
A/B Testing in Dense Large-Scale Networks: Design and Inference

Preetam Nandy, Kinjal Basu, Shaunak Chatterjee, Ye Tu

LinkedIn Corporation

Mountain View, CA 94083

{ pnandy, kbasu, shchatte, ytu } @linkedin.com

Abstract

Design of experiments and estimation of treatment effects in large-scale networks, in the presence of strong interference, is a challenging and important problem. Most existing methods’ performance deteriorates as the density of the network increases. In this paper, we present a novel strategy for accurately estimating the causal effects of a class of treatments in a dense large-scale network. First, we design an approximate randomized controlled experiment by solving an optimization problem to allocate treatments in the presence of competition among neighboring nodes. Then we apply an importance sampling adjustment to correct for any leftover bias (from the approximation) in estimating average treatment effects. We provide theoretical guarantees, verify robustness in a simulation study, and validate the scalability and usefulness of our procedure in a real-world experiment on a large social network.

1 Introduction

Measuring the effect of treatment variants is a fundamental problem in several fields of study [2, 22, 14]. A/B testing is a commonly used method in the internet industry [18, 19, 25, 29] and beyond, wherein randomized experiments are run with two or more such variants. Traditional A/B testing depends on the key assumption that the effect of treatment on an experiment unit is independent of the treatment allocation to other experiment units – commonly called “Stable Unit Treatment Value Assumption” (SUTVA) [22]. However, in many important network settings, this assumption is violated due to various forms of interference [3, 4, 13, 16, 17, 23, 26, 28].

We focus on one such class of problems – a marketplace of commodity producers and consumers. An example is a social media platform, where the commodity is content, and the return utility is feedback. Examples include Facebook, Instagram, LinkedIn, Twitter, etc. Other marketplace problems differ primarily on the commodity and return utility in question. Rides are commodities in Uber and Lyft, retail items are commodities in marketplaces like Amazon and eBay, and money is the common return utility. Note that we do not make any assumptions on how the utility affects a producer’s behavior (i.e., the incentive function).

In the marketplace setting, we define a treatment class represented by a **modification of the edge-level probability p_{ij} of showing a commodity of a producer i to a consumer j** . This abstraction represents a very general class of treatments since any redistribution or shift in producer and consumer exposure can be expressed via the change in the edge-level probabilities. In content marketplaces, such treatments can reshape the content creator’s exposure to their audience, whereas, in a transportation marketplace, it can reshape the likelihood of specific drivers and riders being matched.

In such marketplace problems, SUTVA is violated because of interference from the network. For instance, the effect of treatment on a producer (the experimental unit) is influenced by the allocation and effect of the treatment on all potential consumers of that producer, which in turn depend on each

of their producers’ experiences (and hence, their allocated treatments) – thus leading to competition among producers connected to common consumers. When SUTVA is violated, the effect of treatment can be measured by allocating experimental units as well as their first-degree neighbors (and second-degree neighbors, depending upon the interference function) to the same treatment (or control) [10, 23]. The number of such units for which we can successfully allocate treatments in this fashion (i.e., whose neighborhood is appropriately treated) decreases with increasing density in the graph, resulting in lower statistical power in the measurements.

In this paper, we propose a novel technique named OASIS, “Optimal Allocation Strategy and Importance Sampling Adjustment”, that provides a randomized testing framework for large-scale marketplace problems in the presence of interference. The performance of OASIS does not deteriorate (and may improve) as the density of the network increases, and it can be used to obtain high-power measurements. Our approach relies on the existence of an “intervening variable” (a.k.a. mediator) — i.e., the effect of the treatment allocation to an experimental unit’s network on the unit itself is fully captured by a sufficient statistic (cf. [1, 8, 20]). For example, the total feedback received by a content producer (or the total money made by a driver) can be an intervening variable. The producers in a marketplace compete with each other to receive feedback, and we attempt to mimic this *competition effect* in the design of experiment. Furthermore, we construct an unbiased estimator of the treatment effect by applying an importance sampling correction. We use simulation studies to compare our method against a graph cluster based baseline. From both the simulation studies and from a real-world experiment on a large social network, we show that our technique **(i) compares favorably against the baseline method and also (ii) works very well for large-scale dense graphs that cannot be decomposed into reasonably isolated clusters.**

This technique can be used in any marketplace problem where the treatment in question can be expressed via the p_{ij} “edge weight” abstraction, and where the key assumptions (outlined as Assumptions 1 and 2 in later sections) are satisfied. In many applications, the treatment needs to be applied over a time window during which p_{ij} can change periodically (e.g., consumers’ favored producers may change over time) and this dynamism can be handled by updating the experiment design periodically (see the real-world experiment in Section 5).

The rest of the paper is organized as follows. Section 2 formulates the problem and discusses the setup in detail. The key aspects of how we design an experiment via an optimization formulation are described in Section 3. We propose the OASIS estimator and provide theoretical guarantees in Section 4. In Section 5, we provide some empirical results both from an elaborate simulation study and a real-world experiment on a large social network graph, before concluding with a discussion in Section 6. We end this section with a selective literature review.

Related Work: Recently there has been a lot of focus on A/B testing in the presence of interference, especially on large real-world networks, with several proposed approaches. This is due to the ever-growing popularity of several marketplace products mentioned previously, and critical applications in those products violating SUTVA in randomized testing. Under the assumption that the interference between units has some known form, [3] propose inverse probability weighting based methods to estimate effects defined in terms of “effective treatments” and refines the estimator via covariance adjustment, and [26] propose to estimate causal effects under partial interference (i.e., when k neighbors are in treatment) using a sequential design. Unlike [3, 26], we aim to estimate the full treatment effect, i.e., all neighbors are in treatment vs. no neighbor is in treatment, for which the sequential design strategy is expected to break down ([26] used $k = 4$ in the experiments).

Network bucket testing [4, 17] looks at a form of interference where a new feature will take effect only if some minimal number of a treated user’s neighbors are also in treatment. The key idea, proposed in [4] is to use a walk-based sampling method to generate a core set of users who are internally well-connected while being approximately representative of the population. [17] generalizes this further and provides variance bounds for various core set generation functions. Graph cluster randomization is a related approach introduced in [28]. In this work, a notion of “network exposure” is defined, where a unit is “network exposed” to treatment if its behavior under a particular assignment is the same as its behavior if everyone were assigned to treatment. For various definitions of network exposure, a clustering approach is proposed, and randomization is done at the cluster level. In a real-world network, it is very difficult to obtain a large number of reasonably isolated clusters, resulting in low-powered experiments. To mitigate this issue, [23] proposed ego-network randomization, where a cluster is comprised of an “ego” (a focal individual), and her “alters” (the individuals she is

immediately connected to). Our class of treatment definitions can be viewed as a specific instance of network exposure. However, **our randomization is still at the node level and allows us to obtain much higher-confidence measurements with much smaller treatment exposure.**

In the marketplace setting, [30] introduced the bipartite randomized experiment framework, where two distinct groups of units are linked together through a bipartite graph. The treatment is assigned to the first group, called diversion units, and the response is measured on the second group, called outcome units. Under a linear treatment exposure model, [21] proposed a clustering-based approach for estimating the exposure-response function, whereas [9] proposed a generalized-propensity-score [15] based estimator. A treatment exposure model relates the responses of the outcome units to a measure of treatment exposure they receive (cf. Definition 1). The measure of treatment exposure is considered to be a linear function of the treatment assignments of the diversion units in a linear exposure model (cf. Assumption 1). Our framework differs from the bipartite randomized experiment framework in the following two major aspects. First, we consider general directed graphs where each experiment unit can be a diversion unit, an outcome unit, or both. Second, we consider a setting with limited exposure to each outcome unit, leading to a competition among the treatment units for being exposed to an outcome unit. The limited exposure is a common feature of many marketplace settings, where the consumers (e.g., buyers in a commodity marketplace or viewers in a content marketplace) can only view a limited number of items. Finally, we note that our approach does not require extrapolation of the estimated exposure-response function for estimating the average treatment effect, which caused an additional overhead in approaches proposed by [9, 21].

Two other related classical lines of work in estimating causal effect is the instrumental variable approach [1, 8] and the mediation or intervening variable effect estimation [20]. Both approaches make key assumptions about the how the causal effect flows through specific known variables to affect the outcome, and is similar to the assumption we make in order to facilitate replicating a completely treated universe with a much smaller exposure.

2 Problem Setup

We describe the problem in a content marketplace setup, where each member is a node in a directed graph (or network) $\mathcal{G} = (\Omega, E)$, and each member is either a content-consumer, a content-producer or both. If $i \rightarrow j$, then producer i is a parent of j and consumer j is a child of i . We denote the set of all parents of j by $Pa(j)$ and the set of all children of i by $Ch(i)$. When a consumer j visits the marketplace she views a content produced by one of her parents i with probability p_{ij} , where $\sum_{i \in Pa(j)} p_{ij} = 1$. We denote the probabilities in the existing system (i.e., prior to any intervention) by p_{ij}^{base} and an intervention (or a treatment) on member j replaces her consumer-side experience $\{p_{ij}^{base} : i \in Pa(j)\}$ by $\{p_{ij}^{(r)} : i \in Pa(j)\}$ satisfying $\sum_{i \in Pa(j)} p_{ij}^{(r)} = 1$.

