



Marketing Science Institute Working Paper Series 2025

Report No. 25-141

Estimating Treatment Effects under Recommender Interference: A Structured Neural Networks Approach

Ruohan Zhan, Shichao Han, Yuchen Hu and Zhenling Jiang

“Estimating Treatment Effects under Recommender Interference: A Structured Neural Networks Approach” © 2025

Ruohan Zhan, Shichao Han, Yuchen Hu and Zhenling Jiang

MSI Working Papers are Distributed for the benefit of MSI corporate and academic members and the general public. Reports are not to be reproduced or published in any form or by any means, electronic or mechanical, without written permission.

Estimating Treatment Effects under Recommender Interference: A Structured Neural Networks Approach

Ruohan Zhan Shichao Han Yuchen Hu Zhenling Jiang*
ruohan.zhan@ucl.ac.uk shichaohan@tencent.com yuchenhu@mit.edu zhenling@wharton.upenn.edu
University College London Tencent Inc. Massachusetts Institute of Technology University of Pennsylvania

Recommender systems are essential for content-sharing platforms by curating personalized content. To improve recommender systems, platforms frequently rely on creator-side randomized experiments to evaluate algorithm updates. We show that commonly adopted difference-in-means estimators can lead to severely biased estimates due to recommender interference, where treated and control creators compete for exposure. This bias can result in incorrect business decisions. To address this, we propose a “recommender choice model” that explicitly represents the interference pathway. The approach combines a structural choice framework with neural networks to account for rich viewer-content heterogeneity. Building on this foundation, we develop a debiased estimator using the double machine learning (DML) framework to adjust for errors from nuisance component estimation. We show that the estimator is \sqrt{n} -consistent and asymptotically normal, and we extend the DML theory to handle correlated data, which arise in our context due to overlapped items. We validate our method with a large-scale field experiment on Weixin short-video platform, using a costly double-sided randomization design to obtain an interference-free ground truth. Our results show that the proposed estimator successfully recovers this ground truth, whereas benchmark estimators exhibit substantial bias, and in some cases, yield reversed signs.

Key words: Recommender Interference; Treatment Effect Estimation; Online Content Platforms; Creator-side Randomization; Semiparametric Choice Model; Double/Debiased Estimation and Inference

We are grateful to Susan Athey, Tat Chan, Wanning Chen, Ryan Dew, Yue Fang, Bryan Graham, David Holtz, Ganesh Iyer, Ramesh Johari, Jussi Keppo, Hannah Li, Ilan Lobel, Xiaojie Mao, Tu Ni, Nian Si, K Sudhir, Zikun Ye, Dennis Zhang, Jinglong Zhao, conference participants at CODE@MIT 2023, Conference on AI/ML and Business Analytics 2023, ACIC 2024, QME 2024, China India Insights Conference 2025, and seminar attendees at University of California Berkeley, The Chinese University of Hong Kong, Carnegie Mellon University, New York University, University of Washington, National University of Singapore, Tsinghua University, The Chinese University of Hong Kong (Shenzhen) for their helpful comments. The short version of this paper has appeared in the ACM Conference on Economics and Computation (EC’24), and we are especially thankful to the anonymous reviewers for their constructive feedback. We are grateful to the WeChat Experimentation Team at Tencent for their generous support, particularly Darwin Yong Wang, who has helped shape and advance this work’s applications in the online platform.

*Corresponding author.

1. Introduction

Recommender systems (RecSys) are at the heart of online content platforms, curating viewer-specific selections that drive content consumption (Kiros 2022). To improve their RecSys, platforms regularly develop and test new algorithms by running experiments. In this paper, we focus on a particularly important and common type of RecSys experiments that *randomize creators* rather than viewers. This creator-centric approach is essential for two reasons. Platforms rely on creators to generate a diverse and vibrant content pool, making it crucial to evaluate how the recommender change impacts content creators (Mehrotra et al. 2018, Rosen 1981). Moreover, many business-relevant levers, such as new content boosting strategies, are targeted directly at creators. These creator-focused experiments are practically important and inherently require creator-level randomization for their evaluation.

A common approach to analyze these experiments is the difference-in-means (DIM) estimator, which compares outcomes between treated and control creators. However, we show that this standard method is unreliable and can lead to severely biased estimates. This bias can lead to wrong business decisions. In a large-scale field experiment we conducted on Weixin Channels, a leading short-video platform, we show that a standard DIM estimator concludes a new algorithm has a significant positive effect on an outcome. In reality, the ground truth estimate reveals the algorithm has a significant negative effect. Relying on the standard, flawed method would have caused the platform to make the wrong decision and deploy a worse recommendation algorithm.

The failure of the DIM estimator occurs because it ignores the fact that treated and control creators compete for exposure through the recommendation algorithm. Consequently, an outcome for one creator depends on the treatment status of others, which violates the Stable Unit Treatment Value Assumption (SUTVA) (Imbens 2004), and leads to interference bias (Bajari et al. 2021, Johari et al. 2022, Bright et al. 2022, Goli et al. 2023). This interference manifests in two distinct forms of bias. First, if the treatment successfully boosts the scores of treated items, they will be shown more often, “crowding out” control items and leading to *item exposure bias*. Second, highly personalized algorithms can cause *viewer selection bias*, where the characteristics of viewers exposed to treated content systematically differ from those exposed to control content, a phenomenon recently documented in advertising A/B testing (Braun and Schwartz 2025, Burtch et al. 2025). Because the standard DIM estimator fails to account for these biases, it produces unreliable results.

The goal of this paper is to develop a reliable approach for estimating treatment effects in creator-side experiments that correctly accounts for this interference. The business question is to evaluate whether to adopt the new recommender algorithm for all creators. Therefore, the relevant causal estimand is the average treatment effect, which measures the system-wide impact of such a global rollout. While a natural first thought might be to simply simulate the new algorithm’s behavior

offline, this approach is notoriously unreliable as it can neither replicate the complex, stochastic behavior of a live recommender system nor the true counterfactual responses of viewers. A method that correctly analyzes data from a live experiment is therefore essential.

Our approach involves modeling two key components: which item the recommender exposes and the expected viewer response. For the first component, we propose a “recommender choice model” to explicitly model the interference pathway where treated and control units compete for exposure. This model is semi-parametric: its structural component enables counterfactual evaluation, while a flexible neural network learns the complex mapping from viewer-content pairs to scores. By conditioning exposure probabilities on viewer characteristics, this model captures the systematic differences in audiences reached by treated versus control units. For the second component, a “viewer response model” uses another flexible neural network to predict the outcome for any given item-viewer pair once exposure occurs.

With the choice and response models serving as nuisance components for the average treatment effect, we propose a *debiased estimator* grounded in the Double/Debiased machine learning (DML) framework (Chernozhukov et al. 2018, Farrell et al. 2021a). Besides the plug-in estimate using the choice and viewer response models, a debiasing term adjusts for errors in neural network nuisance estimates. This procedure ensures our estimator satisfies the Neyman orthogonality condition, which makes it robust to the imperfections of the choice and response models and allows it to achieve \sqrt{n} -consistency and asymptotic normality under standard regularity conditions.

A key methodological contribution of our paper is extending the DML framework from its canonical i.i.d. samples to handle the correlated data in our setting, which arises from overlapped items in consideration sets. To do so, our proof models the data generating process sequentially and apply martingale limiting theorems (Hall and Heyde 2014) to establish the estimator’s asymptotics properties of an oracle estimator with known nuisances. Because of the Neyman orthogonality condition, the difference between the debiased and oracle estimators is shown to be small. This extension not only validates our approach but also broadens the applicability of the DML framework to other important settings with sample correlation, such as panel data.

We first validate our approach using Monte Carlo simulations, which allow us to evaluate performance where the ground truth is known. We test our debiased estimator against several benchmarks, including difference-in-means (DIM) estimators, regression-based methods, and propensity-based methods. The results confirm our theoretical predictions: while our proposed estimator successfully recovers the true treatment effect and provides valid statistical inference, the benchmark estimators are biased and produce unreliable confidence intervals.

Beyond Monte Carlo studies, we conduct a large-scale experiment on the Weixin short-video platform to validate the performance of the proposed method. Besides running a creator-side experiment,

our validation relies on a double-sided experiment (Ye et al. 2023a, Su et al. 2024). Specifically, we partition viewers and creators into three equal-sized distinct sub-universes, where viewers can only access creators within their own sub-universe. One sub-universe runs the creator-side experiment, while the other two implement the treatment and control algorithms, respectively. While this double-sided design provides an interference-free ground truth because treated and control units do not compete, it is very costly for routine use due to several practical challenges. Partitioning the platform for hundreds of concurrent experiments destroys statistical power. The resulting “market thinning” degrades the experience for both viewers and creators. Furthermore, implementing and maintaining these parallel sub-universes requires significant and ongoing engineering resources.

We find that our proposed estimator, using only data from the standard creator-side experiment, yields results comparable to the ground-truth estimates from the costly double-sided design. In contrast, benchmark estimators exhibit significant bias and, in some cases, even produce effects with the wrong sign, leading to incorrect business decisions. This paper’s contributions are therefore twofold: substantively, we provide a reliable solution to a common and costly problem in creator experiments; and methodologically, we extend the DML framework to handle correlated data.

1.1. Related Literature

There is a growing literature studying interference in randomized experiments, where an individual’s outcome is affected by others’ treatments, violating the Stable Unit Treatment Value Assumption (SUTVA). Interference can lead to biased treatment effect estimates under standard A/B testing framework (Sävje et al. 2021, Hu et al. 2021, Farias et al. 2022, Johari et al. 2022, Dhaouadi et al. 2023, Zhu et al. 2024). To correctly estimate treatment effect, most of the existing literature leverages innovative *experimental designs* to address the bias from interference. For example, one can use a clustered randomization design if interference primarily occurs within a cluster (e.g., on social networks) and then randomizes treatment at the cluster level (Holtz and Aral 2020, Holtz et al. 2023, Ugander et al. 2013, Hudgens and Halloran 2008). Without a clear cluster structure, a switchback design can be applied that assigns a treatment condition randomly at the market level across different time periods (Bojinov et al. 2023, Hu and Wager 2022, Ni et al. 2023, Xiong et al. 2023). These experimental designs can be very useful but do not apply to all settings, including ours, where the content markets are too interconnected to form clusters and there exist significant time varying factors influencing the outcomes.

Our work aligns more closely with an alternative approach that proposes innovative *estimators* while leveraging standard experimental design. This requires domain-specific analysis of interference types. For example, in experiments for ranking algorithms, Goli et al. (2023) address interference by leveraging items with positions close to those under the counterfactual ranking using historical

A/B test data. In marketplaces where the platform matches supply and demand via linear programming, [Bright et al. \(2022\)](#) propose a shadow-price based estimator to mitigate interference bias. For interference bias that arises from consumers impacting others through limited inventory, [Farias et al. \(2023\)](#) model consumer behavior via a Markovian structure to address interference. Our paper proposes an estimator that applies to a common source of interference where treated and control units compete for exposure in online platforms.

This paper draws from recent advances in double machine learning and semiparametric inference ([Newey 1994](#), [Chernozhukov et al. 2018](#), [Farrell et al. 2021a](#), [Chernozhukov et al. 2019](#)). In particular, the proposed treatment effect estimator directly builds on the doubly robust estimators for semi-parametric models proposed by [Farrell et al. \(2021a\)](#), where the parametric outcome model is enriched by non-parametric components. We show that the inference results of the debiased estimator in [Farrell et al. \(2021a\)](#) apply to correlated samples, thus extending results from [Chernozhukov et al. \(2018\)](#), which deals with i.i.d. samples. While in our setting, sample correlation arises from the overlapped items in consideration sets that share treatment status, our inference results can apply to other settings with correlated samples, such as panel settings with repeated observations from the same customer. Since many empirical data can be correlated, this inference result can broaden the applicability of the doubly robust estimator, which has already seen wide adoption in marketing and business research (e.g., [Mummalaneni et al. 2022](#), [Ye et al. 2023b](#), [Kim et al. 2023](#), [Cheng et al. 2023](#), [Ellickson et al. 2024](#)).

2. Background on Recommender System Experiments

We start by providing background on recommender system experiments, which aim to assess whether a new recommender algorithm should be adopted. We focus on creator-side experiments, which are crucial to evaluating updates targeted at creators and measuring creator outcomes. We then discuss double-sided randomization as a conceptual “oracle” design and highlight its various limitations.

2.1. Recommender Systems and Algorithmic Treatment Effect

Recommender systems are central to digital platforms, delivering personalized content from libraries that may contain billions of items. Recommender systems rely on a multi-stage pipeline, typically retrieval, ranking, and re-ranking, that progressively narrows the number of items in the candidate set. Each stage applies machine learning models that score content based on predicted user engagement and platform objectives, with only a subset advancing to the next stage. In the final re-ranking stage, recommendations are further adjusted to account for additional factors such as advertising or promotional content and diversity. The re-ranking stage is typically the focus of algorithm updates and evaluation. In our short-video platform setting, the re-ranking stage selects a single item in

response to each viewer query, where users watch one video at a time in immersive mode and swipe to request the next.¹

Recommender systems are frequently updated, ranging from hyperparameter tuning to the deployment of new algorithms. The central managerial question is whether to adopt a new algorithm (*treatment*) compared to the status quo (*control*). For creator-side experiments, the focus of the paper, randomization is carried out at the creator level: each creator (and their content items²) is assigned to either treatment (scored by the new algorithm) or to control (scored by the status quo). The *treatment effect* of interest, which informs the adoption decision, is the difference in outcomes when the new algorithm is applied to all creators (*global treatment*) vs. when the status quo is applied to all creators (*global control*).

We formalize this objective using the potential outcomes framework (Imbens and Rubin 2015). Let \mathcal{C} denote the set of all content items. For each content item c , let $w_c \in \{0, 1\}$ denote its treatment status,³ and let \mathbf{w}_{-c} denote the treatment assignments of all other items. The full treatment assignment vector is \mathbf{w} . For each content item c , the potential outcome $r_c(\mathbf{w})$ is determined by a viewer query v that generates a response y (e.g., like or comment). We set $r_c(\mathbf{w}) = y$ if content c is exposed to viewer v , and $r_c(\mathbf{w}) = 0$ otherwise.⁴

The value of a policy π , which specifies a treatment assignment rule, given viewer population \mathbb{P}_v , is defined as the expected outcome across all content items (Johari et al. 2022, Goli et al. 2023):

$$Q(\pi; \mathbb{P}_v) := \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v, \mathbf{w} \sim \pi} [r_c(\mathbf{w})]. \quad (1)$$

Let $\pi_{\mathbf{1}}$ denote the global treatment policy where $(w_c, \mathbf{w}_{-c}) = \mathbf{1}$ and $\pi_{\mathbf{0}}$ denote the global control policy where $(w_c, \mathbf{w}_{-c}) = \mathbf{0}$. The platform’s adoption decision depends on whether the policy value under global treatment exceeds that under global control. We define the treatment effect as:

$$\tau := Q(\pi_{\mathbf{1}}; \mathbb{P}_v) - Q(\pi_{\mathbf{0}}; \mathbb{P}_v). \quad (2)$$

This estimand τ is the average treatment effect of the new algorithm, as it represents the expected impact of a global rollout across all creators, which is often referred to as the Total Average Treatment Effect (TATE) (Goli et al. 2023) or the Global Treatment Effect (GTE) (Johari et al. 2022).

¹ Our framework can be extended to slate recommendations via multiple-item choice models, where the core idea of interference applies. We focus on the single-item case, consistent with our empirical context.

² In practice, randomization occurs at the creator level, with content inheriting the same status. Throughout the paper, we use “creator” and “content items” interchangeably.

³ For notational simplicity, we restrict attention to a binary treatment; the framework extends naturally to multiple treatments.

⁴ We set 0 as the default outcome because in empirical applications outcomes typically measure engagement (e.g., watch time or interactions), which is naturally zero when the item is not exposed to viewers. Other default values could be adopted in different applications without affecting the analysis.

2.2. Double-Sided Randomization as an Oracle Design

A natural way to estimate the treatment effect τ is to design an experiment that directly mimics the global treatment and global control policies. One way to achieve this is a *blocked double-sided design* that partitions the market into two separate sub-universes by randomizing both creators and viewers. For example, one sub-universe would contain one group of viewers and creators using the treatment algorithm, and the other would contain the other creators and viewers using the control algorithm. By ensuring content from one sub-universe is excluded from the other, this design eliminates cross-group interference by construction. Comparing outcomes between these sub-universes therefore yields a direct, unbiased estimate of the treatment effect, making it a conceptual “oracle” for identification.

Despite being an interference-free benchmark, the double-sided design is very costly and rarely used for routine experiments for three main reasons. First is the statistical cost: platforms run hundreds of concurrent orthogonal experiments, and a double-sided design for each would require splitting the market into ever-smaller sub-universes. This causes sample sizes to decline exponentially, destroying the statistical power of all experiments. Second is the economic cost: partitioning the market leads to “market thinning”, which degrades the experience for both parties. Creators receive less exposure and viewers are exposed to potentially less relevant content. Third is the engineering cost: implementing and maintaining multiple parallel versions of the entire recommendation platform requires significant and ongoing operational resources, introducing complexity and system latency. Aside from these practical limitations, it is also unclear whether double-sided experiments are perfectly suited to recover the true global treatment effect. Because equilibrium outcomes may depend on the overall market size, estimates derived from these artificially partitioned small sub-universes may not generalize to the entire platform. For these reasons, platforms often rely on the more practical creator-side experiments.

3. Interference Bias from Creator Experiments

In this section, we start by describing creator-side experiments and the standard difference-in-means (DIM) estimators that compare treated and control outcomes. We then describe how interference arises and highlight two sources of bias: item exposure bias and viewer selection bias. We illustrate how these biases distort DIM estimators through both numerical examples and theoretical analysis.

3.1. Creator-Side Experiments and DIM Estimators

We start by describing the data structure from creator-side experiments. Each observation i corresponds to a viewer query in which one content item is exposed, represented by the tuple $(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i)$. V_i denotes the viewer characteristics. The consideration set $\vec{C}_i = \{C_{i,1}, \dots, C_{i,K}\}$ contains the set of K content items that reach the stage before randomization and are therefore

unaffected by treatment assignment.⁵ $\vec{W}_i = \{W_{i,1}, \dots, W_{i,K}\}$ records treatment assignment, with each $W_{i,k}$ drawn i.i.d. from a Bernoulli distribution with treatment probability q . The recommender exposes item k_i^* , and Y_i is the viewer’s response. We have a sample of n viewer queries, where each one is treated as independent.

Difference-in-means (DIM) estimators are commonly used to estimate treatment effects by comparing outcomes between treated and control items. We focus on two DIM estimators commonly used in the literature (Horvitz and Thompson 1952, Hájek 1971). Given n samples $(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i)_{i=1}^n$ from a creator-side experiment with treatment probability q :

$$\text{Horvitz-Thompson Estimator: } \hat{\tau}_n^{HT-DIM} := \frac{\sum_{i=1}^n W_{i,k_i^*} Y_i}{nq} - \frac{\sum_{i=1}^n (1 - W_{i,k_i^*}) Y_i}{n(1 - q)}, \quad (3)$$

$$\text{Hájek Estimator: } \hat{\tau}_n^{HA-DIM} := \frac{\sum_{i=1}^n W_{i,k_i^*} Y_i}{\sum_{i=1}^n W_{i,k_i^*}} - \frac{\sum_{i=1}^n (1 - W_{i,k_i^*}) Y_i}{\sum_{i=1}^n (1 - W_{i,k_i^*})}, \quad (4)$$

where Y_i is the outcome for the exposed item k_i^* , and W_{i,k_i^*} denotes its treatment status. The two estimators differ only in their normalization. The Horvitz–Thompson estimator normalizes by the assignment probability q and thus expected exposures, while the Hájek estimator normalizes by realized exposures. Next, we show why DIM estimators are biased under creator-side experiments.

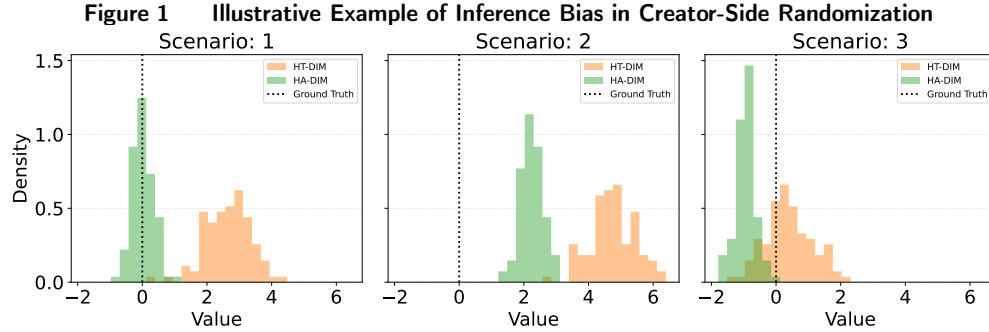
3.2. Interference and Sources of Bias

In creator-side experiments, interference arises because treated and control items compete for exposure within the same consideration set. This happens because viewers have access to content items from both groups. As a consequence, an item’s outcome depends not only on its own treatment status but also on the treatment status of other items in the set, violating the Stable Unit Treatment Value Assumption (SUTVA).

More specifically, interference leads to two forms of bias. First, suppose the treatment algorithm boosts the scores of treated items. Then these higher-scoring treated items are now more likely to be exposed, “crowding out” the control items. This leads to *item exposure bias*: the share of treated items actually exposed to viewers differs from the treatment assignment probability. Second, since recommendation algorithm is highly personalized, it might learn that the treatment is particularly effective for certain types of viewers (e.g., highly engaged or more valuable viewers). This leads to *viewer selection bias*: the compositions of viewers exposed to treated and control items differ systematically. As a result, the outcomes of treated and control items are not directly comparable because they are drawn from different viewer populations.

Such interference directly affects both DIM estimators. With treated and control items exposed to systematically different viewers, both the Horvitz–Thompson and Hájek estimators are subject

⁵ The model easily accommodates consideration sets of varying size; we assume a fixed K for notational convenience.



Notes: HT-DIM = Horvitz–Thompson DIM estimator; HA-DIM = Hájek DIM estimator.

Left figure: uplift independent of viewer type \rightarrow positive exposure bias only. Middle figure: uplift larger for high-baseline viewers \rightarrow positive exposure bias and positive viewer selection bias. Right figure: uplift larger for low-baseline viewers \rightarrow positive exposure bias and negative viewer selection bias.

to viewer selection bias. In addition, the Horvitz–Thompson estimator is affected by item exposure bias, since it normalizes by the assignment probability q rather than realized exposures. As we will show later, this does not imply that the Hájek estimator is less biased: the magnitude of bias depends on whether the two sources of bias act in the same or opposite directions.

We use a simple example to illustrate the two sources of interference bias. Suppose each item is assigned to the treatment or control with equal probability, and treated items receive a positive score uplift. Because treated and control items compete for exposure within the same set, treated items crowd out controls and are exposed more than 50% of the time, despite 50/50 treatment assignment. This gap between assignment and realized exposure generates item exposure bias. Not accounting for the actual exposure leads to an upward bias in the treatment effect.

Now consider a scenario where the treatment uplift varies across viewers. Viewers differ in their baseline tendency to engage with promotional content (e.g., higher income viewers are more likely to purchase from ads). Consider viewers with high vs. low baseline tendencies. When uplift is larger for high-baseline viewers (e.g., the algorithm prioritizes conversions), treated items are more likely to be exposed to high-baseline viewers relative to control items. Not accounting for the viewer selection bias leads to an upward bias of the treatment effect. The opposite happens when treatment uplift is larger for low-baseline viewers. Treated items are more likely to reach viewers with lower baseline tendency, and viewer selection bias leads to a downward bias in the treatment effect.

Figure 1 summarizes these patterns using a numerical example. In the left figure, when the uplift does not depend on viewer types, the Hájek estimator recovers the true effect by accounting for realized exposures, while the Horvitz–Thompson estimator overestimates due to item exposure bias. In the middle figure, when the uplift is larger for high-baseline viewers, both biases lead to upward-biased estimates. The Horvitz–Thompson estimator shows a larger bias since it reflects both item exposure and viewer selection effects, while Hájek estimator reflects only viewer selection bias. In

the right figure, when the uplift is higher for low-baseline viewers, viewer selection bias and item exposure bias act in opposite directions. Hájek, affected only by viewer selection bias, underestimates the effect, while Horvitz–Thompson, affected by both, ends up slightly overestimating the effect.

