

Treatment Effects in Market Equilibrium

Evan Munro

Stefan Wager

Kuang Xu

December 31, 2022

Stanford Graduate School of Business

Abstract

When randomized trials are run in a marketplace equilibrated by prices, interference arises. To understand the impact on RCT analysis, we build a stochastic model of treatment effects in equilibrium. We characterize the average direct (ADE) and indirect treatment effect (AIE) asymptotically. A standard RCT can consistently estimate the ADE, but confidence intervals and estimation of the AIE require price elasticity estimates, which we provide using a novel experimental design. We define heterogeneous treatment effects and derive an optimal targeting rule that meets an equilibrium stability condition. We illustrate our results using a freelance labor market simulation and data from a cash transfer experiment.

Keywords: Program Evaluation, Experimental Design, Interference

*We thank Isaiah Andrews, Joshua Angrist, Dmitry Arkhangelsky, Susan Athey, Lanier Benkard, Han Hong, Guido Imbens, Michael Kosorok, Whitney Newey, Fredrik Sävje, Paulo Somaini, Ali Yurukoglu, and seminar participants at Columbia, Harvard, Michigan State University, Princeton, Stanford, UT Austin, University of Chicago, University of Montreal, CODE, ICSDS, JSM, and the Online Causal Inference Seminar for helpful comments and discussions. This research is partially supported by a gift from Facebook. Code for the simulations in this paper is available at <https://github.com/evanmunro/market-interference>.

1 Introduction

Modern computational infrastructure has enabled an enormous growth of randomized trials on websites and mobile applications (Kohavi et al. 2020). On many online platforms, interactions between individuals occur as individuals buy and sell items at the prevailing market prices. The goal of running a standard randomized control trial (RCT), where treatment is assigned with constant probability to all individuals, is often to evaluate the impact of a new policy on some outcome of interest. The standard analysis of treatment effects in a randomized trial relies on an assumption that the treatment status of one individual does not affect the outcomes of other individuals (Fisher 1935, Banerjee & Duflo 2011, Imbens & Rubin 2015), also known as the Stable Unit Treatment Value Assumption (SUTVA). When a randomized trial affects supply or demand, and market prices affect the outcome of interest, then general equilibrium effects lead to a violation of SUTVA, and measurements of the Average Treatment Effect (ATE) do not correspond to a meaningful measure of policy impact (Heckman et al. 1998).

A primary goal of this paper is to provide a framework for the analysis of randomized trials in an equilibrium setting without relying on strong parametric, distributional, or homogeneity assumptions. We integrate a price equilibrium into a stochastic model of treatment effects under interference. We use it to provide asymptotic results for various definitions of policy-relevant average and conditional average treatment effects, and to design new forms of randomized trials and targeting rules. In the process, we overcome a variety of challenges unaddressed by the existing literature on treatment effects under dense patterns of interference, including providing asymptotically exact confidence intervals for the indirect effect, and proposing general definitions of conditional estimands that can be used to construct targeting rules.

There is existing work that addresses interference by clustering units at a higher level at which SUTVA holds (Baird et al. 2018, Hudgens & Halloran 2008). In the market context, this involves finding clusters of buyers and sellers who trade only within the cluster and not across clusters. However, if a marketplace is very connected, finding a good clustering may not be feasible. Furthermore, clustering can be costly, both because it reduces power (Abadie et al. 2017) and because cluster-level designs require modifying the equilibrium in different clusters. In this paper, we focus on settings where clustering is not desirable, so that the solution must rely on unit-level randomized experimentation only. This approach results in increased power, but comes at a cost of making assumptions based on ex-ante knowledge about the structure of the market equilibrium.

Structural modeling is used in a variety of fields to analyze the effects of policies when individuals interact with one another in economic environments (Heckman et al. 1998, Duflo 2004, Abbring & Heckman 2007). Some structural models are population-level and are designed for theoretical analysis, as in Johnson (1980). Others are designed for estimating demand systems from observational data by imposing exclusion restrictions or parametric assumptions; see Akerberg et al. (2007), Matzkin (2015), Berry et al. (1995), Berry & Haile (2021). Compared to observational data that is often at an aggregate level, experiments can be used to collect individual-level data that are much richer. As suggested in Manski (1993), collecting richer data is helpful for analyz-

ing endogenous social effects without imposing strict assumptions on the data generating process. We build a non-parametric structural model that is designed for treatment effect inference using individual-level data from randomized experiments. Our model builds connections between the non-parametric program evaluation literature and the modeling-based approaches that have traditionally been used to analyze equilibrium effects of policies. When marketplace effects are removed, it reduces to a standard Neyman-Rubin potential outcomes model (Imbens & Rubin 2015). We are able to derive estimators, standard errors and targeted policies in equilibrium that are robust to a variety of modeling choices and can be compared directly to related results in the causal inference literature.

The starting point for our model is the potential outcomes framework that allows for cross-unit interference (Hudgens & Halloran 2008, Manski 2013, Aronow & Samii 2017). We assume that each individual’s production and consumption choices are determined by latent supply and demand curves (Angrist et al. 2000, Heckman & Vytlacil 2005). Units interact via a marketplace modeled as a classical multiple goods economy in general equilibrium, where individuals make choices on production and consumption based on their own private information as well as their expectations of what the market price equilibrium will be (Walras 1900, Mas-Colell et al. 1995). Our two main non-parametric structural assumptions are that prices equalize latent supply and demand, and the potential outcomes for each individual depend only on their own treatment and market prices. Under this structure, we show that interference is restricted in that it operates only through the J -dimensional equilibrium price. That is, equilibrium prices induce an exposure mapping in the sense of Manski (2013) and Aronow & Samii (2017). When individual supply, demand and outcomes are stochastic, we show that the random equilibrium price that arises in finite samples converges to a fixed and unique price in the limit as the number of suppliers and consumers grows. We use this mean-field limit to characterize the asymptotic behavior of treatment effects.

Under interference, a meaningful policy-relevant counterfactual is the Global Treatment Effect (GTE), which is the average effect of treating all individuals in the sample compared to treating no individuals in the sample. Hu et al. (2022) show that under any pattern of interference, the sum of the average direct and indirect effects of a binary treatment is a first-order approximation to the GTE. Thus, we start by analyzing the average direct and indirect effects of a binary treatment in equilibrium (Halloran & Struchiner 1995, Hu et al. 2022, Sävje et al. 2021). We show that the direct and indirect effects of a binary treatment converge to simple and interpretable mean-field limits in large samples. The direct effect converges to the difference in expected outcomes for treated and control individuals, holding market prices fixed. The indirect effect converges to the product of the effect of the treatment on excess demand, the derivative of expected outcomes with respect to prices, and the inverse of the sensitivity of excess demand with respect to prices.

The direct effect of the treatment can be consistently estimated using a Horvitz-Thompson estimator and data from a traditional RCT, as shown in Sävje et al. (2021). We also provide a Central Limit Theorem for the Horvitz-Thompson estimator and characterize the asymptotic

variance explicitly. Confidence intervals built for the direct effect using a standard RCT approach that ignores equilibrium effects are not exact. In order to perform asymptotically exact inference for the direct effect and to consistently estimate the indirect effect, which is the missing component of the policy counterfactual, estimates of price elasticities are required.

We propose an augmented unit-level randomized experiment that adds small price perturbations to a traditional Bernoulli randomized experiment. This is feasible, for example, in online marketplaces, where there are often fees for buyers and sellers that can be perturbed randomly. We show that adding small price perturbations does not affect the limiting equilibrium of the system. In previous work that characterizes the indirect effect under dense patterns of interference, estimators for the variance of the indirect effect were not available (Li & Wager 2022, Sävje et al. 2021). Under this augmented randomized experiment, we provide a central limit theorem and characterize the asymptotic variance for the indirect effect, and show that consistent estimators for the direct and indirect effect and their asymptotic variances are available.

Given that the model we’ve been using so far allows for unrestricted heterogeneity across agents, it is also interesting to see whether we can estimate how treatment effects vary with pre-treatment characteristics—and exploit this heterogeneity to learn improved policies. Estimating heterogeneity in individual responses to treatment both on outcomes and on excess demand can benefit a policymaker. For example, subsidies improve supplier welfare directly, but indirectly reduce welfare through decreased market-clearing prices. Characterizing heterogeneity in production responses is helpful in optimizing supplier welfare through targeted interventions.

In Section 4, we analyze heterogeneous effects and optimal treatment rules under equilibrium interference. In settings where SUTVA holds, there is a large literature on characterizing heterogeneity in treatment effects and treatment targeting (Manski 2004, Athey & Imbens 2016, Wager & Athey 2018, Van der Weele et al. 2019). Optimal policy in this setting depends on the sign and magnitude of the conditional average treatment effect (CATE) (Manski 2004). However, in the presence of interference, the CATE alone does not explain who should be treated to optimize outcomes; and in fact the standard definition of the CATE is ill-specified (since it relies on SUTVA). Given this background, our first contribution is to define conditional estimands under general patterns of interference that connect to a meaningful policy counterfactual and are useful in defining optimal targeting rules. We propose the conditional average direct effect (CADE), which is the expected effect of changing the treatment of an individual with covariate x on their own outcomes, and the conditional average indirect effect (CAIE), which is the expected effect of changing the treatment of an individual with covariate x on everyone else’s outcomes. We show that the sum of the CADE and the CAIE, scaled appropriately, is equal to the expected increase in average outcomes for a marginal increase in the treatment probability of an individual with covariate x . Just as the ADE is equal to the ATE when interference is removed, the CADE is equal to the CATE when interference is removed. Under the model of Section 2, the CADE and the CAIE converge to mean-field limits, that depend on a combination of average price elasticities and conditional average direct effects of the treatment on outcomes and excess demand.

The sum of the CADE and the CAIE can be used to estimate the optimal unconstrained treatment rule using the experiment of Section 3, which includes price perturbations. The structure of the conditional estimands under equilibrium interference, however, is such that it is possible to use data from a standard RCT, without price elasticity estimates, to estimate the optimal targeting rule when it is constrained to have the same equilibrium effect as the RCT. The optimal treatment rule takes the form of a hyperplane in the space of the conditional average direct effects of the treatment on both outcomes and excess demand. The gain of this targeting rule over a uniform rule is estimable using a data splitting approach. This provides a metric of heterogeneity under equilibrium interference that is estimable using data from a standard RCT and readily available software for CATE estimation.

Finally, in Section 5, we demonstrate the consistency and coverage properties of our treatment effect and variance estimators using simulations of marketplaces with a supply-side nudge intervention that affect individuals directly and indirectly through market prices. We also include an empirical illustration of the equilibrium-stable targeting rule using data from a large cash transfer experiment in Mexico published by [Gertler et al. \(2012\)](#). The targeting rule results in significantly higher estimated average agricultural income for households in the data sample, when average land usage under the targeting policy is constrained to be equal to that of the uniform policy from the original RCT.

1.1 Related Work

As discussed in the introduction, most existing work on treatment effect estimation under equilibrium effects (or other types of interference) adopts a standard randomization-based framework, where SUTVA holds for higher-level clusters of units ([Baird et al. 2018](#), [Basse et al. 2019](#), [Hudgens & Halloran 2008](#), [Karrer et al. 2021](#), [Liu & Hudgens 2014](#), [Tchetgen & VanderWeele 2012](#)). Another related approach is the network interference model, which posits that interference operates along a network and the connections between units are sparse ([Athey et al. 2018](#), [Leung 2020](#), [Sävje et al. 2021](#)). The sparsity of connections then enables randomization-based methods for studying cluster-randomized experiments to be extended to this setting. In our setting, however, the interference pattern produced by marketplace price effects is dense and simultaneously affects all units, so cluster- or sparsity-based methods are not applicable.

There is very little available work on targeting treatments under interference. One notable exception is [Viviano \(2019\)](#), who estimates treatment allocation rules in a network interference model. However, his results rely on sparsity of the network structure and so they are not directly comparable to ours. Furthermore, [Viviano \(2019\)](#) proceeds via direct empirical welfare maximization over a restricted policy class as in, e.g., [Kitagawa & Tetenov \(2018\)](#) and [Athey & Wager \(2021\)](#), and so he doesn't define CATE-like quantities in his model. We also note recent work by [Sahoo & Wager \(2022\)](#) on policy learning in a model where agents can strategically modify covariates used for treatment assignment and the policymaker has a budget constraint. They show that this setting results in a type of interference, and consider methods that account for this. Again, however,

their approach relies on a form of direct welfare maximization and they do not consider CATE-like quantities.