Let $T_E^{(r)} = \{p_{ij}^{(r)} : (i, j) \in E\}$ be an intervention over the entire population and $\tau_r = \mathbb{E}[Y(T_E^{(r)})]$ be the expected response of the population under $T_E^{(r)}$. We assume that the responses of members i and j , denoted by $Y_i(T_E^{(r)})$ and $Y_j(T_E^{(r)})$ respectively, are independently distributed for a given $T_E^{(r)}$. We consider m treatments $T_E^{(1)}, \dots, T_E^{(m)}$ and a control environment $T_E^{(0)} = \{p_{ij}^{base} : (i, j) \in E\}$. Our goal is to estimate the average treatment effects $\tau_1 - \tau_0, \dots, \tau_m - \tau_0$ in a marketplace setup (see Figure 1b and 1a) defined below.

Definition 1. A directed network $\mathcal{G} = (\Omega, E)$ of members is said to be a marketplace if member i 's response $Y_i(T_E^{(r)})$ depends on the treatment $T_E^{(r)}$ through

1. **Consumer-side experience:** the exposure of producers to consumer i : $\{p_{ki}^{(r)} : k \in Pa(i)\}$, and
2. **Producer-side experience:** the exposure of producer i to consumers: $\{p_{ij}^{(r)} : j \in Ch(i)\}$.

In a marketplace setup, $T_E^{(r)}$ not only has a direct effect on $Y_i(T_E^{(r)})$ through the modification of the existing consumer-side experience $\{p_{ri}^{base} : r \in Pa(i)\}$ (i.e., what contents have been shown to i), but also has an indirect effect on $Y_i(T_E^{(r)})$ through the modification of the producer-side

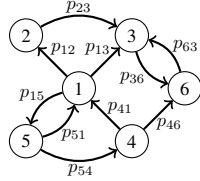
experience $\{p_{ij}^{base} : j \in Ch(i)\}$ (i.e., other members liking, sharing or commenting on the contents produced by i). The latter violates the SUTVA assumption and is known as the network effect. Further, each producer competes with its second-degree neighbors, to obtain higher $p_{ij}^{(r)}$ values as $\sum_{i \in Pa(j)} p_{ij}^{(r)} = 1$ for all j . This was referred as the competition effect in Section 1. In Figure 1a, the producer-side experience of member 1 depends on the feedback received from consumers 2, 3 and 5, while the consumer-side depends on the incoming edges from producers 4 and 5. Due to this fact, there is a competition between member 4 and member 5 for getting the feedback from consumer 1.

Recall that a classical A/B testing setting would consider randomly selected non-overlapping subsets $\Omega_0, \dots, \Omega_m$ of Ω and assign treatment $\{p_{ki}^{(r)} : k \in Pa(i)\}$ to each $i \in \Omega_r$ and estimate $\tau_r - \tau_0$ by

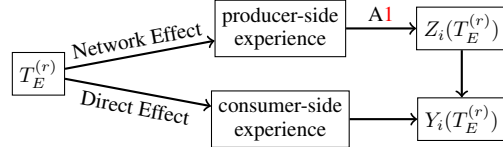
$$\hat{\tau}_r - \hat{\tau}_0 = \frac{1}{|\Omega_r|} \sum_{i \in \Omega_r} Y_i^{(exp)} - \frac{1}{|\Omega_0|} \sum_{i \in \Omega_0} Y_i^{(exp)}, \quad (1)$$

where $Y_i^{(exp)}$ is the response of member i under the aforementioned experimental design. It is easy to see that the estimator defined in (1) is not unbiased as the producer-side experience of member $i \in \Omega_r$ is not the same in the classical A/B testing design and under $T_E^{(r)}$, since all $j \in Ch(i)$ is not guaranteed to be in Ω_r . The bias of the estimator in (1) is known as the “spillover effect” in the literature. A commonly used technique for mitigating the spillover effect is to choose $\Omega_0, \dots, \Omega_m$ as disjoint from each other as possible. This is called the cluster-based approach. Even for moderately dense networks, a cluster-based method does not provide a reasonable solution, since it either induces a large bias by ignoring inter-cluster edges and/or suffers from the loss of power by considering a small sample consisting of nearly perfect clusters. In this paper, we propose a complementary approach under the following assumption.

Assumption 1. *The producer-side experience of member i depends on the treatment condition $T_E^{(r)}$ only through a weighted total exposure $Z_i(T_E^{(r)}) = \sum_{j \in Ch(i)} \alpha_{ij} p_{ij}^{(r)}$ for some constants α_{ij} ’s representing the strength of the relationship between producer i and consumer j (see Figure 1b).*



(1a) A producer-consumer network.



(1b) A graphical representation of Assumption 1.

In the content marketplace setup, Assumption 1 corresponds to the situation where the producer-side experience of member i depends only on an aggregated feedback to member i ’s content. Note that as the potential number of consumers for member i ’s content increases, it becomes more likely that the total feedback determines the producer-side experience. Thus, Assumption 1 becomes more reasonable, as the density of the network increases.

In Section 3, we construct an experiment T_E^* where we exactly match p_{ki}^* with $p_{ki}^{(r)}$ for all $i \in \Omega_r$ and $k \in Pa(i)$ (as in the classical A/B testing design), and attempt to match $Z_i(T_E^*)$ with $Z_i(T_E^{(r)})$ as much as possible by solving a joint optimization problem for all $i \in \Omega_r$ and $r \in [m]$ (where $[m] = 0, \dots, m$). In practice, it is reasonable to assume that α_{ij} ’s are unknown or approximately known. In Section 3.1, we show the robustness of our estimator under the violation of the assumption that α_{ij} ’s are known. Finally, we propose an importance sampling based adjustment to correct for experimental bias and describe the OASIS estimator.

3 Design of Experiment

We present our design in Algorithm 1 and an example in Figure 2. Let Ω_r, Λ_r , and C' be disjoint subsets of Ω (as in Algorithm 1), and let $\Omega' = \cup_{r=0}^m \Omega_r$ and $\Lambda' = \cup_{r=0}^m \Lambda_r$. The three disjoint subsets represent the following:

1. Ω_r : This is the population which will be used to measure $Y(T_E^{(r)})$. $|\Omega_r|$ determines the statistical power of the measurement. We allocate the exact consumer-side experience for each member $i \in \Omega_r$. The producer side experience is approximately matched by an optimization formulation (see below).
2. Λ_r : This is a shadow population for which we also allocate the exact consumer-side experience. As shown in Figure 2, since Ω_r matches both the consumer-side and producer-side experience while Λ_r matches only the consumer side experience, we can compare Ω_r and Λ_r to detect the presence of network effect.
3. C' : Members whose incoming edges from Ω' are assigned weights to match the producer experience for members in Ω' .

In order to achieve the above objectives, we assign edges in the following manner:

1. $p_{ij}^{(r)}$ to all incoming edges to members in $\Omega_r \cup \Lambda_r$, hence exactly matching their consumer side experience.
2. a new treatment p_{ij}^* to all incoming edges to members in C' , to approximately match the producer side experience of members in Ω' , and
3. p_{ij}^{base} to all incoming edges to members not in $\Omega_r \cup \Lambda_r \cup C'$, to control the overall risk.

We solve an optimization problem to define the new treatment p_{ij}^* for all edges from Ω' to C' and line (6) of Algorithm 1 defines p_{ij}^* of all other incoming edges to C' such that $\sum_{i \in Pa(j)} p_{ij}^* = 1$ for all $j \in C'$. The selection probability q in Algorithm 1 controls the number of members $|C'|$ that are receiving the new treatment.

Algorithm 1 Optimal Allocation Strategy (OAS)

- 1: Randomly select disjoint subsets $\Omega_0, \Omega_1, \dots, \Omega_m$ of Ω and define $\Omega' = \cup_{r=0}^m \Omega_r$;
 - 2: Randomly select disjoint subsets $\Lambda_0, \Lambda_1, \dots, \Lambda_m$ of $\Omega \setminus \Omega'$ and define $\Lambda' = \cup_{r=0}^m \Lambda_r$;
 - 3: Set $p_{ij}^* = p_{ij}^{(r)}$ for all $j \in \Omega_r \cup \Lambda_r, r = 0, \dots, m, i \in Pa(j)$;
 - 4: Construct a subset of children C' by selecting each member in $\cup_{i \in \Omega'} Ch(i) \setminus (\Omega' \cup \Lambda')$ with probability q ;
 - 5: Obtain $\{p_{ij}^* : j \in C', i \in Pa(j) \cap \Omega'\}$ by solving the optimization problem given in (2);
 - 6: Set $p_{ij}^* = p_{ij}^{base} \times \frac{1 - \sum_{k \in Pa(j) \cap \Omega'} p_{kj}^*}{1 - \sum_{k \in Pa(j) \cap \Omega'} p_{kj}^{base}}$ for $j \in C'$ and $i \in Pa(j) \setminus \Omega'$;
 - 7: For all other $(i, j) \in E$, set $p_{ij}^* = p_{ij}^{base}$;
-

The sets Λ_r will also be used in the importance sampling step described in Section 4. Thus, the existence of Λ_r increases the accuracy of our final estimator without increasing the time complexity of the optimization problem. Further, as mentioned above, Ω_r and Λ_r can be compared to detect the presence of the network effect in a real-world experiment.

