M. Position bias estimation for unbiased learning to rank in personal search. In *Proceedings of the eleventh ACM International Conference on Web search and data mining*, pp. 610–618, 2018.

Zhou, T., Liu, J., Jiao, Y., Dong, C., Chen, Y., Gao, Y., and Sun, Y. Bandit learning to rank with position-based click models: Personalized and equal treatments. *arXiv preprint arXiv:2311.04528*, 2023.

A. Appendix

A.1. Figure of Data-generating Process in Section 3

We sketch the data-generating process for an admin in Figure 2.

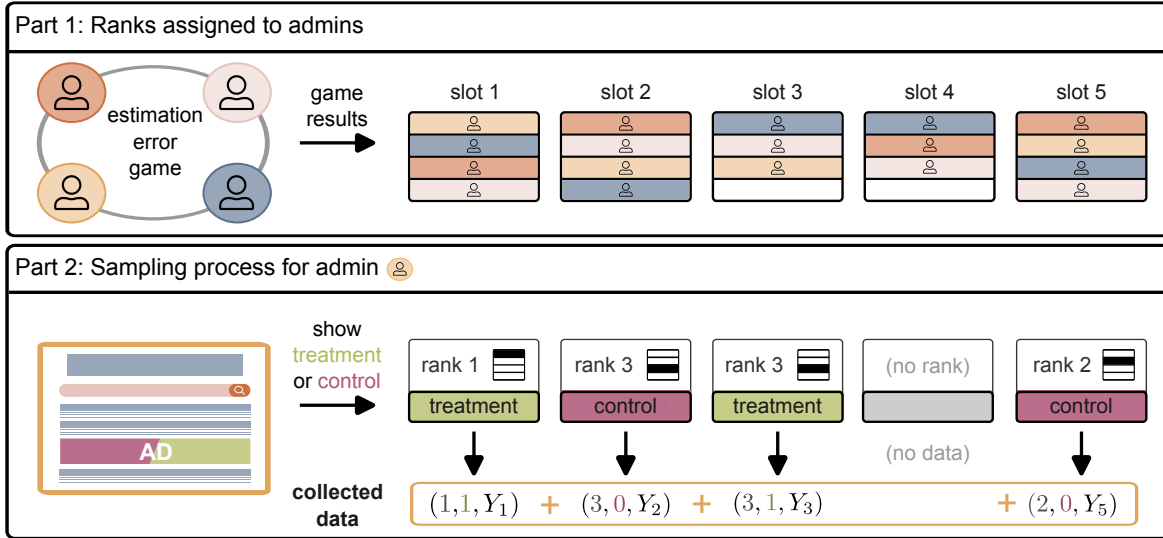


Figure 2. Illustration of the data-generating process for an admin. In Part 1, the admins compete with each other to obtain different ranks for subject slots. After placing the bids, a rank allocation rule is applied. Each admin receives a rank for all slots for which they have bid. For example, the admin depicted in yellow received rank 1 for subject slots 1, rank 3 for subject slots 2 and 3, no rank for subject slot 4 (which happens if an admin does not bid on a slot), and rank 2 for subject slot 5. In Part 2, individuals are sampled at the allocated ranks, with treatment and control assigned in a randomized manner. The collected data for the admin depicted in yellow consists of tuples (r_i, T_i, Y_i) for all individuals, where r_i represents the rank, T_i the binary treatment-control variable, and Y_i the outcome.

A.2. Proofs for Section 2 Results

Proof of Theorem 2.2. We assume that every admin a can only bid up to some constant C units of its budget $B^{(a)}$ on any particular subject slot t . This assumption is motivated by throttling behavior on online advertising markets (Charles et al., 2013) or advertising channel constraints in media (Maehara et al., 2015). We also assume that the budget of any admin is at most the number of subject slots: $B^{(a)} \leq n$. We note that this assumption is not restrictive, as results hold easily for larger budgets.

We aim to show that for \mathbf{x} a Nash equilibrium for the \mathcal{G}_{SV} game, \mathbf{x} is an (ϵ, η) -approximate Nash equilibrium for the \mathcal{G}_{MSE} game. For ease of notation, denote the sample value by a variable S :

$$S := \sum_{r=1}^k n_r \cdot \alpha_r^2. \quad (19)$$

The sample value S is determined by the allocation \mathbf{x} . Note that S is a positive random variable, as the variables $(n_r)_r$ counts of the rank allocation \bar{r} , $n_r = \sum_i \mathbb{1}\{r_i = r\}$ and the variables $(\alpha_r)_r$ are nonnegative. The rank allocation \bar{r} is sampled from the distribution given by the allocation rule. We aim to show the following result, as a stepping stone in proving the existence of an ϵ -approximate Nash equilibrium for the game \mathcal{G}_{MSE} :

Proposition A.1. *For allocations $\mathbf{x} \in \mathcal{D}$ in which all admins use all their budget, $\sum_{t=1}^n \mathbf{x}_a(t) = B^{(a)}$, $\forall a \in [k]$, with high probability $S \in [(1 - \epsilon) \cdot \mathbb{E}[S], (1 + \epsilon) \cdot \mathbb{E}[S]]$, with $\epsilon > 0$.*

Proof of Proposition A.1. The intuition is that the sample value concentrates around its expectation with high probability given that it satisfies a bounded differences property, as further shown. Our aim is to show that with high probability

$$\begin{aligned} S \in [(1 - \epsilon) \cdot \mathbb{E}[S], (1 + \epsilon) \cdot \mathbb{E}[S]] &\Leftrightarrow \\ (1 + \epsilon) \cdot \mathbb{E}[S] \geq S \geq (1 - \epsilon) \cdot \mathbb{E}[S] &\Leftrightarrow \\ \epsilon \cdot \mathbb{E}[S] \geq S - \mathbb{E}[S] \geq -\epsilon \cdot \mathbb{E}[S] &\Leftrightarrow \\ |S - \mathbb{E}[S]| \leq \epsilon \cdot \mathbb{E}[S] & \end{aligned} \quad (20)$$

In order to show this, we upper bound the probability $\mathbb{P}(|S - \mathbb{E}[S]| \geq \epsilon \cdot \mathbb{E}[S])$ using McDiarmid's inequality by writing the random variable S as the output of a function that satisfies the bounded differences property. An admin a only obtains samples from subject slots on which they bid, i.e. if they bid 0 budget on a subject slot, they do not get *any* rank for that subject slot, as per the definition of the allocation rule. Given the assumption that an admin can only bid up to some constant C on each subject, there are at least $B^{(a)}/C$ slots where the admin gets any rank; hence, there are at least $B^{(a)}/C$ of the r_i variables. We denote by m the number of the subject slots on which an admin gets ranks, noting that $B^{(a)}/C \leq m \leq B^{(a)}$.

For each sample profile $(n_r)_r$, we can write S as the output of a function of the random variables $(r_i)_i$, $g : [k]^m \rightarrow \mathbb{R}$, with $g(r_1, \dots, r_m) = \sum_{r=1}^k \sum_{i=1}^m \mathbb{1}(r_i = r) \cdot \alpha_r^2$ (in this notation, n_r becomes $n_r = \sum_{i=1}^m \mathbb{1}(r_i = r)$).

We next show that the function g satisfies the bounded differences property in the variables $(r_i)_i : \exists$ constants $c_i, \forall i \in [m]$ such that $\forall i \in [m]$ and $r_i \in [k]$, the following property is satisfied:

$$\sup_{r'_i \in [k]} |g(r_1, \dots, r_i, \dots, r_m) - g(r_1, \dots, r'_i, \dots, r_m)| \leq c_i \quad (21)$$

To show this, we denote

$$\begin{aligned} g(r_1, \dots, r_i, \dots, r_m) &= \sum_{r=1}^k n_r \cdot \alpha_r^2, \\ g(r_1, \dots, r'_i, \dots, r_m) &= \sum_{r=1}^k n'_r \cdot \alpha_r^2, \end{aligned} \quad (22)$$

where $n'_{r_i} = n_{r_i} - 1$, $n'_{r'_i} = n_{r'_i} + 1$, and for $r \neq r_i, r'_i$, $n_r = n'_r$, since changing one variable r_i to r'_i only changes two counts of the ranks, n_{r_i} and $n_{r'_i}$, leaving the rest unchanged. Thus, we get that

$$\sup_{r'_i \in [k]} |g(r_1, \dots, r_i, \dots, r_m) - g(r_1, \dots, r'_i, \dots, r_m)| = \sup_{r'_i \in [k]} |\alpha_{r_i}^2 - \alpha_{r'_i}^2| \leq 1, \quad (23)$$

since $\alpha_r \in [0, 1]$, $\forall r \in [k]$. Thus, the function g satisfies the bounded differences property with right-hand-side constants equal to 1 for all $i \in [n]$. From McDiarmid's inequality (McDiarmid et al., 1989), we get that for independent random variables $R_i \in [k]$, $\forall i \in [m]$,

$$\begin{aligned} \mathbb{P}(|g(R_1, \dots, R_m) - \mathbb{E}[g(R_1, \dots, R_m)]| \geq \eta) &\leq 2 \exp\left(-\frac{2\eta^2}{m}\right) \Leftrightarrow \\ \mathbb{P}(|S - \mathbb{E}[S]| \geq \eta) &\leq 2 \exp\left(-\frac{2\eta^2}{m}\right) \end{aligned} \quad (24)$$

Thus, the probability that the even in equation 20 happens is bounded below:

$$\mathbb{P}(|S - \mathbb{E}[S]| \leq \epsilon \cdot \mathbb{E}[S]) \geq 1 - 2 \exp\left(-\frac{2\epsilon^2 \mathbb{E}[S]^2}{m}\right) \quad (25)$$

As $\epsilon \rightarrow 0$, the right hand side of equation 25 approaches 1. □

Next, we will use the properties of the allocation rule $\mathcal{A}_{\text{prob}}$ that allow us to bound $\mathbb{E}[S]$:

Proposition A.2. *If $B^{(a)} \leq n$ for an admin a , then under the allocation rule $\mathcal{A}_{\text{prob}}$ and a Nash equilibrium \mathbf{x} ,*

$$\mathbb{E}\left[\sum_{r=1}^k n_r \cdot \alpha_r^2\right] = \Omega(B^{(a)}/k).$$

Corollary A.3. *If $B^{(a)} \leq n$ for an admin a , then under the allocation rule $\mathcal{A}_{\text{prob}}$ and an allocation \mathbf{x} with $\sum_{t=1}^n \mathbf{x}_a(t) = B^{(a)}$,*

$$\mathbb{E}\left[\sum_{r=1}^k n_r \cdot \alpha_r^2\right] = \Omega(B^{(a)}/C \cdot k).$$

Proof of Proposition A.2. First, we note by linearity of expectation and the fact that all variables are non-negative that $\mathbb{E}\left[\sum_{r=1}^k n_r \cdot \alpha_r^2\right] = \sum_{r=1}^k \mathbb{E}[n_r] \cdot \alpha_r^2 \geq \mathbb{E}[n_1] \cdot \alpha_1^2 = \mathbb{E}[n_1]$, as we set $\alpha_1 = 1$. The allocation rule $\mathcal{A}_{\text{prob}}$ gives us a closed form formula for $\mathbb{E}[n_1]$,

$$\mathbb{E}[n_1] = \sum_{t=1}^n \text{PA}_a(\mathbf{x}_a, t) \cdot \frac{1}{k!} \sum_{\sigma \in \mathcal{P}([k])} \text{PA}_a(\mathbf{x}_a, t) \cdot \prod_{j <_{\sigma} a} (1 - \text{PA}_j(\mathbf{x}_j, t)), \quad (26)$$

where $\mathcal{P}(K)$ denotes the set of permutation of elements of a set K . We note that each summand in equation (26) is the probability that admin a wins rank 1 for a subject slot t , which we write as

$$p_{a,t}(\mathbf{x}) = \text{PA}_a(\mathbf{x}_a, t) \cdot \frac{1}{k!} \sum_{\sigma \in \mathcal{P}([k])} \text{PA}_a(\mathbf{x}_a, t) \cdot \prod_{j <_{\sigma} a} (1 - \text{PA}_j(\mathbf{x}_j, t)), \quad (27)$$

thus giving

$$\mathbb{E}[n_1] = \sum_{t=1}^n p_{a,t}(\mathbf{x}) \quad (28)$$

As $p_{a,t}$ is essentially averaging over all possible orderings of the admins, we note that

$$p_{a,t}(\mathbf{x}) = \sum_{s=0}^{k-1} p_{a,t}(\mathbf{x} | \text{admin } a \text{ appears after } s \text{ admins in } \sigma) \cdot \mathbb{P}(\text{admin } a \text{ appears after } s \text{ admins in } \sigma), \quad (29)$$

where σ is a random ordering of the admins. Thus,

$$p_{a,t}(\mathbf{x}) \geq p_{a,t}(\mathbf{x} | \text{admin } a \text{ appears first in } \sigma) \cdot \mathbb{P}(\text{admin } a \text{ appears first in } \sigma) \quad (30)$$