We use total differentiation of the mean-field equilibrium condition to characterize the sensitivity of the mean-field equilibrium price with respect to the treatment in terms of estimable elasticities and treatment effects. This type of analysis has a long history in economics. In a theoretical analysis, [Johnson \(1980\)](#) uses a low-skilled labor market-clearing condition to characterize the effects of immigration on domestic employment. In non-equilibrium settings, similar techniques are used to express policy effects in terms of estimable elasticities; see [Chetty \(2009\)](#). In this paper, we start with a non-parametric stochastic model of a finite-sample marketplace, which has an equilibrium condition that is not differentiable or amenable to traditional techniques. We show that average effects in this stochastic model have limits which correspond to population quantities that are amenable to “sufficient statistics”-type analysis. This leads to valid inference strategies and an analysis of heterogeneity that are not available when working with a population model directly.

In empirical work with observational data, instrumental variables approaches have been used in both parametric and non-parametric settings to estimate price elasticities or a weighted average of elasticities ([Angrist et al. 1996](#), [Berry & Haile 2021](#)). [Angrist & Krueger \(2001\)](#) reviews IV approaches and their relation to causal inference for observational data. There is a variety of more recent work on using IV approaches to identify causal effects in a variety of settings beyond a price equilibrium, an incomplete list includes [Heckman & Pinto \(2018\)](#), [Chesher & Rosen \(2017\)](#), and [Vazquez-Bare \(2022\)](#). In our paper, we require an estimate of price elasticities to perform inference on average treatment effects. The price elasticities required are an average over all individuals in the market, rather than a weighted average as in the Local Average Treatment Effect literature. Our introduction of price perturbations can be interpreted as creating an ideal instrumental variable for the market price, where by construction everyone in the market is a complier so that the required price elasticities are estimated exactly. In settings where it is not possible to randomize individual-level fees, then with additional assumptions, a more structured IV-based approach could be used instead for estimating elasticities. Given there is already a large existing literature on price elasticity estimation in settings where there is more limited variation available in the data, we do not explore this direction further in this paper.

In development economics, there is concern that interventions can have general equilibrium effects that have a meaningful welfare impact but are not captured by standard randomized control trials. [Banerjee et al. \(2021\)](#) uses a cluster-randomized experiment to show that, in certain remote villages, in-kind food subsidies can have an impact of market prices. Other recent work on general equilibrium effects in development settings includes the structural approach of [Duflo \(2004\)](#) and the cluster-based approach of [Muralidharan et al. \(2017\)](#) and [Egger et al. \(2022\)](#). Our results suggest that it is possible for researchers to evaluate general equilibrium effects without running a cluster-randomized experiment, if the market of interest has a large number of buyers and sellers and convincing estimates of relevant price elasticities are available.

There is some existing work characterizing policy effects in equilibrium using a semi-structural approach. [Wolf \(2019\)](#) uses a semi-structural approach to decompose the effect of macroeconomic shocks into partial equilibrium effects and general equilibrium effects (e.g., price effects), but uses time series methods to estimate the price effects, in contrast to the unit-level approach presented in this paper. [Heckman et al. \(1998\)](#) uses a calibrated overlapping generations model to estimate equilibrium treatment effects in a model of school choice in a labor market equilibrium. [Neilson et al. \(2019\)](#) uses a randomized control trial to estimate the direct effect of information provision on school choice but simulates a fully parametric structural model to evaluate equilibrium effects of the treatment. Compared to this existing work, our approach is unique in that it takes a fully non-parametric approach to analyzing general equilibrium effects based on unit-level experiments that reduces to the standard Neyman-Rubin potential outcomes framework when equilibrium effects are removed.

Our use of latent choice models—specifically latent supply and demand curves—to construct potential outcomes builds on a long tradition in econometrics going back to the work of [Roy \(1951\)](#) and [Heckman \(1979\)](#) on endogenous selection models. [Heckman & Vytlacil \(2005\)](#) offer a more recent synthesis line of work, and use it to derive instrumental variable estimators for a number of policy-relevant variables. The specific latent choice model used here is most closely related to that used in [Angrist et al. \(2000\)](#), who used it to characterize the behavior of two-stage least squares estimators when used to estimate supply and demand elasticities.

Our approach relies on positing a stochastic model where the types of individuals are generated by an underlying probability distribution, and analyzing effects in the limit as the number of individuals in the model grows large. This approach is inspired by mean-field modeling. There is a large literature that uses the mean-field limit to study large-scale stochastic systems in operations research, engineering and physics ([Mézard et al. 1987](#), [Tsitsiklis & Xu 2013](#), [Vvedenskaya et al. 1996](#)). Also related are economic models where players respond to the average behavior in the system ([Weintraub et al. 2008](#), [Hopenhayn 1992](#), [Krusell & Smith 1998](#)).

We also note a handful of recent papers that leverage stochastic modeling and mean-field type results in a causal inference setting. [Johari et al. \(2020\)](#) use a mean-field model to examine the benefits of a two-sided randomization design in two-sided market platforms and quantify the bias due to interference between market participants. [Li & Wager \(2022\)](#) study large sample results for treatment effect estimation under network interference, where units are placed at vertices of an “exposure graph” and a unit’s potential outcome may depend on another unit’s treatment if and only if they are connected by an edge in the graph; they further posit that the exposure graph is randomly generated from a graphon model. Finally, [Wager & Xu \(2021\)](#) consider a model of a marketplace that matches exogenous demand to endogenous supply, with interference via supply cannibalization. They propose using zero-mean perturbations to estimate a revenue gradient and optimize a continuous decision variable dynamically. Our work also uses zero-mean perturbations as part of a unit-level experiment, but is otherwise distinct, given our focus is on treatment effects of a binary intervention in a general price equilibrium.

2 Treatment Effects under Market Interference

2.1 Treatment Effects Under Interference

We first review a definition of treatment effects in the potential outcomes framework when individuals are sampled from a population, which will allow us to model general interference patterns between units. Consider a system comprised of $i = 1, \dots, n$ individuals who are sampled independently from a population, where individual i has features $X_i \in \mathcal{X}$. We will work in a setting where there is a binary treatment $W_i \in \{0, 1\}$ assigned to each individual with Bernoulli probability $\pi_i = \pi(X_i)$. The treatment allocation function is $\pi(\cdot) : \mathcal{X} \rightarrow [0, 1]$. In this paper, we will study both a “standard” RCT where the treatment probability $\pi_i = q$ is constant for every individual, and more general settings where the treatment probability can vary with pre-treatment covariates. We assume that overlap holds.

Assumption 1. *Treatments are sampled from a Bernoulli distribution, $W_i \sim \text{Bernoulli}(\pi(X_i))$. For every $x \in \mathcal{X}$, $0 < \pi(x) < 1$.*

The vector of treatments for a sample of n individuals is \mathbf{W} . The effect of treatment vector on an individual’s outcome, conditional on all other randomness in the system, is captured by the notion of potential outcomes. We will define the potential outcomes for the sample as the family of mappings $\{Y_{in}(\cdot)\}_{i=1, \dots, n}$, such that $Y_{in}(\mathbf{W})$ is the outcome associated with individual i under the treatment assignment \mathbf{W} . Under interference, potential outcomes for a vector of treatments \mathbf{W} are only well-defined conditional on the realization of the sample of individuals, which is indicated by the subscript n . Since both the sample and the treatment vector are drawn randomly, $Y_{in}(\mathbf{W})$ is a random variable that depends on both the treatment and the characteristics of other individuals with $j \neq i$ in the sample. Note that this formulation is in contrast to the setting under SUTVA where an individual’s potential outcome $Y_i(W_i)$ depends only on their own treatment W_i , and does not depend on the treatments or characteristics of other individuals in the sample (Hudgens & Halloran 2008, Manski 2013, Aronow & Samii 2017).

Let $\mathbb{E}_\pi[A(\mathbf{W})]$ be the expectation of A over the random treatment vector \mathbf{W} , conditional on all other sources of randomness in the variable $A(\mathbf{W})$. We can define the average direct effect of a binary treatment as the average marginal effect on an individual’s outcome by changing their own treatment:

$$\tau_{\text{ADE}} = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\pi[Y_{in}(W_i = 1; \mathbf{W}_{-i}) - Y_{in}(W_i = 0; \mathbf{W}_{-i})].$$

Here, the notation $(W_i = a; \mathbf{W}_{-i})$ represents the random vector we obtain by first sampling \mathbf{W} , and then setting its i -th coordinate to a fixed value a . Next, we define the average indirect effect of a binary treatment to be the average marginal effect on everyone else’s outcomes by changing an individual’s treatment:

$$\tau_{\text{AIE}} = \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} \mathbb{E}_\pi[Y_{jn}(W_i = 1; \mathbf{W}_{-i}) - Y_{jn}(W_i = 0; \mathbf{W}_{-i})].$$

Note also that τ_{ADE} and τ_{AIE} depend on potential outcomes $\{Y_{in}(\cdot)\}$, which are stochastic, and therefore are random variables themselves. Furthermore, they depend on the treatment allocation rule $\pi(\cdot)$, so the ADE for one treatment rule may differ from the ADE under a treatment distribution induced by a different allocation rule. The direct effect τ_{ADE} defined above is a standard quantity in the causal inference literature (e.g., [Halloran & Struchiner 1995](#), [Sävje et al. 2021](#)). The indirect effect τ_{AIE} is an analogue to the τ_{ADE} designed to capture the marginal effect of treating i -th unit on other units. [Hu et al. \(2022\)](#) discuss both estimands at length, and in particular show that in a number of parametric models for interference, τ_{ADE} and τ_{AIE} capture natural parametric notions of interference. Furthermore, the sum of the two effects corresponds to a meaningful policy counterfactual for any treatment rule $\pi(\cdot)$. This is the effect of a marginal increase in the treatment probability for each individual on average outcomes:

$$\tau_{\text{INF}} = \frac{1}{n} \sum_{i=1}^n \sum_{k=1}^n \frac{\partial}{\partial \pi_k} \mathbb{E}_{\pi}[Y_{in}(\mathbf{W})] = \tau_{\text{ADE}} + \tau_{\text{AIE}}.$$

This is a first order approximation to the Global Total Effect (GTE), which is the average effect of treating everyone compared to treating nobody on outcomes:

$$\tau_{\text{GTE}} = \frac{1}{n} \sum_{i=1}^n Y_{in}(\mathbf{1}) - Y_{in}(\mathbf{0}).$$

τ_{INF} is the relevant effect for a policymaker making the decision of whether or not to rollout a treatment to a slightly larger fraction of the population. τ_{GTE} is the relevant effect for making the decision to rollout treatment to everybody. Without interference, $\tau_{\text{INF}} = \tau_{\text{GTE}} = \tau_{\text{ATE}}$ and is estimable using data from an RCT where treatment is assigned with constant probability.

Without further restrictions, the ADE and AIE can be difficult to analyze or estimate in a general causal inference model. The goal of this work is to show that we can both characterize and estimate τ_{ADE} , τ_{AIE} and τ_{INF} by leveraging the fact that interference is mediated via equilibrium prices.

2.2 Potential Outcomes in a Stochastic General Equilibrium Model

In this section, we provide a stochastic model of potential outcomes in a market economy. The market contains a total of J types of goods that individuals may buy or sell (an example of our setting is a two-sided market, where some individuals only sell goods and others only buy). The supply and demand of various goods are coordinated via a price equilibrium that approximately clears the market.

We first describe how we can add some structure to the potential outcomes model introduced earlier to capture individuals' interactions in a market equilibrium. The assumption that we make is that all dependence among individuals' outcomes is captured by an equilibrium price. In the language of the literature on treatment effects under interference, this means that a market equilib-

rium statistic along with an individual's treatment defines an exposure for each individual (Manski 2013, Aronow & Samii 2017).

Assumption 2. Interference Operates Through Market-Clearing Prices. *Let two independent samples from the population be identified by subscripts n and n' . If $W_i = W'_i$ and $P_n(\mathbf{W}) = P_{n'}(\mathbf{W}')$, then $Y_{in}(\mathbf{W}) = Y_{in'}(\mathbf{W}')$.*

In other words, individual j influences outcomes of individual i only through their effect on the equilibrium that the market reaches, and not directly through peer effects or some other network interference mechanism. This assumption captures market models where individuals interact through a Nash equilibrium, and choices depend on an individual's (correct) expectation of market prices, rather than the specific strategies of other individuals in the game.