3.3. Theoretical Analysis of DIM Estimators

We now formalize the intuition developed above by characterizing the asymptotic limits of DIM estimators under creator-side experiments. We show how item exposure bias and viewer selection bias affect these estimators in large samples.

Let $e^*(q | v)$ be the probability that viewer v is exposed to a treated item when the treatment assignment probability is q . The population-level probability of exposure is then $e^*(q) = \mathbb{E}_{v \sim \mathbb{P}_v} [e^*(q | v)]$, where \mathbb{P}_v is the distribution of viewers. Let $\mathbb{P}_v^{1,q}$ and $\mathbb{P}_v^{0,q}$ denote the distributions of viewers exposed to treated and control items, respectively. Item exposure bias occurs whenever $e^*(q) \neq q$. Viewer selection bias arises whenever $\mathbb{P}_v^{1,q} \neq \mathbb{P}_v^{0,q}$, indicating that viewers exposed to treated vs. control differ systematically from the overall population.

We assume that no single item appears in an excessively large fraction of consideration sets (formal condition in Appendix A). This assumption is plausible for creator-side experiments that tend to focus on promotional content, because advertising videos are subject to ad budget and traffic allocation constraints that naturally prevent any single item from dominating viewer queries.

THEOREM 1 (Bias of DIM Estimators). *Suppose Assumption 1 holds. Under creator-side randomization with treatment probability $q \in (0, 1)$, the Horvitz-Thompson estimator $\hat{\tau}_n^{HT-DIM}$ converges in probability to τ^{HT} , where*

$$\tau^{HT} := \frac{e^*(q)}{q} \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v^{1,q}, \mathbf{w} \sim \mathcal{B}(q)} [r_c | w_{k^*} = 1] - \frac{1 - e^*(q)}{1 - q} \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v^{0,q}, \mathbf{w} \sim \mathcal{B}(q)} [r_c | w_{k^*} = 0], \quad (5)$$

with w_{k^*} denoting the treatment status of exposed item and $\mathcal{B}(q)$ denoting the Bernoulli trial with assignment probability q . The Hájek estimator $\hat{\tau}_n^{HA-DIM}$ converges in probability to τ^{HA} , where

$$\tau^{HA} := \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v^{1,q}, \mathbf{w} \sim \mathcal{B}(q)} [r_c | w_{k^*} = 1] - \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v^{0,q}, \mathbf{w} \sim \mathcal{B}(q)} [r_c | w_{k^*} = 0]. \quad (6)$$

The proof is given in Appendix A. For comparison, the true treatment effect is

$$\tau = \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v, \mathbf{w} = \mathbf{1}} [r_c | w_{k^*} = 1] - \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v, \mathbf{w} = \mathbf{0}} [r_c | w_{k^*} = 0]. \quad (7)$$

This result formalizes the two sources of bias discussed previously. Relative to the true treatment effect, the bias in the Hájek estimator is driven by the difference between the viewer populations ($\mathbb{P}_v^{1,q} \neq \mathbb{P}_v^{0,q}$), which is precisely viewer selection bias. The Horvitz-Thompson estimator is subject to this same bias, and is further affected by the deviation of the realized exposure rate from the assignment probability ($e^*(q) \neq q$), which is item exposure bias.

4. Modeling Interference

In this section, we introduce a framework to model the interference pathway in creator-side experiments. A recommender choice model captures how treated and control items compete for exposure, while a viewer response model describes outcomes once exposure occurs. Combining the two enables counterfactual analysis of alternative treatment policies to estimate treatment effects. We conclude by discussing the scope and robustness of the modeling approach.

4.1. Recommender Choice Model

Upon receiving a viewer query, the recommender system chooses one content item for viewer V_i from the consideration set \vec{C}_i . The evaluation of an item $C_{i,k}$ depends on its treatment status $W_{i,k}$, where $W_{i,k} = 1$ indicates the treatment algorithm and $W_{i,k} = 0$ the control algorithm.

We approximate the recommender system with a choice model built around a latent score $S_{i,k}$ of each item $C_{i,k}$ for viewer V_i . Intuitively, this score summarizes several stages of the recommender pipeline into a single measure that depends on viewer and content characteristics and the treatment status. The score functions are estimated to rationalize the observed exposure of the recommender system. Specifically, let the latent score $S_{i,k}$ takes the form:

$$S_{i,k} = s_0(V_i, C_{i,k}) + W_{i,k} \cdot s_1(V_i, C_{i,k}) + \epsilon_{i,k}. \quad (8)$$

The score function $s_0(\cdot, \cdot)$ represents baseline evaluation under the control algorithm, while $s_1(\cdot, \cdot)$ captures the treatment uplift. Both functions are parameterized by neural networks to flexibly approximate personalized recommendations across diverse viewer–content pairs (Covington et al. 2016). The error term $\epsilon_{i,k}$ accounts for the randomness in recommender systems and the approximation errors of neural networks in representing the actual pipeline. We assume that the error term $\epsilon_{i,k}$ follows an i.i.d. Type 1 Extreme Value distribution, and the item with the highest score or “utility” in the consideration set is exposed to viewer V_i .

Under this formulation, the probability that item k is exposed to viewer V_i given consideration set \vec{C}_i and treatment allocation \vec{W}_i follows a multinomial logit form:

$$p_k(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1) := \mathbb{P}(k_i^* = k) = \frac{e^{s_0(V_i, C_{i,k}) + W_{i,k} \cdot s_1(V_i, C_{i,k})}}{\sum_{k'=1}^K e^{s_0(V_i, C_{i,k'}) + W_{i,k'} \cdot s_1(V_i, C_{i,k'})}}. \quad (9)$$

The choice model explicitly captures interference in the recommendation process. An item’s exposure probability depends not only on its own treatment status but also on the treatment assignment of the other items in the set.

The recommender choice model is semi-parametric: it combines a structured logit choice model with flexible neural networks. The structural component enables counterfactual analysis under alternative treatment policies, which is difficult for fully nonparametric black-box models that often

fail to generalize to counterfactual settings. The neural networks capture complex viewer–content relationships, allowing the model to approximate the behavior of real-world recommender systems. They also help relax the restrictive independence of irrelevant alternatives (IIA) property of standard logit models. Instead of proportional substitution across items, the neural networks allow a new item to disproportionately affect the exposure probabilities of items depending on content similarity and viewer-specific tastes.

Beyond enabling treatment effect estimation, the recommender choice model can be practically valuable for platforms. By approximating the complex pipeline with a distilled model (Hinton et al. 2015), it can reduce latency when serving viewers online during peak traffic (Tang and Wang 2018). Moreover, the learned score functions summarize the overall value of items to viewers, which can serve as a useful metric for offline evaluation.

4.2. Viewer Response Model

To evaluate counterfactual treatment policies, we also need to model how viewers respond once an item is exposed. Outcomes of interest may include view time, likes, or conversions. We model the response Y_i of viewer V_i to exposed item C_{i,k_i^*} as

$$Y_i = z(V_i, C_{i,k_i^*}) + \zeta_i, \quad (10)$$

where $z(\cdot, \cdot)$ is a flexible function, parameterized by a neural net, and ζ_i is i.i.d. noise. Separate response models can be trained for different outcomes.

This specification relies on two simplifying assumptions. First, conditional on the exposed item, outcomes do not depend on the item’s treatment status. This is reasonable since viewers are typically not aware of changes in recommendation algorithms. Second, it assumes that outcomes do not depend on prior content exposure. The viewer response model can be extended to allow previous content interactions, but doing so will introduce considerable complexity.

In practice, platforms often maintain pre-trained response predictors, which we leverage in our empirical setting. It’s worth noting that our framework does not require these predictions to be unbiased. As detailed in Section 5.2, we introduce a debiased estimator to correct for potential errors in the nuisance predictions. As we will discuss in more details later, the debiasing step is made possible by the treatment randomization in creator-side experiments.

4.3. Counterfactual Analysis for Treatment Effect

The recommender choice model in Equation (9) and the viewer response model in Equation (10) together allow us to evaluate counterfactual treatment assignment policies. For any policy π , the expected outcome, or policy value, is (see Appendix B.1 for detailed derivation):

$$Q(\pi) = \mathbb{E}_{(V_i, \vec{C}_i, \vec{W}_i \sim \pi)} \left[\sum_{k=1}^K z(V_i, C_{i,k}) \cdot p_k \left(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1 \right) \right]. \quad (11)$$

The inner term calculates the sum of each item’s exposure probability multiplied by its expected response. The outer expectation is taken over the distribution of the viewer, consideration set, and treatment assignment.

The treatment effect can be estimated by comparing the policy values of global treatment $Q(\pi_1)$ and global control $Q(\pi_0)$:

$$\tau = Q(\pi_1) - Q(\pi_0) = \mathbb{E}_{(V_i, \vec{C}_i)} \left[\sum_{k=1}^K z(V_i, C_{i,k}) \cdot \left(p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1) - p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{0}; s_0, s_1) \right) \right]. \quad (12)$$

Here the difference in exposure probabilities quantifies how the likelihood of showing each content item $C_{i,k}$ changes when going from global control to global treatment.

Finally, the framework is flexible to extend beyond the average treatment effect by comparing global treatment versus control. It can be applied to estimate heterogeneous effects on different types of creators. For a subgroup \mathcal{C}_0 , the corresponding heterogeneous treatment effect can be obtained by summing over items in the subgroup of interest for $C_{i,k} \in \mathcal{C}_0$. We can evaluate policies that allocate a fraction of traffic to the treatment algorithm (e.g., as the platform scales up the proportion of treatment algorithm). One can also compute the treatment effect in relative terms with $\frac{Q(\pi_1) - Q(\pi_0)}{Q(\pi_0)}$.

4.4. Scope and Assumptions of the Proposed Approach

Our framework is well suited for creator experiments occurring in the later, re-ranking stage of a recommender system. At this stage, items compete for exposure within a relatively small consideration set, which allows for tractable choice model estimation. The approach is less suited for early retrieval stages that involve very large candidate pools where the large number of options makes the choice model difficult to estimate reliably. Furthermore, our analysis focuses on short-run outcomes, taking content and queries as given. We abstract away from other potential forms of interference, such as viewer-side temporal dynamics (Farias et al. 2023) or long-run content supply effects. Capturing these important but distinct dynamics would require different modeling approaches and are left as avenues for future work.

The framework’s choice model relies on two parametric assumptions. First, we assume the choice probabilities follow a multinomial logit form. This is a standard workhorse model in the discrete choice literature, arising from the assumption of Gumbel-distributed error terms on the each item’s latent utility. While parametric, the resulting softmax function is highly flexible and can approximate any categorical distribution (Cervera et al. 2021, Ye et al. 2023b), making it a reasonable choice given the stochastic nature of online recommendations.

Second, we assume that the score of each item has an additive structure that depends only on its own characteristics and treatment status: $\bar{S}_{i,k} = s_0(V_i, C_{i,k}) + W_{i,k} \cdot s_1(V_i, C_{i,k})$. This additive

separability allows us to isolate the treatment uplift s_1 as a distinct function. A fully general model, where an item’s score depends on the treatment status of all other items, would be computationally infeasible, with the complexity scaling exponentially with the size of the consideration set. Crucially, this does not imply a homogeneous treatment effect. Because we parameterize both the baseline score s_0 and the uplift s_1 with flexible neural networks, our model can capture complex, non-linear treatment effect heterogeneity across different viewers and content items.

From a practical standpoint, our framework is designed for the platform experimentation team. Its implementation requires access to the standard data logged during a creator-side experiment: the consideration set, the exposed item, the treatment assignment, and viewers and content features. A key premise of our approach is a simple offline simulation, even with the model’s code, is fundamentally unreliable for answering the business question. An offline simulation fails to replicate a live recommendation system’s decisions because of factors missing from offline data, such as the inherent stochasticity from live exploration (e.g., testing new items) and concurrent experiments, and transient, real-time features of items, viewers, and system latency. Furthermore, evaluating recommender systems requires outcome variables from real viewer interactions. Even if we are willing to use the imperfect simulated exposure and predicted responses, these proxies will likely introduce errors and biases, making it impossible to perform valid inference for treatment effect estimation. These limitations are precisely why platforms invest in costly live experiments.

A potential concern lies in possible errors in representing the recommender choice model and the viewer response model. These models may be imperfect due to factors such as misspecification, omitted features, or neural network approximation errors. As we explain in Section 5.2, our framework is designed to be robust to such imperfections. The debiasing procedure exploits the randomness in treatment assignment to correct first-order bias arising from imperfectly estimated nuisance functions, thereby preserving the validity of the treatment effect estimate. However, this robustness depends on the nuisance models being “good enough” approximations. If they are severely misspecified, they may fail to achieve the convergence rate required for the debiasing theory to hold.

5. Estimation and Inference Procedure

In this section, we describe the estimation and inference procedure of the proposed model. We start with the estimation of the recommender choice and viewer response models. We then construct a debiased estimator for the average treatment effect that remains valid even when the nuisances converge at slower-than- \sqrt{n} rates.

5.1. Model Estimation

We start by estimating the nuisance components that characterize the data-generating process as described in Section 4. These nuisance components include the recommender score components (s_0 ,

s_1), which capture how the platform assigns exposure probabilities to content items under control and treatment, and viewer response z , which captures how viewers react once an item is exposed. We use deep feedforward neural networks with ReLU activation for their approximation guarantees (Farrell et al. 2021a,b). Below we describe the estimation procedure, the conditions required for identification, and the convergence guarantees.

The recommender score components (s_0, s_1) are estimated by minimizing cross-entropy loss, which is equivalent to maximizing the multinomial logit likelihood. Intuitively, this procedure learns the score functions so that the choice model best replicates which item the recommender actually exposes. For a viewer query i , we observe viewer V_i with consideration set \vec{C}_i and treatment status \vec{W}_i with realized exposure k_i^* . The loss is:

$$\ell_1(V_i, \vec{C}_i, \vec{W}_i, k_i^*; \tilde{s}_0, \tilde{s}_1) := -\log \left(p_{k_i^*} \left(V_i, \vec{C}_i, \vec{W}_i; \tilde{s}_0, \tilde{s}_1 \right) \right) \quad (13)$$

We estimate the nuisance components $\hat{s}_0, \hat{s}_1 \in \mathcal{F}_{DNN}$ by minimizing the empirical loss:

$$(\hat{s}_0, \hat{s}_1) \in \arg \min_{\tilde{s}_0, \tilde{s}_1 \in \mathcal{F}_{DNN}} \frac{1}{n} \sum_{i=1}^n \ell_1(\cdot). \quad (14)$$

For the viewer response model z , we estimate the mapping from exposed viewer–item pairs to outcomes. When viewer V_i gets exposed to content item C_{i,k_i^*} , we observe the outcome Y_i . Estimating the viewer response function is a standard machine learning task. For continuous outcomes, we use the mean square error loss for any candidate function \tilde{z} :

$$\ell_2(V_i, C_{i,k_i^*}, Y_i; \tilde{z}) = \left(\tilde{z}(V_i, C_{i,k_i^*}) - Y_i \right)^2. \quad (15)$$

For categorical outcomes, one can use the cross entropy loss. With either type of loss function, we get the estimated $\hat{z} \in \mathcal{F}_{DNN}$ by minimizing the loss:

$$\hat{z} \in \arg \min_{\tilde{z} \in \mathcal{F}_{DNN}} \frac{1}{n} \sum_{i=1}^n \ell_2(\cdot). \quad (16)$$

To ensure that these nuisance functions are well defined, we impose a boundedness condition (Assumption 2 in Appendix C.1), which ensures that each item in the consideration set has a positive probability of being exposed under both treatment and control. Under this assumption, the nuisance functions are identified, as formalized in the following proposition and proved in Appendix C.2.

PROPOSITION 1 (Identification). *Suppose the recommender choice behavior follows the semi-parametric form in Equation (9), and the viewer response model follows the nonparametric form in Equation (10). Under Assumption 2, the nuisance functions (s_0, s_1, z) can be nonparametrically identified, up to a location normalization of the baseline score $s_0(V_i, C_{i,1}) \equiv 0$.*

5.2. Debiased Estimator

Having estimated the nuisance components, we next construct an estimator of the treatment effect. We begin by describing a direct plug-in approach and then introduce our debiased (DB) estimator. Because the nuisance components estimated via neural networks generally converge at rates slower than \sqrt{n} , the direct plug-in approach fails to deliver valid inference. The debiased estimator, by contrast, corrects for the bias introduced by nuisance estimation errors through Neyman orthogonality. We describe the asymptotic results in Section 5.3.

We start with a straightforward plug-in estimator. With the estimated recommender choice and viewer response models, one can directly plug in the estimated components into the ATE formulation in Equation (12). For each viewer-consideration set pair (V_i, \vec{C}_i) , the direct plug-in estimate μ is:

$$\mu(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z}) = \sum_{k=1}^K \hat{z}(V_i, C_{i,k}) [p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; \hat{s}_0, \hat{s}_1) - p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{0}; \hat{s}_0, \hat{s}_1)] \quad (17)$$

Averaging $\mu(\cdot)$ across the sample produces a plug-in estimator of the average treatment effect. If the components (s_0, s_1, z) are fully parametric, their estimates can achieve \sqrt{n} -consistency and we can directly rely on the plug-in estimator for treatment effect estimation. However, with these components being approximated by neural networks, the direct plug-in is generally not \sqrt{n} -consistent and often converges at slower rates, resulting in biased inference.

To overcome the limitation of the direct plug-in estimator, we propose a debiased estimator that is \sqrt{n} -consistent and asymptotically normal. The key idea is to adjust the plug-in estimator with a correction term that removes first-order bias from nuisance estimation. Following Farrell et al. (2021a), for each observation, the debiased estimator ψ is:

$$\psi_i^{DB} = \mu(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z}) - \underbrace{\nabla \mu(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z})^T H(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z})^{-1} \nabla \ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; \hat{s}_0, \hat{s}_1, \hat{z})}_{\text{debiasing term}}. \quad (18)$$

The first term is the plug-in estimate based on the fitted nuisance functions, and the second term removes the first-order bias introduced by imperfect nuisance estimation. $\nabla \mu$ is the gradient of the plug-in estimator μ with respect to the nuisance components. $\nabla \ell$ is the gradient of the total loss ℓ , which combines ℓ_1 from the recommender choice model and ℓ_2 from the viewer response model, with respect to the nuisance functions. H is the expected Hessian of the loss ℓ with respect to the nuisances, with the expectation taken over treatment assignment \vec{W}_i under the creator-side randomization. Explicit expressions and derivations are given in Appendix B.2 and Appendix B.3.

Intuitively, the gradient $\nabla \mu$ captures the sensitivity of the treatment effect estimate to small perturbations in nuisance components. The gradient $\nabla \ell$ reflects how nuisance estimation errors

affect the total loss, which is scaled by the curvature of the loss function H^{-1} .⁶ Together, these terms form the correction $\nabla\mu^\top H^{-1}\nabla\ell$, which offsets the first-order bias in the plug-in estimator and ensures that the final estimate is locally robust to nuisance estimation errors.

We next describe how the components of the debiasing term can be computed in practice. Both $\nabla\mu$ and $\nabla\ell$ can be explicitly written down and calculated (see Appendix B.3 for details). When an explicit form is not available, these derivatives can be obtained through numeric differentiation. The expected Hessian H can be explicitly computed in our case, unlike the general setup of Farrell et al. (2021a), because the platform’s treatment assignment mechanism is known. But the computational complexity of explicit calculation scales exponentially with the size of the consideration set. When the size of the consideration set is manageable, as in the simulation studies in Section 6, we compute the exact expectation. When the consideration set is large, as in our empirical study in Section 7, we approximate the expected Hessian via Monte Carlo simulations. Specifically, we sample the treatment assignment vector \vec{W}_i from the creator-side randomization and take the empirical mean of the resulting Hessians to estimate the expectation.

The overall debiased (DB) estimator $\hat{\tau}_n^{DB}$ takes the average of ψ_i^{DB} across all observations:

$$\hat{\tau}_n^{DB} := \frac{1}{n} \sum_{i=1}^n \psi_i^{DB}. \quad (19)$$

The standard error of the DB estimator $\hat{\tau}_n^{DB}$ is estimated as $n^{-1/2}(\hat{V}_n^{DB})^{1/2}$, where the estimated variance \hat{V}_n^{DB} is computed as:

$$\hat{V}_n^{DB} = \frac{1}{n} \sum_{i=1}^n (\psi_i^{DB} - \hat{\tau}_n^{DB})^2. \quad (20)$$

Following Chernozhukov et al. (2018) and Farrell et al. (2021a), we employ sample splitting and cross-fitting to estimate the nuisance components. Specifically, the data are partitioned into K folds. For each fold k , the nuisance functions $(\hat{s}_0^{(-k)}, \hat{s}_1^{(-k)}, \hat{z}^{(-k)}, \hat{H}^{(-k)})$ are estimated using all observations except those in fold k . The treatment effect is then estimated on the held-out fold to obtain $\hat{\tau}_n^{(k)}$. The cross-fitted treatment effect estimator is the average across K folds: $\hat{\tau}_n^{DB} = \frac{1}{K} \sum_{k=1}^K \hat{\tau}_n^{(k)}$, with estimated variance $\hat{V}_n = \frac{1}{K} \sum_{k=1}^K \hat{V}_n^{(k)}$.

5.3. Asymptotic Results under Correlated Samples

We now turn to the asymptotic properties of the debiased estimator. We start with the convergence rates of the estimated nuisance components, then establish the Neyman orthogonality of the debiased estimator. We derive the \sqrt{n} -consistency and asymptotic normality, which together guarantee valid

⁶ The validity of the debiasing step relies on the randomization of treatment assignment in creator-side experiments, which guarantees invertibility of the Hessian H under Assumption 2 (bounded scores). Formal results and proofs on Hessian invertibility are provided in Appendix D.1.

inference of the debiased estimator. A key contribution of our theoretical analysis is to extend double machine learning asymptotics to correlated data.

We start with the convergence rates of the nuisance components estimated with neural networks. Leveraging recent results on the approximation power of deep neural networks (Farrell et al. 2021a,b), we provide convergence guarantees. The proof is provided in Appendix C.3.

PROPOSITION 2 (Convergence). *Suppose the recommender choice follows the semi-parametric form in Equation (9), and the viewer response model follows the nonparametric form in Equation (10) and is uniformly bounded. Suppose Assumption 2 and Assumption 3 (in Appendix C) hold. Let p be the smoothness parameter of the true nuisance functions (s_0, s_1, z) and d the dimension of the viewer-content covariates. If the neural networks $\hat{s}_0, \hat{s}_1, \hat{z}$ have width $J \asymp O(n^{d/2(p+d)} \log^2 n)$ and depth $L \asymp \log n$, then for sufficiently large n , with probability $1 - \exp(-n^{d/(p+d)} \log^8 n)$,*

$$\|\hat{s}_0 - s_0\|_{L_2} + \|\hat{s}_1 - s_1\|_{L_2} + \|\hat{z} - z\|_{L_2} = O\left(n^{-\frac{p}{p+d}} \log^8 n + \frac{\log \log n}{n}\right).$$

Proposition 2 characterizes the convergence rate of the nuisance components. For typical levels of smoothness p and input dimension d , the convergence rate $n^{-\frac{p}{p+d}}$ lies between the parametric rate $n^{-\frac{1}{2}}$ and the slower nonparametric rate $n^{-\frac{1}{4}}$. This result is useful for two reasons. First, the direct plug-in estimator would inherit the slower convergence and therefore fail to achieve \sqrt{n} consistency. Second, the debiased estimator is constructed to be orthogonal to first-order nuisance estimation error, so the established convergence rate ensures that it retains \sqrt{n} consistency and valid asymptotic inference even when the nuisance components converge at a slower rate.

We next establish that the debiased estimator ψ satisfies universal Neyman orthogonality (Chernozhukov et al. 2019, Foster and Syrgkanis 2023). This property ensures that small imperfections in the nuisance estimates have only a second-order impact on the debiased estimate. Neyman orthogonality is a key property for establishing the asymptotic normality of the debiased estimator. The proof is given in Appendix D.2.