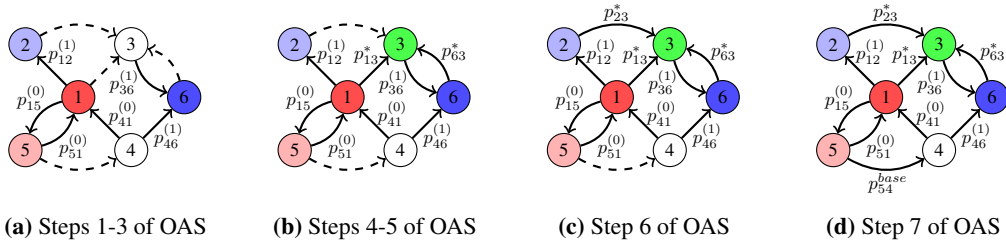


Figure 2: Consider Algorithm 1 with the network in Figure 1a. Let $\Omega_0 = \{1\}, \Omega_1 = \{6\}, \Lambda_0 = \{5\}, \Lambda_1 = \{2\}$ and $C' = \{3\}$. In (a), we assign the color red or blue to a node j with consumer-side experience $\{p_{ij}^{(0)} : i \in Pa(j)\}$ or $\{p_{ij}^{(1)} : i \in Pa(j)\}$ respectively. In (b) we assume $\alpha_{ij} = 1$ for all $(i, j) \in E$ and obtain $\{p_{ij}^*, p_{63}^*\}$ by minimizing $(p_{12}^{(0)} + p_{13}^{(0)} - p_{12}^{(1)} - p_{13}^{(1)})^2 + (p_{63}^{(1)} - p_{63}^*)^2$. Next, we assign $p_{23}^* = p_{23}^{base} \times \frac{1 - (p_{13}^* + p_{63}^*)}{1 - (p_{13}^{base} + p_{63}^{base})}$ in (c). Finally, we assign p_{54}^{base} to the edge $5 \rightarrow 4$ in (d).

3.1 Optimization Problem

Let $E^* \subseteq E$ be the set of all edges from Ω' to C' . Our goal is to obtain $\{p_{ij}^* : (i, j) \in E^*\}$ such that $J(E^*) = \sum_{r=0}^m \sum_{i \in \Omega_r} (Z_i(T_E^{(r)}) - Z_i(T_E^*))^2$ is minimized. Recall that $Z_i(T_E^{(r)}) = \sum_{j \in Ch(i)} \alpha_{ij} p_{ij}^{(r)}$. By Algorithm 1, we first assign all edge weights that won't be determined by the optimization. This includes assigning $p_{ij}^{(s)}$ to all edges $E_{i,s}$ from i to members in $\Omega_s \cup \Lambda_s$ and assign p_{ij}^{base} to all edges $E_{i,base}$ from i to members not in $\Omega_s \cup \Lambda_s \cup C'$. To capture this already assigned exposure of a member $i \in \Omega'$, we define

$$Z'_i(T_E^*) = \sum_{s=0}^m \sum_{(i,j) \in E_{i,s}} \alpha_{ij} p_{ij}^{(s)} + \sum_{(i,j) \in E_{i,base}} \alpha_{ij} p_{ij}^{base}.$$

Next, we denote the set of all edges from i to C' by $E_{i,C'}$. All outgoing edges from i is the union of $\{E_{i,s} : s = 0, \dots, m\}$, $E_{i,base}$ and $E_{i,C'}$ and we adjust the exposure level corresponding to $E_{i,C'}$ to minimize $(Z_i(T_E^{(r)}) - Z_i(T_E^*))^2$ simultaneously for all $i \in \Omega_r$, $r \in \{0, \dots, m\}$. Also, note that $\cup_{i \in \Omega'} E_{i,C'} = E^*$. Therefore, we find $\{p_{ij}^* : (i, j) \in E^*\}$ by solving

$$\begin{aligned} & \underset{p_{ij}}{\text{Minimize}} && \sum_{r=0}^m \sum_{i \in \Omega_r} (Z_i(T_E^{(r)}) - Z'_i(T_E^*) - \sum_{(i,j) \in E_{i,C'}} \alpha_{ij} p_{ij})^2 \\ & \text{subject to} && \sum_{i \in Pa(j) \cap \Omega'} p_{ij} \leq 1, \\ & && p_{ij} \geq 0. \end{aligned} \tag{2}$$

Risk Control: The risk of an experiment is defined as the total negative impact of an online experiment on members' experience. The risk of OAS tend to be higher than a classical A/B testing design due the modification in the consumer-side experience of members in C' . We can control the risk by restricting the size of C' , and by imposing additional constraints to the optimization problem to control the deviation of p_{ij}^* from p_{ij}^{base} (see Appendix A.1).

Scaling the Optimization: The quadratic programming (QP) defined in (2) does not scale well with the number of edges $n := |E^*|$. Even in moderately sized experiments in social network graphs, we can expect n to range in billions, making it almost impossible to solve the QP. We propose an iterative approximation that iterates over K overlapping sub-problems of roughly equal size n_{sub} , where each sub-problem respects the constraints of the full optimization and the solution of each sub-problem can improve the potential solution obtained in the previous step. By following this strategy, we achieve an $\mathcal{O}(K \times n_{sub}^3)$ time complexity, which is much better than the $\mathcal{O}(n^3)$ time complexity of QP when $K^{1/3} \times n_{sub} \ll n$. We use the Operator Splitting method to solve each QP [5, 6, 24]. The detailed algorithm and simulation results are provided in Appendix A.2.

Robustness: In the Appendix A.3, we show the robustness of the OAS design in the sense that we can work with partially known α_{ij} 's (see Assumption 1) or assume them to be equal to one without inducing any bias, as long as certain assumptions of independence are satisfied. More precisely, if we design a treatment T_E^* and a member i such that $p_{ki}^* = p_{ki}^{(r)}$ for all $k \in Pa(i)$ and

$$\sum_{j \in Ch(i)} \beta_{ij} p_{ij}^* = \sum_{j \in Ch(i)} \beta_{ij} p_{ij}^{(r)},$$

then $\mathbb{E}[Y_i(T_E^*)] = \mathbb{E}[Y_i(T_E^{(r)})]$ holds under much less strict condition than the equality of β_{ij} 's and α_{ij} 's. In our simulation study in Section 5, consider further violations of Assumption 1, where $Z_i(T_E^{(r)})$'s are chosen to be non-linear functions of $p_{ij}^{(r)}$'s.

4 Post-experiment Bias Correction

Importance Sampling Adjustment: Under the following assumption, we propose an unbiased estimator of the average treatment effect $\tau_r - \tau_0$ under a weaker requirement than a "perfect" design of experiment that facilitates the unbiased estimation using (1).

Assumption 2. For any treatment $T_E^{(r)}$, the producer-side experience of member i depends on $T_E^{(r)}$ only through the total exposure $Z_i(T_E^{(r)}) = \sum_{j \in Ch(i)} Z_{ij}(T_E^{(r)})$ for some observable i.i.d. random variables Z_{ij} 's. We assume that $Z_i(T_E^{(r)})$ and $\{p_{ki}^{(r)} : k \in Pa(i)\}$ are independently distributed.

The first part of Assumption 2 is weaker than Assumption 1 where we additionally assume $Z_{ij}(T_E^{(r)}) = \alpha_{ij}p_{ij}^{(r)}$. Also, note that the existence of an intermediate variable or mediator $Z_i(T_E^{(r)})$ need not be unique. For example, in the content marketplace setup, $Z_{ij}(T_E^{(r)})$ can be a function of $p_{ij}^{(r)}$, as well as $Z_{ij}(T_E^{(r)})$ can be a more downstream variable representing the feedback of member j on the content of member i . In the latter case, the second part of Assumption 2 corresponds to the assumption that the feedback on member i 's content is independent of the contents shown to member i , which is a reasonable assumption for most real-world recommendation systems.

Theorem 1. Let f_r denote the density of $Z_i(T_E^{(r)})$, and let f_r^* denote the density of $Z_i(T_E^*)$ conditionally on $i \in \Omega_r$, where T_E^* is the output of Algorithm 1. We define

$$w_i = \frac{f_r(Z_i(T_E^*))}{f_r^*(Z_i(T_E^*))}, \text{ and } \hat{\tau}_r = \frac{1}{|\Omega_r|} \sum_{i \in \Omega_r} w_i Y_i(T_E^*)$$

Then under Assumption 2, we have $\mathbb{E}[\hat{\tau}_r] = \tau_r, \forall r \in \{0, \dots, m\}$.

The correctness of Theorem 1 does not rely on Assumption 1 and the experimental design. However, from a practical standpoint, if we (approximately) know the dependency of $Z_i(T_E)$, we should use it to do the matching while designing the experiment. It is crucial for reducing the variance of the importance sampling weights as well as the corresponding estimator. Further, we recommend to use the self-normalized importance sampling $\frac{1}{\sum w_i} \sum Y_i w_i$, as it is known to be more stable in practice.

The densities f_r^* and f_r in Theorem 1 can be estimated using the data $\{Z_i(T_E^*) : i \in \Omega_r\}$ and the data $\{Z_{ij}(T_E^{(r)}) : i \in \Omega_r, j \in Ch(i) \cap (\Omega_r \cup \Lambda_r)\}$ respectively, where Ω_r and Λ_r are as in Algorithm 1. For further the details, please see Appendix A.5.

OASIS: We conclude this section by gluing all the pieces together in Algorithm 2, called Optimal Allocation Strategy and Importance Sampling Adjustment (OASIS). Under certain assumptions, an asymptotically correct $(1 - \alpha)100\%$ confidence interval for τ_r is given by $\hat{\tau}_r \pm \Phi^{-1}(1 - \alpha/2) \times \hat{\sigma}_r$, where Φ is the CDF of the standard normal distribution, where $\hat{\sigma}_r$ is a bootstrap variance estimator [11] (see Appendix A.6). The result follows from the consistency of density estimation, the central limit theorem and the consistency of the bootstrap variance estimator. Note that the estimator in Theorem 1 is an empirical average when the densities are known, implying the consistency of the bootstrap variance estimators under mild assumptions. When the densities are unknown, the consistency of the bootstrap variance relies on the consistency of the density estimation (which holds under some regularity conditions).