Since σ is a random ordering of the admins, the probability that admin a is the first one in σ is simply $1/k$, whereas the probability that admin a wins rank 1 for subject slot 1 conditioned on the fact that it the first one to play is just $\text{PA}_a(\mathbf{x}_a, t) = 1 - (1-p)^{\mathbf{x}_a(t)}$ given the allocation rule $\mathcal{A}_{\text{prob}}$. Thus, we obtain

$$\mathbb{E}[n_1] \geq \sum_{t=1}^n \frac{\text{PA}_a(\mathbf{x}_a, t)}{k} \quad (31)$$

Under a Nash equilibrium \mathbf{x} , we know from Lemma 4.1 that admin a has exhausted their budget $B^{(a)}$ and has bid only 1s and 0s. Thus,

$$\mathbb{E}[n_1] \geq \sum_{\substack{t \in [n], \\ \mathbf{x}_a(t)=1}} \frac{\text{PA}_a(\mathbf{x}_a, t)}{k} + \sum_{\substack{t \in [n], \\ \mathbf{x}_a(t)=0}} \frac{\text{PA}_a(\mathbf{x}_a, t)}{k} \quad (32)$$

Since $\text{PA}_a(\mathbf{x}_a, t) = 1 - (1-p)^{\mathbf{x}_a(t)}$, we get that $\text{PA}_a(\mathbf{x}_a, t) = p$ for $\mathbf{x}_a(t) = 1$ and $\text{PA}_a(\mathbf{x}_a, t) = 0$ for $\mathbf{x}_a(t) = 0$, thus giving

$$\begin{aligned} \mathbb{E}[n_1] &\geq \sum_{\substack{t \in [n], \\ \mathbf{x}_a(t)=1}} \frac{\text{PA}_a(\mathbf{x}_a, t)}{k} \Rightarrow \\ &\mathbb{E}[n_1] \geq \frac{B^{(a)} \cdot p}{k} \end{aligned} \quad (33)$$

Since p is defined as a constant with respect to all other model parameters, we get that $\mathbb{E}[n_1] = \Omega\left(\frac{B^{(a)}}{k}\right)$. We conclude that $\mathbb{E}\left[\sum_{r=1}^k n_r \cdot \alpha_r^2\right] = \Omega(B^{(a)}/k)$, which follows from the simple property that if a function f is $\Omega(n)$, then for function g for which $g(x) \geq f(x), \forall x$ in the domain of f and g, g also grows as $\Omega(n)$. We note that there may be other allocations rule that satisfy this property, and thus our analysis is not contingent on the specific choice of the allocation rule, as long as it preserves this property. \square

Proof for Corollary A.3. Corollary A.3 follows easily from the above proof with a minimal modification of equations (32)–(33), knowing that an admin can bid at most C budget units on every subject slot and they have exhausted their budget:

$$\mathbb{E}[n_1] \geq \sum_{\substack{t \in [n], \\ \mathbf{x}_a(t) \geq 1}} \frac{\text{PA}_a(\mathbf{x}_a, t)}{k} + \sum_{\substack{t \in [n], \\ \mathbf{x}_a(t)=0}} \frac{\text{PA}_a(\mathbf{x}_a, t)}{k} \quad (34)$$

Since $\text{PA}_a(\mathbf{x}_a, t) = 1 - (1-p)^{\mathbf{x}_a(t)}$, we get that $\text{PA}_a(\mathbf{x}_a, t) \geq p$ for $\mathbf{x}_a(t) \geq 1$ and $\text{PA}_a(\mathbf{x}_a, t) = 0$ for $\mathbf{x}_a(t) = 0$, thus giving

$$\begin{aligned}\mathbb{E}[n_1] &\geq \sum_{\substack{t \in [n], \\ \mathbf{x}_a(t) \geq 1}} \frac{\text{PA}_a(\mathbf{x}_a, t)}{k} \Rightarrow \\ \mathbb{E}[n_1] &\geq \frac{B^{(a)} \cdot p}{C \cdot k}\end{aligned}\tag{35}$$

under the assumption that $B^{(a)} \leq n$. Just as before, we get that $\mathbb{E}[n_1] = \Omega\left(\frac{B^{(a)}}{C \cdot k}\right)$ and conclude that $\mathbb{E}\left[\sum_{r=1}^k n_r \cdot \alpha_r^2\right] = \Omega(B^{(a)}/(C \cdot k))$. □

Corollary A.3 simply states that the expectation of the variable S grows at at least the rate of $B^{(a)}/k$ (since C is a constant) when an admin uses their entire budget, whether in a Nash equilibrium or not (intuitively, subject to using all budget, any permutation of how the budget is allocated across subject slots, the expectation of S does not get ‘too small’). Take a constant $c > 0$ such that

$$\mathbb{E}[S] \geq c \cdot \frac{B^{(a)}}{C \cdot k}.\tag{36}$$

Since $B^{(a)}/C \leq m \leq B^{(a)}$, we also know that there exists a constant $c' > 0$ such that $\mathbb{E}[S] \geq c' \cdot \frac{m}{k}$. We take η in equation (24) to be equal to $\eta = c_0 \cdot \sqrt{m}$, with c_0 to be determined. Then, equation (24) becomes

$$\begin{aligned}\mathbb{P}(|S - \mathbb{E}[S]| \leq \eta) &\geq 1 - 2 \exp\left(-\frac{2\eta^2}{m}\right) \Leftrightarrow \\ \mathbb{P}(|S - \mathbb{E}[S]| \leq \sqrt{m} \cdot c_0) &\geq 1 - 2 \exp(-2c_0^2)\end{aligned}\tag{37}$$

Taking $c_0 = \epsilon \cdot c' \cdot \frac{\sqrt{m}}{k}$, we get from equation (36):

$$\begin{aligned}\mathbb{P}\left(|S - \mathbb{E}[S]| \leq \epsilon \cdot c' \cdot \frac{m}{k}\right) &\geq 1 - 2 \exp(-2\epsilon^2 c'^2 m/k^2) \Rightarrow \\ \mathbb{P}(|S - \mathbb{E}[S]| \leq \epsilon \cdot \mathbb{E}[S]) &\geq 1 - 2 \exp(-2\epsilon^2 c'^2 m/k^2)\end{aligned}\tag{38}$$

The right hand side tends to 1 when the exponential vanishes, i.e. in the regime where $B^{(a)} = \omega(k^2)$. Various modeling assumptions may satisfy this requirement, for example when the budget and the number of administrators grow sublinearly in n , with $B^{(a)} = \Omega(n^b)$ and $k = \mathcal{O}(n^{b/2-\delta})$, for $b \in (0, 1)$ and $\delta \in (0, b/2)$.

A simple corollary follows from Proposition A.1:

Corollary A.4. For allocations $\mathbf{x} \in \mathcal{D}$ in which all admins use all their budget, $\sum_{t=1}^n \mathbf{x}_a(t) = B^{(a)}, \forall a \in [k]$, with high probability $\frac{1}{S} \in \left[(1 - \epsilon) \cdot \frac{1}{\mathbb{E}[S]}, (1 + \epsilon) \cdot \frac{1}{\mathbb{E}[S]}\right]$, with $\epsilon > 0$.

Proof. The proof follows easily by constructing ϵ and using Proposition A.1. Since we know that with high probability, $S \geq (1 - \epsilon) \cdot \mathbb{E}[S]$ for an allocation \mathbf{x} and S is positive,

$$\begin{aligned}S \geq (1 - \epsilon) \cdot \mathbb{E}[S] &\Rightarrow \frac{1}{S} \leq \frac{1}{(1 - \epsilon) \cdot \mathbb{E}[S]}, \\ S \leq (1 + \epsilon) \cdot \mathbb{E}[S] &\Rightarrow \frac{1}{S} \geq \frac{1}{(1 + \epsilon) \cdot \mathbb{E}[S]}\end{aligned}\tag{39}$$

Set $\frac{1}{1-\epsilon} = 1 + \epsilon_0 \Rightarrow \epsilon = \frac{\epsilon_0}{1+\epsilon_0} \Rightarrow \frac{1}{1+\epsilon} = 1 - \frac{\epsilon_0}{1+2\epsilon_0}$. Choosing $\epsilon' = \max\left(\epsilon_0, \frac{\epsilon_0}{1+2\epsilon_0}\right)$, we get that

$$\mathbb{P}\left(\left|\frac{1}{S} - \frac{1}{\mathbb{E}[S]}\right| \geq \epsilon' \cdot \frac{1}{\mathbb{E}[S]}\right) = \mathbb{P}(|S - \mathbb{E}[S]| \geq \epsilon \cdot \mathbb{E}[S]) \geq 2 \exp(-2\epsilon^2 c'^2 m/k^2), \quad (40)$$

concluding our proof. For $m/k^2 \rightarrow \infty$, the right hand side vanishes. \square

We return to the proof of Theorem 2.2. We denote by f_a^{MSE} the utility function of the game \mathcal{G}_{MSE} and by f_a^{SV} the utility function of the game \mathcal{G}_{SV} , for each admin a .

Note that f_a^{SV} is a non-negative function on its entire domain for all admins a , whereas f_a^{MSE} is a non-positive function on its entire domain for all admins a . Note that $f_a^{\text{MSE}}(\mathbf{x}) = -\mathbb{E}\left[\frac{1}{\sum_{r=1}^k n_r \cdot \alpha_r^2}\right] = -\mathbb{E}\left[\frac{1}{S}\right]$ and $f_a^{\text{SV}}(\mathbf{x}) = \mathbb{E}\left[\sum_{r=1}^k n_r \cdot \alpha_r^2\right] = \mathbb{E}[S]$, by definition of the random variable S from equation (19).

According to Proposition A.1, we know that S concentrates around its expectation. For $\epsilon > 0$, denote the event that $S \in [(1 - \epsilon) \cdot \mathbb{E}[S], (1 + \epsilon) \cdot \mathbb{E}[S]]$ by \mathcal{E} for ease of notation. Thus,

$$\mathbb{E}\left[\frac{1}{S}\right] = \mathbb{E}\left[\frac{1}{S}|\mathcal{E}\right] \cdot \mathbb{P}[\mathcal{E}] + \mathbb{E}\left[\frac{1}{S}|\mathcal{E}^c\right] \cdot \mathbb{P}[\mathcal{E}^c], \quad (41)$$

where A^c denotes the complement of the event A . We know that, conditioned on \mathcal{E} , $\frac{1}{S}$ is concentrated around $\frac{1}{\mathbb{E}[S]}$, whereas conditioned on \mathcal{E}^c , expectation of $\frac{1}{S}$ is bounded by 1 given that the error estimation objective outputs $\min\left(\frac{1}{S}, 1\right)$. Denote by η the probability of \mathcal{E} (from Proposition A.1, $\eta = 2e^{-2\epsilon^2 c'^2 m/k^2}$). Thus, we obtain

$$\begin{aligned} \mathbb{E}\left[\frac{1}{S}\right] &\leq \mathbb{E}\left[\frac{1}{S}|\mathcal{E}\right] + 1 \cdot \eta \leq (1 + \epsilon) \cdot \frac{1}{\mathbb{E}[S]} + \eta \Rightarrow \\ f_a^{\text{MSE}}(\mathbf{x}) &\geq (1 + \epsilon) \left(-\frac{1}{f_a^{\text{SV}}(\mathbf{x})}\right) - \eta \end{aligned} \quad (42)$$

which holds for all allocations \mathbf{x} such that all admins have allocated all their budget, under the assumption that an admin can bid at most C budget units on any subject slot.