Assumption 2 implies that potential outcomes in our model can be expressed as $Y_{in}(\mathbf{W}) = Y_i(W_i, P_n(\mathbf{W}))$, where the function $Y_i(w, \mathbf{p})$ does not depend on the realized sample. It captures the i -th unit's response to its own treatment and to market prices.

We are now ready to describe the mechanism that generates the equilibrium price, $P_n(\mathbf{W})$. In addition to a outcome function $Y_i(w, \mathbf{p})$, we assume that each individual has a J -dimensional excess demand function $\mathbf{Z}_i(w, \mathbf{p}) = \mathbf{D}_i(w, \mathbf{p}) - \mathbf{S}_i(w, \mathbf{p})$, which also depends on an individual's treatment and market prices. $\mathbf{D}_i(w, \mathbf{p})$ is a demand function and $\mathbf{S}_i(w, \mathbf{p})$ is a supply function. For each individual, their outcome, excess demand functions, and characteristics $\{Y_i(\cdot), \mathbf{Z}_i(\cdot), \mathbf{X}_i\}$ are sampled i.i.d. from some distribution Θ . $P_n(\mathbf{W})$ is the J -dimensional price that ensures supply and demand are approximately matched in the resulting stochastic and non-parametric general equilibrium model.

$$P_n = P_n(\mathbf{W}) = \arg \min_{\mathbf{p}} \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p}) \right\|. \quad (1)$$

For any vector of treatment assignments \mathbf{W} , this defines a potential market price $P_n(\mathbf{W})$ conditional on the realization of the excess demand functions for each individual in the sample. The observed market clearing-price is a random variable that depends on the set of realized excess demand functions $\{\mathbf{Z}_i(\cdot)\}_{i=1}^n$ and the treatment vector \mathbf{W} . Realized outcomes and excess demand are

$$Y_i = Y_i(W_i, P_n(\mathbf{W})), \quad \mathbf{Z}_i = \mathbf{Z}_i(W_i, P_n(\mathbf{W})).$$

Up until this point, the model corresponds to a stochastic Walrasian general equilibrium model (Mas-Colell et al. 1995) with heterogeneous agents and very few assumptions on agent behavior.

In our equilibrium model, we can now write τ_{ADE} and τ_{AIE} as

$$\begin{aligned} \tau_{\text{ADE}} &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\pi} [Y_i(W_i = 1, P_n(W_i = 1; \mathbf{W}_{-i})) - Y_i(W_i = 0, P_n(W_i = 0; \mathbf{W}_{-i}))], \\ \tau_{\text{AIE}} &= \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} \mathbb{E}_{\pi} [Y_j(W_j, P_n(W_i = 1; \mathbf{W}_{-i})) - Y_j(W_j, P_n(W_i = 0; \mathbf{W}_{-i}))]. \end{aligned} \quad (2)$$

In the next section, we will use this stochastic model to examine the behavior of the equilibrium as the market size grows large, which is a key step towards characterizing treatment effects.

2.3 Mean-Field Prices

In this section, we introduce an important concept in our analysis, the mean-field equilibrium price, and demonstrate that the finite sample market-clearing price converges to it at an appropriate rate as the market size n grows. To do so, we first set out several regularity assumptions that will ensure that the mean-field equilibrium price is unique, differentiable, and that the convergence to the mean-field price occurs. This requires some additional notation. In the previous section, we described our sampling model, which is that $\{Y_i(\cdot), \mathbf{Z}_i(\cdot), X_i\}$ are drawn i.i.d. from a distribution Θ . These functions are in a class, which we define as $Y_i(\cdot) \in \mathcal{F}^y$ and $\mathbf{Z}_i(\cdot) \in \mathcal{F}^z$. An equivalent representation for the random functions is $Y_i(w, p) = f^y(w, p, \theta_i)$ and $\mathbf{Z}_i(w, p, \theta_i) = f^z(w, p, \theta_i)$ for random θ_i .

Assumption 3. Regularity at a Sample Level.

1. \mathcal{F}^z is such that the equilibrium price is always an element of some compact set $\mathcal{S} \subset \mathbb{R}^J$.
2. \mathcal{F}^z and \mathcal{F}^y are uniformly bounded.
3. $\{f^y(w, \mathbf{p}, \theta) : \mathbf{p} \in \mathcal{S}\}$ and $\{f^z(w, p, \theta) : \mathbf{p} \in \mathcal{S}\}$ are each a universal Donsker class of functions.
4. For any $p \in \mathcal{S}$, $Pr(Y_i(w, \mathbf{p}) \text{ continuous at } \mathbf{p}) = 1$ and $Pr(\mathbf{Z}_i(w, \mathbf{p}) \text{ continuous at } \mathbf{p}) = 1$.
5. In expectation, the equilibrium price sets aggregate excess demand close to zero. For each $j \in \{1, \dots, J\}$,

$$\mathbb{E} [(Z_{nj}(\mathbf{P}_n(\mathbf{W})))^4] = o\left(\frac{1}{n^2}\right),$$

$$\text{where } Z_{nj}(p) = \frac{1}{n} \sum_{i=1}^n Z_{ij}(W_i, \mathbf{p}).$$

The first regularity assumption indicates that we can restrict attention to a compact set in \mathbb{R}^J when looking for an equilibrium price. A sufficient condition for this to hold is that for all vectors of prices larger than some upper bound, then all excess demand functions in \mathcal{F}^z are weakly negative, and for all prices smaller than some lower bound, then all excess demand functions in \mathcal{F}^z are weakly positive. The second assumption rules out infinite potential outcomes, supply or demand. The third assumption says that for any probability distribution F governing the random treatments and the random functions, then the function classes indexed by the price p are F -Donsker. This is a weak assumption that does not rule out typical choices of outcome or excess demand function. Sufficiently smooth functions, such as Lipschitz functions in 1 dimension, are a Donsker class, as are functions with some discontinuities, including step functions and monotonic functions. More generally, [Benkeser & Van Der Laan \(2016\)](#) indicates the class of all cadlag multivariate functions of bounded variation are Donsker classes; see [van der Vaart & Wellner \(1997\)](#) and [van der Vaart \(1998\)](#)

for a variety of other examples. Continuity almost everywhere requires that for any given market price, the set of outcome or excess demand functions that are discontinuous there must have 0 probability. If demand functions are step functions, for example, then the location of the step must have a continuous distribution. We don't require that the finite sample price exactly sets excess demand equal to zero; the final assumption is a restriction on the approximation error of the finite sample price. This assumption implies both that $\mathbb{E}[Z_{nj}(p)^2] = o(\frac{1}{n})$ and that $Z_n(P_n) = o_p(n^{-0.5})$.

We next define expected outcome and excess demand functions for individuals in treatment and control. For $w \in \{0, 1\}$,

$$y(w, \mathbf{p}) = \mathbb{E}[Y_i(w, \mathbf{p})], \quad z(w, \mathbf{p}) = \mathbb{E}[Z_i(W_i, \mathbf{p})].$$

The mean-field market price is defined as the price that sets the expected excess demand to zero when the treatment is allocating according to π :

$$\mathbf{p}_\pi^* = \{\mathbf{p} : \mathbb{E}[Z_i(W_i, \mathbf{p})] = 0\}.$$

We then impose some additional assumptions on the population outcome and excess demand functions.

Assumption 4. Regularity at a Population Level.

1. $y(w, \mathbf{p})$, $z(w, \mathbf{p})$, and $\mathbb{E}[\mathbf{V}_i(w, \mathbf{p})\mathbf{V}_i(w, \mathbf{p})^\top]$ are twice continuously differentiable in \mathbf{p} for all $\mathbf{p} \in \mathcal{S}$ and $w \in \{0, 1\}$ with bounded derivatives, where $\mathbf{V}_i(W_i, \mathbf{p})$ is a column vector stacking $Y_i(w, \mathbf{p})$ and $Z_i(W_i, \mathbf{p})$.
2. For all $\mathbf{p} \in \mathcal{S}$, $\|\nabla_{\mathbf{p}}[\mathbb{E}[Z_i(W_i, \mathbf{p})] + \mathbf{p}]\| < 1$.
3. The $J \times J$ Jacobian $\boldsymbol{\xi}_z = \nabla_{\mathbf{p}}\mathbb{E}[Z_i(W_i, \mathbf{p}_\pi^*)]$ is full rank.
4. The $J \times J$ matrix $\mathbb{E}[Z_i(w, \mathbf{p}_\pi^*)Z_i(w, \mathbf{p}_\pi^*)^\top]$ is full rank.

While we allow for some non-differentiability at the individual level, the first assumption indicates that we require more smoothness in the moments of the outcomes and excess demand functions. Next, in order for our policy effects to be well defined, we require the mean-field price to be unique and to exactly clear the market. In a single good market, this would only require continuity and strict monotonicity of $\mathbb{E}[Z_i(W_i, \mathbf{p})]$ in p . However, in a market with multiple goods, uniqueness and existence of the equilibrium price can be shown by ensuring the mean-field excess demand function is a contraction. Contraction mapping approaches are commonly used to prove equilibrium existence and uniqueness in models of strategic behavior; see for example [Cornes et al. \(1999\)](#) and [Van Long & Soubeyran \(2000\)](#). Other approaches that rely on more primitive assumptions are also possible; under a zero-degree homogeneity assumption in excess demand then a gross-substitutes condition ensures uniqueness ([Arrow & Hahn 1971](#)). The last two parts of Assumption 4 ensure that the asymptotic distribution of the finite sample equilibrium price is well

defined. The proof of Proposition 1, given in Appendix A.2, is an application of the Banach fixed point theorem.

Proposition 1. *Assumption 4 implies the existence and uniqueness of \mathbf{p}_π^* .*

We can now show that as the number of participants in the market grows large, the sample equilibrium price converges to the unique mean-field price; furthermore, this convergence occurs at a \sqrt{n} rate. The proofs of Theorem 2, given in Appendix A.3, rely on general results for Z -estimators discussed in, e.g., van der Vaart (1998). Agarwal & Somaini (2018) use a similar technique to characterize the asymptotic distributions of equilibrium cutoffs in school choice mechanisms.

Theorem 2. *Under Assumptions 2-4, the finite sample equilibrium price converges in distribution to the mean-field price and is asymptotically normal:*

$$\sqrt{n}(\mathbf{P}_n - \mathbf{p}_\pi^*) \rightarrow_d N(0, \boldsymbol{\xi}_z^{-1} \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^*) \mathbf{Z}_i(W_i, \mathbf{p}_\pi^*)^\top] \boldsymbol{\xi}_z^{-1}).$$

2.4 Mean-Field Treatment Effects

In this section, we show that the random variables τ_{ADE} and τ_{AIE} , defined in (2), converge to meaningful mean-field quantities. These mean-field quantities are useful for two reasons. First, they are interpretable, and provide some insight on how interference occurs in a market, and in what settings the indirect effect will be stronger or weaker. Second, they will provide the foundation for our estimation strategy in Section 3. The following theorem is one of our main results and is proved in Appendix A.5.

Theorem 3. *Under Assumptions 1-4, the average direct, and average indirect effect converge in probability to interpretable mean-field quantities:*

$$\begin{aligned} \tau_{ADE} &\rightarrow_p y(1, \mathbf{p}_\pi^*) - y(0, \mathbf{p}_\pi^*) = \tau_{ADE}^*, \\ \tau_{AIE} &\rightarrow_p -\boldsymbol{\xi}_y^\top \boldsymbol{\xi}_z^{-1} [z(1, \mathbf{p}_\pi^*) - z(0, \mathbf{p}_\pi^*)] = \tau_{AIE}^*, \end{aligned} \tag{3}$$

where $\boldsymbol{\xi}_y = \nabla_p \mathbb{E}[Y_i(W_i, \mathbf{p}_\pi^*)]$ is a $J \times 1$ vector.

For the τ_{ADE} result, a key component of the proof is to leverage the convergence of the finite sample equilibrium price to the mean-field price (Theorem 2). We use it to show that the agents' outcomes become asymptotically independent as n grows, which subsequently allows us to evoke the law of large numbers. Showing convergence for τ_{AIE} is more challenging; the proof relies on an asymptotically linear representation of $\mathbf{P}_n(\mathbf{W})$ and concentration results for Donsker classes of functions.