Algorithm 2 OASIS

- 1: Obtain T_E^* by applying Algorithm 1;
 - 2: Run experiment T_E^* to collect data;
 - 3: Obtain estimated densities \hat{f}_r^* and \hat{f}_r of $\{Z_i(T_E^*) : i \in \Omega_r\}$ and $\{Z_i(T_E^{(r)}) : i \in \Omega_r\}$ using the aforementioned technique, for $r = 0, \dots, m$;
 - 4: Obtain $\hat{\tau}_r = \frac{1}{|\Omega_r|} \sum_{i \in \Omega_r} Y_i(T_E^*) \frac{\hat{f}_r(Z_i(T_E^*))}{\hat{f}_r^*(Z_i(T_E^*))}$ for $r = 0, \dots, m$;
 - 5: Return $\hat{\tau}_1 - \hat{\tau}_0, \dots, \hat{\tau}_m - \hat{\tau}_0$.
-

5 Experiments

Simulation Study: We compare OASIS with a graph-cluster randomization method that randomly assigns treatment (or control) to all members in a cluster and uses (1) for estimating treatment effects. The cluster-based (CB) method yields an unbiased estimator when the clusters are disjoint, and the performance of CB is expected to deteriorate as the number of inter-cluster edges increase. We

generate graphs with 50000 nodes from some well studied graphical models [7, 12] with 5 equal-sized clusters and varying density (i.e. average degree) and varying inter-cluster edge proportion (see Appendix A.7.1). We directly use these clusters (instead of estimating them) for graph-cluster randomization to avoid the dependency on the performance of a graph clustering technique (see Appendix A.7.2). While the chosen network topology (i.e. the presence of clusters) and the use of the known clusters is somewhat favorable to CB, the presence of clusters does not provide any advantage to the OASIS design and estimates.

We generate data under moderate violations of Assumption 1. More precisely, we set $Z_i(T_E^{(r)})$ as $Z_i(\delta, T_E^{(r)}) = \sum_{j \in Ch(i)} \alpha_{ij} (p_{ij})^\delta$ for $r \in \{0, 1\}$, $\delta \in \{1/4, 1/2, 1\}$. **Observe that the strength of the network effect $Z_i^{(r)}(\delta, T_E^{(r)})$ decreases as δ increases.** We choose $\alpha_{ij} = 1$ for constructing T_E^* using Algorithm 1. The details on the graph and data generation and the treatment definitions are given in the Appendix.

\bar{d}_G	τ_0	$\tau_1 - \tau_0$	OASIS	CB	τ_0	$\tau_1 - \tau_0$	OASIS	CB	τ_0	$\tau_1 - \tau_0$	OASIS	CB
21	9.81	0.042	0.958	0.974	9.10	0.165	0.968	0.976	6.76	0.262	0.928	0.946
55	9.73	0.064	0.950	0.948	8.32	0.233	0.950	0.946	5.84	0.137	0.954	0.938
120	9.61	0.089	0.956	0.914	7.55	0.234	0.938	0.592	5.50	0.078	0.932	0.880

(a) $\delta = 1/4$ (b) $\delta = 1/2$ (c) $\delta = 1$

Table 1: Empirical coverage probability for 95% confidence intervals, computed based on $n = 500$ experiments. Higher coverage corresponds to higher accuracy in measurement. \bar{d}_G denotes the average degree of the graph.

Table 1 shows the empirical coverage probabilities defined as the proportion of times $\tau_1 - \tau_0$ lies in the corresponding 95% confidence interval. Ideally, the empirical coverage probability should match 0.95, a lower value would indicate a higher Type-1 error and a higher value would indicate a higher Type-2 error. The setting \bar{d}_G is the most unfavorable setting for CB with the proportion of inter-cluster edges, and among all settings with \bar{d}_G , the setting with $\delta = 1/2$ has the largest treatment effect. We suspect that CB failed to capture a significant portion of the network effect under these extreme conditions.

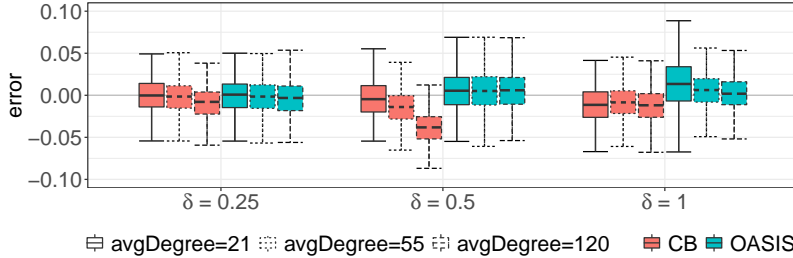


Figure 3: Box-plots of the estimation errors under different settings corresponding to $\delta \in \{1/4, 1/2, 1\}$ and $\text{avgDegree} = \bar{d}_G \in \{21, 55, 120\}$ based on 500 repeats.

We compare the errors of estimating $\tau_1 - \tau_0$ with the OASIS and the cluster-based estimators in Figure 3. We observe two important patterns:

1. The performance of the cluster-based method deteriorates as we move from a sparse and low impurity setting to a dense and high impurity setting, while the performance of OASIS remains the same or gets better as the density of the network increases. Note that with increasing density in the graph, the experiment design part of OASIS gets more flexibility to redistribute the exposure optimally, and the post-experiment adjustment part gets a larger sample size for density estimation (see Section A.5 in the supplement).
2. Performance is better, and the difference between OASIS and the cluster-based method is higher when the treatment effect is larger (e.g., the middle set of experiments, where $\tau_1 - \tau_0$ is large).

Real World Experiments: We demonstrate an application of our method on the LinkedIn newsfeed, a content recommender system with 675M+ members. A content recommender system is often optimized for content consumers' experience. However, there has been a growing interest in developing

recommender systems that additionally enhance content producers' experience [27]. It is easy to see that the classical A/B testing framework won't be able to capture the enhancement in the producers' experience. We apply OASIS for A/B testing a new recommender system $M^{(1)}$ and a baseline model $M^{(0)}$. For details on the setup, please see Appendix A.8. To estimate the $p_{ij}^{(0)}$ and $p_{ij}^{(1)}$ values, we ran a preliminary experiment for one week where we assigned $M^{(0)}$ and $M^{(1)}$ to members in $\Omega_0 \cup \Lambda_0 \cup \Gamma$ and $\Omega_1 \cup \Lambda_1$ respectively. The assignment of $M^{(r)}$ to member i means that the LinkedIn newsfeed of member i generated by $M^{(r)}$ as follows. An item-generator model creates a list of items for each member i , then each *item* in the list receives a score $M^{(r)}(item, i)$, and finally the items are shown to member i ranked by their scores.

Let's denote the sum of the scores of all items of producer k that are shown to member i by $S^{(r)}(k, i)$. Then,

$$p_{ij}^{(r)} := \frac{S^{(r)}(i, j)}{\sum_{k \in Pa(j)} S^{(r)}(k, j)}.$$

Moreover, we define

$$Z_i(T_E^{(r)}) := \frac{|\Omega|}{|\Omega_r \cup \Lambda_r|} \sum_{j \in Ch(i) \cap (\Omega_r \cup \Lambda_r)} p_{ij}^{(r)} \text{ for } r = 0, 1.$$

We assume $\alpha_{ij} = 1$ and we obtain $\{p_{ij}^* : (i, j) \in E^*\}$ by minimizing $J(E^*)$ defined in Section 3, where E^* is the set of all edges from Ω' to C' . We started running the main experiment by modifying the scores generated by $M^{(0)}$ with a multiplicative boost factors b_{ij} 's (see Appendix A.8.1) such that

$$\frac{b_{ij} S^{(0)}(i, j)}{\sum_{k \in Pa(j)} b_{kj} S^{(0)}(k, j)} = p_{ij}^*, \text{ for all incoming edges to members in } C'. \quad (3)$$

We ran the main experiment for three weeks while updating p_{ij}^* 's and b_{ij} 's periodically (twice a week) to incorporate the dynamic nature of the underlying network. The OASIS estimates (based on the main experiment) showed significant impact in the following metrics that are related to the producer side experience of a member:

- (i) *Daily Unique Contributors*: Number of unique members in a day who contribute on (i.e., either like, comment or share) items on the newsfeed, showed a lift of **+0.52%** with a p -value of 0.041.
- (ii) *Members in Active Community*: Number of members who contribute to or receive contributions from at least 5 members, showed a lift of **+1.14%** with a p -value of 0.016.
- (iii) *Feed Interaction received Uniques*: Number of unique members who receive at least one contribution, showed a lift of **+1.34%** with a p -value of 0.001.

The above results compares the responses of members in Ω_1 with the members in Ω_0 . A similar comparison for Λ_1 with the members in Λ_0 turns out to be statistically insignificant for these metrics, indicating the presence of network effect.

6 Discussion

We have presented a two-step method, called OASIS, for estimating the average treatment effect for a class of treatments represented by modifications of edge weights in networks with interference. First, we design an experiment by optimally allocating a "modified" treatment exposure to a set of randomly selected consumer-producer pairs by solving a large-scale quadratic program. Secondly, we apply an importance sampling correction for estimating the average treatment effect, which corrects for the design bias induced by the violations of assumptions and/or the restrictions applied for risk control. The main advantage of OASIS compared to cluster-based methods is that it tends to perform better for dense networks that cannot be decomposed into reasonably isolated clusters, while the cluster-based methods would have an advantage over OASIS for sparse networks. An interesting future work could be to combine a cluster-based approach with OASIS in order to gain additional robustness and efficiency. Finally, note that we can handle dynamic networks where the edge weights p_{ij} change over the experiment time window by updating the OASIS design periodically. In some cases, p_{ij} may change very quickly with time or depend on the current context (e.g., drivers and riders are matched based on current proximity). Therefore, considering the distribution of p_{ij} (instead of a point estimate) for designing the OASIS treatment allocation can be another interesting future work.