PROPOSITION 3 (Universal Orthogonality). *The debiased estimator ψ , defined in Equation (18), is universally orthogonal with respect to the nuisance components. Specifically, for any nuisance components $(\tilde{s}_0, \tilde{s}_1, \tilde{z}, \tilde{H})$,*

$$\mathbb{E}[\nabla\psi(V, \vec{C}, \vec{W}, k^*, Y; \tilde{s}_0 = s_0, \tilde{s}_1 = s_1, \tilde{z} = z, \tilde{H} = H) \mid V, \vec{C}] = 0,$$

where $(V, \vec{C}, \vec{W}, k^*, Y)$ is drawn from creator-side experiments, and $\nabla\psi$ is the gradient with respect to the nuisances.

Leveraging the Neyman orthogonality property, we now show our main asymptotic results. The debiased estimator achieves \sqrt{n} -consistency and asymptotic normality even when the nuisance components converge at rates slower than $n^{-1/2}$. This central limit theorem thereby enables valid statistical inference and confidence interval construction for the treatment effect. The formal result is stated in Theorem 2 below.

THEOREM 2 (Asymptotic Normality of the Debiased Estimator). *Suppose Assumptions 1 & 2 hold. Assume that the data-generating process follows the recommender choice model in Equation (9) and the viewer response model in Equation (10). Suppose that the estimated nuisance functions are all bounded by the constant C in Assumption 2 and satisfy the convergence rate: $\|\hat{s}_0 - s_0\|_{L_2} + \|\hat{s}_1 - s_1\|_{L_2} + \|\hat{z} - z\|_{L_2} = o(n^{-1/4})$. Then the debiased estimator $\hat{\tau}_n^{DB}$ defined in Equation (19) is \sqrt{n} -consistent and asymptotically normal with:*

$$n^{1/2} (\hat{\tau}_n^{DB} - \tau) / (\widehat{V}_n^{DB})^{1/2} \Rightarrow \mathcal{N}(0, 1),$$

implying that $\hat{\tau}_n^{DB} - \tau = O_p(n^{-1/2})$.

Our main theoretical contribution lies in extending the asymptotic results of the debiased estimator to correlated samples. This extension is essential in our empirical application, where sample correlation arises from overlapping consideration sets. Even when each viewer query i is treated as independent, samples are correlated because once an item appears in a consideration set, its treatment status in subsequent overlapping consideration sets becomes deterministic. Only items unique to a given consideration set follow independent Bernoulli randomization.

To account for sample correlation, we model the data-generating process sequentially and employ martingale limit theorems to characterize the asymptotic behavior of the debiased estimator. The proof proceeds by comparing the empirical debiased estimator $\hat{\tau}_n^{DB}$ with an oracle counterpart $\tilde{\tau}_n^{DB}$ that assumes the nuisance components are known without error. The oracle estimator forms a martingale difference sequence and hence converges to a normal distribution. By Neyman orthogonality, the discrepancy between the empirical and oracle estimators is asymptotically negligible, and the bound on their difference is shown to be sufficiently tight to establish the desired limit distribution. We provide detailed proof in Appendix E.

Sample correlation can also impact the cross-fitting procedure. Overlapping consideration sets induce dependence between observations, violating the i.i.d. assumption underlying standard cross-fitting procedure. To address this issue, one could apply the network cross-fitting of Viviano (2025), that extends cross-fitting to dependent or networked samples. In practice, we adopt the standard cross-fitting procedure for estimating nuisance components, which is much easier to implement and does not materially affect inference performance in our empirical application.

6. Monte Carlo Simulation

In this section, we use Monte Carlo simulations to evaluate the performance of the proposed debiased estimator. We compare it against several benchmark estimators and show that, unlike the benchmarks, the debiased estimator delivers unbiased estimates and valid inference.

We simulate data for $n = 3000$ viewers. Each viewer has a two-dimensional feature $V_i \sim \mathcal{U}(0, 1)^2$ and a consideration set of $K = 5$ items drawn from a pool of 500 content items. Each item C_k has a two-dimensional continuous feature $C_{1,k} \sim \mathcal{U}(0, 1)^2$ and a binary feature $C_{2,k} \sim \text{Bernoulli}(0.5)$. Let the baseline score be $s_0(V_i, C_k) = (V_i + C_{1,k}) + 0.05(V_i + C_{1,k})^2$, and the score uplift by the treatment algorithm be $s_1(V_i, C_k) = \delta \cdot C_{2,k} \cdot (V_i + C_{1,k})$. One item k_i^* is exposed according to the choice model in Equation (9). The viewer outcome model is $Y_i = (V_i + C_{1,k_i^*}) + \zeta_i$, where $\zeta_i \sim \mathcal{N}(0, 0.1)$.

This setup generates the two forms of interference bias. Consider a case when $\delta > 0$. Item exposure bias arises because treated items receive higher scores on average and therefore are more likely to be exposed than the treatment assignment probability. Moreover, viewer selection bias emerges because the treatment uplift $s_1(V_i, C_{i,k})$ depends on viewer features V_i . When $\delta > 0$, viewers with higher baseline engagement (larger V_i) are systematically more likely to see treated items, leading to differences in the viewer populations exposed to treatment versus control.

6.1. Benchmark Estimators

Besides the proposed debiased (DB) estimator $\hat{\tau}_n^{DB}$, we evaluate six other benchmark estimators.

Difference-in-means estimators. As described in Section 3.3, we include two difference-in-means (DIM) estimators commonly used in practice: the *Horvitz–Thompson DIM* estimator $\hat{\tau}_n^{HT-DIM}$, and the *Hájek DIM* estimator $\hat{\tau}_n^{HA-DIM}$. Both estimators compare average outcomes between treated and control items but differ in their normalization.

Regression-based estimators. We next consider regression-based approaches, which extend DIM estimators by adjusting for covariates of the exposed viewer–item pair. We estimate

$$Y_i = \alpha + \tau W_{i,k_i^*} + \beta^\top X_{i,k_i^*}, \quad (21)$$

where Y_i is the observed outcome, W_{i,k_i^*} indicates whether the exposed item is treated, and X_{i,k_i^*} represents covariates for the exposed viewer–item pair. Ordinary least squares estimation yields the *regression-based (RB)* estimator $\hat{\tau}_n^{RB}$, which is the coefficient on W_{i,k_i^*} . We also consider a weighted variant that applies inverse exposure probabilities, $1/p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; \hat{s}_0, \hat{s}_1)$, as observation weights. These weights adjust for unequal exposure likelihoods induced by the recommender system. This weighted OLS produces the *weighted regression-based (W-RB)* estimator $\hat{\tau}_n^{W-RB}$.

Propensity-based estimators. Finally, we consider estimators based on importance weighting (Hahn 1998). The idea is to reweight each observation by the propensity under a target environment

(e.g., global treatment) relative to the experimental environment. A naïve implementation would use the full treatment assignment vector \vec{W}_i , yielding the inverse propensity $1\{\vec{W}_i = \mathbf{1}\}/q^K$, where q is the treatment probability and K the size of the consideration set. The naïve implementation is infeasible because this weight grows exponentially in K , resulting in very high variance.

To address this issue, we leverage the structure of the recommender choice model and express the propensity in terms of exposure probabilities rather than the full treatment assignment vector. Let $f^{(E)}$, $f^{(T)}$, and $f^{(C)}$ denote the probability densities (or mass functions) under the experimental, global treatment, and global control environments, respectively. The importance weight for global treatment can be written as

$$\frac{f^{(T)}(Y_i)}{f^{(E)}(Y_i)} = \frac{f^{(T)}(k_i^* | V_i, \vec{C}_i)}{f^{(E)}(k_i^* | V_i, \vec{C}_i)}, \quad (22)$$

which represents the relative likelihood of an observation under global treatment versus the experimental environment. The equality holds because the outcome Y_i depends only on the exposed item and viewer features, while the distribution of viewers and their consideration sets, which are determined before the experiment, remains identical across environments. Hence, differences arise solely from the exposure mechanism.

We estimate this ratio using the exposure probabilities from the fitted choice model:

$$r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) = \frac{p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; \hat{s}_0, \hat{s}_1)}{\mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)}[p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; \hat{s}_0, \hat{s}_1)]},$$

and estimate $r^{(C)}(\cdot)$ analogously with $\vec{W}_i = \mathbf{0}$.

Using these weights, we construct the *Generalized Inverse Propensity Weighted* (GIPW) estimator:

$$\hat{\tau}_n^{GIPW} = \frac{1}{n} \sum_{i=1}^n [r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) - r^{(C)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1)] Y_i. \quad (23)$$

To account for potential prediction error in the outcome model, we extend GIPW by adding an augmentation term, yielding the *Generalized Augmented Inverse Propensity Weighted* (GAIPW) estimator $\hat{\tau}_n^{GAIPW}$:

$$\hat{\tau}_n^{GAIPW} = \frac{1}{n} \sum_{i=1}^n \left\{ [\mu^{(T)}(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z}) - \mu^{(C)}(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z})] - [r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) - r^{(C)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1)] [Y_i - \hat{z}(V_i, C_{i,k_i^*})] \right\}. \quad (24)$$

Here, $\mu^{(T)}$ and $\mu^{(C)}$ are plug-in estimates of expected outcomes under global treatment and control, obtained using the fitted nuisance components $(\hat{s}_0, \hat{s}_1, \hat{z})$. This formulation combines reweighting by exposure probabilities with an outcome model correction term.

Similar to the direct plug-in estimator, both GIPW and GAIPW estimators are consistent when the nuisance components are correctly estimated. However, because these estimators lack full Neyman orthogonality, when nuisance components converge more slowly with flexible neural networks, first-order estimation errors can introduce substantial bias and lead to invalid inference. The GAIPW estimator includes an augmentation term that partially mitigates this sensitivity, achieving partial orthogonality. These properties are summarized below.

THEOREM 3. *Assume that the data-generating process follows the recommender choice model in Equation (9), and the viewer response model in Equation (10). Suppose Assumptions 1 and 2 hold, and that the nuisance estimates are consistent: $\|\hat{s}_0 - s_0\|_{L_2} + \|\hat{s}_1 - s_1\|_{L_2} + \|\hat{z} - z\|_{L_2} = o(1)$. Then both the GIPW estimator $\hat{\tau}_n^{GIPW}$ and the GAIPW estimator $\hat{\tau}_n^{GAIPW}$ are consistent, $\hat{\tau}_n^{GIPW} - \tau = o_p(1)$ and $\hat{\tau}_n^{GAIPW} - \tau = o_p(1)$.*

Moreover, GAIPW is Neyman orthogonal with respect to the nuisance functions \hat{z} and $\mathbb{E}_{\vec{W} \sim \mathcal{B}(q)}[p_k(\cdot, \cdot, \vec{W}; \hat{s}_0, \hat{s}_1)]$, but not with respect to $p_k(\cdot, \cdot, \vec{W} = \mathbf{1}; \hat{s}_0, \hat{s}_1)$ or $p_k(\cdot, \cdot, \vec{W} = \mathbf{0}; \hat{s}_0, \hat{s}_1)$.

The proof of Theorem 3 is provided in Appendix F. Here we outline the key intuition behind the partial orthogonality result. In standard AIPW, the outcome model and the treatment propensity score are separate nuisance functions, and the augmentation term corrects errors in either. In our setting, however, the outcome model itself depends on the same exposure probabilities that govern the weighting mechanism. Specifically, the expected outcome under treatment is

$$\mu^{(T)}(V_i, \vec{C}_i) = \sum_{k=1}^K z(V_i, C_{i,k}) p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}),$$

where $p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1})$ is the exposure probability of item k under global treatment. Since the probability enters both the outcome and the weighting terms, its estimation errors affect GAIPW through both channels, preventing the augmentation term from fully canceling their effect. Indeed, we show that GAIPW cannot eliminate first-order bias from errors in the exposure propensities.

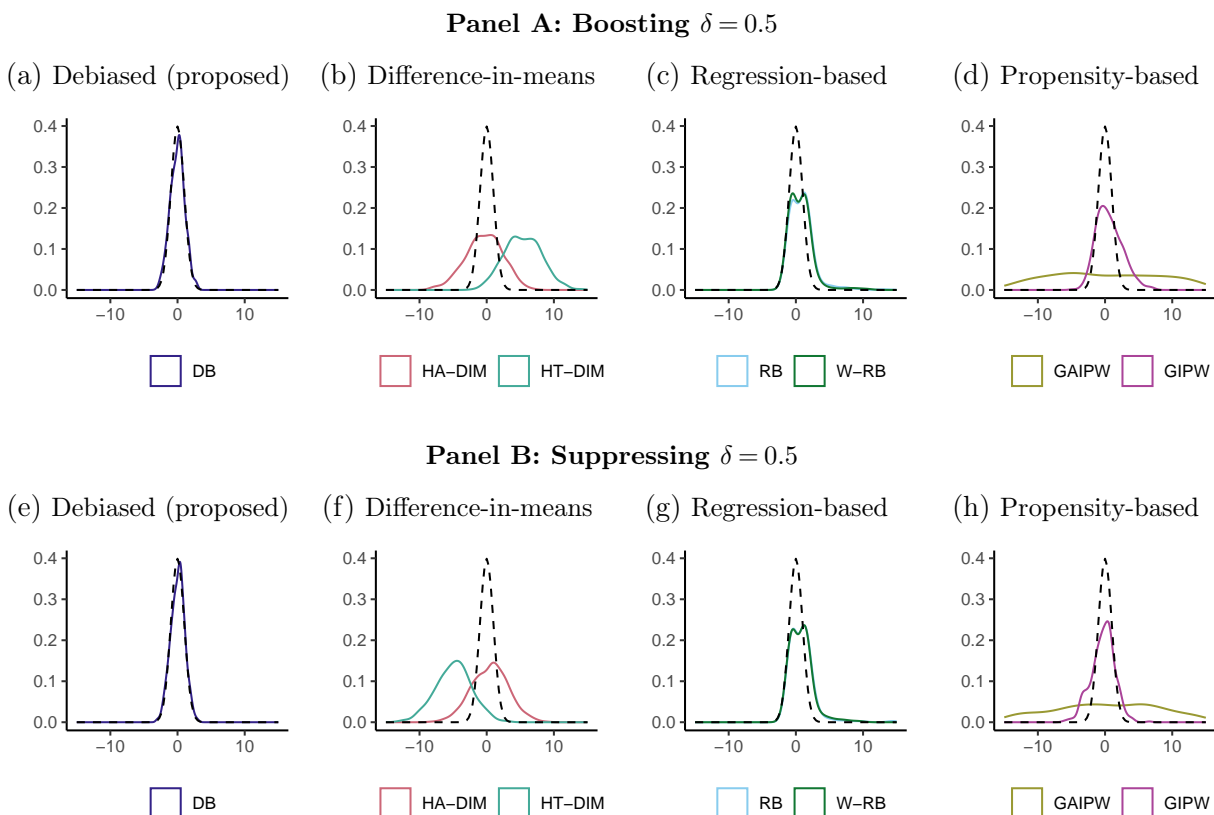
6.2. Simulation Results

We conduct 500 Monte Carlo simulations under two scenarios: one with $\delta = 0.5$ to mimic a boosting treatment and the other with $\delta = -0.5$ to mimic a suppressing treatment. We compare the proposed debiased (DB) estimator $\hat{\tau}_n^{DB}$ with the six benchmark estimators. For each simulation, we calculate both the point estimate and the standard error.⁷

⁷ The standard error for the DB estimator is calculated as $n^{-1/2} \widehat{V}_n^{1/2}$ following Theorem 2, using 3-fold cross-fitting for nuisance estimation. For the Hájek DIM estimator, we use the plug-in variance from Theorem 1.2 of Wager (2024). For the (weighted) regression-based estimators, we use the classical OLS formula. For the remaining benchmarks, standard errors are obtained as the sample standard deviation of the estimates.

Figure 2 presents the distribution of standardized estimates from all estimators. Specifically, we plot $(\hat{\tau} - \tau)/\hat{\sigma}$, where τ is the true ATE, $\hat{\tau}$ the point estimate, and $\hat{\sigma}$ the estimated standard error. This normalization expresses estimator bias relative to its estimated variability. For valid inference, the distribution of standardized estimates should approximate the standard normal distribution, shown by the dashed black curve. Across both boosting and suppressing scenarios, the DB estimator's distribution closely aligns with the standard normal. This provides empirical evidence, consistent with theoretical predictions, that the estimator is unbiased (centered at zero) and its standard errors are correctly estimated (variance is one). In contrast, the benchmark estimators exhibit noticeable bias or overly flat distributions, reflecting underestimated standard errors.

Figure 2 Monte Carlo Results: Standardized Distribution



Note: Density plots of the standardized estimates $(\hat{\tau} - \tau)/\hat{\sigma}$, where $\hat{\tau}$ is the estimator, τ is the true ATE, and $\hat{\sigma}$ is the estimated standard error.

Table 1 reports the bias and uncertainty estimates for the debiased and benchmark estimators under both boosting ($\delta = 0.5$) and suppressing ($\delta = -0.5$) scenarios. The first two columns summarize estimator bias: Column (1) reports the average of the estimate minus the true value, and Column (2) reports its standard error. An unbiased estimator should have a mean bias not statistically different from zero. The DB estimator satisfies this criterion, showing nearly zero bias

with small sampling variability, confirming its unbiasedness in finite samples. All other estimators exhibit varying degrees of bias, with the mean bias significantly different from zero. The HT-DIM estimator shows the largest bias, positive in the boosting scenario and negative in the suppressing scenario, consistent with item exposure bias. Regression-based estimators perform reasonably well here, likely because the data-generating process is relatively simple and linear adjustment captures most of the bias. However, their performance would likely deteriorate under a more complex, non-linear data-generating process. For propensity-based estimators, bias can arise in either direction, reflecting their lack of Neyman orthogonal and resulting sensitivity to nuisance estimation errors.

Table 1 Point Estimate and Uncertainty Estimate

	Bias ($\hat{\tau} - \tau$)		Uncertainty $\hat{\sigma}$	
	Mean (1)	(Std. Err.) (2)	True (MC) (3)	Model Est. (4)
<i>Boosting $\delta = 0.5$</i>				
Debiased (proposed)				
DB	-0.0002	(0.0004)	0.0097	0.0091
Difference-in-means				
HT-DIM	0.4167	(0.0094)	0.2105	0.0774
HA-DIM	-0.0042	(0.0013)	0.0297	0.0107
Regression-based				
RB	0.0007	(0.0002)	0.0037	0.0030
W-RB	0.0007	(0.0002)	0.0041	0.0033
Propensity-based				
GIPW	0.0026	(0.0006)	0.0133	0.0078
GAIPW	0.0021	(0.0004)	0.0098	0.0007
<i>Suppressing $\delta = 0.5$</i>				
Debiased (proposed)				
DB	0.0004	(0.0004)	0.0094	0.0090
Difference-in-means				
HT-DIM	-0.3801	(0.0093)	0.2073	0.0774
HA-DIM	0.0053	(0.0013)	0.0294	0.0107
Regression-based				
RB	0.0011	(0.0002)	0.0039	0.0031
W-RB	0.0011	(0.0002)	0.0042	0.0034
Propensity-based				
GIPW	-0.0009	(0.0005)	0.0119	0.0075
GAIPW	-0.0008	(0.0004)	0.0090	0.0007

Table 1 also evaluates how well each estimator captures uncertainty: Column (3) reports the true Monte Carlo standard deviation, and Column (4) reports the average estimated standard error. For valid inference, the model-based standard errors should closely match the Monte Carlo counterparts. The debiased estimator satisfies this criterion, with estimated standard error closely aligned with

the true sampling variability. Regression-based estimators also produce reasonable standard error estimates, but since they remain biased, they still do not support valid inference. Most benchmark estimators underestimate uncertainty, largely because they ignore sample correlation arising from overlapping items in consideration sets that share the same treatment status. In particular, as shown in Theorem 3, the GAIPW estimator is only partially orthogonal and therefore underestimates variance, leading to overconfident inference. In contrast, the DB estimator accurately quantifies uncertainty and achieves the asymptotic properties of the oracle estimator through Neyman orthogonality (Section 5.2).

7. Empirical Application

In this section, we present results from a large-scale field experiment. The results show that the treatment effect estimates from the proposed debiased estimator closely align with the approximate ground truth, obtained from the costly double-sided design, whereas estimates from other methods deviate more substantially and in some cases even have the opposite sign.

7.1. Institutional Background

We conducted a large-scale field experiment in collaboration with *Weixin Channels*, the short-video platform embedded in China’s largest mobile messaging app. Similar to other short-video platforms (e.g., TikTok), viewers watch one video at a time and swipe to request the next recommendation. Each viewer query thus results in the exposure of a single video, making it directly consistent with our recommender choice model.

The organic and promoted content items operate in separate recommendation channels. When a viewer reaches an advertising “slot”, the recommender system selects exclusively from the pool of promoted videos. Since promoted content is an important source of platform revenue, the platform frequently experiments with different algorithms for serving promoted content and evaluates their impact on advertisers’ outcomes. Because these recommender experiments directly target advertisers, randomization is naturally conducted at the advertiser (i.e., creator) level.

We conduct one such experiment focusing on promoted content. The treatment recommendation uses a new re-ranking algorithm, while the control recommendation maintains the status quo. The intervention occurs at the re-ranking stage where the consideration sets, consisting of videos that have reached the stage prior to the intervention, are relatively small (see Section 4.4). Interference arises when treated videos, using treatment algorithm for scoring, and control videos, using control algorithm, appear in the same consideration set competing for exposure.

We discuss the plausibility of the set of assumptions in our empirical context. The controlled appearances Assumption 1 ensures that the same item does not repeatedly appear in many consideration sets. The bounded scores Assumption 2 ensures that every video has a positive probability

of being exposed. In our empirical setting, both assumptions are reasonable since promoted videos are subject to ad budget and delivery constraints, so that each item has positive exposure and no item can dominate a very large number of user queries.⁸ Furthermore, viewers rarely encounter two promoted videos in close succession, so temporal interference across queries is likely limited.

7.2. Experimental Design

To evaluate the validity of treatment effect estimators, we simultaneously conduct a creator-side experiment and a blocked double-sided experiment, with the latter serving as an interference-free benchmark. We track three key outcomes of interest. While we cannot disclose their exact definitions for confidentiality reasons, they are all binary measures related to viewer engagement.

We randomly divide viewers and creators into three equal-sized sub-universes to mitigate potential market-size effects. Each sub-universe is blocked from the others, so viewers can only see videos from creators within the same sub-universe. The first sub-universe hosts the creator experiment, where half of the creators are randomly assigned to the treatment algorithm and half to the control. The second and third sub-universes implement the double-sided randomized experiment: one uses the treatment algorithm, while the other uses the control algorithm. Because the two sub-universes are blocked, treated and control videos never appear in the same consideration set, eliminating interference. As a result, the double-sided design provides an interference-free benchmark for the true treatment effect, denoted as $\hat{\tau}^{DS}$. While effective, the double-sided design is costly to implement and thus impractical for evaluating all creator-focused treatments at scale (see Section 2.2).

The experiments from all three sub-universes were conducted simultaneously to avoid time-varying differences in advertiser outcomes. The experiment ran for eight days, from May 27 to June 3, 2024, encompassing approximately 32.7 million viewer queries and 0.1 million videos, evenly split across the three sub-universes. For each viewer query i , we observe $(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i)$, where V_i and $C_{i,k}$ include both raw features and learned embeddings from historical data. The treatment assignment vector $\vec{W}_i = (W_{i,1}, \dots, W_{i,K})$ indicates whether the video was scored using the treatment ($W_{i,k} = 1$) or control ($W_{i,k} = 0$) algorithm.

For randomization check, we compare the distributions of several key covariates identified as strong predictors of the outcomes of interest in historical data. While we cannot disclose their exact definitions, these covariates are commonly examined by practitioners to ensure that randomization was implemented properly. The set includes video features, users' historical engagement, pre-treatment values of the three outcome variables, and a system-level performance metric. Table 2 reports p-values from balance tests on six categorical covariates and four continuous covariates. The

⁸ In our data, the most popular item appears in only 2.33% of consideration sets, while about 90% of items appear in no more than 5.28×10^{-5} of total user queries.

“creator experiment” and “double-sided experiment” rows test covariate balance between treatment and control groups within each experiment, while the “across experiments” row compares covariates between the two experimental designs. None of the covariates show significant imbalance, confirming that randomization was implemented successfully.