By imposing some non-parametric and stochastic assumptions on the structure of interference, we have shown that the complex expression for the indirect effect in (2) converges to a simple mean-field expression. Our first observation is that the direct effect converges to the difference in the mean-field outcome function evaluated at $W_i = 1$ and $W_i = 0$. In the next section, we show that

this can be estimated with a differences in means estimate. This finding is in line with [Sävje et al. \(2021\)](#), who show that estimators that target the average treatment effect under no-interference settings generally recover the average direct effect under interference. However, our results will go beyond those of [Sävje et al. \(2021\)](#), since we provide a central limit theorem and asymptotically exact confidence intervals for the direct effect.

The indirect effect, which is a key component of the policy counterfactual τ_{INF} , is not estimable using variation in treatment only. The mean-field indirect effect is the product of the gradient of the expected outcome function with respect to prices, and the treatment effect on excess demand scaled by the gradient of excess demand with respect to prices. By the implicit function theorem, this is equivalent to the product of the gradient of the outcome function with respect to prices and the sensitivity of market prices to raising everyone’s treatment assignment probability. In settings where the treatment impacts market prices through excess demand, and outcomes are sensitive to market prices, then interference effects and the indirect effect are stronger. In the next section, we will show how each component of τ_{ADE}^* and τ_{AIE}^* can be estimated using unit-level experiments.

3 Estimation and Inference

In this section we analyze various estimators of treatment effects that use data from unit-level randomized experiments. We first derive the limiting distribution of a Horvitz-Thompson estimator when data is generated from a standard RCT where $\pi(x) = q$ for all $x \in \mathcal{X}$. This estimator is consistent for the Average Direct Effect, but its asymptotic variance depends on price-elasticity terms that cannot be estimated using data from an RCT that randomizes treatment only. We show that if confidence intervals are constructed using a variance estimator based on the asymptotic variance without price interference, then coverage will not be asymptotically exact.

In order to perform inference for the direct effect and to perform estimation and inference for the indirect effect, we introduce an augmented randomized trial, with Bernoulli randomization for the treatment and local perturbation of market prices to estimate price elasticities. The estimator for the ADE remains unchanged from the standard RCT. Algorithmically, the estimator for the AIE looks like a combination of a Horvitz-Thompson estimator and a weak instrumental variables estimator. Under the augmented randomized experiment, we show how to construct asymptotically valid confidence intervals for both effects.

3.1 Estimation for the Direct Effect in a Standard RCT

In this section, we derive and analyze the asymptotic distribution of differences-in-means types estimators that use data from a standard RCT, described below.

Design 1. *Standard RCT.* $W_i \sim \text{Bernoulli}(q)$. q is the treatment probability, which is constant.

We use the notation \mathbf{p}_q^* to represent the mean-field equilibrium price when treatment is assigned according to Definition 1. Following a number of recent papers, including [Sävje et al. \(2021\)](#) and

Li & Wager (2022), we consider estimating the direct effect using a Horvitz-Thompson Estimator:

$$\hat{\tau}_{ADE} = \hat{\tau}_y^{HT} = \frac{1}{n} \sum_{i=1}^n \frac{W_i Y_i}{q} - \frac{(1 - W_i) Y_i}{1 - q}. \quad (4)$$

Theorem 4. *Under the assumptions of Theorem 3, when treatment is allocated according to Design 1, we have that*

$$\sqrt{n} (\hat{\tau}_{ADE} - \tau_{ADE}^*) \rightarrow_d N(0, \sigma_D^2),$$

where

$$\begin{aligned} \sigma_D^2 &= \text{Var} \left[\left(\frac{W_i}{q} - \frac{1 - W_i}{1 - q} \right) Y_i(W_i, \mathbf{p}_q^*) + A_i \right] \\ A_i &= -\nabla_p [y(1, \mathbf{p}_q^*) - y(0, \mathbf{p}_q^*)]^\top \boldsymbol{\xi}_z^{-1} \mathbf{Z}_i(W_i, \mathbf{p}_q^*) \end{aligned} \quad (5)$$

Our first result, proven in Appendix A.7, is that—even under interference that occurs through market prices—the Horvitz-Thompson estimator converges to the mean-field direct effect at a \sqrt{n} rate and is asymptotically normal. In other words, mirroring the finding of Sävje et al. (2021), we see that an analyst who ignored marketplace effects and used a Horvitz-Thompson estimator designed for RCT with SUTVA would still consistently target a well defined estimand, namely the ADE. However, unlike in an RCT with SUTVA, this estimand no longer captures the total effect of the intervention.

Furthermore, the asymptotic variance of the Horvitz-Thompson estimator in our model has an additional term due to interference effects. Specifically, this additional term is due to stochastic fluctuations in $\mathbf{P}_n - \mathbf{p}_q^*$, and depends on price elasticities. The asymptotic variance of the Horvitz-Thompson estimator in the no-interference case is $\text{Var} \left[\left(\frac{W_i}{q} - \frac{1 - W_i}{1 - q} \right) Y_i(W_i) \right]$. A consistent estimator for this is $\widehat{\text{Var}} \left[\left(\frac{W_i}{q} - \frac{1 - W_i}{1 - q} \right) Y_i \right]$, where

$$\widehat{\text{Var}}[Y_i] = \frac{1}{n} \sum_{i=1}^n (Y_i - \bar{Y})^2$$

is the empirical variance for n -length vector \mathbf{Y} and \bar{Y} is the empirical mean. If we use the no-interference variance estimator in settings with marketplace interference to construct confidence intervals, then the confidence intervals are not asymptotically exact unless the treatment does not impact the derivative of expected outcomes with respect to prices. Whether they under or overcover depends on the sign of the covariance between $\left(\frac{W_i}{q} - \frac{1 - W_i}{1 - q} \right) Y_i$ and A_i . In order to construct confidence intervals that are always asymptotically exact, we require estimates of price sensitivity of outcomes and excess demand. In the next section, we develop an augmented randomized experiment that allows us to estimate these price sensitivities and to estimate τ_{AIE} , which is not estimable with treatment randomization only.

3.2 Augmented Randomized Experiment

In order to estimate the AIE or the build confidence intervals for the ADE, we need to get a handle on price elasticities. Here, we do so by considering an augmented experiment that introduces small, unit level price perturbation into our randomization design. The excess demand and outcome processes remain the same as in the previous section, but prices are shifted at an individual level by U_i for individual i . This affects the equilibrium in finite samples. The equilibrium price now depends on the realized sample of n individuals, the vector of treatments \mathbf{W} , and a vector of price perturbations \mathbf{U} :

$$\tilde{P}_n = P_n(\mathbf{W}, \mathbf{U}) = \arg \min_{\mathbf{p}} \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p} + \mathbf{U}_i) \right\|.$$

Under the perturbed equilibrium price, the realized outcomes and excess demand are

$$Y_i = Y_i(W_i, P_n(\mathbf{W}, \mathbf{U}) + U_i), \quad \mathbf{Z}_i = \mathbf{Z}_i(W_i, P_n(\mathbf{W}, \mathbf{U}) + U_i).$$

In order to estimate price elasticities, we restrict ourselves to price experiments that meet Definition 1 below. These are price experiments which are unobtrusive in the sense that they don't disturb the price equilibrium in the limit. Although we allow some experimentation with price, we cannot simply randomize prices globally and estimate the entire excess demand function, for example. There are two reasons for this. First, while it may be possible to make very small changes to individual-level fees on a platform in order to estimate price elasticities, it would be too costly in terms of customer experience to run global price experiments where prices fluctuated more widely. Second, from a technical perspective, extending our results on estimation for the direct effect to the setting with price experimentation requires that we measure outcomes close to the mean-field equilibrium where there is no price experimentation present.

Definition 1. Unobtrusive Experiment. The distribution of U_i is such that $\sup |U_i| \rightarrow_p 0$ and the difference between $P_n(\mathbf{W})$ and $P_n(\mathbf{W}, \mathbf{U})$ converges in probability to 0 at a \sqrt{n} rate,

$$\sqrt{n}(P_n(\mathbf{W}, \mathbf{U}) - P_n(\mathbf{W})) \rightarrow_p 0. \tag{6}$$

A price experiment that follows Definition 1 generates a finite sample market price with the same asymptotic distribution as $P_n(\mathbf{W})$. Without any further assumptions on outcomes or excess demand, this condition limits our experimentation to mean-zero perturbations, whose size decreases with sample size.

Design 2. Augmented Randomized Experiment

1. For $i \in [n]$, W_i is drawn i.i.d. from $\{0, 1\}$ with $\mathbb{P}(W_i = 1) = q$.
2. For $i \in [n]$, $j \in [J]$, $U_{i,j}$ is drawn i.i.d. and uniformly at random from $\{-h_n, +h_n\}$.
3. The perturbation size decreases as the sample size increases. $h_n = cn^{-\alpha}$ with $\frac{1}{4} < \alpha < \frac{1}{2}$.

The following Proposition shows that Theorem 2 and Theorem 4 still hold in the perturbed system, which implies that the augmented price experiment introduced meets Definition 1.

Proposition 5. *Suppose we have a marketplace governed by Assumptions 1–4 and an experiment following Design 2 is run, where we replace Assumption 3.5 with an appropriate assumption for the perturbed system: for each $j \in \{1, \dots, J\}$, assume*

$$Z_{nj}(\mathbf{P}_n(\mathbf{W}, \mathbf{U})) = o_p(n^{-0.5}),$$

where $Z_{nj}(p) = \frac{1}{n} \sum_{i=1}^n Z_{ij}(W_i, \mathbf{p} + \mathbf{U}_i)$.

Then, Equation 6 holds, which implies the market price in the perturbed system converges to the same distribution as the market price $\mathbf{P}_n(\mathbf{W})$ in the unperturbed system. Furthermore, the distribution of the Horvitz-Thompson estimator for the direct effect is unchanged:

$$\sqrt{n}(\hat{\tau}_{ADE} - \tau_{ADE}^*) \rightarrow_d N(0, \sigma_D^2),$$

where σ_D^2 is defined in Theorem 4.

The mean-field equilibrium under Design 2 has price \mathbf{p}_q^* , which is unchanged from the equilibrium under Design 1. We now develop a point estimator for the indirect effect. Our proposal consists of a combination of Horvitz-Thompson estimators and estimators of price elasticities. The strategy is plug-in estimation based on the functional form for τ_{AIE}^* derived in Theorem 3. The estimator is

$$\hat{\tau}_{AIE} = -\hat{\gamma}^\top \hat{\tau}_z^{HT}. \quad (7)$$

Let \mathbf{Y} be the n -length vector of outcomes, \mathbf{U} is the $n \times J$ matrix of price perturbations, and \mathbf{Z} is the $n \times J$ matrix of excess demand observations. Then, $\hat{\gamma} = (\mathbf{U}^\top \mathbf{Z})(\mathbf{U}^\top \mathbf{Y})$ is a $J \times 1$ vector that estimates $\boldsymbol{\xi}_y^\top \boldsymbol{\xi}_z^{-1}$. The direct effects of the treatment on excess demand are estimated via a Horvitz-Thompson estimator,

$$\hat{\tau}_z^{HT} = \frac{1}{n} \sum_{i=1}^n \frac{W_i \mathbf{Z}_i}{q} - \frac{(1 - W_i) \mathbf{Z}_i}{1 - q}. \quad (8)$$

Our next result, proved in Appendix A.7, provides a central limit theorem for our indirect effect estimator. This estimator converges at a rate that depends on the magnitude of the price perturbations h_n , and is always slower than the \sqrt{n} -rate obtained for the direct effect. In deriving this result, it is helpful to note that algorithmically $\hat{\tau}_{AIE}$ is the product of a weak instrumental variables estimator and a differences in mean estimator. We can then use standard results on the asymptotic variance of IV estimators to guide our analysis.

