

A Design-Based Minimax Theory for Network Experiments

Vardis Kandiros¹, Christopher Harshaw¹, and Fredrik Sävje²

¹*Columbia University*

²*Uppsala University*

August 5, 2026

Abstract

Network experiments are used throughout the social and medical sciences to investigate causal effects under the presence of interference. While a large body of work has developed improved statistical procedures, the fundamental limits of statistical estimation in these settings is less well understood. In this paper, we develop and investigate a design-based theory of minimax risk for network experiments under an arbitrary neighborhood interference model. Our notion of minimax risk describes the optimal precision among all statistical procedures for investigating a particular causal effect on the observed interference network. We show that the minimax risk is a function of the corresponding conflict graph, which captures inherent unobservability of estimand-relevant potential outcomes given the observed interference network. Our main contribution is a series of upper and lower bounds on the minimax rate in terms of local and global connectivity properties of the conflict graph. To illustrate their utility, we apply these general results to obtain minimax analyses for two commonly studied effects: the direct treatment effect and global average treatment effect.

We thank James M. Robins for insightful discussions which helped to shape this work. The authors gratefully acknowledge support from the following agencies: NSF Grant MMS-2316335; Jan Wallander, Tom Hedelius & Tore Browaldh foundations Grant P25-0067; Swedish Research Council Grant 2025-04794. Part of this research was conducted at the Isaac Newton Institute for Mathematical Sciences.

Table of Contents

1	Introduction	1
1.1	Illustration of Main Results	2
1.2	Related Work	4
2	Minimax Theory for Network Experiments	5
2.1	Preliminaries: Model and Estimands	5
2.2	Minimax Risk for Estimating Causal Effects	7
2.3	Conflict Graph	9
3	Lower Bounds on the Minimax Risk	11
3.1	Le Cam’s Method for Design-Based Inference	11
3.2	Lower Bound using Global Information	12
3.3	Lower Bound using Local Information	14
4	Upper Bounds on the Minimax Risk	17
4.1	Unbiased Statistical Procedures	17
4.2	Bias-Variance Tradeoff	19
5	Tightness of Minimax Analysis	20
5.1	Direct Treatment Effect	21
5.2	Global Average Treatment Effect	22
6	Conjectured Computational Hardness	22
7	Conclusion	24
	Appendix	28

1 Introduction

Network interference experiments are used in the social and medical sciences to investigate spillover effects and other causal interactions between units. An important problem for the execution of such an experiment is the choice of the experimental design and the causal effect estimator, which we collectively refer to as the *statistical procedure*. The choice is crucial because it determines how much can be learned from the experiment. An ill-advised choice can introduce both bias and imprecision, often rendering the estimator inconsistent unless the underlying network is very sparse.

A growing body of literature has sought to shine light on this question. The typical line of attack is the development of a new statistical procedure for network experiments that can be shown to provide improved rates of convergence under suitable circumstances and restrictions on the underlying interference network. Examples include experimental designs based on cluster randomization (Ugander et al., 2013; Ugander and Yin, 2023; Viviano et al., 2026), independent sets (Karwa and Airolidi, 2018; Jagadeesan, Pillai, and Volfovsky, 2020; Fatemi and Zheleva, 2020), and conflict-resolution schemes (Kandiros et al., 2024).

While significant progress has been made with this approach, the literature is stifled by a lack of a theoretical benchmark: what are the statistically optimal rates of estimation for network experiments? The question of rate optimality is not only of theoretical interest, but it also has practical implications. Without an understanding of the fundamental statistical limits of what can be learned from these experiments, both statisticians and practitioners remain unaware of how near or far existing statistical procedures are from being optimal. If it is found that existing procedures are far from the fundamental optimal limit, we should be incentivized to continue working on new procedures, possibly trying novel approaches.

In this paper, we develop and investigate a design-based theory of minimax optimality for network experiments. The proposed notion of minimax optimality describes the best precision among all statistical procedures for investigating a particular causal effect τ on an observed interference network G , providing a fundamental statistical limit of the empirical problem at hand. Our notion of minimax optimality has two key features. First, optimality is considered with respect to a single observed interference network G , considering all potential outcomes obeying the restrictions imposed by the network. Second, optimality is with respect to all possible choices of experimental design and estimator without any restrictions. We believe that these two features reflect the challenges faced by practitioners, who are often able to select any statistical procedure after measuring the interference network. Moreover, this fills a gap in the design-based literature; to the best of our knowledge, minimax analyses which account for both experimental design and estimator have not previously been investigated in design-based causal inference.

The main contribution of the paper is an analysis of the minimax rate of causal effect estimation under the assumption of arbitrary neighborhood interference (ANI), where the outcome of a unit depends arbitrarily on the treatment of its neighbors. Central to our analysis is the notion of an exposure-level conflict graph, which is closely related to the unit-level conflict graph described by Kandiros et al. (2024). The vertices in the conflict graph $\mathcal{H}$ are all potential outcomes relevant for the estimand τ at hand, and an edge is drawn between two vertices if the corresponding potential outcomes cannot be observed simultaneously given the interference network G . We show that the minimax risk of the network experiment problem is in fact a function of the conflict graph $\mathcal{H}$, albeit a complicated one. Thus, the central question becomes: what aspects of the conflict graph determine the minimax rate?

Our primary analysis appears in a series of lower and upper bounds on the minimax risk presented

throughout the paper. We collect these results here, presented in a slightly simplified version.

Main Result (Informal). *Let G be an interference network, τ be a causal estimand, and $\mathcal{H}$ be the corresponding conflict graph. The minimax risk for estimating τ on network G is bounded as*

$$\max\left(\frac{1}{|\mathcal{I}(\mathcal{H})|}, \frac{\sqrt{d^*(\mathcal{H})}}{n}\right) \lesssim \inf_{\mathcal{D}, \hat{\tau}} \sup_y \mathbb{E}[(\hat{\tau} - \tau)^2] \lesssim \frac{\min_S Q(S, \mathcal{H})}{n},$$

where the infimum is taken over all possible designs $\mathcal{D}$ and estimators $\hat{\tau}$, and the supremum is taken over all potential outcomes y having bounded second moments and satisfying the network interference assumption. The expectation is taken with respect to the experimental design $\mathcal{D}$.

The lower bounds on minimax risk depend on two quantities describing the conflict graph: its largest independent set $\mathcal{I}(\mathcal{H})$ and its critical degree $d^*(\mathcal{H})$, which represent global and local connectivity information of the graph. More precisely, $\mathcal{I}(\mathcal{H})$ is the largest subset of vertices for which no pair is adjacent in $\mathcal{H}$ and $d^*(\mathcal{H})$ is the largest integer k such that at least $\sqrt{k}$ vertices have degree at least k . The upper bound is the solution to a regularized induced subgraph problem, where for any subset S of vertices, the objective function $Q(S, \mathcal{H}) = \lambda(\mathcal{H} \setminus S) + |S|$ has two terms: the largest eigenvalue of the induced subgraph obtained by removing S and the size of S . This upper bound reflects a bias-variance trade-off and improves over the rate attained by Kandiros et al. (2024). The full analysis considers general moment restrictions, and it uncovers an interplay between the strength of the moment restrictions and degree heterogeneity for the minimax rates, where stronger moment restrictions can accommodate more heterogeneity.

The primary technical innovation facilitating this work is the development of techniques to derive statistical lower bounds in a design-based setting. Our approach begins by adopting Le Cam’s method within the design-based framework, where a mixture distribution over potential outcome functions is constructed to certify the statistical lower bound. In the setting of network experiments, several novel challenges arise necessitating the construction of mixture distributions which are both network-specific and design-adaptive. The lower bound techniques we develop may be of independent interest, and may be useful in other high-dimensional network problems.

1.1 Illustration of Main Results

To illustrate the utility of our general minimax rate analyses, we apply the main result to two causal estimands previously considered in the literature. In what follows, we denote the minimax rate as $\mathcal{R}(G, \tau)$ for a network G and causal effect τ . For simplicity, we here focus on second moment restrictions and present results without considering the aforementioned bias-variance trade-off, effectively setting $S = \emptyset$. All formal definitions are described in Section 2 and proofs of the corollaries in this subsection appear in Section D of the appendix.

We first consider estimation of the Direct Treatment Effect (DTE). The DTE estimand is the effect of treating a single experimental subject when all other subjects are assigned control treatment, averaged over all subjects. The estimand isolates the effect that treatment has directly on the treated unit, removing any potential spillover and interference effects. The following corollary uses the main results above to establish bounds on the minimax optimal rate of estimating DTE in terms of the underlying interference graph G .

Corollary 1.1. *The minimax rate of estimating the Direct Treatment Effect τ_{DTE} on a network G is bounded as*

$$\max\left(\frac{1}{|\mathcal{I}(G)|}, \frac{\sqrt{d_{\text{avg}}(G)}}{n}\right) \lesssim \mathcal{R}(G, \tau_{\text{DTE}}) \lesssim \frac{\lambda(G)}{n},$$

where $\mathcal{I}(G)$ is the maximum independent set in G , $d_{\text{avg}}(G)$ is its average degree, and $\lambda(G)$ is the largest eigenvalue of its adjacency matrix.

The bounds are similar to those that appeared in the general main results, but with the underlying interference network G in place of the conflict graph $\mathcal{H}$. This simplification is possible because $\mathcal{H}$ is tightly coupled with G for the DTE.

The minimax rate is precisely characterized when the upper and lower bounds match; however, there are graphs G for which these bounds do not match. In Section 5, we show that both bounds are tight even within the restricted class of d -regular graphs, in the sense that there exist graphs such that the true minimax rate coincides with the lower bound, and other graphs such that it coincides with the upper bound. In addition, we show that these bounds characterize the minimax risk for random d -regular graphs, up to logarithmic factors.

We next consider estimation of the Global Average Treatment Effect (GATE). The GATE estimand is the effect of treating all subjects versus treating no one. The estimand captures both the direct effect of treatment and spillover effects. The following corollary uses the main results above to establish bounds for the GATE in terms of the two-hop network G^2 , which is obtained by connecting all subjects that are within distance 2 in the original interference network G .

Corollary 1.2. *The minimax rate of estimating the Global Average Treatment Effect τ_{GATE} on a network G is bounded as*

$$\frac{\sqrt{d_{\text{avg}}(G^2)}}{n} \lesssim \mathcal{R}(G, \tau_{\text{GATE}}) \lesssim \frac{\lambda(G^2)}{n},$$

where G^2 is the two-hop interference graph, $d_{\text{avg}}(G^2)$ is the average degree of the two-hop graph, and $\lambda(G^2)$ is the largest eigenvalue of its adjacency matrix.

Because the two-hop network G^2 is typically much more densely connected than the original network G , the minimax analyses of Corollary 1.2 aligns with the commonly held belief that the DTE is easier to estimate than the GATE. In Section 5, we formalize this intuition by showing that the minimax rate of estimating the DTE is at most that of estimating the GATE for random d -regular graphs.

The fact that the derived upper and lower bounds on the minimax rates for network experiments do not always match, in the sense that they are the same up to constant terms, means that they do not fully characterize the fundamental limit of the statistical problem. Although the minimax risk is a function of the conflict graph $\mathcal{H}$, this exact relationship depends on the structure of maximal independent sets of $\mathcal{H}$, which are likely too complex to be succinctly described. In Section 6, we formalize this possibility through a set of well-motivated conjectures about the computational hardness of deriving sharper analyses of the minimax risk and discuss their practical consequences. Nevertheless, we believe the bounds in this paper provide an important step towards the understanding of network experiments as a statistical problem.

1.2 Related Work

We consider experimental design and estimation in network experiments under *arbitrary neighborhood interference*, allowing the outcome of a unit to depend on the treatments of its neighbors in a known graph. This is an instance of the design-based exposure mapping framework developed by Aronow and Samii (2017), which generalizes the ideas developed by Sobel (2006) and Hudgens and Halloran (2008) to allow for outcomes to depend on arbitrary summaries of the treatment vector. The concept of effective treatments described by Manski (2013) is an alternative formulation.

Experimental design under arbitrary neighborhood interference has been considered previously. Ugander et al. (2013) and Ugander and Yin (2023) consider estimation of the Global Average Treatment Effect (GATE) under uniformly bounded outcomes and describe designs that split the graph in clusters used for the assignment of treatment. The work closest to the current paper is Kandiros et al. (2024), who consider arbitrary contrastive effects under bounded second outcome moments and describe an experimental design based on a conflict resolution scheme. Compared to this previous work, our contribution is to derive lower bounds on the minimax rate, which has previously been absent from the literature, and improved upper bounds. We also formalize the study of the minimax rate in the setting of network experiments.

Several authors have considered network experiments under restrictions on the potential outcomes beyond those imposed by the neighborhood interference model itself. For example, Toulis and Kao (2013) describe two designs for assigning treatments in graphs under a linear outcome model. A long line of work has explored other types of additional outcome assumptions in network experiments (Sussman and Airoldi, 2017; Fatemi and Zheleva, 2020; Jagadeesan, Pillai, and Volfovsky, 2020; Cai, Zhang, and Airoldi, 2024; Viviano et al., 2026). The purpose of these assumptions is to make potential outcomes that are irrelevant under the neighborhood interference model alone informative of the causal effect of interest, simplifying the experimental design and estimation problems.

Our concept of minimax risk and optimality differs from previously considered concepts in the experimental design literature. Previously, analyses have typically fixed the estimator and considered a restricted class of designs (Wu, 1981; Baird et al., 2018; Kallus, 2018; Leung, 2022; Bojinov, Simchi-Levi, and Zhao, 2023; Candogan, Chen, and Niazadeh, 2024). Some analyses instead restrict the design and consider a larger class of estimators (see, e.g., Cortez-Rodriguez, Eichhorn, and Yu, 2023). This type of analysis demonstrates optimality within the restricted class of designs or estimators, but not optimality in a universal sense. We consider the choice of estimator and design simultaneously without restrictions, making the analysis informative of universal optimality. Concepts of optimality not building on minimax performance have also been considered. For example, Basse and Airoldi (2018) consider designs that are optimal with respect to a generative working model, and Karwa and Airoldi (2023) consider admissibility of linear exposure effect estimators under interference.

Within the superpopulation framework for causal inference, statistical optimality and minimax analysis have been studied extensively. The data generating process is beyond the researcher’s control here, so optimality is considered only over possible estimators. The predominate focus in this literature is on *local minimax risk* (van der Vaart, 2000), which considers the best possible asymptotic variance that can be achieved by an asymptotically unbiased estimator. In semiparametric efficiency theory, this is done using the efficient influence curve (Bickel et al., 1993). The approach is local because it considers perturbations of the parameter of interest locally in the parameter space. We instead use a global notion of minimax risk, considering the worst-case performance of a statistical procedure for all potential outcome functions in some class of functions. The global notion is suitable

to capture the optimal *rate* of estimating an effect, without consideration for the exact constants involved. A global notion of minimax risk has been considered also in the superpopulation setting to study optimal rates (see, e.g., Robins et al., 2008; Rotnitzky, Smucler, and Robins, 2021; Foster and Syrgkanis, 2023; Hu and Wager, 2023; Kennedy et al., 2024).

Within the survey sampling literature, several minimax optimality results for finite populations have been developed (Rinott, 2009). For example, Aggarwal (1959) and Bickel and Lehmann (1981) established that the sample mean is minimax optimal under the constraint that the variance of the outcomes is bounded, and Hodges and Lehmann (1982) showed that a slightly biased version of the sample mean is minimax optimal under the constraint that all outcomes are uniformly bounded. More recently, Aronow and Lopatto (2026) derived lower bounds in a setting with unequal inclusion probabilities under uniformly bounded outcomes. While the underlying structure of the statistical problems are similar, the techniques for establishing minimax optimality in survey sampling are not directly applicable to causal inference problems. The survey sampling techniques rely on classical decision theory arguments based on symmetry of the problem with respect to the unit, allowing the proof to be divided into two parts, establishing the optimal design separately from the optimal estimator (Blackwell and Girshick, 1979; Gabler, 2012). No such symmetry exists in the causal inference setting, and we must consider the choice of design and estimator simultaneously. For this reason, there is no obvious way to apply these approaches to the problems we consider in this paper. The techniques we describe here are developed from first principles and unrelated to those used in the survey sampling literature.

2 Minimax Theory for Network Experiments

2.1 Preliminaries: Model and Estimands

We consider an experimental study with n subjects indexed by integers $\{1, 2, \dots, n\} \triangleq [n]$. The experimenter selects a random intervention z from a set of interventions Ω and the randomized selection mechanism is referred to as the experimental design. Formally, the *interventions* are a measurable space $(\Omega, \mathcal{F})$ and the *experimental design* is a probability measure $\mathcal{D} : \mathcal{F} \rightarrow [0, 1]$. We focus on the setting where intervention is a binary treatment to each experimental subject, i.e. $z = (z_1, z_2, \dots, z_n)$ with $z_i \in \{0, 1\}$ so that $\Omega = \{0, 1\}^n$.

Each subject $i \in [n]$ has an associated *potential outcome function* $y_i : \Omega \rightarrow \mathbb{R}$ which describes its responses under each intervention. For notational convenience, we collect these n functions into an *entire potential outcome function* $y : \Omega \rightarrow \mathbb{R}^n$ defined as $y(z) = (y_1(z), \dots, y_n(z))$. In the randomized experiment, the experimenter draws a random intervention Z according to the design $\mathcal{D}$ and observes the n outcomes $y_1(Z), \dots, y_n(Z)$.

In order to facilitate meaningful inference of causal effects, the experimenter must place further restrictions on the potential outcome functions. A common choice is the arbitrary neighborhood interference (ANI) model, which depends on an observed interference network. We presume that the experimenter has access to an undirected network $G = (V, E)$ whose vertices V are experimental subjects with edge set E . For a subject $i \in [n]$, we define its *neighborhood* in the network to be the other subjects to which it is connected: $N(i) = \{j \in [n] : (i, j) \in E\}$. For notational convenience, we define the *extended neighborhood* as $\tilde{N}(i) = N(i) \cup \{i\}$, which is just a unit's neighbors and itself. The ANI model stipulates that a subject's potential outcome function can depend arbitrarily on treatments assigned to subjects in its extended neighborhood, but is not affected by treatment assigned to other subjects.




Figure 1: An illustration of the relevant exposures that define the global and direct effects. Red and blue correspond to treated and untreated units.

Definition 1. The *arbitrary neighborhood interference* model $\mathcal{M}(G)$ contains all functions $y : \Omega \rightarrow \mathbb{R}^n$ satisfying the following condition: $y_i(z) = y_i(z')$ for all subjects $i \in [n]$ and interventions $z, z' \in \Omega$ with $z_j = z'_j$ for all $j \in \tilde{N}(i)$.

The ANI model is an example of an exposure mapping model (Aronow and Samii, 2017). It will sometimes be convenient to work with the underlying exposure mappings, so we define them now. For each subject $i \in [n]$, its *exposure set* Δ_i is defined to contain all subsets of its extended neighborhood, i.e. $\Delta_i = \mathcal{P}(\tilde{N}(i))$. The exposure mapping $h_i : \Omega \rightarrow \Delta_i$ maps an intervention z to the subset of extended neighbors which were treated under z , i.e. $h_i(z) = \{j \in \tilde{N}(i) : z_j = 1\}$. In an exposure mapping model, a unit's potential outcome function is constant across level sets of the exposure mapping, i.e. if $h_i(z) = h_i(z')$ then $y_i(z) = y_i(z')$. In this way, we can see that the ANI model is an exposure mapping model, where the exposure is the specific configuration of extended neighbors which are treated. In line with the previous literature, we overload the potential outcome notation and write $y_i(e)$ to denote the potential outcome that would be observed for unit i under exposure $e \in \Delta_i$.

We consider the class of contrastive causal effects, which are defined as the average over subjects of the contrast between two different potential outcomes. Formally, a *contrastive effect* τ is a linear functional $\tau : \mathcal{M}(G) \rightarrow \mathbb{R}$ which is specified by $2n$ unit-specific exposures $\{(e_{i,1}, e_{i,0})\}_{i=1}^n$ and defined as

$$\tau(y) = \frac{1}{n} \sum_{i=1}^n y_i(e_{i,1}) - y_i(e_{i,0}) .$$

To avoid notational clutter, we do not parameterize τ by the choice of exposures. We review two of the most well-studied examples of contrastive effects below:

- **Global Average Treatment Effect:** The *global average treatment effect* (GATE) is defined by the following two exposures: $e_{i,1} = \tilde{N}(i)$, which corresponds to all subjects in the extended neighborhood receiving treatment, and $e_{i,0} = \emptyset$, which corresponds to all subjects in the extended neighborhood receiving control. The GATE is the effect of treating everyone versus treating no one. This may be relevant for policy makers which intend to roll out treatment to everyone.
- **Direct Treatment Effect:** The *direct treatment effect* (DTE) is defined by the following two exposures: $e_{i,1} = \{i\}$, which corresponds to the subject i receiving treatment while all of its neighbors remain untreated, and $e_{i,0} = \emptyset$, which corresponds to both the subject and its neighbors receiving control. The DTE is the effect of treatment when a subject is the only

one to get it. This may be more relevant for an individual’s decision if they believe that the crowd will not adopt treatment.

See Figure 1 for visual depictions of these two effects and their constituent exposures. Without loss of generality, we only consider estimands for which $e_{i,1} \neq e_{i,0}$. To see this is without loss of generality, observe that if $e_{i,1} = e_{i,0}$ then $y_i(e_{i,1}) = y_i(e_{i,0})$ so that subject i does not contribute to the causal effect and may be removed from consideration.

The causal estimand depends on the entire potential outcome function which is unknown to the experimenter even after the experiment is run, a challenge often referred to as the “fundamental problem of causal inference” (Holland, 1986). For this reason, the experimenter must estimate the causal effect from the observed data. Recall that the experimenter draws a random intervention $Z \in \Omega$ according to the design $\mathcal{D}$ and observes the n outcomes $y_1(Z), \dots, y_n(Z)$. Formally, we define an *estimator* as a measurable function $\hat{\tau} : \Omega \times \mathbb{R}^n \rightarrow \mathbb{R}$ which takes as input the realized intervention and observed outcomes and produces an estimate of the causal effect. Because the researcher can select both the design $\mathcal{D}$ and estimator $\hat{\tau}$, we refer to the pair $(\mathcal{D}, \hat{\tau})$ jointly as a *statistical procedure*.

When analyzing the precision of a statistical procedure, it is necessary to place restrictions on the magnitude of the potential outcomes, because the precision can be arbitrarily poor if the magnitude of the potential outcomes is sufficiently large. To this end, we impose further moment restrictions on the ANI model, which may be of arbitrary order. For any $q \in [2, \infty]$, define the q th *moment restricted ANI model* as

$$\mathcal{M}(G, \tau, q) = \left\{ y \in \mathcal{M}(G) : \left(\frac{1}{n} \sum_{i=1}^n |y_i(e_{i,k})|^q \right)^{1/q} \leq 1 \text{ for } k \in \{0, 1\} \right\} .$$

The q th moment restricted model $\mathcal{M}(G, \tau, q)$ only contains potential outcome functions in the ANI model $\mathcal{M}(G)$ for which the q th moments of the estimand-relevant potential outcomes are bounded. In defining this restricted model, we have made two particular choices: we have assumed that the q th moments are bounded by 1, and we have placed no restrictions on the potential outcomes which are not relevant for the estimand. In Appendix A.4, we show that both of these choices are without loss of generality when studying the minimax risk. First, we show that replacing the moment restriction 1 with any fixed constant $C > 0$ scales the minimax risk by a multiplicative factor of C^2 , so that we may consider $C = 1$ without loss of generality. Second, we show that the minimax risk is unaffected if we were to place extreme restrictions on the estimand-irrelevant outcomes, so that they were effectively presumed to be known. Because the outcomes that are irrelevant for the estimand provide no information about the causal effect under an ANI model, the extent to which they are restricted is irrelevant for the minimax risk.

2.2 Minimax Risk for Estimating Causal Effects

In this paper, we use the mean squared error to measure the precision, or risk, of a statistical procedure. Formally, the *mean squared error* (MSE) of a design-estimator pair $(\mathcal{D}, \hat{\tau})$ for a causal effect τ under potential outcome function y is defined as

$$\text{MSE}(\mathcal{D}, \hat{\tau}, \tau, y) = \mathbb{E}_{Z \sim \mathcal{D}} \left[(\hat{\tau}(Z, y(Z)) - \tau(y))^2 \right] .$$

The MSE depends not only on the design $\mathcal{D}$ and estimator $\hat{\tau}$, but also the potential outcome function y under consideration. Because the underlying potential outcome function y is unknown to the

experimenter, the MSE of a statistical procedure cannot be evaluated before the experiment is run. In order to evaluate different statistical procedures without exact knowledge of the potential outcome functions, we consider its worst case MSE within a set of potential outcome functions, in our case the moment-restricted interference model. More precisely, the *maximum risk* of a statistical procedure within a q -th moment restricted ANI model is defined as

$$\sup_{y \in \mathcal{M}(G, \tau, q)} \mathbb{E}_{Z \sim \mathcal{D}} \left[(\hat{\tau}(Z, y(Z)) - \tau(y))^2 \right] .$$

The maximum risk evaluates a design-estimator pair by comparing its worst-case performance among all potential outcome functions in the q -th moment restricted ANI model, $y \in \mathcal{M}(G, \tau, q)$. Unlike MSE, the maximum risk can, in principle, be evaluated before the experiment is run and can therefore be used as a criterion for selecting a design-estimator pair.

We define the minimax risk as the smallest achievable maximum risk when considering all possible choices of design and estimator. Formally, the *minimax risk* depends on the causal effect τ , the underlying network G , and the moment restriction q , and is defined as

$$\mathcal{R}(\tau, G, q) = \inf_{\mathcal{D}, \hat{\tau}} \mathcal{R}(\mathcal{D}, \hat{\tau}, \tau, G, q) = \inf_{\mathcal{D}, \hat{\tau}} \sup_{y \in \mathcal{M}} \mathbb{E}_{Z \sim \mathcal{D}} \left[(\hat{\tau}(Z, y(Z)) - \tau(y))^2 \right] .$$

The minimax risk $\mathcal{R}(\tau, G, q)$ characterizes the fundamental statistical limits of the experimental problem at hand: no statistical procedure can attain a maximum risk which is smaller than $\mathcal{R}(\tau, G, q)$. In addition to its theoretical motivation, a clear understanding of the minimax MSE can also provide practical guidance. For example, an applied researcher might be satisfied using a statistical procedure whose performance is close to minimax optimal because no other procedure can offer significant improvements. In Section A.1 of the appendix, we show how this notion of minimax risk for ANI models and contrastive effects can be extended to the more general design-based framework of Harshaw, Sävje, and Wang (2022).

The minimax risk obeys several basic monotonicity properties: it increases if the moments are made to be less restrictive or if additional edges are added to the graph. These properties follow directly from the fact that these modifications enlarge the restricted ANI model under consideration. For two graphs G, G' on the same set of vertices, we write $G' \subseteq G$ when every edge of G' is also an edge of G . We present the monotonicity properties below:

Proposition 2.1. *The minimax risk $\mathcal{R}(G, \tau, q)$ obeys the following monotonicity properties:*

- **Monotonicity in q :** *If $q' \leq q$, then $\mathcal{R}(G, \tau, q) \leq \mathcal{R}(G, \tau, q')$.*
- **Monotonicity in G :** *If $G' \subseteq G$, then $\mathcal{R}(G', \tau, q) \leq \mathcal{R}(G, \tau, q)$.*

We have defined the minimax risk in terms of deterministic estimators which use only the data from the experiment; however, it is possible to consider randomized statistical procedures which use an additional source of randomness other than the treatment assignment, e.g. a random number generator. In Section A.3 of the appendix, we show that the minimax risk among all procedures, including randomized ones, is the same as the minimax risk among the restricted class of deterministic procedures, mirroring well-known results from classical decision theory (Lehmann and Casella, 1998). Nevertheless, additional randomness may still be beneficial for constructing simpler and computationally efficient statistical procedures. Indeed, the statistical procedures introduced in

this work are randomized, and their “de-randomizations” appear computationally challenging to implement.

Finally, it may be beneficial to draw comparisons between the design-based perspective and the sampling-based perspective, particularly from the viewpoint of semi-parametric inference. In the classical superpopulation setting, one observes independent samples from some distribution P_θ , parametrized by a high dimensional nuisance component θ . The goal is to estimate some low dimensional functional $\psi(\theta)$ of the nuisance parameter. In design-based inference, one can view the entire potential outcome function y as the high-dimensional nuisance parameter and $\tau(y)$ as the low-dimensional quantity of interest. One major distinction between these perspectives arises in the context of experimental design: the semi-parametric theory typically takes the data generating distribution P_θ as given, and this includes the random assignment of treatment. In the design-based framework, the experimental design $\mathcal{D}$ is decoupled from the nuisance parameter $\tau(y)$ and needs to be chosen by the experimenter. As we will see, this selection of both design and estimator will introduce additional challenges in establishing lower bounds on the minimax risk in the design-based setting.

2.3 Conflict Graph

Our analysis of the minimax risk will prominently feature the conflict graph $\mathcal{H}$, originally introduced by Kandiros et al. (2024) for the purpose of constructing statistical procedures. In this paper, we show that the conflict graph is in fact a central quantity when studying statistical complexity of these network experiment problems.

At a high level, the conflict graph $\mathcal{H}$ encodes the fundamental unobservability between estimand-relevant exposures. Recall the exposures $\{(e_{i,1}, e_{i,0})\}_{i=1}^n$ which define the contrastive estimand. We say that two exposures $e_{i,k}$ and $e_{j,\ell}$ are *in conflict* if there is no intervention under which they can both be observed, i.e. there does not exist $z \in \Omega$ for which $h_i(z) = e_{i,k}$ and $h_j(z) = e_{j,\ell}$. This happens exactly when the two exposures $e_{i,k}$ and $e_{j,\ell}$ specify different treatment assignments to at least one of the units in the intersection of extended neighborhoods $\tilde{N}(i) \cap \tilde{N}(j)$. The conflict graph is formally defined below:

Definition 2. Given a network G and contrastive effect τ , the corresponding *conflict graph* $\mathcal{H} = (V_{\mathcal{H}}, E_{\mathcal{H}})$ contains the estimand-relevant exposures as vertices $V_{\mathcal{H}} = \{(e_{i,1}, e_{i,0})\}_{i=1}^n$ with an edge if two exposures are in conflict.

Note that the conflict graph $\mathcal{H}$ depends on both the underlying network G and the causal effect of interest τ . In this way, any analysis of the minimax risk in terms of the conflict graph $\mathcal{H}$ reflects both G and τ . However, we suppress this dependency to avoid unnecessary notational clutter. We typically consider a generic contrastive effect throughout the paper and make it clear from context when we focus on a specific effect.

To better understand how the structure of the conflict graph $\mathcal{H}$ depends on both the underlying network G and the causal effect τ , we focus on several examples below:

- **No Interference:** Consider the setting of no-interference, in which case we may see the underlying graph G as containing no edges. Suppose we want to estimate the average treatment effect, where $e_{i,1}$ is the exposure that subject i receives treatment and $e_{i,0}$ is the exposure that subject i receives the control. In this case, if $i \neq j$ then there are no conflicts between $e_{i,k}$ and $e_{j,\ell}$ because the individual treatments can be selected separately. On the other hand,




Figure 2: An illustration of a network G and conflict graphs for the DTE and GATE.

there will be a conflict between $e_{i,1}$ and $e_{i,0}$ because a subject cannot be treated and untreated under any intervention. Thus, the conflict graph $\mathcal{H}$ is simply a matching of the $2n$ exposures.

- **Direct Treatment Effect:** Recall that $e_{i,1}$ is the exposure that subject i is treated and all of its neighbors are untreated while $e_{i,0}$ is the exposure that the entire extended neighborhood of subject i receives the control. In this case, there is no edge in the conflict graph $\mathcal{H}$ between all-control exposures $e_{i,0}$ and $e_{j,0}$ since we can observe both exposures by setting all units to control. On the other hand, if at least one of k, ℓ is equal to 1 (i.e. representing the direct exposure), then there is an edge between $e_{i,k}$ and $e_{j,\ell}$ if and only if i, j are connected in G . In this way, the conflict graph $\mathcal{H}$ has the underlying graph G as a subgraph.
- **Global Average Treatment Effect:** Recall that $e_{i,1}$ is the exposure that all extended neighbors are treated while $e_{i,0}$ is the exposure that no extended neighbors are treated. In this case, there are no edges in the conflict graph $\mathcal{H}$ between all-treatment exposures $e_{i,1}$ and $e_{j,1}$ for subjects i and j because they can both be realized by the intervention $z = \mathbf{1}$. Likewise, there are no edges in the conflict graph $\mathcal{H}$ between all-control exposures $e_{i,0}$ and $e_{j,0}$. Instead, the edges in the conflict graph arise between exposures $e_{i,k}$ and $e_{j,\ell}$ where $k \neq \ell$ and subjects i and j have common extended neighbors, i.e. $\tilde{N}(i) \cap \tilde{N}(j) \neq \emptyset$. These exposures are in conflict because both i and j have a common neighbor who would have to be treated and not treated to observe $e_{i,k}$ and $e_{j,\ell}$, which is impossible. In this way, the conflict graph $\mathcal{H}$ is a bipartite graph that resembles the two hop graph G^2 of the underlying graph G .

See Figure 2 for an illustration of how for the same underlying graph G , the resulting conflict graph will typically be different for these two effects. As seen in these examples, the edges in the conflict graph encode the fundamental problem of causal inference and its extension to network interference.

The conflict graph is a summary of the experimental settings, capturing exactly which pairs of exposures are impossible to realize simultaneously under any design. It abstracts away from other details of the experiment, such as the exact definition of the estimand and the interference structure that gave rise to these conflicts. The following proposition demonstrates that this abstraction retains enough information about the statistical problem to exactly characterize the minimax risk of the corresponding network experiment problem. In what follows, we let $\mathbb{H}$ denote the set of all labelled graphs (see Section A.5 in the appendix for a detailed proof).

Proposition 2.2. *Let G be an interference network, τ be a contrastive effect. For each moment parameter $q \geq 2$, there exists a function $f_q : \mathbb{H} \rightarrow \mathbb{R}_{\geq 0}$ such that $\mathcal{R}(G, \tau, q) = f_q(\mathcal{H})$.*

Proposition 2.2 demonstrates that the minimax risk can be expressed as a function of the conflict graph $\mathcal{H}$. A key aspect of its proof is the translation of the network experiment minimax problem into an equivalent minimax problem defined solely in terms of $\mathcal{H}$ through its independence sets. The translation relies on the fact that every experimental design can be identified with a probability measure over independent sets in $\mathcal{H}$. In fact, we show that it suffices to consider probability measures over maximal independent sets in $\mathcal{H}$, corresponding to experimental designs which provide as much information as possible about potential outcomes. The structure of the maximal independent sets in $\mathcal{H}$ therefore determines the minimax risk $\mathcal{R}(G, \tau, q)$, albeit in a complex way. In particular, the function $f_q : \mathbb{H} \rightarrow \mathbb{R}_{\geq 0}$ in Proposition 2.2 is not explicitly constructed, but rather defined through a minimax problem involving $\mathcal{H}$ alone. Our primary analysis in this paper therefore may be understood as providing upper and lower bounds on $f_q(\mathcal{H})$ in terms of more accessible graph quantities.

The conflict graph presented here is slightly different than that considered in Kandiros et al. (2024). In that paper, the vertices of $\mathcal{H}$ were subjects in the experiment whereas in this paper the vertices of $\mathcal{H}$ are the estimand-relevant exposures. While the two are related, we prefer to use the definition in this paper as it contains exactly the information to characterize the minimax risk.

3 Lower Bounds on the Minimax Risk

In this section, we present statistical lower bounds on the minimax risk in terms of the conflict graph. In Section 3.1, we begin by demonstrating how Le Cam’s method can be adapted from classical minimax theory to the design-based setting. Next, we derive two types of statistical lower bounds using global and local information from the conflict graph. These are presented in Sections 3.2 and 3.3, respectively.

3.1 Le Cam’s Method for Design-Based Inference

Le Cam’s method is one of the most prominent approaches for obtaining statistical lower bounds on the minimax risk (Yu, 1997; Duchi, 2023; Wu, 2017). The method itself proceeds in two steps: the first step is a reduction from estimation to testing and the second step is the construction of a mixture which quantifies the limits of the testing problem. In this section, we show how the first step can be adapted to the design-based setting.

Le Cam’s method requires the construction of probability measures supported on potential outcome functions in the interference model. To facilitate this, let $\mathcal{A}$ be the σ -algebra generated by all finite subsets of the restricted model $\mathcal{M}(G, \tau, q)$, i.e. the countable-cocountable σ -algebra. A *distribution* over potential outcome functions is a probability measure on $(\mathcal{M}(G, \tau, q), \mathcal{A})$.

Let H_0, H_1 be two distributions over $\mathcal{M}(G, \tau, q)$, supported on submodels $\mathcal{M}_0, \mathcal{M}_1 \subseteq \mathcal{M}(G, \tau, q)$ respectively. In order to quantify the difference between the evaluation of the causal effect on the support of H_0 and H_1 , we define their *effect separation* as

$$d(H_0, H_1) = \inf_{y \in \mathcal{M}_0, y' \in \mathcal{M}_1} |\tau(y) - \tau(y')| .$$

For any distribution H over potential outcome functions and design $\mathcal{D}$, let $Q_H^{(\mathcal{D})}$ be the pushforward measure of the product measure $\mathcal{D} \otimes H$ under the mapping $(z, y) \mapsto (z, y(z))$. In other words, $Q_H^{(\mathcal{D})}$ is the distribution of the observed data when the potential outcome function is drawn according to H and the intervention is drawn according to $\mathcal{D}$.

Proposition 3.1. *For any network G , causal effect τ , and moment parameter q , the minimax risk is lower bounded as*

$$\mathcal{R}(G, \tau, q) \geq \inf_{\mathcal{D}} \sup_{H_0, H_1} \frac{d^2(H_0, H_1)}{8} \cdot \left(1 - d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})}) \right),$$

where d_{TV} is the total variation distance and the infimum is taken over all designs $\mathcal{D}$.

Proposition 3.1 provides a lower bound that increases with the effect separation between H_1 and H_0 , but decreases with their total variation distance. Applying the proposition requires constructing distributions H_0 and H_1 to trade-off these quantities, which is a challenge shared by all statistical lower bounds of this form.

The aspect of Proposition 3.1 that is unique to the design-based setting is the outer infimum over experimental designs $\mathcal{D}$, reflecting the fact that the choice of $\mathcal{D}$ is in the researcher’s control. This does not appear in the classical setting, where the data generating process is fixed. This means that when applying Proposition 3.1, we must consider a choice of distributions H_0 and H_1 for each possible experimental design $\mathcal{D}$. If we were to pick only choice of H_0 and H_1 , it would need to produce good lower bounds for all possible designs to be informative. Otherwise, the distributions H_0 and H_1 must be chosen with respect to the design, which we refer to as a *design-adaptive* construction. Design-adaptivity plays an important role in the development of our local-information lower bounds, whereas our global-information lower bounds are non-adaptive.

Very recently, Gao et al. (2026) established statistical lower bounds for misspecification testing within the design-based framework. Their approach may be understood as carrying out the second step of Le Cam’s method. In this work, we focus on minimax optimal rates of estimation and the mixtures we construct are considerably more complex.

3.2 Lower Bound using Global Information

Before presenting our lower bound, we must first define several graph-theoretic notions. For the moment, let us consider a generic graph which (in a slight abuse of notation) we denote as $G = (V, E)$. A subset of vertices $S \subseteq V$ is said to be *independent* if no two vertices in S are adjacent in G and the *independence number* of G is the size of its largest independent set, denoted $\mathcal{I}(G)$. Given a subset of vertices $T \subseteq V$, the *T -induced subgraph* is $G_T = (V_T, E_T)$ where $V_T = T$ and the edges E_T are those edges in the original graph G between vertices in T . We use $\mathcal{I}(G, T)$ to denote the independence number of the induced subgraph G_T .

The statistical lower bound in this section depends on the induced independence numbers within the conflict graph. These quantities represent global information about the conflict graph in the sense that they depend on its entire topology and cannot be determined from local information such as degree sequences. To gain intuition for these results, observe that each intervention $z \in \Omega$ reveals a set of exposures $S \subseteq V_{\mathcal{H}}$ which, by the definition of conflicts, form an independent set in $\mathcal{H}$. Thus, if all independent sets in $\mathcal{H}$ are small then each intervention can only provide information on a similarly small number of exposures. The following theorem formalizes this intuition into a precise statistical lower bound:

Theorem 3.2. *Let G be an interference network, τ be a contrastive effect. If $n \geq 3$, then for all moment parameters $q \geq 2$*

$$\mathcal{R}(G, \tau, q) \geq 1.05 \cdot 10^{-5} \cdot \sup_{T \subseteq V_{\mathcal{H}}, |T| \geq 5} \frac{1}{\mathcal{I}(\mathcal{H}, T)} \cdot \left(\frac{|T|}{n} \right)^{2 - \frac{2}{q}}.$$

For a fixed subset of exposures T in the conflict graph $\mathcal{H}$, the lower bound in Theorem 3.2 has two parts: the induced independence number $\mathcal{I}(\mathcal{H}, T)$ and the ratio of $(T/n)^{2-2/q}$. Roughly speaking, the former arises as an effective sample size while the latter reflects the amount of mass that can be “hidden” depending on the moment q that is being restricted. The intuition behind these quantities is that if the induced independence number is large, we can construct a mixture that hides mass among more potential outcomes. The overall lower bound is obtained by taking the supremum over sufficiently large subsets of exposures T . We have not made an attempt to optimize the constant appearing in Theorem 3.2.