Broader Impact

Online A/B testing plays the most crucial role in product development in the internet industry by deciding which recommender system is an optimal choice. The optimality criterion often focuses on the experience of the direct users of the recommender system, e.g., the viewers of a social media newsfeed or the recruiters searching for potential candidates. This choice can have a potentially negative impact on the people indirectly affected by these recommendation choices. For example, the less popular content creators or a certain group of jobseekers might be getting less exposure than the already popular creators or candidates, leading to a “rich get richer” ecosystem. There has been a growing interest in enhancing the experience of the people that are indirectly affected. However, a significant obstacle to achieving this goal is that these indirect effects are challenging to measure in an online A/B testing experiment. In this paper, we take a step toward solving this problem, which can potentially lead to product developments focusing on the holistic improvement of the underlying ecosystem.

Acknowledgment

We would sincerely like to thank Guillaume Saint-Jacques, Iavor Bojinov, Ya Xu, Lingjie Weng, Shipeng Yu, Hema Raghavan, Romer Rosales and Deepak Agarwal for their support and detailed insightful feedback during the development of this system. We would also like to thank the anonymous reviewers for their helpful comments which has significantly improved the paper.

Finally, none of the authors received any third-party funding for this submission and there is no competing interest other than LinkedIn Corporation to which all authors are affiliated.

References

- [1] M. Arellano and O. Bover. Another look at the instrumental variable estimation of error-components models. *Journal of Econometrics*, 68:29–51, 1995.
- [2] P. Armitage, G. Berry, and J. N. S. Matthews. *Statistical Methods in Medical Research*. Blackwell Science Ltd, 2002.
- [3] P. M. Aronow and C. Samii. Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics*, 11:1912–1947, 2017.
- [4] L. Backstrom and J. Kleinberg. Network bucket testing. In *Proceedings of the 20th International Conference on World Wide Web*, pages 615–624. ACM, 2011.
- [5] G. Banjac, P. Goulart, B. Stellato, and S. Boyd. Infeasibility detection in the alternating direction method of multipliers for convex optimization. In *2018 UKACC 12th International Conference on Control (CONTROL)*, pages 340–340, 2018.
- [6] G. Banjac, B. Stellato, N. Moehle, P. Goulart, A. Bemporad, and S. Boyd. Embedded code generation using the OSQP solver. In *IEEE Conference on Decision and Control (CDC)*, 2017.
- [7] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.
- [8] J. G. Cragg and S. G. Donald. Testing identifiability and specification in instrumental variable models. *Econometric Theory*, 9:222–240, 1993.
- [9] N. Doudchenko, M. Zhang, E. Drynkin, E. Airolidi, V. Mirrokni, and J. Pouget-Abadie. Causal inference with bipartite designs. arXiv: 2010.02108, 2020.
- [10] D. Eckles, B. Karrer, and J. Ugander. Design and analysis of experiments in networks: Reducing bias from interference. *Journal of Causal Inference*, 5(1), 2017.
- [11] B. Efron. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7:1–26, 1979.
- [12] P. Erdős and A. Rényi. On random graphs I. *Publicationes Mathematicae*, 6:290–297, 1959.
- [13] H. Gui, Y. Xu, A. Bhasin, and J. Han. Network A/B testing: From sampling to estimation. In *Proceedings of the 24th International Conference on World Wide Web*, pages 399–409, 2015.

- [14] M. A. Hernán and J. M. Robins. Estimating causal effects from epidemiological data. Journal of Epidemiology & Community Health, 60:578–586, 2006.
- [15] K. Hirano and G. Imbens. The Propensity Score with Continuous Treatments, pages 73–84. Wiley Blackwell, 2004.
- [16] M. G. Hudgens and M. E. Halloran. Toward causal inference with interference. Journal of the American Statistical Association, 103:832–842, 2008.
- [17] L. Katzir, E. Liberty, and O. Somekh. Framework and algorithms for network bucket testing. In Proceedings of the 21st International Conference on World Wide Web, pages 1029–1036, 2012.
- [18] R. Kohavi, A. Deng, B. Frasca, T. Walker, Y. Xu, and N. Pohlmann. Online controlled experiments at large scale. In Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 1168–1176. ACM, 2013.
- [19] R. Kohavi, A. Deng, R. Longbotham, and Y. Xu. Seven rules of thumb for web site experimenters. In Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 1857–1866. ACM, 2014.
- [20] D. P. MacKinnon, C. M. Lockwood, J. M. Hoffman, S. G. West, and V. Sheets. A comparison of methods to test mediation and other intervening variable effects. Psychological Methods, 7:83–104, 2002.
- [21] J. Pouget-Abadie, K. Aydin, W. Schudy, K. Brodersen, and V. Mirrokni. Variance reduction in bipartite experiments through correlation clustering. In Advances in Neural Information Processing Systems 32, pages 13309–13319. Curran Associates, Inc., 2019.
- [22] D. B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. Journal of Educational Psychology, 66:688–701, 1974.
- [23] G. Saint-Jacques, M. Varshney, J. Simpson, and Y. Xu. Using ego-clusters to measure network effects at LinkedIn. arXiv: 1903.08755, 2019.
- [24] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd. OSQP: An operator splitting solver for quadratic programs. Math. Prog. Comp., 12:637–672, 2020.
- [25] D. Tang, A. Agarwal, D. O’Brien, and M. Meyer. Overlapping experiment infrastructure: More, better, faster experimentation. In Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 17–26, 2010.
- [26] P. Toulis and E. Kao. Estimation of causal peer influence effects. In International Conference on Machine Learning, pages 1489–1497, 2013.
- [27] Y. Tu, C. Lo, Y. Yuan, and S. Chatterjee. Feedback shaping: A modeling approach to nurture content creation. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 2241–2250, 2019.
- [28] J. Ugander, B. Karrer, L. Backstrom, and J. Kleinberg. Graph cluster randomization: Network exposure to multiple universes. In Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 329–337, 2013.
- [29] Y. Xu, N. Chen, A. Fernandez, O. Sinno, and A. Bhasin. From infrastructure to culture: A/b testing challenges in large scale social networks. In Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 2227–2236, 2015.
- [30] C. M. Zigler and G. Papadogeorgou. Bipartite causal inference with interference. arXiv: 1807.08660, 2018.

A Appendix

A.1 Risk Control

We find $\{p_{ij}^* : (i, j) \in E^*\}$ by minimizing

$$J(E^*) = \sum_{r=0}^m \sum_{i \in \Omega_r} (Z_i(T_E^{(r)}) - Z'_i(T_E^*) - \sum_{(i,j) \in E_{i,C'}} \alpha_{ij} p_{ij})^2.$$

under some constraints controlling the risk of overexposure or underexposure of parents of $j \in C'$. We choose $0 \leq R_{\min} \leq 1 \leq R_{\max}$ and $0 \leq S_{\min} \leq 1 \leq S_{\max}$ to control the changes in $\{p_{ij}^{base} : (i, j) \in E^*\}$ and $\{p_{ij}^{base} : j \in C', (i, j) \in E \setminus E^*\}$ respectively by defining $\{p_{ij}^* : (i, j) \in E^*\}$ as a solution of the following quadratic optimization problem with linear constraints:

$$\begin{aligned} \text{Min}_{p_{ij}} J(E^*) \text{ such that: } & p_{ij}^{base} R_{\min} \leq p_{ij} \leq p_{ij}^{base} R_{\max}, \\ & S_{\min} \leq \frac{1 - \sum_{i \in Pa(j) \cap \Omega'} p_{ij}}{1 - \sum_{i \in Pa(j) \cap \Omega'} p_{ij}^{base}} \leq S_{\max}, 0 \leq \sum_{i \in Pa(j) \cap \Omega'} p_{ij} \leq 1. \end{aligned} \quad (4)$$

Note that the last two constraints can be combined as $\ell_j \leq \sum_{i \in Pa(j) \cap \Omega'} p_{ij} \leq u_j$ for $\ell_j = \max\{0, 1 - S_{\max}(1 - \sum_{i \in Pa(j) \cap \Omega'} p_{ij}^{base})\}$ and $u_j = \min\{1, 1 - S_{\min}(1 - \sum_{i \in Pa(j) \cap \Omega'} p_{ij}^{base})\}$.

A.2 Scaling the Optimization

We split the set of consumers C' into disjoint sets S_k of roughly equal sizes, such that $C' = \cup_{k=1}^K S_k$. Since this induces a natural partition in the constraint space of (4), it is easy to see that any candidate solution $\{p_{ij}^{old} : (i, j) \in E^*\}$ of (4) can be improved by updating $\{p_{ij}^{old} : (i, j) \in E^*, j \in S_k\}$ with the solution of the following optimization problem for all $k \in \{1, \dots, K\}$. For $i \in \Omega_r$ let

$$\Delta_{ik,r} := Z_i(T_E^{(r)}) - Z'_i(T_E^*) - \sum_{(i,j) \in E_{i,C'}, j \notin S_k} \alpha_{ij} p_{ij}^{old}.$$

Minimize
 p_{ij}

$$\sum_{r=0}^m \sum_{i \in \Omega_r \cap (\cup_{j \in S_k} Pa(j))} \left(\Delta_{ik,r} - \sum_{(i,j) \in E_{i,C'}, j \in S_k} \alpha_{ij} p_{ij} \right)^2$$

such that

$$\begin{aligned} R_{\min} \times p_{ij}^{base} &\leq p_{ij} \leq R_{\max} \times p_{ij}^{base}, \\ S_{\min} &\leq \frac{1 - \sum_{i \in Pa(j) \cap \Omega'} p_{ij}}{1 - \sum_{i \in Pa(j) \cap \Omega'} p_{ij}^{base}} \leq S_{\max}, \text{ and} \\ 0 &\leq \sum_{i \in Pa(j) \cap \Omega'} p_{ij} \leq 1. \end{aligned} \quad (5)$$

Note that (5) and (4) have the same set of constraints for each $j \in S_k$.