Furthermore, applying Jensen's inequality on the function $f(x) = 1/x$, we get that for all allocations \mathbf{x}_a

$$\begin{aligned} \mathbb{E}\left[\frac{1}{\sum_{r=1}^k n_r \cdot \alpha_r^2}\right] &\geq \frac{1}{\mathbb{E}\left[\sum_{r=1}^k n_r \cdot \alpha_r^2\right]} \Leftrightarrow \\ -\mathbb{E}\left[\frac{1}{\sum_{r=1}^k n_r \cdot \alpha_r^2}\right] &\leq -\frac{1}{\mathbb{E}\left[\sum_{r=1}^k n_r \cdot \alpha_r^2\right]} \Leftrightarrow \\ f_a^{\text{MSE}}(\mathbf{x}_a, \mathbf{x}_{-a}) &\leq -\frac{1}{f_a^{\text{SV}}(\mathbf{x}_a, \mathbf{x}_{-a})} \end{aligned} \quad (43)$$

Now, take \mathbf{x} to be a Nash equilibrium for the game, and we will show that \mathbf{x} has to be an (ϵ, η) -approximate Nash equilibrium for the game \mathcal{G}_{MSE} . Thus, we would like to show that $\forall \mathbf{x}'_a \in \mathcal{D}_a$,

$$\begin{aligned} f_a^{\text{SV}}(\mathbf{x}_a, \mathbf{x}_{-a}) &\geq f_a^{\text{SV}}(\mathbf{x}'_a, \mathbf{x}_{-a}) \Rightarrow \\ f_a^{\text{MSE}}(\mathbf{x}_a, \mathbf{x}_{-a}) &\geq (1 + \epsilon) \cdot f_a^{\text{MSE}}(\mathbf{x}'_a, \mathbf{x}_{-a}) - \eta \end{aligned} \quad (44)$$

with $\epsilon, \eta > 0, \forall a \in [k]$.

We will first show that the above statement holds for all $\mathbf{x}'_a \in \mathcal{D}_a$ such that $\sum_{t=1}^n \mathbf{x}'_a(t) = B^{(a)}$ (in other words, admin a has exhausted their budget with strategy \mathbf{x}'_a). Given that \mathbf{x} is a Nash equilibrium for \mathcal{G}_{SV} , we get

$$\begin{aligned} f_a^{SV}(\mathbf{x}_a, \mathbf{x}_{-a}) &\geq f_a^{SV}(\mathbf{x}'_a, \mathbf{x}_{-a}) \Rightarrow \\ -\frac{1}{f_a^{SV}(\mathbf{x}_a, \mathbf{x}_{-a})} &\geq -\frac{1}{f_a^{SV}(\mathbf{x}'_a, \mathbf{x}_{-a})} \Rightarrow \\ f_a^{\text{MSE}}(\mathbf{x}_a, \mathbf{x}_{-a}) &\geq (1 + \epsilon) \left(-\frac{1}{f_a^{SV}(\mathbf{x}_a, \mathbf{x}_{-a})} \right) - \delta \geq (1 + \epsilon) \left(-\frac{1}{f_a^{SV}(\mathbf{x}'_a, \mathbf{x}_{-a})} \right) - \delta \geq (1 + \epsilon) \cdot f_a^{\text{MSE}}(\mathbf{x}'_a, \mathbf{x}_{-a}) - \eta \end{aligned} \quad (45)$$

where the last inequality follows from Jensen's inequality, as described above. Since $\epsilon, \eta > 0$, we get the (ϵ, η) -approximation condition. Thus, we are left to show inequality (44) for allocations \mathbf{x}'_a that may not exhaust the entire budget $B^{(a)}$. For each such allocation, create another allocation $\mathbf{x}_a^{\text{full}}$ by using the rest of the budget randomly across subject slots that had no bid on \mathbf{x}'_a . Thus, $\mathbf{x}_a^{\text{full}}$ satisfies the conditions necessary for Proposition A.1. We note that the utility function f_a^{SV} can only increase from bidding a unit of budget or more on additional subjects which had a previous bid of 0 budget, since f_a^{SV} is a summation of the utility coming from all subjects $t \in [n]$ with all terms being non-negative, and $\text{PA}_a(\mathbf{x}_a, t) = 0$ if $\mathbf{x}_a(t) = 0$ and $\text{PA}_a(\mathbf{x}_a, t) \geq p > 0$ if $\mathbf{x}_a(t) \geq 1$. Thus, we get that

$$f_a^{SV}(\mathbf{x}_a^{\text{full}}, \mathbf{x}_{-a}) \geq f_a^{SV}(\mathbf{x}'_a, \mathbf{x}_{-a}) \Leftrightarrow -\frac{1}{f_a^{SV}(\mathbf{x}_a^{\text{full}}, \mathbf{x}_{-a})} \geq -\frac{1}{f_a^{SV}(\mathbf{x}'_a, \mathbf{x}_{-a})} \quad (46)$$

From equations (45)–(46) and using the Jensen inequality again, we obtain that

$$f_a^{\text{MSE}}(\mathbf{x}_a^{\text{full}}, \mathbf{x}_{-a}) \geq (1 + \epsilon) \left(-\frac{1}{f_a^{SV}(\mathbf{x}_a^{\text{full}}, \mathbf{x}_{-a})} \right) - \eta \geq (1 + \epsilon) \left(-\frac{1}{f_a^{SV}(\mathbf{x}'_a, \mathbf{x}_{-a})} \right) - \eta \geq (1 + \epsilon) \cdot f_a^{\text{MSE}}(\mathbf{x}'_a, \mathbf{x}_{-a}) - \eta, \quad (47)$$

thus proving that a Nash equilibrium for the game \mathcal{G}_{SV} is an (ϵ, η) -approximate Nash equilibrium for the game \mathcal{G}_{MSE} .

We note that ϵ and η are

$$\epsilon = \mathcal{O} \left(\frac{k}{\sqrt{B}} \right) \quad \text{and} \quad \eta = \mathcal{O} \left(\exp(-B/k^2) \right), \quad (48)$$

where $B = \min_a B^{(a)}$ and the maximum bid per subject slot is bounded. These come from equation (38), in which the concentration bounds links the concentration ‘closeness’ ϵ with the probability of concentration η . □

A.3. Proofs for Section 3 Results

We define a population distribution \mathcal{X} on the subjects, which are fully described by the random variables (rank, T, Y) . Individuals are drawn independently for each slot, according to a fixed rank allocation \bar{r} . In particular, we sample the n_r individuals from their slots at rank r from the do-interventional distribution $\mathcal{X}^{\text{do}(\text{rank}=r)}$.

We highlight that the outcomes Y_i for $i \in [n^{(a)}]$ are *not* identically distributed. Their distribution differs according to their associated ranks r_i . The ranks r_i are assigned by the rank allocation \bar{r} , and therefore, deterministic in i . They merely indicate from which do-interventional distribution the individual was drawn. We denote the joint distribution of the outcomes and treatment assignment variables $Z_i = (Y_i, T_i)$, $i \in [n^{(a)}]$ by P .

We explicitly define the set of considered model instances \mathcal{M} . For a given sequence $(\alpha_r)_{r=1,\dots,k}$, we set

$$\mathcal{M} = \left\{ \left(c_{i,0}, c_{i,1}^{(r_i)} \right)_{i \in [n^{(a)}]} : |\tau| \leq 1, \tau_r = \alpha_r \cdot \tau \right\}, \quad (49)$$

with the quantities $\tau = \frac{1}{n^{(a)}} \sum_{i=1}^{n^{(a)}} \frac{c_{i,1}^{(r_i)}}{\alpha_{r_i}}$ and

$$\tau_r = \begin{cases} \frac{1}{n_r} \sum_{i:r_i=r} c_{i,1}^{(r)} & \text{for } n_r \geq 1 \\ \alpha_r \cdot \tau & \text{for } n_r = 0 \end{cases}$$

for $r \in [k]$.⁶

Moreover, note that we also identify the causal effect of interest in equation (7) through the potential outcomes model as

$$\tau = \frac{1}{n^{(a)}} \sum_{i=1}^{n^{(a)}} \frac{c_{i,1}^{(r_i)}}{\alpha_{r_i}}.$$

Proof of Theorem 3.4. We follow the proof of Theorem 2 in Cortez-Rodriguez et al. (2023) and apply Le Cam's method (Le Cam, 1973) to obtain the minimax bound. For the computation, we make use of the independence between individuals, and the fact that we can sort them according to their assigned ranks.

We use Le Cam's method (Le Cam, 1973) to obtain the proof of Theorem 3.4.

Lemma A.5 (Le Cam's method). *Let \mathcal{M} be a set of model instances. Then, we have for any two instances $M_1, M_2 \in \mathcal{M}$ such that $|\tau(M_1) - \tau(M_2)| \geq 2\delta$ with $\delta > 0$:*

$$\inf_{\hat{\tau}} \sup_{M \in \mathcal{M}} \mathbb{E} [(\hat{\tau} - \tau)^2] \geq \frac{\delta^2}{2} \left(1 - \sqrt{\frac{D_{KL}(P_1 || P_2)}{2}} \right), \quad (50)$$

where P_1 and P_2 denote the population distribution with respect to M_1 and M_2 , respectively, and $D_{KL}(\cdot || \cdot)$ denotes the Kullback-Leibler divergence.

We begin the proof by constructing instances M_1 and M_2 . We set $c_{i,0} = 0$ and $c_{i,1}^{r_i} = \delta \cdot \alpha_{r_i}$ for M_1 and $c_{i,0} = 0$ and $c_{i,1}^{r_i} = -\delta \cdot \alpha_{r_i}$ for M_2 . Note that $\alpha_1 = 1$ by definition. Hence, we get

$$\tau(M_1) = \frac{1}{n^{(a)}} \sum_{i=1}^{n^{(a)}} \frac{c_{i,1}^{(r_i)}}{\alpha_{r_i}} = \frac{1}{n^{(a)}} \sum_{i=1}^{n^{(a)}} \frac{\delta \cdot \alpha_{r_i}}{\alpha_{r_i}} = \delta \quad (51)$$

$$\tau(M_2) = \dots = -\delta \quad (52)$$

So, we have $|\tau(M_1) - \tau(M_2)| \geq 2\delta$. Thus, the instances M_1 and M_2 satisfy the coverage condition in Lemma A.5.

⁶We note that when there are no datapoints at rank r , we cannot identify the effect directly, but we can identify it from data at other ranks since the structural parameters $(\alpha_r)_r$ are known.

Next, we compute $D_{KL}(P_1||P_2)$ with P_1 and P_2 are the joint distribution of $(Y_i, T_i), i \in [n^{(a)}]$ for the instances M_1 and M_2 . We will choose a δ to upper bound $D_{KL}(P_1||P_2)$ by $1/2$. As all samples are independently drawn, the joint distributions P_1 and P_2 factorize as:

$$P_1(\{(Y_i, T_i)\}_{i=1}^{n^{(a)}}) = \prod_{i=1}^{n^{(a)}} P_1(Y_i, T_i)$$

$$P_2(\{(Y_i, T_i)\}_{i=1}^{n^{(a)}}) = \prod_{i=1}^{n^{(a)}} P_2(Y_i, T_i).$$

Note that the outcomes Y_i and Y_j are differently distributed when the subjects i and j are sampled from different ranks, $r_i \neq r_j$. Hence, we have in general that $P_1(Y_i, T_i) \neq P_1(Y_j, T_j)$ and $P_2(Y_i, T_i) \neq P_2(Y_j, T_j)$ for $i \neq j$. We obtain that

$$D_{KL}(P_1||P_2) = \mathbb{E}_{M_1} \left[\log \left(\frac{p_1(\{(Y_i, T_i)\}_{i=1}^{n^{(a)}})}{p_2(\{(Y_i, T_i)\}_{i=1}^{n^{(a)}})} \right) \right] = \sum_{i=1}^{n^{(a)}} \mathbb{E}_{M_1} \left[\log \left(\frac{p_1(Y_i, T_i)}{p_2(Y_i, T_i)} \right) \right]$$

with p_1 and p_2 the densities of P_1 and P_2 , respectively. We sort the samples according to their allocated ranks and apply the Law of Total Expectation,

$$\begin{aligned} D_{KL}(P_1||P_2) &= \sum_{r=1}^k \sum_{j:r_j=r} \mathbb{E}_{M_1} \left[\log \left(\frac{p_1(Y_j, T_j)}{p_2(Y_j, T_j)} \right) \right] \\ &= \sum_{r=1}^k \sum_{j:r_j=r} P_1(T_j = 1) \cdot \mathbb{E}_{M_1} \left[\log \left(\frac{p_1(Y_j, T_j)}{p_2(Y_j, T_j)} \right) | T_j = 1 \right] \\ &\quad + P_1(T_j = 0) \cdot \mathbb{E}_{M_1} \left[\log \left(\frac{p_1(Y_j, T_j)}{p_2(Y_j, T_j)} \right) | T_j = 0 \right] \end{aligned} \quad (53)$$