In typical applications of Theorem 3.2, it suffices to consider T which are of order $\Omega(n)$. In general, it is always possible to set $T = V_{\mathcal{H}}$ in which case the lower bound becomes $1/\mathcal{I}(\mathcal{H})$. In the specific case of the direct effect, the conflict graph $\mathcal{H}$ contains G as an induced subgraph of size n , so that by setting $T = \{e_{i,1} : i \in [n]\}$, we obtain that $\mathcal{I}(\mathcal{H}, T) = \mathcal{I}(G)$. On the other hand, there are cases where the subset T that maximizes the lower bound contains a vanishingly small fraction of vertices in $\mathcal{H}$. For example, consider a conflict graph $\mathcal{H}$ that is the union of a clique R of size k_n and $2n - k_n$ isolated nodes, where $k_n = o(n)$. In this case, the optimal choice is to set $T = R$ in which case $\mathcal{I}(\mathcal{H}, T) = 1$ and the resulting lower bound is $\Omega(k_n/n)$ for $q = 2$ and $\Omega((k_n/n)^2)$ for $q = \infty$, both of which can be shown to be tight up to constant factors. In general, the lower bound increases as q decreases.

As stated above, the main idea for proving Theorem 3.2 comes from the relationship between independent sets in $\mathcal{H}$ and sets of exposures which can be simultaneously observed. In the remainder of this section, we provide a proof sketch together with the necessary ingredients. For a simplified proof sketch, we consider the case where $T = V_{\mathcal{H}}$ which yields the weaker $\Omega(1/|\mathcal{I}(\mathcal{H})|)$ bound. This simplification highlights the key features of the proof while avoiding some technical complications. For the purpose of notational simplicity (and without loss of generality), we consider the casual estimand $\tau = (1/n) \sum_{i=1}^n \{y_i(e_{i,1}) + y_i(e_{i,0})\}$ where the contrast is now a sum.

To apply Le Cam’s method, we will construct distributions H_1 and H_0 over potential outcome functions in the restricted ANI model. To this end, we consider a class of distributions $H(p)$ on the potential outcome functions which are defined by a single parameter $p \in (0, 1)$. Under $H(p)$, we select a subset S of the $2n$ exposures uniformly at random among all subsets of size $\lceil 2pn \rceil$ and assign outcome 1 to exposures in S and outcome 0 to exposures not in S . For a distribution H over outcome functions and intervention $z \in \Omega$, let $Q_H(\cdot|z)$ be the conditional distribution of the observed outcomes conditioned on the intervention being equal to z . Note that this distribution does not depend on the design $\mathcal{D}$ at all, since we are fixing the assignment z . For any $z \in \{0, 1\}^n$, let R_z be the subset of the $2n$ exposures that we get to observe when the treatment assignment is z .

We will consider mixtures of the form $H_0 = H(1/2)$ and $H_1 = H(1/2 + \delta)$ where the mixture parameter $\delta > 0$ controls the separation between the two distributions. We provide an auxiliary lemma about distinguishing H_0 and H_1 from the observed outcomes.

Lemma 3.3. *Suppose that $n \geq 3$ and fix an intervention $z \in \Omega$. If the mixture parameter satisfies $\delta \leq 0.026/\sqrt{|R_z|}$, then the total variation distance between the conditional observed distributions is bounded as $d_{\text{TV}}(Q_{H_0}^{(z)}, Q_{H_1}^{(z)}) \leq \frac{7}{8}$.*

Lemma 3.3 quantifies how large the mixture parameter δ can be so that the observed outcomes under the two mixtures H_0 and H_1 , conditioned on the intervention z , are hard to distinguish. Together with an application of Le Cam’s method as stated in Proposition 3.1, we are ready to provide a proof sketch of Theorem 3.2.

Proof Sketch of Theorem 3.2. Let $H_0 = H(1/2)$ and $H_1 = H(1/2 + \delta)$ where $\delta = 0.026/\sqrt{|\mathcal{I}(\mathcal{H})|}$. We will show that these two distributions satisfy the conditions of Proposition 3.1, which will then imply the desired lower bound. Note that the choice of H_0, H_1 is independent of the design $\mathcal{D}$.

First, by definition of H_0, H_1 we have $d(H_0, H_1) = 2\delta = 0.052/\sqrt{|\mathcal{I}(\mathcal{H})|}$. Second, by definition of the observed outcome distributions and the law of total probability, we have that for any design $\mathcal{D}$

$$d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})}) = \sum_{z \in \{0,1\}^n} \mathcal{D}(z) \cdot d_{\text{TV}}(Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z)) . \quad (1)$$

Therefore, if we can show that $Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z)$ are close for any value of z , then the same would be true for $Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})}$ as well, for any design $\mathcal{D}$. For a fixed z , the set R_z is an independent set in $\mathcal{H}$. Indeed, if two exposures in R_z were connected by an edge in $\mathcal{H}$, then by definition of the conflict graph we would not be able to observe both with the same assignment z . Thus, we have that $|R_z| \leq |\mathcal{I}(\mathcal{H})|$, which implies that

$$\delta = \frac{0.026}{\sqrt{|\mathcal{I}(\mathcal{H})|}} \leq \frac{0.026}{\sqrt{|R_z|}} .$$

Thus, by Lemma 3.3, we have that $d_{\text{TV}}(Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z)) \leq 7/8$ for all z , which implies that $d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})}) \leq 7/8$ for all designs $\mathcal{D}$ by (1). Therefore, using Proposition 3.1, we can lower bound the minimax risk as

$$\mathcal{R}(G, \tau, q, C) \geq \frac{d(H_0, H_1)^2}{8} \cdot \inf_{\mathcal{D}} \left(1 - d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})})\right) \geq \frac{0.052^2}{8 \cdot |\mathcal{I}(\mathcal{H})|} \cdot (1 - 7/8) = \Omega\left(\frac{1}{|\mathcal{I}(\mathcal{H})|}\right) . \quad \square$$

3.3 Lower Bound using Local Information

In this section, we provide a statistical lower bound on the minimax risk in terms of the degrees of the conflict graph $\mathcal{H}$. In contrast to global properties like induced independent sets, the degrees of the $\mathcal{H}$ describe local connectivity information. The lower bound presented in this section is complimentary to the one in the previous section in the sense that there are situations where one or the other is more informative.

For a vertex $e_{i,k} \in V_{\mathcal{H}}$ in the conflict graph $\mathcal{H}$, its *degree* $d_{i,k}$ is the number of vertices to which it is connected. For a subset of vertices T , we denote the minimum degree in T as $d_{\min}(\mathcal{H}, T) = \min_{e_{i,k} \in T} d_{i,k}$. Finally, we define the *generalized harmonic mean degree* of vertices in T as $d_{\text{harm}}(\mathcal{H}, T, r) = |T|/\{\sum_{e_{i,k} \in T} 1/d_{i,k}^r\}^{1/r}$. Note that for $r = 1$, the generalized harmonic mean reduces to the usual harmonic mean.

The minimax rate will increase as the conflict graph becomes more densely connected and the degrees of the vertices in $\mathcal{H}$ provide local information regarding its connectivity. Thus, it seems natural that the the minimax rate can be lower bounded in terms of the degrees. This can be formalized thru the following theorem, whose proof appears in Section B.3 of the appendix.

Theorem 3.4. *Let G be an interference network, τ be a contrastive effect. If $n \geq 5$ then for all moment parameters $q \geq 2$,*

$$\mathcal{R}(G, \tau, q) \geq \frac{6.6 \cdot 10^{-7}}{n^{2-\frac{2}{q}}} \sup_{T \subseteq V_{\mathcal{H}}, |T| \geq 5} \min\left(d_{\min}(\mathcal{H}, T), \sqrt{d_{\text{harm}}(\mathcal{H}, T, 2 - 2/q)}, |T|\right)^{2-\frac{2}{q}} .$$

For a fixed subset of vertices T in the conflict graph, the lower bound in Theorem 3.4 involves the minimum of three terms: the minimum degree in T , the generalized harmonic mean degree in T , and the size of T . In order for the lower bound to be large, there must exist a subset T for which all three of these quantities are large. In particular, the second term indicates that we must choose T such that its (generalized harmonic) average degree is large. At this point, the first and third terms are in tension: we need to select a large number of vertices T , but we do not want the minimum degree to be small. This means the supremum will typically not be achieved by the entire subset of vertices $T = V_{\mathcal{H}}$ because the resulting lower bound will be (at most) the minimum degree in $\mathcal{H}$, which could be quite small, i.e. constant order. There is less tension between these three terms when many exposures have approximately the same degree in the conflict graph.

The intuition that the “bulk” of the degrees in the conflict graph governs the minimax risk can be formalized using the critical degree, which is a quantile of the degree distribution. More precisely, we define the *degree distribution* of $\mathcal{H}$ as $F_{\mathcal{H}} : \mathbb{R} \rightarrow [0, 1]$ where

$$F_{\mathcal{H}}(t) = \frac{1}{2n} \sum_{e_{i,k} \in V_{\mathcal{H}}} \mathbf{1}[d_{i,k} \leq t] .$$

From the degree distribution, we can define for any $q \geq 2$ the *critical degree* $d_q^*(\mathcal{H})$ of $\mathcal{H}$ as the largest integer k such that there are at least $k^{\frac{2q-2}{3q-2}}$ nodes in $\mathcal{H}$ with degree at least k , i.e.

$$d_q^*(\mathcal{H}) = \sup \left\{ k \in [n] : 1 - F_{\mathcal{H}}(k) \geq \frac{k^{\frac{2q-2}{3q-2}}}{n} \right\} .$$

The supremum can be understood as follows: keep searching for vertices with degree at least k and we stop the first time there are fewer than $k^{\frac{2q-2}{3q-2}}$ of them. In this way, the critical degree describes a way to resolve the trade-off between finding a subset of vertices that is large in size but also has a large (generalized harmonic) mean of degrees, which is a central tension in the lower bound of Theorem 3.4. Note that the critical degree is decreasing as q increases.

The following corollary of Theorem 3.4 provides a lower bound based on the critical degree of $\mathcal{H}$, which is the bound mentioned in the Main Result of Section 1.

Corollary 3.5. *Let G be an interference network, τ be a contrastive effect. For all moment parameters $q \geq 2$, we have*

$$\mathcal{R}(G, \tau, q) \geq 6.6 \cdot 10^{-7} \cdot \frac{(d_q^*(\mathcal{H}))^{\frac{4(q-1)^2}{q(3q-2)}}}{n^{2-\frac{2}{q}}} .$$

The critical degree provides a resolution of the tension between the three terms of Theorem 3.4 by subsets T consisting of nodes with the largest degree. This simplifies the lower bound, but hides the fundamental tension. To understand the implications of Corollary 3.5, it is informative to consider the lower bound for d -regular conflict graphs, for which the critical degree is $d_q^*(\mathcal{H}) = d$. Under $q = 2$ moment restrictions, the minimax lower bound in Corollary 3.5 is of rate $d^{1/2}/n$. Under more restrictive $q = \infty$ conditions, the lower bound becomes $d^{4/3}/n^2$. As anticipated from monotonicity of the minimax risk as a function of q , we have that $d^{4/3}/n^2 \leq d^{1/2}/n$ because $d \leq n$.

If the average degree in the conflict graph is large, then a standard argument using Markov’s inequality implies that there should be a sufficiently large subset of vertices with sufficiently large

degree. Indeed, we can show that the critical degree is always lower bounded by the average degree in the conflict graph (see Proposition B.6 in Section B.3 of the appendix).

We now present a proof sketch for Theorem 3.4. For simplicity, we focus on second moment restrictions ($q = 2$) and make the simplifying assumption that all degrees in the conflict graph $\mathcal{H}$ are lower bounded by d , in which case the lower bound will be $\sqrt{d}/n$, as realized by $T = [n]$. This case captures the essential obstacles for establishing the general claim, while avoiding some technical details of the argument. A notable difference of this proof compared to the global lower bound of Section 3.2 is that the mixture distributions H_0 and H_1 will now be selected *adaptively* depending on the experimental design $\mathcal{D}$ under consideration.

Proof Sketch of Theorem 3.4. Assume $d_i \geq d$ for all $i \in V_{\mathcal{H}}$. For any exposure $e = e_{i,k} \in V(\mathcal{H})$, associated with a unit i in the graph G , let $N(e)$ be the set of neighboring exposures in $\mathcal{H}$ and let A_e be the event that we observe the outcome of i for exposure e , i.e.

$$A_e = \left\{ z \in \{0, 1\}^n : h_i(z) = e \right\} .$$

Let R_z denote the set of exposures whose outcomes we observe under z . By definition, we have

$$|R_z| = \sum_{e \in V(\mathcal{H})} \mathbf{1}[A_e](z) .$$

Consider an arbitrary design $\mathcal{D}$ and let $p \in (0, 1/2)$ be a parameter that will depend on $\mathcal{D}$, but will be chosen later. Our construction will be design-adaptive, depending on the size of exposure probabilities. If at least one exposure has large probability, we can exploit that exposure to hide a large causal effect in its conflicting exposures which are unlikely to be observed. If all exposure probabilities are small, then it is unlikely that we will observe many potential outcomes at all, in which case we can hide the causal effect among these outcomes. The parameter p represents what is to be considered small and large exposure probabilities and is chosen at the end to optimize the resulting lower bound. Let us now formally distinguish between these two cases.

- **Case 1:** Suppose there exists an exposure $e \in V(\mathcal{H})$ such that $\mathcal{D}(E_e) \geq p$. In that case, we define H_0 and H_1 to be supported on potential outcome functions $y^{(0)}, y^{(1)}$ respectively, where

$$y_i^{(1)}(e_{i,k}) = \begin{cases} \sqrt{\frac{n}{|N(e)|}} & \text{if } e_{i,k} \in N(e) \\ 0 & \text{otherwise} \end{cases} , \quad y_i^{(0)}(e_{i,k}) = 0 \quad \text{for all } e_{i,k} \in V(\mathcal{H}) .$$

By construction, we have that

$$d(H_0, H_1) = \frac{1}{n} \sum_{e' \in N(e)} \left(\sqrt{\frac{n}{|N(e)|}} - 0 \right) = \sqrt{\frac{|N(e)|}{n}} \geq \sqrt{\frac{d}{n}} .$$

Also, if event E_e occurs, then we get to observe the outcome $y_i(e)$. By definition of the conflict graph, this implies that we do not observe any of the outcomes of exposures $e' \in N(e)$. Since H_0, H_1 only differ in the outcomes of $N(e)$, they are impossible to distinguish under that event. Thus,

$$d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})}) \leq 1 - \mathcal{D}(E_e) \leq 1 - p$$

Therefore, using Proposition 3.1 for this particular design $\mathcal{D}$, we have

$$\inf_{\hat{\tau}} \text{Risk}(\mathcal{D}, \hat{\tau}) \geq \frac{d(H_0, H_1)^2}{8} \cdot (1 - d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})})) \geq \frac{dp}{8n} . \quad (2)$$

- **Case 2:** Suppose for all exposures $e \in V(\mathcal{H})$ we have that $\mathcal{D}(E_e) < p$. In that case, define the event A that at most $4pn$ exposures are observed in total, i.e.

$$A = \{z \in \{0, 1\}^n : |R_z| \leq 4pn\} ,$$

By Markov's inequality,

$$\mathcal{D}(A^c) = \mathcal{D}\left(|R_z| \geq 4pn\right) \leq \frac{\mathbb{E}[|R_z|]}{4pn} = \frac{\sum_{e \in V(\mathcal{H})} \mathcal{D}(E_e)}{4pn} \leq \frac{p \cdot |V(\mathcal{H})|}{4pn} = \frac{1}{2} .$$

Define $H_0 = H(1/2)$ and $H_1 = H(1/2 + \delta)$ where $\delta = 0.026/\sqrt{4pn}$. Then, if $z \in A$, by definition $|R_z| \leq 4pn$. By Lemma 3.3, for any $z \in A$, we have

$$d_{\text{TV}}(Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z)) \leq \frac{7}{8} .$$

The law of total probability then implies that

$$1 - d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})}) \geq \mathcal{D}(A) \cdot (1 - c) \geq \frac{1}{2} \left(1 - \frac{7}{8}\right) = \frac{1}{16} .$$

Therefore, using Proposition 3.1 for this particular design $\mathcal{D}$, we get

$$\inf_{\hat{\tau}} \text{Risk}(\mathcal{D}, \hat{\tau}) \geq \frac{d(H_0, H_1)^2}{8} \cdot (1 - d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})})) \geq \frac{4}{8 \cdot 2pn} \cdot \frac{1}{16} = \Omega\left(\frac{1}{pn}\right) \quad (3)$$

Since any design $\mathcal{D}$ falls in exactly one of the two cases above, the minimax risk is lower bounded by the minimum of (2) and (3). By selecting $p = \Theta(1/\sqrt{d})$ to optimize this minimum, we get $\mathcal{R}(G, \tau, 2) = \Omega(\sqrt{d}/n)$, as desired. $\square$

4 Upper Bounds on the Minimax Risk

In this section, we present general upper bounds on the minimax risk of network experiment problems. One can view the analysis of previously proposed statistical procedures as implicitly deriving upper bounds on the minimax risk. However, most analyses in the literature are appropriate for providing sufficient conditions for consistency and may not accurately describe the minimax risk itself. Additionally, most works consider only a single moment restriction, usually $q = 2$ or $q = \infty$. In contrast, the upper bounds we present in this section are improved (i.e. smaller) and hold for arbitrary moment restrictions.

4.1 Unbiased Statistical Procedures

In order to upper bound the minimax risk, we must construct a statistical procedure and analyze its maximum risk. We begin by considering a class of unbiased statistical procedures, which we then refine in the next section thru a bias-variance trade-off. We summarize these upper bounds below in Theorem 4.1 and defer a discussion of the underlying statistical procedure until the end of this section.

Theorem 4.1. *Let G be an interference network, τ be a contrastive effect. For any moment parameter $q \geq 2$, the minimax risk is bounded as*

$$\mathcal{R}(G, \tau, q) \leq 8.92 \cdot \frac{\lambda(\mathcal{H})^{\frac{2}{q}} \cdot d_{\text{avg}}(\mathcal{H})^{1-\frac{2}{q}}}{n},$$

where $\lambda(\mathcal{H})$ is the largest eigenvalue of the adjacency matrix of $\mathcal{H}$ and $d_{\text{avg}}(\mathcal{H})$ is the average degree of $\mathcal{H}$. Moreover, this bound on the minimax risk is achieved by an unbiased statistical procedure.

Theorem 4.1 provides upper bounds on the minimax risk which depend on the moment restriction under consideration. Under second moment restrictions ($q = 2$), we recover the $\lambda(\mathcal{H})/n$ bound of Kandiros et al. (2024). For uniformly bounded outcomes ($q = \infty$), we obtain an improved rate of $d_{\text{avg}}(\mathcal{H})/n$. For intermediate moments $q \in (2, \infty)$, the resulting upper bound interpolates between these two quantities. Like the true minimax risk, our upper bounds are decreasing as the network experiment problem becomes easier, i.e. q increases so the moment condition becomes more restrictive.

The difference between $d_{\text{avg}}(\mathcal{H})$ and $\lambda(\mathcal{H})$ becomes more pronounced under degree heterogeneity. For example, the largest adjacency eigenvalue always lies between the average and maximum degrees: $d_{\text{avg}}(\mathcal{H}) \leq \lambda(\mathcal{H}) \leq d_{\text{max}}(\mathcal{H})$. When there is low degree heterogeneity, meaning that all vertices have roughly the same degree, then all three quantities are roughly equal. On the other hand, the quantities can be of different asymptotic orders in the presence of degree heterogeneity. An illustrative example which demonstrates this separation is when $\mathcal{H}$ is a star graph, where one “center” node is connected to all remaining “leaf” nodes and no leaf nodes are connected. Then we would have $d_{\text{avg}}(\mathcal{H}) \approx 2$, $\lambda(\mathcal{H}) \approx \sqrt{n}$, and $d_{\text{max}}(\mathcal{H}) \approx n$. Degree heterogeneity is commonly observed in social networks where degree distributions are heavy-tailed, i.e. a small number of people are more popular than the average person. In this case, the degree heterogeneity of the underlying network G will induce degree heterogeneity in the conflict graph $\mathcal{H}$ and, as a result, the quantities $\lambda(\mathcal{H})$ and $d_{\text{avg}}(\mathcal{H})$ will have different asymptotic orders.

The interplay between moment restrictions and degree heterogeneity in the minimax rate can be understood by considering how much of the total “mass” of the potential outcomes can be concentrated on a few number of subjects. The key idea can be explained as follows: under less restrictive moment conditions, the “mass” of the potential outcomes is allowed to be more concentrated on a few number of subjects, which in turn means that the maximal risk will depend on the smallest exposure probabilities among all subjects. If the conflict graph has heterogeneous degrees, then increasing the probability of observing an exposure with many conflicts is only possible by decreasing other exposure probabilities. On the other hand, under more restrictive moment conditions, the “mass” of potential outcomes cannot be hidden in just a few subjects. This means that it can be advantageous to select an experimental design under which a few exposures are observed with small probability if doing so allows most exposures to be observed with larger probability.

Underlying Theorem 4.1 is a new family of experimental designs which are extensions of the Conflict Graph Design of Kandiros et al. (2024) to accommodate general moment restrictions. At a high level, the Conflict Graph Design works by sampling independent auxiliary “desired inclusion” variables and then resolving conflicts according to a particular importance ordering. The distribution of the auxiliary variables and the importance ordering depend on the underlying moment restriction q . For the setting of $q = 2$ considered by Kandiros et al. (2024), auxiliary variables have identical distributions and the importance ordering is given by the largest eigenvector of the adjacency matrix

$\mathcal{H}$. For the case of $q = \infty$, the distribution of each auxiliary variable depends on the exposure's degree and the importance ordering is simply the degree ordering. Detailed description and analyses of this new family of experimental designs may be found in Section C.1 of the supplement.

4.2 Bias-Variance Tradeoff

In this section, we consider bias-variance trade-offs which can further improve the upper bounds on the minimax rate. The key idea is simple: if a few “central” or “high degree” vertices in $\mathcal{H}$ cause $\lambda(\mathcal{H})$ or $d_{\text{avg}}(\mathcal{H})$ to be large, then we can choose to ignore their outcomes when constructing our experimental design and estimator so that the resulting variance is improved at the cost of some additional bias. Similar ideas have been explored in the context of AIPW estimation in network experiments (Zöcklein, 2025). The challenge will be to identify the optimal set S of vertices to ignore.

More formally, let $S \subseteq V_{\mathcal{H}}$ be a subset of the vertices in the conflict graph. Let $\mathcal{H}(V_{\mathcal{H}} \setminus S)$ denote the subgraph of $\mathcal{H}$ induced by the vertices in $V_{\mathcal{H}} \setminus S$. Let τ_S be the contrastive treatment effect restricted to the exposures in $V_{\mathcal{H}} \setminus S$, i.e.

$$\tau_S = \frac{1}{n} \sum_{i=1}^n \left(y_i(e_{i,1}) \mathbf{1}[e_{i,1} \notin S] - y_i(e_{i,0}) \mathbf{1}[e_{i,0} \notin S] \right).$$

If we were to target the estimand τ_S by running the Conflict Graph Design from Theorem 4.1 on the subgraph $\mathcal{H}(V_{\mathcal{H}} \setminus S)$, then our estimator $\hat{\tau}$ is unbiased for the restricted estimand τ_S , i.e. $\mathbb{E}[\hat{\tau}] = \tau_S$. Thus, its variance and bias for the actual estimand τ can be bounded as

$$\text{Var}(\hat{\tau}) = \mathcal{O}\left(\frac{\lambda(\mathcal{H}(V_{\mathcal{H}} \setminus S))^{\frac{2}{q}} \cdot d_{\text{avg}}(\mathcal{H}(V_{\mathcal{H}} \setminus S))^{1-\frac{2}{q}}}{n}\right) \quad \text{and} \quad (\mathbb{E}[\hat{\tau}] - \tau)^2 = \mathcal{O}\left(\left(\frac{|S|}{n}\right)^{2-\frac{2}{q}}\right).$$

Based on these bounds, we can define the following proxy for the mean squared error of the estimator $\hat{\tau}$, as it depends on the subset S that is to be removed:

$$Q_q(S, \mathcal{H}) := \lambda(\mathcal{H}(V_{\mathcal{H}} \setminus S))^{\frac{2}{q}} \cdot d_{\text{avg}}(\mathcal{H}(V_{\mathcal{H}} \setminus S))^{1-\frac{2}{q}} + \frac{|S|^{2-\frac{2}{q}}}{n^{1-\frac{2}{q}}}.$$

The best possible MSE using this approach would then be given by the subset S that minimizes $Q_q(S, \mathcal{H})$. This is the content of the main result of this section (for the proof, see Appendix C.2.1).

Theorem 4.2. *Let G be an interference network, τ be a contrastive effect. For any moment parameter $q \geq 2$, the corresponding minimax risk is bounded as*

$$\mathcal{R}(G, \tau, q) \leq 8.92 \cdot \frac{\min_{S \subseteq V_{\mathcal{H}}} Q_q(S, \mathcal{H})}{n}.$$

The upper bound on the minimax rate in Theorem 4.2 is an improvement upon the one presented in Theorem 4.1. To see this, observe that selecting $S = \emptyset$ results in the same rate, and indeed the same statistical procedure. Whether the optimal set $S^* = \arg \min_{S \subseteq V_{\mathcal{H}}} Q_q(S, \mathcal{H})$ offers a substantially different rate depends on the heterogeneity of the conflict graph. If the vertices of the conflict graph all have degree roughly d , then the variance term is $\mathcal{O}(d)$ because in this case $\lambda(\mathcal{H}) \approx d_{\text{avg}}(\mathcal{H}) \approx d$. In this case, we would need to remove $S = \Omega(n)$ vertices before the variance

term would become $o(d)$, at which point the squared bias term would already be $\Omega(1)$. Thus, the bias-variance trade-off does not improve the minimax rate when $\mathcal{H}$ has homogeneous degrees.

On the other hand, the bias-variance trade-off offers a large improvement when the conflict graph has significant degree heterogeneity. The reason is that removing a small number of vertices will substantially decrease the variance term whereas the small number of them ensures the increase in bias is minimal. Consider the illustrative example where $\mathcal{H}$ is the star graph and let us focus on $q = 2$. In this case, the unbiased procedure of Theorem 4.1 yields a rate of $1/\sqrt{n}$ whereas by removing the center node $S = \{1\}$, we have that $\lambda(\mathcal{H}(V_{\mathcal{H}} \setminus S)) = 1$, so that the resulting rate from Theorem 4.2 is $1/n$. Moreover, the lower bounds establish that $1/n$ is indeed the exact minimax-rate for the star graph.

Finding the optimal subset S^* that minimizes $Q_q(S, \mathcal{H})$ is a combinatorial optimization problem. The problem is similar to the Partial Vertex Cover (Hochbaum, 1997) which is NP-Hard, so a computationally feasible algorithm to find S^* is unlikely to exist. However, building upon approximation algorithms for Partial Vertex Cover, we show in the following proposition that there is a polynomial-time algorithm which computes a set S that is within a logarithmic factor of optimal, i.e. $Q_q(S, \mathcal{H}) \leq \log(n) \cdot Q_q(S^*, \mathcal{H})$. This yields a computationally efficient approach to the bias-variance trade-off that incurs only a logarithmic increase in the risk. This is the content of the proposition below and we refer to Section C.3 in the supplement for more details.

Proposition 4.3. *Let G be an interference network, τ be a contrastive effect. For any moment parameter $q \geq 2$, there exists a design-estimator pair $(\mathcal{D}, \hat{\tau})$ that can be implemented in polynomial time whose maximum risk is bounded as*

$$\sup_{y \in \mathcal{M}(G, \tau, q)} \text{MSE}(\mathcal{D}, \hat{\tau}, \tau, y) \leq \mathcal{O}\left((\log n)^{2-\frac{2}{q}}\right) \cdot \frac{\min_{S \subseteq V(\mathcal{H})} Q_q(S, \mathcal{H})}{n}.$$

5 Tightness of Minimax Analysis

The minimax risk introduced in Section 2.2 is a function of the underlying network G and the contrastive estimand τ . While we present upper and lower bounds that use different information from the conflict graph $\mathcal{H}$, the bounds themselves do not always match in the usual sense, i.e. they are not equal up to a constant. In this section, we show that the upper and lower bounds are both tight within the class of regular graphs. Moreover, these bounds are nearly matching for a uniformly random regular graph, characterizing the minimax rate in an average-case sense.

To demonstrate the tightness of our upper and lower bounds, we consider the class of d -regular graphs, where each vertex has exactly d neighbors. We denote the class of d -regular graphs on n nodes by $\mathcal{G}_{d,n}$, though we simply write $\mathcal{G}_d$ when n is clear from context. Using triangular array asymptotics, we consider the degree parameter d to be possibly growing with n .

Restricting our attention to the class of regular graphs is useful because it is parametrized by a single parameter d . This has the benefit of simplifying the presentation of our upper and lower bounds on the minimax rate. On the other hand, the class is rich enough to contain a variety of different graphs. We further focus our attention on the bounds obtained for second moment restrictions, $q = 2$. Our goal is two fold: we aim to show (1) how the upper and lower bounds depend on d and that (2) they provide tight bounds on the minimax risk within this class $\mathcal{G}_{d,n}$.

5.1 Direct Treatment Effect

We first focus on estimating the Direct Effect τ_{DTE} , which was defined in Section 2.1. As discussed there, the conflict graph $\mathcal{H}$ for this estimand behaves roughly like the original interference graph G . In particular, as we have seen in Corollary 1.1, the following upper and lower bounds on the minimax risk hold for any d -regular graph G :

$$\frac{\sqrt{d+1}}{n} \lesssim \mathcal{R}(G, \tau_{\text{DTE}}, 2) \lesssim \frac{d+1}{n} .$$

Note that we have removed the independent set $\mathcal{I}(G)$ term from the lower bound because it is not governed by d , though it will play a role in our later analysis. When the degree d is constant relative to the number of units n , then the bounds match up to a constant factor. A particular case of interest is the Stable Unit Treatment Value Assumption (SUTVA), where there is no interference and the graph G is empty. In that case, we obtain that the minimax risk of estimating the Average Treatment Effect (ATE) is $\Theta(1/n)$, which is the classical parametric rate of estimation.

When d is growing with n , there is a multiplicative gap of $\sqrt{d}$ between the upper and lower bounds. We will demonstrate that there are families of graphs in $\mathcal{G}_d$ for which the minimax rate coincides with the upper bound as well as families for which the minimax rate coincides with the lower bound. The next proposition establishes tightness of the upper bound.

Proposition 5.1. *For any sequence $d = \{d^{(n)}\}_{n=1}^\infty$ with $d = o(n)$,*

$$\sup_{G \in \mathcal{G}_d} \mathcal{R}(G, \tau_{\text{DTE}}, 2) \asymp \frac{d}{n} .$$

For the proof, we consider the family of d -chorded cycles, where the maximum independent set is $\mathcal{O}(d/n)$. We then use the global lower bound of Theorem 3.2 to establish the matching lower bound of $\Omega(d/n)$. For further details, see Section E.2 of the supplementary.

Next, we show that the lower bound of $\sqrt{d}/n$ is also attainable up to constant factors for a family of graphs in $\mathcal{G}_d$ with $d = \Theta(n^{2/3})$.

Proposition 5.2. *There exists a sequence $d = \{d^{(n)}\}_{n=1}^\infty$ such that $d = \Theta(n^{2/3})$ and*

$$\inf_{G \in \mathcal{G}_d} \mathcal{R}(G, \tau_{\text{DTE}}, 2) \asymp \frac{\sqrt{d}}{n} .$$

For the proof, we focus on a family of graphs called the *Generalized Quadrangles*, which are well studied in finite geometry (Tits, 1959; Payne and Thas, 2009). The crucial property of these graphs is the existence of large independent sets that have very well-behaved intersections with each other. Combined with other strong connectivity properties, this allows us to construct an improved design tailored for this graph family that achieves the $\sqrt{d}/n$ estimation rate. More details can be found in Section E.3 of the supplementary. Together, these results imply that both the upper and lower bounds from 5.1 in terms of the degree parameter are tight.

Finally, we investigate the behavior of the minimax risk for the “average” graph in $\mathcal{G}_d$, by considering a graph sampled uniformly at random from $\mathcal{G}_d$.

Corollary 5.3. *Suppose G is a uniformly random d -regular graph with $d = o(n)$. Then, with probability at least $1 - o(1)$ over the random choice of G ,*

$$\frac{d}{n \cdot \log d} \lesssim \mathcal{R}(G, \tau_{\text{DTE}}, 2) \lesssim \frac{d}{n} .$$

Corollary 5.3 shows that for the overwhelming majority of d -regular graphs, the minimax risk of estimating the DTE is equal to the upper bound of Corollary 1.1 up to a $\log d$ factor. To prove that, we utilize the fact that if $d = o(n)$, then with high probability the largest independent set in G has size at most $\mathcal{O}(n \log d/d)$ (see e.g. Cooper et al., 2002), together with our global lower bound from Theorem 3.2. See Section E.4 of the supplementary for more details.

5.2 Global Average Treatment Effect

When the estimand of interest is the GATE τ_{GATE} , the conflict graph $\mathcal{H}$ is similar in structure to the two hop graph G^2 of the underlying network G , as demonstrated in Section 2.3. In particular, Corollary 1.2 provides the following upper and lower bounds on the minimax risk that hold for any graph G :

$$\frac{\sqrt{d_{\text{avg}}(G^2)}}{n} \lesssim \mathcal{R}(G, \tau_{\text{GATE}}, 2) \lesssim \frac{\lambda(G^2) + 1}{n} .$$

Even when the graph G is d -regular, the quantities $\sqrt{d_{\text{avg}}(G^2)}$ and $\lambda(G^2)$ are unfortunately not determined by the degree parameter d . The reason is that the size of a two-hop neighborhood can range from d to d^2 . However, the following theorem shows that when the network G is random d -regular, $\sqrt{d_{\text{avg}}(G^2)}$ and $\lambda(G^2)$ are well-approximated by d and d^2 so that there is a multiplicative gap of order d between them:

Corollary 5.4. *Suppose G is a random d -regular graph with $d = o(\sqrt{n})$. Then, with probability at least $1 - o(1)$ over the random choice of G ,*

$$\frac{d}{n} \lesssim \mathcal{R}(G, \tau_{\text{GATE}}, 2) \lesssim \frac{d^2}{n} .$$

Corollary 5.4 together with Corollary 5.3 provides evidence for the widely held belief that the DTE is easier to estimate than the GATE. To see this, observe that when G is a random d -regular graph, then these two corollaries together yield that $\mathcal{R}(G, \tau_{\text{DTE}}) \lesssim d/n$ whereas $\mathcal{R}(G, \tau_{\text{GATE}}) \gtrsim d/n$. To the best of our knowledge, this provides the first theoretical evidence of the difference between the rates at which different network effects can be estimated.

6 Conjectured Computational Hardness

In the previous section, we have shown that while the derived upper and lower bounds on the minimax risk do not match in all cases, they are both tight for the class of regular graphs and are in fact matching for random regular graphs. However, this still leaves open the question of whether a simple characterization of the minimax risk is possible. In this section, we present two conjectures against the existence of an efficient algorithm which can approximately compute the minimax risk and discuss their consequences. For the purposes of conciseness, we gloss over certain technical aspects of computational complexity theory.

At the heart of our conjectures is an intimate connection between the experimental design problem and computational problems regarding independent sets in graphs. As noted in Section 2.3, the fact that $\mathcal{R}(G, \tau, q)$ can be expressed as $f_q(\mathcal{H})$ for some function f_q relies on the translation of the experimental design problem into a problem regarding distributions over maximal independent

sets of $\mathcal{H}$. The central question considered in this section is whether this function $f_q : \mathbb{H} \rightarrow \mathbb{R}_{\geq 0}$ is efficiently computable in any meaningful sense.

The computer science literature has established that computational problems involving the independent sets of a graph are provably hard. For example, it is NP-Hard to calculate the size of the largest independent set $\mathcal{I}(G)$ of a graph G (Karp, 1972). In fact, even approximating $\mathcal{I}(G)$ to within a multiplicative $n^{1-\epsilon}$ factor is NP-Hard for all $\epsilon > 0$ (Hastad, 1996; Zuckerman, 2006). This $n^{1-\epsilon}$ inapproximability result presumes that the input graph is a priori unrestricted so that the size of the maximal independent set could range from 1 to n . By focusing specifically on d -regular graphs where the size of the largest independent set is known to lie between n/d and n , the inapproximability can be strengthened further still: approximating $\mathcal{I}(G)$ to within a $\mathcal{O}(d/\log^2 d)$ multiplicative factor for a d -regular graph G is hard assuming the Unique Games Conjecture (Austrin, Khot, and Safra, 2011). Based on these results and the connection between the minimax rate of network experiment problems and the structure of independent sets in the conflict graph, we present our two conjectures:

Conjecture 1. *Given a network G , contrastive effect τ , and moment parameter q , the minimax risk $\mathcal{R}(\tau, G, q) = f_q(\mathcal{H})$ is NP-Hard to approximate within a multiplicative factor n^c for some $c \in (0, 1)$.*

Conjecture 2. *Given a d -regular graph G and the direct effect τ_{DTE} , the minimax risk $\mathcal{R}(G, \tau_{\text{DTE}}, q) = f_q(\mathcal{H})$ is NP-Hard to approximate within a multiplicative factor of $d^{1/2-\epsilon}$ for all $\epsilon > 0$.*

Conjecture 1 posits that even an approximation of the minimax rate is computationally challenging for general network experiment problems. Conjecture 2 is a strengthening of Conjecture 1 specifically for the direct effect, which states that even determining whether the true minimax rate is close to our derived upper bound d/n or lower bound $\sqrt{d}/n$ is computationally infeasible. As discussed above, our primary evidence for these conjectures is the close relationship between the minimax rate $\mathcal{R}(G, \tau, q)$ and the structure of maximal independent sets in the conflict $\mathcal{H}$ as well as the general inapproximability results for maximum independent set. Note that the computational hardness of approximating the minimax rate would not preclude an explicit construction of a minimax optimal statistical procedure; however, implementing such a procedure would likely be computationally inefficient, otherwise one could use Monte Carlo simulations of this procedure to approximate the minimax rate.

These conjectures posit that the minimax rate does not admit a simple description. If true, these conjectures would have significant consequences for the practical relevance of minimax theory in network experiments. Perhaps the most valuable aspect of minimax theory is to provide a theoretical benchmark that certain statistical procedures should pass, e.g. being minimax rate-optimal. If the minimax rate of a network experiment is not efficiently computable, then the rate optimality of a particular statistical procedure cannot be assessed on an arbitrary network and effect. In this case, we see two possible alternatives. The first is to show rate optimality of statistical procedures within a certain family of graphs in a worst-case or average-case sense. For example, Corollary 5.3 on the rate optimality of Conflict Graph Design for random d -regular graphs may be seen as an example of this type of analysis. The second alternative is to formulate a notion of computationally efficient minimax risk, which restricts the statistical procedure to be implementable in polynomial time. From this perspective, understanding both the optimal computationally efficient statistical procedure and the rate of its corresponding risk would be natural next steps.

7 Conclusion

In this paper, we have studied design-based minimax optimal rates for network experiments. Our upper and lower bounds provide further understanding of how the minimax rate depends on the conflict graph $\mathcal{H}$. Generally speaking, our results show that the minimax rates are governed by its “bulk connectivity”, in the sense that a few well-connected vertices will not affect the rates. Moreover, the analysis suggests an interplay between the strength of moment restrictions and the extent to which small but well-connected subgraphs can affect the minimax rates.

We believe that the most interesting open directions posed by this work are further investigation into the computational complexity of network experiments. Is approximately computing the minimax risk $\mathcal{R}(G, \tau, q)$ computationally hard? In Section 6, we have provided several well-motivated conjectures that suggest the answer is yes. Progress towards resolving these conjectures in any direction would provide deeper understanding into the fundamental obstructions in constructing optimal statistical procedures. If statistical procedures which attain the minimax risk require exponential computational resources, then what is the best rates achievable by a computationally efficient experimental design and estimator? Questions like these lie in the intersection of statistics, computational complexity, and graph theory. Computational considerations notwithstanding, an exact characterization of the minimax rates in terms of the conflict graph (i.e. a description of $f_q(\mathcal{H})$) would be valuable in further understanding the essence of these statistical limits. Doing this would ostensibly require both improved statistical lower bounds and experimental designs, which would be interesting in their own right regardless of computational tractability.