Theorem 6. *Under the assumptions of Theorem 3 and under Design 2, the average indirect effect*

has an asymptotically normal distribution:

$$\sqrt{nh_n}(\hat{\tau}_{AIE} - \tau_{AIE}^*) \rightarrow_d N(0, \sigma_I^2),$$

where

$$\sigma_I^2 = \mathbb{E} \left[(Y_i(W_i, \mathbf{p}^*) - \mathbf{Z}_i(W_i, \mathbf{p}^*)^\top \boldsymbol{\gamma})^2 \boldsymbol{\nu}^\top \boldsymbol{\nu} \right],$$

$$\boldsymbol{\gamma} = \boldsymbol{\xi}_z^{-1} \boldsymbol{\xi}_y \text{ and } \boldsymbol{\nu} = -\boldsymbol{\xi}^z [z(1, \mathbf{p}_q^*) - z(0, \mathbf{p}_q^*)].$$

Under general patterns of interference, inference on the indirect effect is challenging (Sävje et al. 2021, Li & Wager 2022). Under dense patterns of interference, there are not consistent estimators for the variance of the indirect effect available in previous work. There are two reasons why our paper overcomes this difficulty. Although our interference pattern is dense, it is structured in that all interference happens through the market price, and the market price forms by satisfying a score condition. This structure leads to an analytical functional form for the variance of the indirect effect, provided in Theorem 6. Second, with data from a richer randomized experiment that includes small price perturbations, we are able to estimate each component of this variance, and the variance for the indirect effect, described in Equations 9 and 10.

$$\begin{aligned} \hat{\sigma}_D^2 &= \widehat{\text{Var}}[G_i], \\ G_i &= \frac{W_i}{q} Y_i - \frac{1 - W_i}{1 - q} Y_i - [\hat{\boldsymbol{\xi}}_{y1} - \hat{\boldsymbol{\xi}}_{y0}]^\top \hat{\boldsymbol{\xi}}_z^{-1} \mathbf{Z}_i, \end{aligned} \quad (9)$$

where $\hat{\boldsymbol{\xi}}_{y1}$ and $\hat{\boldsymbol{\xi}}_{y0}$ are $J \times 1$ vectors estimated from regressions of treated and control outcomes on price perturbations. $\hat{\boldsymbol{\xi}}_z$ is a $J \times J$ matrix which is computed via regressions of each component of excess demand on price perturbations. Meanwhile, for the variance of the indirect effect, we consider

$$\hat{\sigma}_I^2 = [\hat{\boldsymbol{\xi}}_z^{-1} \hat{\boldsymbol{\tau}}_z^{HT}]^\top [\hat{\boldsymbol{\xi}}_z^{-1} \hat{\boldsymbol{\tau}}_z^{HT}] \hat{v}^2, \quad (10)$$

where $\hat{v}^2 = \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{Z}_i^\top \hat{\boldsymbol{\gamma}})^2$. The asymptotic normal approximation can be used to build confidence intervals for the direct and indirect estimates. For example, for a confidence level of 95%, the confidence intervals are constructed as

$$\hat{\tau}_D \pm 1.96 \cdot \frac{\hat{\sigma}_D}{\sqrt{n}}, \quad \hat{\tau}_I \pm 1.96 \cdot \frac{\hat{\sigma}_I}{\sqrt{nh_n}}.$$

The proofs of the following results are given in Appendix A.8.

Theorem 7. *Under the assumptions of Theorem 3 and Design 2, $\hat{\sigma}_D^2 \rightarrow_p \sigma_D^2$ and $\hat{\sigma}_I^2 \rightarrow_p \sigma_I^2$.*

Remark. To derive the asymptotic variance of the indirect effect, we can use the delta method to decompose it into one term that depends on $\boldsymbol{\tau}_Z^{HT}$ and the asymptotic variance of $\hat{\boldsymbol{\gamma}}$, and one term that depends on $\boldsymbol{\gamma}$ and the asymptotic variance of $\hat{\boldsymbol{\tau}}_Z^{HT}$, see Appendix A.7 for details. The

second term is asymptotically negligible, since $\hat{\gamma}$ converges at a slower rate than the differences in means estimator. However, we found in simulations that including a second order correction that retains the asymptotically negligible second term leads to better coverage under a wider variety of choices for the perturbation size when the sample size is smaller. $\hat{\sigma}_I^2$ is a plug-in estimator for the first term depending only on the asymptotic variance of $\hat{\gamma}$ and τ_Z^{HT} . $\tilde{\sigma}_I^2$ adds a plug-in estimator for the second term, including an estimator for the variance of $\hat{\tau}_Z^{HT}$: $\tilde{\sigma}_I^2 = \hat{\sigma}_I^2 + h_n^2 \hat{\gamma}^\top \hat{\Omega} \hat{\gamma}$, where $\hat{\Omega} = \widehat{\text{Var}}[\mathbf{B}_i]$ and $\hat{\mathbf{B}}_i = \frac{W_i \mathbf{Z}_i}{q} - \frac{(1-W_i) \mathbf{Z}_i}{1-q} - [\hat{\xi}_{z1} - \hat{\xi}_{z0}]^\top \hat{\xi}_z^{-1} \mathbf{Z}_i$.

4 Heterogeneous Treatment Effects

So far, we defined a potential outcomes model that captured equilibrium interference but reduced to the Neyman-Rubin model without interference. We then used this model to define average direct and indirect treatment effects under general treatment allocation rules. In the previous section, we proposed estimators for these effects that relied on data generated from randomized trials where the treatment was assigned with constant probability. In this section, we return to general treatment rules and discuss heterogeneous effects and optimal targeting when there is interference through an equilibrium statistic.

The planner controls the treatment allocation function $\pi(\cdot) : \mathcal{X} \rightarrow [0, 1]$ where $\pi(x) = Pr(W_i = 1 | X_i = x)$. The conditional expectation functions are defined as $\mu(w, p, x) = \mathbb{E}[Y_i(w, p) | X_i = x]$ and $\mathbf{z}(w, p, x) = \mathbb{E}[\mathbf{Z}_i(w, p) | X_i = x]$. We are interested both in quantifying how relevant treatment effects vary with x , and how this information can be used to guide choices of $\pi(\cdot)$ that achieve better outcomes.

4.1 Definitions of Conditional Treatment Effects Under Interference

Under SUTVA, the CATE is defined as $\mathbb{E}[Y_i(1) - Y_i(0) | X_i = x]$ (Imbens 2004). Without interference, the optimal unconstrained targeting rule allocates treatments only to those with a positive CATE (Manski 2004). In budget-constrained settings, treatments are allocated to those with the largest CATEs above a threshold (Bhattacharya & Dupas 2012). When an individual's outcome depends on the entire vector of treatments, then the CATE is not well-defined, so it is not possible to immediately extend the existing optimal policy results. There is a much broader space of possible conditional estimands under interference than with SUTVA. Should the estimand be defined under changes to a single individual's treatment or multiple? Should the estimand be conditional on the treated individual's covariate, or the covariate of individuals whose treatment response is measured? We begin this section by proposing two definitions of conditional estimands that play a similar role to the ADE and the AIE, but for targeted treatments. Their sum corresponds to a global policy counterfactual that is useful for optimizing covariate-dependent treatment assignment.

Definition 2 proposes the Conditional Average Direct Effect (CADE) and the Conditional Average Indirect Effect (CAIE) under interference. The CADE is the expected effect of treating an individual with covariate value x on their own outcomes in a sample of n individuals. The CAIE is

is the expected effect of treating an individual with covariate value x on everyone else’s outcomes in a sample of n individuals. When interference is removed, the CADE is equal to the CATE, and the CAIE is zero.

Definition 2. The Conditional Average Direct Effect is defined as

$$\tau_{\text{CADE}}(x) = \mathbb{E}[Y_{in}(W_i = 1; W_{-i}) - Y_{in}(W_i = 0; W_{-i}) | X_i = x]$$

The Conditional Average Indirect Effect is defined, where $j \neq i$, as

$$\tau_{\text{CAIE}}(x) = (n - 1)\mathbb{E}[Y_{jn}(W_i = 1; W_{-i}) - Y_{jn}(W_i = 0; W_{-i}) | X_i = x],$$

The CATE, defined under SUTVA, is a population-level quantity that does not depend on the sample size. Both the CADE and the CAIE are non-random quantities defined in terms of expectations over random outcome functions and treatments. However, unlike the CATE, the conditional average treatment effects defined under interference do depend on the sample size n .

We next connect these definitions of heterogeneous effects to a counterfactual of average outcomes with respect to the policy $\pi(\cdot)$. This result is valid under general patterns of interference. In Section 2, we reported the results of [Hu et al. \(2022\)](#), showing the sum of τ_{ADE} and τ_{AIE} is equal to the effect on average outcomes of an infinitesimal increase in each individual’s treatment probability. This is a local approximation to τ_{GTE} , which is relevant for making a decision on whether or not to treat everyone. For heterogenous effects, comparable results are not currently available in the literature. In Proposition 8, we show that the expected derivative of average outcomes with respect to the treatment probability of an individual who has covariate x is the sum of $\tau_{\text{CADE}}^Y(x)$ and $\tau_{\text{CAIE}}(x)$.

Proposition 8. *Under the assumptions of Theorem 3,*

$$n\mathbb{E} \left[\frac{\partial}{\partial \pi_k} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\pi[Y_i(\mathbf{W})] \Big| X_k = x \right] = \tau_{\text{CADE}}(x) + \tau_{\text{CAIE}}(x).$$

A policymaker may be interested in taking advantage of heterogeneous responses to treatment in the population by implementing a targeting rule, rather than a treatment rule with uniform probability. Proposition 8 shows that the sum of the CADE and the CAIE is relevant for making the decision on which group’s treatment probability to increase, and which to decrease, when covariates are discrete, and the objective is maximizing expected outcomes.

The next step, in Theorem 9, is to derive limits of the heterogeneous effects in our stochastic equilibrium model. This allows us to interpret what influences the CADE and CAIE in settings with interference through market prices.

Theorem 9. *Under the Assumptions of Theorem 3, the limit of the conditional average direct effect*

is:

$$\lim_{n \rightarrow \infty} \tau_{CADE}(x) = \tau_{CADE}^*(x) = y(1, \mathbf{p}_\pi^*, x) - y(0, \mathbf{p}_\pi^*, x).$$

The limit of the conditional average indirect effect is:

$$\lim_{n \rightarrow \infty} \tau_{CAIE}(x) = \tau_{CAIE}^*(x) = -\boldsymbol{\xi}_y^\top \boldsymbol{\xi}_z^{-1} [\mathbf{z}(1, \mathbf{p}_\pi^*, x) - \mathbf{z}(0, \mathbf{p}_\pi^*, x)].$$

These limits are interpretable and suggest estimation strategies for $\tau_{CADE}(x)$ and $\tau_{CAIE}(x)$. The limit of the CADE is the direct treatment effect on outcomes conditional on x , holding the equilibrium price fixed. The limit of the CAIE is the direct treatment effect on excess demand conditional on x , multiplied by an elasticity correction that does not depend on x . The augmented randomized experiment from Design 2 can be used to estimate the conditional average direct effects and the elasticity corrections required to estimate $\tau_{CADE}(x)$ and $\tau_{CAIE}(x)$. Given the results in Proposition 8 and Theorem 9, we have shown that these estimates can then be used to optimize expected outcomes without constraints by (repeatedly) adjusting $\pi(x)$ in the direction of the estimated derivative of expected outcomes with respect to $\pi(x)$.

The structure of the price equilibrium leads to a simple form for the mean-field conditional effects. The market price is an aggregate statistic, so individuals with different covariates all respond to the same market prices. A change in excess demand of a given size always has the same impact on the market price. As a result, although individuals responses to the treatment through outcomes or excess demand are heterogeneous, the elasticity correction that transforms the $\tau_{CADE}(x)$ to $\tau_{CAIE}(x)$ is unconditional. This implies that a group of individuals' effect on the system depends on their covariates only through conditional direct effects.

4.2 Equilibrium-Stable Targeting

In the previous section, we briefly discussed how estimates of $\tau_{CADE}(x)$ and $\tau_{CAIE}(x)$ could be used to locally optimize an unconstrained targeting policy. There are a variety of constrained targeting rules that may be of interest. Here, we focus on one specific question of this type, namely what is the optimal treatment assignment policy that does not move equilibrium prices relative to those seen in the experiment? There are two reasons to consider this class of targeting rules. First, a conservative policymaker may be reluctant to significantly modify the equilibrium—even if it is beneficial on average to individuals—and so they may want to know how much they can improve outcomes without changing the equilibrium. Second, from a practical point of view, we find that the answer to this question admits a simple econometric strategy, and can be answered without needing to estimate price elasticities and without recourse to an augmented experimental design. Instead, the optimal equilibrium-stable targeting rule (and its performance) can be estimated using data from a baseline RCT following Design 1, as long as the RCT collects outcome and relevant supply and demand data at an individual level.