Table 2 P-values from Randomization Check

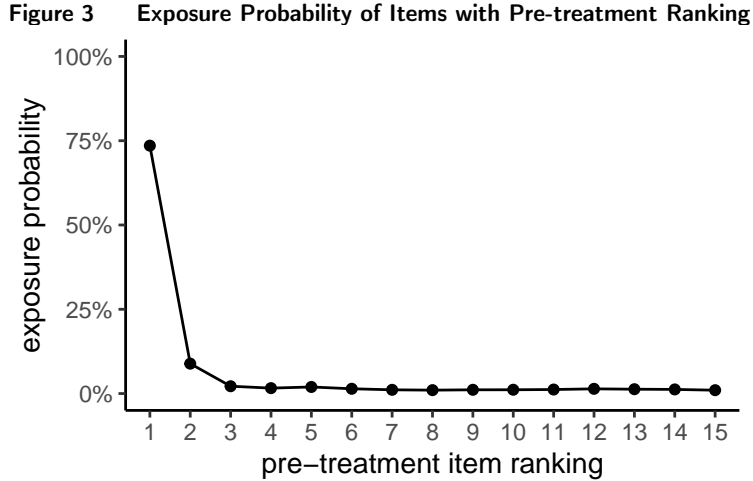
	Categorical Covariates					
	Cat.1	Cat.2	Cat.3	Cat.4	Cat.5	Cat.6
Creator experiment	0.279	0.789	0.462	0.770	0.689	0.393
Double-sided experiment	0.137	0.102	0.976	0.200	0.165	0.067
Across experiments	0.792	0.168	0.047	0.540	0.056	0.461
	Continuous Covariates					
	Cont.1	Cont.2	Cont.3	Cont.4		
Creator experiment	0.580	0.403	0.521	0.549		
Double-sided experiment	0.186	0.225	0.771	0.630		
Across experiments	0.537	0.921	0.461	0.819		

7.3. Estimation Details

We describe the implementation details of the proposed debiased estimator using data from the creator-side experiment. Estimating the debiased estimator in Equation (18) requires several components: the nuisance functions \hat{s}_0 and \hat{s}_1 from the recommender choice model, the viewer response model \hat{z} , and the gradients and Hessian matrix to construct the debiasing term. We discuss each component in turn.

We start with the recommender choice model. To construct the loss function in Equation (13), we first define the consideration set. This is done by empirically identifying the smallest set of items that could realistically be exposed prior to treatment. Figure 3 shows the empirical exposure probability of items based on their “pre-treatment” ranking by the recommendation algorithm. Exposure probability is highest for the top-ranked item and quickly declines to nearly zero. Based on this pattern, we set the consideration set size to $K = 15$, which captures nearly all items with non-negligible exposure probability. This choice ensures that the consideration sets cover relevant items while not being too large.

The recommender choice model is estimated using data $(V_i, \vec{C}_i, \vec{W}_i, k^*)$. Figure 4 illustrates the architecture of the semiparametric model. The viewer and content embeddings $(V_i, C_{i,k})$ pass through a black-box component that outputs two quantities: the baseline score s_0 and the treatment uplift s_1 . These scores are then passed to a model layer that computes the exposure probabilities for all items in the consideration set.



To compute the scores, we leverage the platform’s pre-trained representation model $\phi(\cdot)$ to generate features for each viewer–content pair $X_{i,k} := \phi(C_{i,k}, V_i)$. The representation model ϕ , which is trained on historical data, captures a range of expected evaluations such as predicted viewer engagement. These features are then passed through two fully connected layers that output the baseline score and the treatment uplift. During estimation, the pre-trained representation model is held fixed, while the fully connected layers are updated. This approach can be viewed as fine-tuning the pre-trained model for our specific estimation task. Leveraging these pre-trained features substantially reduces both data and computational demands relative to training from raw embeddings.

In the model layer, the treatment assignment vector \vec{W}_i enters explicitly, and exposure probabilities for all items in the consideration set are computed following Equation (9). The nuisance components \hat{s}_0 and \hat{s}_1 are estimated by maximizing the likelihood that item k^* is selected for viewer V_i from the consideration set \vec{C}_i , conditional on treatment assignment \vec{W}_i . The model is trained with categorical cross-entropy loss as specified in Equation (13).

We evaluate the fit of the recommender choice model by comparing its predictions with the actual exposed items. Figure 5 presents the ROC curve. The dashed diagonal line corresponds to a non-informative classifier with an AUC of 0.5, whereas our fitted model achieves an AUC of 0.97, demonstrating excellent out of sample predictive accuracy. It is worth noting that our network structure, which builds upon a pre-trained representation model, may not strictly satisfy the width and depth conditions specified in Proposition 2. These conditions guarantee that the estimated nuisance components converge faster than $n^{-1/4}$, a key requirement for the asymptotic normality and inference guarantees established in Theorem 2. Nonetheless, the model’s strong empirical performance suggests that the convergence rate is likely sufficiently fast in practice.

For the view response model, we leverage the platform’s pre-trained models that predict various outcomes for each viewer–content pair. Since the experiment tracks three outcomes of interest,

Figure 4 Semiparametric Recommender Choice Model

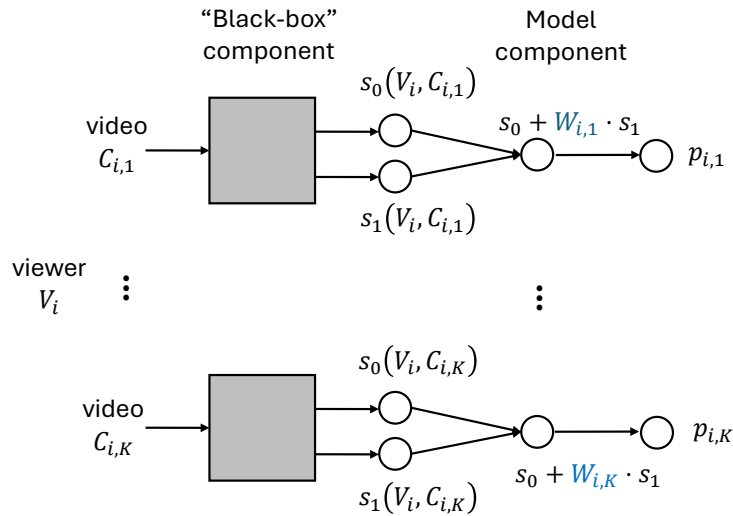

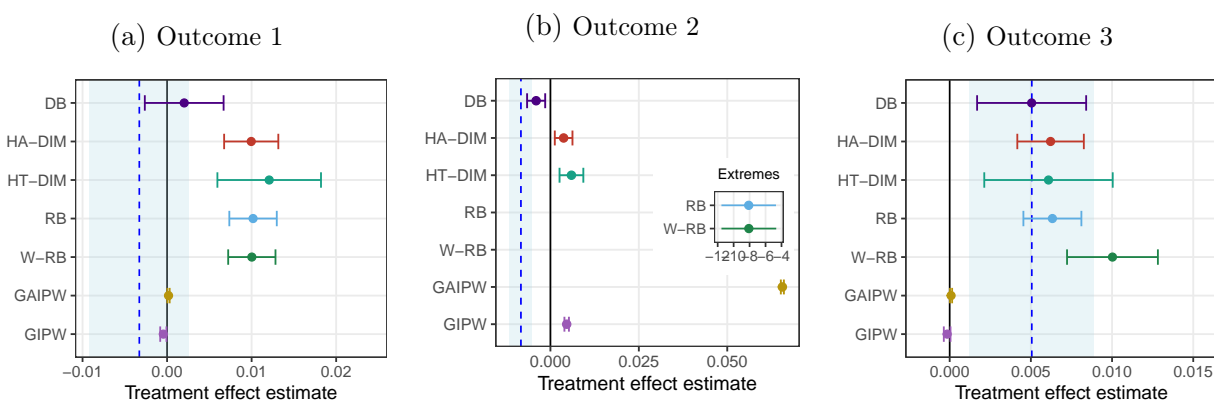


Figure 5 ROC Curve of the Fitted Recommender Choice Model



each uses a separate viewer response model, while they all share the same recommender choice model. The observed viewer responses from the experiment are used to compute the loss specified in Equation (15), which contributes to the construction of the debiasing term.

The debiasing term contains the gradients of the plug-in estimator and the loss function with respect to the nuisance components, as well as the Hessian of the loss function. The gradients can be computed directly in our setting, as detailed in Section 5.2. The expected Hessian matrix of the loss function is taken with respect to the treatment assignment. To approximate this expectation, we draw 500 realizations of the treatment assignment and compute the empirical mean of the corresponding Hessians across these samples.

Figure 6 Treatment Effect Estimation


Note: The dashed blue line indicates the treatment effect estimated from the (approximate) ground truth using the double-sided design, and the blue shaded area covers the 95% confidence interval.

7.4. Empirical Results

We evaluate the debiased estimator using data from the field experiment. Specifically, we estimate the treatment effect from the creator-side experiment using the debiased estimator alongside all benchmark estimators. To ensure a fair comparison, the viewer response models and recommender choice model employed in the debiased estimator are also leveraged in the benchmark methods where applicable, including the construction of the propensities of the GIPW and GAIPW estimators.

Figure 6 presents the estimated treatment effects, with each panel corresponding to one of the three outcomes of interest. The (approximate) ground-truth measures obtained from the double-sided experiment are shown as dashed blue lines, with shaded regions representing the 95% confidence intervals. The solid black lines at zero indicate the natural benchmark for assessing whether the new algorithm should be adopted. For the proposed and benchmark estimators, we evaluate how closely their estimated treatment effects align with the ground-truth measures.

The practical value and economic significance of our de-biased (DB) estimator are most starkly illustrated by the results for outcome 2. For this metric, the ground truth measure (dashed blue line) from the double-sided experiment reveals a statistically significant negative treatment effect. Our DB estimator successfully recovers this result with comparable magnitude and significance, correctly identifying the new algorithm as the worse option. In sharp contrast, all benchmark estimators fail. The standard DIM and propensity-based estimators report a statistically significant positive effect, while the regression-based estimators are wildly inaccurate. Relying on these estimators would have led to an incorrect managerial decision regarding which recommender algorithm to adopt. This result highlights the most critical danger of using naive estimators: the potential for reaching directionally wrong conclusions.

The unreliability of the benchmark estimators is also apparent in the other two outcomes. For outcome 1, the ground-truth measure indicates a null effect, and our DB estimator correctly cannot

reject the null hypothesis. In contrast, the DIM and regression-based estimators both report a statistically significant positive effect. The GIPW and GAIPW estimators align with the ground truth in terms of point estimates but are overly confident, consistent with simulation results in Section 6.2. While not as severe as getting the sign wrong, this would still lead the platform to dedicate resources to a costly rollout for a change with no actual benefit.

For outcome 3, while some of the benchmarks (the DIM and regression-based estimators) happen to align with the positive ground truth, their demonstrated failure in the other cases makes this result untrustworthy. The platform should not rely on an estimator that is correct only by chance. Across all three key metrics, our DB estimator is the only approach that consistently provides reliable results aligned with the ground truth, making it a trustworthy guide for high-stakes business decisions. In contrast, all other estimators yield confidence intervals that fail to overlap with the ground truth in at least two of the three outcomes.

Comparing across the three outcome measures, treatment effect on outcome 2 is particularly challenging to estimate. All benchmark methods perform poorly in this case, with regression-based estimators yielding extreme estimates and the DIM and propensity-based estimators producing effects with the wrong sign. Although it is difficult to pinpoint the exact source of difficulty, outcome 2 exhibits a highly skewed distribution with a large mass at zero,⁹ which generally makes the estimation of nuisance components more difficult. The DB estimator mitigates the issue through its Neyman orthogonality property, a feature not shared by the other estimators in our setting.

8. Conclusion

Creator-side experiments are important for platforms to evaluate creator-focused recommendation system updates. This paper demonstrates that interference from competition between content creators makes standard DIM estimators dangerously misleading and leads to biased treatment effect estimation. We show empirically that relying on standard DIM estimators can lead to the wrong business decision by deploying the worse algorithm.

To address this challenge, we develop a new approach grounded in the Double/Debiased Machine Learning framework that provides a reliable estimate using data from a standard creator-side experiment. Our method explicitly models the interference pathway with a semi-parametric choice model and uses a debiased estimator to ensure valid inference. We validate our estimator against ground-truth estimates obtained from a costly double-sided experimental design, showing that it successfully recovers the true treatment effect where common benchmarks fail.

This paper’s contributions are therefore both practical and theoretical. For practitioners, we offer a trustworthy tool to make better decisions about their algorithms. For academics, our key methodological contribution is the extension of the DML framework to handle the non-i.i.d., correlated

⁹ The Fisher–Pearson coefficients of skewness for the three outcomes are 10.46, 21.21, and 11.09, respectively.

data. While our analysis focuses on short-run effects in a re-ranking context, future work could extend this framework to model ranked slates recommendation or long-term ecosystem effects. More broadly, this framework demonstrates how combining structural and machine learning components in semi-parametric models can address complex interference and identification challenges in modern digital marketplaces.

References

- Bajari P, Burdick B, Imbens GW, Masoero L, McQueen J, Richardson T, Rosen IM (2021) Multiple randomization designs. *arXiv preprint arXiv:2112.13495* .
- Bojinov I, Simchi-Levi D, Zhao J (2023) Design and analysis of switchback experiments. *Management Science* 69(7):3759–3777.
- Braun M, Schwartz EM (2025) Where a/b testing goes wrong: How divergent delivery affects what online experiments cannot (and can) tell you about how customers respond to advertising. *Journal of Marketing* 89(2):71–95.
- Bright I, Delarue A, Lobel I (2022) Reducing marketplace interference bias via shadow prices. *arXiv preprint arXiv:2205.02274* .
- Burtch G, Moakler R, Gordon BR, Zhang P, Hill S (2025) Characterizing and minimizing divergent delivery in meta advertising experiments. *arXiv preprint arXiv:2508.21251* .
- Cervera MR, Dätwyler R, D’Angelo F, Keurti H, Grewe BF, Henning C (2021) Uncertainty estimation under model misspecification in neural network regression. *arXiv preprint arXiv:2111.11763* .
- Cheng Y, Wang J, Cao X, Shen ZJM, Zhang Y (2023) Selecting creators to sign on a content-sharing platform: A deep-did approach. *Available at SSRN 4622422* .
- Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W, Robins J (2018) Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* 21(1):C1–C68.
- Chernozhukov V, Demirer M, Lewis G, Syrgkanis V (2019) Semi-parametric efficient policy learning with continuous actions. *Advances in Neural Information Processing Systems* 32.
- Covington P, Adams J, Sargin E (2016) Deep neural networks for youtube recommendations. *Proceedings of the 10th ACM conference on recommender systems*, 191–198.
- Dhaouadi W, Johari R, Weintraub GY (2023) Price experimentation and interference in online platforms. *arXiv preprint arXiv:2310.17165* .
- Ellickson PB, Kar W, Reeder III JC, Zeng G (2024) Using contextual embeddings to predict the effectiveness of novel heterogeneous treatments. *Available at SSRN 4845956* .
- Farias V, Li H, Peng T, Ren X, Zhang H, Zheng A (2023) Correcting for interference in experiments: A case study at douyin. *Proceedings of the 17th ACM Conference on Recommender Systems*, 455–466.
- Farias VF, Li AA, Peng T, Zheng AT (2022) Markovian interference in experiments. *arXiv preprint arXiv:2206.02371* .
- Farrell MH, Liang T, Misra S (2021a) Deep learning for individual heterogeneity: An automatic inference framework. *arXiv preprint arXiv:2010.14694* .
- Farrell MH, Liang T, Misra S (2021b) Deep neural networks for estimation and inference. *Econometrica* 89(1):181–213.

- Foster DJ, Syrgkanis V (2023) Orthogonal statistical learning. *The Annals of Statistics* 51(3):879–908.
- Goli A, Lambrecht A, Yoganarasimhan H (2023) A bias correction approach for interference in ranking experiments. *Marketing Science* .
- Hahn J (1998) On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica* 315–331.
- Hájek J (1971) Comment on “an essay on the logical foundations of survey sampling, part one”. *The foundations of survey sampling* 236.
- Hall P, Heyde CC (2014) *Martingale limit theory and its application* (Academic press).
- Hinton G, Vinyals O, Dean J (2015) Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* .
- Holtz D, Aral S (2020) Limiting bias from test-control interference in online marketplace experiments. *arXiv preprint arXiv:2004.12162* .
- Holtz D, Lobel F, Lobel R, Liskovich I, Aral S (2023) Reducing interference bias in online marketplace experiments using cluster randomization: Evidence from a pricing meta-experiment on airbnb. *Management Science* .
- Horvitz DG, Thompson DJ (1952) A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association* 47(260):663–685.
- Hu Y, Li S, Wager S (2021) Average treatment effects in the presence of interference. *arXiv preprint arXiv:2104.03802* .
- Hu Y, Wager S (2022) Switchback experiments under geometric mixing. *arXiv preprint arXiv:2209.00197* .
- Hudgens MG, Halloran ME (2008) Toward causal inference with interference. *Journal of the American Statistical Association* 103(482):832–842.
- Imbens GW (2004) Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and statistics* 86(1):4–29.
- Imbens GW, Rubin DB (2015) *Causal Inference in Statistics, Social, and Biomedical Sciences* (Cambridge University Press).
- Johari R, Li H, Liskovich I, Weintraub GY (2022) Experimental design in two-sided platforms: An analysis of bias. *Management Science* 68(10):7069–7089.
- Kim D, Jiang Z, Thomadsen R (2023) Tv advertising effectiveness with racial minority representation: Evidence from the mortgage market. *Available at SSRN 4521178* .
- Kiros H (2022) Hated that video? YouTube’s algorithm might push you another just like it. URL <https://www.technologyreview.com/2022/09/20/1059709/youtube-algorithm-recommendations/#:~:text=YouTube’s%20recommendation%20algorithm%20drives%2070,adjust%20what%20it%20shows%20them>.

- Mehrotra R, McInerney J, Bouchard H, Lalmas M, Diaz F (2018) Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. *Proceedings of the 27th acm international conference on information and knowledge management*, 2243–2251.
- Mummalaneni S, Yoganarasimhan H, Pathak V (2022) How do content producers respond to engagement on social media platforms? *Available at SSRN 4173537* .
- Newey WK (1994) The asymptotic variance of semiparametric estimators. *Econometrica: Journal of the Econometric Society* 62(6):1349–1382.
- Ni T, Bojinov I, Zhao J (2023) Design of panel experiments with spatial and temporal interference. *Available at SSRN 4466598* .
- Rosen S (1981) The economics of superstars. *The American economic review* 71(5):845–858.
- Sävje F, Aronow P, Hudgens M (2021) Average treatment effects in the presence of unknown interference. *Annals of statistics* 49(2):673.
- Su Y, Wang X, Le EY, Liu L, Li Y, Lu H, Lipshitz B, Badam S, Heldt L, Bi S, et al. (2024) Long-term value of exploration: Measurements, findings and algorithms. *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, 636–644.
- Tang J, Wang K (2018) Ranking distillation: Learning compact ranking models with high performance for recommender system. *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 2289–2298.
- Ugander J, Karrer B, Backstrom L, Kleinberg J (2013) Graph cluster randomization: Network exposure to multiple universes. *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, 329–337.
- Viviano D (2025) Policy targeting under network interference. *Review of Economic Studies* 92(2):1257–1292.
- Wager S (2024) Causal inference: A statistical learning approach *Accessed on Sep 4, 2025*. URL https://web.stanford.edu/~swager/causal_inf_book.pdf.
- Xiong R, Chin A, Taylor S (2023) Bias-variance tradeoffs for designing simultaneous temporal experiments. *The KDD'23 Workshop on Causal Discovery, Prediction and Decision*, 115–131 (PMLR).
- Ye Z, Zhang DJ, Zhang H, Zhang R, Chen X, Xu Z (2023a) Cold start to improve market thickness on online advertising platforms: Data-driven algorithms and field experiments. *Management Science* 69(7):3838–3860.
- Ye Z, Zhang Z, Zhang D, Zhang H, Zhang RP (2023b) Deep learning based causal inference for large-scale combinatorial experiments: Theory and empirical evidence. *Available at SSRN 4375327* .
- Zhu Z, Cai Z, Zheng L, Si N (2024) Seller-side experiments under interference induced by feedback loops in two-sided platforms. *arXiv preprint arXiv:2401.15811* .

Appendix A: Proof of Theorem 1 (Bias of DIM Estimators)

We adopt the following condition to control the extent of interference, in line with prior work on interference (Viviano 2025, Sävje et al. 2021).

ASSUMPTION 1 (Controlled Item Appearance). *Let n be the sample size. Each item appears in at most d_n consideration sets, with $\mathbb{E}[d_n^3] = o(n)$, where the expectation is taken with respect to the viewer population.*

Let n be the number of recommendations and $m = |\mathcal{C}|$ be the number of content items. We assume $m = O(n)$ without loss of generality since at most nK contents are considered. Let q be the probability of treatment assignment. For a viewer v , we use $\vec{c}_v = (c_{v,c})$ and $\vec{w}_v = (w_{v,c})$ to denote the content consideration set and the treatment statuses there in, and use $\vec{p}_v = (p_{v,c})$ to denote the exposure probabilities. Let $e^*(q | v)$ be the probability of a viewer v seeing treated content under the creator randomization with treatment assignment probability q . Mathematically, we have

$$e^*(q | v) = \mathbb{E}_{\mathbf{w} \sim \mathcal{B}(q)} [w_{k^*} | v]. \quad (25)$$

Similarly, define the population level probability of a viewer seeing a treated content item under the design as:

$$e^*(q) = \mathbb{E}_{v \sim \mathbb{P}_v} [e^*(q | v)], \quad (26)$$

where \mathbb{P}_v is the total viewer distribution.

Next, define the distributions of viewers exposed to treated and control content items under the randomization as $\mathbb{P}_v^{1,q}$ and $\mathbb{P}_v^{0,q}$ respectively, with:

$$\frac{d\mathbb{P}_v^{1,q}}{dv} = \frac{d\mathbb{P}_v}{dv} \cdot \frac{e^*(q | v)}{e^*(q)} \quad \text{and} \quad \frac{d\mathbb{P}_v^{0,q}}{dv} = \frac{d\mathbb{P}_v}{dv} \cdot \frac{1 - e^*(q | v)}{1 - e^*(q)}. \quad (27)$$

A.1. Convergence of Horvitz-Thompson estimator ($\hat{\tau}_n^{HT-DIM}$)

Given n samples $\{(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i)\}_{i=1}^n$ collected from a creator-side randomization experiment, with treatment probability q , the Horvitz-Thompson estimator ($\hat{\tau}_n^{HT-DIM}$) is defined as:

$$\hat{\tau}_n^{HT-DIM} := \frac{\sum_{i=1}^n W_{i,k_i^*} Y_i}{nq} - \frac{\sum_{i=1}^n (1 - W_{i,k_i^*}) Y_i}{n(1 - q)}. \quad (3)$$

We first show that $\mathbb{E}[\hat{\tau}_n^{HT-DIM}] = \tau^{HT}$, where τ^{HT} is defined as:

$$\tau^{HT} := \frac{e^*(q)}{q} \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v^{1,q}, \mathbf{w} \sim \mathcal{B}(q)} [r_c | w_{k^*} = 1] - \frac{1 - e^*(q)}{1 - q} \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v^{0,q}, \mathbf{w} \sim \mathcal{B}(q)} [r_c | w_{k^*} = 0], \quad (5)$$

For each sample $(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i)$ collected from the experiment, we have:

$$\begin{aligned} \mathbb{E} \left[\frac{1}{q} W_{i,k_i^*} Y_i \right] &= \frac{1}{q} \mathbb{E}_{v \sim \mathbb{P}_v} [\mathbb{E}_{\mathbf{w} \sim \mathcal{B}(q)} [w_{k^*} y | v]] \\ &= \frac{1}{q} \mathbb{E}_{v \sim \mathbb{P}_v^{1,q}} \left[\frac{e^*(q)}{e^*(q | v)} \mathbb{E}_{\mathbf{w} \sim \mathcal{B}(q)} [w_{k^*} y | v] \right] \\ &= \frac{e^*(q)}{q} \mathbb{E}_{v \sim \mathbb{P}_v^{1,q}} \left[\frac{1}{e^*(q | v)} \mathbb{E}_{\mathbf{w} \sim \mathcal{B}(q)} [w_{k^*} y | v] \right] \\ &= \frac{e^*(q)}{q} \mathbb{E}_{v \sim \mathbb{P}_v^{1,q}} \left[\frac{1}{\mathbb{E}_{\mathbf{w} \sim \mathcal{B}(q)} [w_{k^*} | v]} \mathbb{E}_{\mathbf{w} \sim \mathcal{B}(q)} [w_{k^*} y | v] \right] \end{aligned}$$

$$\begin{aligned}
&= \frac{e^*(q)}{q} \mathbb{E}_{v \sim \mathbb{P}_v^1, q} \left[\frac{1}{\mathbb{E}_{\mathbf{w} \sim \mathcal{B}(q)} [w_{k^*} | v]} \mathbb{E}_{\mathbf{w} \sim \mathcal{B}(q)} [w_{k^*} | v] \mathbb{E}_{\mathbf{w} \sim \mathcal{B}(q)} [y | w_{k^*} = 1, v] \right] \\
&= \frac{e^*(q)}{q} \mathbb{E}_{v \sim \mathbb{P}_v^1, q} \left[\mathbb{E}_{\mathbf{w} \sim \mathcal{B}(q)} [y | w_{k^*} = 1, v] \right] \\
&= \frac{e^*(q)}{q} \mathbb{E}_{v \sim \mathbb{P}_v^1, q, \mathbf{w} \sim \mathcal{B}(q)} [y | w_{k^*} = 1] \\
&= \frac{e^*(q)}{q} \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v^1, q, \mathbf{w} \sim \mathcal{B}(q)} [r_c | w_{k^*} = 1],
\end{aligned}$$

where the last line uses the definition of r_c in Section 2.1 such that $r_c = y$ if content c is exposed to viewer v , otherwise $r_c = 0$.