References

- Aggarwal, Om P (1959). “Bayes and minimax procedures in sampling from finite and infinite populations—I”. *The Annals of Mathematical Statistics*, pp. 206–218.
- Aronow, P.M. and Patrick Lopatto (2026). “Minimax unbiased estimation for finite populations with bounded outcomes”. *arXiv preprint arXiv:2605.20572*.
- Aronow, P.M. and Cyrus Samii (2017). “Estimating average causal effects under general interference, with application to a social network experiment”. *The Annals of Applied Statistics* 11.4, pp. 1912–1947.
- Austrin, Per, Subhash Khot, and Muli Safra (2011). “Inapproximability of vertex cover and independent set in bounded degree graphs”. *Theory of Computing* 7.1, pp. 27–43.
- Baird, Sarah, J. Aislinn Bohren, Craig McIntosh, and Berk Özler (Dec. 2018). “Optimal Design of Experiments in the Presence of Interference”. *The Review of Economics and Statistics* 100.5, pp. 844–860.
- Basse, Guillaume W and Edoardo M Airoidi (Dec. 2018). “Model-assisted design of experiments in the presence of network-correlated outcomes”. *Biometrika* 105.4, pp. 849–858.
- Bickel, Peter J, Chris AJ Klaassen, Peter J Bickel, Ya’acov Ritov, J Klaassen, Jon A Wellner, and YA’Acov Ritov (1993). *Efficient and adaptive estimation for semiparametric models*. Vol. 4. Springer.
- Bickel, Peter J and Erich L Lehmann (1981). “A Minimax Property of the Sample Mean in Finite Populations”. *The Annals of Statistics* 9.5, pp. 1119–1122.
- Blackwell, David A and Meyer A Girshick (1979). *Theory of games and statistical decisions*. Courier Corporation.

- Bojinov, Iavor, David Simchi-Levi, and Jinglong Zhao (2023). “Design and Analysis of Switchback Experiments”. *Management Science* 69.7, pp. 3759–3777.
- Cai, Chencheng, Xu Zhang, and Edoardo Airoldi (2024). “Independent-Set Design of Experiments for Estimating Treatment and Spillover Effects under Network Interference”. *The Twelfth International Conference on Learning Representations*.
- Candogan, Ozan, Chen Chen, and Rad Niazadeh (2024). “Correlated Cluster-Based Randomized Experiments: Robust Variance Minimization”. *Management Science* 70.6, pp. 4069–4086.
- Clark, Pete L (2013). “Quadratic forms chapter I: Witt’s theory”. *University of Georgia* 3.
- Cooper, Colin, Alan Frieze, Bruce Reed, and Oliver Riordan (2002). “Random regular graphs of non-constant degree: Independence and chromatic number”. *Combinatorics, Probability and Computing* 11.4, pp. 323–341.
- Cortez-Rodriguez, Mayleen, Matthew Eichhorn, and Christina Lee Yu (2023). “Exploiting neighborhood interference with low-order interactions under unit randomized design”. *Journal of Causal Inference* 11.1, p. 20220051.
- Duchi, John (2023). “Lecture notes on statistics and information theory”. URL: <https://web.stanford.edu/class/stats311/lecture-notes.pdf>.
- Fatemi, Zahra and Elena Zheleva (2020). “Network Experiment Design for Estimating Direct Treatment Effects”. KDD Workshop on Mining and Learning with Graphs.
- Foster, Dylan J and Vasilis Syrgkanis (2023). “Orthogonal statistical learning”. *The Annals of Statistics* 51.3, pp. 879–908.
- Frieze, Alan M and TOMASZ Łuczak (1992). “On the independence and chromatic numbers of random regular graphs”. *Journal of Combinatorial Theory, Series B* 54.1, pp. 123–132.
- Gabler, Siegfried (2012). *Minimax solutions in sampling from finite populations*. Vol. 64. Springer Science & Business Media.
- Gandhi, Rajiv, Samir Khuller, and Aravind Srinivasan (2004). “Approximation algorithms for partial covering problems”. *Journal of Algorithms* 53.1, pp. 55–84.
- Gao, Chao, Christopher Harshaw, Fredrik Sävje, and Yitan Wang (2026). “On the Impossibility of Specification Testing of Interference Models Based on Exposure Mappings”. *arXiv preprint arXiv:2605.09726*.
- Godsil, Chris and Gordon F Royle (2013). *Algebraic graph theory*. Springer Science & Business Media.
- Harshaw, Christopher, Fredrik Sävje, and Yitan Wang (2022). “A design-based Riesz representation framework for randomized experiments”. *arXiv preprint arXiv:2210.08698*.
- Hastad, Johan (1996). “Clique is hard to approximate within n^{ϵ} ”. *Proceedings of 37th Conference on Foundations of Computer Science*. IEEE, pp. 627–636.
- Hochbaum, Dorit S (1997). “Approximation algorithms for NP-hard problems”. *ACM Sigact News* 28.2, pp. 40–52.
- Hodges, JL and EL Lehmann (1982). “Minimax estimation in simple random sampling”. *Statistics and probability: essays in honor of CR Rao*. North-Holland Publishing Company, pp. 325–327.
- Holland, Paul W (1986). “Statistics and causal inference”. *Journal of the American statistical Association* 81.396, pp. 945–960.
- Hu, Yuchen and Stefan Wager (2023). “Off-policy evaluation in partially observed Markov decision processes under sequential ignorability”. *The Annals of Statistics* 51.4, pp. 1561–1585.
- Hudgens, Michael G and M Elizabeth Halloran (2008). “Toward causal inference with interference”. *Journal of the american statistical association* 103.482, pp. 832–842.

- Jagadeesan, Ravi, Natesh S. Pillai, and Alexander Volfovsky (2020). “Designs for estimating the treatment effect in networks with interference”. *Annals of Statistics* 48.2, pp. 679–712.
- Kallus, Nathan (2018). “Optimal A Priori Balance in the Design of Controlled Experiments”. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 80.1, pp. 85–112.
- Kandiros, Vardis, Charilaos Pipis, Constantinos Daskalakis, and Christopher Harshaw (2024). “The Conflict Graph Design: Estimating causal effects under arbitrary neighborhood interference”. *arXiv preprint arXiv:2411.10908*.
- Karp, Richard M (1972). “Reducibility among combinatorial problems”. *Complexity of computer computations*. Springer, pp. 85–103.
- Karwa, Vishesh and Edoardo M. Airoidi (2018). “A Systematic Investigation of Classical Causal Inference Strategies under Mis-specification Due to Network Interference”. *arXiv:1810.08259v1*. *arXiv: 1810.08259v1 [stat.ME]*.
- Karwa, Vishesh and Edoardo M Airoidi (2023). “On the admissibility of horvitz-thompson estimator for estimating causal effects under network interference”. *arXiv preprint arXiv:2312.01234*.
- Kennedy, Edward H, Sivaraman Balakrishnan, James M Robins, and Larry Wasserman (2024). “Minimax rates for heterogeneous causal effect estimation”. *Annals of statistics* 52.2, p. 793.
- Krivelevich, Michael, Benny Sudakov, Van H Vu, and Nicholas C Wormald (2001). “Random regular graphs of high degree”. *Random Structures & Algorithms* 18.4, pp. 346–363.
- Lam, Tsit-Yuen (2005). *Introduction to quadratic forms over fields*. Vol. 67. American Mathematical Soc.
- Lehmann, Erich Leo and George Casella (1998). *Theory of point estimation*. Springer.
- Leung, Michael P. (2022). “Rate-optimal cluster-randomized designs for spatial interference”. *The Annals of Statistics* 50.5, pp. 3064–3087.
- Manski, Charles F (2013). “Identification of treatment response with social interactions”. *The Econometrics Journal* 16.1, S1–S23.
- McKay, Brendan D, Nicholas C Wormald, and Beata Wysocka (2004). “Short cycles in random regular graphs”. *the electronic journal of combinatorics*, R66–R66.
- Payne, Stanley E and Joseph Adolf Thas (2009). *Finite generalized quadrangles*. Vol. 9. European Mathematical Society.
- Rinott, Yosef (2009). “Chapter 41 - Some Decision-Theoretic Aspects of Finite Population Sampling”. *Handbook of Statistics*. Ed. by C.R. Rao. Vol. 29. Handbook of Statistics. Elsevier, pp. 523–558.
- Robins, James, Lingling Li, Eric Tchetgen, Aad van der Vaart, et al. (2008). “Higher order influence functions and minimax estimation of nonlinear functionals”. *Probability and statistics: essays in honor of David A. Freedman*. Vol. 2. Institute of Mathematical Statistics, pp. 335–422.
- Rotnitzky, Andrea, Ezequiel Smucler, and James M Robins (2021). “Characterization of parameters with a mixed bias property”. *Biometrika* 108.1, pp. 231–238.
- Saha, Sudip, Abhijin Adiga, B Aditya Prakash, and Anil Kumar S Vullikanti (2015). “Approximation algorithms for reducing the spectral radius to control epidemic spread”. *Proceedings of the 2015 SIAM International Conference on Data Mining*. SIAM, pp. 568–576.
- Sobel, Michael E (2006). “What Do Randomized Studies of Housing Mobility Demonstrate?” *Journal of the American Statistical Association* 101.476, pp. 1398–1407.
- Sussman, Daniel L. and Edoardo M. Airoidi (2017). “Elements of estimation theory for causal effects in the presence of network interference”. *arXiv:1702.03578*. *arXiv: 1702.03578v1 [stat.ME]*.
- Tits, Jacques (1959). “Sur la transitivité et certains groupes qui s’ en déduisent”. *Publications Mathématiques de l’IHÉS* 2, pp. 13–60.

- Toulis, Panos and Edward Kao (2013). “Estimation of Causal Peer Influence Effects”. *Proceedings of the 30th International Conference on Machine Learning*. Ed. by Sanjoy Dasgupta and David McAllester. Vol. 28. Proceedings of Machine Learning Research 3. Atlanta, Georgia, USA: PMLR, pp. 1489–1497.
- Ugander, Johan, Brian Karrer, Lars Backstrom, and Jon Kleinberg (2013). “Graph cluster randomization: Network exposure to multiple universes”. *KDD '13*. ACM, pp. 329–337.
- Ugander, Johan and Hao Yin (2023). “Randomized graph cluster randomization”. *Journal of Causal Inference* 11.1, p. 20220014.
- van der Vaart, Aad W (2000). *Asymptotic statistics*. Vol. 3. Cambridge university press.
- Viviano, Davide, Lihua Lei, Guido Imbens, Brian Karrer, Okke Schrijvers, and Liang Shi (2026). *Causal clustering: design of cluster experiments under network interference*. arXiv: 2310.14983 [econ.EM].
- Wilf, Herbert S (1967). “The eigenvalues of a graph and its chromatic number”. *Journal of the London mathematical Society* 1.1, pp. 330–332.
- Wu, Chien-Fu (1981). “On the Robustness and Efficiency of Some Randomized Designs”. *The Annals of Statistics* 9.6, pp. 1168–1177.
- Wu, Yihong (2017). “Lecture notes on information-theoretic methods for high-dimensional statistics”. *Lecture Notes for ECE598YW (UIUC)* 16, p. 15.
- Yu, Bin (1997). “Assouad, fano, and le cam”. *Festschrift for Lucien Le Cam: research papers in probability and statistics*. Springer, pp. 423–435.
- Zöcklein, Gernot (2025). “Regression Adjustment for Randomized Network Experiments”. Master’s thesis. ETH Zurich.
- Zuckerman, David (2006). “Linear degree extractors and the inapproximability of max clique and chromatic number”. *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pp. 681–690.

Appendix

Table of Contents

A	Minimax Theory for Design-Based Inference	29
A.1	A General Definition of Minimax Risk	29
A.2	Monotonicity of the Minimax Risk (Proposition 2.1)	29
A.3	Randomized Estimators	30
A.4	Variations of Outcome Restrictions	32
A.5	Conflict Graph Characterizes Minimax Risk (Proposition 2.2)	35
B	Lower Bounds on the Minimax Risk	41
B.1	Le Cam’s Method (Proposition 3.1)	41
B.2	Lower Bound using Global Information (Theorem 3.2)	43
B.3	Lower Bound using Local Information (Theorem 3.4)	52
C	Upper Bounds on the Minimax Risk	63
C.1	Analysis of the Conflict Graph Design (Theorem 4.1)	63
C.2	Bias-Variance Tradeoff	77
C.3	Polynomial Time Approximation Algorithm (Proposition 4.3)	80
D	Illustrations of Main Results	83
D.1	Direct Treatment Effect (Corollary 1.1)	83
D.2	Global Average Treatment Effect (Corollary 1.2)	84
E	Tightness of Bounds	85
E.1	The Structure of the Conflict Graph	85
E.2	Tightness of Upper Bound for Direct Effect (Proposition 5.1)	86
E.3	Tightness of Lower Bound for Direct Effect (Proposition 5.2)	87
E.4	Random Regular Graphs (Corollaries 5.3 and 5.4)	94

A Minimax Theory for Design-Based Inference

In this section, we present the proofs of our general results involving the minimax risk for estimating arbitrary causal effects. The subsections are organized as follows:

- In Section A.1, we present a general definition of minimax risk for arbitrary restrictions on the potential outcome functions.
- In Section A.2, we establish various monotonicity properties of the minimax risk.
- In Section A.3, we show that additional randomness does not improve estimation.
- In Section A.4, we prove that the minimax risk is invariant to some natural restrictions of the potential outcome functions.
- Finally, in Section A.5, we prove that the minimax risk is characterized by the conflict graph.

A.1 A General Definition of Minimax Risk

In this section, we briefly describe how to extend the definition of minimax optimality to more general design-based causal inference. We focus on the design-based framework proposed by setting introduced by Harshaw, Sävje, and Wang (2022).

In the framework of Harshaw, Sävje, and Wang (2022), the researcher will posit, for each subject $i \in [n]$, a *subject-level model space* $\mathcal{M}_i$ which is a linear subspace of measurable functions. The model spaces serve to encode the structure of the potential outcome functions. The *entire model* is a subset $\mathcal{M} \subseteq \mathcal{M}_1 \times \cdots \times \mathcal{M}_n$, which may incorporate additional moment restrictions. The *causal estimand* is a linear functional $\tau : \mathcal{M} \rightarrow \mathbb{R}$.

Generalizing the definition in Section 2.2, we define the *minimax risk* of estimating a causal estimand τ over a model $\mathcal{M}$ on the potential outcomes as follows:

$$\mathcal{R}(\mathcal{M}, \tau) = \inf_{\mathcal{D}, \hat{\tau}} \sup_{y \in \mathcal{M}} \mathbb{E}_{Z \sim \mathcal{D}} [(\hat{\tau}(Z, y(Z)) - \tau(y))^2] .$$

When τ is a contrastive effect and $\mathcal{M}(G)$ is an ANI model, then the definition of minimax risk above reduces to the one presented in Section 2.2.

A.2 Monotonicity of the Minimax Risk (Proposition 2.1)

We first establish a monotonicity result which holds for arbitrary restrictions on the potential outcome functions and for arbitrary estimands. We then use it to prove each of the monotonicity properties in Proposition 2.1.

Lemma A.1. *If $\mathcal{M} \subseteq \mathcal{M}'$ then $\mathcal{R}(\mathcal{M}, \tau) \leq \mathcal{R}(\mathcal{M}', \tau)$.*

Proof. Let us fix any estimator $\hat{\tau}$ and a design P . Then, since $\mathcal{M}$ is contained in $\mathcal{M}'$, we have that

$$\text{Risk}(P, \hat{\tau}) = \sup_{y \in \mathcal{M}} \mathbb{E}_{Z \sim P} [(\hat{\tau}(Z, Y(Z)) - \tau(y))^2] \leq \sup_{y \in \mathcal{M}'} \mathbb{E}_{Z \sim P} [(\hat{\tau}(Z, Y(Z)) - \tau(y))^2] = \text{Risk}(P, \hat{\tau}) .$$

Taking the infimum over all estimators $\hat{\tau}$ and designs P on both sides gives the desired result. $\square$

Now let us prove Proposition 2.1 using the above general monotonicity result. For convenience, we restate it below.

Proposition 2.1. *The minimax risk $\mathcal{R}(G, \tau, q)$ obeys the following monotonicity properties:*

- **Monotonicity in q :** *If $q' \leq q$, then $\mathcal{R}(G, \tau, q) \leq \mathcal{R}(G, \tau, q')$.*
- **Monotonicity in G :** *If $G' \subseteq G$, then $\mathcal{R}(G', \tau, q) \leq \mathcal{R}(G, \tau, q)$.*

Proof. The proof of each of the monotonicity properties relies on applying Lemma A.1 to the appropriate sets of potential outcome functions. We now show how to do that for each statement separately.

1. **Monotonicity in q .** This will follow immediately from Lemma A.1 if we could show that $\mathcal{M}(G, q, C) \subseteq \mathcal{M}(G, q')$ for any $q' \leq q$. Consider any tuple of potential outcome functions $y \in \mathcal{M}(G, q)$. For any $k \in \{0, 1\}$, we have that

$$\frac{1}{n} \sum_{i=1}^n |y_i(\mathbf{z}_{i,(k)})|^{q'} = \frac{1}{n} \sum_{i=1}^n (|y_i(\mathbf{z}_{i,(k)})|^q)^{\frac{q'}{q}} \leq \left(\frac{1}{n} \sum_{i=1}^n |y_i(\mathbf{z}_{i,(k)})|^q \right)^{\frac{q'}{q}} \leq 1 ,$$

where the first inequality follows from concavity of $x \mapsto x^{q'/q}$ and Jensen's inequality and the final inequality follows from the definition of $y \in \mathcal{M}(G, q)$. Therefore, $y \in \mathcal{M}(G, q')$, which implies that $\mathcal{M}(G, q) \subseteq \mathcal{M}(G, q')$.

2. **Monotonicity in G .** Finally, for the monotonicity in G , it suffices to establish that if $G' \subseteq G$, then $\mathcal{M}(G', q) \subseteq \mathcal{M}(G, q)$, since then Lemma A.1 applies. To prove this, it suffices to show that $\mathcal{M}(G') \subseteq \mathcal{M}(G)$, i.e. the potential outcome functions that satisfy ANI with respect to G' also satisfy ANI with respect to G . But this also follows from the definition of graph inclusion. To distinguish between the two different graphs, we use $\tilde{N}_H(i)$ to denote the extended neighborhood of a node i in graph H . In particular, suppose $y \in \mathcal{M}(G')$. Then, let $i \in [n]$ and $z, z' \in \{0, 1\}^n$ be two treatment assignments such that for all i the exposure of i under z and z' is the same with respect to G , i.e.

$$\{j \in \tilde{N}_G(i) : z_j = 1\} = \{j \in \tilde{N}_G(i) : z'_j = 1\} .$$

Since $G' \subseteq G$, we have that $\tilde{N}_{G'}(i) \subseteq \tilde{N}_G(i)$, which implies that

$$\{j \in \tilde{N}_{G'}(i) : z_j = 1\} = \{j \in \tilde{N}_{G'}(i) : z'_j = 1\} .$$

Since $y \in \mathcal{M}(G')$, the last equality implies that $y_i(z) = y_i(z')$. Since this holds for all such i and all z, z' , we have that $y \in \mathcal{M}(G)$, which completes the proof. □

A.3 Randomized Estimators

In Section 2.1, we defined our setup so that the randomness only came from the treatment assignment. However, since the experimenter controls the data collection process, they could also utilize additional randomness in choosing the treatment and in constructing the estimator for the causal estimand of

interest. In this section, we show that allowing a statistical procedure to use additional randomness does not improve the minimax risk.

In particular, let S be the sample space of the additional randomness and let $\mathcal{S}$ be a σ -algebra on S . Since the arguments below work for any choice of S and $\mathcal{S}$, we suppress the dependence on them in the following definitions.

We define a *design with randomness* as a probability measure $\mathcal{D}_{\text{rand}} : \mathcal{F} \otimes \mathcal{S} \rightarrow [0, 1]$ over the product space of treatment assignments and the additional randomness. Drawing a sample from $\mathcal{D}_{\text{rand}}$ produces a pair (Z, X) , where $Z \in \Omega$ is the treatment assignment that is used by the experimenter and $X \in S$ is the additional randomness, which could be used by the estimator. We emphasize again that $\mathcal{D}_{\text{rand}}$ is the only source of randomness in this problem and it is controlled by the experimenter. Also, the experimenter could correlate the additional randomness with the treatment assignment Z , which implies that $\mathcal{D}_{\text{rand}}$ is not necessarily a product measure over Ω and S . A *randomized estimator* $\hat{\tau}_{\text{rand}} : \Omega \times S \times \mathbb{R}^n \rightarrow \mathbb{R}$ is a measurable function with respect to the product σ -algebra $\mathcal{F} \otimes \mathcal{S} \otimes \mathcal{B}(\mathbb{R}^n)$, where $\mathcal{B}(\mathbb{R}^n)$ is the Borel σ -algebra on $\mathbb{R}^n$. The pair $(\mathcal{D}_{\text{rand}}, \hat{\tau}_{\text{rand}})$ is referred to as a *randomized statistical procedure*. In this setting, the experimenter first draws a random pair (Z, X) according to the design with randomness $\mathcal{D}_{\text{rand}}$, where $Z \in \Omega$ and $X \in S$. After applying the treatment Z , they observe outcomes $Y(Z)$ and output the estimate $\hat{\tau}_{\text{rand}}(Z, X, Y(Z))$ of the effect τ .

In analogy with the definitions in Section A.1, we can define the *Randomized Minimax Risk* for a given causal effect τ and model $\mathcal{M}$ as the minimum risk over all designs with randomness and all randomized estimators, i.e.

$$\mathcal{R}_{\text{rand}}(\mathcal{M}, \tau) = \inf_{\mathcal{D}_{\text{rand}}, \hat{\tau}_{\text{rand}}} \sup_{y \in \mathcal{M}} \mathbb{E}_{(Z, X) \sim \mathcal{D}_{\text{rand}}} [(\hat{\tau}_{\text{rand}}(Z, X, Y(Z)) - \tau(y))^2] .$$

We now show that the minimax risk of randomized statistical procedures is the same as that of deterministic procedures. In particular, we establish that for any pair of design with randomness and randomized estimator, there exists a deterministic estimator that uses only the treatment part of the randomness and has the same or smaller MSE for all potential outcomes functions.

Proposition A.2. *For any model $\mathcal{M}$ and any effect τ , $\mathcal{R}_{\text{rand}}(\mathcal{M}, \tau) = \mathcal{R}(\mathcal{M}, \tau)$.*

Proof. Since deterministic estimators are a special case of randomized estimators, we have by definition that

$$\mathcal{R}_{\text{rand}}(\mathcal{M}, \tau) \leq \mathcal{R}(\mathcal{M}, \tau) .$$

To establish the reverse inequality, let's fix any randomized experimental design $\mathcal{D}_{\text{rand}}$ and any randomized estimator $\hat{\tau}_{\text{rand}}$. Since $\mathcal{D}_{\text{rand}}$ is a probability measure over the product space $\Omega \times S$, let's denote the marginal distribution of $\mathcal{D}_{\text{rand}}$ over Ω by $\mathcal{D}$. Now, define the estimator $\hat{\tau} : \Omega \times \mathbb{R}^n \rightarrow \mathbb{R}$ as

$$\hat{\tau}(Z, y(Z)) = \mathbb{E}_{X \sim \mathcal{D}_{\text{rand}}(\cdot | Z)} [\hat{\tau}_{\text{rand}}(Z, X, y(Z))] .$$

Clearly, $\hat{\tau}$ is a deterministic estimator that only depends on the treatment assignment Z and the observed outcomes $y(Z)$.

Let $y \in \mathcal{M}$ be a potential outcome function in model $\mathcal{M}$. Then, we have that

$$\mathbb{E}_{(Z, X) \sim \mathcal{D}_{\text{rand}}} \left[\left(\hat{\tau}_{\text{rand}}(Z, X, y(Z)) - \tau(y) \right)^2 \right] = \mathbb{E}_{Z \sim \mathcal{D}} \left[\mathbb{E}_{X \sim \mathcal{D}_{\text{rand}}(\cdot | Z)} \left[\left(\hat{\tau}_{\text{rand}}(Z, X, y(Z)) - \tau(y) \right)^2 \right] \right]$$

$$\geq \mathbb{E}_{Z \sim \mathcal{D}} \left[\left(\mathbb{E}_{X \sim \mathcal{D}_{\text{rand}}(\cdot|Z)} [\hat{\tau}_{\text{rand}}(Z, X, y(Z))] - \tau(y) \right)^2 \right]$$

by using Jensen's inequality for the convex function $x \mapsto x^2$

$$= \mathbb{E}_{Z \sim \mathcal{D}} \left[\left(\hat{\tau}(Z, y(Z)) - \tau(y) \right)^2 \right] .$$

Therefore, the statistical procedure $(\mathcal{D}, \hat{\tau})$ has smaller or equal MSE than the procedure $(\mathcal{D}_{\text{rand}}, \hat{\tau}_{\text{rand}})$ that operates with extra randomness, for all potential outcome functions $y \in \mathcal{M}$. Therefore, taking supremum over all $y \in \mathcal{M}$ and then infimum over all pairs $(\mathcal{D}_{\text{rand}}, \hat{\tau}_{\text{rand}})$, we have

$$\mathcal{R}_{\text{rand}}(\mathcal{M}, \tau) \geq \mathcal{R}(\mathcal{M}, \tau) ,$$

which completes the proof. $\square$

A.4 Variations of Outcome Restrictions

In this section, we show that the minimax risk remains unaffected under some natural variations of the outcome restrictions introduced in Section 2.1.

We start by showing that if we change the scaling of the constraint on the outcomes, the minimax risk scales accordingly. In particular, for an interference network G , effect τ , moment parameter $q \geq 2$ and positive real number $C > 0$, we define the following set of potential outcome functions:

$$\mathcal{M}(G, \tau, q, C) = \left\{ y \in \mathcal{M}(G) : \left(\frac{1}{n} \sum_{i=1}^n |y_i(e_{i,k})|^q \right)^{1/q} \leq C \text{ for } k \in \{0, 1\} \right\} .$$

This generalizes the q th moment restricted model $\mathcal{M}(G, \tau, q)$ defined in Section 2.1, which is defined with $C = 1$. Likewise, we can define the minimax risk of estimating the effect τ over this set of potential outcome functions as

$$\mathcal{R}(G, \tau, q, C) = \inf_{\mathcal{D}, \hat{\tau}} \sup_{y \in \mathcal{M}(G, \tau, q, C)} \mathbb{E}_{Z \sim \mathcal{D}} [(\hat{\tau}(Z, y(Z)) - \tau(y))^2] .$$

We now show that changing the bound on the outcomes results in a simple scaling of the minimax risk. Therefore, it is sufficient to consider the case $C = 1$ without loss of generality.

Proposition A.3. *Let G be an interference network, τ be a contrastive effect and $q \geq 2$ be a moment parameter. Then, for any positive real number $C > 0$, we have that*

$$\mathcal{R}(G, \tau, q, C) = C^2 \cdot \mathcal{R}(G, \tau, q, 1) .$$

Proof. For any $L > 0$ and potential outcome function y , we denote by Ly the potential outcome function defined as $(Ly)_i(z) = L \cdot y_i(z)$ for all $i \in [n]$ and $z \in \{0, 1\}^n$. Let $U, V > 0$ be any positive real numbers. Let $\mathcal{D}$ be any design and $\hat{\tau}$ be any estimator. We define the following estimator $\hat{\tau}_{U,V}$ as

$$\hat{\tau}_{U,V}(Z, y(Z)) = \frac{U}{V} \cdot \hat{\tau}(Z, (V/U) \cdot y(Z)) .$$

In particular, $\hat{\tau}_{U,V}$ scales the observed outcomes by V/U , applies the estimator $\hat{\tau}$, and then scales the output by U/V .

Now consider an arbitrary potential outcome function $y \in \mathcal{M}(G, \tau, q, U)$. Then, we have that $(V/U) \cdot y \in \mathcal{M}(G, \tau, q, V)$, since for any $k \in \{0, 1\}$, we have that

$$\left(\frac{1}{n} \sum_{i=1}^n |((V/U) \cdot y)_i(e_{i,k})|^q \right)^{1/q} = \frac{V}{U} \left(\frac{1}{n} \sum_{i=1}^n |y_i(e_{i,k})|^q \right)^{1/q} \leq \frac{V}{U} \cdot U = V .$$

Also, by linearity we have

$$\tau((V/U) \cdot y) = (V/U) \cdot \tau(y) .$$

Therefore, for any $y \in \mathcal{M}(G, \tau, q, U)$, we have that

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} \left[(\hat{\tau}_{U,V}(Z, y(Z)) - \tau(y))^2 \right] &= \mathbb{E}_{\mathcal{D}} \left[\left(\frac{U}{V} \hat{\tau}(Z, (V/U) \cdot y(Z)) - \frac{U}{V} \tau((V/U) \cdot y) \right)^2 \right] \\ &= \frac{U^2}{V^2} \cdot \mathbb{E}_{\mathcal{D}} \left[\left(\hat{\tau}(Z, (V/U) \cdot y(Z)) - \tau((V/U) \cdot y) \right)^2 \right] \\ &\leq \frac{U^2}{V^2} \cdot \sup_{y \in \mathcal{M}(G, \tau, q, V)} \mathbb{E}_{\mathcal{D}} \left[\left(\hat{\tau}(Z, y(Z)) - \tau(y) \right)^2 \right] , \end{aligned}$$

where in the last step we used the fact that $(V/U)y \in \mathcal{M}(G, \tau, q, V)$. Therefore, for any statistical procedure $(\mathcal{D}, \hat{\tau})$, we can find another statistical procedure $(\mathcal{D}, \hat{\tau}_{U,V})$ whose maximal error under $\mathcal{M}(G, \tau, q, V)$ is upper bounded by the scaled maximal error of $(\mathcal{D}, \hat{\tau})$ under $\mathcal{M}(G, \tau, q, U)$. It follows that

$$\mathcal{R}(G, \tau, q, U) \leq \frac{U^2}{V^2} \cdot \mathcal{R}(G, \tau, q, V)$$

We can now apply inequality A.4 twice, one with $U = C, V = 1$ and one with $U = 1, V = C$ to obtain the result. $\square$

We now show that the outcomes of units corresponding to exposures that do not appear in the estimand of interest do not affect the minimax risk. In particular, we show that even if we restrict these outcomes to have value 0, the minimax risk remains the same.

Specifically, for any graph G and estimand τ with relevant exposures $e_{i,1}, e_{i,0}$ for $i \in [n]$, let $\mathcal{M}^0(G, \tau)$ be the model of outcomes that satisfy ANI with respect to G and have zero outcome in the non-relevant exposures, i.e.

$$\mathcal{M}^0(G, \tau) = \left\{ y \in \mathcal{M}(G) : y_i(e) = 0, \text{ if } e \neq e_{i,0}, e_{i,1}, \forall i \in [n] \right\}$$

Then, for any graph G , estimand τ and $q \geq 2$, we can define the model of potential outcome functions with bounded q moments and zero non-relevant outcomes as follows

$$\mathcal{M}^0(G, \tau, q) = \mathcal{M}^0(G, \tau) \cap \mathcal{M}(G, \tau, q) .$$

Analogously to the previous definitions, we can define the minimax risk with respect to outcome functions in $\mathcal{M}^0(G, \tau, q)$ as follows

$$\mathcal{R}^0(G, \tau, q) = \inf_{\mathcal{D}, \hat{\tau}} \sup_{y \in \mathcal{M}^0(G, \tau, q)} \mathbb{E}_{Z \sim \mathcal{D}} [(\hat{\tau}(Z, y(Z)) - \tau(y))^2] .$$

In the next proposition, we show that the minimax risk is unaffected by the restrictions in the non-relevant exposures.

Proposition A.4. *For any graph G , estimand τ and $q \geq 2$,*

$$\mathcal{R}^0(G, \tau, q) = \mathcal{R}(G, \tau, q) .$$

Proof. Since $\mathcal{M}^0(G, \tau, q) \subseteq \mathcal{M}(G, \tau, q)$, we trivially have

$$\mathcal{R}^0(G, \tau, q) \leq \mathcal{R}(G, \tau, q) .$$

We thus focus on showing the reverse inequality. For any outcome function $y \in \mathcal{M}(G)$, we denote by y^0 the corresponding outcome function that agrees with y in the relevant exposures and has 0 outcome everywhere else, i.e.

$$y_i^0(z) = \begin{cases} y_i(z) & \text{if } h_i(z) \in \{e_{i,0}, e_{i,1}\} \\ 0 & \text{otherwise} \end{cases} .$$

Clearly, if $y \in \mathcal{M}(G, \tau, q)$, then $y^0 \in \mathcal{M}(G, \tau, q)$, since the moment constraint is not affected by outcomes of exposures that do not appear in τ . Also, we always have $\tau(y) = \tau(y^0)$, since the estimand doesn't depend on the outcomes we are zeroing out.

Now, let $(\mathcal{D}, \hat{\tau}_0)$ be an arbitrary statistical procedure. We will construct an estimator $\hat{\tau}$, such that the worst case performance of $(\mathcal{D}, \hat{\tau})$ under $\mathcal{M}(G, \tau, q)$ is no worse than that of $(\mathcal{D}, \hat{\tau}_0)$ under $\mathcal{M}^0(G, \tau, q)$. Indeed, we define $\hat{\tau}$ in terms of $\hat{\tau}_0$ as follows

$$\hat{\tau}(z, y(z)) = \hat{\tau}_0(z, y^0(z))$$

In words, $\hat{\tau}$ works by zeroing out the non-relevant outcomes and then applying $\hat{\tau}_0$ to the result. Then, for any $y \in \mathcal{M}(G, \tau, q)$

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} \left[(\hat{\tau}(Z, y(Z)) - \tau(y))^2 \right] &= \mathbb{E}_{\mathcal{D}} \left[(\hat{\tau}_0(Z, y^0(Z)) - \tau(y))^2 \right] \\ &= \mathbb{E}_{\mathcal{D}} \left[(\hat{\tau}_0(Z, y^0(Z)) - \tau(y^0))^2 \right] \end{aligned}$$

Taking the supremum over $y \in \mathcal{M}(G, \tau, q)$ on both sides yields

$$\begin{aligned} \sup_{y \in \mathcal{M}(G, \tau, q)} \mathbb{E}_{\mathcal{D}} \left[(\hat{\tau}(Z, y(Z)) - \tau(y))^2 \right] &= \sup_{y \in \mathcal{M}(G, \tau, q)} \mathbb{E}_{\mathcal{D}} \left[(\hat{\tau}_0(Z, y^0(Z)) - \tau(y^0))^2 \right] \\ &= \sup_{y \in \mathcal{M}^0(G, \tau, q)} \mathbb{E}_{\mathcal{D}} \left[(\hat{\tau}_0(Z, y(Z)) - \tau(y))^2 \right] . \end{aligned}$$

The last step follows by the following observation, which follows by the definitions.

$$\mathcal{M}^0(G, \tau, q) = \{y^0 : y \in \mathcal{M}(G, \tau, q)\} .$$

Since the above holds for an arbitrary statistical procedure $(\mathcal{D}, \hat{\tau})$, we conclude that

$$\mathcal{R}(G, \tau, q) \leq \mathcal{R}^0(G, \tau, q) .$$

The result follows. □

A.5 Conflict Graph Characterizes Minimax Risk (Proposition 2.2)

In this section, we show that while the conflict graph abstracts away many features of the experimental design problem, it contains the essential information to characterize the minimax risk. More precisely, we will show that the minimax risk is a function of the conflict graph. In the main body, this is formalized as Proposition 2.2, which we reproduce below for completeness. Recall that $\mathbb{H}$ is the set of all labeled graphs.

Proposition 2.2. *Let G be an interference network, τ be a contrastive effect. For each moment parameter $q \geq 2$, there exists a function $f_q : \mathbb{H} \rightarrow \mathbb{R}_{\geq 0}$ such that $\mathcal{R}(G, \tau, q) = f_q(\mathcal{H})$.*

Our proof strategy will proceed by defining two minimax problems on an arbitrary labeled graph $\mathcal{H}$. In these minimax problems, each of the vertices in $\mathcal{H}$ has an associated outcome and the target estimand is the average of the outcomes. A random independent set in $\mathcal{H}$ is selected and the outcomes of its vertices is revealed. The minimax problem is to construct both the distribution over independent sets and the estimator. We first consider distributions over all independent sets, and then the restriction to distributions supported on maximal independent sets. Through a series of reductions, we will show that the minimax risk of these graph theoretic problems exactly corresponds to the minimax risk of the original network experiment problem.

We begin by formally defining these minimax problems. Consider a labeled graph $\mathcal{H} = (V_{\mathcal{H}}, E_{\mathcal{H}})$ with $2n$ vertices. Note that we have assumed the number of vertices is even, for simplicity. Consider a function $u : V_{\mathcal{H}} \rightarrow \mathbb{R}$ which maps each vertex to a real number and may be identified with a vector $u \in \mathbb{R}^{|V_{\mathcal{H}}|}$. For a moment parameter $q \geq 2$, define the q th moment restricted model

$$\mathcal{M}(\mathcal{H}, q) = \{u \in \mathbb{R}^{V_{\mathcal{H}}} : \left(\frac{1}{n} \sum_{i=1}^n |u(i)|^q\right)^{1/q} \leq 1 \text{ and } \left(\frac{1}{n} \sum_{i=n+1}^{2n} |u(i)|^q\right)^{1/q} \leq 1\} ,$$

which restricts the q th moment for the first half and second half of the outcomes. Note that the labelling of the graph $\mathcal{H}$ is required to give meaning to “first half” and “second half”. The labelling will also be required to define the contrast between two groups of vertices. Formally, we define the *contrastive functional* $\tau : \mathcal{M}(\mathcal{H}, q) \rightarrow \mathbb{R}$ as

$$\tau(u) = \frac{1}{n} \sum_{i=n+1}^{2n} u(i) - \frac{1}{n} \sum_{i=1}^n u(i) .$$

In these minimax problems, we must select a probability measure over independent sets in $\mathcal{H}$, which we formalize below. Recall that a subset of vertices $S \subseteq V_{\mathcal{H}}$ is *independent* if no pair of vertices $i, j \in S$ is adjacent in $\mathcal{H}$. We denote the independent sets of $\mathcal{H}$ as $\mathcal{I}(\mathcal{H})$. We have that the independent sets form a measurable space $(\mathcal{I}(\mathcal{H}), \mathcal{P}(\mathcal{I}(\mathcal{H})))$, where $\mathcal{P}(\mathcal{I}(\mathcal{H}))$ is the power set. Let $\mathbb{D}(\mathcal{I}(\mathcal{H}))$ be the set of all probability measures on independent sets $\mathcal{I}(\mathcal{H})$.

What is observed in these minimax problems is the outcomes for each of the vertices in the (randomly) selected independent set and this is the data which must be used to estimate the average outcome. In order to simplify the formal notation, we introduce an auxiliary variable $*$ which will denote that no outcome has been observed. Given an independent set $S \in \mathcal{I}(\mathcal{H})$, define the vector $u_S \in \mathbb{R}^{V_{\mathcal{H}}}$ where $u_S(i) = u(i)$ if $i \in S$ and $u_S(i) = *$ if $i \notin S$. Thus, the experimenter observes $(S, u_S) \in \mathcal{I}(\mathcal{H}) \times (\mathbb{R}_*)^{V_{\mathcal{H}}}$, where $\mathbb{R}_* = \mathbb{R} \cup \{*\}$. An *estimator* is a measurable function $\hat{\tau} : \mathcal{I}(\mathcal{H}) \times (\mathbb{R}_*)^{V_{\mathcal{H}}} \rightarrow \mathbb{R}$. We denote the set of all real-valued measurable functions on $\mathcal{I}(\mathcal{H}) \times (\mathbb{R}_*)^{V_{\mathcal{H}}}$ as $\Gamma(\mathcal{I}(\mathcal{H}))$.

With this notation in place, we can now define the first minimax problem for labelled graphs. For a graph $\mathcal{H}$ and a moment parameter $q \geq 2$, define the *independent set minimax risk* $\mathcal{R}(\mathcal{H}, q)$ as

$$\mathcal{R}(\mathcal{H}, q) = \inf_{\substack{\mathcal{D} \in \mathbb{D}(\mathcal{I}(\mathcal{H})) \\ \hat{\tau} \in \Gamma(\mathcal{I}(\mathcal{H}))}} \sup_{u \in \mathcal{M}(\mathcal{H}, q)} \mathbb{E}_{S \sim \mathcal{D}} \left[(\tau(u) - \hat{\tau}(S, u_S))^2 \right] .$$

The independent minimax risk describes the best achievable mean squared error in this graph estimation problem, defined solely in terms of the labelled graph $\mathcal{H}$.