Focusing on samples drawn from a fixed rank r , the outcomes Y_j conditional on the treatment T_j , are distributed as:

$$Y_j|T_j \sim N \left(c_{i,1}^{(r)} \cdot T_j, \sigma^2 \right) \quad (54)$$

as $Y_j = c_{j,0} + c_{j,1}^{(r)} \cdot T_j + \varepsilon_j$ with $c_{j,0} = 0$ and $\varepsilon_j \sim N(0, \sigma^2)$. Hence, we have for M_1

$$Y_j|T_j \sim N(\delta \cdot \alpha_r \cdot T_j, \sigma^2)$$

and for M_2

$$Y_j|T_j \sim N(-\delta \cdot \alpha_r \cdot T_j, \sigma^2).$$

We have $Y_j|(T_j = 0) \sim N(0, \sigma^2)$ for both instances M_1 and M_2 . Therefore, the second term in equation (53) simplifies to zero. We remain with

$$\begin{aligned} D_{KL}(P_1||P_2) &= \sum_{r=1}^k \sum_{j:r_j=r} P_1(T_j = 1) \cdot \mathbb{E}_{M_1} \left[\log \left(\frac{p_1(Y_j, T_j)}{p_2(Y_j, T_j)} \right) | T_j = 1 \right] \\ &= \sum_{r=1}^k \sum_{j:r_j=r} P_1(T_j = 1) \cdot \mathbb{E}_{M_1} \left[\log \left(\frac{\varphi_{\delta \cdot \alpha_r, \sigma^2}(Y_j, T_j)}{\varphi_{-\delta \cdot \alpha_r, \sigma^2}(Y_j, T_j)} \right) | T_j = 1 \right] \end{aligned}$$

where φ_{μ, σ^2} denotes the density function of $N(\mu, \sigma^2)$. Note that the difference in distribution across ranks r is now made explicit in the densities, e.g., $\varphi_{\delta \cdot \alpha_r, \sigma^2}$ for model instance M_1 . We input the explicit formula for the normal density and use




Figure 3. Example of minimax lower bound with $q = 0.5$ and $\sigma^2 = 4$.

that $T_j \sim \text{Ber}(q_j)$, for $q_j \in [q, 1 - q]$,

$$D_{KL}(P_1||P_2) = \sum_{r=1}^k \sum_{j:r_j=r} q_j \cdot \mathbb{E}_{M_1} \left[-\frac{1}{2\sigma^2} \left((Y_i - \delta \cdot \alpha_r)^2 - (Y_i + \delta \cdot \alpha_r)^2 \right) | T_j = 1 \right] \quad (55)$$

$$= \sum_{r=1}^k \sum_{j:r_j=r} q_j \cdot \frac{2}{\sigma^2} \cdot \delta \cdot \alpha_r \cdot \mathbb{E}_{M_1} [Y_j | T_j = 1] \quad (56)$$

$$= \sum_{r=1}^k \sum_{j:r_j=r} q_j \cdot \frac{2}{\sigma^2} \cdot \delta \cdot \alpha_r \cdot (0 + \delta \cdot \alpha_r + \mathbb{E}[\varepsilon_j]) \quad (57)$$

$$= \sum_{r=1}^k \sum_{j:r_j=r} q_j \cdot \frac{2}{\sigma^2} \cdot \delta^2 \cdot \alpha_r^2 \quad (58)$$

$$\leq \delta^2 \cdot (1 - q) \cdot \frac{2}{\sigma^2} \cdot \sum_{r=1}^k n_r \cdot \alpha_r^2. \quad (59)$$

We can bound $D_{KL}(P_1||P_2)$ by $1/2$ if we set

$$\delta^2 = \frac{\sigma^2}{4 \cdot (1 - q) \cdot \sum_{r=1}^k n_r \cdot \alpha_r^2}.$$

Then, Le Cam's method gives,

$$\inf_{\hat{\tau}} \sup_{M \in \mathcal{M}} \mathbb{E} [(\hat{\tau} - \tau)^2] \geq \frac{\delta^2}{2} \left(1 - \sqrt{\frac{D_{KL}(P_1||P_2)}{2}} \right) \quad (60)$$

$$= \frac{\sigma^2}{16 \cdot (1 - q) \cdot \sum_{r=1}^k n_r \cdot \alpha_r^2} \quad (61)$$

Finally, note that the sample value $\sum_r n_r \cdot \alpha_r^2$ is lower bounded by one if the admin attains at least one slot at rank one, i.e., $n_1 = 1$. However, it might happen that no slot is won at rank one (in fact, it may happen that $n_r = 0$ for some other rank r too). In this case, the sample value is sometimes lower than one, and as it decreases, the inverse of $\left(\sum_{r=1}^k n_r \cdot \alpha_r^2 \right)^{-1}$ may explode. However, the minimum error is still upper bounded by 1 since the admin is better off using a constant estimator, as the effect is bounded, $|\tau| \leq 1$. Therefore, the minimax lower bound is the minimum of the error achieved using the samples and the error obtained by a constant estimator. See Figure 3 for an illustration.

□

Proof of Proposition 3.5. We note that $\tau_r = \frac{1}{n_r} \sum_{i:r_i=r} c_{i,1}^{(r)}$ from the potential outcomes model.

Then, we compute the expectation of our estimator as:

$$\begin{aligned} \mathbb{E}(\widehat{\tau}_r) &= \mathbb{E} \left[\frac{1}{n_r} \sum_{i:r_i=r} \left(\frac{Y_i \mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{Y_i \mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right) \right] \\ &= \frac{1}{n_r} \sum_{i:r_i=r} \mathbb{E} \left[\frac{Y_i \mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{Y_i \mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right] \end{aligned} \quad (62)$$

For every i , note that, by linearity of expectation and the potential outcomes model for subjects i at rank $r_i = r$,

$$\begin{aligned} &\mathbb{E} \left[\frac{Y_i \mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{Y_i \mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right] \\ &= c_{i,0} \cdot \mathbb{E} \left[\frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{\mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right] \\ &+ c_{i,1}^{(r)} \cdot \mathbb{E} \left[T_i \cdot \left(\frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{\mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right) \right] \\ &+ \mathbb{E}[\varepsilon_i] \cdot \mathbb{E} \left[\frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{\mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right] \\ &= c_{i,0} \cdot \mathbb{E} \left[\frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{\mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right] \\ &+ c_{i,1}^{(r)} \cdot \mathbb{E} \left[T_i \cdot \left(\frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{\mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right) \right] \end{aligned} \quad (63)$$

Assuming that the actual treatment T_i is administered through a Bernoulli randomized design, $T_i \sim \text{Ber}(q_i)$, for $q_i \in [q, 1-q]$, we get that the first term of equation (63) reduces to 0 and the second one to $c_{i,1}^{(r)}$. To see this briefly, note that the first term is equal to

$$\begin{aligned} &c_{i,0} \cdot \mathbb{E} \left[\frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{\mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right] \\ &= c_{i,0} \cdot \mathbb{E} \left[\frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{\mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right] \\ &= c_{i,0} \cdot \left(\frac{\mathbb{P}(T_i = 1)}{\mathbb{P}(T_i = 1)} - \frac{\mathbb{P}(T_i = 0)}{\mathbb{P}(T_i = 0)} \right) = 0 \end{aligned} \quad (64)$$

using the law of total expectation. Similarly, the second term is equal to

$$\begin{aligned} &c_{i,1}^{(r)} \cdot \mathbb{E} \left[T_i \cdot \left(\frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{\mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right) \right] \\ &= c_{i,1}^{(r)} \cdot \mathbb{E} \left[T_i \cdot \left(\frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} - \frac{\mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right) \right] \\ &= c_{i,1}^{(r)} \cdot \left(\mathbb{E} \left[T_i \cdot \frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} \right] - \mathbb{E} \left[T_i \cdot \frac{\mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right] \right) \\ &= c_{i,1}^{(r)} \cdot \left(\mathbb{E} \left[\mathbb{1}\{T_i = 1\} \cdot \frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} \right] - \mathbb{E} \left[\mathbb{1}\{T_i = 1\} \cdot \frac{\mathbb{1}\{T_i = 0\}}{\mathbb{P}(T_i = 0)} \right] \right) \\ &= c_{i,1}^{(r)} \cdot \left(\mathbb{E} \left[\frac{\mathbb{1}\{T_i = 1\}}{\mathbb{P}(T_i = 1)} \right] - 0 \right) \\ &= c_{i,1}^{(r)} \cdot \frac{\mathbb{P}(T_i = 1)}{\mathbb{P}(T_i = 1)} = c_{i,1}^{(r)} \end{aligned} \quad (65)$$

as $T_i = \mathbb{1}\{T_i = 1\}$ by definition, and we know that $\mathbb{1}\{T_i = 1\} \cdot \mathbb{1}\{T_i = 1\} = \mathbb{1}\{T_i = 1\}$ and $\mathbb{1}\{T_i = 1\} \cdot \mathbb{1}\{T_i = 0\} = 0$.

Coming back to the sum over all subjects i at rank r , we get that

$$\mathbb{E}(\widehat{\tau}_r) = \frac{1}{n_r} \cdot \sum_{i:r_i=r} c_{i,1}^{(r)} = \tau_r \quad (66)$$

□

Proof for Proposition 3.6. Denote by

$$\widehat{\tau}_{r,d} = \frac{1}{n_r} \cdot \sum_{i:r_i=r} \frac{Y_i \mathbb{1}\{T_i = d\}}{\mathbb{P}(T_i = d)}, \quad (67)$$

for $d \in \{0, 1\}$. Then, we have $\widehat{\tau}_r = \widehat{\tau}_{r,1} - \widehat{\tau}_{r,0}$. We begin by computing the variance of $\widehat{\tau}_{r,d}$ for $d \in \{0, 1\}$.

Step 1:

$$\text{Var}(\widehat{\tau}_{r,d}) = \frac{1}{n_r^2} \sum_{i:r_i=r} \sum_{j:r_j=r} \frac{\text{Cov}(Y_i \mathbb{1}\{T_i = d\}, Y_j \mathbb{1}\{T_j = d\})}{\mathbb{P}(T_i = d)\mathbb{P}(T_j = d)} \quad (68)$$

Let's first take the case $d = 1$. Case 1: For $i = j$,

$$\text{Cov}(Y_i \mathbb{1}\{T_i = 1\}, Y_j \mathbb{1}\{T_j = 1\}) = \text{Var}(Y_i \mathbb{1}\{T_i = 1\}) \quad (69)$$

Note that

$$\begin{aligned} \text{Var}(Y_i \mathbb{1}\{T_i = 1\}) &= \text{Var}(\varepsilon_i) \cdot \text{Var}(\mathbb{1}\{T_i = 1\}) + \text{Var}(\varepsilon_i) \cdot \mathbb{E}(\mathbb{1}\{T_i = 1\})^2 \\ &\quad + \mathbb{E}(Y_i | T_i)^2 \cdot \text{Var}(\mathbb{1}\{T_i = 1\}) \\ &\leq \sigma^2 \cdot q_i(1 - q_i) + \sigma^2 \cdot q_i^2 + \left(c_{i,0} + c_{i,1}^{(r)} + \mathbb{E}(\varepsilon_i)\right)^2 \cdot q_i(1 - q_i) \\ &= \sigma^2 \cdot q_i + Y_{\max}^2 \cdot q_i(1 - q_i) \end{aligned} \quad (70)$$

where the last inequality holds under the bounded effect assumption from equation 3.1.

Case 2: For $i \neq j$, under the Bernoulli design, treatment assignments are independent and so no treatment affects both units i and j . Hence, $\text{Cov}(Y_i \mathbb{1}\{T_i = 1\}, Y_j \mathbb{1}\{T_j = 1\}) = 0$.