Similarly, we have

$$\mathbb{E} \left[\frac{1}{1-q} (1 - W_{i, k_i^*}) Y_i \right] = \frac{1 - e^*(q)}{1 - q} \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v^0, q, \mathbf{w} \sim \mathcal{B}(q)} [r_c | w_{k^*} = 0]. \quad (28)$$

Recall that $\hat{\tau}_n^{HT-DIM} = \frac{1}{n} \sum_{i=1}^n \frac{1}{q} W_{i, k_i^*} Y_i - \frac{1}{1-q} (1 - W_{i, k_i^*}) Y_i$. Therefore we have $\mathbb{E}[\hat{\tau}_n^{HT-DIM}] = \tau^{HT}$.

Next, we characterize the variance of $\hat{\tau}_n^{HT-DIM}$. By the law of total variance, we have

$$\text{Var}(\hat{\tau}_n^{HT-DIM}) = \mathbb{E} \left[\text{Var} \left(\hat{\tau}_n^{HT-DIM} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) \right] + \text{Var} \left(\mathbb{E} \left[\hat{\tau}_n^{HT-DIM} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right] \right).$$

Conditioning on $\{(V_i, \vec{C}_i)\}_{i=1}^n$, for each content c , let w_c denote its treatment status, d_c denote the number of consideration set including content c , and $R_c := \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{c = C_{i, k_i^*}\} Y_i$ to denote the average viewer-outcome for content c . Denote $d = \max_c d_c$. We have:

$$\begin{aligned}
&\text{Var} \left(\frac{w_c R_c}{q} - \frac{(1 - w_c) R_c}{1 - q} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) \\
&= \frac{\text{Var}(w_c R_c \mid \{(V_i, \vec{C}_i)\}_{i=1}^n)}{q^2} + \frac{\text{Var}((1 - w_c) R_c \mid \{(V_i, \vec{C}_i)\}_{i=1}^n)}{(1 - q)^2} - \frac{2 \text{Cov}(w_c R_c, (1 - w_c) R_c \mid \{(V_i, \vec{C}_i)\}_{i=1}^n)}{q(1 - q)} \\
&= \frac{\mathbb{E}[w_c R_c^2 \mid \{(V_i, \vec{C}_i)\}_{i=1}^n] - \mathbb{E}[w_c R_c \mid \{(V_i, \vec{C}_i)\}_{i=1}^n]^2}{q^2} + \frac{\mathbb{E}[(1 - w_c) R_c^2 \mid \{(V_i, \vec{C}_i)\}_{i=1}^n] - \mathbb{E}[(1 - w_c) R_c]^2}{(1 - q)^2} \\
&\quad + \frac{2 \mathbb{E}[w_c R_c \mid \{(V_i, \vec{C}_i)\}_{i=1}^n] \mathbb{E}[(1 - w_c) R_c \mid \{(V_i, \vec{C}_i)\}_{i=1}^n]}{q(1 - q)} \\
&= \frac{\mathbb{E}[R_c^2 \mid w_c = 1, \{(V_i, \vec{C}_i)\}_{i=1}^n]}{q} - \mathbb{E}[R_c \mid w_c = 1, \{(V_i, \vec{C}_i)\}_{i=1}^n]^2 \\
&\quad + \frac{\mathbb{E}[R_c^2 \mid w_c = 0, \{(V_i, \vec{C}_i)\}_{i=1}^n]}{1 - q} - \mathbb{E}[R_c \mid w_c = 0, \{(V_i, \vec{C}_i)\}_{i=1}^n]^2 \\
&\quad + 2 \mathbb{E}[R_c \mid w_c = 1, \{(V_i, \vec{C}_i)\}_{i=1}^n] \mathbb{E}[R_c \mid w_c = 0, \{(V_i, \vec{C}_i)\}_{i=1}^n] \\
&= \frac{\mathbb{E}[R_c^2 \mid w_c = 1, \{(V_i, \vec{C}_i)\}_{i=1}^n]}{q} + \frac{\mathbb{E}[R_c^2 \mid w_c = 0, \{(V_i, \vec{C}_i)\}_{i=1}^n]}{1 - q} \\
&\quad - \left(\mathbb{E}[R_c \mid w_c = 1, \{(V_i, \vec{C}_i)\}_{i=1}^n] - \mathbb{E}[R_c \mid w_c = 0, \{(V_i, \vec{C}_i)\}_{i=1}^n] \right)^2.
\end{aligned}$$

Recall that $R_c := \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{c = C_{i, k_i^*}\} Y_i$. By the boundedness of Y_i , we have that $R_c = O(\frac{d_c}{n})$, and thus:

$$\text{Var} \left(\frac{w_c R_c}{q} - \frac{(1 - w_c) R_c}{1 - q} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) = O \left(\frac{d_c^2}{n^2} \right) = O \left(\frac{d^2}{n^2} \right).$$

Also conditioning $\{(V_i, \vec{C}_i)\}_{i=1}^n$, define $I_{c_1, c_2} = 1$ if item c_1 and item c_2 are present at one consideration set;¹⁰ otherwise set $I_{c_1, c_2} = 0$. Note that if $I_{c_1, c_2} = 0$, there is no interference among the items c_1 & c_2 and thus R_{c_1} and R_{c_2} are independent.

$$\begin{aligned} & \text{Cov} \left(\frac{w_{c_1} R_{c_1}}{q} - \frac{(1-w_{c_1}) R_{c_1}}{1-q}, \frac{w_{c_2} R_{c_2}}{q} - \frac{(1-w_{c_2}) R_{c_2}}{1-q} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) \\ &= \begin{cases} 0 & \text{if } I_{c_1, c_2} = 0, \\ O \left(\sqrt{\text{Var} \left(\frac{w_{c_1} R_{c_1}}{q} - \frac{(1-w_{c_1}) R_{c_1}}{1-q} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right)} \right. & \text{o.w.} \\ \quad \left. \times \sqrt{\text{Var} \left(\frac{w_{c_2} R_{c_2}}{q} - \frac{(1-w_{c_2}) R_{c_2}}{1-q} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right)} \right) & \end{cases} \end{aligned}$$

Together, we have

$$\begin{aligned} \text{Var} \left(\hat{\tau}_n^{HT-DIM} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) &\lesssim \sum_{c \in \mathcal{C}} \text{Var} \left(\frac{w_c R_c}{q} - \frac{(1-w_c) R_c}{1-q} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) \\ &\quad + \sum_{c_1 \neq c_2 \in \mathcal{C}} \text{Cov} \left(\frac{w_{c_1} R_{c_1}}{q} - \frac{(1-w_{c_1}) R_{c_1}}{1-q}, \frac{w_{c_2} R_{c_2}}{q} - \frac{(1-w_{c_2}) R_{c_2}}{1-q} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) \\ &= O \left(\frac{d^2}{n^2} \right) + \sum_{c_1 \neq c_2 \in \mathcal{C}} I_{c_1, c_2} O \left(\frac{d^2}{n^2} \right) \\ &= O_p \left(\frac{md^3}{n^2} \right) = O_p \left(\frac{d^3}{n} \right). \end{aligned}$$

Thus,

$$\mathbb{E} \left[\text{Var} \left(\hat{\tau}_n^{HT-DIM} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) \right] = O \left(\frac{\mathbb{E}[d^3]}{n} \right). \quad (29)$$

On the other hand, we have

$$\begin{aligned} \mathbb{E} \left[\hat{\tau}_n^{HT-DIM} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right] &= \sum_{c \in \mathcal{C}} \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \frac{\mathbf{1}\{c = C_{i, k_i^*}\} w_c}{q} Y_i - \frac{1}{n} \sum_{i=1}^n \frac{\mathbf{1}\{c = C_{i, k_i^*}\} (1-w_c)}{1-q} Y_i \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right] \\ &= \frac{1}{n} \sum_{i=1}^n \sum_{c \in \mathcal{C}} \left\{ \mathbb{E} \left[\frac{\mathbf{1}\{c = C_{i, k_i^*}\} w_c}{q} Y_i \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right] - \mathbb{E} \left[\frac{\mathbf{1}\{c = C_{i, k_i^*}\} (1-w_c)}{1-q} Y_i \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right] \right\} \\ &= \frac{1}{n} \sum_{i=1}^n Z_i, \end{aligned}$$

$$\text{where } Z_i = \mathbb{E} \left[\left(\frac{W_{i, k_i^*}}{q} - \frac{1-W_{i, k_i^*}}{1-q} \right) Y_i \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right] = \mathbb{E} \left[\left(\frac{W_{i, k_i^*}}{q} - \frac{1-W_{i, k_i^*}}{1-q} \right) Y_i \mid (V_i, \vec{C}_i) \right].$$

Note that the randomness in Z_i comes solely from (V_i, \vec{C}_i) , which are i.i.d. across i . Thus,

$$\text{Var} \left(\mathbb{E} \left[\hat{\tau}_n^{HT-DIM} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right] \right) = \text{Var} \left(\frac{1}{n} \sum_{i=1}^n Z_i \right) = \frac{\sum_{i=1}^n \text{Var}(Z_i)}{n^2} = O \left(\frac{1}{n} \right). \quad (30)$$

Combining Equation (29) & Equation (30), we have

$$\begin{aligned} \text{Var} \left(\hat{\tau}_n^{HT-DIM} \right) &= \mathbb{E} \left[\text{Var} \left(\hat{\tau}_n^{DIM} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) \right] + \text{Var} \left(\mathbb{E} \left[\hat{\tau}_n^{DIM} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right] \right) \\ &= O \left(\frac{\mathbb{E}[d^3]}{n} \right) + O \left(\frac{1}{n} \right) = O \left(\frac{\mathbb{E}[d^3]}{n} \right) = o(1), \end{aligned}$$

where the last equation is by Assumption 1. Recall that we have proved $\mathbb{E}[\hat{\tau}_n^{HT-DIM}]$ equals to τ^{HT} ; then by Markov inequality, we have $\hat{\tau}_n^{HT-DIM} \xrightarrow{p} \tau^{HT}$.

¹⁰ By convention, we say that c_1 and c_2 are in the same consideration set if both refer to the same item.

A.2. Convergence of Hájek estimator ($\hat{\tau}_n^{HA-DIM}$)

Given n samples $\{(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i)\}_{i=1}^n$ collected from a creator-side randomization experiment, with treatment probability q , the Hájek estimator ($\hat{\tau}_n^{HA-DIM}$) is defined as:

$$\hat{\tau}_n^{HA-DIM} := \frac{\sum_{i=1}^n W_{i,k_i^*} Y_i}{\sum_{i=1}^n W_{i,k_i^*}} - \frac{\sum_{i=1}^n (1 - W_{i,k_i^*}) Y_i}{\sum_{i=1}^n (1 - W_{i,k_i^*})}, \quad (4)$$

We here show that $\hat{\tau}_n^{HA-DIM} \xrightarrow{p} \tau^{HA}$, where τ^{HA} is defined as:

$$\tau^{HA} := \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v^{1,q}, \mathbf{w} \sim \mathcal{B}(q)} [r_c \mid w_{k^*} = 1] - \sum_{c \in \mathcal{C}} \mathbb{E}_{v \sim \mathbb{P}_v^{0,q}, \mathbf{w} \sim \mathcal{B}(q)} [r_c \mid w_{k^*} = 0], \quad (6)$$

By Slutsky's theorem, the missing component to show the convergence of $\hat{\tau}_n^{HA-DIM}$ is the convergence of $\frac{1}{n} \sum_{i=1}^n W_{i,k_i^*}$ to $e^*(q)$. Then with the convergence of $\hat{\tau}_n^{HT-DIM} \xrightarrow{p} \tau^{HT}$, we finish the proof.

We now show the convergence of $\frac{1}{n} \sum_{i=1}^n W_{i,k_i}$. Note that by definition of $e^*(q)$, we have:

$$\mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n W_{i,k_i^*} \right] = e^*(q).$$

Now we characterize its variance.

$$\begin{aligned} \text{Var} \left(\frac{1}{n} \sum_{i=1}^n W_{i,k_i^*} \right) &= \mathbb{E} \left[\text{Var} \left(\frac{1}{n} \sum_{i=1}^n W_{i,k_i^*} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) \right] + \text{Var} \left[\mathbb{E} \left(\frac{1}{n} \sum_{i=1}^n W_{i,k_i^*} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) \right] \\ &= \frac{1}{n^2} \mathbb{E} \left[\sum_{i \neq j} \text{Cov} \left(W_{i,k_i^*}, W_{j,k_j^*} \mid \{(V_i, \vec{C}_i)\}_{i=1}^n \right) \right] + O \left(\frac{1}{n} \right) \\ &= O \left(\frac{\mathbb{E}[d^3]}{n} \right) + O \left(\frac{1}{n} \right) = o(1). \end{aligned} \quad (31)$$

By Markov inequality, we have

$$\frac{1}{n} \sum_{i=1}^n W_{i,k_i^*} \xrightarrow{p} e^*(q). \quad (32)$$

Similarly, we have:

$$\frac{1}{n} \sum_{i=1}^n (1 - W_{i,k_i^*}) \xrightarrow{p} 1 - e^*(q). \quad (33)$$

Note that,

$$\begin{aligned} \hat{\tau}_n^{HA-DIM} &= \frac{n}{\sum_{i=1}^n W_{i,k_i^*}} \frac{\sum_{i=1}^n W_{i,k_i^*} Y_i}{n} - \frac{n}{\sum_{i=1}^n (1 - W_{i,k_i^*})} \frac{\sum_{i=1}^n (1 - W_{i,k_i^*}) Y_i}{n} \\ \hat{\tau}_n^{HT-DIM} &= \frac{\sum_{i=1}^n W_{i,k_i^*} Y_i}{nq} - \frac{\sum_{i=1}^n (1 - W_{i,k_i^*}) Y_i}{n(1-q)}. \end{aligned}$$

By Slutsky's theorem, we have $\hat{\tau}_n^{HA-DIM} \xrightarrow{p} \tau^{HA}$.

Appendix B: Explicit Expressions

B.1. Counterfactual Policy Value

We write out the explicit expression of counterfactual policy values using the recommender choice model. Given a policy π , the policy value is

$$\begin{aligned}
Q(\pi) &= \mathbb{E}_{\mathbf{w} \sim \pi} \left[\sum_{c \in \mathcal{C}} r(c; w_c, \mathbf{w}_{-c}) \right] = \mathbb{E}_{\mathbf{w} \sim \pi} \left[\sum_{c \in \mathcal{C}} \mathbb{E}_v [y(v, c; w_c, \mathbf{w}_{-c})] \right] \\
&= \mathbb{E}_v \left[\mathbb{E}_{\mathbf{w} \sim \pi} \left[\sum_{c \in \mathcal{C}} y(v, c; w_c, \mathbf{w}_{-c}) \right] \right] = \mathbb{E}_{V_i, \vec{C}_i} \left[\mathbb{E}_{\vec{W}_i \sim \pi} \left[\sum_{c \in \mathcal{C}_i} y(V_i, c; \vec{W}_i) \right] \right] \\
&= \mathbb{E}_{V_i, \vec{C}_i, \vec{W}_i \sim \pi} \left[Y_i(V_i, \vec{C}_i; \vec{W}_i) \right] \\
&= \mathbb{E}_{V_i, \vec{C}_i, \vec{W}_i \sim \pi} \left[\sum_{k=1}^K z(V_i, C_{i,k}) \cdot p_k(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1) \right] \\
&= \mathbb{E}_{V_i, \vec{C}_i, \vec{W}_i \sim \pi} \left[\sum_{k=1}^K z(V_i, C_{i,k}) \cdot \frac{e^{s_0(V_i, C_{i,k}) + W_{i,k} \cdot s_1(V_i, C_{i,k})}}{\sum_{k'=1}^K e^{s_0(V_i, C_{i,k'}) + W_{i,k'} \cdot s_1(V_i, C_{i,k'})}} \right].
\end{aligned}$$

B.2. Debiased Estimator

We now write out the explicit debiased estimate $\psi(\cdot)$ for each observation, where we drop the subscript i and write the notation as $(V, \vec{C}, \vec{W}, k^*, Y)$ for succinctness. For estimated baseline score function $\hat{s}_0 \in \mathcal{F}_s$, treatment score uplift function $\hat{s}_1 \in \mathcal{F}_s$, and viewer response function $\hat{z} \in \mathcal{F}_z$, we have the estimated exposure probability be

$$p_k(V, \vec{C}, \vec{W}; \hat{s}_0, \hat{s}_1) = \frac{e^{\hat{s}_0(V, C_k) + W_k \cdot \hat{s}_1(V, C_k)}}{\sum_{k'=1}^K e^{\hat{s}_0(V, C_{k'}) + W_{k'} \cdot \hat{s}_1(V, C_{k'})}} \stackrel{(i)}{=} \frac{e^{\hat{s}_0(V, C_k) - \hat{s}_0(V, C_1) + W_k \cdot \hat{s}_1(V, C_k)}}{\sum_{k'=1}^K e^{\hat{s}_0(V, C_{k'}) - \hat{s}_0(V, C_1) + W_{k'} \cdot \hat{s}_1(V, C_{k'})}},$$

where the equation (i) is obtained by normalizing both the numerator and denominator by the exponential of the first content item's baseline score. In other words, for any baseline score vector $(\hat{s}_0(V, C_1), \hat{s}_0(V, C_2), \dots, \hat{s}_0(V, C_K))$, if we replace it by $(0, \hat{s}_0(V, C_2) - \hat{s}_0(V, C_1), \dots, \hat{s}_0(V, C_K) - \hat{s}_0(V, C_1))$, we will get the same exposure probability result.

This implies that, for any nuisance estimates $(\hat{s}_0, \hat{s}_1, \hat{z})$, the value $\mu(V, \vec{C}; \hat{s}_0, \hat{s}_1, \hat{z})$ can be fully recovered by the vectors

- $\vec{\hat{S}}_0 = (\hat{S}_{0,2}, \dots, \hat{S}_{0,K}) \in \mathbb{R}^{K-1}$ with $\hat{S}_{0,k} = \hat{s}_0(V, C_k) - \hat{s}_0(V, C_1)$;
- $\vec{\hat{S}}_1 = (\hat{S}_{1,1}, \dots, \hat{S}_{1,K}) \in \mathbb{R}^K$ with $\hat{S}_{1,k} = \hat{s}_1(V, C_k)$;
- $\vec{\hat{Z}} = (\hat{Z}_{1,1}, \dots, \hat{Z}_{1,K}) \in \mathbb{R}^K$ with $\hat{Z}_{1,k} = \hat{z}(V, C_k)$.

We thus abuse notations and write $\mu(V, \vec{C}; \tilde{s}_0, \tilde{s}_1, \tilde{z})$ as $\mu(\vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{\hat{Z}})$.

The bias of $\mu(\vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{\hat{Z}})$ comes from the deviation of $(\vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{\hat{Z}})$ to the true vectors $(\vec{S}_0, \vec{S}_1, \vec{Z})$ that are defined similarly under the true model (s_0, s_1, z) .

We next follow the double machine learning literature and use the outcome (k^*, Y) to correct the bias of $\mu(\vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{\hat{Z}})$ due to the bias of $(\vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{\hat{Z}})$ approximating the true $(\vec{S}_0, \vec{S}_1, \vec{Z})$.

Under the estimates $(\vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{\hat{Z}})$, we reload the loss function notation and write it as:

$$\ell(\vec{W}, k^*, Y; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{\hat{Z}}) = \ell_1(\vec{W}, k^*; \vec{\hat{S}}_0, \vec{\hat{S}}_1) + \ell_2(k^*, Y; \vec{\hat{Z}}),$$

where

$$\ell_1(\vec{W}, k^*; \vec{S}_0, \vec{S}_1) = \begin{cases} -W_1 \hat{S}_{1,1} + \log\left(e^{W_1 \hat{S}_{1,1}} + \sum_{i=2}^K e^{\hat{S}_{0,k} + W_k \hat{S}_{1,k}}\right) & \text{if } k^* = 1, \\ -(\hat{S}_{0,k} + W_k \hat{S}_{1,k}) + \log\left(e^{W_1 \hat{S}_{1,1}} + \sum_{i=2}^K e^{\hat{S}_{0,k} + W_k \hat{S}_{1,k}}\right) & \text{o.w.;} \end{cases}$$

and

$$\ell_2(k^*, Y; \vec{Z}) = (\hat{Z}_{k^*} - Y)^2.$$

We are now ready to introduce the debiased term ψ :

$$\psi(\vec{W}, k^*, Y; \vec{S}_0, \vec{S}_1, \vec{Z}, \hat{H}) = \mu(\vec{S}_0, \vec{S}_1, \vec{Z}) - \nabla \mu^T \hat{H}^{-1} \nabla \ell.$$

Above, $\nabla \mu$ and $\nabla \ell$ are gradients of μ and ℓ with respect to the nuisance estimates $(\vec{S}_0, \vec{S}_1, \vec{Z})$, and \hat{H} estimates the expected Hessian of ℓ regarding $(\vec{S}_0, \vec{S}_1, \vec{Z})$, where the expectation is taken with respect to the treatments \vec{W} that are assigned following the specified creator-side randomization. The explicit expressions for these derivatives are deferred to Appendix B.3.

B.3. Expressions of Gradients and Hessian

B.3.1. Gradient of μ . We have

$$\nabla \mu = \left(\frac{\partial \mu}{\partial \vec{S}_0}, \frac{\partial \mu}{\partial \vec{S}_1}, \frac{\partial \mu}{\partial \vec{Z}} \right)^T = \left(\frac{\partial \mu}{\partial \hat{S}_{0,2}}, \dots, \frac{\partial \mu}{\partial \hat{S}_{0,K}}, \frac{\partial \mu}{\partial \hat{S}_{1,1}}, \dots, \frac{\partial \mu}{\partial \hat{S}_{1,K}}, \frac{\partial \mu}{\partial \hat{Z}_1}, \dots, \frac{\partial \mu}{\partial \hat{Z}_K} \right)^T,$$

where

- for each $k = 2, \dots, K$,

$$\begin{aligned} \frac{\partial \mu}{\partial \hat{S}_{0,k}} &= P(k^* = k; \vec{S}_0, \vec{S}_1, \vec{W} \equiv 1) \left\{ \hat{Z}_k - \mathbb{E}[Y | \vec{S}_0, \vec{S}_1, \vec{Z}, \vec{W} \equiv 1] \right\} \\ &\quad - P(k^* = k; \vec{S}_0, \vec{S}_1, \vec{W} \equiv 0) \left\{ \hat{Z}_k - \mathbb{E}[Y | \vec{S}_0, \vec{S}_1, \vec{Z}, \vec{W} \equiv 0] \right\}. \end{aligned}$$

- for each $k = 1, \dots, K$,

$$\begin{aligned} \frac{\partial \mu}{\partial \hat{S}_{1,k}} &= P(k^* = k; \vec{S}_0, \vec{S}_1, \vec{W} \equiv 1) \left\{ \hat{Z}_k - \mathbb{E}[Y | \vec{S}_0, \vec{S}_1, \vec{Z}, \vec{W} \equiv 1] \right\}, \\ \frac{\partial \mu}{\partial \hat{Z}_k} &= P(k^* = k; \vec{S}_0, \vec{S}_1, \vec{W} \equiv 1) - P(k^* = k; \vec{S}_0, \vec{S}_1, \vec{W} \equiv 0). \end{aligned}$$

Above, $\langle \cdot, \cdot \rangle$ denotes inner product between two vectors.