The second minimax problem is similar, but considers the restriction to observe only maximal independent sets in $\mathcal{H}$. This restriction is natural—given that we get to observe all outcomes in an independent set, restricting the choice of independent set to be maximal ensures that we observe as much information as possible. Formally, an independent set $S \in \mathcal{I}(\mathcal{H})$ is *maximal* if for any $A \in \mathcal{I}(\mathcal{H})$ with $S \subseteq A$, we have that $A = S$. In other words, S is not contained in a strictly larger independent set. We denote the set of maximal independent sets in $\mathcal{H}$ by $\mathcal{MI}(\mathcal{H})$. In an analogous fashion, we denote the set of probability measures on $\mathcal{MI}(\mathcal{H})$ by $\mathbb{D}(\mathcal{MI}(\mathcal{H}))$ and the set of real-valued measurable functions on $\mathcal{MI}(\mathcal{H}) \times (\mathbb{R}_*)^{V_{\mathcal{H}}}$ as $\Gamma(\mathcal{MI}(\mathcal{H}))$. For a moment parameter $q \geq 2$, we define the *maximal independent set minimax risk* as

$$\mathcal{R}_{\mathcal{MI}}(\mathcal{H}, q) = \inf_{\substack{\mathcal{D} \in \mathbb{D}(\mathcal{MI}(\mathcal{H})) \\ \hat{\tau} \in \Gamma(\mathcal{MI}(\mathcal{H}))}} \sup_{u \in \mathcal{M}(\mathcal{H}, q)} \mathbb{E}_{S \sim \mathcal{D}} \left[(\tau(u) - \hat{\tau}(S, u_S))^2 \right] .$$

Because this minimax risk considers a restricted class of probability distributions and estimators, it follows directly from the definition that $\mathcal{R}(\mathcal{H}, q) \leq \mathcal{R}_{\mathcal{MI}}(\mathcal{H}, q)$.

In order to show that the conflict graph $\mathcal{H}$ characterizes the minimax risk of the original network experiment problem $\mathcal{R}(G, \tau, q)$, we will establish a series of reductions between these three minimax problems:

$$\mathcal{R}(G, \tau, q) \geq \mathcal{R}(\mathcal{H}, q) \geq \mathcal{R}_{\mathcal{MI}}(\mathcal{H}, q) \geq \mathcal{R}(G, \tau, q) ,$$

which will in turn establish equality amongst all three minimax risks. Each inequality will be established in its own lemma.

In order to aid in these reductions, we establish some basic propositions which draw the necessary connections between outcomes which are observed under network interventions $z \in \Omega$ and independent sets in the corresponding conflict graph $\mathcal{H}$. For an intervention $z \in \Omega$, define $R_z \subseteq V_{\mathcal{H}}$ to be the subset of exposures which are observed under z , i.e.

$$R_z = \left\{ e_{i,k} : h_i(z) = e_{i,k} \right\} .$$

First, we show that each intervention $z \in \Omega$ corresponds to an independent set in the conflict graph $\mathcal{H}$.

Proposition A.5. *Fix an interference network G and contrastive estimand τ . For each intervention $z \in \Omega$, the observed exposures R_z form an independent set in the conflict graph $\mathcal{H}$, i.e. $R_z \in \mathcal{I}(\mathcal{H})$.*

Proof. Consider two exposures $e_{i,k}, e_{j,\ell}$ which are both observed under intervention $z \in \Omega$, i.e. $e_{i,k}, e_{j,\ell} \in R_z$. Because they are both observed under z , they cannot be in conflict, by definition. Thus, they are not adjacent in the conflict graph $\mathcal{H}$. This establishes that none of the exposures in R_z are adjacent, therefore it is in an independent set. $\square$

It is worth noting that the reverse direction is not necessarily true, in the following sense: there are independent sets $A \in \mathcal{I}(\mathcal{H})$ for which there is not intervention $z \in \Omega$ such that $R_z = A$. However, the following proposition shows that this reverse direction is true when we restrict A to be a maximal independent set.

Proposition A.6. *Fix an interference network G and contrastive estimand τ . There exists an injective function $f : \mathcal{MI}(\mathcal{H}) \rightarrow \Omega$ with the following property: each maximal independent set $A \in \mathcal{MI}(\mathcal{H})$ has a corresponding intervention $z = f(A)$ satisfying $A = R_z$.*

Proof. Fix a maximal independent set A . We will construct an intervention $z \in \Omega$ such that $R_z = A$. Our construction will proceed in an iterative process. Suppose that A has r vertices, which we arbitrarily order as $e_{i_1, k_1}, e_{i_2, k_2} \dots e_{i_r, k_r}$. Initialize an intervention $z^{(0)} = (0, \dots, 0)$, which corresponds to all subject receiving control. For $t = 1 \dots r$, define the intervention $z^{(t)}$ from $z^{(t-1)}$ by the update: for all $j \in [n]$,

$$z^{(t)}(j) \leftarrow \begin{cases} z^{(t-1)}(j) & \text{if } j \notin \tilde{N}(i_t) \\ 1 & \text{if } j \in \tilde{N}(i_t) \text{ and } j \in e_{i_t, k_t} \end{cases}$$

In other words, we iterate over the exposures in A , updating the intervention so that all exposures $e_{i_1, k_1} \dots e_{i_t, k_t}$ are realized by intervention $z^{(t)}$. An important aspect in this construction is that if $z^{(t)}(j)$ is set to 1 at some iteration t , then it remains set to 1 throughout the process. This is due to the fact that A forms an independent set, and thus all exposures $e_{i_t, k_t} \in A$ agree on the treatment assignments of their extended neighborhoods. Once the iterative procedure finishes, we have that all exposures in A are observed simultaneously under intervention $z = z^{(r)}$.

The argument above establishes that $A \subseteq R_z$ and requires only that A is an independent set in $\mathcal{H}$. In order to establish the reverse containments, i.e. that $R_z \subseteq A$, we need to invoke maximality of A . But this is now simple: because $A \subseteq R_z$, the definition of maximality implies that $A = R_z$.

The argument above shows that the function $f : \mathcal{MI}(\mathcal{H}) \rightarrow \Omega$ with the desired property exists, but it remains to be shown that it is injective. Consider two unique maximal independent set A and A' . By maximality of A and A' , there must exist vertices $e_{i, k} \in A \setminus A'$ and $e_{j, \ell} \in A' \setminus A$ which are adjacent and thus in conflict. In this case, the construction above would assign different individual treatments to at least one common neighbor of i and j . It follows that $f(A) \neq f(A')$ and thus f is injective. $\square$

Finally, we need to endow the conflict graph with a labelling scheme so that the moment restrictions and contrastive estimand agree between the network experiment problem and the independent set problems. Given an interference network G and a contrastive effect τ , we will always assign the following labelling to its vertices: exposure $e_{i, k}$ gets the label $i + n * k$. The following proposition shows that this corresponds to a suitable bijection between the classes $\mathcal{M}(G, \tau, q)$ and $\mathcal{M}(\mathcal{H}, q)$, under which the estimand of interest is also preserved. In a slight abuse of notation, we will always be using the more restricted potential outcome model $\mathcal{M}^0(G, \tau, q)$ considered in Section A.4, though we suppress this notation for simplicity since the minimax risks are equal.

Proposition A.7. *Fix an interference network G , a contrastive estimand τ , and let $\mathcal{H}$ be the correspondingly labelled conflict graph. The function $\phi : \mathcal{M}(G, \tau, q) \rightarrow \mathcal{M}(\mathcal{H}, q)$ given by*

$$\phi(y)(i, k) = y_i(e_{i, k})$$

is a well-defined bijection satisfying $\tau(y) = \tau(\phi(y))$.

Proof. The proof is essentially a bookkeeping exercise. Both claims follow because the bijection maps the potential outcomes $y_i(e_{i,0})$ to the first half of vertices labelled $1 \dots n$ and the potential outcomes $y_i(e_{i,1})$ to the second half of vertices labelled $n + 1 \dots 2n$. $\square$

We are now ready to proceed with the three reductions. Each of these reductions relies on a similar proof strategy, the key structure of which we provide now. Abstractly, we seek to show that the minimax risk for Problem A is greater than that of Problem B. Given a statistical procedure $(\mathcal{D}, \hat{\tau})$ for Problem A, we construct a statistical procedure $(\mathcal{D}^*, \hat{\tau}^*)$ for Problem B so that the maximum risk of the second procedure is upper bounded by that of the first, establishing the minimax reduction. Each construction has essentially the following components:

- **Natural Mapping:** We show the existence of a natural mapping f from the sample space in Problem A to the sample space in Problem B. Importantly, the observed information in Problem B completely determines the observed information in the pre-image of f for Problem A.
- **New Probability Measure:** The new probability measure $\mathcal{D}^*$ is defined as the pushforward measure of the original probability measure $\mathcal{D}$ under this mapping, i.e. roughly speaking $\mathcal{D}^*(E) = \mathcal{D}(f^{-1}(E))$.
- **New Estimator:** The new estimator $\hat{\tau}^*$ is defined as the conditional expectation of the original estimator $\hat{\tau}$ under the pre-image of f , i.e. roughly speaking, $\hat{\tau}^*(E) = \mathbb{E}[\hat{\tau}(x) \mid x \in f^{-1}(E)]$. Because f is such that observed information in Problem B completely determines the observed information in the pre-image of f for Problem A, this is a well-defined estimator for Problem B.
- **Jensen's Inequality:** An application of Jensen's inequality shows that the maximum risk of the original procedure $(\mathcal{D}, \hat{\tau})$ for Problem A provides an upper bound on the maximum risk of the new procedure $(\mathcal{D}^*, \hat{\tau}^*)$ for Problem B. Thus, the minimax risk is bounded as desired.

We will start with the first reduction, which shows that the minimax risk of the network experiment problem is at most the minimax risk of the independent set problem.

Lemma A.8. *Fix an interference network G , a contrastive estimand τ , and let $\mathcal{H}$ be the corresponding conflict graph. For every moment parameter $q \geq 2$, $\mathcal{R}(G, \tau, q) \geq \mathcal{R}(\mathcal{H}, q)$.*

Proof. Fix a design-estimator pair $(\mathcal{D}, \hat{\tau})$ for the network experiment problem, i.e. $\mathcal{D} \in \mathcal{D}(\Omega)$ is a probability measure on interventions and $\hat{\tau} \in \Gamma(\Omega)$ is a real-valued measureable function on network outcomes. Our goal is to construct a procedure $(\mathcal{D}^*, \hat{\tau}^*)$ for the independent set problem, where $\mathcal{D}^* \in \mathbb{D}(\mathcal{I}(\mathcal{H}))$ is a probability measure on independent sets and $\hat{\tau}^* \in \Gamma(\mathcal{I}(\mathcal{H}))$ is a measureable function on the observations from an independent set.

We consider the natural mapping $f : \Omega \rightarrow \mathcal{I}(\mathcal{H})$ from network interventions to independent sets given by Proposition A.5; namely, that the intervention z can be mapped to the independent set R_z of observed exposures in the conflict graph $\mathcal{H}$. Crucially, all interventions z which map to the same independent set R_z provide the same information about the outcomes.

Let $\mathcal{D}^*$ be the pushforward measure of $\mathcal{D}$ thru f . In particular, the new probability measure assigns the following mass to each independent set:

$$\mathcal{D}^*(S) = \sum_{z \in \Omega} \mathbf{1}[f(z) = S] \cdot \mathcal{D}(z) = \mathcal{D}(\{z : f(z) = S\}) .$$

The new estimator $\hat{\tau}^*$ is defined via taking a conditional expectation of the original estimator $\hat{\tau}$ under the original design $\mathcal{D}$ where the conditioning event is that $\{f(z) = S\}$, i.e.

$$\hat{\tau}^*(S, u_S) = \mathbb{E}_{Z \sim \mathcal{D}} \left[\hat{\tau}(Z, \phi^{-1}(u)(Z)) \mid f(z) = S \right] ,$$

where $\phi : \mathcal{M}(G, \tau, q) \rightarrow \mathcal{M}(\mathcal{H}, q)$ is the bijection described in Proposition A.7. Note that $\hat{\tau}^*(S, u_S)$ is in fact a well-defined estimator because the outcomes observed in u_S provide enough information to evaluate $\phi^{-1}(u)(z)$ for each intervention z such that $f(z) = S$.

Finally, let us bound the risk of the new procedure in terms of the old one. Fix functions $y \in \mathcal{M}(G, \tau, q)$ and $u \in \mathcal{M}(\mathcal{H}, q)$ satisfying $\phi(y) = u$. Using Jensen's inequality, we have that

$$\begin{aligned} \mathbb{E}_{z \sim \mathcal{D}} \left[(\tau(y) - \hat{\tau}(z, y(z)))^2 \right] &\geq \mathbb{E}_{z \sim \mathcal{D}} \left[\left(\tau(y) - \mathbb{E}[\hat{\tau}(z, y(z)) \mid f(z) = S] \right)^2 \right] \\ &= \mathbb{E}_{S \sim \mathcal{D}^*} \left[(\tau(u) - \hat{\tau}^*(S, u_S))^2 \right] , \end{aligned}$$

where the last line used the definition of the new procedures $(\mathcal{D}^*, \hat{\tau}^*)$ together with the fact that $\tau(y) = \tau(u)$ for $u = \phi(y)$ (Proposition A.7). Because ϕ is a bijection, we have that the maximum risk of the first procedure $(\mathcal{D}, \hat{\tau})$ is an upper bound on the second procedure $(\mathcal{D}^*, \hat{\tau}^*)$, i.e.

$$\sup_{y \in \mathcal{M}(G, \tau, q)} \mathbb{E}_{z \sim \mathcal{D}} [(\tau(y) - \hat{\tau}(z, y(z)))^2] \geq \sup_{u \in \mathcal{M}(\mathcal{H}, q)} \mathbb{E}_{S \sim \mathcal{D}^*} [(\tau(u) - \hat{\tau}^*(S, u_S))^2] .$$

To establish the claim, we consider the infimum over all procedures $(\mathcal{D}, \hat{\tau})$ and their mappings to $(\mathcal{D}^*, \hat{\tau}^*)$. $\square$

We are now ready to proceed to the next reduction. This time, we show that the independent set minimax risk is at least as large as the restricted independent set minimax risk.

Lemma A.9. *Fix a labeled graph $\mathcal{H}$. For any moment parameter $q \geq 2$, we have that $\mathcal{R}(\mathcal{H}, q) \geq \mathcal{R}_{\mathcal{MI}}(\mathcal{H}, q)$.*

Proof. Fix a pair $(\mathcal{D}, \hat{\tau})$ for the independent set problem where $\mathcal{D} \in \mathbb{D}(\mathcal{I}(\mathcal{H}))$ is a probability measure on independent sets and $\hat{\tau} \in \Gamma(\mathcal{I}(\mathcal{H}))$ is a measurable function of observed outcomes. Our goal is to construct a procedure $(\mathcal{D}^*, \hat{\tau}^*)$ for maximal independent set problem where $\mathcal{D}^* \in \mathbb{D}(\mathcal{MI}(\mathcal{H}))$ is a probability measure on maximal independent sets and $\hat{\tau}^* \in \Gamma(\mathcal{MI}(\mathcal{H}))$ is a measurable function on the observations from maximal independent sets.

To this end, consider any mapping $f : \mathcal{I}(\mathcal{H}) \rightarrow \mathcal{MI}(\mathcal{H})$ that maps independent sets to maximal independent sets and also satisfies the additional inclusion principle: for all independent sets S , we have that $S \subseteq f(S)$. For our purposes, any mapping satisfying the inclusion principle will suffice. Crucially, any such mapping satisfies the following property: the outcomes we observe in a maximal independent set A completely determine all the outcome information we would observe in a subset $S \subseteq A$.

Let $\mathcal{D}^*$ be the pushforward measure of $\mathcal{D}$ through f . In particular, the new probability measure assigns the following mass to each maximal independent set:

$$\mathcal{D}^*(A) = \sum_{S \in \mathcal{I}(\mathcal{H})} \mathbf{1}[f(S) = A] \cdot \mathcal{D}(S) = \mathcal{D}(\{S : f(S) = A\}) .$$

The new estimator is defined via taking a conditional expectation of the original estimator $\hat{\tau}$ under $\mathcal{D}$ where the conditioning event is that $\{f(S) = A\}$, i.e.

$$\hat{\tau}^*(A, u_A) = \mathbb{E}_{S \sim \mathcal{D}} [\hat{\tau}(S, u_S) \mid f(S) = A] .$$

Note that this is indeed a well-defined estimator because the observed outcomes in the larger independent set A provide enough information to evaluate u_S for each subset $S \subseteq A$.

Finally, upper bounding the risk of the new procedure in terms of the old one is an application of Jensen's inequality together with the definition of the newly constructed procedure. For an arbitrary $u \in \mathcal{M}(\mathcal{H}, q)$, we have that

$$\begin{aligned} \mathbb{E}_{S \sim \mathcal{D}} [(\tau(u) - \hat{\tau}(S, u_S))^2] &\geq \mathbb{E}_{S \sim \mathcal{D}} \left[\left(\tau(u) - \mathbb{E}[\hat{\tau}(S, u_S) \mid f(u_S) = A] \right)^2 \right] \\ &= \mathbb{E}_{A \sim \mathcal{D}^*} \left[\left(\tau(u) - \hat{\tau}^*(A, u_A) \right)^2 \right] \end{aligned}$$

The inequality between minimax risks follows by taking the supremum over outcomes and infimum over procedures. $\square$

Finally, we arrive at the third reduction. This time, we show that the minimax risk of the network experiment problem is at most the minimax risk of the associated maximal independent set problem. This reduction is more direct and so we do not need to use any conditional expectation nor Jensen's inequality.

Lemma A.10. *Fix an interference network G , a contrastive estimand τ , and let $\mathcal{H}$ be the corresponding conflict graph. For every moment parameter $q \geq 2$, $\mathcal{R}_{\mathcal{MI}}(\mathcal{H}, q) \geq \mathcal{R}(G, \tau, q)$.*

Proof. Fix a procedure $(\mathcal{D}, \hat{\tau})$ for the maximal independent set problem where $\mathcal{D} \in \mathcal{D}(\mathcal{MI}(\mathcal{H}))$ is a probability measure over maximal independent sets and $\hat{\tau} \in \Gamma(\mathcal{MI}(\mathcal{H}))$ is a real-valued measurable function on maximal independent sets. Our goal is to construct a statistical procedure $(\mathcal{D}^*, \hat{\tau}^*)$ for the network experimental design problem, where $\mathcal{D}^* \in \mathbb{D}(\Omega)$ is a measure on network interventions and $\hat{\tau}^* \in \Gamma(\Omega)$ is a measurable function on the observations.

Consider the injective mapping $f : \mathcal{MI}(\mathcal{H}) \rightarrow \Omega$. The function has two key properties: first, it is injective so that maximal independent sets are in bijection to a subset of interventions. Second, given a maximal independent set A with intervention $f(A) = z$, it holds that A is exactly the observed outcomes under intervention z . Thus, the potential outcomes observed under the intervention z exactly correspond to the vertex outcomes that would be observed under A . This is the crucial property that we need for this final reduction.

The new design and estimator are simple: they are essentially reconstructing exactly the procedure $(\mathcal{D}, \hat{\tau})$. Fix a network intervention $z \in \Omega$. If there exists a maximal independent set A such that $f(A) = z$, then assign probability mass $\mathcal{D}^*(z) = \mathcal{D}(A)$. Otherwise, if no such A exists then $\mathcal{D}^*(z) = 0$. We only need to define the estimator on interventions in the image of f . For these interventions $z \in f(\mathcal{MI}(\mathcal{H}))$, define the new estimator as

$$\hat{\tau}^*(z, y(z)) = \hat{\tau}(A, \phi(y)_A) ,$$

where $A = f(z)$ and ϕ is the bijection in Proposition A.7. Note that $\phi(y)_A$ can be reconstructed exactly from the information observed in $y(z)$, so that $\hat{\tau}^*$ is indeed a well-defined estimator.

Now we show that the risks are equivalent. Consider potential outcome function $y \in \mathcal{M}(G, \tau, q)$ with $u = \phi(y)$. By the definitions above, we have the equality

$$\mathbb{E}_{A \sim \mathcal{D}} \left[(\tau(u) - \hat{\tau}(A, u_A))^2 \right] = \mathbb{E}_{z \sim \mathcal{D}^*} \left[(\tau(y) - \hat{\tau}(z, y(z)))^2 \right] ,$$

where we have used that $\tau(y) = \tau(u)$ from Proposition A.7. Because ϕ is a bijection, the maximum risks of the two procedures are equal:

$$\sup_{u \in \mathcal{M}(\mathcal{H}, q)} \mathbb{E}_{A \sim \mathcal{D}} \left[(\tau(u) - \hat{\tau}(u, u_A))^2 \right] = \sup_{y \in \mathcal{M}(G, \tau, q)} \mathbb{E}_{z \sim \mathcal{D}^*} \left[(\tau(y) - \hat{\tau}(z, y(z)))^2 \right] .$$

The final inequality between the minimax risks follows by taking the infimum of both sides and using the mapping from $(\mathcal{D}, \hat{\tau}) \mapsto (\mathcal{D}^*, \hat{\tau}^*)$. $\square$

After completing these three reductions, we are now ready to prove Proposition 2.2.

Proof of Proposition 2.2. Fix an interference network G and a contrastive estimand τ and let $\mathcal{H}$ be the resulting conflict graph. Fix a moment parameter $q \geq 2$. The lemmas A.8, A.9, and A.10 demonstrate the series of inequalities of the three minimax risks

$$\mathcal{R}(G, \tau, q) \geq \mathcal{R}(\mathcal{H}, q) \geq \mathcal{R}_{\mathcal{MI}}(\mathcal{H}, q) \geq \mathcal{R}(G, \tau, q) ,$$

which in turn establishes that they are all equal. Thus, we have that the minimax risk of the network experiment $\mathcal{R}(G, \tau, q)$ is equal to the minimax risk of the maximum independent set problem $\mathcal{R}_{\mathcal{MI}}(\mathcal{H}, q)$. In this sense, we can take $f_q = \mathcal{R}_{\mathcal{MI}}(\mathcal{H}, q)$. $\square$

The proof above show that the minimax risk of the network experiment problem can be completely characterized by the conflict graph, albeit through a similarly complex minimax formulation. The reductions themselves show that the structure of the maximal independent sets in $\mathcal{H}$ contain the relevant information.

B Lower Bounds on the Minimax Risk

We start by presenting the proof of LeCam's method in Section B.1. We then continue with the proof of the global lower bound in Section B.2, followed by the proof of the local lower bound in Section B.3.

B.1 Le Cam's Method (Proposition 3.1)

For convenience, we restate the proposition here.

Proposition 3.1. *For any network G , causal effect τ , and moment parameter q , the minimax risk is lower bounded as*

$$\mathcal{R}(G, \tau, q) \geq \inf_{\mathcal{D}} \sup_{H_0, H_1} \frac{d^2(H_0, H_1)}{8} \cdot \left(1 - d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})}) \right) ,$$

where d_{TV} is the total variation distance and the infimum is taken over all designs $\mathcal{D}$.

As we discussed in Section 3.1, the proof proceeds in two steps. In the first step, we argue that if we could solve the estimation problem, we could also test between two different hypotheses about the potential outcome functions. In the second step, we connect the error of the best testing procedure for this problem with the total variation distance between the observed distributions under these two hypotheses.

We start by presenting a Lemma regarding the first step.

Lemma B.1. *Let H_0, H_1 be two distributions over $\mathcal{M}(G, \tau, q)$ supported on finite subsets M_0, M_1 respectively. Let*

$$\tau_0 = \inf_{y \in M_0} \tau(y) \quad \text{and} \quad \tau_1 = \sup_{y \in M_1} \tau(y) .$$

Assume that $\tau_1 < \tau_0$. Let $(\mathcal{D}, \hat{\tau})$ be an arbitrary design-estimator pair for effect τ . Define the test function $\psi : \{0, 1\}^n \times \mathbb{R}^n \rightarrow \{0, 1\}$ as

$$\psi(z, y(z)) = \begin{cases} 0 & \text{if } \hat{\tau}(z, y(z)) > \frac{\tau_0 + \tau_1}{2} \\ 1 & \text{otherwise} \end{cases}$$

Then, we have that

- *If $y \in M_0$,*

$$\mathcal{D}(\psi(Z, y(Z)) = 1) \leq \frac{4}{d(H_0, H_1)^2} \cdot \mathbb{E}_{Z \sim \mathcal{D}}[(\hat{\tau}(Z, y(Z)) - \tau(y))^2] .$$

- *If $y \in M_1$,*

$$\mathcal{D}(\psi(Z, y(Z)) = 0) \leq \frac{4}{d(H_0, H_1)^2} \cdot \mathbb{E}_{Z \sim \mathcal{D}}[(\hat{\tau}(Z, y(Z)) - \tau(y))^2] .$$

Proof. We only prove the first statement, since the second one follows by an analogous argument. First of all, by definition of τ_0 and τ_1 , we have that

$$d(H_0, H_1) = \inf_{y \in M_0, y' \in M_1} |\tau(y) - \tau(y')| = \tau_0 - \tau_1 .$$

By definition of ψ , if $\psi(z, y(z)) = 1$, then

$$\hat{\tau}(z, y(z)) \leq \frac{\tau_0 + \tau_1}{2} .$$

For any $y \in M_0$, we have that $\tau(y) \geq \tau_0$. Therefore, in the event that $\psi(z, y(z)) = 1$, we have that

$$(\hat{\tau}(z, y(z)) - \tau(y)) \geq \left(\frac{\tau_0 + \tau_1}{2} - \tau_0 \right) = \frac{d(H_0, H_1)^2}{4} .$$

Thus, using Markov's inequality, we have that

$$\begin{aligned} \mathbb{E}_{Z \sim \mathcal{D}}[(\hat{\tau}(Z, y(Z)) - \tau(y))^2] &\geq \mathcal{D}(\psi(Z, y(Z)) = 1) \cdot \mathbb{E}_{Z \sim \mathcal{D}}[(\hat{\tau}(Z, y(Z)) - \tau(y))^2 \mid \psi(Z, y(Z)) = 1] \\ &\geq \mathcal{D}(\psi(Z, y(Z)) = 1) \cdot \frac{d(H_0, H_1)^2}{4} . \end{aligned}$$

□

Note that in Lemma B.1 the randomness comes only from the treatment assignment Z , while the outcomes are fixed.

We are now ready to present the proof of Proposition 3.1, which includes the second step of connecting the testing error with the total variation distance between the observed distributions.

Proof of Proposition 3.1. We have that

$$\begin{aligned}
\mathcal{R}(G, \tau, q, C) &= \inf_{\mathcal{D}, \hat{\tau}} \sup_{y \in \mathcal{M}(G, \tau, q)} \mathbb{E}_{Z \sim \mathcal{D}} \left[(\hat{\tau}(Z, y(Z)) - \tau(y))^2 \right] \\
&\geq \inf_{\mathcal{D}, \hat{\tau}} \left(\frac{1}{2} \max_{y \in M_0} \mathbb{E}_{Z \sim \mathcal{D}} \left[(\hat{\tau}(Z, y(Z)) - \tau(y))^2 \right] + \frac{1}{2} \max_{y \in M_1} \mathbb{E}_{Z \sim \mathcal{D}} \left[(\hat{\tau}(Z, y(Z)) - \tau(y))^2 \right] \right) \\
&\geq \inf_{\mathcal{D}, \hat{\tau}} \left(\frac{1}{2} \mathbb{E}_{y \sim H_0} \left[\mathbb{E}_{Z \sim \mathcal{D}} \left[(\hat{\tau}(Z, y(Z)) - \tau(y))^2 \right] \right] + \frac{1}{2} \mathbb{E}_{y \sim H_1} \left[\mathbb{E}_{Z \sim \mathcal{D}} \left[(\hat{\tau}(Z, y(Z)) - \tau(y))^2 \right] \right] \right) \\
&\geq \frac{d(H_0, H_1)^2}{8} \inf_{\mathcal{D}, \hat{\tau}} \left(\Pr_{y \sim H_0, Z \sim \mathcal{D}} (\psi(Z, y(Z)) = 1) + \Pr_{y \sim H_1, Z \sim \mathcal{D}} (\psi(Z, y(Z)) = 0) \right),
\end{aligned}$$

where in the last step we used Lemma B.1. Now, define the event $E = \{(z, y) : \psi(z, y) = 1\}$. This event depends on the observed data. Recally that $Q_{H_0}^{(\mathcal{D})}$ and $Q_{H_1}^{(\mathcal{D})}$ are the distributions of the observed data under H_0 and H_1 , respectively (see Section 3.1 for the definition). Then, rewriting the last inequality yields

$$\begin{aligned}
\mathcal{R}(G, \tau, q, C) &\geq \frac{d(H_0, H_1)^2}{8} \inf_{\mathcal{D}, \hat{\tau}} \left(Q_{H_0}^{(\mathcal{D})}(E) + Q_{H_1}^{(\mathcal{D})}(E^c) \right) \\
&= \frac{d(H_0, H_1)^2}{8} \inf_{\mathcal{D}, \hat{\tau}} \left(1 - (Q_{H_1}^{(\mathcal{D})}(E) - Q_{H_0}^{(\mathcal{D})}(E)) \right) \\
&\geq \frac{d(H_0, H_1)^2}{8} \inf_{\mathcal{D}} \left(1 - \sup_{A \in \mathcal{F} \otimes \mathcal{B}(\mathbb{R}^n)} (Q_{H_1}^{(\mathcal{D})}(A) - Q_{H_0}^{(\mathcal{D})}(A)) \right)
\end{aligned}$$

where the supremum is taken over all measurable sets A

$$= \frac{d(H_0, H_1)^2}{8} \inf_{\mathcal{D}} \left(1 - d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})}) \right),$$

where in the last step we used the definition of the total variation distance. $\square$

B.2 Lower Bound using Global Information (Theorem 3.2)

Throughout this section, when we refer to distributions over the outcome functions, they will always be supported on a finite number of elements of $\mathcal{M}(G, \tau, q)$, as noted in Section 3.1. For a distribution H over outcome functions, we use the notation $y \sim H$ to denote that y is drawn according to H . We recall the notation used in Section 3.2. In particular, the following definition is useful.

Definition 3. Let $p \in (0, 1)$ and $T \subseteq V_{\mathcal{H}}$. We define the distribution $H(p, T)$ to be the distribution over outcome functions in $\mathcal{M}(G, \tau, q)$ that is sampled in the following way.

- First, we sample a uniformly random subset of $S \subseteq T$, of size $\lfloor p|T| \rfloor$.
- Then, we set all outcomes for all exposures that are not used to define the effect τ to be 0, i.e.

$$y_i(z) = 0 \quad \text{for all } i \in [n] \text{ and } z \in \{0, 1\}^n \text{ such that } h_i(z) \notin \{e_{i,0}, e_{i,1}\} .$$

- For the remaining exposures, we set

$$y_i(e_{i,k}) = \begin{cases} \left(\frac{n}{|T|}\right)^{\frac{1}{q}} & \text{if } e_{i,k} \in S \text{ and } k = 1 \\ -\left(\frac{n}{|T|}\right)^{\frac{1}{q}} & \text{if } e_{i,k} \in S \text{ and } k = 0 \\ 0 & \text{if } e_{i,k} \notin S \end{cases}$$

In the sequel, we will always assume that p is such that $p|T|$ is an integer, which can always be done by rounding off p .

We will denote by $\mathcal{M}(p, T)$ the support of $H(p, T)$, for any $p \in (0, 1), T \subseteq V_{\mathcal{H}}$. We start by presenting a few immediate properties of the distribution $H(p, T)$.

Lemma B.2. *The following are true for all $p, r \in (0, 1), T \subseteq V_{\mathcal{H}}$, causal effect τ , conflict graph $\mathcal{H}$ and integer $q \geq 2$:*

(i) $\mathcal{M}(p, T) \subseteq \mathcal{M}(G, \tau, q)$.

(ii) For all $y \in \mathcal{M}(p, T)$,

$$\tau(y) = p \cdot \left(\frac{|T|}{n}\right)^{1-\frac{1}{q}} .$$

(iii)

$$d(H(p, T), H(r, T)) = |p - r| \cdot \left(\frac{|T|}{n}\right)^{1-\frac{1}{q}} .$$

Proof. Property (i) follows by noticing that for every $y \in \mathcal{M}(p)$ we have $|y_i(e_{i,k})| \leq (n/|T|)^{1/q}$ for all $i \in [n]$ and $k \in \{0, 1\}$. Therefore, we have that

$$\frac{1}{n} \sum_{i=1}^n (|y_i(e_{i,1})|^q + |y_i(e_{i,0})|^q) \leq \frac{|T|}{n} \cdot \frac{n}{|T|} = 1 .$$

Property (ii) follows by noticing that for every $y \in \mathcal{M}(p, T)$, we have that

$$\begin{aligned} \tau(y) &= \frac{1}{n} \sum_{i=1}^n (y_i(e_{i,1}) + y_i(e_{i,0})) \\ &= \frac{1}{n} \sum_{i=1}^n (\mathbf{1}[e_{i,1} \in S] \cdot \left(\frac{n}{|T|}\right)^{1/q} - \mathbf{1}[e_{i,0} \in S] \cdot \left(-\left(\frac{n}{|T|}\right)^{1/q}\right)) \\ &= \frac{1}{n} \cdot |S| \cdot \left(\frac{n}{|T|}\right)^{1/q} \end{aligned}$$

$$= p \cdot \left(\frac{|T|}{n} \right)^{1-\frac{1}{q}}.$$

Property (iii) follows immediately from the second. $\square$

For an arbitrary distribution on the outcome functions H and an arbitrary design $\mathcal{D}$, we denote by $Q_H^{(\mathcal{D})}$ the distribution of the observed outcomes and treatment assignment when the potential outcomes are drawn from H and the treatment assignment is drawn from $\mathcal{D}$ independently, i.e. the pushforward of the product measure $H \otimes \mathcal{D}$ under the map $(z, y) \mapsto (z, y(z))$. Note that the distribution $Q_H^{(\mathcal{D})}$ depends on the design $\mathcal{D}$. Also, for an intervention $z \in \{0, 1\}^n$, we denote by $Q_H^{(\mathcal{D})}(\cdot|z)$ the conditional distribution of the observed outcomes conditioned on the treatment assignment being equal to z . The next proposition shows that this conditional distribution does not depend on the design $\mathcal{D}$, as long as z has positive probability under $\mathcal{D}$.

Proposition B.3. *For any distribution H over the outcome functions, any designs $\mathcal{D}$ and any intervention $z \in \{0, 1\}^n$ with $\mathcal{D}(z) > 0$, we have that for any $u \in \mathbb{R}^n$,*

$$Q_H^{(\mathcal{D})}(u|z) = \Pr_{y \sim H}(y(z) = u) .$$

Proof. By definition, we have that for any $u \in \mathbb{R}^n$,

$$\begin{aligned} Q_H^{(\mathcal{D})}(u|z) &= \Pr(y(Z) = u \mid Z = z) \\ &= \frac{\Pr(y(Z) = u, Z = z)}{\Pr(Z = z)} \\ &= \frac{\Pr_{y \sim H}(y(z) = u) \mathcal{D}(z)}{\mathcal{D}(z)} \end{aligned}$$

using the independence of the sampled y from the sampled Z

$$= \Pr_{y \sim H}(y(z) = u) .$$

$\square$

Since the conditional distribution does not depend on the design, we will simply denote it by $Q_H(\cdot|z)$ from now on, i.e.

$$Q_H(u|z) \triangleq \Pr_{y \sim H}(y(z) = u) .$$

Property (ii) of Lemma B.2 shows that each distribution $H(p, T)$ is supported on outcomes that have a fixed value of τ . Therefore, if we could observe all counterfactual outcomes, we could easily distinguish between $H(p, T)$ and $H(q, T)$ for any $p \neq q$. However, the fundamental problem of causal inference is that we always observe a subset of the outcomes. As defined in Section 3.2, for any $z \in \{0, 1\}^n$, we denote by R_z the subset of the $2n$ exposures that we get to observe when the treatment assignment is z , i.e.

$$R_z = \left\{ h_i(z) : h_i(z) \in \{e_{i,0}, e_{i,1}\}, i \in [n] \right\} .$$

We now state and prove a lemma, which quantifies the tradeoff between separating the two distributions $H(p, T), H(q, T)$ and distinguishing between them from the observed outcomes. This is a generalization of Lemma 3.3 that appeared in the main text for the case $T = V_{\mathcal{H}}$.

Lemma B.4. *Let $z \in \{0, 1\}^n$ be an intervention, $T \subseteq V_{\mathcal{H}}$ an arbitrary subset with $|T| \geq 5$ and suppose $|R_z \cap T| \leq |T|/2$. Let $H_0 = H(1/2, T)$ and $H_1 = H(1/2 + \delta, T)$, where*

$$\delta \leq \frac{0.026}{\sqrt{|R_z \cap T|}} .$$

Then,

$$d_{\text{TV}}(Q_{H_0}^{(z)}, Q_{H_1}^{(z)}) \leq \frac{7}{8} .$$

Proof. We will assume that

$$\delta \leq \frac{C}{\sqrt{|R_z \cap T|}}$$

where $C > 0$ is a constant that will be determined later according to the constraints.

Any point $u \in \mathbb{R}^n$ that is in the support of $Q_{H_0}^{(z)}$ or $Q_{H_1}^{(z)}$ must have the property that $u_i \in \{-(n/|T|)^{1/q}, 0, (n/|T|)^{1/q}\}$ for all $i \in [n]$, since these are all the possible outcomes under both distributions. Let $U = y(z)$ be the random vector of observed outcomes under $Q_{H_0}^{(z)}$ or $Q_{H_1}^{(z)}$. Let S be the random set that is sampled under H_0 or H_1 . Then, U is uniquely determined by $S \cap (R_z \cap T)$. Indeed, if we know $S \cap (R_z \cap T)$, then we have

$$y_i(z) = \begin{cases} (n/|T|)^{1/q} & \text{if } e_{i,1} \in S \cap (R_z \cap T) \\ -(n/|T|)^{1/q} & \text{if } e_{i,0} \in S \cap (R_z \cap T) \\ 0 & \text{otherwise} \end{cases}$$

Hence, U is determined by $S \cap (R_z \cap T)$. Conversely, if we know U , then we can determine $S \cap (R_z \cap T)$ as follows:

$$S \cap (R_z \cap T) = \left\{ e_{i,1} : i \in [n] \text{ s.t. } y_i(z) = (n/|T|)^{1/q} \right\} \cup \left\{ e_{i,0} : i \in [n] \text{ s.t. } y_i(z) = -(n/|T|)^{1/q} \right\} .$$

Therefore, the distribution of U is the same as the distribution of $S \cap (R_z \cap T)$. This allows us to compute the total variation distance between $Q_{H_0}(\cdot|z)$ and $Q_{H_1}(\cdot|z)$ as follows:

$$\begin{aligned} d_{\text{TV}}(Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z)) &= \frac{1}{2} \sum_{A \subseteq R_z \cap T} \left| \Pr_{H_1}(S \cap (R_z \cap T) = A) - \Pr_{H_0}(S \cap (R_z \cap T) = A) \right| \\ &= \frac{1}{2} \sum_{A \subseteq R_z \cap T} \Pr_{H_0}(S \cap (R_z \cap T) = A) \left| \frac{\Pr_{H_1}(S \cap (R_z \cap T) = A)}{\Pr_{H_0}(S \cap (R_z \cap T) = A)} - 1 \right| . \end{aligned}$$

Let us now compute these probabilities. For convenience we denote $s = |R_z \cap T|$ and $t = |T|$. Suppose $|A| = k$ and we will restrict our attention to k such that $|k - s/2| \leq \sqrt{s/2}$. We will show later that it is enough to restrict to such subsets A .

Then, under H_0 we have $|S| = t/2$. There are $\binom{t-s}{t/2-k}$ ways to choose the remaining $t/2 - k$ elements of S that are not in $R_z \cap T$, which leads to

$$\Pr_{H_0}(S \cap (R_z \cap T) = A) = \frac{\binom{t-s}{t/2-k}}{\binom{t}{t/2}} .$$

By assumption, $s \leq t/2$. This means that $t - s \geq t/2$, and hence the above probability is positive for all A with $|A| \leq s$.

On the other hand, under H_1 we have $|S| = t/2 + \delta \cdot t = t(1/2 + \delta)$. There are $\binom{t-s}{t(1/2+\delta)-k}$ ways to choose the remaining $t(1/2 + \delta) - k$ elements of S that are not in $R_z \cap T$, which leads to

$$\Pr_{H_1}(S \cap (R_z \cap T) = A) = \frac{\binom{t-s}{t(1/2+\delta)-k}}{\binom{t}{t(1/2+\delta)}} .$$

Note that for this probability to be nonzero, we need $t(1/2 + \delta) - k \leq t - s$, since we need to have enough elements outside of $R_z \cap T$ to choose from. Using the fact that $k \geq s/2 - \sqrt{s/2}$, it is sufficient to have

$$\frac{s}{2} + \sqrt{\frac{s}{2}} \leq t\left(\frac{1}{2} - \delta\right)$$

Also, using the assumption $s \leq t/2$, it is sufficient to have

$$\frac{t}{4} + \sqrt{\frac{t}{4}} \leq t\left(\frac{1}{2} - \delta\right) .$$

Since $\delta \leq C$, for the above to be satisfied, it is sufficient to have

$$C \leq \frac{1}{4} - \frac{1}{2\sqrt{t}} . \quad (4)$$

For the right hand side to be positive, we need $t \geq 5$. We will choose C such that (4) holds, so for now we can assume that for all k such that $|k - s/2| \leq \sqrt{s/2}$, both probabilities under H_0 and H_1 are nonzero.