We apply this strategy in Algorithm 3 to solve the overall optimization problem. We start with a feasible candidate $p_{ij}^{old} = p_{ij}^{base}$ and run an iterative scheme to update p_{ij} using (5) as we loop over each $k = \{1, \dots, K\}$. Once this inner loop completes, we get the full next best $\{p_{ij} : (i, j) \in E^*\}$. We continue the outer loop till convergence. By doing this iterative scheme we are able to solve much larger problems, since size of the each optimization problem $n_k = \sum_{j \in S_k} |Pa(j) \cap \Omega'|$ is much smaller than $n = \sum_{j \in C'} |Pa(j) \cap \Omega'|$. In fact, the worst-case complexity of the iterative method with $nIter$ outer iterations is given by $\mathcal{O}(maxIter \times \sum_{k=1}^K n_k^3)$. This is much better than the $\mathcal{O}(n^3)$ worst-case complexity of (4) when $maxIter = \mathcal{O}(1)$ and $K^{1/3} \times (\max_k n_k) \ll n$. Although the convergence cannot be guaranteed for $maxIter = \mathcal{O}(1)$, we were able to obtain reasonable solutions

Algorithm 3 Solving for Optimal p_{ij}

Input: Number of Groups K , Outer Iteration Limit $maxIter$

- 1: Create \mathcal{S}_k s.t. $\mathcal{S} = \cup_{k=1}^K \mathcal{S}_k$ s.t. $|\mathcal{S}_k| \approx |\mathcal{S}|/K$ for all k ;
 - 2: Set $p_{ij}^{old} = p_{ij}^{base}$ and $p_{ij}^* = p_{ij}^{base}$;
 - 3: **for** $t = 1, \dots, maxIter$ **do**
 - 4: **for** $k = 1, \dots, K$ **do**
 - 5: Solve the optimization problem as in (5) to obtain p_{ij}^* for all $j \in \mathcal{S}_k, i \in \text{Ne}(j) \cap \Omega'$;
 - 6: Set $p_{ij}^{old} = p_{ij}^*$ for all $j \in \mathcal{S}_k, i \in \text{Ne}(j) \cap \Omega'$;
 - 7: **end for**
 - 8: **end for**
 - 9: Return $\{p_{ij}^* : i \in \Omega', j \in \mathcal{S}\}$.
-

in simulations and real applications by using $maxIter = 10$. Finally, note that for a fixed value of $maxIter$, K represents the speed-accuracy trade-off, i.e. the quality of the solution is expected to be better for a smaller value of K but the computational efficiency is expected to be better for a larger value of K .

We illustrate the computational efficiency and the accuracy of the iterative algorithm with the following simulation setup. We consider 5 settings with $n = 10000, 20000, 30000, 40000$ and 50000 variables. For each setting, we generate n producers and n consumers by sampling with replacement from $\Omega = \{1, \dots, 100\}$. Let Ω' denote the set of all selected producers and let E^* denote the set of selected producer-consumer pairs. For each producer $i \in \Omega'$, we generate h_i from $Uniform[0, n/100]$ and for each producer-consumer pair $(i, j) \in E^*$, we generate $p_{ij}^{base} \sim Uniform[0, 1]$. We solve the following optimization problem:

$$\begin{aligned} & \text{Minimize } \sum_{i \in \Omega'} \left(h_i - \sum_{j \in Ch(i)} p_{ij} \right)^2 \text{ such that:} \\ & 0.1 \times p_{ij}^{base} \leq p_{ij} \leq 10 \times p_{ij}^{base}, \quad \text{and} \quad 0.2 \times \sum_{i \in Pa(j)} p_{ij}^{base} \leq \sum_{i \in Pa(j)} p_{ij} \leq 5 \times \sum_{i \in Pa(j)} p_{ij}^{base}, \end{aligned} \quad (6)$$

where the $Pa(i)$ and $Ch(i)$ are defined on the producer-consumer graph (Ω, E^*) .

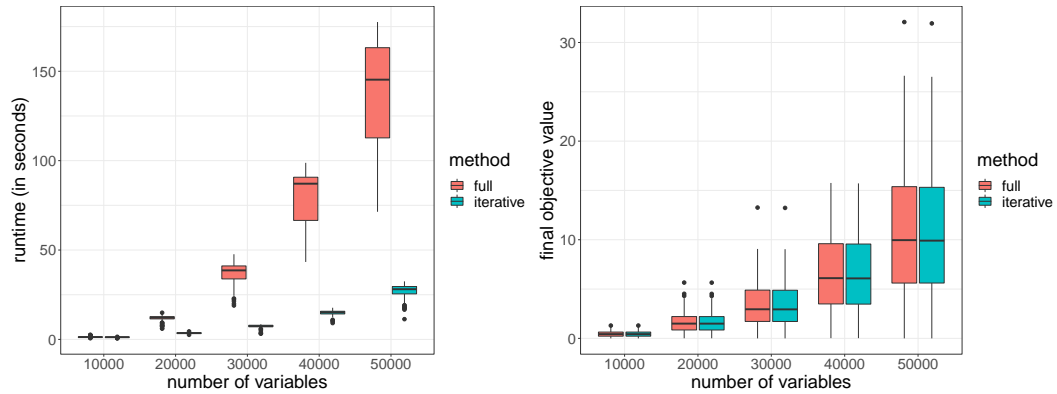


Figure 4: Boxplots of runtimes and the final objective values for solving (6) directly using a QP solver (method: full) and for solving (6) using Algorithm 3 with $K = 10$ and $maxIter = 2$ (method: iterative). We used the Operator Splitting method [24] in R package **rosqp** for solving each subproblem in Algorithm 3 as well as for solving the full QP.

Figure 4 shows that our iterative QP algorithm (Algorithm 3) can outperform the full QP solver in terms of computational efficiency without sacrificing the solution quality (measured by the final objective value).

A.3 Robustness

Lemma 1. Let i be a randomly chosen member in Ω and let $\{p_{ki}^{(r)} : k \in Pa(i)\}$ and $\{\alpha_{ij}, p_{ij}^{(r)} : j \in Ch(i)\}$ be well-defined random variables. Assume

$$Y_i(T_E^{(r)}) = g\left(\sum_{j \in Ch(i)} \alpha_{ij} p_{ij}^{(r)}, \mathbf{W}_i(T_E^{(r)})\right) + \epsilon_i,$$

where g is a differentiable function with respect to the first coordinate, $\mathbf{W}_i(T_E^{(r)}) := \{p_{ki}^{(r)} : k \in Pa(i)\}$, and ϵ_i does not depend on $T_E^{(r)}$. Let $\{U_{ij} : j \in Ch(i)\}$ be a set of i.i.d. random variables independent of $\{(\beta_{ij} p_{ij}^{(r)}, \alpha_{ij}/\beta_{ij}) : j \in Ch(i)\}$. Define

$$U_{ij}^* = \frac{U_{ij} \sum_{k \in Ch(i)} \beta_{ik} p_{ik}^{(r)}}{\sum_{k \in Ch(i)} \beta_{ik} p_{ik}^{(r)} U_{ik}}$$

such that $\{U_{ij}^* : j \in Ch(i)\}$ are identically distributed random variables satisfying

$$\sum_{j \in Ch(i)} \beta_{ij} p_{ij}^{(r)} U_{ij}^* = \sum_{j \in Ch(i)} \beta_{ij} p_{ij}^{(r)}. \quad (7)$$

If there exist an experimental design T_E^* such that $\mathbf{W}_i(T_E^*) = \mathbf{W}_i(T_E^{(r)})$ and $p_{ij}^* = p_{ij}^{(r)} U_{ij}^*$, then $\mathbb{E}[Y_i(T_E^*)] = \mathbb{E}[Y_i(T_E^{(r)})]$.

Proof. Since g is differentiable with respect to the first coordinate, and $\mathbf{W}_i(T_E^*) = \mathbf{W}_i(T_E^{(r)})$, by applying mean value theorem, we obtain

$$\begin{aligned} & \mathbb{E}[Y_i(T_E^*)] - \mathbb{E}[Y_i(T_E^{(r)})] \\ &= \mathbb{E} \left[g\left(\sum_{j \in Ch(i)} \alpha_{ij} p_{ij}^{(r)} U_{ij}^*, \mathbf{W}_i(T_E^{(r)})\right) \right. \\ & \quad \left. - g\left(\sum_{j \in Ch(i)} \alpha_{ij} p_{ij}^{(r)}, \mathbf{W}_i(T_E^{(r)})\right) \right] \\ &= \mathbb{E} \left[g'(x_i^*, \mathbf{W}_i(T_E)) \sum_{j \in Ch(i)} \alpha_{ij} p_{ij}^{(r)} (U_{ij}^* - 1) \right], \end{aligned}$$

where g' denotes the partial derivative of g with respect to the first coordinate, and x_i^* depends on $\{(\alpha_{ij}, p_{ij}^{(r)}, U_{ij}^*) : j \in Ch(i)\}$. We complete the proof by showing that the conditional expectation of

$$V_{ij}(T_E^{(r)}) := g'(x_i^*, \mathbf{W}_i(T_E))(U_{ij}^* - 1)$$

given $\mathbf{Z}_i(T_E^{(r)}) := \{\alpha_{ij} p_{ij}^{(r)} : j \in Ch(i)\}$ is zero.