B.3.2. Gradient of ℓ . We have

$$\nabla \ell = \left(\frac{\partial \ell_1}{\partial \vec{S}_0}, \frac{\partial \ell_1}{\partial \vec{S}_1}, \frac{\partial \ell_2}{\partial \vec{Z}} \right)^T = \left(\frac{\partial \ell_1}{\partial \hat{S}_{0,2}}, \dots, \frac{\partial \ell_1}{\partial \hat{S}_{0,K}}, \frac{\partial \ell_1}{\partial \hat{S}_{1,1}}, \dots, \frac{\partial \ell_1}{\partial \hat{S}_{1,K}}, \frac{\partial \ell_2}{\partial \hat{Z}_1}, \dots, \frac{\partial \ell_2}{\partial \hat{Z}_K} \right)^T,$$

where

- for each $k = 2, \dots, K$:

$$\frac{\partial \ell_1}{\partial \hat{S}_{0,k}} = P(k^* = k; \vec{S}_0, \vec{S}_1, \vec{W}) - \mathbb{I}[k^* = k].$$

- for each $k = 1, \dots, K$:

$$\begin{aligned} \frac{\partial \ell_1}{\partial \hat{S}_{1,k}} &= \mathbb{I}[W_k = 1] \left(P(k^* = k; \vec{S}_0, \vec{S}_1, \vec{W}) - \mathbb{I}[k^* = k] \right), \\ \frac{\partial \ell_2}{\partial \hat{Z}_k} &= \mathbb{I}[k^* = k] (\hat{Z}_k - Y). \end{aligned}$$

B.3.3. Hessian of ℓ . We have

$$\nabla^2 \ell = \begin{pmatrix} \frac{\partial^2 \ell_1}{\partial \hat{S}_0^2} & \frac{\partial^2 \ell_1}{\partial \hat{S}_0 \partial \hat{S}_1} & 0 \\ \frac{\partial^2 \ell_1}{\partial \hat{S}_1 \partial \hat{S}_0} & \frac{\partial^2 \ell_1}{\partial \hat{S}_1^2} & 0 \\ 0 & 0 & \frac{\partial^2 \ell_2}{\partial \hat{Z}^2} \end{pmatrix},$$

where the Hessian of loss function ℓ_1 follows;

$$\begin{aligned} \frac{\partial^2 \ell_1}{\partial \hat{S}_{0,k}^2} &= P(k^* = k; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}) \left(1 - P(k^* = k; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}) \right), & k \in \{2 : K\} \\ \frac{\partial^2 \ell_1}{\partial \hat{S}_{1,k}^2} &= W_k P(k^* = k; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}) \left(1 - P(k^* = k; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}) \right), & k \in \{1 : K\} \\ \frac{\partial^2 \ell_1}{\partial \hat{S}_{0,k} \partial \hat{S}_{1,k}} &= W_k P(k^* = k; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}) \left(1 - P(k^* = k; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}) \right), & k \in \{2 : K\} \\ \frac{\partial^2 \ell_1}{\partial \hat{S}_{0,k_1} \partial \hat{S}_{0,k_2}} &= -P(k^* = k_1; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}) P(k^* = k_2; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}), & k_1 \neq k_2 \in \{2 : K\} \\ \frac{\partial^2 \ell_1}{\partial \hat{S}_{0,k_1} \partial \hat{S}_{1,k_2}} &= -W_{k_2} P(k^* = k_1; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}) P(k^* = k_2; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}), & k_1 \neq k_2, k_1 \in \{2 : K\}, k_2 \in \{1 : K\} \\ \frac{\partial^2 \ell_1}{\partial \hat{S}_{1,k_1} \partial \hat{S}_{1,k_2}} &= -W_{k_1} W_{k_2} P(k^* = k_1; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}) P(k^* = k_2; \vec{\hat{S}}_0, \vec{\hat{S}}_1, \vec{W}), & k_1 \neq k_2 \in \{1 : K\}. \end{aligned}$$

and the Hessian of loss function ℓ_2 follows:

$$\begin{aligned} \frac{\partial \ell_2^2}{\partial \hat{Z}_k^2} &= \mathbb{I}[k^* = k], & k \in \{1 : K\} \\ \frac{\partial \ell_2^2}{\partial \hat{Z}_{k_1} \partial \hat{Z}_{k_2}} &= 0, & k_1 \neq k_2 \in \{1 : K\}. \end{aligned}$$

Appendix C: Supporting Results for Estimating Nuisance Components

C.1. Bounded Scores and Positive Exposure

ASSUMPTION 2 (**Bounded scores**). *There exists a universal constant C such that the true score functions are bounded: $\|s_0\|_\infty \leq C$ and $\|s_1\|_\infty \leq C$.*

This assumption plays the same role as the overlap condition in standard causal inference. It ensures that each item in the consideration set has a positive probability of being exposed under both treatment and control. This property guarantees that counterfactual exposure probabilities and nuisance functions are well defined.

LEMMA 1. *Under Assumption 2, there exists a universal constant $\delta > 0$ such that each content item in the consideration set has at least δ probability to be recommended:*

$$\mathbb{P}\left(k_i^* = k \mid V_i, \vec{C}_i, \vec{W}_i; s_0, s_1\right) \geq \delta, \quad \forall (V_i, \vec{C}_i, \vec{W}_i, k).$$

We have

$$\mathbb{P}\left(k_i^* = k \mid V_i, \vec{C}_i, \vec{W}_i; s_0, s_1\right) = \frac{e^{s_0(V_i, C_{i,k}) + W_{i,k} \cdot s_1(V_i, C_{i,k})}}{\sum_{k'=1}^K e^{s_0(V_i, C_{i,k'}) + W_{i,k'} \cdot s_1(V_i, C_{i,k'})}} \geq \frac{e^{-2C}}{\sum_{k'=1}^K e^{2C}} = e^{-2C - 2 \log(K)C},$$

where the inequality is by Assumption 2. Set $\delta = e^{-2C - 2 \log(K)C}$, concluding the proof.

C.2. Score Identification

We now prove Proposition 1 and show that (s_0, s_1, z) can be nonparametrically identified up to a location normalization $s_0(V_i, C_{i,1}) \equiv 0$. Under Assumption 2, there exists an $\delta > 0$ such that $p_k(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1) \geq \delta$ for all $V_i, \vec{C}_i, \vec{W}_i$ by Lemma 1.

Identifiability of (s_0, s_1) . First set $\vec{W}_i = \mathbf{0}$, which happens with probability of $(1 - q)^K$ by the creator randomization, we have s_0 be uniquely identified as:

$$s_0(V_i, C_{i,k}) = \log \frac{p_k(V_i, \vec{C}_i, \mathbf{0}; s_0, s_1)}{p_1(V_i, \vec{C}_i, \mathbf{0}; s_0, s_1)}, \quad (34)$$

with the normalization $s_0(V_i, C_{i,1}) \equiv 0$. Subsequently,

$$s_1(V_i, C_{i,1}) = \log \left\{ \left(\sum_{k=2}^K e^{s_0(V_i, C_{i,k})} \right) \cdot \frac{p_1(V_i, \vec{C}_i, (1, \mathbf{0}); s_0, s_1)}{1 - p_1(V_i, \vec{C}_i, (1, \mathbf{0}); s_0, s_1)} \right\},$$

where $(1, \mathbf{0})$ means that $W_{i,1} = 1$ and $W_{i,2:K} = \mathbf{0}$, which happens with probability $q(1 - q)^{K-1}$ by the experimental design. Finally, set $\vec{W}_i = \mathbf{1}$, which happens with probability q^K , we then have s_1 being uniquely identified as:

$$s_1(V_i, C_{i,k}) = \log \frac{p_k(V_i, \vec{C}_i, \mathbf{1}; s_0, s_1)}{p_1(V_i, \vec{C}_i, \mathbf{1}; s_0, s_1)} - s_0(V_i, C_{i,k}) + s_1(V_i, C_{i,1}). \quad (35)$$

Identifiability of z . For any (V_i, \vec{C}_i) , $k_i^* = k$ happens with probability at least δ by Lemma 1, we thus have $z(V_i, C_{i,k}) = \mathbb{E}[Y_i \mid k_i^* = k, V_i, \vec{C}_i]$ uniquely identified from the data distribution.

C.3. Score Convergence

ASSUMPTION 3. (i) Both viewer covariates and content covariates have compact, connected support, with each entry taking values in $[-1, 1]$. (ii) As functions of a viewer-content pair (v, c) , $s_0(\cdot, \cdot)$, $s_1(\cdot, \cdot)$, $z(\cdot, \cdot) \in \mathcal{W}^{p, \infty}([-1, 1]^{d_v + d_c})$, where for positive integers p and q , define the Hölder ball $\mathcal{W}^{p, \infty}([-1, 1]^q)$ of function $h: \mathbb{R}^q \rightarrow \mathbb{R}$ with smoothness $p \in \mathbb{N}_+$ as

$$\mathcal{W}^{p, \infty}([-1, 1]^q) := \left\{ h : \max_{r, |r| \leq p} \text{ess sup}_{v \in [-1, 1]^q} |D^r h(v)| \leq 1 \right\},$$

where $r = (r_1, \dots, r_q)$, $|r| = r_1 + \dots + r_q$ and $D^r h$ is the weak derivative.

We prove Proposition 2 by verifying conditions of Theorem 1 in Farrell et al. (2021a). We first verify Assumption 1 in Farrell et al. (2021a).

- The nuisance parameters (s_0, s_1, z) are nonparametrically identified by Proposition 1 and uniformly bounded by the model assumptions.

- Appendix B.3.2 shows that the loss function has bounded gradients with bounded z and outcome Y . Therefore it satisfies the Lipschitz condition.

- Write $\theta := (s_0, s_1, z)$ and $\tilde{\theta} := (\tilde{s}_0, \tilde{s}_1, \tilde{z})$ for any nuisance candidate $\tilde{s}_0, \tilde{s}_1, \tilde{z}$. By Taylor expansion, for some $\check{\theta}$ between $\tilde{\theta}$ and θ , we have

$$\begin{aligned} & \mathbb{E} \left[\ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; \tilde{\theta}) \mid V_i, \vec{C}_i \right] - \mathbb{E} \left[\ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; \theta) \mid V_i, \vec{C}_i \right] \\ &= \nabla_{\theta} \mathbb{E} \left[\ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; \theta) \mid V_i, \vec{C}_i \right] \|\tilde{\theta}(V_i, \vec{C}_i) - \theta(V_i, \vec{C}_i)\|_2 \\ & \quad + \frac{1}{2} \nabla_{\theta}^2 \mathbb{E} \left[\ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; \check{\theta}) \mid V_i, \vec{C}_i \right] \|\tilde{\theta}(V_i, \vec{C}_i) - \theta(V_i, \vec{C}_i)\|_2^2 \\ &= \frac{1}{2} \nabla_{\theta}^2 \mathbb{E} \left[\ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; \check{\theta}) \mid V_i, \vec{C}_i \right] \|\tilde{\theta}(V_i, \vec{C}_i) - \theta(V_i, \vec{C}_i)\|_2^2, \end{aligned}$$

where the last equality is by the first-order optimality of θ in ℓ . Lemma 2 shows that the Hessian $\nabla_{\theta}^2 \mathbb{E} \left[\ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; \check{\theta}) \mid V_i, \vec{C}_i \right]$ is universally invertible with bounded inverse. Therefore, there exists a $c_1 > 0$ such that:

$$\mathbb{E} \left[\ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; \tilde{\theta}) \mid V_i, \vec{C}_i \right] - \mathbb{E} \left[\ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; \theta) \mid V_i, \vec{C}_i \right] \geq c_1 \|\tilde{\theta}(V_i, \vec{C}_i) - \theta(V_i, \vec{C}_i)\|_2^2.$$

On the other hand, Appendix B.3.3 writes out the explicit form of the Hessian, which has bounded entries, implying bounded eigenvalues. Therefore, there exists a $c_2 > 0$ such that:

$$\mathbb{E} \left[\ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; \tilde{\theta}) \mid V_i, \vec{C}_i \right] - \mathbb{E} \left[\ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; \theta) \mid V_i, \vec{C}_i \right] \leq c_2 \|\tilde{\theta}(V_i, \vec{C}_i) - \theta(V_i, \vec{C}_i)\|_2^2.$$

We next verify Assumption 2 in Farrell et al. (2021a). The random variable are bounded with compact, connected support by our assumptions. The nuisance parameters are also uniformly bounded. Together with smoothness Assumption 3, we have Assumption 2 in Farrell et al. (2021a) hold. Then we apply Theorem 1 in Farrell et al. (2021a) and yield Proposition 2.

Appendix D: Supporting Results for the Debiased Estimator

D.1. Invertible Hessian

LEMMA 2. *Suppose Assumptions 2 hold. Suppose that the estimated scores are bounded: $\|\hat{s}_0\|_\infty \leq C$ and $\|\hat{s}_1\|_\infty \leq C$ for the C in Assumption 2. Under our modeling framework, the expected hessian $H = \mathbb{E} \left[\nabla^2 \ell | V, \vec{C} \right]$ is universally invertible with bounded inverse.*

Note that H is a two-block diagonal matrix with the first diagonal block being $H_1 := \mathbb{E} \left[\nabla^2 \ell_1 | V, \vec{C} \right]$ and the second diagonal block being $H_2 := \mathbb{E} \left[\nabla^2 \ell_2 | V, \vec{C} \right]$. It suffices to show that both H_1 and H_2 are universally invertible with bounded inverses.

Regularity of H_1 . Consider the sample (V, \vec{C}, \vec{W}) and estimated nuisance \hat{s}_0, \hat{s}_1 . For notation convenience, we use p_k to represent $P(k^* = k; V, \vec{C}, \vec{W}, \hat{s}_0, \hat{s}_1)$. Define the vector

$$\beta = \begin{pmatrix} p_2 \\ \vdots \\ p_K \\ W_1 p_1 \\ \vdots \\ W_k p_K \end{pmatrix} \in \mathbb{R}^{2K-1},$$

the matrix $A \in \mathbb{R}^{(2K-1) \times (2K-1)}$ with its (K, K) -th entry being $\mathbf{1}[W_1 = 1]p_1$ and others zero and matrix

$$B = \begin{pmatrix} \text{diag}(\sqrt{p_2}, \dots, \sqrt{p_K}) \\ 0, \dots, 0, \\ \text{diag}(\sqrt{p_2}W_2, \dots, \sqrt{p_K}W_k) \end{pmatrix} \in \mathbb{R}^{(2K-1) \times (2K-1)},$$

We thus have

$$H_1 = \mathbb{E} \left[A + BB^T - \beta\beta^T \mid V, \vec{C} \right].$$

We now show that H_1 is positive definite. For any vector $x = (x_1, \dots, x_{2K-1})$ and any treatment assignment \vec{W} , we have

$$\begin{aligned} & x^T (A + BB^T - \beta\beta^T) x & (36) \\ &= p_1 W_1 x_K^2 + \sum_{j=1}^{K-1} p_{j+1} (x_j + W_{j+1} x_{K+j})^2 - \left(p_1 W_1 x_K + \sum_{j=1}^{K-1} p_{j+1} (x_j + W_{j+1} x_{K+j}) \right)^2 \\ &= \left(p_1 W_1 x_K^2 + \sum_{j=1}^{K-1} p_{j+1} (x_j + W_{j+1} x_{K+j})^2 \right) \left(\sum_{j=1}^K p_j \right) \\ &\quad - \left(p_1 W_1 x_K + \sum_{j=1}^{K-1} p_{j+1} (x_j + W_{j+1} x_{K+j}) \right)^2 \stackrel{(i)}{\geq} 0, \end{aligned}$$

where (i) is by Cauchy-Swartz inequality. That is, for any treatment assignment, the Hessian is positive semi-definite, and thus H_1 (the expected Hessian) is positive semi-definite.

We next consider specific assignments of \vec{W} to lower bound the smallest eigenvalue of H_1 . Also note that under the assumption of the boundedness of score functions, there exists a constant $\delta > 0$, such that for any $(U, \vec{V}, \vec{W}, k^*)$, we have $P(k^* \mid U, \vec{V}, \vec{W}, \hat{s}_0, \hat{s}_1) \geq \delta$.

Consider $\vec{W}^{(a)} = (0, \dots, 0)$. Then we have Equation (36) instantiate as

$$\begin{aligned} x^T(A + BB^T - \beta\beta^T)x &= \sum_{j=1}^{K-1} p_{j+1}^{(a)} x_j^2 - \left(\sum_{j=1}^{K-1} p_{j+1}^{(a)} x_j \right)^2 \\ &= p_1^{(a)} \sum_{j=1}^{K-1} p_{j+1}^{(a)} x_j^2 + \left(\sum_{j=1}^{K-1} p_{j+1}^{(a)} \right) \sum_{j=1}^{K-1} p_{j+1}^{(a)} x_j^2 - \left(\sum_{j=1}^{K-1} p_{j+1}^{(a)} x_j \right)^2 \geq p_1 \sum_{j=1}^{K-1} p_{j+1}^{(a)} x_j^2 \geq \delta^2 \sum_{j=1}^{K-1} x_j^2, \end{aligned} \quad (37)$$

where the last inequality is due to Cauchy-Schwartz inequality.

Consider $\vec{W}^{(b)} = (0, 1, \dots, 1)$. Then we have Equation (36) instantiate as

$$\begin{aligned} x^T(A + BB^T - \beta\beta^T)x &= \sum_{j=1}^{K-1} p_{j+1}^{(b)} (x_j + x_{K+j})^2 - \left(\sum_{j=1}^{K-1} p_{j+1}^{(b)} (x_j + x_{K+j}) \right)^2 \\ &= p_1^{(b)} \sum_{j=1}^{K-1} p_{j+1}^{(b)} (x_j + x_{K+j})^2 + \left(\sum_{j=1}^{K-1} p_{j+1}^{(b)} \right) \sum_{j=1}^{K-1} p_{j+1}^{(b)} (x_j + x_{K+j})^2 - \left(\sum_{j=1}^{K-1} p_{j+1}^{(b)} (x_j + x_{K+j}) \right)^2 \\ &\stackrel{(i)}{\geq} p_1^{(b)} \sum_{j=1}^{K-1} p_{j+1}^{(b)} (x_j + x_{K+j})^2 \geq \delta^2 \sum_{j=1}^{K-1} (x_j + x_{K+j})^2 = \delta^2 \sum_{j=1}^{K-1} (x_j^2 + x_{K+j}^2 + 2x_j x_{K+j}) \\ &\stackrel{(ii)}{\geq} \delta^2 \sum_{j=1}^{K-1} (x_j^2 + x_{K+j}^2 - \frac{2}{3}x_{K+j}^2 - \frac{3}{2}x_j^2) = \delta^2 \sum_{j=1}^{K-1} \left(-\frac{1}{2}x_j^2 + \frac{1}{3}x_{K+j}^2 \right), \end{aligned} \quad (38)$$

where both (i) and (ii) are due to Cauchy-Schwartz inequality.

Consider $\vec{W}^{(c)} = (1, 0, \dots, 0)$. We have Equation (36) instantiate as

$$x^T(A + BB^T - \beta\beta^T)x = p_1^{(c)} x_K^2 + \sum_{j=1}^{K-1} p_{j+1}^{(c)} x_j^2 - \left(p_1^{(c)} x_K + \sum_{j=1}^{K-1} p_{j+1}^{(c)} x_j \right)^2 \quad (39)$$

Note that under creator-side randomization, each \vec{W} has assignment probability bounded away from zero, and we denote this lower bound as $\eta > 0$. Also recall that each item exposure probability is lower bounded by $\delta > 0$, with $K\delta \leq 1$. Combining Equation (37), (38), (39), with the fact that the Hessian for any \vec{W} is semi-definite, we have

$$\begin{aligned} x^T H_1 x &= \mathbb{E}[x^T(A + BB^T - \beta\beta^T) | V, \vec{C}] \\ &\geq P(\vec{W} = \vec{W}^{(a)}) \delta^2 \sum_{j=1}^{K-1} x_j^2 + P(\vec{W} = \vec{W}^{(b)}) \delta^2 \sum_{j=1}^{K-1} \left\{ -\frac{1}{2}x_j^2 + \frac{1}{3}x_{K+j}^2 \right\} \\ &\quad + P(\vec{W} = \vec{W}^{(c)}) \left\{ p_1^{(c)} x_K^2 + \sum_{j=1}^{K-1} p_{j+1}^{(c)} x_j^2 - \left(p_1^{(c)} x_K + \sum_{j=1}^{K-1} p_{j+1}^{(c)} x_j \right)^2 \right\} \\ &\geq \eta \delta^2 \sum_{j=1}^{K-1} x_j^2 + \frac{\eta \delta^2}{2} \sum_{j=1}^{K-1} \left\{ -\frac{1}{2}x_j^2 + \frac{1}{3}x_{K+j}^2 \right\} + \frac{\eta \delta^2}{2(1 - (K-1)\delta)^2} \left\{ p_1^{(c)} x_K^2 + \sum_{j=1}^{K-1} p_{j+1}^{(c)} x_j^2 - \left(p_1^{(c)} x_K + \sum_{j=1}^{K-1} p_{j+1}^{(c)} x_j \right)^2 \right\} \\ &= \frac{\eta \delta^2}{2} \sum_{j=1}^{K-1} \left\{ \frac{1}{2}x_j^2 + \frac{1}{3}x_{K+j}^2 \right\} + I, \end{aligned} \quad (40)$$

where $I = \frac{\eta \delta^2}{2(1 - (K-1)\delta)^2} \left\{ p_1^{(c)} x_K^2 + \sum_{j=1}^{K-1} (p_{j+1}^{(c)} + (1 - (K-1)\delta)^2) x_j^2 - \left(p_1^{(c)} x_K + \sum_{j=1}^{K-1} p_{j+1}^{(c)} x_j \right)^2 \right\}$. We next lower bound term I.

$$I \geq \frac{\eta \delta^2}{2(1 - (K-1)\delta)^2} \left\{ p_1^{(c)} x_K^2 + \sum_{j=1}^{K-1} (1 + p_1^{(c)}) p_{j+1}^{(c)} x_j^2 - \left(p_1^{(c)} x_K + \sum_{j=1}^{K-1} p_{j+1}^{(c)} x_j \right)^2 \right\}$$

$$\begin{aligned}
&= \frac{\eta\delta^2}{2(1-(K-1)\delta)^2} \left\{ \left(p_1^{(c)} x_K^2 + \sum_{j=1}^{K-1} (1+p_1^{(c)}) p_{j+1}^{(c)} x_j^2 \right) \left(p_1^{(c)} + \sum_{j=1}^{K-1} \frac{p_{j+1}^{(c)}}{1+p_1^{(c)}} + \frac{p_1^{(c)} - (p_1^{(c)})^2}{p_1^{(c)} + 1} \right) - \left(p_1^{(c)} x_K + \sum_{j=1}^{K-1} p_{j+1}^{(c)} x_j \right)^2 \right\} \\
&\stackrel{(i)}{\geq} \frac{\eta\delta^2}{2(1-(K-1)\delta)^2} \left(p_1^{(c)} x_K^2 + \sum_{j=1}^{K-1} (1+p_1^{(c)}) p_{j+1}^{(c)} x_j^2 \right) \left(\frac{p_1^{(c)} - (p_1^{(c)})^2}{p_1^{(c)} + 1} \right) \\
&\geq \frac{\eta\delta^2}{2(1-(K-1)\delta)^2} \left(p_1^{(c)} x_K^2 + \sum_{j=1}^{K-1} (1+p_1^{(c)}) p_{j+1}^{(c)} x_j^2 \right) \frac{p_1^{(c)}(1-p_1^{(c)})}{2} \\
&\geq \frac{\eta\delta^2 (p_1^{(c)})^2 (1-p_1^{(c)})}{4(1-(K-1)\delta)^2} x_K^2 \geq \frac{\eta\delta^5}{4(1-(K-1)\delta)^2} x_K^2, \tag{41}
\end{aligned}$$

where (i) is by Cauchy-Schwarz. Putting Equation (40) and (41) together, for any $x = (x_1, \dots, x_{2K-1})$, we have

$$x^T H_1 x \geq \frac{\eta\delta^2}{2} \sum_{j=1}^{K-1} \left\{ \frac{1}{2} x_j^2 + \frac{1}{3} x_{K+j}^2 \right\} + \frac{\eta\delta^5}{4(1-(K-1)\delta)^2} x_K^2. \tag{42}$$

Therefore, H_1 has the smallest eigenvalue of greater than or equal to $\frac{\eta\delta^2}{2} \min\left(\frac{1}{3}, \frac{\delta^3}{2(1-(K-1)\delta)^2}\right)$ and thus is invertible with bounded inverse.