For such a subset A of size k , we have

$$\begin{aligned} \left| \frac{\Pr_{H_1}(S \cap (R_z \cap T) = A)}{\Pr_{H_0}(S \cap (R_z \cap T) = A)} - 1 \right| &= \left| \frac{\binom{t-s}{t(1/2+\delta)-k}}{\binom{t}{t(1/2+\delta)}} \cdot \frac{\binom{t}{t/2-k}}{\binom{t}{t/2-k}} - 1 \right| \\ &= \left| \frac{\prod_{l=0}^{s-k-1} (t(1/2 - \delta) - l) \cdot \prod_{l=0}^{k-1} (t(1/2 + \delta) - l)}{\prod_{l=0}^{s-k-1} (t/2 - l) \cdot \prod_{l=0}^{k-1} (t/2 - l)} - 1 \right| \\ &= \left| \prod_{l=0}^{s-k-1} \left(1 - \frac{2\delta t}{t - 2l}\right) \cdot \prod_{l=0}^{k-1} \left(1 + \frac{2\delta t}{t - 2l}\right) - 1 \right| \end{aligned}$$

We would like to turn these products into a single exponential factor. To do that, we will utilize the following elementary inequality, which can be proved via taking derivatives.

$$1 + x = \exp(x + b(x)) \quad \text{where} \quad |b(x)| \leq 2x^2 \quad \text{for all } |x| \leq 1/2 . \quad (5)$$

To apply this inequality, we need to ensure that the absolute value of each factor in the products is at most $1/2$. For all $l \leq k - 1$ we have

$$\begin{aligned} \left| \frac{2\delta t}{t - 2l} \right| &\leq \frac{2\delta t}{t - 2(k - 1)} \\ &\leq \frac{2\delta t}{t - 2(s/2 + \sqrt{s/2} - 1)} \end{aligned}$$

$$\begin{aligned}
&\leq \frac{2\delta t}{t/2 - \sqrt{t} + 2} \\
&\leq C \frac{2t}{t/2 - \sqrt{t} + 2}
\end{aligned}$$

Therefore, we have the following constraint on C .

$$C \leq \frac{1}{8} - \frac{1}{4\sqrt{t}} + \frac{1}{2t} . \quad (6)$$

We will choose $C > 0$ so that (6) is satisfied. The same exact constraint arises in the case $l \leq s - k - 1$. Therefore, applying inequality (5) to each factor in the products, we have that

$$\begin{aligned}
&\prod_{l=0}^{s-k-1} \left(1 - \frac{2\delta t}{t-2l}\right) \cdot \prod_{l=0}^{k-1} \left(1 + \frac{2\delta t}{t-2l}\right) \\
&= \exp(\delta_1(t) + \delta_2(t)) ,
\end{aligned}$$

where

$$\begin{aligned}
\delta_1(t) &= - \sum_{l=0}^{s-k-1} \frac{2\delta t}{t-2l} + \sum_{l=0}^{k-1} \frac{2\delta t}{t-2l} \\
\delta_2(t) &= \sum_{l=0}^{s-k-1} b\left(\frac{2\delta t}{t-2l}\right) + \sum_{l=0}^{k-1} b\left(\frac{2\delta t}{t-2l}\right)
\end{aligned}$$

Let us bound $|\delta_1(t)|$ first. Assume without loss of generality that $k \geq s/2$, so that $k-1 \geq s-k-1$. The other case is symmetric. Then, we have

$$\begin{aligned}
|\delta_1(t)| &= \left| - \sum_{l=0}^{s-k-1} \frac{2\delta t}{t-2l} + \sum_{l=0}^{k-1} \frac{2\delta t}{t-2l} \right| \\
&= \left| \sum_{l=s-k}^{k-1} \frac{2\delta t}{t-2l} \right| \\
&\leq |2k - s| \cdot \frac{2\delta t}{t - 2(k-1)} \\
&\leq \frac{2\sqrt{2s}\delta t}{t - 2(s + s/2 - 1)}
\end{aligned}$$

using $k \leq s/2 + \sqrt{s/2}$

$$\leq \frac{2\sqrt{2}t}{t/2 - \sqrt{t} + 2} \cdot C .$$

We now bound the second error term.

$$|\delta_2(t)| = \left| \sum_{l=0}^{s-k-1} b\left(\frac{2\delta t}{t-2l}\right) + \sum_{l=0}^{k-1} b\left(\frac{2\delta t}{t-2l}\right) \right|$$

$$\begin{aligned}
&\leq 2 \cdot \sum_{l=0}^{s-k-1} \left(\frac{2\delta t}{t-2l} \right)^2 + 2 \cdot \sum_{l=0}^{k-1} \left(\frac{2\delta t}{t-2l} \right)^2 \\
&\leq 8(s-k)\delta^2 t^2 \cdot \frac{1}{(t-2(s-k-1))^2} + 8k\delta^2 t^2 \cdot \frac{1}{(t-2(k-1))^2} \\
&\leq 8s\delta^2 \cdot \frac{t^2}{(t/2 - \sqrt{t} + 2)^2} \\
&\leq C^2 \cdot \frac{8t^2}{(t/2 - \sqrt{t} + 2)^2} ,
\end{aligned}$$

where we have used again the facts $k \leq s/2 + \sqrt{s/2}$ and $s \leq t/2$. As a result, we have that for any k such that $|k - s/2| \leq \sqrt{s}$,

$$\begin{aligned}
\left| \frac{\Pr_{H_1}(S \cap (R_z \cap T) = A)}{\Pr_{H_0}(S \cap (R_z \cap T) = A)} - 1 \right| &= \left| \exp\left(\delta_1(t) + \delta_2(t)\right) - 1 \right| \\
&\leq \exp\left(|\delta_1(t)| + |\delta_2(t)|\right) - 1
\end{aligned}$$

using the fact that $|e^x - 1| \leq e^{|x|} - 1$ for all x . We would like to establish that

$$\left| \frac{\Pr_{H_1}(S \cap (R_z \cap T) = A)}{\Pr_{H_0}(S \cap (R_z \cap T) = A)} - 1 \right| < \frac{1}{2} , \tag{7}$$

which by the previous calculations is implied by the following inequality

$$|\delta_1(t)| + |\delta_2(t)| < \log(3/2) .$$

Let us denote

$$D(t) = \frac{t}{t/2 - \sqrt{t} + 2} .$$

Then, combining the bounds on $|\delta_1(t)|$ and $|\delta_2(t)|$, it is sufficient to have

$$2\sqrt{2}D(t)C + 8D(t)^2C^2 < \log(3/2) .$$

Solving this quadratic inequality for C , we have that it is sufficient to have

$$C < \frac{-\sqrt{2} + \sqrt{2 + 8\log(3/2)}}{8D(t)} . \tag{8}$$

Thus, inequality (7) will hold as long as the three constraints (4), (6) and (8) are satisfied for $C > 0$ and $t \geq 5$. Since all the constraints are improving as t increases, it is sufficient to check that they are satisfied for $t = 5$. Substituting $t = 5$ into the right hand sides of the three constraints, we get the conditions

$$C \leq 0.026$$

$$C \leq 0.113$$

$$C \leq 0.049$$

Among these, we choose the most restrictive one, which is $C \leq 0.026$.

We next show that the contribution of the sets A with $|k - s/2| \leq \sqrt{s}$ to the total variation distance is non-negligible. To do that, it suffices to show that the total probability of such sets under H_0 is lower bounded by a constant. We can prove this is the case using Chebyshev's inequality. In particular, we have that

$$|A| = \sum_{i \in R_z \cap T} \mathbf{1}[i \in S]$$

Let us calculate the expectation and variance of this random variable under H_0 . By linearity of expectation

$$\mathbb{E}_{H_0}[|A|] = \sum_{i \in R_z \cap T} \Pr_{H_0}(i \in S) = |R_z \cap T| \cdot \frac{1}{2} = \frac{s}{2},$$

using the fact that $\Pr_{H_0}(i \in S) = 1/2$ for all $i \in R_z \cap T$. We also have

$$\begin{aligned} \text{Var}_{H_0}(|A|) &= \sum_{i \in R_z \cap T} \text{Var}_{H_0}(\mathbf{1}[i \in S]) + \sum_{i \neq j \in R_z \cap T} \text{Cov}_{H_0}(\mathbf{1}[i \in S], \mathbf{1}[j \in S]) \\ &= \frac{s}{4} + s(s-1) \left(\frac{t/2(t/2-1)}{t(t-1)} - \frac{1}{4} \right) \\ &\leq \frac{s}{4}. \end{aligned}$$

Therefore, by Chebyshev's inequality, we have that

$$\Pr_{H_0} \left(\left| |A| - \frac{s}{2} \right| \geq \sqrt{\frac{s}{2}} \right) \leq \frac{\text{Var}_{H_0}(|A|)}{\frac{s}{2}} \leq \frac{1}{2}.$$

We are now ready to upper bound the total variation distance. In particular, define the events E, F as follows

$$\begin{aligned} E &= \left\{ A : \left| |A| - \frac{s}{2} \right| \leq \sqrt{\frac{s}{2}} \right\} \\ F &= \left\{ A : \Pr_{H_0}(S \cap (R_z \cap T) = A) \geq \Pr_{H_1}(S \cap (R_z \cap T) = A) \right\} \end{aligned}$$

We have established that $\Pr_{H_0}(E) \geq 1/2$. This means that $\Pr_{H_0}(E \cap F) \geq 1/4$ or $\Pr_{H_0}(E \cap F^c) \geq 1/4$. Without loss of generality, let's assume the first is true, otherwise a similar argument applies. Then, we can write the total variation distance as

$$\begin{aligned} d_{\text{TV}}(Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z)) &= \sum_{A \in F} \Pr_{H_0}(S \cap (R_z \cap T) = A) - \Pr_{H_1}(S \cap (R_z \cap T) = A) \\ &= \sum_{A \in E \cap F} \left(\Pr_{H_0}(S \cap (R_z \cap T) = A) - \Pr_{H_1}(S \cap (R_z \cap T) = A) \right) \\ &\quad + \sum_{A \in E^c \cap F} \left(\Pr_{H_0}(S \cap (R_z \cap T) = A) - \Pr_{H_1}(S \cap (R_z \cap T) = A) \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{A \in E \cap F} \Pr_{H_0}(S \cap (R_z \cap T) = A) \cdot \left(1 - \frac{\Pr_{H_1}(S \cap (R_z \cap T) = A)}{\Pr_{H_0}(S \cap (R_z \cap T) = A)}\right) \\
&\quad + \sum_{A \in E^c \cap F} \left(\Pr_{H_0}(S \cap (R_z \cap T) = A) - \Pr_{H_1}(S \cap (R_z \cap T) = A)\right) \\
&< \frac{1}{2} \cdot \Pr_{H_0}(E \cap F) + \Pr_{H_0}(E^c \cap F) \\
&= \Pr_{H_0}(F) - \frac{1}{2} \cdot \Pr_{H_0}(E \cap F) \\
&\leq 1 - \frac{1}{2} \cdot \frac{1}{4} \\
&= \frac{7}{8} .
\end{aligned}$$

This concludes the proof. $\square$

We are now ready to prove Theorem 3.2.

Proof of Theorem 3.2. Let's fix an arbitrary subset $T \subseteq V_{\mathcal{H}}$ of size $t = |T| \geq 5$. Recall that we denote by $\mathcal{I}(T, \mathcal{H})$ the largest independent set in the subgraph of the conflict graph $\mathcal{H}$ induced by T . Assume for now that $|\mathcal{I}(T, \mathcal{H})| \leq |T|/2$. We define $H_0 = H(1/2, T)$ and $H_1 = H(1/2 + \delta, T)$ for $\delta = 0.026/\sqrt{|\mathcal{I}(T, \mathcal{H})|}$. By property (i) of Lemma B.2, we have that the support of H_0 and H_1 is contained in $\mathcal{M}(G, \tau, q)$ for any $q \geq 2$.

Next, we establish that for any $z \in \{0, 1\}^n$, we have $|R_z \cap T| \leq |\mathcal{I}(T, \mathcal{H})|$. It suffices to show that the exposures in $R_z \cap T$ is an independent set in the induced subgraph $\mathcal{H}(T)$. To see that, note that for any $e_{i,k}, e_{j,l} \in R_z \cap T$, by definition we have $h_i(z) = e_{i,k}$ and $h_j(z) = e_{j,l}$. This means that $e_{i,k}, e_{j,l}$ are not connected in the conflict graph $\mathcal{H}$, thus also not connected in the induced subgraph $\mathcal{H}(T)$. Thus $R_z \cap T$ is an independent set of $\mathcal{H}(T)$.

Let $\mathcal{D}$ be an arbitrary design and let $\text{supp}(\mathcal{D}) \subseteq \{0, 1\}^n$ be the support of $\mathcal{D}$. As we have remarked earlier, all outcomes in the support of H_0 and H_1 take values in the set

$$\{-(|T|/n)^{-1/q}, 0, (|T|/n)^{-1/q}\} .$$

Therefore, the observed outcomes $y(Z)$ under H_0 or H_1 take values in the set

$$U_n = \{-(|T|/n)^{-1/q}, 0, (|T|/n)^{-1/q}\}^n .$$

We can write the total variation distance between the observed distributions under H_0 and H_1 as follows

$$\begin{aligned}
d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})}) &= \frac{1}{2} \sum_{z \in \{0,1\}^n} \sum_{u \in U_n} \left| \Pr_{H_1}(Z = z, y(Z) = u) - \Pr_{H_0}(Z = z, y(Z) = u) \right| \\
&= \frac{1}{2} \sum_{z \in \text{supp}(\mathcal{D})} \mathcal{D}(z) \cdot \sum_{u \in U_n} \left| \Pr_{H_1}(y(z) = u | Z = z) - \Pr_{H_0}(y(z) = u | Z = z) \right| \\
&= \sum_{z \in \text{supp}(\mathcal{D})} \mathcal{D}(z) \cdot d_{\text{TV}}(Q_{H_0}(\cdot | z), Q_{H_1}(\cdot | z))
\end{aligned}$$

$$\leq \sum_{z \in \text{supp}(\mathcal{D})} \mathcal{D}(z) \cdot \frac{7}{8}$$

using Lemma 3.3 and the fact that $|R_z \cap T| \leq |\mathcal{I}(T)|$ and $|\mathcal{I}(T)| \leq |T|/2$

$$= \frac{7}{8} .$$

By property (iii) of Lemma B.2 we have that

$$d(H_0, H_1) = \delta\left(\frac{|T|}{n}\right)^{1-\frac{1}{q}} = \frac{0.026}{\sqrt{|\mathcal{I}(T)|}} \cdot \left(\frac{|T|}{n}\right)^{1-\frac{1}{q}} .$$

We will use LeCam's method (Proposition 3.1) with the same choice of the two distributions H_0 and H_1 for all designs $\mathcal{D}$. Applying the bound yields

$$\mathcal{R}(G, \tau, q) \geq \frac{d(H_0, H_1)^2}{8} \cdot \inf_{\mathcal{D}} \left(1 - d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})})\right) \geq 1.05 \cdot 10^{-5} \cdot \frac{1}{|\mathcal{I}(\mathcal{H})|} \cdot \left(\frac{|T|}{n}\right)^{2-\frac{2}{q}} .$$

This proves the desired lower bound for the case $|\mathcal{I}(T, \mathcal{H})| \leq |T|/2$. We now handle the case $|\mathcal{I}(T, \mathcal{H})| > |T|/2$. First of all, by assumption we have that $(e_{i,0}, e_{i,1}) \in E(\mathcal{H}(T))$ for all $i \in [n]$, which means that any independent set in $\mathcal{H}$ contains at most one of the two exposures for each unit. Therefore, we have $|R_z| \leq n$. This means that $T = V_{\mathcal{H}}$ satisfies $|\mathcal{I}(T, \mathcal{H})| \leq |T|/2$, so applying the previously obtained bound with $T = V_{\mathcal{H}}$ yields

$$\mathcal{R}(G, \tau, q) \geq 1.05 \cdot 10^{-5} \cdot 2^{2-\frac{2}{q}} \cdot \frac{1}{n} .$$

Now, consider an arbitrary subset $T \subseteq V_{\mathcal{H}}$ with $|\mathcal{I}(T, \mathcal{H})| > |T|/2$. We have that

$$\begin{aligned} 1.05 \cdot 10^{-5} \cdot \frac{1}{|\mathcal{I}(T, \mathcal{H})|} \cdot \left(\frac{|T|}{n}\right)^{2-\frac{2}{q}} &\leq 1.05 \cdot 10^{-5} \cdot \frac{2}{|T|} \cdot \left(\frac{|T|}{n}\right)^{2-\frac{2}{q}} \\ &= 1.05 \cdot 10^{-5} \cdot \frac{2|T|^{1-\frac{2}{q}}}{n^{2-\frac{2}{q}}} \\ &\leq 1.05 \cdot 10^{-5} \cdot 2^{2-\frac{2}{q}} \cdot \frac{1}{n} \\ &\leq \mathcal{R}(G, \tau, q) , \end{aligned}$$

where in the last step we used the previously obtained bound for $T = V_{\mathcal{H}}$. Thus, the bound holds for all subsets $T \subseteq V_{\mathcal{H}}$ with $|T| \geq 5$. The proof is now complete. $\square$

B.3 Lower Bound using Local Information (Theorem 3.4)

In this section, we present the proof of Theorem 3.4. Let's recall that for any exposure $e_{i,k} \in T$, we denote by $d_{i,k}$ the degree of $e_{i,k}$ in the conflict graph $\mathcal{H}$. We also denote by $d_{\min}(T)$ the minimum degree of any exposure in T , and by $d_{\text{harm}}(T, r)$ the harmonic mean of the degrees of the exposures in T with parameter r , i.e.

$$d_{\min}(T) = \min_{e_{i,k} \in T} d_{i,k} \quad , \quad d_{\text{harm}}(T, r) = \frac{|T|}{\left(\sum_{e_{i,k} \in T} 1/d_{i,k}^r\right)^{\frac{1}{r}}} .$$

We omit the dependence on $\mathcal{H}$ since it is clear from the context. We repeat the statement for completeness.

Theorem 3.4. *Let G be an interference network, τ be a contrastive effect. If $n \geq 5$ then for all moment parameters $q \geq 2$,*

$$\mathcal{R}(G, \tau, q) \geq \frac{6.6 \cdot 10^{-7}}{n^{2-\frac{2}{q}}} \sup_{T \subseteq V_{\mathcal{H}}, |T| \geq 5} \min \left(d_{\min}(\mathcal{H}, T), \sqrt{d_{\text{harm}}(\mathcal{H}, T, 2 - 2/q)}, |T| \right)^{2-\frac{2}{q}}.$$

Before proving Theorem 3.4, we establish an intermediate Lemma that contains the main idea of the argument. It establishes a lower bound on the minimax risk in terms of some tunable parameters $c_{i,k}$.

Lemma B.5. *Let G be a graph on n vertices, τ an arbitrary effect and $\mathcal{H}$ the conflict graph. Then, for any subset of exposures $T \subseteq V_{\mathcal{H}}$, any integer $q \geq 2$, any integers $r_{i,k}$ that satisfy $r_{i,k} \leq d_{i,k}$ for all $i \in [n], k \in \{0, 1\}$ and any positive numbers $c_{i,k} \in (0, 1)$ for $i \in [n], k \in \{0, 1\}$ that satisfy*

$$\sum_{e_{i,k} \in T} c_{i,k} \geq 1.36 \cdot 10^{-3},$$

we have

$$\mathcal{R}(G, \tau, q) \geq \frac{2.64 \cdot 10^{-6}}{n^{2-\frac{2}{q}}} \cdot \min \left(\min_{e_{i,k} \in T} c_{i,k} r_{i,k}^{2-\frac{2}{q}}, \frac{|T|^{2-\frac{2}{q}}}{\sum_{e_{i,k} \in T} c_{i,k}} \right).$$

Proof. We define the event $A_{i,k}$ as the event that i gets exposure $e_{i,k}$, i.e.

$$A_{i,k} = \{h_i(z) = e_{i,k}\}$$

Recall that for each exposure $e_{i,k} \in T$, $\mathcal{N}(e_{i,k})$ is the set of neighboring exposures of $e_{i,k}$ in $\mathcal{H}$. By definition, $|\mathcal{N}(e_{i,k})| = d_{i,k}$. Consider an arbitrary subset $U_{i,k} \subseteq \mathcal{N}(e_{i,k})$ of size $r_{i,k}$ for each $e_{i,k} \in T$.

Let $\mathcal{D}$ be an arbitrary design. We distinguish between two cases.

Case 1: There exists $i \in [n], k \in \{0, 1\}$ with $e_{i,k} \in T$ such that $\mathcal{D}(A_{i,k}) \geq c_{i,k}$. We define H_0 and H_1 to be supported on potential outcome functions $y^{(0)}, y^{(1)}$ respectively, where

$$y_i^{(1)}(e) = \begin{cases} \left(\frac{n}{r_{i,k}}\right)^{1/q} & \text{if } e = e_{j,1} \in U_{i,k} \\ -\left(\frac{n}{r_{i,k}}\right)^{1/q} & \text{if } e = e_{j,0} \in U_{i,k} \\ 0 & \text{otherwise} \end{cases}, \quad y_i^{(0)}(e) = \begin{cases} -\left(\frac{n}{r_{i,k}}\right)^{1/q} & \text{if } e = e_{j,1} \in U_{i,k} \\ \left(\frac{n}{r_{i,k}}\right)^{1/q} & \text{if } e = e_{j,0} \in U_{i,k} \\ 0 & \text{otherwise} \end{cases}.$$

Let us first verify that these outcome functions satisfy the moment restriction. We have for $l \in \{0, 1\}$,

$$\sum_{j=1}^n (|y_j^{(l)}(e_{j,1})|^q + |y_j^{(l)}(e_{j,0})|^q) = \sum_{e \in U_{i,k}} \left(\left(\frac{n}{r_{i,k}} \right)^{1/q} \right)^q = n.$$

Therefore, $y^{(0)}, y^{(1)} \in \mathcal{M}(G, q)$. Since H_0 and H_1 are supported on single potential outcome functions, their separation is simply

$$d(H_0, H_1) = |\tau(y^{(0)}) - \tau(y^{(1)})|$$

$$\begin{aligned}
&= \frac{1}{n} \sum_{e \in U_{i,k}} \left(\left(\frac{n}{r_{i,k}} \right)^{1/q} - \left(- \left(\frac{n}{r_{i,k}} \right)^{1/q} \right) \right) \\
&= 2 \cdot \left(\frac{r_{i,k}}{n} \right)^{1 - \frac{1}{q}}.
\end{aligned}$$

Suppose event $A_{i,k}$ occurs, i.e. we observe the outcome of unit i for exposure $e_{i,k}$. By definition of the conflict graph, this implies that we do not observe the outcomes of any of the exposures in $U_{i,k}$. Since H_0, H_1 only differ in the outcomes of exposures in $U_{i,k}$, they are impossible to distinguish under that event. In other words, for any $z \in A_{i,k}$, we have $Q_{H_0}(\cdot|z) = Q_{H_1}(\cdot|z)$ and hence

$$d_{\text{TV}}(Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z)) = 0,$$

which implies that

$$\begin{aligned}
d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})}) &= \sum_{z \in \text{supp}(\mathcal{D})} \mathcal{D}(z) \cdot d_{\text{TV}}(Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z)) \\
&= \sum_{z \notin A_{i,k}} \mathcal{D}(z) \cdot d_{\text{TV}}(Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z)) \\
&\leq 1 - \mathcal{D}(A_{i,k}) \leq 1 - c_{i,k}.
\end{aligned}$$

Following the proof of Proposition 3.1 for this particular design $\mathcal{D}$, we have

$$\begin{aligned}
\inf_{\hat{\tau}} \text{Risk}(\mathcal{D}, \hat{\tau}) &\geq \frac{d^2(H_0, H_1)}{8} \cdot (1 - d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})})) \\
&\geq \frac{1}{2} \cdot \left(\frac{r_{i,k}}{n} \right)^{2 - \frac{2}{q}} \cdot c_{i,k}.
\end{aligned}$$

Case 2: For all $i \in [n], k \in \{0, 1\}$ with $e_{i,k} \in T$ we have $\mathcal{D}(A_{i,k}) < c_{i,k}$. In that case, define the event A_T as

$$A_T = \left\{ z \in \{0, 1\}^n : \sum_{e_{i,k} \in T} \mathbf{1}[A_{i,k}](z) \leq 2 \cdot \sum_{e_{i,k} \in T} c_{i,k} \right\} = \left\{ z \in \{0, 1\}^n : |R_z \cap T| \leq 2 \cdot \sum_{e_{i,k} \in T} c_{i,k} \right\}.$$

By Markov's inequality, we have

$$\begin{aligned}
\mathcal{D}(A_T^c) &= \mathcal{D} \left(\sum_{e_{i,k} \in T} \mathbf{1}[A_{i,k}] > 2 \cdot \sum_{e_{i,k} \in T} c_{i,k} \right) \\
&\leq \frac{\mathbb{E} \left[\sum_{e_{i,k} \in T} \mathbf{1}[A_{i,k}] \right]}{2 \cdot \sum_{e_{i,k} \in T} c_{i,k}} \\
&= \frac{\sum_{e_{i,k} \in T} \mathcal{D}(A_{i,k})}{2 \cdot \sum_{e_{i,k} \in T} c_{i,k}} \\
&< \frac{1}{2}.
\end{aligned}$$

Thus, $\mathcal{D}(A_T) \geq 1/2$. Now, define $H_0 = H(1/2, T)$ and $H_1 = H(1/2 + \delta, T)$ where

$$\delta = \frac{0.026}{\sqrt{2 \cdot \sum_{e_{i,k} \in T} c_{i,k}}} ,$$

where 0.026 is the constant from Lemma B.4. Note that for this construction to be well-defined, we need $\delta < 1/2$, which is equivalent to

$$\sum_{e_{i,k} \in T} c_{i,k} > 2 \cdot 0.026^2 \approx 1.352 \cdot 10^{-3} .$$

Thus, we choose the constant $1.36 \cdot 10^{-3}$ in the statement of the lemma to ensure the above inequality is satisfied.

By Lemma B.2

$$d(H_0, H_1) = \delta \cdot \left(\frac{|T|}{n} \right)^{1 - \frac{1}{q}} = \frac{0.026}{\sqrt{2 \sum_{e_{i,k} \in T} c_{i,k}}} \cdot \left(\frac{|T|}{n} \right)^{1 - \frac{1}{q}} .$$

For any $z \in A_T$, we have

$$|R_z \cap T| \leq 2 \cdot \sum_{e_{i,k} \in T} c_{i,k} \leq \frac{|T|}{2} .$$

Therefore, applying Lemma B.4, we have that for any $z \in A_T$

$$d_{\text{TV}}(Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z)) \leq \frac{7}{8} .$$

By the law of total probability, we have

$$\begin{aligned} 1 - d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})}) &= \sum_{z \in \text{supp}(\mathcal{D})} \mathcal{D}(z) \cdot (1 - d_{\text{TV}}(Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z))) \\ &\geq \sum_{z \in A_T} \mathcal{D}(z) \cdot (1 - d_{\text{TV}}(Q_{H_0}(\cdot|z), Q_{H_1}(\cdot|z))) \\ &\geq \mathcal{D}(A_T) \cdot \left(1 - \frac{7}{8}\right) \\ &\geq \frac{1}{16} . \end{aligned}$$

Therefore, following the proof of Proposition 3.1 for this particular design $\mathcal{D}$, we have

$$\begin{aligned} \inf_{\hat{\tau}} \text{Risk}(\mathcal{D}, \hat{\tau}) &\geq \frac{d^2(H_0, H_1)}{8} \cdot (1 - d_{\text{TV}}(Q_{H_0}^{(\mathcal{D})}, Q_{H_1}^{(\mathcal{D})})) \\ &\geq \frac{0.026^2}{2^8 \cdot \sum_{e_{i,k} \in T} c_{i,k}} \cdot \left(\frac{|T|}{n} \right)^{2 - \frac{2}{q}} \\ &\geq \frac{2.64 \cdot 10^{-6}}{\sum_{e_{i,k} \in T} c_{i,k}} \left(\frac{|T|}{n} \right)^{2 - \frac{2}{q}} . \end{aligned}$$

Since any possible design $\mathcal{D}$ falls into either Case 1 or Case 2, we can combine the lower bounds from both cases (inequalities B.3 and B.3) to obtain a lower bound on the minimax risk as follows

$$\mathcal{R}(G, \tau, q) \geq 2.64 \cdot 10^{-6} \cdot \min \left(\min_{e_{i,k} \in T} c_{i,k} \left(\frac{r_{i,k}}{n} \right)^{2-\frac{2}{q}}, \frac{1}{\sum_{e_{i,k} \in T} c_{i,k}} \left(\frac{|T|}{n} \right)^{2-\frac{2}{q}} \right)$$

□

We are now ready to prove Theorem 3.4. The general strategy will be to invoke Lemma B.5 with the appropriate parameters. As we shall see, this choice depends crucially on the properties of each subset T of exposures.

Proof of Theorem 3.4. We distinguish between various cases based on the properties of the subset T .

Case 1: Suppose $d_{\min}(T) < |T|$. In that case, it suffices to prove that

$$\mathcal{R}(G, \tau, q) \geq \frac{2.64 \cdot 10^{-6}}{n^{2-\frac{2}{q}}} \cdot \min \left(\frac{d_{\min}(T)^{2-\frac{2}{q}}}{4}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}}} \right).$$

Suppose $e_{\min} \in T$ is an exposure of minimum degree $d_{\min}(T)$ in T . Set $r_{i,k} = d_{i,k}$ for all $e_{i,k} \in T$. Let

$$t^* = \frac{1}{n^{2-\frac{2}{q}}} \cdot \min \left(\frac{\min_{e_{i,k} \in T} r_{i,k}^{2-\frac{2}{q}}}{4}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}}} \right)$$

$$c_{i,k} = t^* \left(\frac{n}{r_{i,k}} \right)^{2-\frac{2}{q}}$$

We will invoke Lemma B.5 with the above choice of $r_{i,k}$ and $c_{i,k}$. Let us first check that the $c_{i,k}$ satisfy the required assumptions of Lemma B.5. We have

$$c_{i,k} \leq \frac{1}{n^{2-\frac{2}{q}}} \cdot \left(\frac{n}{r_{i,k}} \right)^{2-\frac{2}{q}} \cdot \frac{\min_{e_{i,k} \in T} r_{i,k}^{2-\frac{2}{q}}}{4} \leq 1.$$

Furthermore, we have

$$\sum_{e_{i,k} \in T} c_{i,k} = t^* \cdot n^{2-\frac{2}{q}} \cdot \sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}$$

$$= \min \left(\frac{\min_{e_{i,k} \in T} r_{i,k}^{2-\frac{2}{q}}}{4}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}}} \right) \cdot \sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}$$

$$= \min \left(\frac{\min_{e_{i,k} \in T} r_{i,k}^{2-\frac{2}{q}}}{4} \sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}, |T|^{1-\frac{1}{q}} \cdot \sqrt{\sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}} \right)$$

From here, we can lower bound the quantity as follows.

$$\sum_{e_{i,k} \in T} c_{i,k} \geq \frac{1}{4} \geq 1.36 \cdot 10^{-3} ,$$

where we lower bounded the sum by the term corresponding to the exposure $e_{\min}$.

We can also upper bound the sum as follows.

$$\begin{aligned} \sum_{e_{i,k} \in T} c_{i,k} &\leq \frac{\min_{e_{i,k} \in T} r_{i,k}^{2-\frac{2}{q}}}{4} \sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}} \\ &\leq \frac{1}{4} \cdot |T| . \end{aligned}$$

Therefore, all assumptions of Lemma B.5 are satisfied. Invoking the lemma, we get

$$\begin{aligned} \mathcal{R}(G, \tau, q) &\geq \frac{2.64 \cdot 10^{-6}}{n^{2-\frac{2}{q}}} \cdot \min \left(\min_{e_{i,k} \in T} c_{i,k} r_{i,k}^{2-\frac{2}{q}}, \frac{|T|^{2-\frac{2}{q}}}{\sum_{e_{i,k} \in T} c_{i,k}} \right) \\ &= \frac{2.64 \cdot 10^{-6}}{n^{2-\frac{2}{q}}} \cdot \min \left(t^* n^{2-\frac{2}{q}}, \frac{|T|^{2-\frac{2}{q}}}{t^* n^{2-\frac{2}{q}} \cdot \sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}} \right) \\ &\geq 2.64 \cdot 10^{-6} \cdot \min \left(t^*, \frac{|T|^{1-\frac{1}{q}}}{n^{2-\frac{2}{q}} \cdot \sqrt{\sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}}} \right) \\ &\geq 2.64 \cdot 10^{-6} \cdot t^* , \end{aligned}$$

where we used the definition of t^* in the last two steps. Thus, our claim is complete in that case.

Case 2: Suppose $d_{\min}(T) \geq |T|$. In that case, it suffices to prove that

$$\mathcal{R}(G, \tau, q) \geq \frac{6.6 \cdot 10^{-7}}{n^{2-\frac{2}{q}}} \min \left(|T|^{2-\frac{2}{q}}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}}} \right) . \quad (9)$$

We distinguish between two sub-cases.

- **Case 2.1:** Suppose

$$\frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}}} \leq |T|^{2-\frac{2}{q}} ,$$

which can equivalently written as

$$|T|^{1-\frac{1}{q}} \sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}} \geq 1 . \quad (10)$$

In that case, set $r_{i,k} = d_{i,k}$ for all $e_{i,k} \in T$ and define

$$t^* = \frac{1}{n^{2-\frac{2}{q}}} \min \left(\frac{|T|^{2-\frac{2}{q}}}{4}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}}} \right)$$

$$c_{i,k} = t^* \left(\frac{n}{d_{i,k}} \right)^{2-\frac{2}{q}}$$

We will show that the $c_{i,k}$ satisfy the assumptions of Lemma B.5. We have

$$c_{i,k} \leq \frac{1}{n^{2-\frac{2}{q}}} \left(\frac{n}{d_{i,k}} \right)^{2-\frac{2}{q}} \cdot \frac{|T|^{2-\frac{2}{q}}}{4} \leq \frac{d_{\min}(T)^{2-\frac{2}{q}}}{4d_{i,k}^{2-\frac{2}{q}}} \leq 1 ,$$

where we have used the fact that $|T| \leq d_{\min}(T)$. Furthermore,

$$\sum_{e_{i,k} \in T} c_{i,k} = t^* \cdot n^{2-\frac{2}{q}} \cdot \sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}$$

$$= \min \left(\frac{|T|^{2-\frac{2}{q}}}{4} \cdot \sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}, |T|^{1-\frac{1}{q}} \cdot \sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}} \right) .$$

We can lower bound the sum as

$$\sum_{e_{i,k} \in T} c_{i,k} \geq \frac{1}{4} \geq 1.36 \cdot 10^{-3} ,$$

where in the last step we used (10). We can also upper bound the sum as follows.

$$\sum_{e_{i,k} \in T} c_{i,k} \leq \frac{|T|^{2-\frac{2}{q}}}{4} \sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}$$

$$\leq \frac{d_{\min}(T)^{2-\frac{2}{q}}}{4} \sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}$$

$$\leq \frac{1}{4} \cdot |T| ,$$

where we used the fact that $|T| \leq d_{\min}(T)$ in the last step. Thus, all assumptions of Lemma B.5 are satisfied. Invoking the lemma, we get

$$\mathcal{R}(G, \tau, q) \geq \frac{2.64 \cdot 10^{-6}}{n^{2-\frac{2}{q}}} \cdot \min \left(\min_{e_{i,k} \in T} c_{i,k} d_{i,k}^{2-\frac{2}{q}}, \frac{|T|^{2-\frac{2}{q}}}{\sum_{e_{i,k} \in T} c_{i,k}} \right)$$

$$\geq \frac{2.64 \cdot 10^{-6}}{n^{2-\frac{2}{q}}} \cdot \min \left(t^* n^{2-\frac{2}{q}}, \frac{|T|^{2-\frac{2}{q}}}{t^* n^{2-\frac{2}{q}} \cdot \sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}} \right)$$

using the fact that $d_{\min}(T) \geq |T|$

$$\begin{aligned} &\geq 2.64 \cdot 10^{-6} \cdot \min \left(t^*, \frac{|T|^{1-\frac{1}{q}}}{n^{2-\frac{2}{q}} \cdot \sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}}} \right) \\ &\geq 2.64 \cdot 10^{-6} \cdot t^* , \end{aligned}$$

where we used the definition of t^* in the last two steps. Inequality (9) now follows.

- **Case 2.2:** Suppose

$$\frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}}} > |T|^{2-\frac{2}{q}} ,$$

which can equivalently be written as

$$|T|^{1-\frac{1}{q}} \sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}} < 1 .$$

We will show that we can pick $r_{i,k} \leq d_{i,k}$ with $r_{i,k} \geq |T|$ for all $e_{i,k} \in T$ such that

$$1 \leq |T|^{2-\frac{2}{q}} \sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}} \leq 2 \quad (11)$$

Indeed, we can achieve this by iteratively picking $e_{i,k} \in T$ with $r_{i,k} > |T|$ and setting $r_{i,k} = r_{i,k} - 1$ until the left-hand side of (11) is satisfied. We first prove that this process will terminate in a finite number of steps and $r_{i,k} \geq |T|$ at termination. Indeed, suppose we reach the point where $r_{i,k} = |T|$ for all $e_{i,k} \in T$. Then, we have

$$|T|^{2-\frac{2}{q}} \sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}} = |T|^{2-\frac{2}{q}} \sum_{e_{i,k} \in T} \frac{1}{|T|^{2-\frac{2}{q}}} = |T| \geq 1 .$$

Therefore, the process will terminate at some point and at termination we have $r_{i,k} \geq |T|$ for all $e_{i,k} \in T$. Furthermore, consider a single update step where we pick $e_{i,k} \in T$ and set $r_{i,k} = r_{i,k} - 1$. The value of the objective changes by at most

$$|T|^{2-\frac{2}{q}} \cdot \left(\frac{1}{(r_{i,k} - 1)^{2-\frac{2}{q}}} - \frac{1}{r_{i,k}^{2-\frac{2}{q}}} \right) \leq 1$$

since $r_{i,k} > |T|$ by definition of the process. Thus, the objective increases by at most 1 at every update. Therefore, the first time the objective exceeds 1, (11) will be satisfied. From now on, let $r_{i,k}$ be the values obtained at termination of the above process. We set

$$t^* = \frac{1}{n^{2-\frac{2}{q}}} \min \left(\frac{\min_{e_{i,k} \in T} r_{i,k}^{2-\frac{2}{q}}}{4}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}}} \right)$$

$$c_{i,k} = t^* \left(\frac{n}{r_{i,k}} \right)^{2-\frac{2}{q}}$$

We will show that the $c_{i,k}$ satisfy the assumptions of Lemma B.5. We have

$$c_{i,k} \leq \frac{1}{n^{2-\frac{2}{q}}} \left(\frac{n}{r_{i,k}} \right)^{2-\frac{2}{q}} \cdot \frac{\min_{e_{i,k} \in T} r_{i,k}^{2-\frac{2}{q}}}{4} \leq 1 .$$

Furthermore,

$$\begin{aligned} \sum_{e_{i,k} \in T} c_{i,k} &= t^* \cdot n^{2-\frac{2}{q}} \cdot \sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}} \\ &= \min \left(\frac{\min_{e_{i,k} \in T} r_{i,k}^{2-\frac{2}{q}}}{4} \cdot \sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}, |T|^{1-\frac{1}{q}} \cdot \sqrt{\sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}} \right) \\ &\geq \min \left(|T|^{2-\frac{2}{q}} \cdot \sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}, |T|^{1-\frac{1}{q}} \cdot \sqrt{\sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}} \right) \\ &\geq 1.36 \cdot 10^{-3} , \end{aligned}$$

where in the last step we used (11). We can also upper bound the sum as follows.

$$\begin{aligned} \sum_{e_{i,k} \in T} c_{i,k} &\leq \frac{\min_{e_{i,k} \in T} r_{i,k}^{2-\frac{2}{q}}}{4} \sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}} \\ &\leq \frac{1}{4} \cdot |T| , \end{aligned}$$

Thus, all assumptions of Lemma B.5 are satisfied. Invoking the lemma, we get

$$\begin{aligned} \mathcal{R}(G, \tau, q) &\geq \frac{2.64 \cdot 10^{-6}}{n^{2-\frac{2}{q}}} \cdot \min \left(\min_{e_{i,k} \in T} c_{i,k} r_{i,k}^{2-\frac{2}{q}}, \frac{|T|^{2-\frac{2}{q}}}{\sum_{e_{i,k} \in T} c_{i,k}} \right) \\ &\geq 2.64 \cdot 10^{-6} \cdot \min \left(t^*, \frac{|T|^{1-\frac{1}{q}}}{n^{2-\frac{2}{q}} \cdot \sqrt{\sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}}} \right) \\ &\geq 2.64 \cdot 10^{-6} \cdot t^* . \end{aligned} \tag{12}$$

Now all we have to do is connect t^* to the desired lower bound of (9). This is because the lower bound is in terms of the original degrees, whereas in t^* some degrees have been “truncated”. We have

$$\begin{aligned}
t^* &= \frac{1}{n^{2-\frac{2}{q}}} \min \left(\frac{\min_{e_{i,k} \in T} r_{i,k}^{2-\frac{2}{q}}}{4}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}}} \right) \\
&\geq \frac{1}{n^{2-\frac{2}{q}}} \min \left(\frac{|T|^{2-\frac{2}{q}}}{4}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{r_{i,k}^{2-\frac{2}{q}}}}} \right) \\
&\geq \frac{1}{n^{2-\frac{2}{q}}} \min \left(\frac{|T|^{2-\frac{2}{q}}}{4}, \frac{|T|^{2-\frac{2}{q}}}{\sqrt{2}} \right)
\end{aligned}$$

where we use (11)

$$\begin{aligned}
&= \frac{1}{4} \frac{1}{n^{2-\frac{2}{q}}} |T|^{2-\frac{2}{q}} \\
&\geq \frac{1}{4} \frac{1}{n^{2-\frac{2}{q}}} \min \left(|T|^{2-\frac{2}{q}}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}}} \right)
\end{aligned}$$

Plugging this back to (12), we obtain

$$\mathcal{R}(G, \tau, q) \geq \frac{2.64 \cdot 10^{-6}}{4} \cdot \frac{1}{n^{2-\frac{2}{q}}} \min \left(|T|^{2-\frac{2}{q}}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}}} \right),$$

which establishes the desired lower bound in that case.