From $\sum_{j \in Ch(i)} \beta_{ij} p_{ij}^{(r)} U_{ij}^* = \sum_{j \in Ch(i)} \beta_{ij} p_{ij}^{(r)}$, we have

$$\sum_{j \in Ch(i)} \beta_{ij} p_{ij}^{(r)} V_{ij}(T_E) = 0. \quad (8)$$

It follows from the definition of U_{ij}^* and the independence of $\{U_{ij} : j \in Ch(i)\}$ and $\mathbf{Z}_i^*(T_E) = \{(\beta_{ij} p_{ij}^{(r)}, \alpha_{ij}/\beta_{ij}) : j \in Ch(i)\}$ that $\{U_{ij}^* : j \in Ch(i)\}$ is a set of identically distributed random variables conditionally on $\mathbf{Z}_i^*(T_E)$. Therefore, the conditional expectation of $V_{ij}(T_E)$ given $\mathbf{Z}_i^*(T_E)$ is identical for all $j \in Ch(i)$. Thus it follows from (8) that

$$\mathbb{E}[V_{ij}(T_E) \mid \mathbf{Z}_i^*(T_E)] = 0 \text{ for all } j \in Ch(i).$$

Since $\alpha_{ij}p_{ij}^{(r)} = (\alpha_{ij}/\beta_{ij}) \times \beta_{ij}p_{ij}^{(r)}$, we have

$$\begin{aligned} & \mathbb{E}[V_{ij}(T_E) \mid \mathbf{Z}_i(T_E)] \\ &= \mathbb{E}[\mathbb{E}[V_{ij}(T_E) \mid \mathbf{Z}_i^*(T_E)] \mid \mathbf{Z}_i(T_E)] = 0. \end{aligned}$$

This completes the proof, since

$$\begin{aligned} & \mathbb{E}[Y_i(T_E^*)] - \mathbb{E}[Y_i(T_E^{(r)})] \\ &= \mathbb{E} \left[\sum_{j \in Ch(i)} \alpha_{ij} p_{ij}^{(T_E)} V_{ij}(T_E) \right] \\ &= \mathbb{E} \left[\sum_{j \in Ch(i)} \alpha_{ij} p_{ij}^{(T_E)} \mathbb{E}[V_{ij}(T_E) \mid \mathbf{Z}_i(T_E)] \right] = 0. \end{aligned}$$

□

A.4 Proof of Theorem 1

It follows from Assumption 2 that

$$\mathbb{E}[Y_i(T_E^*) \mid Z_i(T_E^*) = z] = \mathbb{E}[Y_i(z, \mathbf{W}_i(T_E^*))],$$

where $\mathbf{W}_i(T_E^*) = \{p_{ki}^* : k \in Pa(i)\}$. By design, $\mathbf{W}_i(T_E^*) = \mathbf{W}_i(T_E^{(r)})$ for all $i \in \Omega_r$. Therefore, for $i \in \Omega_r$,

$$\begin{aligned} & \mathbb{E} \left[Y_i(T_E^*) \frac{f_r(Z_i(T_E^*))}{f_r^*(Z_i(T_E^*))} \right] \\ &= \int \mathbb{E} \left[Y_i(z, \mathbf{W}_i(T_E^{(r)})) \frac{f_r(z)}{f_r^*(z)} \mid Z_i(T_E^*) = z \right] f_r^*(z) dz \\ &= \int \mathbb{E} \left[Y_i(z, \mathbf{W}_i(T_E^{(r)})) \right] f_r(z) dz \\ &= \mathbb{E}[Y_i(T_E^{(r)})] = \tau_r. \end{aligned}$$

This completes the proof, since we have

$$\mathbb{E}[\hat{\tau}_r] = \frac{1}{\Omega_r} \sum_{i \in \Omega_r} \mathbb{E}[Y_i(T_E^{(r)})] = \tau_r.$$

□

A.5 Density Estimation

Since we observe $\{Z_i(T_E^*) : i \in \Omega_r\}$, any parametric or non-parametric density estimation method can be applied for estimating f_r^* . For estimating f_r , we make the following additional assumptions:

Assumption 3. Assume that $f_r(z) \propto f_r^*((\sigma_r^*(z + \mu_r^*) - \mu_r)/\sigma_r)$ for all $r \in [m]$, where $\mu_r = \mathbb{E}[Z_i(T_E^{(r)})]$, $\mu_r^* = \mathbb{E}[Z_i(T_E^*)]$, $\sigma_r^2 = \text{Var}[Z_i(T_E^{(r)})]$ and $\sigma_r^{*2} = \text{Var}[Z_i(T_E^*)]$.

Assumption 3 states that the functional form of the densities of $Z_i(T_E^*)$ and $Z_i(T_E^{(r)})$ are identical, except for the mean and the variance. This allows us to obtain an estimate of f_r by adjusting the mean and the variance of an estimated f_r^* . In Lemma 2, we provide consistent estimators of the first two moments of $Z_i(T_E^{(r)})$ from the data $\{Z_{ij}(T_E^{(r)}) : i \in \Omega_r, j \in Ch(i) \cap (\Omega_r \cup \Lambda_r)\}$, where Ω_r and Λ_r are as in Algorithm 1. Finally, note that Assumption 3 is not a theoretical necessity, but a practical recommendation for avoiding overfitting in estimating the target density.

Lemma 2. Assume that

$$\rho_i := \frac{|Ch(i) \cap (\Omega_r \cup \Lambda_r)|}{|Ch(i)|} \text{ and } \rho_i' := \frac{(|Ch(i) \cap (\Omega_r \cup \Lambda_r) - 1|)}{(|Ch(i)| - 1)}$$

are independent of $|Ch(i)|$, $\mu_{r|i} := \mathbb{E}[Z_i(T_E^{(r)}) \mid i]$ and $\sigma_{r|i}^2 := \text{Var}[Z_i(T_E^{(r)}) \mid i]$. We define

$$\begin{aligned} V_1^{(r)} &:= \sum_{i \in \Omega_r} \sum_{j \in Ch(i) \cap (\Omega_r \cup \Lambda_r)} Z_{ij}(T_E^{(r)}), \\ V_2^{(r)} &:= \sum_{i \in \Omega_r} \sum_{j \in Ch(i) \cap (\Omega_r \cup \Lambda_r)} Z_{ij}(T_E^{(r)})^2 \text{ and} \\ V_3^{(r)} &:= \sum_{i \in \Omega_r} \left(\sum_{j \in Ch(i) \cap (\Omega_r \cup \Lambda_r)} Z_{ij}(T_E^{(r)}) \right)^2. \end{aligned}$$

Then

$$\frac{V_1^{(r)}}{\sum_{i \in \Omega_r} \rho_i} \xrightarrow{a.s.} \mathbb{E}[Z_i(T_E^{(r)})] \text{ and } \frac{V_2^{(r)}}{\sum_{i \in \Omega_r} \rho_i} + \frac{V_3^{(r)} - V_2^{(r)}}{\sum_{i \in \Omega_r} \rho_i \rho'_i} \xrightarrow{a.s.} \mathbb{E}[Z_i(T_E^{(r)})^2].$$

Proof. Let $d_i = |Ch(i)|$ denote the out-degree of member i . Then the expected value of $Z_i(T_E^{(r)})$ is given by

$$\mathbb{E}[Z_i(T_E^{(r)})] = \lim_{|\Omega| \rightarrow \infty} \frac{1}{|\Omega|} \sum_{i \in \Omega} d_i \mu_{r|i} = \mathbb{E}[d_i \mu_{r|i}]. \quad (9)$$

Therefore,

$$\begin{aligned} \lim_{|\Omega_r| \rightarrow \infty} \frac{V_1^{(r)}}{\sum_{i \in \Omega_r} \rho_i} &= \lim_{|\Omega_r| \rightarrow \infty} \frac{\frac{1}{|\Omega_r|} \sum_{i \in \Omega_r} \rho_i d_i \mu_{r|i}}{\frac{1}{|\Omega_r|} \sum_{i \in \Omega_r} \rho_i} \\ &= \frac{\mathbb{E}[\rho_i d_i \mu_{r|i}]}{\mathbb{E}[\rho_i]} = \mathbb{E}[Z_i(T_E^{(r)})] \end{aligned}$$

where the last equality follows from (9) and the fact that ρ_i and $\mu_{r|i} d_i$ are independently distributed.