Regularity of H_2 . We have H_2 is a diagonal matrix with its k -th diagonal entry being

$$H_2(k, k) = \mathbb{E} \left[\mathbb{E} \left[P(k^* = k \mid U, \vec{V}, \vec{W}, \hat{s}_0, \hat{s}_1) \mid U, \vec{V} \right] \right] \geq \delta.$$

As a result, $H_2 \geq \delta \cdot I$ and thus is invertible with bounded inverse.

D.2. Universal Neyman Orthogonality

We show that the debiased estimate ψ , defined in Equation (18), satisfies the universal Neyman orthogonality (Chernozhukov et al. 2019, Foster and Syrgkanis 2023). This property means that the nuisance estimation error only has a second order effect on the debiased estimate – a key property for achieving the asymptotic normality of the debiased estimator.

PROPOSITION 4 (Universal Orthogonality). *The debiased estimator ψ defined in Equation (18) is universally orthogonal with respect to the nuisances in the sense that, for any nuisance components $(\tilde{s}_0, \tilde{s}_1, \tilde{z}, \tilde{H})$,*

$$\mathbb{E}[\nabla\psi(V, \vec{C}, \vec{W}, k^*, Y; \tilde{s}_0 = s_0, \tilde{s}_1 = s_1, \tilde{z} = z, \tilde{H} = H) \mid V, \vec{C}] = 0,$$

where $(V, \vec{C}, \vec{W}, k^*, Y)$ is sampled from the creator-side randomization experiment, and $\nabla\psi$ is the gradient with respect to the nuisances.

Recall that

$$\begin{aligned}
\psi(V, \vec{C}, \vec{W}, k^*, Y; \tilde{s}_0, \tilde{s}_1, \tilde{z}, \tilde{H}) &= \mu(V, \vec{C}; \tilde{s}_0, \tilde{s}_1, \tilde{z}) \\
&\quad - \nabla\mu(V, \vec{C}; \tilde{s}_0, \tilde{s}_1, \tilde{z})^T \tilde{H}(V, \vec{C}; \tilde{s}_0, \tilde{s}_1, \tilde{z})^{-1} \nabla\ell(V, \vec{C}, \vec{W}, k^*, Y; \tilde{s}_0, \hat{s}_1, \tilde{z}),
\end{aligned}$$

Let $\tilde{h} = (\tilde{s}_0, \tilde{s}_1, \tilde{z}, \tilde{H}^{-1})$ with the ground truth $h = (s_0, s_1, z, H^{-1})$. It suffices to show that

$$\mathbb{E}[\nabla_{\tilde{h}}\psi(V, \vec{C}, \vec{W}, k^*, Y; \tilde{h} = h) \mid V, \vec{C}] = 0.$$

We have

$$\begin{aligned} \frac{\partial \psi(\tilde{h} = h)}{\partial(\tilde{s}_0, \tilde{s}_1, \tilde{z})} &= \nabla \mu(V, \vec{C}; s_0, s_1, z) - \nabla^2 \mu(V, \vec{C}; s_0, s_1, z) H(V, \vec{C}; s_0, s_1, z)^{-1} \nabla \ell(V, \vec{C}, \vec{W}, k^*, Y; s_0, s_1, z) \\ &\quad - \nabla \mu(V, \vec{C}; s_0, s_1, z) H(V, \vec{C}; s_0, s_1, z)^{-1} \nabla^2 \ell(V, \vec{C}, \vec{W}, k^*, Y; s_0, s_1, z). \end{aligned}$$

Then taking the expectation with respect to (\vec{W}, k^*, Y) , we have

$$\begin{aligned} \mathbb{E} \left[\frac{\partial \psi(\tilde{h} = h)}{\partial(\tilde{s}_0, \tilde{s}_1, \tilde{z})} \mid V, \vec{C} \right] &= \nabla \mu(V, \vec{C}; s_0, s_1, z) \left(I - H(V, \vec{C}; s_0, s_1, z)^{-1} \mathbb{E} \left[\nabla^2 \ell(V, \vec{C}, \vec{W}, k^*, Y; s_0, s_1, z) \mid V, \vec{C} \right] \right) \\ &\quad - \nabla^2 \mu(V, \vec{C}; s_0, s_1, z) H(V, \vec{C}; s_0, s_1, z)^{-1} \mathbb{E} \left[\nabla \ell(V, \vec{C}, \vec{W}, k^*, Y; s_0, s_1, z) \mid V, \vec{C} \right] \\ &\stackrel{(i)}{=} \nabla \mu(V, \vec{C}; s_0, s_1, z) \left(I - H(V, \vec{C}; s_0, s_1, z)^{-1} H(V, \vec{C}; s_0, s_1, z) \right) = 0, \end{aligned}$$

where (i) is because the ground truth (s_0, s_1, z) satisfies the first-order optimality condition of loss function ℓ . Similarly, we have

$$\mathbb{E} \left[\frac{\partial \psi(\tilde{h} = h)}{\partial(\tilde{H}^{-1})} \mid V, \vec{C} \right] = -\mathbb{E} \left[\nabla \ell(V, \vec{C}, \vec{W}, k^*, Y; s_0, s_1, z) \mid V, \vec{C} \right] \nabla \mu(V, \vec{C}; s_0, s_1, z)^T = 0.$$

Appendix E: Proof of the Asymptotics of the Debiased Estimator

We prove the below theorem to show the asymptotic regularity of the debiased estimator.

THEOREM 2 (Asymptotic Normality of the Debiased Estimator). *Suppose Assumptions 1 & 2 hold. Assume that the data-generating process follows the recommender choice model in Equation (9) and the viewer response model in Equation (10). Suppose that the estimated nuisance functions are all bounded by the constant C in Assumption 2 and satisfy the convergence rate: $\|\hat{s}_0 - s_0\|_{L_2} + \|\hat{s}_1 - s_1\|_{L_2} + \|\hat{z} - z\|_{L_2} = o(n^{-1/4})$. Then the debiased estimator $\hat{\tau}_n^{DB}$ defined in Equation (19) is \sqrt{n} -consistent and asymptotically normal with:*

$$n^{1/2} (\hat{\tau}_n^{DB} - \tau) / (\widehat{V}_n^{DB})^{1/2} \Rightarrow \mathcal{N}(0, 1),$$

implying that $\hat{\tau}_n^{DB} - \tau = O_p(n^{-1/2})$.

Define the oracle estimator $\tilde{\tau}_n^{DB}$ as

$$\tilde{\tau}_n^{DB} = \frac{1}{n} \sum_{i=1}^n \psi \left(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; s_0, s_1, z, H \right), \quad (43)$$

with that

$$\begin{aligned} \psi(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; s_0, s_1, z, H) &= \mu(V_i, \vec{C}_i; s_0, s_1, z) \\ &\quad - \nabla \mu(V_i, \vec{C}_i; s_0, s_1, z)^T H(V_i, \vec{C}_i; s_0, s_1, z)^{-1} \nabla \ell(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; s_0, s_1, z). \end{aligned} \quad (44)$$

Note that

$$\mathbb{E}[\psi(V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i; s_0, s_1, z, H) \mid V_i, \vec{C}_i] = \mu \left(V_i, \vec{C}_i; s_0, s_1, z \right), \quad (45)$$

by the first order condition of (s_0, s_1, z) in ℓ when conditioning on $(V_i, \vec{C}_i, \vec{W}_i)$.

Our proof follows three main steps:

- (i) Show $\hat{\tau}_n^{DB}$ and $\tilde{\tau}_n^{DB}$ have similar asymptotic behavior.
- (ii) Show asymptotic normality of $\tilde{\tau}_n^{DB}$.
- (iii) Show asymptotically consistent variance estimator \widehat{V}_n .

E.1. Step I: Show $\hat{\tau}_n^{DB}$ and $\tilde{\tau}_n^{DB}$ have similar asymptotic behavior.

This part's proof is largely motivated by the proof pattern for Theorem 3.1 in Chernozhukov et al. (2018), except that we need to additionally deal with the correlation among samples due to the shared items in their consideration sets. For notation convenience, define $\theta_0 = (s_0, s_1, z, H)$ and $\hat{\theta}_0 = (\hat{s}_0, \hat{s}_1, \hat{z}, \hat{H})$. Write $Z_i := (V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i)$. Note that

$$\mathbb{E} \left[\psi(Z_i; \theta_0) \mid V_i, \vec{C}_i, \vec{W}_i \right] \stackrel{(i)}{=} \mathbb{E} \left[\mu \left(V_i, \vec{C}_i; s_0, s_1, z \right) \mid V_i, \vec{C}_i, \vec{W}_i \right] = \mu \left(V_i, \vec{C}_i; s_0, s_1, z \right), \quad (46)$$

where (i) is by the first-order optimality of (s_0, s_1, z) in ℓ . We have

$$\sqrt{n} |\hat{\tau}_n^{DB} - \tilde{\tau}_n^{DB}| \leq I_1 + I_2, \quad (47)$$

where

$$\begin{aligned} I_1 &:= \left\| \frac{1}{\sqrt{n}} \sum_i \left(\psi(Z_i; \hat{\theta}_0) - \mathbb{E}[\psi(Z_i; \hat{\theta}_0)] \right) - \frac{1}{\sqrt{n}} \sum_i \left(\psi(Z_i; \theta_0) - \mathbb{E}[\psi(Z_i; \theta_0)] \right) \right\|, \\ I_2 &:= \left\| \sqrt{n} \left(\mathbb{E}[\psi(Z_i; \hat{\theta}_0)] - \mathbb{E}[\psi(Z_i; \theta_0)] \right) \right\|. \end{aligned}$$

We now bound I_1 and I_2 respectively.

Bounding I_1 . Write $Q_i := \left(\psi(Z_i; \hat{\theta}_0) - \mathbb{E}[\psi(Z_i; \hat{\theta}_0)] \right) - \left(\psi(Z_i; \theta_0) - \mathbb{E}[\psi(Z_i; \theta_0)] \right)$. We have

$$\begin{aligned} \mathbb{E}[I_1^2] &= \frac{1}{n} \mathbb{E} \left[\left(\sum_i Q_i \right)^2 \right] = \frac{1}{n} \sum_i \mathbb{E}[Q_i^2] + \frac{1}{n} \sum_i \sum_{j \neq i} \mathbf{1}(\vec{W}_i \cap \vec{W}_j \neq \emptyset) \mathbb{E}[Q_i Q_j] \\ &= \frac{O(a_n)}{n} \sum_i \mathbb{E}[Q_i^2] = O\left(\frac{a_n}{n}\right) \sum_i \mathbb{E} \left[\left\{ \left(\psi(Z_i; \hat{\theta}_0) - \mathbb{E}[\psi(Z_i; \hat{\theta}_0)] \right) - \left(\psi(Z_i; \theta_0) - \mathbb{E}[\psi(Z_i; \theta_0)] \right) \right\}^2 \right] \end{aligned} \quad (48)$$

$$\leq O\left(\frac{a_n}{n}\right) \sum_i \mathbb{E} \left[\left(\psi(Z_i; \hat{\theta}_0) - \psi(Z_i; \theta_0) \right)^2 \right] = O(a_n \epsilon_n^2). \quad (49)$$

Therefore, by Markov inequality, we have $I_1 = O_p(a_n^{1/2} \epsilon_n)$.

Bounding I_2 . Define the function

$$f(r) := \mathbb{E}[\psi(Z_i; \theta_0 + r(\hat{\theta}_0 - \theta_0))] - \mathbb{E}[\psi(Z_i; \theta_0)], \quad r \in (0, 1). \quad (50)$$

By Taylor expansion, we have

$$f(1) = f(0) + f'(0) + f''(\tilde{r})/2, \quad \text{for some } \tilde{r} \in (0, 1). \quad (51)$$

Note that by the Neyman orthogonality (shown in [D.2](#)), we have $f'(0) = 0$. With the bounded inverse of Hessian, we have

$$\mathbb{E}[\|f''(\tilde{r})\|] \leq \sup_{r \in (0, 1)} \|f''(r)\| = O(\epsilon_n^2). \quad (52)$$

We have

$$I_{2,k} = O_p(\sqrt{n} \epsilon_n^2). \quad (53)$$

Combining the bound for I_1 and I_2 , we have

$$\sqrt{n} (\hat{\tau}_n^{DB} - \tilde{\tau}_n^{DB}) = O_p(a_n^{1/2} \epsilon_n + n^{1/2} \epsilon_n^2) = o_p(1), \quad (54)$$

with $a_n = O(n^{1/4})$ and $\epsilon_n = o(n^{-1/4})$.

E.2. Step II: Show asymptotic normality of $\tilde{\tau}_n^{DB}$.

Recall the data generating process: when a new viewer V_i arrives, the back-end retrieval system firstly generates the consideration set \vec{C}_i . For content items that have shown in previous samples, the treatment status remains unchanged. For content items that haven't appear, we sample the treatment status from i.i.d. Bernoulli randomized trials. This procedure constructs the treatment collection \vec{W}_i . Then given $(V_i, \vec{C}_i, \vec{W}_i)$, the recommender chooses item k_i^* to expose, yielding the viewer outcome Y_i and the observation tuple $Z_i := (V_i, \vec{C}_i, \vec{W}_i, k_i^*, Y_i)$. We now apply the martingale theorem to analyze the asymptotic behavior of $\tilde{\tau}_n$. Denote the σ -field $\mathcal{F}_i := \sigma(Z_1, \dots, Z_i)$. We have that

$$\begin{aligned} \mathbb{E}[\psi(Z_i; \theta_0) | \mathcal{F}_{i-1}] &= \mathbb{E} \left[\mathbb{E} \left[\psi(Z_i; \theta_0) | V_i, \vec{C}_i, \vec{W}_i \right] | \mathcal{F}_{i-1} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\mu(V_i, \vec{C}_i; \theta_0) | V_i, \vec{C}_i, \vec{W}_i \right] | \mathcal{F}_{i-1} \right] = \mathbb{E} \left[\mu(V_i, \vec{C}_i; \theta_0) \right] = \tau. \end{aligned}$$

Therefore $\{\psi(Z_i; \theta_0) - \tau\}$ forms a martingale difference sequence with respect to filtration $\{\mathcal{F}_i\}$. We now apply the following result from [Hall and Heyde \(2014\)](#).

PROPOSITION 5 (Martingale Central Limit Theorem, Theorem 3.2, Hall and Heyde (2014)).

Let $\{\xi_i\}$ be a martingale difference sequence with respect to filtration $\{\mathcal{F}_i\}$, and let η^2 be an a.s. finite random variable. Suppose that:

$$\max_i |\xi_i| \xrightarrow{p} 0, \quad (55)$$

$$\sum_i \xi_i^2 \xrightarrow{p} \eta^2, \quad (56)$$

$$\mathbb{E}[\max_i \xi_i^2] \text{ is bounded.} \quad (57)$$

Then $\sum_{i=1}^n X_i \xrightarrow{d} Z$ (stably), where the random variable Z has characteristic function $\mathbb{E}[\exp(-\frac{1}{2}\eta^2 t^2)]$.

Now let's verify the above conditions for $\sqrt{n}(\hat{\tau} - \tau)$. Define

$$\xi_i = \frac{1}{\sqrt{n}} \{\psi(Z_i; \theta_0) - \tau\}.$$

We have $\mathbb{E}[\xi_i | \mathcal{F}_{i-1}] = 0$. By the bounded inverse of Hessian, we have $|\xi_i| = O(n^{-1/2})$, implying Equation (55) and (57). Section E.3 shows that:

$$\tilde{V}_n := \sum_i \xi_i^2 = \frac{1}{n} \sum_i (\psi_i(Z_i; \theta_0) - \tau)^2 \xrightarrow{p} \eta^2,$$

where $\eta^2 := \text{Var}(\psi_i(Z_i; \theta_0))$ that is a finite number, verifying Equation (56). We thus have

$$\sqrt{n}(\hat{\tau}_n^{DB} - \tau) = \sum_i \xi_i \xrightarrow{d} \mathcal{N}(0, \eta^2), \quad (58)$$

where z is a Gaussian random variable $\mathcal{N}(0, \eta^2)$ that has characteristic function $\mathbb{E}[\exp(-\frac{1}{2}\eta^2 t^2)]$. Connecting Equation (58) with Equation (54), by Slutsky's theorem, we have

$$\sqrt{n}(\hat{\tau}_n^{DB} - \tau) = \sqrt{n}(\hat{\tau}_n^{DB} - \tilde{\tau}_n^{DB}) + \sqrt{n}(\tilde{\tau}_n^{DB} - \tau) \xrightarrow{d} z, \quad (59)$$

yielding that $\sqrt{n}(\hat{\tau}_n^{DB} - \tau)/\eta$ is asymptotically normal and converges to $\mathcal{N}(0, 1)$.

E.3. Step III: Variance convergence and connecting \tilde{V}_n to \hat{V}_n .

We start by showing that \tilde{V}_n converges in probability to $\eta^2 := \text{Var}(\psi_i(Z_i; \theta_0))$. We first argue that $\mathbb{E}[(\psi_i(Z_i; \theta_0) - \tau)^2]$ has the same value, denoted as η^2 . Note that $\mathbb{E}[\psi_i(Z_i; \theta_0)] = \mathbb{E}[\mathbb{E}[\psi_i(Z_i; \theta_0) | \mathcal{F}_{i-1}]] = \tau$. Therefore $\mathbb{E}[(\psi_i(Z_i; \theta_0) - \tau)^2]$ represents the variance of $\psi_i(Z_i; \theta_0)$. Given n samples, let \vec{w} denote the treatment status of all items involved in the n samples. We have that each entry of \mathbf{w} is i.i.d. sampled from Bernoulli distribution with the treated probability being q , for which we abuse notation and denote as $\mathbf{w} \sim \mathcal{B}(q)$. Therefore,

$$\begin{aligned} \text{Var}(\psi_i(Z_i; \theta_0)) &:= \mathbb{E}[(\psi_i(Z_i; \theta_0) - \tau)^2] = \mathbb{E}_{\mathbf{w} \sim \mathcal{B}(q)} [\mathbb{E}[(\psi_i(Z_i; \theta_0) - \tau)^2 | \mathbf{w}]] \\ &= \mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} [\mathbb{E}[(\psi_i(Z_i; \theta_0) - \tau)^2 | \vec{W}_i]] \\ &= \mathbb{E}_{(V_i, \vec{C}_i), \vec{W}_i \sim \mathcal{B}(q)} [\mathbb{E}[(\psi_i(Z_i; \theta_0) - \tau)^2 | V_i, \vec{C}_i, \vec{W}_i]]. \end{aligned}$$

Recall that $\{(V_i, \vec{C}_i)\}_{i=1}^n$ are sampled i.i.d. from i.i.d. viewer queries, we thus obtain that $\text{Var}(\psi_i(Z_i; \theta_0))$ across i have the same value, denoted as η^2 . Similarly, we have that for any m -moment $\mathbb{E}[(\psi_i(Z_i; \theta_0) - \tau)^m]$,

they have the same bounded value across sample i , where the boundedness comes from the boundedness of ψ_i .

We thus have $\mathbb{E}[\tilde{V}_n] = \eta^2$. We proceed to show $\tilde{V}_n \xrightarrow{p} \eta^2$. By Markov inequality, given any $\epsilon > 0$, we have

$$P\left(|\tilde{V}_n - \eta^2| \geq \epsilon\right) \leq \frac{\text{Var}(\tilde{V}_n)}{\epsilon^2}. \quad (60)$$

Note that

$$\begin{aligned} \text{Var}(\tilde{V}_n) &= \frac{\sum_{i=1}^n \mathbb{E}[(\psi_i(Z_i; \theta_0) - \tau)^4]}{n^2} + \frac{\sum_{i \neq j} \mathbf{1}(\vec{W}_i \cap \vec{W}_j \neq \emptyset) \text{Cov}\{(\psi_i(Z_i; \theta_0) - \tau)^2, (\psi_j(Z_j; \theta_0) - \tau)^2\}}{n^2} \\ &\stackrel{(i)}{=} O\left(\frac{n}{n^2}\right) + \frac{na_n}{n^2} = O\left(\frac{a_n}{n}\right) = o(n^{-3/4}) = o(1), \end{aligned}$$

where (i) uses the boundedness of $\mathbb{E}[(\psi_i(Z_i; \theta_0) - \tau)^4]$ and the Cauchy-Schwartz inequality such that $\text{Cov}\{(\psi_i(Z_i; \theta_0) - \tau)^2, (\psi_j(Z_j; \theta_0) - \tau)^2\} \leq \sqrt{\mathbb{E}[(\psi_i(Z_i; \theta_0) - \tau)^4] \mathbb{E}[(\psi_j(Z_j; \theta_0) - \tau)^4]}$. Continuing Equation (60), we have

$$P\left(|\tilde{V}_n - \eta^2| \geq \epsilon\right) \leq \frac{\text{Var}(\tilde{V}_n)}{\epsilon^2} = o(1),$$

implying that $\tilde{V}_n \xrightarrow{p} \eta^2$. Therefore, $\sqrt{n}(\hat{\tau}_n^{DB} - \tau)/\eta$ converges in distribution to $\mathcal{N}(0, 1)$.

We now show that \hat{V}_n also converges in probability to η^2 . We have

$$\begin{aligned} |\hat{V}_n - \tilde{V}_n| &= \left| \frac{1}{n} \sum_i \left(\psi(Z_i; \hat{\theta}_0) - \psi(Z_i; \theta_0) - \hat{\tau}_n^{DB} + \tau \right) \left(\psi(Z_i; \hat{\theta}_0) + \psi(Z_i; \theta_0) - \hat{\tau}_n^{DB} - \tau \right) \right| \\ &\stackrel{(i)}{\leq} O(1) \cdot \left\{ \frac{1}{n} \sum_i \left| \psi(Z_i; \hat{\theta}_0) - \psi(Z_i; \theta_0) \right| + \frac{1}{n} \sum_i \left| \hat{\tau}_n^{DB} - \tau \right| \right\} \\ &= O_p(\|\hat{\theta}_0 - \theta_0\|_{L_2}) + O_p(n^{-1/2}) = O_p(\epsilon_n) + O_p(n^{-1/2}) = o_p(1), \end{aligned}$$

where (i) is by the boundedness of $\psi(Z_i; \hat{\theta}_0)$ and $\psi(Z_i; \theta_0)$.

Therefore,

$$|\hat{V}_n - \eta^2| \leq |\hat{V}_n - \tilde{V}_n| + |\eta^2 - \tilde{V}_n| = o_p(1).$$

Combining the above with $\sqrt{n}(\hat{\tau}_n^{DB} - \tau)/\eta \xrightarrow{d} \mathcal{N}(0, 1)$, by Slutsky's theorem, we have $\sqrt{n}(\hat{\tau}_n^{DB} - \tau)/\sqrt{\hat{V}_n} \xrightarrow{d} \mathcal{N}(0, 1)$, concluding the proof.

Appendix F: Proof of the Asymptotics of the Generalized Propensity Weighted Estimators

We prove the below theorem to show the asymptotic regularity of the debiased estimator.