Thus,

$$\begin{aligned} \text{Var}(\widehat{\tau}_{r,1}) &\leq \frac{1}{n_r^2} \sum_{i:r_i=r} \frac{\sigma^2 \cdot q_i + Y_{\max}^2 \cdot q_i(1 - q_i)}{t^2} \\ &= \frac{1}{n_r^2} \sum_{i:r_i=r} \frac{\sigma^2 + Y_{\max}^2 \cdot (1 - q_i)}{q_i} \\ &\leq \frac{1}{n_r \cdot q} (\sigma^2 + Y_{\max}^2 \cdot (1 - q)), \end{aligned} \quad (71)$$

given that the second equation is a decreasing function in q_i and $q_i \in [q, 1 - q]$. In a similar way we obtain a bound for $d = 0$:

$$\text{Var}(\widehat{\tau}_{r,0}) \leq \frac{1}{n_r \cdot (1 - q)} (\sigma^2 + Y_{\max}^2 \cdot q) \quad (72)$$

Step 2: Next, we bound the covariance between $\widehat{\tau}_{r,0}$ and $\widehat{\tau}_{r,1}$:

$$\text{Cov}(\widehat{\tau}_{r,1}, \widehat{\tau}_{r,0}) = \mathbb{E}[\widehat{\tau}_{r,1} \cdot \widehat{\tau}_{r,0}] - \mathbb{E}[\widehat{\tau}_{r,1}] \cdot \mathbb{E}[\widehat{\tau}_{r,0}] \quad (73)$$

A simple calculation shows that

$$\mathbb{E}[\widehat{\tau}_{r,t}] = \frac{1}{n_r} \sum_{i:r_i=t} Y_i(T_i = t) \quad (74)$$

and

$$\mathbb{E}[\widehat{\tau}_{r,1} \cdot \widehat{\tau}_{r,0}] = \mathbb{E} \left[\frac{1}{n_r^2} \left(\sum_{i:r_i=1} \frac{Y_i \mathbb{1}(T_i = 1)}{\mathbb{P}(T_i = 1)} \right) \cdot \left(\sum_{i:r_i=0} \frac{Y_i \mathbb{1}(T_i = 0)}{\mathbb{P}(T_i = 0)} \right) \right] \quad (75)$$

Using equations (74) and (75) in the covariance expression and simplifying, we get

$$\begin{aligned} \text{Cov}(\widehat{\tau}_{r,1}, \widehat{\tau}_{r,0}) &= \frac{1}{n_r^2} \left(\sum_{i \neq j} \frac{\mathbb{P}(T_i = 1 \& T_j = 0)}{\mathbb{P}(T_i = 1)\mathbb{P}(T_j = 0)} \cdot Y_i(T_i = 1)Y_j(T_j = 0) - \sum_{i:r_i=1} Y_i(T_i = 1) \sum_{i:r_i=0} Y_i(T_i = 0) \right) \\ &= \frac{1}{n_r^2} \left(\sum_{i \neq j} Y_i(T_i = 1)Y_j(T_j = 0) - \sum_{i,j} Y_i(T_i = 1)Y_j(T_j = 0) \right) \\ &= -\frac{1}{n_r^2} \sum_{i:r_i=0} Y_i(T_i = 1)Y_i(T_i = 0) \end{aligned} \quad (76)$$

Given the bounded effect assumption from equation 3.1, we can bound

$$-\text{Cov}(\widehat{\tau}_{r,1}, \widehat{\tau}_{r,0}) \leq \frac{Y_{\max}^2}{n_r} \quad (77)$$

Step 3: Putting equations (71), (72), and (77) together, we get

$$\begin{aligned} \text{Var}(\widehat{\tau}_r) &= \text{Var}(\widehat{\tau}_{r,1}) + \text{Var}(\widehat{\tau}_{r,0}) - 2 \cdot \text{Cov}(\widehat{\tau}_{r,1}, \widehat{\tau}_{r,0}) \Rightarrow \\ \text{Var}(\widehat{\tau}_r) &\leq \frac{1}{n_r} \left[\frac{\sigma^2}{q(1-q)} + Y_{\max}^2 \left(\frac{(1-q)^2 + q^2}{q(1-q)} + 2 \right) \right] \end{aligned} \quad (78)$$

□

A.4. A Discussion on Estimating the Discount Factors

When the discount factors are not known, an admin is put in the following position: should he discard the (potentially useful) data at lower ranks and just use the data at rank 1 for estimating the treatment effect, or should he try to engage the lower rank data by using part of it to estimate the discount factors and part of it for estimating the treatment? At a first glance, the latter might sound appealing, especially if there was significantly more data at lower ranks than at rank 1. However, this latter option does not actually decrease the estimation error, even in idealized cases. We showcase this claim in an ideal situation below, where the lower rank data exhibits no error in the τ_r estimator (i.e. $\mathbb{E}[(\tau_r - \hat{\tau}_r)^2] = 0$ for $r > 1$). This result uses a naïve estimator for the discount parameters $(\alpha_r)_r$, leaving the problem of using more complex estimators for future work.

To formalize our set-up, we define:

- Scenario (1): an admin uses just the data where their treatment was shown on the first rank, and thus $\tau = \tau_1$. We denote an estimator for this scenario by $\hat{\tau}_{(1)}$, noting that it uses n_1 samples.
- Scenario (2): an admin splits the data at all positions, using part of it for estimating the position effect parameters $(\alpha_r)_r$, and then using these estimates together with the rest of the data for estimating the treatment effect τ . We denote an estimator for this scenario by $\hat{\tau}_{(2)}$.

Lemma A.6. *Assume a naïve estimator for the discount factors: $\hat{\alpha}_r = \hat{\tau}_r / \hat{\tau}_1, \forall r$. For all valid ways of splitting the data between estimating $(\alpha_r)_r$ and estimating τ ,*

$$\mathbb{E}[(\hat{\tau}_{(1)} - \tau)^2] \leq \mathbb{E}[(\hat{\tau}_{(2)} - \tau)^2]$$

Proof. First, splitting the data in a valid way means using some amount of samples n' from both the rank 1 data and rank r data to estimate α_r with the naïve estimator, and using the remaining (non-overlapping) samples to estimate τ . One must have $n' \leq n_1$ and $n' \leq n_r$. Since we have $k - 1$ such estimation problems, for every rank $r \in [k]$ except the first one, one must use $k - 1$ disjoint sample sets from the rank 1 data, totalling at most the number of samples at rank 1 (since we cannot use more data than available).

In scenario (1), the admin can only use the data at rank 1 in order to estimate τ (since it is the treatment effect had the ads been only shown at rank 1 to all subjects). Without any knowledge of $(\alpha_r)_r$, the admin cannot use the data at rank 2, 3, etc to identify the treatment effect through τ_r . Thus, $\tau = \tau_1$, which, as we have showed in Theorem 3.4 and Proposition 3.6, has error $\mathbb{E}[(\tau - \hat{\tau})^2] = \Theta(n_1^{-1})$ (recall that n_1 is the sample size at position 1).

In scenario (2), the admin uses the samples at all positions, relating the effect τ from data at rank r through the $\hat{\alpha}_r$ estimate: $\hat{\tau}^{(r)} = \frac{\hat{\tau}_r}{\hat{\alpha}_r}$. Note that the data used for estimating $\hat{\alpha}_r$ must be separate from the data used for estimating $\hat{\tau}^{(r)}$. Let's say that an admin uses n'_r samples from the n_r samples at each rank for estimating $\hat{\alpha}_r$. The rest of $n_r - n'_r$ samples will be used to estimate the r -estimand. In an ideal case where the estimator at rank r is a perfect estimator, we have $\hat{\tau}_r = \tau_r$. In this case, we only have to worry about the contribution of the $\hat{\alpha}_r$ estimates to the estimation error. For example, one can imagine that we have an abundance of data at lower ranks, $n_r \gg 1$ for $r > 1$, and thus the estimation error of τ_r vanishes with increased number of samples.

Note that $\hat{\tau}_{(2)}$ depends on the values $(n'_r)_r$. However, the result holds no matter what these values are, i.e. no matter how many samples the admin uses from each rank. This is at first unintuitive. The sample size of the (final) estimator in scenario (2), $\sum_r n'_r$, may be much larger than n_1 , and yet, that does not help in reducing the estimation error. At a closer look, we'll see that the error contributed by the discount factors estimation, $(\hat{\alpha}_r)_r$, is a dominant term in the error of the treatment effect.

Take as a naïve estimator for α_r the ratio of the treatment effects at rank r and rank 1: $\hat{\alpha}_r = \frac{\hat{\tau}_r}{\hat{\tau}_1}, \forall r > 1$, using n'_r samples from the n_r samples from rank r and n'_r samples from the n_1 samples at rank 1; i.e., we use $2n'_r$ samples to estimate α_r , split equally between rank 1 and rank r data. The remaining $n_r - n'_r$ samples from the data at rank r are kept for contributing to estimating τ . To better keep track of the number of samples used, we denote an estimator $\hat{\theta}$ based on n samples by $\hat{\theta}(n)$. Then, $\hat{\alpha}_r = \frac{\hat{\tau}_r(n'_r)}{\hat{\tau}_1(n'_r)}, \forall r > 1$. So, $\frac{1}{\hat{\alpha}_r} = \frac{\hat{\tau}_1(n'_r)}{\hat{\tau}_r(n'_r)} = \frac{\hat{\tau}_1(n'_r)}{\tau_r}$.

Using the remaining samples, the r -estimand can be estimated through $\frac{\hat{\tau}_r(n_r - n'_r)}{\hat{\alpha}_r} = \frac{\tau_r}{\hat{\alpha}_r} = \hat{\tau}_1(n'_r)$. Thus, the r -estimand uses the n'_r samples used for the estimator of the discount factors. The $n_r - n'_r$ samples used in the τ_r estimation do not factor in because we assumed it is a perfect estimator.

Thus, the available data at all ranks can be used to estimate τ through the identified effects at each rank, using $n_1 - \sum_r n'_r$ samples for rank 1, n'_2 samples for rank 2, and so on until n'_k samples for rank k . Just as in the main text, $\alpha_1 = 1$, so we can use the entire remaining samples at rank 1 directly, i.e., $n'_1 = 0$.

Theorem 3.4 computes a lower bound based on the combination of the r -estimands. The lower bound obtain is a summation of the errors each of them contributes, namely over the sequence $(n_r \cdot \alpha_r^2)_r$. This connects to the sample value definition from Section 2.1. In our case, the same r -estimands contribute a sequence of errors $(n_1 - \sum_r n'_r, n'_2, \dots, n'_k)$. Applying Theorem 3.4 with the same proof but this updated sequence, we obtain an error lower bound of $\Omega\left(\left((n_1 - \sum_r n'_r) + n'_2 + \dots + n'_k\right)^{-1}\right) = \Omega(n_1^{-1})$. (Note that $n'_1 = 0$ since $\alpha_1 = 1$.) The constants are, in fact, the same as in the original Theorem 3.4. This updated bound does not depend on the amount of samples used to estimate α_r , namely, on $(n'_r)_r$. Thus, the lower bound remains the same no matter how many or how few samples an admin has used to estimate the discount factors.

To see this exemplified, we use the Horvitz-Thompson estimator for $\hat{\tau}_1$, and define the estimand as a weighted sum of the r -estimands: $\hat{\tau} = \sum_r w_r \cdot \hat{\tau}^{(r)}$. As $\hat{\tau}^{(r)}$ are estimated through $\hat{\tau}_1$ applied on either $n_1 - \sum_r n'_r$ or n'_r samples for each rank, it is still unbiased. We can then apply the optimal weighting argument from the main text that dictates that error achieved by $\hat{\tau}$ through optimal weighting is $\mathcal{O}\left(\left(\sum_r 1/\text{Var}\left(\hat{\tau}^{(r)}\right)\right)^{-1}\right)$. We know from Proposition 3.6 that $\text{Var}\left(\hat{\tau}^{(r)}\right) = C \cdot \frac{1}{n_r}$ when there are n_r data samples for a constant C . We thus obtain an upper bound on the treatment effect estimation error of

$$\mathcal{O}\left(\left(\sum_r 1/\text{Var}\left(\hat{\tau}^{(r)}\right)\right)^{-1}\right) = \mathcal{O}\left(\left((n_1 - \sum_r n'_r) + n'_2 + \dots + n'_k\right)^{-1}\right) = \mathcal{O}\left(n_1^{-1}\right), \quad (79)$$

matching the lower bound (up to constants).

One would have imagined that having more samples used from rank 2, 3, and so on that contribute to estimating the treatment effect would have helped reduced the treatment effect estimation error. However, even in this idealized scenario where one has perfect estimation of τ_r at each rank r , the contribution of the discount factors error to the treatment effect error trumps using more low-ranked samples. In realistic scenarios where $\hat{\tau}_r$ is not a perfect estimator for $r > 1$, one would expect that their error contributes to the estimation error in scenario (2). In this case, an admin can only increase their error by attempting to estimate $(\alpha_r)_r$ from an ad campaign, and is better off (up to constants) just using the n_1 samples it has from rank 1, scenario (1). By how much does an admin increases their estimation error is a complex question, depending on the modeling assumptions and estimators chosen for the discount factors. We leave this for future work. Related work has proposed various methods for estimating the discount factors (Chuklin et al., 2016; Joachims et al., 2017), but to our knowledge, none answer this particular question of when is it actually worth doing that to reduce the treatment effect estimation error.