Next, it follows from similar calculations that

$$\begin{aligned} \mathbb{E}[Z_i(T_E^{(r)})^2] &= \mathbb{E}[d_i \sigma_{r|i}^2 + d_i^2 \mu_{r|i}^2] \\ \lim_{|\Omega_r| \rightarrow \infty} \frac{1}{|\Omega_r|} V_2^{(r)} &= \mathbb{E}[\rho_i d_i \sigma_{r|i}^2 + \rho_i d_i^2 \mu_{r|i}^2], \text{ and} \\ \lim_{|\Omega_r| \rightarrow \infty} \frac{1}{|\Omega_r|} V_3^{(r)} &= \mathbb{E}[\rho_i d_i \sigma_{r|i}^2 + \rho_i^2 d_i^2 \mu_{r|i}^2] \end{aligned}$$

Furthermore, it is easy to verify that $\rho_i^2 d_i^2 - \rho_i d_i = \rho_i \rho'_i (d_i^2 - d_i)$. Therefore,

$$\lim_{|\Omega_r| \rightarrow \infty} \frac{1}{|\Omega_r|} [V_3^{(r)} - V_2^{(r)}] = \mathbb{E}[\rho_i \rho'_i (d_i^2 - d_i) \mu_{r|i}^2].$$

Using the independence of (ρ_i, ρ'_i) and $(d_i, \mu_{r|i}, \sigma_{r|i})$, we obtain

$$\begin{aligned} \lim_{|\Omega_r| \rightarrow \infty} \frac{\frac{1}{|\Omega_r|} V_3^{(r)} - V_2^{(r)}}{\frac{1}{|\Omega_r|} \sum_{i \in \Omega_r} \rho_i \rho'_i} &= \mathbb{E}[d_i^2 \mu_{r|i}^2 - d_i \mu_{r|i}^2] \text{ and} \\ \lim_{|\Omega_r| \rightarrow \infty} \frac{\frac{1}{|\Omega_r|} V_2^{(r)}}{\frac{1}{|\Omega_r|} \sum_{i \in \Omega_r} \rho_i} &= \mathbb{E}[d_i \sigma_{r|i}^2 + d_i \mu_{r|i}^2]. \end{aligned}$$

Hence, we have,

$$\begin{aligned} &\lim_{|\Omega_r| \rightarrow \infty} \frac{V_2^{(r)}}{\sum_{i \in \Omega_r} \rho_i} + \frac{V_3^{(r)} - V_2^{(r)}}{\sum_{i \in \Omega_r} \rho_i \rho'_i} \\ &= \mathbb{E}[d_i \sigma_{r|i}^2 + d_i \mu_{r|i}^2] + \mathbb{E}[d_i^2 \mu_{r|i}^2 - d_i \mu_{r|i}^2] \\ &= \mathbb{E}[d_i \sigma_{r|i}^2 + d_i^2 \mu_{r|i}^2] = \mathbb{E}[Z_i(T_E^{(r)})^2]. \end{aligned}$$

□

A.6 Bootstrap Variance Estimation

We draw B random samples with replacement $\{\Omega'^{(1)}, \dots, \Omega'^{(B)}\}$ of size $|\Omega'|$ from Ω' . For each $t \in \{1, \dots, B\}$ and $r = 0, \dots, m$, we obtain $\Omega_r^{(t)}$ by selecting all elements of $\Omega'^{(t)}$ that are in Ω_r . Then we obtain $\hat{\tau}_r^{(t)}$ by applying the density estimation and the estimate computation of Algorithm 2 with $\Omega_r^{(t)}$ instead of Ω_r . Finally, we obtain the bootstrap variance $\hat{\sigma}^2(T_E^{(r)})$ by computing the sample variance of $\hat{\tau}_r^{(1)}, \dots, \hat{\tau}_r^{(B)}$. Finally, note that the bootstrap parameter B should be large enough to have a stable variance estimator. We found $B=1000$ to be good enough for our simulations and real-world experiments.

A.7 Simulation Setup Details

A.7.1 Graph Construction

For our simulation, we want a graph composed of a set of pre-determined clusters with significantly high intra-cluster connectivity, and sparse inter-cluster connectivity. To that end, we first obtain an undirected graph $\mathcal{G}_{BA}^{undir} = (\Omega, E_{BA}^{undir})$ by combining 10 randomly generated graphs with 5000 vertices and average degree equals \bar{d}_{BA} , where each graph is generated according to the Barabasi-Albert model [7] with the power of the preferential attachment equals 0.25. Then we generate an Erdős-Rényi graph [12] $\mathcal{G}_{ER}^{undir} = (\Omega, E_{ER}^{undir})$ with 50000 vertices and average degree equals \bar{d}_{ER} , where all pairs of nodes have an equal probability of being connected. We obtain $\mathcal{G}^{undir} = (\Omega, E_{BA}^{undir} \cup E_{ER}^{undir})$. Finally, we transform \mathcal{G}^{undir} to a directed graph \mathcal{G} by replacing each undirected edge $i - j$ by two directed edges $i \rightarrow j$ and $j \rightarrow i$. Note that the average degree (of the incoming edges) in \mathcal{G} is given by $\bar{d}_G := \frac{1}{|\Omega|} \sum_{i \in \Omega} |Pa(i)| = \bar{d}_{BA} + \bar{d}_{ER}$.

For the simulation study, we choose the following three settings:

d_{BA}	d_{ER}	\bar{d}_G	\bar{d}_{ER}/\bar{d}_G	Sparsity	Impurity (inter-cluster edge proportion)
20	1	21	1/20	High	Low
50	5	55	1/10	Medium	Medium
80	40	120	1/3	Low	High

Table 2: Simulation Settings.

A.7.2 Cluster-based method (CB)

We take advantage of the presence of the disjoint connected components in \mathcal{G}_{BA} to define a "oracle" cluster-based method. More precisely, we randomly choose two clusters to assign them $T_E^{(0)}$ and $T_E^{(1)}$ by defining

$$p_{ij}^* := \begin{cases} p_{ij}^{(r)} & \text{if } i \in Pa(j), j \in \mathcal{H}_r, r = 0, 1 \\ p_{ij}^{base} & \text{if } i \in Pa(j), j \in \Omega \setminus (\mathcal{H}_0 \cup \mathcal{H}_1). \end{cases}$$

where \mathcal{H}_0 and \mathcal{H}_1 are randomly chosen disjoint connected components of \mathcal{G}_{BA} .

A.7.3 Data Generation Mechanism

For each directed edge (i, j) in \mathcal{G} , we independently generate U_{ij} and V_{ij} from $Uniform[10, 100]$ and $Uniform[1, 2]$ respectively. We define $\alpha_{ij} = U_{ij}/d_j$ and $p_{ij}^{base} = V_{ij} / \sum_{i \in Pa(j)} V_{ij}$, where $d_j = |Pa(j)|$ denotes the degree of node j . We define $T_E^{(0)} = \{p_{ij}^{(0)} : (i, j) \in E\}$ and $T_E^{(1)} = \{p_{ij}^{(1)} : (i, j) \in E\}$ where $p_{ij}^{(0)} = p_{ij}^{base}$ and $p_{ij}^{(1)} \propto p_{ij}^{base} \times \left(\frac{\alpha_{ij}}{\log(1+d_i d_j)}\right)^{1/2}$, where d_i is the degree of member i . Finally, we define

$$W_i^{(r)} = \frac{1}{d_i} \sum_{k \in Pa(i)} \alpha_{ki} p_{ki}^{(r)}, \quad Z_i^{(r)}(\delta) = \sum_{j \in Ch(i)} \alpha_{ij} (p_{ij}^{(r)})^\delta \text{ and} \\ Y_i(T_E^{(r)}, \delta) = g(W_i^{(r)} + Z_i^{(r)}(\delta)(1 + W_i^{(r)})) + \epsilon_i$$

where $g(x) = 10/(1 + \exp(-x/10))$ and $\epsilon_i \stackrel{\text{i.i.d}}{\sim} N(0, 1)$. Note that Assumption 1 is satisfied if and only if $\delta = 1$ and α_{ij} 's are known. To verify robustness of our method, we choose $\delta \in \{1/4, 1/2, 1\}$ and we choose $\alpha_{ij} = 1$ for constructing T_E^* using Algorithm 1.

To match with sample size of the cluster-based method, we choose $|\Omega_0| = |\Omega_1| = 0.1 \times |\Omega|$. Furthermore, we choose $|\Lambda_0| = |\Lambda_1| = 0.1 \times |\Omega|$. For choosing C' in Algorithm 1, we set $q = 0.5$. For setting the optimization constraints, we choose $R_{\min} = 0$, $R_{\max} = 10$, $S_{\min} = 0.2$, $S_{\max} = 5$. For running the iterative QP, we set $K = 1000$ and $maxIter = 10$.

A.8 Real-World Experimental Setup

We consider a directed network $\mathcal{G} = (\Omega, E)$ of all LinkedIn members (675+ million) with edges $i \rightarrow j$ if at least one content of member i is considered by the baseline recommender system to be shown to member j in the last 7 days. First, we randomly choose five disjoint sets subsets $\Omega_0, \Omega_1, \Lambda_0, \Lambda_1, \Gamma$ such that $|\Omega_0| = |\Omega_1| = 0.03 \times |\Omega|$ and $|\Lambda_0| = |\Lambda_1| = 0.04 \times |\Omega|$ and $|\Gamma| = 0.06 \times |\Omega|$. Next, we choose $C' = \Gamma \cap (\cup_{i \in \Omega_0 \cup \Omega_1} Ch(i))$. The reason for choosing C' in a slightly different manner is to make C' well-defined even if the underlying network is dynamic over the experiment time window.

A.8.1 Boost factors b_{ij}

It follows from a straightforward calculation that the following definition of b_{ij} 's for $j \in C'$ will satisfy (3):

$$b_{ij} := \begin{cases} \frac{p_{ij}^* (1 - \sum_{k \in \Omega' \cap Pa(j)} p_{kj}^{(0)})}{p_{ij}^{(0)} (1 - \sum_{k \in \Omega' \cap Pa(j)} p_{kj}^*)} & \text{if } i \in Pa(i) \cap \Omega' \\ 1 & \text{if } i \in Pa(i) \setminus \Omega'. \end{cases}$$