We have established the desired lower bound in all cases, therefore the proof is complete. $\square$

We now give the proof of Corollary 3.5, which is restated here for completeness. Recall that in Section 3.3 we defined $F_{\mathcal{H}}$ to be the degree distribution function of the conflict graph $\mathcal{H}$. We also defined the *critical degree* of a conflict graph $\mathcal{H}$ with respect to a moment $q \geq 2$ as follows

$$d_q^*(\mathcal{H}) = \sup \left\{ d \in [n] : 1 - F_{\mathcal{H}}(d) \geq \frac{d^{\frac{2q-2}{3q-2}}}{n} \right\}.$$

Corollary 3.5. *Let G be an interference network, τ be a contrastive effect. For all moment parameters $q \geq 2$, we have*

$$\mathcal{R}(G, \tau, q) \geq 6.6 \cdot 10^{-7} \cdot \frac{(d_q^*(\mathcal{H}))^{\frac{4(q-1)^2}{q(3q-2)}}}{n^{2-\frac{2}{q}}}.$$

Proof. By definition of $d_q^*(\mathcal{H})$, there exist a subset of exposures $T \subseteq V_{\mathcal{H}}$ with $|T| \geq (d_q^*(\mathcal{H}))^{\frac{2(q-1)}{3q-2}}$ such that $d_{i,k} \geq d_q^*(\mathcal{H})$ for all $e_{i,k} \in T$. Therefore, we can invoke Theorem 3.4 with the above choice of T to get

$$\begin{aligned}
\mathcal{R}(G, \tau, q) &\geq \frac{6.6 \cdot 10^{-7}}{n^{2-\frac{2}{q}}} \cdot \min \left(|T|^{2-\frac{2}{q}}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\sum_{e_{i,k} \in T} \frac{1}{d_{i,k}^{2-\frac{2}{q}}}}}, d_{\min}(T)^{2-\frac{2}{q}} \right) \\
&\geq \frac{6.6 \cdot 10^{-7}}{n^{2-\frac{2}{q}}} \cdot \min \left(|T|^{2-\frac{2}{q}}, \frac{|T|^{1-\frac{1}{q}}}{\sqrt{\frac{|T|}{(d_q^*(\mathcal{H}))^{2-\frac{2}{q}}}}}, (d_q^*(\mathcal{H}))^{2-\frac{2}{q}} \right) \\
&= \frac{6.6 \cdot 10^{-7}}{n^{2-\frac{2}{q}}} \cdot \min \left(|T|^{2-\frac{2}{q}}, |T|^{\frac{1}{2}-\frac{1}{q}} \cdot (d_q^*(\mathcal{H}))^{1-\frac{1}{q}}, (d_q^*(\mathcal{H}))^{2-\frac{2}{q}} \right) \\
&\geq \frac{6.6 \cdot 10^{-7}}{n^{2-\frac{2}{q}}} \cdot \min \left((d_q^*(\mathcal{H}))^{\frac{2q-2}{3q-2}(2-\frac{2}{q})}, (d_q^*(\mathcal{H}))^{\frac{2q-2}{3q-2}(\frac{1}{2}-\frac{1}{q})+1-\frac{1}{q}}, (d_q^*(\mathcal{H}))^{2-\frac{2}{q}} \right) \\
&= 6.6 \cdot 10^{-7} \cdot \frac{(d_q^*(\mathcal{H}))^{\frac{4(q-1)^2}{q(3q-2)}}}{n^{2-\frac{2}{q}}} .
\end{aligned}$$

□

Finally, we present the proof of Proposition B.6 that lower bounds the critical degree by the average degree of the conflict graph $\mathcal{H}$.

Proposition B.6. *Let G be a graph and τ a contrastive effect, with $\mathcal{H}$ being the associated conflict graph. We denote by $d_{\text{avg}}(\mathcal{H})$ the average degree of $\mathcal{H}$. Then, for any $q \geq 2$, we have*

$$d_q^*(\mathcal{H}) \geq \frac{d_{\text{avg}}(\mathcal{H})}{4} .$$

Proof. Let $A, B \subseteq V_{\mathcal{H}}$ be the subsets of exposures that have degree at least $d_q^*(\mathcal{H}) + 1$ and at most $d_q^*(\mathcal{H})$, respectively, i.e.

$$\begin{aligned}
A &= \left\{ e_{i,k} \in V_{\mathcal{H}} : d_{i,k} \geq d_q^*(\mathcal{H}) + 1 \right\} \\
B &= \left\{ e_{i,k} \in V_{\mathcal{H}} : d_{i,k} \leq d_q^*(\mathcal{H}) \right\} .
\end{aligned}$$

Clearly A and B partition $V_{\mathcal{H}}$. By definition of the critical degree, we have

$$|A| \leq (d_q^*(\mathcal{H}) + 1)^{\frac{2q-2}{3q-2}} .$$

Thus, we have that

$$\begin{aligned}
d_{\text{avg}}(\mathcal{H}) &= \frac{\sum_{e_{i,k} \in V_{\mathcal{H}}} d_{i,k}}{|V_{\mathcal{H}}|} \\
&= \frac{\sum_{e_{i,k} \in A} d_{i,k} + \sum_{e_{i,k} \in B} d_{i,k}}{|V_{\mathcal{H}}|}
\end{aligned}$$

$$\begin{aligned}
&\leq \frac{|A|n + |B|d_q^*(\mathcal{H})}{|V_{\mathcal{H}}|} \\
&\leq \frac{2n(d_q^*(\mathcal{H}) + 1)^{\frac{2q-2}{3q-2}} + 2nd_q^*(\mathcal{H})}{|V_{\mathcal{H}}|} \\
&= (d_q^*(\mathcal{H}) + 1)^{\frac{2q-2}{3q-2}} + d_q^*(\mathcal{H}) \\
&\leq 2^{\frac{2q-2}{3q-2}} \cdot (d_q^*(\mathcal{H}))^{\frac{2q-2}{3q-2}} + d_q^*(\mathcal{H}) \\
&\leq 4d_q^*(\mathcal{H}) ,
\end{aligned}$$

using the fact that $2^{\frac{2q-2}{3q-2}} \leq 3$ for all $q \geq 2$. Rearranging the above inequality, we get the conclusion. $\square$

C Upper Bounds on the Minimax Risk

In this section, we present the proofs of our general upper bounds on the minimax risk. We start by presenting the analysis of the generalized Conflict Graph Design in Section C.1 and then continue to the optimal bias-variance tradeoff in Section C.2.

C.1 Analysis of the Conflict Graph Design (Theorem 4.1)

In this section, we present the proof of Theorem 4.1, which we restate here for convenience.

Theorem 4.1. *Let G be an interference network, τ be a contrastive effect. For any moment parameter $q \geq 2$, the minimax risk is bounded as*

$$\mathcal{R}(G, \tau, q) \leq 8.92 \cdot \frac{\lambda(\mathcal{H})^{\frac{2}{q}} \cdot d_{\text{avg}}(\mathcal{H})^{1-\frac{2}{q}}}{n} ,$$

where $\lambda(\mathcal{H})$ is the largest eigenvalue of the adjacency matrix of $\mathcal{H}$ and $d_{\text{avg}}(\mathcal{H})$ is the average degree of $\mathcal{H}$. Moreover, this bound on the minimax risk is achieved by an unbiased statistical procedure.

For $q = 2$, this upper bound was established in Kandiros et al. (2024). In particular, they introduced the Conflict Graph Design (CGD) and provided an unbiased estimator which achieves a worst case variance of $\mathcal{O}(\lambda(G)/n)$. We will follow the general strategy of the Conflict Graph Design, although as we will see there are important obstacles to overcome for general moments.

We now provide a roadmap for this section.

- In Section C.1.1, we introduce the importance ordering that will be used for our design and establish a crucial property. In Section C.1.2 we describe the design for arbitrary moments q .
- In Section C.1.3 we introduce the estimator and prove some basic properties of it.
- In Section C.1.4 we bound the covariance terms that arise in the variance of the estimator.
- In Section C.1.5 we bound suitable norms of matrices that are related to the covariance matrix of the design.
- In Section C.1.6 we present the proof of Theorem 4.1 using all the above ingredients.

Throughout this section, we will denote by G the interference graph, τ the effect of interest, $\mathcal{H}$ the conflict graph associated with G and τ . For a node $e_{i,k} \in V(\mathcal{H})$, we will denote by $\mathcal{N}(e_{i,k})$ the set of neighbors of $e_{i,k}$ in $\mathcal{H}$. We denote by $d_{i,k} = |\mathcal{N}(e_{i,k})|$ the degree of node $e_{i,k}$ in the conflict graph $\mathcal{H}$. Also, for a unit $i \in V(G)$, we denote by $\tilde{N}(i)$ the set consisting of i and its neighbors in G .

C.1.1 Importance Ordering

An *importance ordering* is a bijection $\pi : V(\mathcal{H}) \mapsto [2n]$ that “ranks” the vertices in the conflict graph $\mathcal{H}$ from the most important to the least important. For an importance ordering π and a node $e_{i,k} \in V(\mathcal{H})$, we denote by $\mathcal{N}_b^\pi(e_{i,k})$ the set of neighbors of $e_{i,k}$ in $\mathcal{H}$ that appear before $e_{i,k}$ in the importance ordering π , i.e.

$$\mathcal{N}_b^\pi(e_{i,k}) = \{e_{j,l} \in \mathcal{N}(e_{i,k}) : \pi(e_{j,l}) < \pi(e_{i,k})\} .$$

In Kandiros et al. (2024), they proposed an importance ordering based on the maximum eigenvector in the graph that ensures every node has a bounded number of more important neighbors. This property was crucial for upper bounding the variance of their design.

We will show that there is a different ordering that behaves more favorably for general moments q . We first provide pseudocode for the procedure that produces the ordering in Algorithm 1.

Algorithm 1: Importance Ordering

Input: Conflict graph $\mathcal{H} = (V_{\mathcal{H}}, E_{\mathcal{H}})$ and moment q .

Output: Mapping $\pi : V_{\mathcal{H}} \rightarrow [2n]$ representing the importance ordering.

```

1  $S \leftarrow V_{\mathcal{H}}$ 
2 for  $e_{i,k} \in V_{\mathcal{H}}$  do
3    $\left[ \text{Set } p_{i,k} \leftarrow \frac{1}{2\lambda(\mathcal{H})^{\frac{2}{q}} d_{i,k}^{1-\frac{2}{q}}} \right]$ 
4 for  $c = 1$  to  $2n$  do
5   for  $e_{i,k} \in S$  do
6      $\left[ w_{i,k} \leftarrow \sum_{e_{j,l} \in \mathcal{N}(e_{i,k}) \cap S} p_{j,l} \right]$ 
7      $e_{i^*,k^*} \leftarrow \arg \min_{e_{i,k} \in S} w_{i,k}$ 
8      $\pi(e_{i^*,k^*}) = 2n - c + 1$ 
9      $S \leftarrow S \setminus \{e_{i^*,k^*}\}$ 
```

The ordering π in Algorithm 1 is constructed iteratively, starting from the least important node and ending with the most important node. At each step, S represents the set of nodes that have not been ordered yet. We pick the node whose neighbors in S have the smallest sum of $p_{i,k}$ values and assign it the next available position in the ordering.

In the case $q = 2$, we observe that the ordering of Algorithm 1 is constructed by iteratively selecting the node with smallest degree in the remaining induced graph, placing them in the last available spot in the ordering and removing them from the graph. This ordering of the nodes is used in the classical proof of Wilf’s Theorem (Wilf, 1967), which states that the chromatic number of a graph is upper bounded by the largest eigenvalue of its adjacency matrix. In Kandiros et al. (2024), they use a different ordering based on the leading eigenvector of the adjacency matrix, even

though they acknowledge that the ordering based on the smallest induced degree also satisfies the desired property.

Since the value of the $p_{i,k}$ variables is different for different values of q , the ordering produced by Algorithm 1 will be different for different values of q . Indeed, the ordering for $q = \infty$ is determined by the degrees in the original graph, while for $q = 2$ it is constructed using the degrees in the induced subgraph that is iteratively produced. Thus, the ordering could be very different for different values of q .

In the next lemma, we establish the crucial property of the ordering produced by Algorithm 1 that will be used in the analysis of our design later.

Lemma C.1. *Let $\mathcal{H}$ be a conflict graph with maximum eigenvalue $\lambda(\mathcal{H})$. Let $q \geq 2$ be an integer. For every node $e_{i,k} \in V(\mathcal{H})$ with degree $d_{i,k}$ in the conflict graph $\mathcal{H}$, let*

$$p_{i,k} = \frac{1}{2\lambda(\mathcal{H})^{\frac{2}{q}} d_{i,k}^{1-\frac{2}{q}}} .$$

Suppose $\pi : V_{\mathcal{H}} \rightarrow [2n]$ is the ordering of the vertices in $\mathcal{H}$ that is produced by Algorithm 1 with input $\mathcal{H}$ and q . Then, for every node $e_{i,k} \in V_{\mathcal{H}}$

$$\sum_{e_{j,l} \in \mathcal{N}_b^{\pi}(e_{i,k})} p_{j,l} \leq 1/2 .$$

Proof. For an arbitrary subset $S \subseteq [2n]$, we denote by $d_{i,k}(S)$ the number of neighbors of $e_{i,k}$ in S , i.e., $d_{i,k}(S) = |\mathcal{N}(e_{i,k}) \cap S|$. Also, for every node $e_{i,k}$, we define the quantity

$$w_{i,k}(S) = \sum_{e_{j,l} \in \mathcal{N}(e_{i,k}) \cap S} p_{j,l}$$

We will show that for every subset S , there exists a node $e_{i,k} \in S$ with $w_{i,k}(S) \leq 1/2$. It suffices to show that the average value of $w_{i,k}(S)$ over all $e_{i,k} \in S$ is at most $1/2$. Indeed, we have

$$\begin{aligned} \frac{1}{|S|} \sum_{e_{i,k} \in S} w_{i,k}(S) &= \frac{1}{|S|} \sum_{e_{i,k} \in S} \sum_{e_{j,l} \in \mathcal{N}(e_{i,k}) \cap S} p_{j,l} \\ &= \frac{1}{|S|} \sum_{e_{i,k} \in S} \sum_{e_{j,l} \in \mathcal{N}(e_{i,k}) \cap S} \frac{1}{2\lambda(G)^{2/q} d_{j,l}^{1-2/q}} \\ &= \frac{1}{|S|} \sum_{e_{j,l} \in S} \sum_{e_{i,k} \in \mathcal{N}(e_{j,l}) \cap S} \frac{1}{2\lambda(G)^{2/q} d_{j,l}^{1-2/q}} \end{aligned}$$

by exchanging the order of summation of i and j

$$\begin{aligned} &= \frac{1}{|S|} \sum_{e_{j,l} \in S} \frac{d_{j,l}(S)}{2\lambda(G)^{2/q} d_{j,l}^{1-2/q}} \\ &\leq \frac{1}{|S|} \sum_{e_{j,l} \in S} \frac{d_{j,l}(S)}{2\lambda(G)^{2/q} d_{j,l}(S)^{1-2/q}} \end{aligned}$$

using the fact that $d_{j,l}(S) \leq d_{j,l}$ for all S by definition

$$\begin{aligned}
&= \frac{1}{|S|} \sum_{e_{j,l} \in S} \frac{d_{j,l}(S)^{2/q}}{2\lambda(G)^{2/q}} \\
&= \frac{1}{2\lambda(G)^{2/q}} \frac{1}{|S|} \sum_{e_{j,l} \in S} d_{j,l}(S)^{2/q} \\
&\leq \frac{1}{2\lambda(G)^{2/q}} \left(\frac{1}{|S|} \sum_{e_{j,l} \in S} d_{j,l}(S) \right)^{2/q}
\end{aligned}$$

using Jensen's inequality for the function $x \mapsto x^{2/q}$, which is concave for $q \in (2, \infty)$

$$\leq \frac{1}{2\lambda(G)^{2/q}} (\lambda(G))^{2/q}$$

using the fact that the average degree of any induced subgraph of G is at most the maximum eigenvalue

$$= 1/2 \ .$$

We now use this property to analyze Algorithm 1. Let $e_{i,k} \in V_{\mathcal{H}}$ be an arbitrary node. We know that $e_{i,k}$ will be selected exactly once as e_{j^*,l^*} in line 7, since this line is executed exactly $|V_{\mathcal{H}}|$ times and every time a different node is selected. Let c be the iteration in which $e_{i,k}$ is selected as e_{j^*,l^*} . When line 8 is executed in that iteration, we have $\pi(e_{i,k}) = |V_{\mathcal{H}}| - i + 1$. Moreover, at that point the set S contains exactly the nodes that are more important than $e_{i,k}$ in the ordering π , since all the less important nodes have already been selected in previous iterations and removed from S . Thus, we have that $\mathcal{N}_b^\pi(e_{i,k}) = \mathcal{N}(e_{i,k}) \cap S$. Therefore,

$$\sum_{e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})} p_{j,l} = \sum_{e_{j,l} \in \mathcal{N}(e_{i,k}) \cap S} p_{j,l} = w_{i,k}(S) \ .$$

By the property we established above, we have

$$\sum_{e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})} p_{j,l} = w_{i,k}(S) = \min_{e_{j,l} \in S} w_{j,l}(S) \leq \frac{1}{|S|} \sum_{e_{j,l} \in S} w_{j,l}(S) \leq 1/2 \ ,$$

as desired. Since $e_{i,k}$ was an arbitrary node, this establishes the desired property for all nodes in the graph. $\square$

In the limit $q = \infty$, Lemma C.1 provides a way of selecting the ordering. However, it turns out that in this special case there is a alternative way of ordering the nodes, which is simply to order them in decreasing order of degree. We present this property as a separate lemma for completeness.

Lemma C.2. For every node $e_{i,k} \in V(\mathcal{H})$ with degree $d_{i,k}$ in the conflict graph $\mathcal{H}$, let

$$p_{i,k} = \frac{1}{2d_{i,k}} .$$

Suppose $\pi : V_{\mathcal{H}} \rightarrow [2n]$ is the ordering of the vertices in decreasing value of their degree in $\mathcal{H}$, i.e. $d_{i,k} \leq d_{j,l}$ if and only if $\pi_{i,k} \geq \pi_{j,l}$. Then, for every node $e_{i,k}$

$$\sum_{e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})} p_{j,l} \leq 1/2 .$$

Proof. Let $e_{i,k}$ be an arbitrary node in $V_{\mathcal{H}}$. By definition of the ordering, we have that for every $e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})$ it holds that $d_{j,l} \geq d_{i,k}$. Thus, we have

$$\sum_{e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})} p_{j,l} = \sum_{e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})} \frac{1}{2d_{j,l}} \leq \sum_{e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})} \frac{1}{2d_{i,k}} \leq \frac{d_{i,k}}{2d_{i,k}} = 1/2 ,$$

as desired. $\square$

C.1.2 Conflict Graph Design for General Moment Restrictions

Since Theorem 4.1 aims to upper bound the minimax risk, it suffices to present a design-estimator pair whose risk is upper bounded by the desired quantity. We start by providing the details of our design for arbitrary moments q , given in Algorithm 2.

Algorithm 2: Conflict Graph Design (CGD)

Input: Maximum eigenvalue $\lambda(\mathcal{H})$, graph G , conflict graph $\mathcal{H}$ and moment q .

Output: Random intervention $Z \in \mathcal{Z} = \{0, 1\}^n$, variables U_i for all $i \in V(\mathcal{H})$.

- 1 Importance ordering $\pi \leftarrow$ Algorithm 1 with input $\mathcal{H}$ and q ;
- 2 **for** $e_{i,k} \in V(\mathcal{H})$ **do**
- 3 $p_{i,k} \leftarrow \frac{1}{2\lambda(\mathcal{H})^{\frac{2}{q}} d_{i,k}^{1-\frac{2}{q}}}$;
- 4 $C \leftarrow 2$;
- 5 Sample desired exposure variables $\{U_{i,k} : e_{i,k} \in V(\mathcal{H})\}$ independently and identically as

$$U_{i,k} \leftarrow \begin{cases} e & \text{with probability } \frac{p_{i,k}}{C} \\ * & \text{with probability } 1 - \frac{p_{i,k}}{C} \end{cases}$$

Initialize the intervention vector $Z \leftarrow \mathbf{0}$;

- 6 **for** $e_{i,k} \in V(\mathcal{H})$ **do**
 - 7 **if** $U_{i,k} = e$ and $U_{r,l} = *$ for all $e_{r,l} \in \mathcal{N}_b^\pi(e_{i,k})$ **then**
 - 8 Update intervention vector: set $Z_j = \mathbf{1}[j \in h_i(e_{i,k})]$ for all $j \in \tilde{N}(i)$;
-

The design takes as input an importance ordering π , which is simply a total ordering of the vertices in the graph G , the maximum eigenvalue $\lambda(\mathcal{H})$ of the adjacency matrix of the conflict graph $\mathcal{H}$, the graph G itself and the moment q . Note that even though the goal is to generate the random

intervention $Z \in \{0, 1\}^n$, the algorithm also outputs the variables $U_{i,k}$ for all $e_{i,k} \in V(\mathcal{H})$. This is because this additional randomness, which is correlated with the treatment assignment, will be useful for defining the estimator. Thus, this falls under the category of a *randomized experimental design*, as was discussed in Section 2.2 and defined in Section A.3.

The general idea of the design is that every exposure $e_{i,k}$ in the conflict graph is assigned some random variable $U_{i,k}$, which determines whether we “desire to be observe” this exposure. As dictated in Line 7, if we desire to observe the outcome under exposure $e_{i,k}$ (i.e. $U_{i,k} = e$), and its more important neighbors in the conflict graph are not desired to be observed (i.e. $U_{r,l} = *$ for all $e_{r,l} \in \mathcal{N}_b^\pi(e_{i,k})$), then we set the intervention Z in such a way that the exposure $e_{i,k}$ is indeed observed (Line 8).

The success of this design relies on the fact that we get to see the outcomes of all exposures with large enough probability. In turn, this depends on the importance ordering π that is chosen. Indeed, if an exposure has many more important neighbors, it is unlikely that it will ever be “allowed” to be observed. We will use the importance ordering produced by Algorithm 1. As we will see, the property established in Lemma C.1 for this ordering will be crucial for showing that every exposure gets observed with large enough probability.

C.1.3 Effect Estimator

As we alluded to earlier, our estimator will use the residual randomness of the variables $U_{i,k}$. To that end, for every node $e_{i,k} \in V(\mathcal{H})$, define the event $E_{i,k}$ as

$$E_{i,k} = \left\{ U_{i,k} = e \text{ and } U_{j,l} = * \text{ for all } e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k}) \right\} .$$

This is the event that we desire to observe exposure $e_{i,k}$ and all of its more important neighbors “allow” it to do so. We note that similar events were used in the definition of the estimator by Kandiros et al. (2024).

The first observation is that if the event $E_{i,k}$ occurs, then in the assignment Z that is produced by the design, the outcome of exposure $e_{i,k}$ will be observed for unit i . A similar fact was established as Lemma 3.2 in Kandiros et al. (2024). We give here the statement and proof for completeness.

Lemma C.3. *For every node $e_{i,k} \in V(\mathcal{H})$, if the event $E_{i,k}$ occurs, then if $Z \in \{0, 1\}^n$ is the assignment that is produced by the design of Algorithm 2, $h_i(Z) = e_{i,k}$, i.e. unit i receives the exposure $e_{i,k}$.*

Proof. Suppose event $E_{i,k}$ occurs. Then, by definition, the conditions in line 7 will be satisfied for node $e_{i,k}$, so that line 8 in Algorithm 2 will be executed for node $e_{i,k}$. This means that after this execution, the intervention vector Z will be such that $Z_j = \mathbf{1}[j \in h_i(e_{i,k})]$ for all $j \in \tilde{N}(i)$. Thus, after this execution vector Z satisfies $h_i(Z) = e_{i,k}$.

We now argue that the values of Z for $j \in \tilde{N}(i)$ will not change for the remainder of the algorithm. To see this, note that in order for Z to change, line 15 would have to be executed for some different node $e_{r,l} \in V(\mathcal{H})$ that comes after $e_{i,k}$ in the importance ordering. If the assignment of Z for some $j \in \tilde{N}(i)$ changes, this means that $j \in \tilde{N}(i) \cap \tilde{N}(r)$ and $\mathbf{1}[j \in h_i(e_{i,k})] \neq \mathbf{1}[j \in h_r(e_{r,l})]$. By the definition of the conflict graph $\mathcal{H}$, this means that $e_{i,k}$ and $e_{r,l}$ are neighbors in $\mathcal{H}$. In order for line 15 to be executed for node $e_{r,l}$, the conditions in line 14 must be met. In particular, we should have $U_{i,k} = *$, since $e_{i,k}$ is a more important neighbor of $e_{r,l}$ in $\mathcal{H}$. However, event $E_{i,k}$ implies that

$U_{i,k} = e$, so this is a contradiction. Therefore, the values of Z for $j \in \tilde{N}(i)$ will not change, therefore at the end of Algorithm 2 we will have $h_i(Z) = e_{i,k}$, as desired. $\square$

Lemma C.3 shows that if event $E_{i,k}$ occurs, then we get to observe the outcome for exposure $e_{i,k}$. Thus, a natural estimator for the effect τ would be to do inverse propensity weighting based on the events $E_{i,k}$. The estimator that arises was referred to as the *modified Horvitz-Thompson estimator* in Kandiros et al. (2024) and is defined as follows.

$$\hat{\tau} = \frac{1}{n} \sum_{i=1}^n \left(\frac{\mathbf{1}[E_{i,1}]}{\Pr(E_{i,1})} - \frac{\mathbf{1}[E_{i,0}]}{\Pr(E_{i,0})} \right) \cdot Y_i . \quad (13)$$

Note that $\hat{\tau}$ crucially depends on the residual randomness of the variables $U_{i,k}$, since the events $E_{i,k}$ are defined in terms of the $U_{i,k}$. Therefore, strictly speaking $\hat{\tau}$ is a randomized estimator.

As in Kandiros et al. (2024), let us first establish that $\hat{\tau}$ is unbiased for τ .

Lemma C.4. *For any $y \in \mathcal{M}(G)$, the estimator $\hat{\tau}$ is unbiased for τ , i.e. $\mathbb{E}[\hat{\tau}] = \tau$.*

Proof. By Lemma C.3 and the fact that $y \in \mathcal{M}(G)$, we have that if $E_{i,k}$ occurs, then $h_i(Z) = e_{i,k}$, so that

$$y_i(Z) = y_i(h_i(Z)) = y_i(e_{i,k}) .$$

Thus, we have

$$\mathbb{E}[\hat{\tau}] = \frac{1}{n} \left(\sum_{i=1}^n \frac{\Pr(E_{i,1})}{\Pr(E_{i,1})} y_i(e_{i,1}) - \sum_{i=1}^n \frac{\Pr(E_{i,0})}{\Pr(E_{i,0})} y_i(e_{i,0}) \right) = \frac{1}{n} \left(\sum_{i=1}^n y_i(e_{i,1}) - \sum_{i=1}^n y_i(e_{i,0}) \right) = \tau . \quad \square$$

C.1.4 Bounds on Correlation Terms

We now present bounds on the diagonal and off-diagonal terms of the covariance matrix of $\hat{\tau}$. Our proof will follow the same general structure as the proof for $q = 2$ in Kandiros et al., 2024, but with some important modifications. A key challenge in Kandiros et al., 2024 was to calculate the marginal and joint probabilities of the events $E_{(i,k)}$. A crucial part of the argument was a linearization step, which used the fact that the importance ordering is such that each i has at most $\lambda(\mathcal{H})$ neighbors that appear before i in the ordering (the so called *more important neighbors*). Here, we will need a generalization of this property, since the probabilities are non-uniform. In particular, we will use the crucial property of the ordering that was presented in Lemma C.1 to justify the linearization that will allow us to calculate the probabilities of the events $E_{i,k}$.

We start with the bound on the diagonal terms of the covariance matrix of $\hat{\tau}$, which amount to lower bounding the probability of the events $E_{i,k}$.

Lemma C.5. *Let $q \geq 2$ be a moment restriction. Suppose we sample the $U_{i,k}, Z$ variables according to the design in Algorithm 2 with importance ordering π given by Lemma C.1. For every node $e_{i,k} \in V(\mathcal{H})$ we have that if $C \geq 2$, then*

$$\text{Var}\left(\frac{\mathbf{1}[E_{i,k}]}{\Pr(E_{i,k})}\right) \leq \frac{Ce^{\frac{1}{C}}}{p_{i,k}} .$$

Proof. We start by noticing that

$$\text{Var}\left(\frac{\mathbf{1}[E_{i,k}]}{\Pr(E_{i,k})}\right) = \frac{1}{\Pr(E_{i,k})} - 1 .$$

By the definition of the design, we have

$$\begin{aligned} \Pr(E_{(i,k)}) &= \Pr(U_{i,k} = e, U_{j,l} = * \text{ for all } e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})) \\ &= \Pr(U_{i,k} = e) \prod_{e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})} \Pr(U_{j,l} = *) \\ &= \frac{p_{i,k}}{C} \prod_{e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})} \left(1 - \frac{p_{j,l}}{C}\right) , \end{aligned}$$

where we used the independence of the variables $U_{i,k}$. Now, consider the function $f(x) = 1 - x - e^{-2x}$ for $x \in [0, 1]$. We have that

$$f'(x) = -1 + 2e^{-2x} \quad , \quad f''(x) = -4e^{-2x}$$

Clearly, $f''(x) < 0$ for all $x \in [0, 1]$, so that f' is decreasing. Also, note that the unique root of f' is $\ln(2)/2 > 1/4$. Thus, as long as $x \leq 1/4$, we have $f'(x) > 0$, which implies that f is increasing in the interval $[0, 1/4]$. Since $f(0) = 0$, it follows that $f(x) \geq 0$ for all $x \in [0, 1/4]$. In particular, we have $1 - x \geq e^{-2x}$ for all $x \in [0, 1/4]$. Now, apply this inequality with $x = p_j/C$. By Lemma C.1, we have for all i

$$\sum_{j \in \mathcal{N}_b^\pi(i)} p_j \leq \frac{1}{2}$$

This implies that $p_{i,k}/C \leq 1/4$ for all i as long as $C \geq 2$. Therefore, we can apply the inequality $1 - x \geq e^{-2x}$ to each term in the product, yielding

$$\prod_{e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})} \left(1 - \frac{p_{j,l}}{C}\right) \geq \prod_{e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})} e^{-2p_{j,l}/C} = e^{-2 \sum_{e_{j,l} \in \mathcal{N}_b^\pi(e_{i,k})} p_{j,l}/C} \geq e^{-1/C} ,$$

from which the Lemma follows. $\square$

Next, we provide bounds on the covariance between two terms corresponding to different exposures. These involve bounding the joint probabilities of events $E_{i,k}$ and $E_{j,l}$.

Lemma C.6. *Let $q \geq 2$ be a moment restriction. Suppose we sample the $U_{i,k}, Z$ variables according to the design in Algorithm 2 with importance ordering π given by Lemma C.1 and $C \geq 1$. For every pair of nodes $e_{i,k}, e_{j,l} \in V(\mathcal{H})$, if $d(e_{i,k}, e_{j,l})$ denotes the graph distance between $e_{i,k}$ and $e_{j,l}$ in $\mathcal{H}$, then we have the following:*

- If $d(e_{i,k}, e_{j,l}) = 1$, then

$$\text{Cov}\left(\frac{\mathbf{1}[E_{i,k}]}{\Pr(E_{i,k})}, \frac{\mathbf{1}[E_{j,l}]}{\Pr(E_{j,l})}\right) = -1 .$$

- If $d(e_{i,k}, e_{j,l}) = 2$, then

$$\left| \text{Cov}\left(\frac{\mathbf{1}[E_{i,k}]}{\Pr(E_{i,k})}, \frac{\mathbf{1}[E_{j,l}]}{\Pr(E_{j,l})}\right) \right| \leq \frac{2}{2C-1} \cdot \sum_{e_{r,s} \in N(e_{i,k}) \cap N(e_{j,l})} p_{r,s} .$$

- If $d(e_{i,k}, e_{j,l}) > 2$, then

$$\text{Cov}\left(\frac{\mathbf{1}[E_{i,k}]}{\Pr(E_{i,k})}, \frac{\mathbf{1}[E_{j,l}]}{\Pr(E_{j,l})}\right) = 0 .$$

Proof. By definition, the event $E_{i,k}$ depends on the values of the variables $U_{i,k}$ and $U_{j,l}$ for all $j \in \mathcal{N}(e_{i,k})$. Thus, if $d(e_{i,k}, e_{j,l}) > 2$, then the events $E_{i,k}$ and $E_{j,l}$ depend on disjoint sets of variables, so they are independent, which implies that the covariance is zero. If $d(e_{i,k}, e_{j,l}) = 1$, then the events $E_{i,k}$ and $E_{j,l}$ are mutually exclusive. Indeed, suppose without loss of generality that $e_{i,k}$ appears before $e_{j,l}$ in the ordering, then if $E_{i,k}$ happens we must have $U_{i,k} = e$, but if $E_{j,l}$ also happens, then we must have $U_{i,k} = *$, which is a contradiction. Thus,

$$\text{Cov}\left(\frac{\mathbf{1}[E_{i,k}]}{\Pr(E_{i,k})}, \frac{\mathbf{1}[E_{j,l}]}{\Pr(E_{j,l})}\right) = \mathbb{E}\left[\frac{\mathbf{1}[E_{i,k}]}{\Pr(E_{i,k})} \frac{\mathbf{1}[E_{j,l}]}{\Pr(E_{j,l})}\right] - 1 = -1 .$$

Finally, let us examine the case where $d(e_{i,k}, e_{j,l}) = 2$. In this case, we can write

$$\begin{aligned} & \text{Cov}\left(\frac{\mathbf{1}[E_{i,k}]}{\Pr(E_{i,k})}, \frac{\mathbf{1}[E_{j,l}]}{\Pr(E_{j,l})}\right) \\ &= \mathbb{E}\left[\frac{\mathbf{1}[E_{i,k}]}{\Pr(E_{i,k})} \frac{\mathbf{1}[E_{j,l}]}{\Pr(E_{j,l})}\right] - 1 \\ &= \frac{\Pr(E_{i,k} \cap E_{j,l})}{\Pr(E_{i,k}) \Pr(E_{j,l})} - 1 \\ &= \frac{\Pr(U_{i,k} = U_{j,l} = e_1, U_{r,s} = *, \forall e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cup \mathcal{N}_b^\pi(e_{j,l}))}{\Pr(U_{i,k} = e_1, U_{r,s} = *, \forall e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k})) \Pr(U_{j,l} = e_1, U_{r,s} = *, \forall e_{r,s} \in \mathcal{N}_b^\pi(e_{j,l}))} - 1 \\ &= \frac{1}{\prod_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cap \mathcal{W}_b^\pi(e_{j,l})} \left(1 - \frac{p_{r,s}}{C}\right)} - 1 \end{aligned}$$

By the inclusion-exclusion principle, we have that

$$\prod_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cap \mathcal{W}_b^\pi(e_{j,l})} \left(1 - \frac{p_{r,s}}{C}\right) \geq 1 - \sum_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cap \mathcal{W}_b^\pi(e_{j,l})} \frac{p_{r,s}}{C} .$$

Hence, we have that

$$\begin{aligned} \frac{1}{\prod_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cap \mathcal{W}_b^\pi(e_{j,l})} \left(1 - \frac{p_{r,s}}{C}\right)} - 1 &= \frac{1 - \prod_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cap \mathcal{W}_b^\pi(e_{j,l})} \left(1 - \frac{p_{r,s}}{C}\right)}{\prod_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cap \mathcal{W}_b^\pi(e_{j,l})} \left(1 - \frac{p_{r,s}}{C}\right)} \\ &\leq \frac{\sum_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cap \mathcal{W}_b^\pi(e_{j,l})} \frac{p_{r,s}}{C}}{1 - \sum_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cap \mathcal{W}_b^\pi(e_{j,l})} \frac{p_{r,s}}{C}} \\ &\leq \frac{2C}{2C - 1} \sum_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cap \mathcal{W}_b^\pi(e_{j,l})} \frac{p_{r,s}}{C} , \end{aligned}$$

where in the last step we used the fact that

$$\sum_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cap \mathcal{W}_b^\pi(e_{j,l})} \frac{p_{r,s}}{C} \leq \sum_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k})} \frac{p_{r,s}}{C} \leq \frac{1}{2C} ,$$

which follows from Lemma C.1. Finally, we trivially have

$$\sum_{e_{r,s} \in \mathcal{N}_b^\pi(e_{i,k}) \cap \mathcal{N}_b^\pi(e_{j,l})} p_{r,s} \leq \sum_{e_{r,s} \in \mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})} p_{r,s} ,$$

which concludes the proof. $\square$

C.1.5 Matrix Norm Bounds

Lemmas C.5 and C.6 suggest that the covariance matrix associated with the estimator $\hat{\tau}$ can be split into three matrices, one diagonal matrix corresponding to $e_{i,k} = e_{j,l}$, one corresponding to pairs $e_{i,k}, e_{j,l}$ with $d(e_{i,k}, e_{j,l}) = 1$ and one for $d(e_{i,k}, e_{j,l}) = 2$. A similar strategy was employed in Kandiros et al. (2024). In particular, we define for all pairs $e_{i,k}, e_{j,l} \in V(\mathcal{H})$

$$C_{(i,k),(j,l)} = \text{Cov}\left(\frac{\mathbf{1}[E_{i,k}]}{\Pr(E_{i,k})}, \frac{\mathbf{1}[E_{j,l}]}{\Pr(E_{j,l})}\right) .$$

Then, let us define the matrices $\mathbf{T}^{(t)} \in \mathbb{R}^{2n \times 2n}$ for $t \in \{0, 1, 2\}$ as follows:

$$\mathbf{T}_{(i,k),(j,l)}^{(t)} = \begin{cases} |C_{(i,k),(j,l)}| & \text{if } d(e_{i,k}, e_{j,l}) = t \\ 0 & \text{otherwise} \end{cases} .$$

The following Lemmas upper bound the quadratic forms of the matrices $\mathbf{T}^{(0)}, \mathbf{T}^{(1)}, \mathbf{T}^{(2)}$, evaluated at vectors whose q -th norm is at most 1. These bounds will be crucial for bounding the variance of $\hat{\tau}$. The proofs are based on careful applications of Hölder's inequality. For a vector $x \in \mathbb{R}^{2n}$, we will sometimes use the double indexing $x_{i,k}$ to refer to the coordinate of x corresponding to node $e_{i,k} \in V(\mathcal{H})$.