□

A.5. Proofs for Section 4 Results

Proof for Lemma 4.1. Our goal is to show that if an allocation \mathbf{x} satisfies the conditions of Lemma 4.1, i.e. $\mathbf{x}_a(t) \in \{0, 1\}, \forall a \in [k], t \in [n], \sum_{t=1}^n x_a(t) = B^{(a)}, \forall a \in [k]$, and $|\sum_{a=1}^k \mathbf{x}_a(t) - \sum_{a=1}^k \mathbf{x}_a(t')| \leq 1, \forall t, t' \in [n]$, then it is an equilibrium. We show this through two properties:

- **Property 1:** an admin a cannot increase his utility by permuting the bids in \mathbf{x} .
- **Property 2:** an admin a cannot increase his utility by aggregating his budget (merging bids from multiple subjects in \mathbf{x}).

In other words, we show that an admin cannot increase his utility by deviating from \mathbf{x} , by analyzing the two possible types of deviations: permutations or aggregation of bids.

Notation: We introduce simplifying notation in order to prove the two aforementioned properties. As the utility function of an admin is the expected sample value, by linearity of expectation, we can write

$$f_a(\mathbf{x}) = \sum_{r=1}^K \mathbb{E}[n_r] \cdot \alpha_r^2 \quad (80)$$

Given our allocation rule, we can write the utility function in combinatorial form, by finding the closed form of $\mathbb{E}[n_r], \forall r \in [k]$,

$$f_a(\mathbf{x}) = \sum_{t=1}^n \sum_{r=1}^k \alpha_r^2 \cdot \frac{1}{\binom{k-1}{k-r}} \cdot \sum_{s \in \mathcal{S}_{k-r}^{(a)}} \left[\frac{1}{(k-r+1)!} \sum_{\sigma \in \mathcal{P}(s \cup \{a\})} \text{PA}_a(\mathbf{x}_a, t) \cdot \prod_{j <_{\sigma} a} (1 - \text{PA}_j(\mathbf{x}_j, t)) \right] \cdot \prod_{\bar{r} < r} \left(1 - \frac{1}{\binom{k-1}{k-\bar{r}}} \cdot \sum_{\bar{s} \in \mathcal{S}_{k-\bar{r}}^{(a)}} \frac{1}{(k-\bar{r}+1)!} \sum_{\bar{\sigma} \in \bar{\mathcal{P}}(s \cup \{a\})} \text{PA}_a(\mathbf{x}_a, t) \cdot \prod_{\bar{j} <_{\bar{\sigma}} a} (1 - \text{PA}_{\bar{j}}(\mathbf{x}_{\bar{j}}, t)) \right), \quad (81)$$

where $\mathcal{S}_{k-r}^{(a)}$ is the set of all choices of $k-r$ elements of $[k] \setminus \{a\}$ and $j <_{\sigma} a$ means that element j appeared before element a in the permutation σ . The summand is the probability of admin a to win rank r for slot t , by summing over all possible permutations at each rank and possible remaining admins (who have not yet won a rank). The probability of admin a to win a rank is the activation probability times the probability that no one else before a has won, and times the probability that a has not won a previous rank.

To make things a bit more readable, take the summand for a particular subject t outside of the factor α_r^2 and denote it by

$$g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) := \text{PA}_a(\mathbf{x}_a, t) \cdot \frac{1}{\binom{k-1}{k-r}} \cdot \frac{1}{(k-r+1)!} \sum_{s \in \mathcal{S}_{k-r}^{(a)}} \sum_{\sigma \in \mathcal{P}(s \cup \{a\})} \prod_{j <_{\sigma} a} (1 - \text{PA}_j(\mathbf{x}_j, t)) \quad (82)$$

Note that the factor outside of $\text{PA}_a(\mathbf{x}_a, t)$ in equation (82) depends on the rank r but does not contain \mathbf{x}_a or $\text{PA}_a(\mathbf{x}_a, t)$. So since we are looking at $f_a(\mathbf{x}_a, \mathbf{x}_{-a})$ and $g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a})$ as a function of \mathbf{x}_a , we consider this factor a constant in \mathbf{x}_a . For an even more simplifying notation:

$$C_{r,a,t} := \frac{1}{\binom{k-1}{k-r}} \cdot \frac{1}{(k-r+1)!} \sum_{s \in \mathcal{S}_{k-r}^{(a)}} \sum_{\sigma \in \mathcal{P}(s \cup \{a\})} \prod_{j <_{\sigma} a} (1 - \text{PA}_j(\mathbf{x}_j, t)), \quad (83)$$

thus giving

$$g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) = \text{PA}_a(\mathbf{x}_a, t) \cdot C_{r,a,t} \quad (84)$$

Coming back to the utility function, we get

$$f_a(\mathbf{x}_a, \mathbf{x}_{-a}) = \sum_{t=1}^n \sum_{r=1}^k \alpha_r^2 \cdot g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) \cdot \prod_{\bar{r} < r} (1 - g_{\bar{r},a,t}(\mathbf{x}_a, \mathbf{x}_{-a})) \quad (85)$$

Denote by $f_{a,t}$ the summand for each $t \in [1, n]$:

$$f_{a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) = \sum_{r=1}^k \alpha_r^2 \cdot g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) \cdot \prod_{\bar{r} < r} (1 - g_{\bar{r},a,t}(\mathbf{x}_a, \mathbf{x}_{-a})), \quad (86)$$

thus giving

$$f_a(\mathbf{x}_a, \mathbf{x}_{-a}) = \sum_{t=1}^n f_{a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) \quad (87)$$

First of all, note that an allocation in which an admin does not bid all their budget is not an equilibrium. This follows from the fact that all terms in equation (82) are non-negative, and $\text{PA}_a(\mathbf{x}_a, t)$ is an increasing function in $\mathbf{x}_a(t)$.

A necessary property: Before proceeding with our proof, we notice that when $\alpha_r = 1, \forall r \in [k]$, we can simplify the expression of the functions $f_{a,t}$ to:

$$f_{a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) = 1 - \prod_{r=1}^k (1 - g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a})) \quad (88)$$

Of course, this simplification does not hold if all the terms α_r are not equal to each other. However, we will argue that it is sufficient to show Properties 1 and 2 when $\alpha_r = 1, \forall r \in [k]$. To see this, we use the following property:

Proposition A.7. *For any $l, v \in \mathbb{N}$ and any functions $h_1, h_2, \dots, h_l : \mathbb{N}^v \rightarrow \mathbb{R}_+$, if for some $x, x' \in \mathbb{N}$ the following conditions hold:*

$$\begin{aligned} h_1(x) &\geq h_1(x') \text{ and} \\ \sum_{j=1}^l h_j(x) &\geq \sum_{j=1}^l h_j(x'), \end{aligned} \quad (89)$$

then for $\omega_1 = 1$ and for any $\omega_2, \dots, \omega_l \in [0, 1]$ the following holds:

$$\sum_{j=1}^l \omega_j \cdot h_j(x) \geq \sum_{j=1}^l \omega_j \cdot h_j(x') \quad (90)$$

Corollary A.8. *We note that the same result holds as in Proposition A.7 if there are two different functions in the first condition: for any $l, v \in \mathbb{N}$ and any functions $h_1, h_2, \dots, h_l, h'_1 : \mathbb{N}^v \rightarrow \mathbb{R}_+$, if for some $x, x' \in \mathbb{N}$ the following conditions hold:*

$$\begin{aligned} h_1(x) &\geq h'_1(x') \text{ and} \\ h_1(x) + \sum_{j=2}^l h_j(x) &\geq h'_1(x') + \sum_{j=2}^l h_j(x'), \end{aligned} \quad (91)$$

then for any $\omega_2, \dots, \omega_l \in [0, 1]$ the following holds:

$$h_1(x) + \sum_{j=2}^l \omega_j \cdot h_j(x) \geq h'_1(x') + \sum_{j=2}^l \omega_j \cdot h_j(x') \quad (92)$$

Remark A.9. As a remark, Proposition A.7 helps us show that the utility function f_a is increasing in the bid of an admin a for a subject t , $\mathbf{x}_a(t)$. This follows from the following properties:

- First, the function $g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a})$ is increasing in $\mathbf{x}_a(t)$ since $\text{PA}_a(\mathbf{x}, t) = 1 - (1-p)^{\mathbf{x}_a(t)}$ is clearly increasing in $\mathbf{x}_a(t)$. This implies that the function $f_{a,t}(\mathbf{x}_a, \mathbf{x}_{-a})$ is also increasing in $\mathbf{x}_a(t)$ from equation (88) in the case when all α_r 's are equal to 1.
- In the case when all α_r 's are not equal to 1, we employ Proposition A.7 by taking $h_j(\mathbf{x}) = g_{j,a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) \cdot \prod_{\bar{j} < j} (1 - g_{\bar{j},a,t}(\mathbf{x}_a, \mathbf{x}_{-a}))$ and $\omega_j = \alpha_j^2$. Then,

$$\sum_{r=1}^k \omega_j \cdot h_j(\mathbf{x}) = \sum_{r=1}^k \alpha_r^2 \cdot g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) \cdot \prod_{\bar{r} < r} (1 - g_{\bar{r},a,t}(\mathbf{x}_a, \mathbf{x}_{-a})) = f_{a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) \quad (93)$$

Proof of Proposition A.7. We prove Proposition A.7 by induction on l . For $l = 1$, the result is trivial given that $h_1(x) \geq h_1(x')$. For $l = 2$, the conditions are that for some $x, x' \in \mathbb{N}^v$, $h_1(x) \geq h_1(x')$ and $h_1(x) + h_2(x) \geq h_1(x') + h_2(x')$. We wish to show that for any $\omega_2 \in [0, 1]$, we also get that $h_1(x) + \omega_2 \cdot h_2(x) \geq h_1(x') + \omega_2 \cdot h_2(x')$. If $h_2(x) \geq h_2(x')$, this clearly holds as $\omega_2 \geq 0$. If $h_2(x) < h_2(x')$, we get

$$\begin{aligned} h_1(x) + \omega_2 \cdot h_2(x) &\geq h_1(x') + \omega_2 \cdot h_2(x') \Leftrightarrow \\ h_1(x) - h_1(x') + \omega_2 \cdot (h_2(x) - h_2(x')) &\geq 0 \Leftrightarrow \\ h_1(x) - h_1(x') + \omega_2 \cdot (h_2(x) - h_2(x')) &\geq h_1(x) - h_1(x') + h_2(x) - h_2(x') \geq 0 \end{aligned} \quad (94)$$

where the last inequality is true from the conditions of the Proposition. Thus, the base case is proved. For the induction case, assume that the statement of the Proposition is true for $l \geq 2$ and we will show that it is also true for $l + 1$. Given the induction hypothesis, we know that for some $x, x' \in \mathbb{N}^v$, $h_1(x) \geq h_1(x')$ and $\sum_{j=1}^l h_j(x) \geq \sum_{j=1}^l h_j(x')$. If for all $j \in \overline{2, l+1}$, $h_j(x) < h_j(x')$, then we get

$$\sum_{j=1}^{l+1} \omega_j \cdot (h_j(x) - h_j(x')) \geq \sum_{j=1}^{l+1} h_j(x) - h_j(x') \geq 0 \quad (95)$$

since $\omega_j \geq 0, \forall j$ and the unweighted sum is non-negative from the conditions of the Proposition. If, however, $\exists j \in \overline{2, l+1}$ such that $h_j(x) \geq h_j(x')$, then we also know that $h_1(x) + \omega_j \cdot h_j(x) \geq h_1(x') + \omega_j \cdot h_j(x')$. Then, denote by $h'_1(y) = h_1(y) + \omega_j \cdot h_j(y), \forall y \in \mathbb{N}^v$. We know that

$$\begin{aligned} h'_1(x) &\geq h'_1(x') \text{ and that} \\ h'_1(x) + \sum_{\substack{j' \in \overline{2, l+1}, \\ j' \neq j}} \omega_{j'} \cdot h_{j'}(x) &\geq h'_1(x') + \sum_{\substack{j' \in \overline{2, l+1}, \\ j' \neq j}} \omega_{j'} \cdot h_{j'}(x'), \end{aligned} \quad (96)$$

and thus we can apply the induction hypothesis, concluding our proof. The corollary follows immediately by the same proof. \square

The proof of Lemma 4.1 follows the following structure: we use Proposition A.7 to show that if an allocation \mathbf{x}_a has better utility than an allocation \mathbf{x}'_a for a subject slot t when $\alpha_r = 1, \forall r \in [k]$ (meaning that $f_{a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) > f_{a,t}(\mathbf{x}'_a, \mathbf{x}_{-a})$), then \mathbf{x}_a has better utility than \mathbf{x}_a for any $\alpha_r \in [0, 1]$. This follows by taking $l = k, \omega_r = \alpha_r^2, \forall r \in [k]$, and $h_r(x) = g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) \cdot \prod_{\bar{r} < r} (1 - g_{\bar{r},a,t}(\mathbf{x}_a, \mathbf{x}_{-a}))$. Lastly, we will verify the first condition of the Proposition A.7, that $h_1(x) \geq h'_1(x')$ for each of the Properties 1 and 2, below. If that holds, then we conclude that if an allocation satisfies Properties 1 and 2 when $\alpha_r = 1, \forall r \in [k]$, then it satisfies Properties 1 and 2 for any $\alpha_r \in [0, 1], \forall r \in \overline{2, k}$ (knowing that we always set $\alpha_1 = 1$). Finally, if an allocation \mathbf{x} satisfies Properties 1 and 2, then it is a Nash equilibrium.