The optimization problem for the equilibrium-stable policy is to maximize expected outcomes

while ensuring the equilibrium price is unchanged from one where individuals are treated with uniform probability $\pi(x) = q$ for all x :

$$\arg \max_{\pi(\cdot)} \mathbb{E}[Y_i(W_i, \mathbf{p}_\pi^*)] \text{ s.t. } \mathbf{p}_\pi^* = \mathbf{p}_q^*. \quad (11)$$

Proposition 10. *Under Assumptions 2–4, optimizing Equation 11 is equivalent to solving the following infinite-dimensional linear optimization problem:*

$$\begin{aligned} \arg \max_{\pi(\cdot)} \mathbb{E} \left[\pi(X_i) \mathbb{E}[Y_i(1, \mathbf{p}_q^*) - Y_i(0, \mathbf{p}_q^*) | X_i = x] \right] \\ \text{s.t. } \mathbb{E} \left[\pi(X_i) \mathbb{E}[\mathbf{Z}_i(1, \mathbf{p}_q^*) - \mathbf{Z}_i(0, \mathbf{p}_q^*) | X_i = x] \right] = \mathbf{b}, \end{aligned}$$

where $\mathbf{b} = q \mathbb{E}[\mathbf{Z}_i(1, \mathbf{p}_q^*) - \mathbf{Z}_i(0, \mathbf{p}_q^*)]$ does not depend on $\pi(\cdot)$.

We can estimate a solution to this using a dataset from a standard RCT by solving a linear program, where for each unit we plug in an estimate of $\tau_{\text{CADE}}^{Y^*}(x) = \mathbb{E}[Y_i(1, \mathbf{p}_q^*) - Y_i(0, \mathbf{p}_q^*) | X_i = x]$ and $\tau_{\text{CADE}}^{Z^*}(x) = \mathbb{E}[\mathbf{Z}_i(1, \mathbf{p}_q^*) - \mathbf{Z}_i(0, \mathbf{p}_q^*) | X_i = x]$; see [Athey et al. \(2019\)](#) for a random forest-based estimator. In contrast to the previous sections of the paper, we do not require estimates of price elasticities to estimate this rule.

Proposition 11 describes how in the single good case the optimal treatment rule takes the form of a separating hyperplane, where individuals are treated if their CADE on Y lies above a line defined in $\tau_{\text{CADE}}^{Z^*}(x) - \tau_{\text{CADE}}^{Y^*}(x)$ space. The absolute value of the coefficient c can be interpreted as a shadow price. It is the incremental value of an incremental relaxation of the constraint that the equilibrium price under the targeted policy is equal to the equilibrium price under the uniform policy. Its sign indicates in which direction the constraint that $\mathbb{E} \left[\pi(X_i) \mathbb{E}[\mathbf{Z}_i(1, \mathbf{p}_q^*) - \mathbf{Z}_i(0, \mathbf{p}_q^*) | X_i = x] \right] = \mathbf{b}$ is relaxed in order to increase the objective.

Proposition 11. *Let X_i have a continuous distribution. Then, the solution to Proposition 10 has the following form,*

$$\pi(x) = \begin{cases} 1 & \text{if } \tau_{\text{CADE}}^{Y^*}(x) \geq c \tau_{\text{CADE}}^{Z^*}(x), \\ 0 & \text{else,} \end{cases} \quad (12)$$

for some $c \in \mathbb{R}$.

This rule can be interpreted as follows. Ignoring the equilibrium effect, the optimal targeting rule allocates any individual with $\tau_{\text{CADE}}^{Y^*}(x) > 0$ to treatment. Let's assume that this results in a larger treatment effect on excess demand than the uniform RCT. In order to respect the equilibrium-stability constraint, this rule is adjusted, so that those with $\tau_{\text{CADE}}^{Y^*}(x)$ negative and small but $\tau_{\text{CADE}}^{Z^*}(x)$ negative and large are treated, and those with $\tau_{\text{CADE}}^{Y^*}(x)$ positive and small but $\tau_{\text{CADE}}^{Z^*}(x)$ positive and large are controlled. The amount of this adjustment depends on the magnitude of c . This form extends to the multiple good case, where an individual is treated if their CADE lies above a linear combination of the vector $\tau_{\text{CADE}}^{Z^*}(X_i)$.

Compared to the rule that allocates treatment to those with a positive CADE on Y and ignores equilibrium effects, the optimal equilibrium-stable rule will always have lower expected outcomes, but will also have a different equilibrium effect. Our preferred benchmark is the uniform rule that assigns treatment with probability q and has the same equilibrium effect in expectation as the targeting rule. We can evaluate the gain of the heterogeneous treatment rule over the uniform rule using data from an RCT and a data-splitting procedure.

We use the following procedure, which operates on a dataset with a vector of outcomes \mathbf{Y} , a vector of randomly assigned treatments \mathbf{W} , a matrix of covariates \mathbf{X} and a vector of excess demand \mathbf{Z} . When $J = 1$, finding \hat{c} is straightforward. The estimated score condition is monotonic in a single c , so we can start at $c = 0$ and adjust it either up or down until the score condition is equal to zero. When $J > 1$, extending this procedure requires solving a plug-in version of the linear program to estimate \hat{c} .

1. On a training dataset of size n_{train} , estimate $\hat{\tau}_{CADE}^Y(x)$ and $\hat{\tau}_{CADE}^Z(x)$ using a consistent estimator for conditional average treatment effects, such as the causal forest of [Athey et al. \(2019\)](#).
2. Let $S(X_i) = \mathbb{1}(\hat{\tau}_{CADE}^Y(X_i) > \hat{c}\hat{\tau}_{CADE}^Z(X_i))$. Choose \hat{c} such that

$$\frac{1}{n_{train}} \sum_{i=1}^{n_{train}} \left(\frac{W_i Z_i}{q} - \frac{(1 - W_i) Z_i}{1 - q} \right) S(X_i) = \frac{q}{n_{train}} \sum_{i=1}^{n_{train}} \left(\frac{W_i Z_i}{q} - \frac{(1 - W_i) Z_i}{1 - q} \right)$$

3. On a test dataset of size n_{test} , evaluate $\hat{\tau}_{CADE}^Y(X_i)$ and $\hat{\tau}_{CADE}^Z(X_i)$ and compute $S(X_i)$ for each $i = 1, \dots, n_{test}$.
4. The gain in the targeting rule compared to the uniform rule is computed as

$$\frac{1}{n_{test}} \sum_{i=1}^{n_{test}} \left(\frac{W_i Y_i}{q} - \frac{(1 - W_i) Y_i}{1 - q} \right) S(X_i) - \frac{q}{n_{test}} \sum_{i=1}^{n_{test}} \left(\frac{W_i Y_i}{q} - \frac{(1 - W_i) Y_i}{1 - q} \right)$$

5. The equilibrium stability of the targeting rule on the test set is computed as

$$\frac{1}{n_{test}} \sum_{i=1}^{n_{test}} \left(\frac{W_i Z_i}{q} - \frac{(1 - W_i) Z_i}{1 - q} \right) S(X_i) - \frac{q}{n_{test}} \sum_{i=1}^{n_{test}} \left(\frac{W_i Z_i}{q} - \frac{(1 - W_i) Z_i}{1 - q} \right).$$

This metric is zero when the targeting rule is equilibrium-stable for the test set.

In the next section, we will use this procedure illustrate the gain of an equilibrium-stable targeting rule over a uniform rule that targets using data from a cash transfer experiment.

5 Simulation and Application

In the first part of this section, we highlight the results from Sections 2 and 3 using a simulation of a “nudge” intervention in a two-sided market for freelance labor. We illustrate how looking at only the direct effect of the intervention (via a difference in means estimator) from an RCT results in a substantial over-estimate of the global treatment effect. We confirm numerically that, when treatments affect market prices, the indirect effect can have a substantial magnitude—but is estimable in finite samples via our proposed unit level augmented experiment. We also find that in finite samples our asymptotic approximations for the distributions of the direct and indirect effect are accurate enough to yield confidence intervals with good coverage properties. In the second part of this section, we estimate the equilibrium-stable targeting rule of Section 4 using data from a cash transfer experiment.

5.1 Simulation of Market for Freelance Labor

We simulate a two-sided market where, depending on the market price and their own preferences, suppliers with heterogeneous costs choose whether to complete 0, 1 or 2 jobs on the platform. On the other side of the market, consumers will book a job if the market price is less than their heterogeneous value for the job. The market price is set to clear the market for freelance labor. In the absence of an intervention, suppliers who have cost less than the market price will always complete one job, but will complete two jobs only 30% of the time (even though it is profitable for them). The treatment is an intervention, such as a reminder, that encourages active suppliers to complete a second job on the platform. We simulate the effect of an intervention that is highly effective; a supplier that receives the message will complete two jobs 70% of the time, as long as it is profitable for them to do so. The platform operator is interested in estimating the effect on average supplier revenue of rolling out this intervention, which is the Global Treatment Effect.

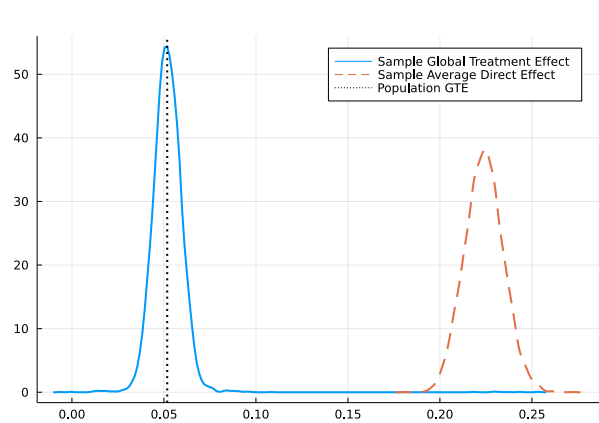
Formally, the demand and supply functions are

$$D_i(W_i, p) = (1 - \theta_i) \mathbb{1}(V_i > p)$$


$$S_i(W_i, p) = \theta_i \mathbb{1}(C_i < p) (1 + W_i \mathbb{1}(G_i > 0.3) + (1 - W_i) \mathbb{1}(G_i > 0.7))$$

$\theta_i = 1$ if an individual is a supplier on the platform. G_i follows a standard uniform distribution, and we have that both V_i and C_i are sampled from a $\text{Uniform}(1, 2)$ distribution. Half of a sample of n individuals are suppliers and the other half are consumers. Supplier revenue is $Y_i(\mathbf{W}) = S_i(W_i, P(\mathbf{W})) \cdot P(\mathbf{W})$. The market price for contract work $P(\mathbf{W})$ sets the average excess demand $Z_i(W_i, p) = D_i(W_i, p) - S_i(W_i, p)$ to zero.

In Figure 1a, we plot the smoothed distribution of the average direct effect and the global treatment effect, computed over 10,000 samples of a marketplace, where each sample is of 2,000 individuals. An RCT estimates the ADE, which is more than triple the Global Treatment Effect. Figure 1b explains why the ADE and GTE are so different in this setting. The figure shows the effect



(a) Smoothed empirical distribution of τ_{ADE} and τ_{GTE} over 10,000 samples of $n = 2,000$ marketplace.



(b) Average excess demand and average excess supply under treatment and control for a single sample of $n = 10,000$ marketplace.

Figure 1: Illustration of the gap between τ_{ADE} and τ_{GTE} in the freelance labor model when the intervention has both a direct effect on quantity supplied and an offsetting indirect effect through decreased market-clearing price.

of the treatment on aggregate supply $S(\mathbf{W}, p) = \frac{1}{n} \sum_{i=1}^n S_i(W_i, p)$ for a sample of 10,000 individuals. Aggregate demand is unchanged from the treatment. The treatment makes it more likely that already profitable suppliers complete a second job, which shifts the slope of the supply curve. This impacts the market-clearing price for the control market compared to the treated market. Holding market prices fixed, the effect of the treatment is a large positive increase in revenue. However, when the effect on market prices is taken into account, which is represented by a large negative indirect effect, then the global treatment effect is much lower.