Lemma C.7. *For any vector $x \in \mathbb{R}^{2n}$ with $\|x\|_q \leq 1$, we have*

$$x^\top \mathbf{T}^{(0)} x \leq 2C e^{\frac{1}{C}} \cdot n^{1-\frac{2}{q}} \lambda(\mathcal{H})^{\frac{2}{q}} d_{\text{avg}}(\mathcal{H})^{1-\frac{2}{q}} .$$

Proof. Let m denote the number of edges in the conflict graph $\mathcal{H}$. Using Hölder's inequality, we have that

$$x^\top \mathbf{T}^{(0)} x \leq C e^{\frac{1}{C}} \cdot \sum_{e_{i,k} \in V(\mathcal{H})} x_{i,k}^2 \frac{1}{p_{i,k}}$$

using Lemma C.5

$$\begin{aligned} &\leq C e^{\frac{1}{C}} \cdot \left(\sum_{e_{i,k} \in V(\mathcal{H})} x_{i,k}^q \right)^{2/q} \left(\sum_{e_{i,k} \in V(\mathcal{H})} \left(\frac{1}{p_{i,k}} \right)^{\frac{q}{q-2}} \right)^{(q-2)/q} \\ &\leq C e^{\frac{1}{C}} \cdot \left(\sum_{e_{i,k} \in V(\mathcal{H})} \left(\frac{1}{p_{i,k}} \right)^{\frac{q}{q-2}} \right)^{(q-2)/q} \end{aligned}$$

by the assumption on x

$$\begin{aligned}
&= Ce^{\frac{1}{C}} \cdot \left(\sum_{e_{i,k} \in V(\mathcal{H})} (2\lambda(\mathcal{H})^{2/q} d_{i,k}^{1-2/q})^{\frac{q}{q-2}} \right)^{(q-2)/q} \\
&= 2Ce^{\frac{1}{C}} \cdot \lambda(\mathcal{H})^{2/q} \left(\sum_{e_{i,k} \in V(\mathcal{H})} d_{i,k} \right)^{(q-2)/q} \\
&= 2Ce^{\frac{1}{C}} \cdot \lambda(\mathcal{H})^{2/q} (2m)^{1-2/q} \\
&= 2Ce^{\frac{1}{C}} \cdot n^{1-\frac{2}{q}} \lambda(\mathcal{H})^{\frac{2}{q}} d_{\text{avg}}(\mathcal{H})^{1-\frac{2}{q}}.
\end{aligned}$$

□

Lemma C.8. For any vector $x \in \mathbb{R}^{2n}$ with $\|x\|_q \leq 1$, we have

$$x^\top \mathbf{T}^{(1)} x \leq n^{1-\frac{2}{q}} \lambda(\mathcal{H})^{\frac{2}{q}} d_{\text{avg}}(\mathcal{H})^{1-\frac{2}{q}}.$$

Proof. We have that

$$\begin{aligned}
x^\top \mathbf{T}^{(1)} x &= \sum_{e_{i,k} \in V_{\mathcal{H}}} x_{i,k} \sum_{e_{j,l} \in V_{\mathcal{H}}} \mathbf{T}_{(i,k),(j,l)}^{(1)} x_{j,l} \\
&\leq \sum_{e_{i,k} \in V_{\mathcal{H}}} |x_{i,k}| \sum_{e_{j,l} \in V_{\mathcal{H}}} |\mathbf{T}_{(i,k),(j,l)}^{(1)}| |x_{j,l}| \\
&\leq \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} |x_{i,k}|^q \right)^{1/q} \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |\mathbf{T}_{(i,k),(j,l)}^{(1)}| |x_{j,l}| \right)^{\frac{q}{q-1}} \right)^{(q-1)/q}
\end{aligned}$$

by Hölder's inequality

$$\leq \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |\mathbf{T}_{(i,k),(j,l)}^{(1)}| |x_{j,l}| \right)^{\frac{q}{q-1}} \right)^{(q-1)/q}$$

by the assumption on x

$$\leq \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |\mathbf{T}_{(i,k),(j,l)}^{(1)}| |x_{j,l}|^{q/2} \right)^{2/q} \cdot d_{i,k}^{1-2/q} \right)^{\frac{q}{q-1}} \right)^{\frac{q-1}{q}}$$

by Hölder's inequality

$$= \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |\mathbf{T}_{(i,k),(j,l)}^{(1)}| |x_{j,l}|^{q/2} \right)^{\frac{2}{q-1}} \cdot d_{i,k}^{\frac{q-2}{q-1}} \right) \right)^{\frac{q-1}{q}}$$

$$\leq \left(\left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |\mathbf{T}_{(i,k),(j,l)}^{(1)}| |x_{j,l}|^{q/2} \right)^2 \right)^{\frac{1}{q-1}} \cdot \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} (d_{i,k}^{\frac{q-2}{q-1}})^{\frac{q-1}{q-2}} \right)^{\frac{q-2}{q-1}} \right)^{\frac{q-1}{q}}$$

by Hölder's inequality

$$\begin{aligned} &= \left(\left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |\mathbf{T}_{(i,k),(j,l)}^{(1)}| |x_{j,l}|^{q/2} \right)^2 \right)^{\frac{1}{q-1}} \cdot \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} d_{i,k} \right)^{\frac{q-2}{q-1}} \right)^{\frac{q-1}{q}} \\ &= \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |\mathbf{T}_{(i,k),(j,l)}^{(1)}| |x_{j,l}|^{q/2} \right)^2 \right)^{\frac{1}{q}} \cdot (2m)^{\frac{q-2}{q}} \\ &\leq \left(\lambda(\mathcal{H})^2 \sum_{e_{i,k} \in V_{\mathcal{H}}} (|x_{j,l}|^{q/2})^2 \right)^{\frac{1}{q}} \cdot (2m)^{\frac{q-2}{q}} \end{aligned}$$

since $|\mathbf{T}_{(i,k),(j,l)}^{(1)}|$ is the adjacency entry of graph $\mathcal{H}$ by Lemma C.6 and using the definition of the maximum eigenvalue property

$$\begin{aligned} &\leq \lambda^{\frac{2}{q}} (2m)^{1-\frac{2}{q}} \\ &= n^{1-\frac{2}{q}} \lambda(\mathcal{H})^{\frac{2}{q}} d_{\text{avg}}(\mathcal{H})^{1-\frac{2}{q}}. \end{aligned}$$

□

Now, let us handle the third matrix.

Lemma C.9. *For any vector $x \in \mathbb{R}^{2n}$ with $\|x\|_q \leq 1$, we have*

$$x^\top \mathbf{T}^{(2)} x \leq \frac{4}{2C-1} \cdot n^{1-\frac{2}{q}} \lambda(\mathcal{H})^{\frac{2}{q}} d_{\text{avg}}(\mathcal{H})^{1-\frac{2}{q}}.$$

Proof. Let $A_{\mathcal{H}}$ be the adjacency matrix of the conflict graph $\mathcal{H}$. We have that

$$\begin{aligned} x^\top \mathbf{T}^{(2)} x &= \sum_{e_{i,k} \in V_{\mathcal{H}}} x_{i,k} \sum_{e_{j,l} \in V_{\mathcal{H}}} \mathbf{T}_{(i,k),(j,l)}^{(2)} x_{j,l} \\ &\leq \sum_{e_{i,k} \in V_{\mathcal{H}}} |x_{i,k}| \sum_{e_{j,l} \in V_{\mathcal{H}}} |\mathbf{T}_{(i,k),(j,l)}^{(2)}| |x_{j,l}| \\ &\leq \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} |x_{i,k}|^q \right)^{1/q} \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |\mathbf{T}_{(i,k),(j,l)}^{(2)}| |x_{j,l}| \right)^{\frac{q}{q-1}} \right)^{\frac{q-1}{q}} \end{aligned}$$

by Hölder's inequality

$$\leq \frac{2}{2C-1} \frac{2}{\lambda(\mathcal{H})^{2/q}} \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |x_{j,l}| \sum_{e_{r,s} \in \mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})} \frac{1}{d_{r,s}^{\frac{q}{q-2}}} \right)^{\frac{q}{q-1}} \right)^{\frac{q-1}{q}}$$

using the bound on $C_{(i,k),(j,l)}$ of Lemma C.6

$$\leq \frac{2}{2C-1} \frac{2}{\lambda(\mathcal{H})^{2/q}} \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |x_{j,l}| |\mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})|^{\frac{2}{q}} \left(\sum_{e_{r,s} \in \mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})} \frac{1}{d_{r,s}} \right)^{\frac{q-2}{q}} \right)^{\frac{q}{q-1}} \right)^{\frac{q-1}{q}}$$

again using Hölder's inequality

$$\leq \frac{2}{2C-1} \frac{2}{\lambda(\mathcal{H})^{2/q}} \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |x_{j,l}|^{\frac{q}{2}} |\mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})| \right)^{\frac{2}{q-1}} \right. \\ \left. \cdot \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} \sum_{e_{r,s} \in \mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})} \frac{1}{d_{r,s}} \right)^{\frac{q-2}{q-1} \frac{q-1}{q}} \right)$$

by another application of Hölder's inequality

$$\leq \frac{2}{2C-1} \frac{2}{\lambda(\mathcal{H})^{2/q}} \left(\left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |x_{j,l}|^{\frac{q}{2}} |\mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})| \right) \right)^2 \right)^{\frac{1}{q-1}} \\ \cdot \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \sum_{e_{j,l} \in V_{\mathcal{H}}} \sum_{e_{r,s} \in \mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})} \frac{1}{d_{r,s}} \right)^{\frac{q-2}{q-1} \frac{q-1}{q}}$$

by yet another application of Hölder's inequality

$$= \frac{4}{(2C-1)\lambda(\mathcal{H})^{2/q}} \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \left(\sum_{e_{j,l} \in V_{\mathcal{H}}} |x_{j,l}|^{\frac{q}{2}} |\mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})| \right) \right)^2 \right)^{\frac{1}{q}} \\ \cdot \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \sum_{e_{j,l} \in V_{\mathcal{H}}} \sum_{e_{r,s} \in \mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})} \frac{1}{d_{r,s}} \right)^{\frac{q-2}{q}} \\ \leq \frac{4}{(2C-1)\lambda(\mathcal{H})^{2/q}} \left(\lambda(\mathcal{H})^2 \sum_{e_{i,k} \in V_{\mathcal{H}}} |x_{i,k}|^q \right)^{\frac{1}{q}} \cdot \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \sum_{e_{j,l} \in V_{\mathcal{H}}} \sum_{e_{r,s} \in \mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})} \frac{1}{d_{r,s}} \right)^{\frac{q-2}{q}}$$

using the definition of the maximum eigenvalue of $A_{\mathcal{H}}$

$$\leq \frac{4}{(2C-1)\lambda(\mathcal{H})^{2/q}} \lambda(\mathcal{H})^{\frac{4}{q}} \left(\sum_{e_{i,k} \in V_{\mathcal{H}}} \sum_{e_{j,l} \in V_{\mathcal{H}}} \sum_{e_{r,s} \in \mathcal{N}(e_{i,k}) \cap \mathcal{N}(e_{j,l})} \frac{1}{d_{r,s}} \right)^{\frac{q-2}{q}}$$

using the constraint for the vector x

$$\begin{aligned}
&= \frac{4}{2C-1} \lambda(\mathcal{H})^{2/q} \left(\sum_{e_{r,s} \in V_{\mathcal{H}}} \frac{1}{d_{r,s}} \sum_{e_{i,k} \in N(e_{r,s})} \sum_{e_{j,l} \in N(e_{r,s})} 1 \right)^{\frac{q-2}{q}} \\
&= \frac{4}{2C-1} \lambda(\mathcal{H})^{2/q} \left(\sum_{e_{r,s} \in V_{\mathcal{H}}} \frac{d_{r,s}^2}{d_{r,s}} \right)^{\frac{q-2}{q}} \\
&= \frac{4}{2C-1} \lambda(\mathcal{H})^{2/q} (2m)^{1-2/q} \\
&= \frac{4}{2C-1} n^{1-\frac{2}{q}} \lambda(\mathcal{H})^{\frac{2}{q}} d_{\text{avg}}(\mathcal{H})^{1-\frac{2}{q}} .
\end{aligned}$$

□

C.1.6 Proof of Theorem 4.1

We are now ready to prove Theorem 4.1, which is a bound on the variance of $\hat{\tau}$.

Proof of Theorem 4.1. To prove the upper bound, it suffices to produce a statistical procedure, whose risk is bounded by the appropriate quantity. We will use as our design $\mathcal{D}$ the Conflict Graph Design of Algorithm 2 with an appropriate value of $C > 0$ that we will specify later and as our estimator $\hat{\tau}$, the modified Horvitz–Thompson estimator (13). Although the Conflict Graph Design and modified Horvitz–Thompson estimator form a randomized statistical procedure (as defined in Section A.3), its maximal risk provides an upper bound for the minimax risk (Proposition A.2, Section A.3).

Our goal is to bound

$$\sup_{y \in \mathcal{M}(G,q)} \mathbb{E}_{(Z,U) \sim \mathcal{D}} [(\hat{\tau}(Z,U,y(Z)) - \tau(y))^2] = \sup_{y \in \mathcal{M}(G,q)} \text{Var}(\hat{\tau}(Z,U,y(Z))) .$$

Let $y \in \mathcal{M}(G,q)$. Let us define the vector $\alpha \in \mathbb{R}^{2n}$ as

$$\alpha_{i,k} = \frac{1}{n^{1/q}} |y_i(e_{i,k})| .$$

We begin by writing the variance of $\hat{\tau}$ as a quadratic form of the covariance matrix with vector α , as in Kandiros et al., 2024. In particular, we have

$$\begin{aligned}
\text{Var}(\hat{\tau}) &= \text{Var} \left(\frac{1}{n} \left(\sum_{i=1}^n \left(\frac{\mathbf{1}[E_{i,1}]}{\Pr(E_{i,1})} y_i(e_{i,1}) - \frac{\mathbf{1}[E_{i,0}]}{\Pr(E_{i,0})} y_i(e_{i,0}) \right) \right) \right) \\
&= \frac{1}{n^2} \sum_{e_{i,k} \in V_{\mathcal{H}}} \sum_{e_{j,l} \in V_{\mathcal{H}}} y_i(e_{i,k}) y_j(e_{j,l}) (-1)^{i+j} \text{Cov} \left(\frac{\mathbf{1}[E_{i,k}]}{\Pr(E_{i,k})}, \frac{\mathbf{1}[E_{j,l}]}{\Pr(E_{j,l})} \right) \\
&\leq \frac{1}{n^2} \sum_{e_{i,k} \in V_{\mathcal{H}}} \sum_{e_{j,l} \in V_{\mathcal{H}}} |y_i(e_{i,k})| |y_j(e_{j,l})| |C_{(i,k),(j,l)}|
\end{aligned}$$

$$= \frac{n^{\frac{2}{q}}}{n^2} \boldsymbol{\alpha}^\top (\mathbf{T}^{(0)} + \mathbf{T}^{(1)} + \mathbf{T}^{(2)}) \boldsymbol{\alpha}$$

Since $y \in \mathcal{M}(G, q)$, the q -th norm of the vector $\boldsymbol{\alpha}$ is bounded as

$$\left(\sum_{i=1}^n \sum_{k=0}^1 |\alpha_{i,k}|^q \right)^{1/q} = \left(\frac{1}{n} \sum_{i=1}^n |y_i(e_{i,0})|^q + \frac{1}{n} \sum_{i=1}^n |y_i(e_{i,1})|^q \right)^{1/q} \leq 1 ,$$

using the definition of $\mathcal{M}(G, q)$. Now, applying Lemmas C.7, C.8 and C.9 to bound the three terms in the sum, we have

$$\begin{aligned} \text{Var}(\hat{\tau}) &\leq \frac{n^{\frac{2}{q}}}{n^2} \left(2Ce^{\frac{1}{C}} + 1 + \frac{4}{2C-1} \right) \lambda(\mathcal{H})^{\frac{2}{q}} (2m)^{1-\frac{2}{q}} \\ &= \left(2Ce^{\frac{1}{C}} + 1 + \frac{4}{2C-1} \right) \cdot \frac{\lambda(\mathcal{H})^{\frac{2}{q}} (2m/n)^{1-\frac{2}{q}}}{n} \\ &= \left(2Ce^{\frac{1}{C}} + 1 + \frac{4}{2C-1} \right) \cdot \frac{\lambda(\mathcal{H})^{\frac{2}{q}} d_{\text{avg}}^{1-\frac{2}{q}}}{n} , \end{aligned}$$

where in the last step we used the fact that $d_{\text{avg}} = 2m/n$. Choosing $C = 2$ yields the desired bound. $\square$

C.2 Bias-Variance Tradeoff

In this section, we first establish in Section C.2.1 the improved upper bound of Theorem 4.2. We then show in Section C.3 that we can find a design-estimator pair that can be implemented in polynomial time, whose risk is within multiplicative logarithmic factors from this improved bound.

C.2.1 Optimal Bias-Variance Tradeoff (Theorem 4.2)

We first introduce some notation that will be particularly useful in this section.

For any subset of nodes $S \subseteq V_{\mathcal{H}}$, we denote by $\mathcal{H}[V_{\mathcal{H}} \setminus S] = (V_{\mathcal{H}} \setminus S, E(V_{\mathcal{H}} \setminus S))$ the subgraph of $\mathcal{H}$ that is *induced* by the nodes in $V_{\mathcal{H}} \setminus S$, i.e. for $i, j \in V_{\mathcal{H}} \setminus S$, $(i, j) \in E(V_{\mathcal{H}} \setminus S)$ if and only if $(i, j) \in E_{\mathcal{H}}$. We denote by $\lambda(\mathcal{H}[V_{\mathcal{H}} \setminus S])$ the largest eigenvalue of the adjacency matrix of $\mathcal{H}[V_{\mathcal{H}} \setminus S]$.

As usual, we will use the notation $\tilde{N}(i)$ for the extended neighborhood of a node i in the original graph G . Also, for convenience, we introduce the following function of a subset $S \subseteq V(\mathcal{H})$ and a moment parameter q .

$$Q_q(S, \mathcal{H}) = \lambda(\mathcal{H}[V_{\mathcal{H}} \setminus S])^{\frac{2}{q}} \cdot d_{\text{avg}}(\mathcal{H}[V_{\mathcal{H}} \setminus S])^{1-\frac{2}{q}} + \frac{|S|^{2-\frac{2}{q}}}{n^{1-\frac{2}{q}}} .$$

In this section, we will prove the following theorem.

Theorem 4.2. *Let G be an interference network, τ be a contrastive effect. For any moment parameter $q \geq 2$, the corresponding minimax risk is bounded as*

$$\mathcal{R}(G, \tau, q) \leq 8.92 \cdot \frac{\min_{S \subseteq V_{\mathcal{H}}} Q_q(S, \mathcal{H})}{n} .$$

Algorithm 3: Biased Conflict Graph Design

Input: maximum eigenvalue $\lambda(\mathcal{H})$, graph G , conflict graph $\mathcal{H}$ and moment q .

Output: Random intervention $Z \in \mathcal{Z} = \{0, 1\}^n$, variables U_i for all $i \in V(\mathcal{H})$.

- 1 $S^* \leftarrow \arg \min_{S \subseteq V(\mathcal{H})} Q_q(S, \mathcal{H})$
- 2 $\mathcal{H}^* \leftarrow \mathcal{H}[V_{\mathcal{H}} \setminus S^*]$
- 3 $\pi^* \leftarrow$ Ordering on $V_{\mathcal{H}} \setminus S^*$ produced by Algorithm 1 with input $\mathcal{H}^*$ and q
- 4 **for** $e_{i,k} \in V(\mathcal{H}^*)$ **do**
- 5 $p_{i,k} \leftarrow \frac{1}{2\lambda(\mathcal{H}^*)^{\frac{2}{q}} d_{i,k}(\mathcal{H}^*)^{1-\frac{2}{q}}};$
- 6 Sample desired exposure variables $\{U_{i,k} : e_{i,k} \in V(\mathcal{H}) \setminus S^*\}$ independently and identically as

$$U_{i,k} \leftarrow \begin{cases} e & \text{with probability } \frac{p_{i,k}}{2} \\ * & \text{with probability } 1 - \frac{p_{i,k}}{2} \end{cases}$$

- 7 Initialize the intervention vector $Z \leftarrow \mathbf{0}$
 - 8 **for** $e_{i,k} \in V(\mathcal{H}) \setminus S^*$ **do**
 - 9 **if** $U_{i,k} = e$ **and** $U_{j,l} = *$ **for all** $e_{j,l} \in \mathcal{N}_b^{\pi^*}(e_{i,k})$ **then**
 - 10 Update intervention vector: set $Z_r = \mathbf{1}[r \in h_i(e_{i,k})]$ for all $r \in \tilde{N}(i)$;
-

We start by introducing the design that will be used in the proof of Theorem 4.2. We give the pseudocode in Algorithm 3.

Algorithm 3 operates similarly to the CGD of Algorithm 2, but with one important difference. In particular, in line 1 we select a particular subset of nodes S^* in the conflict graph $\mathcal{H}$ whose outcomes we will not try to estimate. Then, an ordering π^* of the nodes in $V_{\mathcal{H}} \setminus S^*$ is produced by calling the importance ordering procedure from Algorithm 1 on the subgraph $\mathcal{H}[V_{\mathcal{H}} \setminus S^*]$ of the original conflict graph. The remaining steps mirror the CGD of Algorithm 2, where we iterate only over the exposures in $V_{\mathcal{H}} \setminus S^*$ and assign treatment.

In order to describe our proposed estimator, we first introduce some definitions for arbitrary subsets of nodes in the conflict graph. In particular, for any subset $S \subseteq V_{\mathcal{H}}$, let π_S be the ordering of the nodes in $V_{\mathcal{H}} \setminus S$ produced by the importance ordering procedure from Algorithm 1 with input $\mathcal{H}[V_{\mathcal{H}} \setminus S]$ and q (we omit the dependence on q since it is clear from the context). For all $e_{i,k} \in V_{\mathcal{H}} \setminus S$, define the events

$$E_{i,k}^{(S)} = \{U_{i,k} = e, U_{j,l} = * \text{ for all } e_{j,l} \in \mathcal{N}_b^{\pi_S}(e_{i,k})\} .$$

Note that we used a similar notation for an analogous event in the original CGD in Section C.1.2, where the event depends on all neighbors of $e_{i,k}$ in the conflict graph. This could be thought of as a special case of the above definition, where $S = \emptyset$ and $\pi_{\emptyset}$ is the ordering produced by Algorithm 1 on the entire conflict graph $\mathcal{H}$.

For any subset $S \subseteq V_{\mathcal{H}}$, we define the following estimator.

$$\hat{\tau}_S = \frac{1}{n} \sum_{i=1}^n \left(\frac{\mathbf{1}[E_{i,1}^{(S)}]}{\Pr(E_{i,1}^{(S)})} \mathbf{1}[e_{i,1} \notin S] - \frac{\mathbf{1}[E_{i,0}^{(S)}]}{\Pr(E_{i,0}^{(S)})} \mathbf{1}[e_{i,0} \notin S] \right) y_i(Z) . \quad (14)$$

All probabilities are calculated with respect to the randomness in the design of Algorithm 3. Note that $\hat{\tau}_S$ is a randomized estimator, in the sense defined in Section 2, since it uses the values of the random variables $U_{i,k}$ in its definition.

For an arbitrary subset of exposures $S \subseteq V_{\mathcal{H}}$, denote by τ_S the average treatment effect defined only on exposures in $V_{\mathcal{H}} \setminus S$, i.e.

$$\tau_S = \frac{1}{n} \sum_{i=1}^n (\mathbf{1}[e_{i,1} \notin S] y_i(e_{i,1}) - \mathbf{1}[e_{i,0} \notin S] y_i(e_{i,0})) .$$

We will use $\hat{\tau}_{S^*}$ as our estimator for τ_{S^*} , where S^* is the subset of nodes in the conflict graph that minimizes $Q(S, q)$, as defined in line 1 of Algorithm 3. Our first proposition establishes that $\hat{\tau}_{S^*}$ is an unbiased estimator for τ_{S^*} .

Proposition C.10. *For any outcome function $y \in \mathcal{M}(G)$, we have $E[\hat{\tau}_{S^*}] = \tau_{S^*}$.*

Proof. First, observe that if event $E_{i,k}^{(S^*)}$ holds, then $y_i(Z) = y_i(e_{i,k})$. The argument is essentially identical to that of Lemma C.3 for the original CGD, since Algorithm 3 essentially runs the conflict graph design on the set of nodes $V_{\mathcal{H}} \setminus S^*$. Therefore, we have

$$\begin{aligned} E[\hat{\tau}_{S^*}] &= \frac{1}{n} \sum_{i=1}^n \left(\frac{\Pr(E_{i,1}^{(S^*)})}{\Pr(E_{i,1}^{(S^*)})} \mathbf{1}[e_{i,1} \notin S^*] y_i(e_{i,1}) - \frac{\Pr(E_{i,0}^{(S^*)})}{\Pr(E_{i,0}^{(S^*)})} \mathbf{1}[e_{i,0} \notin S^*] y_i(e_{i,0}) \right) \\ &= \frac{1}{n} \sum_{i=1}^n \left(\mathbf{1}[e_{i,1} \notin S^*] y_i(e_{i,1}) - \mathbf{1}[e_{i,0} \notin S^*] y_i(e_{i,0}) \right) \\ &= \tau_{S^*} . \end{aligned} \quad \square$$

We are now ready to prove Theorem 4.2.

Proof of Theorem 4.2. We will prove the upper bound on the minimax risk by bounding the risk of a particular design-estimator pair. We will use the Biased Conflict Graph Design of Algorithm 3 as the design $\mathcal{D}$ and $\hat{\tau}_{S^*}$ as the estimator, as defined in equation (14). As observed above, $\mathcal{D}$ is a randomized experimental design and $\hat{\tau}_{S^*}$ is a randomized estimator, since they both use the extra randomness of the variables $U_{i,k}$. By Proposition A.2, it suffices to bound the risk of this design-estimator pair to obtain an upper bound on the minimax risk.

By Proposition C.10, we have that $\hat{\tau}_{S^*}$ is an unbiased estimator for τ_{S^*} . Since τ_{S^*} is not necessarily equal to τ , the estimator $\hat{\tau}_{S^*}$ may be biased for τ . Thus, we can decompose the MSE of $\hat{\tau}_{S^*}$ as a sum of bias and variance terms, as follows.

$$E[(\hat{\tau}_{S^*} - \tau)^2] = \text{Var}(\hat{\tau}_{S^*}) + (\tau_{S^*} - \tau)^2 . \quad (15)$$

We will now bound each term separately, starting with the bias term. For any $y \in \mathcal{M}(G, q)$, we have

$$\begin{aligned} (\tau_{S^*} - \tau)^2 &= \left(\frac{1}{n} \sum_{i=1}^n (\mathbf{1}[e_{i,1} \in S^*] y_i(e_{i,1}) - \mathbf{1}[e_{i,0} \in S^*] y_i(e_{i,0})) \right)^2 \\ &\leq \frac{1}{n^2} \left(\sum_{e_{i,k} \in S^*} |y_i(e_{i,k})| \right)^2 \end{aligned}$$

using triangle inequality

$$\leq \frac{1}{n^2} \left(\sum_{e_{i,k} \in S^*} |y_i(e_{i,k})|^q \right)^{\frac{2}{q}} \cdot \left(\sum_{e_{i,k} \in S^*} 1 \right)^{\frac{2(q-1)}{q}}$$

by Hölder's inequality

$$\leq \frac{|S^*|^{2-\frac{2}{q}}}{n^{2-\frac{2}{q}}}, \quad (16)$$

using the fact that $y \in \mathcal{M}(G, q)$ in the last step.

We now focus on the variance term. As we've already discussed, after selecting the subset S^* , Algorithm 3 proceeds identically to the CGD of Algorithm 2 but only on the subgraph $\mathcal{H}[V_{\mathcal{H}} \setminus S^*]$. Moreover, the estimator $\hat{\tau}_{S^*}$ is essentially the exact analogue of the estimator $\hat{\tau}$ defined in equation (13) for the original CGD, but only on the exposures in $V_{\mathcal{H}} \setminus S^*$. Therefore, we can use the analysis of the variance of the original CGD, as given in the proof of Theorem 4.1, to conclude that for any $y \in \mathcal{M}(G, q)$, we have

$$\text{Var}(\hat{\tau}_{S^*}) \leq 8.92 \cdot \frac{\lambda(\mathcal{H}(V_{\mathcal{H}} \setminus S^*))^{\frac{2}{q}} \cdot d_{\text{avg}}(\mathcal{H}[V_{\mathcal{H}} \setminus S^*])^{1-\frac{2}{q}}}{n}. \quad (17)$$

Combining the bounds in equations (16) and (17) with the bias-variance decomposition in equation (15), we have that

$$\begin{aligned} \mathcal{R}(G, \tau, q) &\leq 8.92 \cdot \frac{\lambda(\mathcal{H}(V_{\mathcal{H}} \setminus S^*))^{\frac{2}{q}} \cdot d_{\text{avg}}(\mathcal{H}[V_{\mathcal{H}} \setminus S^*])^{1-\frac{2}{q}}}{n} + \frac{|S^*|^{\frac{2(q-1)}{q}}}{n^{2-\frac{2}{q}}} \\ &= 8.92 \cdot \frac{Q_q(S^*, \mathcal{H})}{n}. \end{aligned}$$

Since S^* was selected in line 1 to optimize the right hand side of the above inequality, the result follows. $\square$

C.3 Polynomial Time Approximation Algorithm (Proposition 4.3)

In this section, we prove that there exists a design-estimator pair that run in polynomial time and achieve estimation rate within logarithmic factors from that of Theorem 4.2. The precise statement follows.

Proposition 4.3. *Let G be an interference network, τ be a contrastive effect. For any moment parameter $q \geq 2$, there exists a design-estimator pair $(\mathcal{D}, \hat{\tau})$ that can be implemented in polynomial time whose maximum risk is bounded as*

$$\sup_{y \in \mathcal{M}(G, \tau, q)} \text{MSE}(\mathcal{D}, \hat{\tau}, \tau, y) \leq \mathcal{O}\left((\log n)^{2-\frac{2}{q}}\right) \cdot \frac{\min_{S \subseteq V(\mathcal{H})} Q_q(S, \mathcal{H})}{n}.$$

In order to establish Theorem 4.3, it will be useful to introduce a few relevant computational problems. The first one is known as the *Partial Vertex Cover* problem.

Definition 4 (Partial Vertex Cover). The *Partial Vertex Cover* problem is defined as follows.

Input: A graph $G = (V, E)$ and an integer k .

Output: A minimum-size subset $S \subseteq V$ such that $|E(G[V \setminus S])| \leq k$.

For a graph G and integer k , we denote by $\text{OPT}_{pvc}(G, k)$ the size of the optimal solution to the Partial Vertex Cover problem on input (G, k) .

The second problem is a variant of the first one. We will refer to it as the *Partial Eigenvalue Cover* problem, although this is not a standard terminology in the literature.

Definition 5 (Partial Eigenvalue Cover). The *Partial Eigenvalue Cover* problem is defined as follows.

Input: A graph $G = (V, E)$ and an integer k .

Output: A minimum-size subset $S \subseteq V$ such that $\lambda(G[V \setminus S]) \leq k$.

For a graph G and integer k , we denote by $\text{OPT}_{pec}(G, k)$ the size of the optimal solution to the Partial Eigenvalue Cover problem on input (G, k) .

We now present two results from the literature, which provide polynomial time approximation algorithms for the above problems.

Theorem C.11 (Theorem 2.1 in Gandhi, Khuller, and Srinivasan (2004)). *There exists a polynomial time algorithm that takes as input a graph G and integer k and outputs a subset $\text{PVC-Approx}(G, k) \subseteq V$ such that $|\text{PVC-Approx}(G, k)| \leq 2 \cdot \text{OPT}_{pvc}(G, k)$.*

Theorem C.12 (Lemma 5.1 in Saha et al. (2015)). *There exists a polynomial time algorithm that takes as input a graph G and integer k and outputs a subset $\text{PEC-Approx}(G, k) \subseteq V(G)$ such that*

$$|\text{PEC-Approx}(G, k)| = \mathcal{O}(\text{OPT}_{pec}(G, k) \cdot \log n)$$

and

$$\lambda(G[V \setminus \text{PEC-Approx}(G, k)]) \leq 2k \ .$$

We will use the above two algorithms as subroutines in order to find a subset S of vertices in the conflict graph $\mathcal{H}$ that approximately minimizes the expression in Theorem 4.2.

In particular, we describe the design in Algorithm 4.

The only difference between Algorithm 4 and the previously described Algorithm 3 is in the selection of the optimal subset of exposures. While Algorithm 3 searches over all subsets of exposures to minimize $Q_q(S, \mathcal{H})$, which takes time exponential in n , Algorithm 4 uses a different procedure that runs in polynomial time.

We next establish that the value $Q_q(\hat{S}, \mathcal{H})$ is only logarithmically larger than the optimal value of $Q_q(S^*, \mathcal{H})$.

Lemma C.13. *Let $S^* = \arg \min_{S \subseteq V(\mathcal{H})} Q_q(S, \mathcal{H})$ as defined in Algorithm 3 and $\hat{S}$ be the subset selected in line 6 of Algorithm 4. Then,*

$$Q_q(\hat{S}, \mathcal{H}) = \mathcal{O}\left((\log n)^{2-\frac{2}{q}}\right) \cdot Q_q(S^*, \mathcal{H}) \ .$$

Algorithm 4: Approximation Design

Input: Interference network G , conflict graph $\mathcal{H}$ and moment q .

Output: Random intervention $Z \in \mathcal{Z} = \{0, 1\}^n$, variables $U_{i,k}$ for all $e_{i,k} \in V(\mathcal{H})$.

```

1 for  $k_1 = 1$  to  $|E(\mathcal{H})|$  do
2   for  $k_2 = 1$  to  $\lceil \lambda(\mathcal{H}) \rceil$  do
3      $S_1 \leftarrow \text{PVC-Approx}(\mathcal{H}, k_1)$ ;
4      $S_2 \leftarrow \text{PEC-Approx}(\mathcal{H}, k_2)$ ;
5      $S_{k_1, k_2} \leftarrow S_1 \cup S_2$ ;
6  $\hat{S} \leftarrow \arg \min_{k_1, k_2} Q_q(S_{k_1, k_2}, \mathcal{H})$ 
7  $\hat{\mathcal{H}} \leftarrow \mathcal{H}[V_{\mathcal{H}} \setminus \hat{S}]$ 
8  $\pi_{\hat{S}} \leftarrow$  Ordering produced by Algorithm 1 with input  $\hat{\mathcal{H}}$  and  $q$ 
9 for  $e_{i,k} \in V_{\mathcal{H}} \setminus \hat{S}$  do
10    $p_{i,k} \leftarrow \frac{1}{2\lambda(\mathcal{H})^{\frac{2}{q}} d_{i,k}(\hat{\mathcal{H}})^{1-\frac{2}{q}}}$ ;
11 Sample desired exposure variables  $\{U_{i,k} : e_{i,k} \in V_{\mathcal{H}} \setminus \hat{S}\}$  independently and identically as

```

$$U_{i,k} \leftarrow \begin{cases} e & \text{with probability } \frac{p_{i,k}}{2} \\ * & \text{with probability } 1 - \frac{p_{i,k}}{2} \end{cases}$$

```

12 Initialize the intervention vector  $Z \leftarrow \mathbf{0}$ 
13 for  $e_{i,k} \in V_{\mathcal{H}} \setminus \hat{S}$  do
14   if  $U_{i,k} = e$  and  $U_{j,l} = *$  for all  $e_{j,l} \in \mathcal{N}_b^{\pi_{\hat{S}}}(e_{i,k})$  then
15      $\text{Update intervention vector: set } Z_r = \mathbf{1}[r \in h_i(e_{i,k})]$  for all  $r \in \tilde{N}(i)$  ;

```

Proof. Let $k_1^* = |E(V_{\mathcal{H}} \setminus S^*)|$ and $k_2^* = \lceil \lambda(\mathcal{H}(V_{\mathcal{H}} \setminus S^*)) \rceil$. We first note that $\text{OPT}_{pvc}(\mathcal{H}, k_1^*) \leq |S^*|$ and $\text{OPT}_{pec}(\mathcal{H}, k_2^*) \leq |S^*|$. This is because S^* is a feasible solution to both the **Partial Vertex Cover** problem with input $(\mathcal{H}, k_1^*)$ and the **Partial Eigenvalue Cover** problem with input $(\mathcal{H}, k_2^*)$.

For convenience, let us denote

$$S_1 = \text{PVC-Approx}(\mathcal{H}, k_1^*) \quad \text{and} \quad S_2 = \text{PEC-Approx}(\mathcal{H}, k_2^*) .$$

Then, by Theorem C.11, we have that $|S_1| \leq 2 \cdot \text{OPT}_{pvc}(\mathcal{H}, k_1^*) \leq 2|S^*|$. Also, by Theorem C.12, we have that

$$|S_2| = \mathcal{O}(\text{OPT}_{pec}(\mathcal{H}, k_2^*) \cdot \log n) = \mathcal{O}(|S^*| \cdot \log n) \quad \text{and} \quad \lceil \lambda(\mathcal{H}[V_{\mathcal{H}} \setminus S_2]) \rceil \leq 2k_2^* .$$

Finally, $S_{k_1^*, k_2^*} = S_1 \cup S_2$ has size $|S_{k_1^*, k_2^*}| \leq |S_1| + |S_2| = \mathcal{O}(|S^*| \cdot \log n)$.

Therefore, at the iteration where $k_1 = k_1^*$ and $k_2 = k_2^*$, we have that

$$\begin{aligned}
n^{1-\frac{2}{q}} \cdot Q_q(S_{k_1^*, k_2^*}, \mathcal{H}) &= \lambda(\mathcal{H}(V_{\mathcal{H}} \setminus S_{k_1^*, k_2^*}))^{\frac{2}{q}} \cdot (2|E(V_{\mathcal{H}} \setminus S_{k_1^*, k_2^*})|)^{1-\frac{2}{q}} + |S_{k_1^*, k_2^*}|^{\frac{2(q-1)}{q}} \\
&\leq (2k_2^*)^{\frac{2}{q}} \cdot (2k_1^*)^{1-\frac{2}{q}} + \mathcal{O}(|S^*| \cdot \log n)^{\frac{2(q-1)}{q}} \\
&= \mathcal{O}\left((\log n)^{2-\frac{2}{q}}\right) \cdot \left(\lambda(\mathcal{H}(V_{\mathcal{H}} \setminus S^*))^{\frac{2}{q}} \cdot |E(V_{\mathcal{H}} \setminus S^*)|^{1-\frac{2}{q}} + |S^*|^{\frac{2(q-1)}{q}}\right)
\end{aligned}$$

$$= \mathcal{O}\left((\log n)^{2-\frac{2}{q}}\right) \cdot n^{1-\frac{2}{q}} Q_q(S^*, \mathcal{H}) .$$

By definition, we have $Q_q(\hat{S}, \mathcal{H}) \leq Q_q(S_{k_1^*, k_2^*}, \mathcal{H})$, which implies the result. $\square$

We are now ready to establish Proposition 4.3.

Proof of Proposition 4.3. As our design $\mathcal{D}$, we will use the design of Algorithm 4 and the estimator $\hat{\tau}_{\hat{S}}$, where $\hat{S}$ is defined in line 6 of Algorithm 4.

We first note that Algorithm 4 runs in polynomial time. This is because the total number of iterations in the for loops in the beginning is $\mathcal{O}(|E(\mathcal{H})| \cdot \lambda(\mathcal{H})) = \mathcal{O}(n^3)$, and each iteration involves running the algorithms of Theorem C.11 and Theorem C.12, which run in polynomial time. The remaining steps of the algorithm also clearly run in polynomial time. After having obtained $\hat{S}$, one can also compute the estimator $\hat{\tau}_{\hat{S}}$ in polynomial time.

Next, we can proceed exactly as in the proof of Theorem 4.2 to show that

$$\text{MMSE}(\mathcal{D}, \hat{\tau}_{\hat{S}}, \mathcal{M}(G, q)) \leq 8.92 \cdot \frac{Q_q(\hat{S}, \mathcal{H})}{n^{2-\frac{2}{q}}} .$$

Finally, by Lemma C.13, we have that

$$Q_q(\hat{S}, \mathcal{H}) = \mathcal{O}\left((\log n)^{2-\frac{2}{q}}\right) \cdot Q_q(S^*, \mathcal{H}) ,$$

where $S^* = \arg \min_{S \subseteq V(\mathcal{H})} Q_q(S, \mathcal{H})$ as defined in Algorithm 3. The result follows. $\square$

D Illustrations of Main Results

In this section, we present the proofs of two illustrations of the main result for estimating two well-studied causal effects, as presented in Section 1.1. Section D.1 contains the proof for the Direct Treatment Effect, while Section D.2 contains the proof for the Global Average Treatment Effect.

D.1 Direct Treatment Effect (Corollary 1.1)

The Direct Treatment Effect (DTE) was formally defined in Section 2.1. We state and prove a more formal version of Corollary 1.1 below.

Corollary D.1. *The minimax rate of estimating the Direct Treatment Effect τ_{DTE} on a network G is bounded as*

$$\max\left(\frac{1}{|\mathcal{I}(G)|}, \frac{\sqrt{d_{\text{avg}}(G)}}{n}\right) \lesssim \mathcal{R}(G, \tau_{\text{DTE}}, 2) \lesssim \frac{\lambda(G) + 1}{n} ,$$

where $\mathcal{I}(G)$ is the maximum independent set in G , $d_{\text{avg}}(G)$ is its average degree, and $\lambda(G)$ is the largest eigenvalue of its adjacency matrix.

Proof. Let $\mathcal{H}$ be the conflict graph corresponding to G and τ_{DTE} . Let $S = \{e_{i,1} : i \in [n]\}$ be the subset of size n of the $e_{i,1}$ exposures. Applying the upper bound of Theorem 4.1 for $q = 2$, as well

as the global lower bound of Theorem 3.2 with $T = S$ and the local lower bound of Corollary 3.5 together with Proposition B.6, we get

$$\max\left(\frac{1}{|\mathcal{I}(\mathcal{H}, S)|}, \frac{\sqrt{d_{\text{avg}}(\mathcal{H})}}{n}\right) \lesssim \mathcal{R}(G, \tau_{\text{DTE}}) \lesssim \frac{\lambda(\mathcal{H})}{n} \quad (18)$$

We now connect the quantities on the upper and lower bounds with the desired quantities of the interference network $G = (V, E)$.