Proof of Property 1: Then, we proceed to show Property 1. Without loss of generality, assume $a = 1$. If there is no subject on which admin 1 has allocated 0 budget or no subject on which admin 1 has allocated 1 budget, there is nothing to show. Therefore, take two subjects, t_0 and t_1 , such that $\mathbf{x}_1(t_0) = 0$ and $\mathbf{x}_1(t_1) = 1$. We would like to show that admin 1 cannot increase their utility by moving the unit of budget from t_1 to t_0 . If the allocations of all other admins on subjects t_0 and t_1 are exactly the same, then admin i has the same utility from allocating 1 unit of budget on t_0 or t_1 . Moreover, given the assumption that the activation probability is the same across admins, if the allocations of all other admins on subjects t_0 and t_1 have the same number of 0s and 1s, then, again, admin i has the same utility from allocating 1 unit of budget on t_0 or t_1 . If, however, the allocations of all other admins on subjects t_0 and t_1 have a different number of 0s and 1s, then we know that the number of 0s in the allocations of t_1 is higher than in t_0 (excluding admin 1). This is because of the property stated in the allocation described in Lemma 4.1: the allocation \mathbf{x} minimizes the number of 0s in each $\mathbf{x}_{\cdot,t}$ (as admin 1 prioritizes subjects with least bids); therefore, the number of 0s in the allocations of t_1 is higher than in t_0 (excluding admin 1). We are almost done at this point, as we only need the following remark:

Remark A.10. Take the following two allocations of all admins over subject t :

$$\mathbf{x}_{\cdot,t} = \begin{pmatrix} \mathbf{x}_{1t} \\ \mathbf{x}_{2t} \\ \vdots \\ \mathbf{x}_{kt} \end{pmatrix} \text{ and } \mathbf{x}'_{\cdot,t} = \begin{pmatrix} \mathbf{x}'_{1t} \\ \mathbf{x}'_{2t} \\ \vdots \\ \mathbf{x}'_{kt} \end{pmatrix}, \quad (97)$$

where $\mathbf{x}_{\cdot,t}$ has the first j coordinates as 1s and the rest zeros, and $\mathbf{x}'_{\cdot,t}$ has the first $j + 1$ coordinates as 1s and the rest zeros, for some $j \geq 0$. Then,

$$f_{a,t}(\mathbf{x}_{\cdot,t}) \geq f_{a,t}(\mathbf{x}'_{\cdot,t}) \quad (98)$$

From the assumption that $p_a = p \in [0, 1]$ is the same across admins and subjects, we get that Remark A.10 is true even for different subjects, and moreover, even for different permutations of the coordinates of $\mathbf{x}_{\cdot,t}$ and $\mathbf{x}'_{\cdot,t}$. (Basically, any allocation that contains j ones and $k - j$ zeros will have better utility for admin a than any allocation that contains $j + 1$ ones and $k - j - 1$ zeros.)

Using Remark A.10 iteratively, this shows that the utility of admin 1 cannot increase if it moves a unit of budget from a subject with fewer allocations of budget units from the other admins to a subject with more allocations of budget units from other admins.

Proof of Remark A.10. Finally, to prove Remark A.10, we simply use the fact that, if we take two allocations of admin j as \mathbf{x}_j and \mathbf{x}'_j that only differ on coordinate t (for some $t \in [n]$), for example $\mathbf{x}_j(t) = 0$ and $\mathbf{x}'_j(t) = 1$, then $\text{PA}_j(\mathbf{x}_j, t) \leq \text{PA}_j(\mathbf{x}'_j, t)$. This is true since $\text{PA}_j(\mathbf{x}_j, t) = 0$ and $\text{PA}_j(\mathbf{x}'_j, t) = p_j \geq 0$. Thus, we get that

$$\begin{aligned} \text{PA}_j(\mathbf{x}_j, t) &\leq \text{PA}_j(\mathbf{x}'_j, t) \Rightarrow \\ 1 - \text{PA}_j(\mathbf{x}_j, t) &\geq 1 - \text{PA}_j(\mathbf{x}'_j, t) \Rightarrow \\ C_{r,a,t}(\mathbf{x}_{-a}) &\geq C_{r,a,t}(\mathbf{x}'_{-a}), \forall r \in [k] \Leftrightarrow \\ g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) &\geq g_{r,a,t}(\mathbf{x}_a, \mathbf{x}'_{-a}), \forall r \in [k], \text{ where } \mathbf{x}_j \in \mathbf{x}_{-a} \text{ and } \mathbf{x}'_j \in \mathbf{x}'_{-a} \Rightarrow \\ 1 - g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) &\leq 1 - g_{r,a,t}(\mathbf{x}_a, \mathbf{x}'_{-a}), \forall r \in [k] \Rightarrow \\ \prod_{r=1}^k (1 - g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a})) &\leq \prod_{r=1}^k (1 - g_{r,a,t}(\mathbf{x}_a, \mathbf{x}'_{-a})) \Rightarrow \\ 1 - \prod_{r=1}^k (1 - g_{r,a,t}(\mathbf{x}_a, \mathbf{x}_{-a})) &\geq 1 - \prod_{r=1}^k (1 - g_{r,a,t}(\mathbf{x}_a, \mathbf{x}'_{-a})) \Rightarrow \\ f_{a,t}(\mathbf{x}_a, \mathbf{x}_{-a}) &\geq f_{a,t}(\mathbf{x}_a, \mathbf{x}'_{-a}), \end{aligned} \quad (99)$$

From the fourth inequality above, we also deduce that the first condition for the Proposition A.7 is satisfied by taking $h_1 = g_{1,a,t}$. Noting that \mathbf{x}_j and \mathbf{x}'_j differ on just one coordinate and we showed that the allocation that contained one more budget unit on subject t has a better utility for admin i , this concludes the proof of the remark. As mentioned above, using Proposition A.7, the remark holds for any $\alpha_r \in [0, 1], \forall r \in [2, k]$ as well. \square

Proof of Property 2: Next, we that an admin does not prefer to aggregate his spent budget across any subject slots. Again, without loss of generality, take admin $a = 1$ and take two allocations over two subject slots, from an allocation \mathbf{x} described in Lemma 4.1 (and we will argue for any number at the end, as it easily generalizes):

$$\mathbf{x}_{\cdot,t} = \begin{pmatrix} \mathbf{x}_{1t} \\ \vdots \\ \mathbf{x}_{kt} \end{pmatrix} \text{ and } \mathbf{x}'_{\cdot,t'} = \begin{pmatrix} \mathbf{x}'_{1t'} \\ \vdots \\ \mathbf{x}'_{kt'} \end{pmatrix} \quad (100)$$

An admin can aggregate his budget spent on t and t' on either of them, or a new subject t'' . For our proof, it does not matter, since the property we use is that the number of 0s among the others admins' bids for any of the subjects t, t', t'' does not differ by more than one. Thus, let's assume that the admin aggregates his budget from t and t' onto subject t (instead of spending \mathbf{x}_{1t} and $\mathbf{x}_{1t'}$, he spends $\mathbf{x}_{1t} + \mathbf{x}_{1t'}$ and 0 on subjects t and t' , respectively). If either \mathbf{x}_{1t} or $\mathbf{x}'_{1t'}$ are equal to 0, then there is nothing to aggregate since we reduce to Property 1, so we only deal with the case $\mathbf{x}_{1t} = \mathbf{x}'_{1t'} = 1$. Since the distribution of 0s and 1s are such that the number of 0s on each subject slot of \mathbf{x} is minimized, a pair of allocations of two subject slots cannot differ by more than 1 in their number of 0s. In addition, given the assumption that the relevance probability is the same across subject slots and admins, the order of the ones and zeros does in the allocations does not matter.

Case 1: $\mathbf{x}_{\cdot,t}$ and $\mathbf{x}'_{\cdot,t'}$ have the same number of zeros and ones. Given that the order does not matter, $f_{1,t}(\mathbf{x}_{\cdot,t}) = f_{1,t'}(\mathbf{x}_{\cdot,t'})$. Aggregating the spent budget into $\mathbf{x}_{\cdot,t}$ would result into an allocation like

$$\mathbf{x}^a_{\cdot,t} = \begin{pmatrix} \mathbf{x}_{1t} + \mathbf{x}'_{1t'} \\ \mathbf{x}_{2t} \\ \vdots \\ \mathbf{x}_{kt} \end{pmatrix} = \begin{pmatrix} 2 \\ \mathbf{x}_{2t} \\ \vdots \\ \mathbf{x}_{kt} \end{pmatrix} \quad (101)$$

as $\mathbf{x}_{1t} = \mathbf{x}'_{1t'} = 1$.

Of course, the admin can aggregate their spent budget into either allocation $\mathbf{x}_{\cdot,t}$ or $\mathbf{x}_{\cdot,t'}$, but that does not change the utility contribution in this case. We are thus left to show that

$$2 \cdot f_{1,t}(\mathbf{x}_{\cdot,t}) \geq f_{1,t}(\mathbf{x}^a_{\cdot,t}), \quad (102)$$

which would also imply that the first condition of Corollary A.8 is satisfied, by taking $h_1 = 2f_{1,t}$ and $h'_1 = f_{1,t}$.

Since all the coordinates except the first one are the same in $\mathbf{x}_{\cdot,t}$ and $\mathbf{x}^a_{\cdot,t}$, all the $C_{r,a,t}$ terms in the utility function expression will be same. We also show that

$$\begin{aligned} 2 \cdot \text{PA}_1(\mathbf{x}_{\cdot,t}, t) &\geq \text{PA}_1(\mathbf{x}^a_{\cdot,t}, t) \Leftrightarrow \\ 2 \cdot p &\geq 1 - (1-p)^2 \Leftrightarrow \\ 2 &\geq 1 + (1-p), \end{aligned} \quad (103)$$

which is true since $p \in [0, 1]$. This also implies that $g_{1,a,t}(\mathbf{x}_{\cdot,t}) \geq g_{1,a,t}(\mathbf{x}^a_{\cdot,t})$. Finally,

$$\begin{aligned} 1 - \text{PA}_1(\mathbf{x}_{\cdot,t}, t) \cdot C_{r,1,t} &\geq 1 - \text{PA}_1(\mathbf{x}^a_{\cdot,t}, t) \cdot C_{r,1,t} \Leftrightarrow \\ \text{PA}_1(\mathbf{x}_{\cdot,t}, t) \cdot C_{r,1,t} &\leq \text{PA}_1(\mathbf{x}^a_{\cdot,t}, t) \cdot C_{r,1,t} \Leftrightarrow \\ \text{PA}_1(\mathbf{x}_{\cdot,t}, t) &\leq \text{PA}_1(\mathbf{x}^a_{\cdot,t}, t) \Leftrightarrow \\ 1 - (1-p) &\leq 1 - (1-p)^2 \Leftrightarrow \\ (1-p)^2 &\leq (1-p), \end{aligned} \quad (104)$$

which is true for any $r \in [k]$ since $p \in [0, 1]$. We note that these last two inequalities remains true if we replace 2 by N , for any $N \geq 2$. Combining inequalities (103) and (104) in the formula for $f_{1,t}$ and noting that all terms are non-negative, we obtain the desired result with Corollary A.8.