In Table 1, we report the results of a Monte Carlo simulation with 10,000 repetitions and a sample size of $n = 2,000$ to evaluate the bias, variance, and coverage properties of the estimators and confidence intervals for $\hat{\tau}_{ADE}$ and $\hat{\tau}_{AIE}$. The ground truth τ_{ADE}^* and τ_{AIE}^* are calculated numerically. The table illustrates that the bias of the average direct effect and indirect effect estimator based on the augmented randomize experiment are low in finite samples. Furthermore, the coverage when $\hat{\sigma}_D$ and $\hat{\sigma}_I$ are used to construct confidence intervals in finite samples is very close to the asymptotic confidence level of 95%. If only data from an RCT was available, then an estimate for the AIE is not available, while the estimator for the ADE remains a Horvitz-Thompson estimator. The naive confidence interval that assumes no-interference for the direct effect can be computed using data from an RCT. In this simulation, the naive approach gives asymptotically conservative confidence intervals, so has slightly wider confidence intervals, and over-covers compared to the asymptotically exact approach. For treatments that increase the derivative of outcomes with respect to prices more drastically than in this simulation, then the naive confidence intervals for the direct effect would be even wider.

	Estimate	Bias	S.D.	Coverage
$\hat{\tau}_{\text{ADE}}$	0.227	0.004	0.06	0.957
$\hat{\tau}_{\text{AIE}}$	-0.173	-0.003	0.05	0.954

Table 1: Monte Carlo Simulation Results for $\hat{\tau}_{\text{ADE}}$ and $\hat{\tau}_{\text{AIE}}$ for a sample of $n = 2,000$ and 10,000 repetitions

5.2 Targeting Cash Transfers

There is a large literature in development economics that randomly allocates cash transfers to households and evaluates the short and long-term impacts on a variety of outcomes, including consumption, savings and investment (Haushofer & Shapiro 2016), health (Haushofer et al. 2020), and criminal behavior (Hidrobo et al. 2016, Blattman et al. 2017).¹

Depending on the characteristics of the local economies where cash transfers are allocated and the type of outcome evaluated, there may be equilibrium interference that impacts the evaluation of the transfer program. Cunha et al. (2019), for example, find meaningful price impacts for in-kind transfers in an evaluation in Mexico, while Egger et al. (2022) find minimal price impacts for a large scale cash transfer in Kenya. The results in Section 3 indicate that estimating the Global Total Effect of an intervention using data from an RCT requires estimating of elasticities of outcomes, supply, and demand with respect to prices. However, our results in Section 4 indicate that even without price elasticities, we can use data from an RCT to analyze the gain from exploiting heterogeneity in the data when we constrain equilibrium effects under a targeted policy to be equal to that under the RCT.

We examine the publicly available data from the replication package of Gertler et al. (2012), which examined the long-term impacts of PROGRESA, a conditional cash transfer program in Mexico which was rolled out randomly at a community-level, beginning in 1998. The community-level assignment was due to logistical and operational constraints, rather than an attempt to control spillovers.

The program included a monthly stipend of 50 pesos per month, and an additional educational stipend, depending on the number of the children in the household, with the total capped at 550 pesos per month. There are 320 treatment communities, which were immediately enrolled in the program, and 186 control communities, which were phased into treatment over time. The data from Gertler et al. (2012) contains a variety of baseline covariates from the 1997 ENCASEH census for households originally classified as poor and eligible for the intervention, as well as measurements from evaluation surveys that were given to all households in treatment and control communities every 6 months, starting in 1998. The evaluation surveys collected a variety of data on household education, consumption, investments, savings, and agricultural income. We follow the original paper in pooling data from October 1998 to November 1999, which was the purely experimental period, before control households began to be enrolled in the program. We are interested in using

¹See the research paper explorer at <https://www.givedirectly.org/cash-evidence-explorer/> for a more comprehensive review of cash transfer RCTs.

a subset of the data to answer a specific question, which is how a policymaker could target a PROGRESA-like subsidy to households to maximize agricultural income, when land usage under the targeted policy is constrained to be equal to that of the RCT. We compute the structure of the targeting rule, and estimate how much the optimal targeted policy raises outcomes compared to the uniform policy.

Before describing our results, we first explain a few caveats that should be kept in mind. First, depending on how quickly prices adjust and the geographic proximity of treated and control households, the community-level assignment could result in different equilibrium conditions for treated and control households. For the purposes of illustrating our method, we assume that treated and control individuals are responding to similar market prices in the experimental period, so that the conditional direct effects we measure do not include an equilibrium effect. Second, a general cash subsidy likely results in equilibrium effects of varying strength that occur through an enormous variety of channels, including the labor market, commodities markets, and even financial markets. For this exercise, we assume that the farm subsidy has an equilibrium effect on outcomes through a single good only, which is the usage of land. Last, it is possible that the usage of land is very slack, so that if the treatment increases usage of land, there is no effect on the actual market demand for land. We assume that the net usage of land does proxy for the net market demand for land in our sample, so that our constraint on land usage corresponds to an equilibrium stability condition.

In Figure 2, we illustrate the optimal rule estimated using a sample of $n = 33,182$ households from Waves 2-4 of the survey. There are 20 covariates in X_i , which are all pre-treatment household demographic characteristics from 1997, such as age and gender of the head of household, the existing landholdings and animal ownership of the household, and whether or not the house had electricity. The outcome Y_i is agricultural income per individual in the household, including crop sales, sales from animal products, and home consumption. Ignoring equilibrium effects, a good targeting rule allocates treatment to those individuals who are estimated to have a positive CADE. This rule, however, results in a larger impact on demand for land compared to the RCT, since on average, those who increase their agricultural income using the subsidy also increase their usage of land. However, there is some heterogeneity in the usage of land. The optimal rule that respects the equilibrium constraint allocates treatment to those with $\hat{\tau}_{\text{CADE}}^Y(X_i) > c\hat{\tau}_{\text{CADE}}^Z(X_i)$. We can see that for the most part, it is those with a positive CADE that are treated. However, those with a large $\frac{-\hat{\tau}_{\text{CADE}}^Z(X_i)}{\hat{\tau}_{\text{CADE}}^Y(X_i)}$ have their treatment status flipped from the naive rule that ignores the equilibrium constraints. Those, for example, who use a lot of additional land but have a small impact on agricultural income, are dropped from the equilibrium-stable targeting rule.

We then evaluate the gain in the targeted rule compared to the uniform rule using the data-splitting approach described at the end of Section 4.2. The results are reported in Table 2. We find that the lift is 2.27 and significantly different from zero at a 5% level. In this example, there is modest heterogeneity in outcomes and demand for land. When this is exploited via the optimal equilibrium-stable targeting rule, the point estimate for the gain over the uniform rule is positive




Figure 2: A scatter-plot of an estimate of $\tau_{\text{CADE}}(X_i)$ and $\tau_{\text{CADE}}^Z(X_i)$ for a sample of 125 households, estimated using the causal forest of [Athey et al. \(2019\)](#). Overlaid is the optimal equilibrium-stable targeting rule, computed by solving a plug-in version of the linear program in [Theorem 9](#) on random training split of $n = 10,822$, which takes the form of the rule in [Proposition 11](#).

	Gain	Equilibrium-Stability	c
Estimate	2.27	-0.006	0.92
Standard Error	(0.44)	(0.021)	(1.12)

Table 2: Estimated out-of-sample performance of equilibrium-stable targeting rule. Standard errors computed using the bootstrap.

and large. The equilibrium stability metric is -0.006 with a standard error of 0.44 , so on average the treatment rule estimated on the training set is disturbing the equilibrium on the test set by a very small amount, compared to the uniform rule.

6 Discussion

Analyzing the performance of randomized control trials in settings with equilibrium effects is needed given the rapid growth of experimentation both in practice and in research studies. The Neyman-Rubin framework that relies on SUTVA rules out interaction effects that can have an important impact on decision-relevant treatment effects. A parametric structural model may capture a variety of complex equilibrium effects, but is not robust to misspecification, which can be problematic when individuals behave in complex and heterogeneous ways. A model of treatment effects under general patterns of interference is intractable without clustering or other assumptions. This paper shows that it is fruitful to marry ex-ante knowledge about the structure of an economic environment with a non-parametric stochastic model of treatment effects under interference. This leads to a characterization of asymptotic properties of treatment effects and estimators of those treatment effects based on new and existing experimental designs that are robust to a wide range of modeling choices. Results on estimation, inference, and optimal targeting in complex environments with some economic structure imposed can then be easily contrasted with the large body of work that studies casual inference under SUTVA.

There are a variety of avenues for future work possible. One limitation of our approach is that we analyze the large sample limit of the market place where the number of suppliers and the number of buyers grow large; an analysis of experiments in settings where firms have significant market power would likely require different techniques. In general, extending our results to a broader class of equilibrium mechanisms would be of considerable interest.

References

- Abadie, A., Athey, S., Imbens, G. W. & Wooldridge, J. (2017), When should you adjust standard errors for clustering?, Technical report, National Bureau of Economic Research.
- Abbring, J. H. & Heckman, J. J. (2007), ‘Econometric evaluation of social programs, part iii: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation’, *Handbook of econometrics* **6**, 5145–5303.
- Ackerberg, D., Benkard, C. L., Berry, S. & Pakes, A. (2007), ‘Econometric tools for analyzing market outcomes’, *Handbook of econometrics* **6**, 4171–4276.
- Agarwal, N. & Somaini, P. (2018), ‘Demand analysis using strategic reports: An application to a school choice mechanism’, *Econometrica* **86**(2), 391–444.
- Angrist, J. D., Graddy, K. & Imbens, G. W. (2000), ‘The interpretation of instrumental variables estimators in simultaneous equations models with an application to the demand for fish’, *The Review of Economic Studies* **67**(3), 499–527.
- Angrist, J. D., Imbens, G. W. & Rubin, D. B. (1996), ‘Identification of causal effects using instrumental variables’, *Journal of the American statistical Association* **91**(434), 444–455.
- Angrist, J. D. & Krueger, A. B. (2001), ‘Instrumental variables and the search for identification: From supply and demand to natural experiments’, *Journal of Economic perspectives* **15**(4), 69–85.
- Aronow, P. M. & Samii, C. (2017), ‘Estimating average causal effects under general interference, with application to a social network experiment’, *The Annals of Applied Statistics* **11**(4), 1912–1947.
- Arrow, K. J. & Hahn, F. (1971), ‘General competitive analysis’.
- Athey, S., Eckles, D. & Imbens, G. W. (2018), ‘Exact p-values for network interference’, *Journal of the American Statistical Association* **113**(521), 230–240.
- Athey, S. & Imbens, G. (2016), ‘Recursive partitioning for heterogeneous causal effects’, *Proceedings of the National Academy of Sciences* **113**(27), 7353–7360.
- Athey, S., Tibshirani, J. & Wager, S. (2019), ‘Generalized random forests’, *The Annals of Statistics* **47**(2), 1148–1178.
- Athey, S. & Wager, S. (2021), ‘Policy learning with observational data’, *Econometrica* **89**(1), 133–161.
- Baird, S., Bohren, J. A., McIntosh, C. & Özler, B. (2018), ‘Optimal design of experiments in the presence of interference’, *Review of Economics and Statistics* **100**(5), 844–860.
- Banerjee, A. & Duflo, E. (2011), *Poor economics: A radical rethinking of the way to fight global poverty*, Public Affairs.
- Banerjee, A., Hanna, R., Olken, B. A., Satriawan, E. & Sumarto, S. (2021), Food vs. food stamps: Evidence from an at-scale experiment in indonesia, Technical report, National Bureau of Economic Research.
- Basse, G. W., Feller, A. & Toulis, P. (2019), ‘Randomization tests of causal effects under interference’, *Biometrika* **106**(2), 487–494.
- Benkeser, D. & Van Der Laan, M. (2016), The highly adaptive lasso estimator, in ‘2016 IEEE international conference on data science and advanced analytics (DSAA)’, IEEE, pp. 689–696.
- Berry, S., Levinsohn, J. & Pakes, A. (1995), ‘Automobile prices in market equilibrium’, *Econometrica: Journal of the Econometric Society* pp. 841–890.
- Berry, S. T. & Haile, P. A. (2021), Foundations of demand estimation, in ‘Handbook of Industrial Organization’, Vol. 4, Elsevier, pp. 1–62.
- Bhattacharya, D. & Dupas, P. (2012), ‘Inferring welfare maximizing treatment assignment under budget constraints’, *Journal of Econometrics* **167**(1), 168–196.
- Blattman, C., Jamison, J. C. & Sheridan, M. (2017), ‘Reducing crime and violence: Experimental evidence from cognitive behavioral therapy in liberia’, *American Economic Review* **107**(4), 1165–1206.
- Boucheron, S., Lugosi, G. & Massart, P. (2013), *Concentration inequalities: A nonasymptotic theory of independence*, Oxford university press.
- Chesher, A. & Rosen, A. M. (2017), ‘Generalized instrumental variable models’, *Econometrica* **85**(3), 959–989.
- Chetty, R. (2009), ‘Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods’, *Annu. Rev. Econ.* **1**(1), 451–488.
- Cornes, R., Hartley, R. & Sandler, T. (1999), ‘Equilibrium existence and uniqueness in public good models: An elementary proof via contraction’, *Journal of Public Economic Theory* **1**(4), 499–509.
- Cunha, J. M., De Giorgi, G. & Jayachandran, S. (2019), ‘The price effects of cash versus in-kind transfers’, *The Review of Economic Studies* **86**(1), 240–281.
- Duflo, E. (2004), ‘The medium run effects of educational expansion: Evidence from a large school construction program in indonesia’, *Journal of Development Economics* **74**(1), 163–197.
- Egger, D., Haushofer, J., Miguel, E., Niehaus, P. & Walker, M. (2022), ‘General equilibrium effects of cash transfers: experimental evidence from kenya’, *Econometrica* **90**(6), 2603–2643.