THEOREM 3. *Assume that the data-generating process follows the recommender choice model in Equation (9), and the viewer response model in Equation (10). Suppose Assumptions 1 and 2 hold, and that the nuisance estimates are consistent: $\|\hat{s}_0 - s_0\|_{L_2} + \|\hat{s}_1 - s_1\|_{L_2} + \|\hat{z} - z\|_{L_2} = o(1)$. Then both the GIPW estimator $\hat{\tau}_n^{GIPW}$ and the GAIPW estimator $\hat{\tau}_n^{GAIPW}$ are consistent, $\hat{\tau}_n^{GIPW} - \tau = o_p(1)$ and $\hat{\tau}_n^{GAIPW} - \tau = o_p(1)$.*

Moreover, GAIPW is Neyman orthogonal with respect to the nuisance functions \hat{z} and $\mathbb{E}_{\vec{W} \sim \mathcal{B}(q)}[p_k(\cdot, \cdot, \vec{W}; \hat{s}_0, \hat{s}_1)]$, but not with respect to $p_k(\cdot, \cdot, \vec{W} = \mathbf{1}; \hat{s}_0, \hat{s}_1)$ or $p_k(\cdot, \cdot, \vec{W} = \mathbf{0}; \hat{s}_0, \hat{s}_1)$.

We show the consistency of GIPW estimator in Appendix F.1, the consistency of GAIPW estimator in Appendix F.2, and the partial orthogonality of GAIPW estimator in Appendix F.3.

F.1. Consistency of GIPW Estimator

Recall the GIPW estimator as:

$$\hat{\tau}_n^{GIPW} = \frac{1}{n} \sum_{i=1}^n \psi_i^{GIPW}, \text{ for } \psi_i^{GIPW} = \left\{ r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) - r^{(C)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) \right\} Y_i, \quad (23)$$

where:

$$r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) := \frac{p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; \hat{s}_0, \hat{s}_1)}{\mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)}[p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; \hat{s}_0, \hat{s}_1)]}, \quad r^{(C)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) := \frac{p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i = \mathbf{0}; \hat{s}_0, \hat{s}_1)}{\mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)}[p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; \hat{s}_0, \hat{s}_1)]}.$$

Define the oracle GIPW estimator $\tilde{\tau}_n^{GIPW}$ as:

$$\tilde{\tau}_n^{GIPW} = \frac{1}{n} \sum_{i=1}^n \tilde{\psi}_i^{GIPW}, \text{ for } \tilde{\psi}_i^{GIPW} = \left\{ r^{(T)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) - r^{(C)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) \right\} Y_i.$$

To show the consistency of $\hat{\tau}_n^{GIPW}$, we show the below two steps:

1. Show the oracle GIPW estimator $\tilde{\tau}_n^{GIPW}$ converges to τ in probability, in Appendix F.1.1;
2. Show the GIPW estimator $\hat{\tau}_n^{GIPW}$ converges to the oracle GIPW estimator $\tilde{\tau}_n^{GIPW}$ in probability, in Appendix F.1.2.

It then directly follows that $\hat{\tau}_n^{GIPW}$ converges to τ in probability.

F.1.1. Show $\tilde{\tau}_n^{GIPW} - \tau \xrightarrow{p} 0$. First note that $\mathbb{E}[\tilde{\tau}_n^{GIPW}] = \tau$ since:

$$\mathbb{E} \left[\left\{ r^{(T)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) - r^{(C)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) \right\} Y_i \right] = \mathbb{E} \left[\left\{ \frac{d\mathbb{P}^{(T)}(Y_i)}{d\mathbb{P}^{(E)}(Y_i)} - \frac{d\mathbb{P}^{(C)}(Y_i)}{d\mathbb{P}^{(E)}(Y_i)} \right\} Y_i \right] = \mathbb{E}_{\mathbb{P}^{(T)}}[Y_i] - \mathbb{E}_{\mathbb{P}^{(C)}}[Y_i] = \tau.$$

We next analyze the variance of $\tilde{\tau}_n^{GIPW}$. Each $\tilde{\psi}_i^{GIPW}$ is bounded and thus has bounded variance. Given a sample j , its correlation with the other samples is bounded by $K\mathbb{E}[d_n]$, where d_n is the maximum appearance each item in n consideration sets in Assumption 1. Therefore:

$$\text{Var}(\tilde{\tau}_n^{GIPW}) = O \left(\frac{n(1 + K\mathbb{E}[d_n])}{n^2} \right) = O(\mathbb{E}[d_n]/n) = o(1).$$

By Markov inequality, we have $\tilde{\tau}_n^{GIPW} \xrightarrow{p} \tau$.

F.1.2. Show $\hat{\tau}_n^{GIPW} - \tilde{\tau}_n^{GIPW} \xrightarrow{P} 0$. Write $\hat{S}_{i,k} = \hat{s}_0(V_i, C_{i,k}) + W_{i,k} \hat{s}_1(V_i, C_{i,k})$ and $S_{i,k} = s_0(V_i, C_{i,k}) + W_{i,k} s_1(V_i, C_{i,k})$. Write $\vec{\hat{S}}_i = (\hat{S}_{i,1}, \dots, \hat{S}_{i,K})$ and $\vec{S}_i = (S_{i,1}, \dots, S_{i,K})$. Write $p_k(\vec{S}_i) := p_k(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1)$, and we have $\nabla p_k(\vec{S}_i) = p_k(\vec{S}_i)(e_k - p(\vec{S}_i))$, where e_k is the one-hot vector with k-th entry being nonzero and $p(\vec{S}_i) := (p_1(\vec{S}_i), \dots, p_K(\vec{S}_i))$. By mean-value theorem, we have $p_k(\vec{S}_i)$ be Lipschitz in \vec{S}_i and:

$$|p_k(\vec{\hat{S}}_i) - p_k(\vec{S}_i)| \leq \sup_{\vec{S}_i} \|p_k(\vec{S}_i)(e_k - p(\vec{S}_i))\|_2 \|\vec{\hat{S}}_i - \vec{S}_i\|_2 \leq \|\vec{\hat{S}}_i - \vec{S}_i\|_2. \quad (61)$$

By the boundedness of $\hat{s}_0, \hat{s}_1, s_0, s_1$ and Assumption 2, Lemma 1 gives that there exists a $\delta > 0$ to lower bound the exposure probability. We have:

$$\begin{aligned} & \left| r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) - r^{(T)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) \right| \\ & \leq \frac{\left| p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; \hat{s}_0, \hat{s}_1) \mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} \left[p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1) \right] - p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1) \mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} \left[p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; \hat{s}_0, \hat{s}_1) \right] \right|}{\mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} \left[p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; \hat{s}_0, \hat{s}_1) \right] \mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} \left[p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1) \right]} \\ & \leq \delta^{-2} \left\{ \left| p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; \hat{s}_0, \hat{s}_1) - p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1) \right| + \mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} \left| p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; \hat{s}_0, \hat{s}_1) - p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1) \right| \right\} \\ & \leq 2\delta^{-2} \|\vec{\hat{S}}_i - \vec{S}_i\|_2, \end{aligned} \quad (62)$$

where the last inequality is by Equation (61). Similarly:

$$\left| r^{(C)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) - r^{(C)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) \right| \leq 2\delta^{-2} \|\vec{\hat{S}}_i - \vec{S}_i\|_2. \quad (63)$$

All together, we have:

$$\begin{aligned} \mathbb{E} |\hat{\tau}_n^{GIPW} - \tilde{\tau}_n^{GIPW}| & \leq \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n |Y_i| \left(\left| r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) - r^{(T)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) \right| \right. \right. \\ & \quad \left. \left. + \left| r^{(C)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) - r^{(C)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) \right| \right) \right] \\ & \stackrel{(i)}{\lesssim} \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \|\vec{\hat{S}}_i - \vec{S}_i\|_2 \right] \leq \|\hat{s}_0 - s_0\|_{L_2} + \|\hat{s}_1 - s_1\|_{L_2} = o(1), \end{aligned}$$

where the inequality (i) uses the boundedness of Y_i , Equation (62) and Equation (63). Therefore $\hat{\tau}_n^{GIPW} - \tilde{\tau}_n^{GIPW} \xrightarrow{P} 0$.

F.2. Consistency of GAIPW Estimator

Recall the GAIPW estimator as:

$$\begin{aligned} \hat{\tau}_n^{GAIPW} & = \frac{1}{n} \sum_{i=1}^n \psi_i^{GAIPW}, \text{ for } \psi_i^{GAIPW} = \left\{ \mu^{(T)}(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z}) - \mu^{(C)}(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z}) \right\} \\ & \quad - \left\{ r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) - r^{(C)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) \right\} (Y_i - \hat{z}(V_i, C_{i,k_i^*})). \end{aligned} \quad (24)$$

where:

$$\begin{aligned} r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) & := \frac{p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; \hat{s}_0, \hat{s}_1)}{\mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} \left[p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; \hat{s}_0, \hat{s}_1) \right]}, \quad r^{(C)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) := \frac{p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i = \mathbf{0}; \hat{s}_0, \hat{s}_1)}{\mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} \left[p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; \hat{s}_0, \hat{s}_1) \right]}, \\ \mu^{(T)}(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z}) & = \sum_{k=1}^K \hat{z}(V_i, C_{i,k}) p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; \hat{s}_0, \hat{s}_1), \quad \mu^{(C)}(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z}) = \sum_{k=1}^K \hat{z}(V_i, C_{i,k}) p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{0}; \hat{s}_0, \hat{s}_1). \end{aligned}$$

Define the semi-oracle GAIPW estimator $\tilde{\tau}_n^{GAIPW}$ as:

$$\begin{aligned} \tilde{\tau}_n^{GAIPW} &= \frac{1}{n} \sum_{i=1}^n \tilde{\psi}_i^{GAIPW}, \text{ for } \tilde{\psi}_i^{GAIPW} = \left\{ \mu^{(T)}(V_i, \vec{C}_i; s_0, s_1, \hat{z}) - \mu^{(C)}(V_i, \vec{C}_i; s_0, s_1, \hat{z}) \right\} \\ &\quad - \left\{ r^{(T)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) - r^{(C)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) \right\} (Y_i - \hat{z}(V_i, C_{i,k_i^*})). \end{aligned}$$

To show the consistency of $\hat{\tau}_n^{GAIPW}$, we show the below two steps:

1. Show the semi-oracle GAIPW estimator $\tilde{\tau}_n^{GAIPW}$ converges to τ in probability, in Appendix F.2.1;
2. Show the GAIPW estimator $\hat{\tau}_n^{GAIPW}$ converges to the semi-oracle GAIPW estimator $\tilde{\tau}_n^{GAIPW}$ in probability, in Appendix F.2.2.

It then directly follows that $\hat{\tau}_n^{GAIPW}$ converges to τ in probability.

F.2.1. Show $\tilde{\tau}_n^{GAIPW} - \tau \xrightarrow{P} 0$. We first show that $\mathbb{E}[\tilde{\tau}_n^{GAIPW}] = \tau$. We have:

$$\begin{aligned} &\mu^{(T)}(V_i, \vec{C}_i; s_0, s_1, \hat{z}) - r^{(T)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) (Y_i - \hat{z}(V_i, C_{i,k_i^*})) \\ &= \sum_{k=1}^K \hat{z}(V_i, C_{i,k}) p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1) + \frac{p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1)}{\mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} [p_{k_i^*}(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1)]} (Y_i - \hat{z}(V_i, C_{i,k_i^*})) \\ &= \sum_{k=1}^K \hat{z}(V_i, C_{i,k}) p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1) + \sum_{k=1}^K \mathbf{1}\{k = k_i^*\} \frac{p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1)}{\mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} [p_k(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1)]} (Y_i - \hat{z}(V_i, C_{i,k})) \\ &= \sum_{k=1}^K p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1) \left(\hat{z}(V_i, C_{i,k}) + \frac{\mathbf{1}\{k = k_i^*\}}{\mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} [p_k(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1)]} (Y_i - \hat{z}(V_i, C_{i,k})) \right). \end{aligned}$$

Therefore,

$$\begin{aligned} &\mathbb{E} \left[\mu^{(T)}(V_i, \vec{C}_i; s_0, s_1, \hat{z}) - r^{(T)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) (Y_i - \hat{z}(V_i, C_{i,k_i^*})) \mid V_i, \vec{C}_i \right] \\ &= \mathbb{E} \left[\sum_{k=1}^K p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1) \mathbb{E} [Y_i \mid k_i^* = k, V_i, \vec{C}_i] \mid V_i, \vec{C}_i \right] \\ &= \sum_{k=1}^K p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1) z(V_i, C_{i,k}). \end{aligned}$$

Similarly,

$$\mathbb{E} \left[\mu^{(C)}(V_i, \vec{C}_i; s_0, s_1, z) - r^{(T)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) (Y_i - z(V_i, C_{i,k_i^*})) \mid V_i, \vec{C}_i \right] = \sum_{k=1}^K p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{0}; s_0, s_1) z(V_i, C_{i,k}).$$

Collectively, we have $\mathbb{E}[\tilde{\psi}_i^{GAIPW}] = \tau$, and thus $\mathbb{E}[\tilde{\tau}_n^{GAIPW}] = \tau$, showing the unbiasedness of $\tilde{\tau}_n^{GAIPW}$.

We next characterize the variance of $\tilde{\tau}_n^{GAIPW}$. Note that each $\tilde{\psi}_i^{GAIPW}$ is bounded with the boundedness of s_0, s_1, z and under Assumption 2. Given a sample j , its correlation with the other samples is bounded by $K\mathbb{E}[d_n]$, where d_n is the maximum appearance each item in n consideration sets in Assumption 1. Therefore:

$$\text{Var}(\tilde{\tau}_n^{GAIPW}) = O\left(\frac{n(1 + K\mathbb{E}[d_n])}{n^2}\right) = O(\mathbb{E}[d_n]/n) = o(1).$$

By Markov inequality, we have $\tilde{\tau}_n^{GAIPW} \xrightarrow{P} \tau$.

F.2.2. Show $\hat{\tau}_n^{GAIPW} - \tilde{\tau}_n^{GAIPW} \xrightarrow{p} 0$. We have:

$$\begin{aligned} & \left\{ \mu^{(T)}(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z}) - r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) (Y_i - \hat{z}(V_i, C_{i,k_i^*})) \right\} \\ & - \left\{ \mu^{(T)}(V_i, \vec{C}_i; s_0, s_1, \hat{z}) - r^{(T)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) (Y_i - \hat{z}(V_i, C_{i,k_i^*})) \right\} \\ & = \sum_{k=1}^K \left\{ \hat{z}(V_i, C_{i,k}) \left(p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; \hat{s}_0, \hat{s}_1) - p_k(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1) \right) \right. \\ & \quad \left. - (Y_i - \hat{z}(V_i, C_{i,k_i^*})) \left(r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) - r^{(T)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) \right) \right\}. \end{aligned}$$

By Equation (61), Equation (62), and the boundedness of Y_i, \hat{z} , we have

$$\left| \left\{ \mu^{(T)}(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z}) - r^{(T)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) (Y_i - \hat{z}(V_i, C_{i,k_i^*})) \right\} \right. \quad (64)$$

$$\left. - \left\{ \mu^{(T)}(V_i, \vec{C}_i; s_0, s_1, \hat{z}) - r^{(T)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) (Y_i - \hat{z}(V_i, C_{i,k_i^*})) \right\} \right| \lesssim \|\vec{S}_i - \vec{S}_i\|_2. \quad (65)$$

Similarly,

$$\begin{aligned} & \left| \left\{ \mu^{(C)}(V_i, \vec{C}_i; \hat{s}_0, \hat{s}_1, \hat{z}) - r^{(C)}(V_i, \vec{C}_i, k_i^*; \hat{s}_0, \hat{s}_1) (Y_i - \hat{z}(V_i, C_{i,k_i^*})) \right\} \right. \\ & \left. - \left\{ \mu^{(C)}(V_i, \vec{C}_i; s_0, s_1, \hat{z}) - r^{(C)}(V_i, \vec{C}_i, k_i^*; s_0, s_1) (Y_i - \hat{z}(V_i, C_{i,k_i^*})) \right\} \right| \lesssim \|\vec{S}_i - \vec{S}_i\|_2. \quad (66) \end{aligned}$$

Combining Equation (65) and Equation (66), we have $|\psi_i^{GAIPW} - \tilde{\psi}_i^{GAIPW}| \lesssim \|\vec{S}_i - \vec{S}_i\|_2$, and so:

$$\mathbb{E} |\hat{\tau}_n^{GAIPW} - \tilde{\tau}_n^{GAIPW}| \leq \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n |\psi_i^{GAIPW} - \tilde{\psi}_i^{GAIPW}| \right] \stackrel{(i)}{\lesssim} \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \|\vec{S}_i - \vec{S}_i\|_2 \right] \leq \|\hat{s}_0 - s_0\|_{L_2} + \|\hat{s}_1 - s_1\|_{L_2} = o(1).$$

Therefore $\hat{\tau}_n^{GAIPW} - \tilde{\tau}_n^{GAIPW} \xrightarrow{p} 0$.

F.3. Partial Neyman Orthogonality of GAIPW Estimator

We show the partial orthogonality of GAIPW estimator. Write

$$\begin{aligned} p^{(E)}(V_i, \vec{C}_i) &:= \mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} \left[p(V_i, \vec{C}_i, \vec{W}_i; s_0, s_1) \right], \\ p^{(T)}(V_i, \vec{C}_i) &= p(V_i, \vec{C}_i, \vec{W}_i = \mathbf{1}; s_0, s_1), \quad p^{(C)}(V_i, \vec{C}_i) = p(V_i, \vec{C}_i, \vec{W}_i = \mathbf{0}; s_0, s_1). \end{aligned}$$

Write their nuisance estimates as $\tilde{p}^{(E)}, \tilde{p}^{(T)}, \tilde{p}^{(C)}$ respectively. With some abuse of notation, define

$$\psi^{GAIPW}(V_i, \vec{C}_i, k_i^*, Y_i; \tilde{p}^{(E)}, \tilde{p}^{(T)}, \tilde{p}^{(C)}) = \sum_{k=1}^K \left(\tilde{p}_k^{(T)}(V_i, \vec{C}_i) - \tilde{p}_k^{(C)}(V_i, \vec{C}_i) \right) \left(\tilde{z}(V_i, C_{i,k}) + \frac{\mathbf{1}\{k = k_i^*\}}{p_k^{(E)}(V_i, \vec{C}_i)} (Y_i - \tilde{z}(V_i, C_{i,k})) \right).$$

Then the GAIPW score $\psi_i^{GAIPW} = \psi^{GAIPW}(V_i, \vec{C}_i, k_i^*, Y_i; \tilde{p}^{(E)}) = \mathbb{E}_{\vec{W}_i \sim \mathcal{B}(q)} [p_k(\cdot, \cdot, \vec{W}_i; \hat{s}_0, \hat{s}_1)]$, $\tilde{p}^{(T)} = p_k(\cdot, \cdot, \vec{W}_i = \mathbf{1}; \hat{s}_0, \hat{s}_1)$, $\tilde{p}^{(C)} = p_k(\cdot, \cdot, \vec{W}_i = \mathbf{0}; \hat{s}_0, \hat{s}_1)$, $\tilde{z} = \hat{z}$. To show the partial orthogonality of GAIPW estimator, we show that:

$$\mathbb{E} [\nabla_{\tilde{p}^{(E)}} \psi^{GAIPW}(V_i, \vec{C}_i, k_i^*, Y_i; \tilde{p}^{(E)} = p^{(E)}, \tilde{p}^{(T)} = p^{(T)}, \tilde{p}^{(C)} = p^{(C)}, \tilde{z} = z) | V_i, \vec{C}_i] = 0, \quad (67)$$

$$\mathbb{E} [\nabla_{\tilde{z}} \psi^{GAIPW}(V_i, \vec{C}_i, k_i^*, Y_i; \tilde{p}^{(E)} = p^{(E)}, \tilde{p}^{(T)} = p^{(T)}, \tilde{p}^{(C)} = p^{(C)}, \tilde{z} = z) | V_i, \vec{C}_i] = 0, \quad (68)$$

but

$$\mathbb{E} [\nabla_{\tilde{p}^{(T)}} \psi^{GAIPW}(V_i, \vec{C}_i, k_i^*, Y_i; \tilde{p}^{(E)} = p^{(E)}, \tilde{p}^{(T)} = p^{(T)}, \tilde{p}^{(C)} = p^{(C)}, \tilde{z} = z) | V_i, \vec{C}_i] \neq 0, \quad (69)$$

$$\mathbb{E} [\nabla_{\tilde{p}^{(C)}} \psi^{GAIPW}(V_i, \vec{C}_i, k_i^*, Y_i; \tilde{p}^{(E)} = p^{(E)}, \tilde{p}^{(T)} = p^{(T)}, \tilde{p}^{(C)} = p^{(C)}, \tilde{z} = z) | V_i, \vec{C}_i] \neq 0. \quad (70)$$

F.3.1. Orthogonality to $\tilde{p}^{(E)}$. It's equivalent to show orthogonality to $g := 1/\tilde{p}^{(E)}$. We have:

$$\begin{aligned} & \mathbb{E} \left[\nabla_g \psi^{GAIPW} (V_i, \vec{C}_i, k_i^*, Y_i; \tilde{p}^{(E)} = p^{(E)}, \tilde{p}^{(T)} = p^{(T)}, \tilde{p}^{(C)} = p^{(C)}, \tilde{z} = z) \mid V_i, \vec{C}_i \right] \\ &= \mathbb{E} \left[\sum_{k=1}^K \left(p_k^{(T)}(V_i, \vec{C}_i) - p_k^{(C)}(V_i, \vec{C}_i) \right) (\mathbf{1}\{k = k_i^*\} (Y_i - z(V_i, C_{i,k}))) \mid V_i, \vec{C}_i \right] \\ &= \sum_{k=1}^K \left(p_k^{(T)}(V_i, \vec{C}_i) - p_k^{(C)}(V_i, \vec{C}_i) \right) \mathbb{E} \left[\mathbf{1}\{k = k_i^*\} (Y_i - z(V_i, C_{i,k})) \mid V_i, \vec{C}_i \right] = 0. \end{aligned}$$

F.3.2. Orthogonality to \tilde{z} . We have:

$$\begin{aligned} & \mathbb{E} \left[\nabla_{\tilde{z}} \psi^{GAIPW} (V_i, \vec{C}_i, k_i^*, Y_i; \tilde{p}^{(E)} = p^{(E)}, \tilde{p}^{(T)} = p^{(T)}, \tilde{p}^{(C)} = p^{(C)}, \tilde{z} = z) \mid V_i, \vec{C}_i \right] \\ &= \sum_{k=1}^K \left(p_k^{(T)}(V_i, \vec{C}_i) - p_k^{(C)}(V_i, \vec{C}_i) \right) \mathbb{E} \left[1 - \frac{\mathbf{1}\{k = k_i^*\}}{p_k^{(E)}(V_i, \vec{C}_i)} \mid V_i, \vec{C}_i \right] = 0. \end{aligned}$$

F.3.3. Non-orthogonality to $\tilde{p}^{(T)}$. We have:

$$\begin{aligned} & \mathbb{E} \left[\nabla_{\tilde{p}^{(T)}} \psi^{GAIPW} (V_i, \vec{C}_i, k_i^*, Y_i; \tilde{p}^{(E)} = p^{(E)}, \tilde{p}^{(T)} = p^{(T)}, \tilde{p}^{(C)} = p^{(C)}, \tilde{z} = z) \mid V_i, \vec{C}_i \right] \\ &= \sum_{k=1}^K \mathbb{E} \left[z(V_i, C_{i,k}) + \frac{\mathbf{1}\{k = k_i^*\}}{p_k^{(E)}(V_i, \vec{C}_i)} (Y_i - z(V_i, C_{i,k})) \mid V_i, \vec{C}_i \right] = \sum_{k=1}^K z(V_i, C_{i,k}) \neq 0. \end{aligned}$$

F.3.4. Non-orthogonality to $\tilde{p}^{(C)}$. We have:

$$\begin{aligned} & \mathbb{E} \left[\nabla_{\tilde{p}^{(C)}} \psi^{GAIPW} (V_i, \vec{C}_i, k_i^*, Y_i; \tilde{p}^{(E)} = p^{(E)}, \tilde{p}^{(T)} = p^{(T)}, \tilde{p}^{(C)} = p^{(C)}, \tilde{z} = z) \mid V_i, \vec{C}_i \right] \\ &= - \sum_{k=1}^K \mathbb{E} \left[z(V_i, C_{i,k}) + \frac{\mathbf{1}\{k = k_i^*\}}{p_k^{(E)}(V_i, \vec{C}_i)} (Y_i - z(V_i, C_{i,k})) \mid V_i, \vec{C}_i \right] = - \sum_{k=1}^K z(V_i, C_{i,k}) \neq 0. \end{aligned}$$