Case 2: $\mathbf{x}_{\cdot,t}$ and $\mathbf{x}'_{\cdot,t}$ differ by 1 in the number of ones they contain. It is sufficient to consider the case where $\mathbf{x}_{\cdot,t}$ has j ones and $\mathbf{x}'_{\cdot,t}$ has $j + 1$ ones, for some $j \geq 1$. Given Remark A.10, it is sufficient to show

$$2 \cdot f_{1,t'}(\mathbf{x}'_{\cdot,t'}) \geq f_{1,t}(\mathbf{x}^a_{\cdot,t}) \quad (105)$$

with $\mathbf{x}^a_{\cdot,t}$ defined as in equation (101) (since $f_{1,t}(\mathbf{x}_{\cdot,t}) \geq f_{1,t'}(\mathbf{x}'_{\cdot,t'})$). We will show that for each $r \in [k]$:

$$2 \cdot \text{PA}_1(\mathbf{x}'_{\cdot,t'}, t') \cdot C'_{r,1,t} \geq \text{PA}_1(\mathbf{x}^a_{\cdot,t}, t) \cdot C_{r,1,t}, \quad (106)$$

where the terms $C_{r,a,t}$ are defined as in equation (83) and that

$$\begin{aligned} 2 \cdot \text{PA}_1(\mathbf{x}'_{\cdot,t'}, t') \cdot C'_{r,1,t} \cdot \prod_{\bar{r} < r} (1 - \text{PA}_1(\mathbf{x}'_{a,t'}, t') \cdot C'_{\bar{r},1,t}) &\geq \\ \text{PA}_1(\mathbf{x}^a_{\cdot,t}, t) \cdot C_{r,1,t} \cdot \prod_{\bar{r} < r} (1 - \text{PA}_1(\mathbf{x}^a_{\cdot,t}, t) \cdot C_{\bar{r},1,t}) & \end{aligned} \quad (107)$$

First of all, we note that the difference between $C_{r,1,t}$ and $C'_{r,1,t}$ is that there is an extra 1 in the allocation of $\mathbf{x}'_{\cdot,t}$ than in $\mathbf{x}_{\cdot,t}$. This extra 1 contributes to the change between $C_{r,1,t}$ and $C'_{r,1,t}$ in the following way: let's say that $C_{r,1,t}$ is composed of two types of terms, the terms in which the extra one shows up in a factor $(1 - \text{PA}_j(\cdot))$, and the terms in which it does not. Thus, let's write $C_{r,1,t}$ as $C_{r,1,t} = C_{r,1,t}^{(j)} + C_{r,1,t}^{(-j)}$. We now note that $C'_{r,1,t} = C_{r,1,t}^{(j)} \cdot (1 - p) + C_{r,1,t}^{(-j)}$, since the terms that did not contain the extra 1 coordinate did not change, while those that did simply changed from 1 to $1 - p$. This also shows that $C'_{r,1,t} \leq C_{r,1,t}$ for $p \in [0, 1]$.

Thus, for any $\bar{r} \in [k]$,

$$\begin{aligned} 1 - \text{PA}_1(\mathbf{x}'_{\cdot,t'}, t') \cdot C'_{\bar{r},1,t} &\geq 1 - \text{PA}_1(\mathbf{x}^a_{\cdot,t}, t) \cdot C_{\bar{r},1,t} \Leftrightarrow \\ \text{PA}_1(\mathbf{x}'_{\cdot,t'}, t') \cdot C'_{\bar{r},1,t} &\leq \text{PA}_1(\mathbf{x}^a_{\cdot,t}, t) \cdot C_{\bar{r},1,t} \Leftrightarrow \\ \text{PA}_1(\mathbf{x}'_{\cdot,t'}, t') &\leq \text{PA}_1(\mathbf{x}^a_{\cdot,t}, t) \Leftrightarrow \\ 1 - (1 - p) &\leq 1 - (1 - p)^2 \Leftrightarrow \\ 1 - p &\leq (1 - p)^2, \end{aligned} \quad (108)$$

which is true since $p \in [0, 1]$. Finally, we also show that

$$\begin{aligned} 2 \cdot \text{PA}_1(\mathbf{x}'_{\cdot,t'}, t') \cdot C'_{r,1,t} &\geq \text{PA}_a(\mathbf{x}^a_{\cdot,t}, t) \cdot C_{r,1,t} \Leftrightarrow \\ 2(1 - (1 - p)) \cdot \left((1 - p) \cdot C_{r,1,t}^{(j)} + C_{r,1,t}^{(-j)} \right) &\geq (1 - (1 - p)^2) \cdot \left(C_{r,1,t}^{(j)} + C_{r,1,t}^{(-j)} \right) \Leftrightarrow \\ C_{r,1,t}^{(-j)} &\geq C_{r,1,t}^{(j)}, \end{aligned} \quad (109)$$

which is true for all $r \in [k]$ since the sum of the terms in permutations where 1 shows up *before* j (j being the coordinate of the extra budget unit in the allocation $\mathbf{x}'_{\cdot,t'}$) is at least equal to the sum of the terms in permutations where 1 shows up *after* j (each permutation σ in which 1 shows up before j can be ‘paired’ with a permutation σ in which 1 is after j by just switching their position. Thus, by the definition of the utility function, the product contribution coming from σ is at most the product contribution coming from σ'). We note that inequality (109) holds when we replace 2 by any $N \geq 2$ since the function $\frac{N \cdot p}{1 - (1 - p)^N}$ is increasing in N (while all other terms do not contain N). Similarly as in Case 1, the conditions for Corollary A.8 hold for $h_1 = 2g_{r,a,t}$ and $h'_1 = g_{r,a,t}$.

□

Proof of Proposition 4.3. Assume by contradiction that there is a Nash equilibrium allocation \mathbf{x} that does not satisfy the conditions of Lemma 4.1. First, we know admins always spend all their budget in a Nash equilibrium from the proof of Lemma 4.1. Second, from the same proof we also know that if the bids are all 0 or 1, then admins also minimize competition (property 1). Therefore, the only way in which an allocation x does not fulfill the conditions of Lemma 4.1 and might *still* be a Nash equilibrium is if some of the bids are not 0 or 1. Thus, assume without loss of generality that there exists a subject of subjects $T \subset [n]$ for which $\mathbf{x}_{at} > 1 \forall t \in T$, for an admin a . Without loss of generality, assume that $a = 2$ (the second admin). We have two cases:

- Case 1: there is a subject $t \in T$ such that $\mathbf{x}_{1t} = 0$.
- Case 2: there is no subject $t \in T$ such that $\mathbf{x}_{1t} = 0$.

Let's first tackle Case 1. In this case, we know that there is a subject $t \in T$ such that $\mathbf{x}_{1t} = 0$. Since $\mathbf{x}_{2t} > 1$ for $t \in T$ and $B_2 \leq n$, there exists a subject $t' \notin T$ such that $\mathbf{x}_{2t'} = 0$. Denote $\mathbf{x}_{2t} := x$ and $\mathbf{x}_{1t'} := 0$. We will show that the configuration $\begin{pmatrix} 0 & x' \\ x & 0 \end{pmatrix}$ is not a Nash equilibrium for $x, x' > 1$. If one of x or x' is equal to 1, then we know this is not Nash because of Property 2 in the proof of Lemma 4.1. Assume without loss of generality that $x \leq x'$. We will show that

$$f_1 \begin{pmatrix} 0 & x' \\ x & 0 \end{pmatrix} < f_1 \begin{pmatrix} 1 & x' - 1 \\ x & 0 \end{pmatrix} \quad (110)$$

We want to show this for any sequence $(\alpha_r)_r$. However, due to Proposition A.7, we will argue that it is sufficient to show it for $\alpha_r = 1, \forall r \in [k]$. To see why Proposition A.7 applies, we take as functions h_r each term corresponding to α_r^2 , for all $r \in [k]$. We also compute the terms corresponding to α_1^2 in the utility function, denoting $\bar{p} := 1 - p$ for ease of notation. Using the combinatorial formula for the utility, we get that the coefficient of α_1^2 in $f_1 \begin{pmatrix} 0 & x' \\ x & 0 \end{pmatrix}$ is $1 - \bar{p}^{x'}$, and the coefficient of α_1^2 in $f_1 \begin{pmatrix} 1 & x' - 1 \\ x & 0 \end{pmatrix}$ is $1 - \bar{p}^{x'-1} + \frac{1}{2}(1 - \bar{p}) \cdot (1 + \bar{p}^x)$. Since $x \leq x'$ and all terms are non-negative, it is sufficient to show that

$$\begin{aligned} 1 - \bar{p}^{x'} &< 1 - \bar{p}^{x'-1} + \frac{1}{2}(1 - \bar{p}) \cdot (1 + \bar{p}^x) \Leftrightarrow \\ \bar{p}^{x'-1} \cdot (1 - \bar{p}) &< \frac{1}{2}(1 - \bar{p}) \cdot (1 + \bar{p}^x) \Leftrightarrow \\ \bar{p}^{x'-1} &< \frac{1}{2} \cdot (1 + \bar{p}^x), \end{aligned} \quad (111)$$

which is true for $x' > 1$. We have used that $\bar{p} \in (0, 1)$. Thus, the conditions for Proposition A.7 would be satisfied. We are left to show that inequality (110) happens for $\alpha_r = 1, \forall r \in [k]$. Writing out the utility functions from their combinatorial form, we get

$$f_1 \begin{pmatrix} 0 & x' \\ x & 0 \end{pmatrix} = (1 - \bar{p}^{x'}) \cdot (1 + \bar{p}^x) \quad (112)$$

and

$$f_1 \begin{pmatrix} 1 & x' - 1 \\ x & 0 \end{pmatrix} = (1 - \bar{p}^{x'-1}) \cdot (1 + \bar{p}^{x'-1}) + (1 - \bar{p}) \cdot \left(1 + \frac{1}{2}\bar{p}(1 + \bar{p}^x)\right) \quad (113)$$

Again, since $x \leq x'$ and all terms are non-negative, it is sufficient to show

$$(1 - \bar{p}^{x'}) \cdot (1 + \bar{p}^x) < (1 - \bar{p}^{x'-1}) \cdot (1 + \bar{p}^{x'-1}) + (1 - \bar{p}) \cdot \left(1 + \frac{1}{2}\bar{p}(1 + \bar{p}^x)\right) \quad (114)$$

Equation (114) simplifies to

$$\bar{p}^{2x'-2}(1 + \bar{p}) < 1 + \frac{1}{2}\bar{p}(1 + \bar{p}^{x'}) \quad (115)$$

This is true for $x' > 2$ since $\bar{p}^{2x'-1} < 1$, $\bar{p}^{2x'-2} < \bar{p}$, and $\bar{p}^{2x'-2} < \bar{p}^{x'+1}$. For $x' = 2$, the equation simplifies to

$$\bar{p}^2 + \frac{1}{2}\bar{p}^3 < 1 + \frac{1}{2}\bar{p}, \quad (116)$$

We next tackle Case 2, for which we know that there is no subject $t \in T$ such that $\mathbf{x}_{1t} = 0$. Now, if there exists a subject $t \notin T$ such that $\mathbf{x}_{1t} > 1$, there must also exist a subject $t' \notin T$ such that $\mathbf{x}_{1t'} = 0$, since $B_1 \leq n$. However, we know that $\mathbf{x}_{2t}, \mathbf{x}_{2t'} \in \{0, 1\}$ as $t, t' \notin T$, by definition of the set T . However, none of the configurations $\begin{pmatrix} \mathbf{x}_{1t} & 0 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} \mathbf{x}_{1t} & 0 \\ 1 & 1 \end{pmatrix}, \begin{pmatrix} \mathbf{x}_{1t} & 0 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} \mathbf{x}_{1t} & 0 \\ 0 & 1 \end{pmatrix}$ is an equilibrium for $\mathbf{x}_{1t} > 1$, from Property 2 of the proof of Lemma 4.1 (when the other bids are 0 or 1, an admin always wants to split their bid equally). That means that for all subjects $t \notin T$, $\mathbf{x}_{1t} \in \{0, 1\}$. We will also show that for all subjects $t \in T$, $\mathbf{x}_{1t} \in \{0, 1\}$. By contradiction, if there is a subject $t \in T$, $\mathbf{x}_{1t} > 1$, take a subject $t' \notin T$ such that $\mathbf{x}_{2t'} = 0$ (we know it exists since admin's 2 bids in T are greater than 1 and $B_2 \leq n$). Then, the configurations $\begin{pmatrix} \mathbf{x}_{1t} & \mathbf{x}_{1t'} \\ \mathbf{x}_{2t} & 0 \end{pmatrix}$ are not equilibria for $\mathbf{x}_{1t}, \mathbf{x}_{2t} > 1$ and $\mathbf{x}_{1t'} \in \{0, 1\}$, easy to show by computing the utility function in its combinatorial form and using $\bar{p} \in (0, 1)$. Thus, we have shown that for all subjects t , $\mathbf{x}_{1t} \in \{0, 1\}$. However, we now have a contradiction, since by Property 2, $\mathbf{x}_{2t} \in \{0, 1\}$ as well, which we assumed otherwise for $t \in T$. \square