As we have discussed in Section 2.3, the conflict graph $\mathcal{H}$ for τ_{DTE} has the following structure: $(e_{i,k}, e_{j,l}) \in E_{\mathcal{H}}$ if and only if $(i, j) \in E$ and $(k, l) \neq (0, 0)$ or $i = j$ and $k \neq l$. Therefore, the induced subgraph $\mathcal{H}(S)$ is isomorphic to G , which implies that

$$|\mathcal{I}(\mathcal{H}, S)| = |E(G)|. \quad (19)$$

Also, this means that $\mathcal{H}$ has at least as many edges as G , which means that

$$d_{\text{avg}}(\mathcal{H}) = \frac{2|E_{\mathcal{H}}|}{2n} \geq \frac{2|E|}{2n} = \frac{d_{\text{avg}}(G)}{2}. \quad (20)$$

Thus, combining (18) together with (19) and (20) yields the lower bound of the statement.

For the upper bound, let $\mathbf{A}_{\mathcal{H}}, \mathbf{A}_G$ denote the adjacency matrices of $\mathcal{H}$ and G , respectively. If we order the nodes in $\mathcal{H}$ starting from the n exposures $e_{i,0}$ and then moving to the n exposures $e_{i,1}$, then by the preceding discussion we can write $\mathbf{A}_{\mathcal{H}}$ in block form as follows.

$$\mathbf{A}_{\mathcal{H}} = \begin{pmatrix} 0 & I \\ I & \mathbf{A}_G \end{pmatrix}.$$

Now, standard matrix bounds imply that

$$\lambda(\mathcal{H}) = \|\mathbf{A}_{\mathcal{H}}\|_{\text{op}} \leq 2\|I\|_{\text{op}} + \|\mathbf{A}_G\|_{\text{op}} = 2 + \lambda(G).$$

This bound combined with (18) yields the upper bound of the statement. $\square$

D.2 Global Average Treatment Effect (Corollary 1.2)

The Global Average Treatment Effect (GATE) was formally defined in Section 2.1. We state and prove a more formal version of Corollary 1.2 below.

Corollary D.2. *The minimax rate of estimating the Global Average Treatment Effect τ_{GATE} on a network G is bounded as*

$$\frac{\sqrt{d_{\text{avg}}(G^2)}}{n} \lesssim \mathcal{R}(G, \tau_{\text{GATE}}, 2) \lesssim \frac{\lambda(G^2) + 1}{n},$$

where G^2 is the two-hop interference graph, $d_{\text{avg}}(G^2)$ is the average degree of the two-hop graph, and $\lambda(G^2)$ is the largest eigenvalue of its adjacency matrix.

Proof. Let $\mathcal{H}$ be the conflict graph corresponding to G and τ_{GATE} . Applying the upper bound of Theorem 4.1 for $q = 2$, as well as the local lower bound of Corollary 3.5 together with Proposition B.6, we get

$$\frac{\sqrt{d_{\text{avg}}(\mathcal{H})}}{n} \lesssim \mathcal{R}(G, \tau_{\text{GATE}}, 2) \lesssim \frac{\lambda(\mathcal{H})}{n} \quad (21)$$

We now connect the quantities on the upper and lower bounds with the desired quantities of the 2-hop graph $G^2 = (V, E(G^2))$.

As we have discussed in Section 2.3, the conflict graph $\mathcal{H}$ for τ_{GATE} has the following structure: $(e_{i,k}, e_{j,l}) \in E_{\mathcal{H}}$ if and only if $(i, j) \in E(G^2)$ and $k \neq l$ or $i = j$ and $k \neq l$. Therefore, for every edge $(i, j) \in E(G^2)$ in the 2-hop graph G^2 , there are two edges $(e_{i,1}, e_{j,0}), (e_{i,0}, e_{j,1}) \in E_{\mathcal{H}}$. It follows that $|E_{\mathcal{H}}| \geq 2|E(G^2)|$, which implies that

$$d_{\text{avg}}(\mathcal{H}) = \frac{2|E_{\mathcal{H}}|}{2n} \geq \frac{4|E(G^2)|}{2n} = d_{\text{avg}}(G^2) .$$

This inequality combined with (21) establishes the lower bound of the statement.

For the upper bound, let $\mathbf{A}_{\mathcal{H}}, \mathbf{A}_{G^2}$ denote the adjacency matrices of $\mathcal{H}$ and G^2 respectively. Then, if we order the $2n$ nodes of $\mathcal{H}$ starting from the $e_{i,0}$ exposures and then continuing with the $e_{i,1}$ exposures, we can write $\mathbf{A}_{\mathcal{H}}$ in block form as

$$\begin{pmatrix} 0 & I + \mathbf{A}_{G^2} \\ I + \mathbf{A}_{G^2} & 0 \end{pmatrix} .$$

The identity part comes from the edges $(e_{i,0}, e_{i,1})$. Therefore, we have

$$\lambda(\mathcal{H}) = \|\mathbf{A}_{\mathcal{H}}\|_{\text{op}} \leq 2\|I + \mathbf{A}_{G^2}\|_{\text{op}} \leq 2 + 2\lambda(G^2)$$

The above inequality combined with (21) yields the lower bound in the statement. $\square$

E Tightness of Bounds

In this section, we present the proofs of the results in Section 5. We start by providing a convenient description of the conflict graph for direct effect when G is an arbitrary d -regular graph. We then continue by establishing that both the upper and lower bounds for estimating the direct effect in regular graphs are tight. Finally, we show that for a random regular graph, the upper bound is tight up to a $\log d$ factor with high probability.

E.1 The Structure of the Conflict Graph

We first provide a Lemma that describes the structure of the conflict graph $\mathcal{H}$ when the estimand of interest is the Direct Effect. We denote by $\mathbf{e}_i \in \mathbb{R}^n$ the vector with 1 in the i -th entry and 0 everywhere else and $\mathbf{0} \in \mathbb{R}^n$ be the all 0 vector. Also, for a symmetric matrix $A \in \mathbb{R}^{n \times n}$, we denote by $\|A\|_2$ its spectral norm. Furthermore, for any graph G , we denote by $\mathbf{A}_G$ its adjacency matrix.

Lemma E.1. *Let $G = (V, E)$ be a graph and $\mathcal{H} = (V_{\mathcal{H}}, E_{\mathcal{H}})$ be the conflict graph associated with G and the direct effect τ . Then, we have that*

- (i) If a node $i \in G$ has degree d_i , then $e_{i,1} \in \mathcal{H}$ has degree $2d_i + 1$ and $e_{i,0} \in \mathcal{H}$ has degree at most $d_i + 1$.
- (ii) Let $T = \{e_{i,1} : i \in [n]\}$ be the subset of vertices of $\mathcal{H}$. Then, $\mathcal{H}[T]$ is a graph isometric to G , where the isometry maps $e_{i,1}$ to i for each $i \in [n]$.

Proof. Let us start with the first part. The vertices of $\mathcal{H}$ are the $2n$ exposures $e_{i,1} = \mathbf{e}_i$ and $e_{i,0} = \mathbf{0}$ for $i \in [n]$. For every $i \in [n]$, there is an edge $(e_{i,1}, e_{i,0}) \in E_{\mathcal{H}}$, since we can not observe both exposures for the same unit. Also, for every $i, j \in [n]$ with $(i, j) \in E$, we have that $(e_{i,1}, e_{j,0}), (e_{i,0}, e_{j,1}), (e_{i,1}, e_{j,1}) \in E_{\mathcal{H}}$. This is because if i receives the $e_{i,1}$ exposure, then i is treated, meaning that j has a neighbor that is treated, so j can not receive the $e_{j,0}$ or the $e_{j,1}$ exposure. Thus, the first property follows.

For the second part, it follows from the previous discussion that for any $i, j \in [n]$, we have that $(e_{i,1}, e_{j,1}) \in E_{\mathcal{H}}$ if and only if $(i, j) \in E$. Therefore, the subgraph of $\mathcal{H}$ induced by T is isometric to G . $\square$

E.2 Tightness of Upper Bound for Direct Effect (Proposition 5.1)

Proposition 5.1. For any sequence $d = \{d^{(n)}\}_{n=1}^{\infty}$ with $d = o(n)$,

$$\sup_{G \in \mathcal{G}_d} \mathcal{R}(G, \tau_{\text{DTE}}, 2) \asymp \frac{d}{n}.$$

Proof. Let d_n be a sequence of degrees such that $d_n = o(n)$ and $d_n n$ is even for all n . For a given n , we will construct a d_n -regular graph G_n on n nodes, which we call the d -regular chorded cycle. We will assume that n is large enough so that $d_n < n/2$, which is sufficient since our statement is asymptotic. The k -regular chorded cycle on n nodes $H_{n,k}$ is defined as follows. We arrange the n nodes in a cycle, labeled $1, 2, \dots, n$. Two nodes $i, j \in [n]$ are connected if and only if $|i - j| \leq k$, where the difference is taken modulo n . In other words, each node is connected to the k nodes on each side of it in the cycle. From the definition it follows immediately that $H_{n,k}$ is $2k$ -regular.

Let us now define G_n . Since nd_n is even, it must be that either n or d_n is even. We distinguish between two cases.

- If d_n is even, we define $G_n = H_{n,d_n/2}$.
- If d_n is odd, then n has to be even. We define G_n to be the graph obtained by taking the union of $H_{n,(d_n-1)/2}$ and a perfect matching on the n nodes.

In either case, G_n is d_n -regular. We now show that $\mathcal{R}(G_n, \tau, 2) = \Theta(d_n/n)$.

For the upper bound, we apply Theorem 4.1. Since G_n is d_n -regular, by (i) of Lemma E.1 it follows that every node in $V_{\mathcal{H}}$ has degree at most $2d_n + 1$. Thus, $\lambda(\mathcal{H}) \leq 2d_n + 1$. Thus, applying Theorem 4.1 with $q = 2$, we get that

$$\mathcal{R}(G_n, \tau, 2) \lesssim \frac{2d_n + 1}{n} = \mathcal{O}\left(\frac{d_n}{n}\right).$$

For the lower bound, we apply Theorem 3.2 on the subset $T = \{e_{i,1} : i \in [n]\}$ of the vertices of $\mathcal{H}$. We will upper bound the size of the largest independent set of $\mathcal{H}[T]$. By the isometry property (ii) of Lemma E.1, it suffices to bound the largest independent set of G . We will assume for simplicity

that n is divided by $\lfloor d_n/2 \rfloor$, since the general case can be handled by a standard modification of this argument. We partition $[n]$ into the subsets $\{T_i\}_{i=1}^{\lfloor d_n/2 \rfloor}$, where

$$T_i = [\lfloor d_n/2 \rfloor(i-1), \lfloor d_n/2 \rfloor i - 1] \cap [n] .$$

Clearly T_i are disjoint and $[n] = \cup_{i=1}^{\lfloor d_n/2 \rfloor} T_i$. Also, within each T_i all nodes are connected to each other in G , since for any $i, j \in T_i$ we have that $|i - j| \leq \lfloor d_n/2 \rfloor \leq d_n$. Therefore, any independent set in $\mathcal{H}[T]$ can contain at most one node from each T_i . Since there are at most

$$\frac{n}{\lfloor \frac{d_n}{2} \rfloor} \leq \frac{n}{d_n/2 - 1} = \mathcal{O}\left(\frac{n}{d_n}\right)$$

such subsets, we have that $|\mathcal{I}(\mathcal{H}[T])| = |\mathcal{I}| = \mathcal{O}(\frac{n}{d_n})$. Thus, applying Theorem 3.2 for the subset T , we get that

$$\mathcal{R}(G_n, \tau, 2) \gtrsim \frac{1}{|\mathcal{I}(\mathcal{H}[T])|} \gtrsim \frac{d_n}{n} .$$

□

E.3 Tightness of Lower Bound for Direct Effect (Proposition 5.2)

For Proposition 5.2, we will use a family of graphs called the *Generalized Quadrangle*. Thus, we first review some of the properties of these graphs that will be relevant for our purposes. To do that, it will be helpful to introduce some basic concepts from the algebraic theory of quadratic forms. For a detailed introduction to this topic, see Lam (2005). Let $s \in \mathbb{N}$ be a prime number and $\mathbb{F}_s$ be the finite field of residues modulo s . Let $r \in \mathbb{N}$ be a positive integer. For every $x \in \mathbb{F}_s^r$, the corresponding *projective point* is the set of all scalar multiples of x , i.e.

$$[x] = \{y \in \mathbb{F}_s^r : y = \lambda x \text{ for some } \lambda \in \mathbb{F}_s\} .$$

The set of all projective points is called the *projective space* $PG(r-1, s)$. Let $B : \mathbb{F}_s^r \times \mathbb{F}_s^r \rightarrow \mathbb{F}_s$ be a bilinear form. Let $Q(r-1, s)$ be the set of all projective points $[x] \in PG(r-1, s)$ such that $B(x, x) = 0$. We call $Q(r-1, s)$ the *quadratic* associated with B on the vector space $\mathbb{F}_s^r$. There are three types of quadratics, depending on the dimension r and the Witt index¹ w of the bilinear form B .

1. If r is even and $w = r/2$, then $Q(r-1, s)$ is called a *hyperbolic quadratic* and denoted $Q^+(r-1, s)$.
2. If r is even and $w = r/2 - 1$, then $Q(r-1, s)$ is called an *elliptic quadratic* and denoted $Q^-(r-1, s)$.
3. If r is odd and $w = (r-1)/2$, then $Q(r-1, s)$ is called a *parabolic quadratic* and denoted $Q(r-1, s)$.

Most of our discussion will focus on the case $r = 5$. Consider the matrix $M \in \mathbb{F}_s^{5 \times 5}$ defined as

$$M = \left(\begin{array}{cc|cc|c} 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 1 \end{array} \right) .$$

¹For the definition of Witt index, see Chapter 1 in Lam (2005).

We associate to M the bilinear form $B : V \times V \rightarrow \mathbb{F}_s$ defined as $B(x, y) = x^T M y$ for all $x, y \in V$. We will denote by $Q(4, s)$ the quadric associated with B on the vector space $\mathbb{F}_s^5$. The Witt index of B is 2, thus by the above discussion $Q(4, s)$ is a parabolic quadric. The points in $Q(4, s)$ will be the vertices of our graph.

A *line* L in $Q(4, s)$ is a 2-dimensional subspace of V such that $B(x, x) = 0$ for all $x \in L$. We say a point $[x]$ is *incident* to a line L if $x \in L$. The *point graph* $\Gamma(Q(4, s))$ associated with $Q(4, s)$ is the graph with vertices corresponding to the points in $Q(4, s)$, where two points are connected by an edge if and only if they are incident to the same line.

We now summarize some basic properties of the graph $\Gamma(Q(4, s))$ related to the number of vertices, edges and spectrum of the adjacency matrix.

Lemma E.2. *Let $s \geq 2$ be a prime number. The following properties hold for the graph $G = \Gamma(Q(4, s))$.*

- (i) *The number of nodes is $n = (s + 1)(s^2 + 1)$.*
- (ii) *The graph is d -regular with $d = s(s + 1)$.*
- (iii) *The eigenvalues of the adjacency matrix $\mathbf{A}_G$ are $s(s + 1), s - 1, -(s + 1)$, with multiplicities $1, s(s + 1), s^2$ respectively.*

Proof. As established in Section 3.1 of Payne and Thas (2009), G satisfies the general properties of a point-line geometry. From these properties, it follows that G is a d -regular graph with $d = s(s + 1)$ and $n = (s + 1)(s^2 + 1)$ nodes (see e.g. Godsil and Royle, 2013, Lemma 10.8.1). This establishes properties (i) and (ii). Property (iii) is exactly Lemma 10.8.2 of Godsil and Royle (2013). $\square$

In order for us to construct an improved design that achieves the $\sqrt{d}/n$ rate for this family of graphs, it is important to understand the structure of independent sets in G . We start by defining a *hyperplane* of $\text{PG}(r, q)$, which is a set of the form

$$H = \{[x] \in \text{PG}(r, q) : \alpha^\top x = 0\},$$

for some $\alpha \in \mathbb{F}_q^r \setminus \{0\}$. Let H be a hyperplane of $\text{PG}(4, s)$ and $I = H \cap Q(4, s)$. If I is isomorphic to the elliptic quadric $Q^-(3, s)$, then we call I an *ovoid* of $Q(4, s)$. By definition, I will have Witt index 1, thus it will contain no lines, meaning that I is an independent set in G . By Lemma 3.4.1 in Payne and Thas (2009), ovoids always exist in $Q(4, s)$. The following lemma establishes some crucial properties of ovoids that will be used in the construction of our design.

Lemma E.3. *Let $G = \Gamma(Q(4, s))$ and $\mathcal{F}$ be the family of ovoids in $Q(4, s)$. Then, the following properties hold.*

- (i) *Every ovoid has size $s^2 + 1$.*
- (ii) *Every node i belongs in the same number of independent sets in the family $\{I_f\}_{f \in \mathcal{F}}$.*
- (iii) *Every pair of nodes i, j for which $(i, j) \notin E$ belongs in the same number of independent sets in the family $\{I_f\}_{f \in \mathcal{F}}$.*

Proof. Property (i) is exactly Lemma 1.8.1 of Payne and Thas (2009). For the second and third properties, we will use Witt's Extension Theorem, as stated in Theorem 22 of Clark (2013). This Theorem states that any isometry between two subspaces of a quadratic space can be extended to an isometry of the whole space. In particular, let $[x], [y] \in Q(4, s)$ be two different vertices of the graph. Consider the one dimensional subspaces $V_1 = [x]$ and $V_2 = [y]$, which are clearly isomorphic. By Witt's Extension Theorem, there exists an isometry of the whole space that maps V_1 to V_2 , thus maps $[x]$ to $[y]$. Since it is an isometry, it preserves the quadratic form, thus all edges between vertices are preserved. This means that ovoids containing $[x]$ are in 1-1 mapping with ovoids containing $[y]$, establishing property (ii).

Also, let $([x], [y])$ and $([z], [u])$ be two pairs of non-adjacent nodes in G . Consider the two-dimensional subspaces $V_1 = \text{span}([x], [y])$ and $V_2 = \text{span}([z], [u])$. Since $[x], [y]$ are non-adjacent, we have that $B(x, y) \neq 0$ and $B(x, x) = B(x, y) = 0$, thus the restriction of the quadratic form to V_1 leads to a hyperbolic quadric. The same is true for V_2 , thus V_1 and V_2 are isometric. By Witt's Extension Theorem, there exists an isometry of the whole space that maps V_1 to V_2 . Since edges are preserved by the isometry, ovoids containing $[x], [y]$ are in 1-1 mapping with ovoids containing $[z], [u]$, establishing property (iii). $\square$

We will use the family $\mathcal{F}$ of independent sets in the graph G of Lemma E.3 to construct a design that achieves the $\sqrt{d}/n$ rate for this family of graphs. The pseudocode for the design is given in Algorithm 5.

Algorithm 5: Independent Set Design

Input: Graph G on n nodes as in Lemma E.2.

Output: Random intervention $Z \in \{0, 1\}^n$.

- 1 Initialize the intervention vector $Z \leftarrow \mathbf{0}$.
- 2 Select $f \in \mathcal{F}$ uniformly at random.
- 3 Sample variables $\{U_i : i \in [n]\}$ independently and identically as

$$U_i \leftarrow \begin{cases} 0 & \text{with probability } 1/2 \\ 1 & \text{with probability } 1/2 \end{cases}$$

```

for  $i \in [n]$  do
4   if  $i \in I_f$  then
5      $Z_i \leftarrow U_i$ 
6   else
7      $Z_i \leftarrow 0$ 

```

The design works by selecting a uniformly random independent set from the family $\{I_f\}_{f \in \mathcal{F}}$ and then treating the nodes in that independent set with probability 1/2 each, while leaving the rest of the nodes untreated. Therefore, we call this the *Independent Set Design*. The properties of the family $\{I_f\}_{f \in \mathcal{F}}$ will enable a precise characterization of the spectrum of the covariance matrix that arises from this design.

Our estimator $\hat{\tau}$ will be a Horvitz-Thompson type of estimator. For $i \in [n], k \in \{0, 1\}$, we define the event $A_{i,k} = \{i \in I_f, Z_i = k\}$, which corresponds to the event that node i is in the independent set that was selected and receives treatment k . We will use the following estimator for τ

$$\hat{\tau} = \frac{1}{n} \sum_{i=1}^n \left(\frac{\mathbf{1}[A_{i,1}]}{\Pr(A_{i,1})} - \frac{\mathbf{1}[A_{i,0}]}{\Pr(A_{i,0})} \right) Y_i . \quad (22)$$

Our next lemma shows that this estimator is unbiased for τ .

Lemma E.4. *The estimator $\hat{\tau}$ defined in (22) is an unbiased estimator for τ , i.e. $\mathbb{E}[\hat{\tau}] = \tau$.*

Proof. We start by observing that if $i \in I_f$ and $Z_i = k$, then all neighbors of i are in control thus $h_i(Z) = e_{i,k}$ and $Y_i = y_i(e_{i,k})$. It also follows by Algorithm 5 that $\Pr(A_{i,k}) > 0$ for all $i \in [n], k \in \{0, 1\}$. Therefore, we can write

$$\begin{aligned} \mathbb{E}[\hat{\tau}] &= \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \left(\frac{\mathbf{1}[A_{i,1}]}{\Pr(A_{i,1})} y_i(e_{i,1}) - \frac{\mathbf{1}[A_{i,0}]}{\Pr(A_{i,0})} y_i(e_{i,0}) \right) \right] \\ &= \frac{1}{n} \sum_{i=1}^n \left(\frac{\Pr(A_{i,1})}{\Pr(A_{i,1})} y_i(e_{i,1}) - \frac{\Pr(A_{i,0})}{\Pr(A_{i,0})} y_i(e_{i,0}) \right) \\ &= \tau . \end{aligned}$$

□

We now provide an auxiliary lemma that contains expressions for the variance and covariance terms arising in the analysis of the variance of the estimator $\hat{\tau}$.

Lemma E.5. *Suppose G is given by Lemma E.2 and Z is generated according to the Independent Set Design of Algorithm 5. Then, for every $i, j \in [n], k \in \{0, 1\}$ we have that*

1.

$$\text{Var} \left(\frac{\mathbf{1}[A_{i,k}]}{\Pr(A_{i,k})} \right) = s .$$

2. If $i \neq j$

$$\text{Cov} \left(\frac{\mathbf{1}[A_{i,k}]}{\Pr(A_{i,k})}, \frac{\mathbf{1}[A_{j,k}]}{\Pr(A_{j,k})} \right) = \begin{cases} \frac{1}{s} & \text{if } (i, j) \notin E \\ -1 & \text{if } (i, j) \in E \end{cases} .$$

Proof. We start with the first part. We have that

$$\text{Var} \left(\frac{\mathbf{1}[A_{i,k}]}{\Pr(A_{i,k})} \right) = \frac{1}{\Pr(A_{i,k})} - 1$$

Thus, it suffices to compute $\Pr(A_{i,k})$. Since the subset I_f is selected uniformly at random from the family $\{I_f\}_{f \in \mathcal{F}}$, independently of the random variables $\{U_i\}_{i=1}^n$, we have that

$$\begin{aligned} \Pr(A_{i,k}) &= \Pr(i \in I_f, Z_i = k) \\ &= \Pr(i \in I_f) \Pr(Z_i = k | i \in I_f) \\ &= \Pr(i \in I_f) \Pr(U_i = k) \\ &= \frac{1}{2} \frac{|\{f : i \in I_f\}|}{|\mathcal{F}|} . \end{aligned}$$

Thus, it suffices to compute $|\{f : i \in I_f\}|$. By property (ii) of Lemma E.3, every node belongs in the same number of independent sets in the family $\{I_f\}_{f \in \mathcal{F}}$, thus this set has the same size M for all $i \in [n]$. We have that

$$\begin{aligned} Mn &= \sum_{i=1}^n |\{f : i \in I_f\}| \\ &= \sum_{f \in \mathcal{F}} |I_f| \end{aligned}$$

by exchanging the order of summation

$$= |\mathcal{F}|(s^2 + 1) ,$$

where in the last step we used property (i) of Lemma E.3. Thus, $M = |\mathcal{F}|(s^2 + 1)/n$. Therefore, we have that

$$\begin{aligned} \Pr(A_{i,k}) &= \frac{1}{2} \frac{|\{f : i \in I_f\}|}{|\mathcal{F}|} \\ &= \frac{1}{2} \frac{|\mathcal{F}|(s^2 + 1)/n}{|\mathcal{F}|} \\ &= \frac{1}{2} \frac{s^2 + 1}{n} \\ &= \frac{1}{2} \frac{s^2 + 1}{(s + 1)(s^2 + 1)} \\ &= \frac{1}{2(s + 1)} , \end{aligned}$$

where we used property (i) of Lemma E.2 in the last step. Thus, we have that

$$\text{Var}\left(\frac{\mathbf{1}[A_{i,k}]}{\Pr(A_{i,k})}\right) = \frac{1}{\Pr(A_{i,k})} - 1 = 2s + 1$$

For the second part, we have that

$$\text{Cov}\left(\frac{\mathbf{1}[A_{i,k}]}{\Pr(A_{i,k})}, \frac{\mathbf{1}[A_{j,k}]}{\Pr(A_{j,k})}\right) = \frac{\Pr(A_{i,k} \cap A_{j,k})}{\Pr(A_{i,k}) \Pr(A_{j,k})} - 1$$

Thus, it suffices to compute $\Pr(A_{i,k} \cap A_{j,k})$. If $(i, j) \in E$, then i and j can not be in the same independent set, thus $\Pr(A_{i,k} \cap A_{j,k}) = 0$. This proves the first case. If $(i, j) \notin E$, then by property (iii) of Lemma E.3 every pair of non-adjacent nodes belongs in the same number of independent sets in the family $\{I_f\}_{f \in \mathcal{F}}$. Let L be this number. Following the same strategy as before, we sum over all non-adjacent pairs of nodes and get that

$$\begin{aligned} L \cdot \left(\binom{n}{2} - \frac{nd}{2} \right) &= \sum_{i,j:(i,j) \notin E} |\{f : i, j \in I_f\}| \\ &= \sum_{f \in \mathcal{F}} \binom{|I_f|}{2} \end{aligned}$$

by exchanging the order of summation

$$= |\mathcal{F}| \frac{(s^2 + 1)s^2}{2} ,$$

where in the last step we used property (i) of Lemma E.3. Thus, we have that

$$\begin{aligned} \Pr(A_{i,k} \cap A_{j,k}) &= \Pr(i, j \in I_f, Z_i = k, Z_j = k) \\ &= \Pr(i, j \in I_f) \Pr(Z_i = k, Z_j = k | i, j \in I_f) \\ &= \Pr(i, j \in I_f) \Pr(U_i = k, U_j = k) \\ &= \frac{1}{4} \frac{L}{|\mathcal{F}|} \\ &= \frac{1}{4} \frac{(s^2+1)s^2}{\binom{n}{2} - \frac{nd}{2}} \\ &= \frac{1}{4} \frac{(s^2 + 1)s^2}{n(n-1-d)} \\ &= \frac{1}{4} \frac{(s^2 + 1)s^2}{(s+1)(s^2+1)((s+1)(s^2+1) - 1 - s(s+1))} \\ &= \frac{1}{4} \frac{s^2}{(s+1)s^3} \\ &= \frac{1}{4s(s+1)} . \end{aligned}$$

Therefore,

$$\begin{aligned} \text{Cov}\left(\frac{\mathbf{1}[A_{i,k}]}{\Pr(A_{i,k})}, \frac{\mathbf{1}[A_{j,k}]}{\Pr(A_{j,k})}\right) &= \frac{\Pr(A_{i,k} \cap A_{j,k})}{\Pr(A_{i,k}) \Pr(A_{j,k})} - 1 \\ &= \frac{\frac{1}{4s(s+1)}}{\frac{1}{4(s+1)^2}} - 1 \\ &= \frac{(s+1)^2}{s(s+1)} - 1 \\ &= \frac{s+1}{s} - 1 \\ &= \frac{1}{s} . \end{aligned}$$

□

We are now ready to prove Proposition 5.2, which we also restate here for convenience.

Proposition 5.2. *There exists a sequence $d = \{d^{(n)}\}_{n=1}^\infty$ such that $d = \Theta(n^{2/3})$ and*

$$\inf_{G \in \mathcal{G}_d} \mathcal{R}(G, \tau_{\text{DTE}}, 2) \asymp \frac{\sqrt{d}}{n} .$$

Proof. Assume for simplicity that there is a prime number s such that $n = (s+1)(s^2+1)$ and $d = s(s+1)$. We will show in the end how to extend to arbitrary values of n and d by adding isolated nodes to the graph. Let G be the graph on n nodes with the properties of Lemma E.2.

Let us start from the lower bound. Since G is d -regular, by (i) of Lemma E.1 it follows that every node in the conflict graph $\mathcal{H}$ has degree at least d . Therefore, $d^*(\mathcal{H}) \geq d$. Thus, applying the lower bound of Corollary 1.1 with $q = 2$, we get that

$$\mathcal{R}(G, \tau, 2) \gtrsim \frac{\sqrt{d}}{n} .$$

Now, let us establish the upper bound. This involves finding a statistical procedure, i.e. a pair of design P and estimator $\hat{\tau}$, whose risk is $\mathcal{O}(\sqrt{d}/n)$. We will use the Independent Set Design of Algorithm 5 as our design P . Below, when we write probabilities, it will always be with respect to the randomness of the Independent Set Design.

Our estimator $\hat{\tau}$ will be the one defined in (22). As established in Lemma E.4, this estimator is unbiased for τ . Therefore, it suffices to upper bound the variance of $\hat{\tau}$. To simplify analysis, we break down the estimator into the estimators for the two exposures, i.e. we write $\hat{\tau} = \hat{\tau}_1 - \hat{\tau}_0$, where

$$\begin{aligned} \hat{\tau}_1 &= \frac{1}{n} \sum_{i=1}^n \frac{\mathbf{1}[A_{i,1}]}{\Pr(A_{i,1})} y_i(e_{i,1}) \\ \hat{\tau}_0 &= \frac{1}{n} \sum_{i=1}^n \frac{\mathbf{1}[A_{i,0}]}{\Pr(A_{i,0})} y_i(e_{i,0}) \end{aligned}$$

We have that

$$\begin{aligned} \text{Var}(\hat{\tau}) &= \text{Var}(\hat{\tau}_1 - \hat{\tau}_0) \\ &= \text{Var}(\hat{\tau}_1) + \text{Var}(\hat{\tau}_0) - 2 \text{Cov}(\hat{\tau}_1, \hat{\tau}_0) \\ &\leq \text{Var}(\hat{\tau}_1) + \text{Var}(\hat{\tau}_0) + 2\sqrt{\text{Var}(\hat{\tau}_1) \text{Var}(\hat{\tau}_0)} \\ &\leq 2 \text{Var}(\hat{\tau}_1) + 2 \text{Var}(\hat{\tau}_0) , \end{aligned}$$

where we used the Cauchy-Schwarz and AM-GM inequalities. Therefore, it suffices to upper bound $\text{Var}(\hat{\tau}_1)$ and $\text{Var}(\hat{\tau}_0)$. We focus on $\hat{\tau}_1$, since the analysis for $\hat{\tau}_0$ is identical. We have

$$\begin{aligned} \text{Var}(\hat{\tau}_1) &= \frac{1}{n^2} \text{Var} \left(\sum_{i=1}^n \frac{\mathbf{1}[i \in I_f, Z_i = 1]}{\Pr(i \in I_f, Z_i = 1)} y_i(e_{i,1}) \right) \\ &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \frac{\text{Cov}(\mathbf{1}[i \in I_f, Z_i = 1], \mathbf{1}[j \in I_f, Z_j = 1])}{\Pr(i \in I_f, Z_i = 1) \Pr(j \in I_f, Z_j = 1)} y_i(e_{i,1}) y_j(e_{j,1}) . \end{aligned}$$

Define the vector $\boldsymbol{\alpha} \in \mathbb{R}^n$ as $\alpha_i = y_i(e_{i,1})/\sqrt{n}$ for $i \in [n]$. Then, the preceding variance can be written as a quadratic form

$$\text{Var}(\hat{\tau}_1) = \frac{1}{n} \boldsymbol{\alpha}^\top \mathbf{C} \boldsymbol{\alpha} ,$$

where $\mathbf{C} \in \mathbb{R}^{n \times n}$ is the matrix with entries

$$C_{i,j} = \frac{\text{Cov}(\mathbf{1}[A_{i,1}], \mathbf{1}[A_{j,1}])}{\Pr(A_{i,1}) \Pr(A_{j,1})} .$$

Since $y \in \mathcal{M}(G, \tau, 2)$, we have that

$$\|\boldsymbol{\alpha}\|^2 = \frac{1}{n} \sum_{i=1}^n y_i (e_{i,1})^2 \leq 1 \ .$$

It follows that

$$\text{Var}(\hat{\tau}_1) \leq \frac{1}{n} \|\mathbf{C}\|_2 \ . \quad (23)$$

We have already calculated all entries of $\mathbf{C}$ in Lemma E.5. In particular, if $\mathbf{I} \in \mathbb{R}^{n \times n}$ is the identity matrix, $\mathbf{1} \in \mathbb{R}^n$ is the all 1 vector and $\mathbf{A}_G$ is the adjacency matrix of G , then we can write

$$\begin{aligned} \mathbf{C} &= 2(s+1)\mathbf{I} + \frac{1}{s}(\mathbf{1}\mathbf{1}^\top - \mathbf{I} - \mathbf{A}_G) - \mathbf{A}_G \\ &= \left(2(s+2) - \frac{1}{s}\right)\mathbf{I} - \left(1 + \frac{1}{s}\right)\mathbf{A}_G + \frac{1}{s}\mathbf{1}\mathbf{1}^\top \ . \end{aligned}$$

The matrix $\mathbf{1}\mathbf{1}^\top/s$ has one eigenvalue equal to n/s with eigenvector $\mathbf{1}$ and $n-1$ eigenvalues equal to 0. The matrix $\mathbf{A}_G$ has leading eigenvalue equal to d with eigenvector $\mathbf{1}$. Thus, the absolute value of the eigenvalue corresponding to the eigenvector $\mathbf{1}$ for matrix $\mathbf{1}\mathbf{1}^\top/s - (1 + 1/s)\mathbf{A}_G$ is equal to

$$\left| \frac{n}{s} - \left(1 + \frac{1}{s}\right)d \right| = \left| \frac{(s+1)(s^2+1)}{s} - \left(1 + \frac{1}{s}\right)s(s+1) \right| = \frac{s^2-1}{s} = \mathcal{O}(s) \ .$$

By property (iii) of Lemma E.2, the remaining eigenvalues of $\mathbf{A}_G$ are $s-1$ and $-(s+1)$, which are both $\mathcal{O}(s)$ in absolute value. Thus, we have established that

$$\left\| \frac{1}{s}\mathbf{1}\mathbf{1}^\top - \left(1 + \frac{1}{s}\right)\mathbf{A}_G \right\|_2 = \mathcal{O}(s) \ .$$

We also trivially have that

$$\left\| \left(2(s+2) - \frac{1}{s}\right)\mathbf{I} \right\|_2 = \mathcal{O}(s) \ .$$

Thus, by the triangle inequality we have that $\|\mathbf{C}\|_2 = \mathcal{O}(s)$. Therefore, by (23) we conclude that $\text{Var}(\hat{\tau}_1) = \mathcal{O}(s/n)$. By the same argument, we also have that $\text{Var}(\hat{\tau}_0) = \mathcal{O}(s/n)$. Thus, we have that $\text{Var}(\hat{\tau}) = \mathcal{O}(s/n)$. Since $s = \Theta(\sqrt{d})$, we have that $\text{Var}(\hat{\tau}) = \mathcal{O}(\sqrt{d}/n)$, which establishes the upper bound. $\square$

E.4 Random Regular Graphs (Corollaries 5.3 and 5.4)

A *random d -regular graph* G is a graph chosen uniformly at random from the set of all d -regular graphs on n vertices $\mathcal{G}_{n,d}$. We denote this as $G \sim \mathcal{U}(\mathcal{G}_{n,d})$. In this section, we identify the minimax rate for random d -regular graphs up to a logarithmic factor. We will use the following fact about independent sets in d -regular graphs. In this section, we identify the minimax rate for random d -regular graphs up to a logarithmic factor. We will use the following fact about independent sets in random d -regular graphs. Recall that for a graph G , we denote by $\mathcal{I}(G)$ the largest independent set in G .

Theorem E.6. *Suppose $G \sim \mathcal{U}(\mathcal{G}_{n,d})$. There exist constants $c_1, c_2 > 0$ such that if $c_1 \leq d < n - c_2 n / \log n$, then with probability $1 - o(1)$ as $n \rightarrow \infty$*

$$|\mathcal{I}(G)| = \mathcal{O}\left(\frac{n \log d}{d}\right) .$$

Proof. Frieze and Łuczak (1992) showed that there exists a constant $c_1 > 0$ such that if $d \geq c_1$ and $d = o(n^{1/3})$, then for $G \sim \mathcal{U}(\mathcal{G}_{n,d})$, with probability $1 - o(1)$ as $n \rightarrow \infty$

$$|\mathcal{I}(G)| = \mathcal{O}\left(\frac{2n \log d}{d}\right) .$$

Furthermore, from the discussion after Theorem 5.2 in Krivelevich et al. (2001), it follows there exists a constant $c_2 > 0$ such that for any $\alpha < 1/2$, if $n^\alpha < d < n - c_2 n / \log n$, then with probability $1 - o(1)$

$$|\mathcal{I}(G)| = \mathcal{O}\left(\frac{n \log d}{d}\right) .$$

We can thus choose any $\alpha < 1/3$ and combine these two results to get the desired conclusion. $\square$

We are now ready to prove Corollary 5.3, which we restate here for convenience.

Corollary 5.3. *Suppose G is a uniformly random d -regular graph with $d = o(n)$. Then, with probability at least $1 - o(1)$ over the random choice of G ,*

$$\frac{d}{n \cdot \log d} \lesssim \mathcal{R}(G, \tau_{\text{DTE}}, 2) \lesssim \frac{d}{n} .$$

Proof. The upper bound follows from Corollary 1.1. For the lower bound, we can apply Theorem E.6 to get that with probability $1 - o(1)$ as $n \rightarrow \infty$

$$|\mathcal{I}(G)| = \mathcal{O}\left(\frac{n \log d}{d}\right) .$$

We can then apply the global lower bound of Theorem 3.2 to get that with probability $1 - o(1)$ as $n \rightarrow \infty$

$$R(\mathcal{G}_{n,d}) = \Omega\left(\frac{\log d}{d}\right) ,$$

which completes the proof. $\square$

Finally, we present the proof of Corollary 5.4 for the global effect, which we restate here for convenience.

Corollary 5.4. *Suppose G is a random d -regular graph with $d = o(\sqrt{n})$. Then, with probability at least $1 - o(1)$ over the random choice of G ,*

$$\frac{d}{n} \lesssim \mathcal{R}(G, \tau_{\text{GATE}}, 2) \lesssim \frac{d^2}{n} .$$

Proof. Corollary 1.2 yields that for any graph G ,

$$\frac{\sqrt{d_{\text{avg}}(G^2)}}{n} \lesssim \mathcal{R}(G, \tau_{\text{GATE}}, 2) \lesssim \frac{\lambda(G^2) + 1}{n} .$$

If G is a d -regular graph, then every node has exactly d neighbors, and each of those neighbors has at most $d-1$ other distinct neighbors. Therefore, every node in G^2 has degree at most $d(d-1) + d = d^2$. It follows that $\lambda(G^2) \leq d^2$ for every d -regular graph G , from which the upper bound follows.

For the lower bound, let X_r denote the number of walks of length r in G . Then, if $d_i(G^2)$ denotes the degree of node i in G^2 , we have that

$$\sum_{i=1}^n d_i(G^2) = nd^2 - 6X_3 - 4X_4 . \quad (24)$$

This is because each triangle of nodes a, b, c contributes two “invalid” paths of length 2 for each of the three vertices. For example, for vertex a , the paths $a \rightarrow b \rightarrow c$ and $a \rightarrow c \rightarrow b$ are invalid, since they end up in a 1-hop neighbor of a . Similarly, each 4-cycle of nodes a, b, c, d contributes one invalid path of length 2 for every vertex. For example, for vertex a , the paths $a \rightarrow b \rightarrow c$ and $a \rightarrow d \rightarrow c$ both end up in the same vertex, so one of them is considered invalid.

Using Theorem 3(a) from McKay, Wormald, and Wysocka (2004), we have the following bounds on the expected number of triangles and 4-cycles in a random d -regular graph G , which hold as long as $d = o(\sqrt{n})$.

$$\mathbb{E}[X_3] = \frac{d^3}{6}(1 + o(1)) \quad \text{and} \quad \mathbb{E}[X_4] = \frac{d^4}{8}(1 + o(1)) .$$

By Markov’s inequality and the fact that $d = o(n)$, it follows that $X_3 = o_p(nd^2)$ and $X_4 = o_p(nd^2)$. Therefore, by equation (24), we have that

$$d_{\text{avg}}(G^2) = \frac{1}{n} \sum_{i=1}^n d_i(G^2) = d^2(1 + o_p(1)) ,$$

from which the lower bound follows. □