- Fisher, R. A. (1935), *The Design of Experiments*, Oliver and Boyd, Edinburgh.
- Gertler, P. J., Martinez, S. W. & Rubio-Codina, M. (2012), ‘Investing cash transfers to raise long-term living standards’, *American Economic Journal: Applied Economics* **4**(1), 164–92.
- Halloran, M. E. & Struchiner, C. J. (1995), ‘Causal inference in infectious diseases’, *Epidemiology* pp. 142–151.
- Haushofer, J., Mudida, R. & Shapiro, J. P. (2020), The comparative impact of cash transfers and a psychotherapy program on psychological and economic well-being, Technical report, National Bureau of Economic Research.
- Haushofer, J. & Shapiro, J. (2016), ‘The short-term impact of unconditional cash transfers to the poor: experimental evidence from kenya’, *The Quarterly Journal of Economics* **131**(4), 1973–2042.
- Heckman, J. J. (1979), ‘Sample selection bias as a specification error’, *Econometrica* pp. 153–161.
- Heckman, J. J., Lochner, L. & Taber, C. (1998), ‘General-equilibrium treatment effects: A study of tuition policy’, *The American Economic Review* **88**(2), 381–386.
- Heckman, J. J. & Pinto, R. (2018), ‘Unordered monotonicity’, *Econometrica* **86**(1), 1–35.
- Heckman, J. J. & Vytlacil, E. (2005), ‘Structural equations, treatment effects, and econometric policy evaluation’, *Econometrica* **73**(3), 669–738.
- Hidrobo, M., Peterman, A. & Heise, L. (2016), ‘The effect of cash, vouchers, and food transfers on intimate partner violence: evidence from a randomized experiment in northern ecuador’, *American Economic Journal: Applied Economics* **8**(3), 284–303.
- Hopenhayn, H. A. (1992), ‘Entry, exit, and firm dynamics in long run equilibrium’, *Econometrica: Journal of the Econometric Society* pp. 1127–1150.
- Hu, Y., Li, S. & Wager, S. (2022), ‘Average treatment effects in the presence of interference’, *Biometrika* (forthcoming).
- Hudgens, M. G. & Halloran, M. E. (2008), ‘Toward causal inference with interference’, *Journal of the American Statistical Association* **103**(482), 832–842.
- Imbens, G. W. (2004), ‘Nonparametric estimation of average treatment effects under exogeneity: A review’, *Review of Economics and Statistics* **86**(1), 4–29.
- Imbens, G. W. & Rubin, D. B. (2015), *Causal Inference in Statistics, Social, and Biomedical Sciences*, Cambridge University Press.
- Johari, R., Li, H. & Weintraub, G. (2020), ‘Experimental design in two-sided platforms: An analysis of bias’, *arXiv preprint arXiv:2002.05670*.
- Johnson, G. E. (1980), ‘The labor market effects of immigration’, *ILR Review* **33**(3), 331–341.
- Karrer, B., Shi, L., Bhole, M., Goldman, M., Palmer, T., Gelman, C., Konutgan, M. & Sun, F. (2021), Network experimentation at scale, in ‘Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining’, pp. 3106–3116.
- Kitagawa, T. & Tetenov, A. (2018), ‘Who should be treated? empirical welfare maximization methods for treatment choice’, *Econometrica* **86**(2), 591–616.
- Kohavi, R., Tang, D. & Xu, Y. (2020), *Trustworthy online controlled experiments: A practical guide to a/b testing*, Cambridge University Press.
- Krusell, P. & Smith, Jr, A. A. (1998), ‘Income and wealth heterogeneity in the macroeconomy’, *Journal of political Economy* **106**(5), 867–896.
- Leung, M. P. (2020), ‘Treatment and spillover effects under network interference’, *Review of Economics and Statistics* **102**(2), 368–380.
- Li, S. & Wager, S. (2022), ‘Random graph asymptotics for treatment effect estimation under network interference’, *The Annals of Statistics* **50**(4), 2334–2358.
- Liu, L. & Hudgens, M. G. (2014), ‘Large sample randomization inference of causal effects in the presence of interference’, *Journal of the American Statistical Association* **109**(505), 288–301.
- Manski, C. F. (1993), ‘Identification of endogenous social effects: The reflection problem’, *The review of economic studies* **60**(3), 531–542.
- Manski, C. F. (2004), ‘Statistical treatment rules for heterogeneous populations’, *Econometrica* **72**(4), 1221–1246.
- Manski, C. F. (2013), ‘Identification of treatment response with social interactions’, *The Econometrics Journal* **16**(1), S1–S23.
- Mas-Colell, A., Whinston, M. D. & Green, J. R. (1995), *Microeconomic theory*, Vol. 1, Oxford university press New York.
- Matzkin, R. L. (2015), ‘Estimation of nonparametric models with simultaneity’, *Econometrica* **83**(1), 1–66.
- Mézard, M., Parisi, G. & Virasoro, M. A. (1987), *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, Vol. 9, World Scientific Publishing Company.
- Muralidharan, K., Niehaus, P. & Sukhtankar, S. (2017), General equilibrium effects of (improving) public employment programs: Experimental evidence from india, Technical report, National Bureau of Economic Research.

- Neilson, C., Allende, C., Gallego, F. et al. (2019), Approximating the equilibrium effects of informed school choice, Technical report.
- Newey, W. K. & McFadden, D. (1994), ‘Large sample estimation and hypothesis testing’, *Handbook of econometrics* **4**, 2111–2245.
- Roy, A. D. (1951), ‘Some thoughts on the distribution of earnings’, *Oxford economic papers* **3**(2), 135–146.
- Sahoo, R. & Wager, S. (2022), ‘Policy learning with competing agents’, *arXiv preprint arXiv:2204.01884* .
- Sävje, F., Aronow, P. M. & Hudgens, M. G. (2021), ‘Average treatment effects in the presence of unknown interference’, *The Annals of Statistics* **49**(2), 673–701.
- Tchetgen, E. J. T. & VanderWeele, T. J. (2012), ‘On causal inference in the presence of interference’, *Statistical methods in medical research* **21**(1), 55–75.
- Tsitsiklis, J. N. & Xu, K. (2013), ‘On the power of (even a little) resource pooling’, *Stochastic Systems* **2**(1), 1–66.
- van der Vaart, A. W. (1998), *Asymptotic statistics*, Cambridge university press.
- van der Vaart, A. & Wellner, J. A. (1997), ‘Weak convergence and empirical processes with applications to statistics’, *Journal of the Royal Statistical Society-Series A Statistics in Society* **160**(3), 596–608.
- Van der Weele, T. J., Luedtke, A. R., van der Laan, M. J. & Kessler, R. C. (2019), ‘Selecting optimal subgroups for treatment using many covariates’, *Epidemiology (Cambridge, Mass.)* **30**(3), 334.
- Van Long, N. & Soubeyran, A. (2000), ‘Existence and uniqueness of cournot equilibrium: a contraction mapping approach’, *Economics Letters* **67**(3), 345–348.
- Vazquez-Bare, G. (2022), ‘Causal spillover effects using instrumental variables’, *Journal of the American Statistical Association* pp. 1–12.
- Viviano, D. (2019), ‘Policy targeting under network interference’, *arXiv preprint arXiv:1906.10258* .
- Vvedenskaya, N. D., Dobrushin, R. L. & Karpelevich, F. I. (1996), ‘Queueing system with selection of the shortest of two queues: An asymptotic approach’, *Problemy Peredachi Informatsii* **32**(1), 20–34.
- Wager, S. & Athey, S. (2018), ‘Estimation and inference of heterogeneous treatment effects using random forests’, *Journal of the American Statistical Association* **113**(523), 1228–1242.
- Wager, S. & Xu, K. (2021), ‘Experimenting in equilibrium’, *Management Science* **11**(67), 6694–6715.
- Walras, L. (1900), *Éléments d’économie politique pure: ou, Théorie de la richesse sociale*, F. Rouge.
- Weintraub, G. Y., Benkard, C. L. & Van Roy, B. (2008), ‘Markov perfect industry dynamics with many firms’, *Econometrica* **76**(6), 1375–1411.
- Wolf, C. (2019), ‘The missing intercept: A demand equivalence approach’, *Job Market Paper* .

A Main Results

A.1 Notation

We first include a brief comment on the asymptotic notations used in the proofs. We use asymptotic notations such as $O(\cdot)$ and $o(\cdot)$ in the conventional sense. In particular, we use $\mathbf{A}_n = O_p(\alpha_n)$ if for every ϵ , there exists an $c < \infty$ such that $\lim_{n \rightarrow \infty} \mathbb{P}\{\|\mathbf{A}_n\| > \alpha_n c\} < \epsilon$. Similarly, we use $\mathbf{A}_n = o_p(\alpha_n)$ if, for all $c > 0$, $\lim_{n \rightarrow \infty} \mathbb{P}\{\|\mathbf{A}_n\| > \alpha_n c\} = 0$.

A.2 Uniqueness of Mean-Field Market Price (Proposition 1)

The proof of uniqueness in this section makes use of the Banach fixed point theorem.

Lemma 12. Banach Fixed Point Theorem: Consider a set $D \subset \mathbb{R}^J$ and a function $g : D \rightarrow \mathbb{R}^J$. If

1. D is closed,
2. $x \in D$ implies that $g(x) \in D$, and
3. g is a contraction on D , meaning there exists a $0 \leq q < 1$ such that for all $\mathbf{x}, \mathbf{y} \in D$, we have that

$$\|g(\mathbf{x}) - g(\mathbf{y})\| \leq q \|\mathbf{x} - \mathbf{y}\|,$$

then, there exists a unique $\mathbf{x}^* \in D$ with $g(\mathbf{x}^*) = \mathbf{x}^*$.

Lemma 13. Contraction Property² Assume that the set $D \subset \mathbb{R}^J$ is convex and the function g has continuous partial derivatives. If $\|\nabla g(x)\| < 1$, then the mapping g is a contraction in D .

Now, we can prove Proposition 1. We first prove that the mean-field equilibrium price exists and is unique. Let $\mathbf{f}(\mathbf{p}) = \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p})] + \mathbf{p}$ be a function from \mathbb{R}^J to \mathbb{R}^J . Then Lemma 13 indicates that $\mathbf{f}(\mathbf{p})$ is a contraction mapping, since we assume that $\|\nabla \mathbf{f}(p)\| < 1$. By Lemma 12, then there is a unique fixed point

$$\mathbf{f}(\mathbf{p}^*) = \mathbf{p}^*,$$

which occurs when $\mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}^*)] = 0$. This means that the set $\{\mathbf{p} : \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p})] = 0\}$ has a unique element, which means that the mean-field equilibrium price \mathbf{p}_π^* is unique.

A.3 Convergence of Market Price

The below lemma is not used directly in the main text, but various steps in this proof are used elsewhere in the proof of the main results.

Lemma 14. Under Assumptions 2-4, the finite sample equilibrium price converges in probability to the mean-field price \mathbf{p}_π^* :

$$\mathbf{P}_n \rightarrow_p \mathbf{p}_\pi^*, \quad \text{as } n \rightarrow \infty.$$

Note that in this section, we prove this for a more general, perturbed system, where individuals receive market price $\mathbf{P}_n(\mathbf{W}, \mathbf{U})$, and prices are perturbed by \mathbf{U}_i . The results hold for any $\mathbf{U}_i \rightarrow 0$. So, they also hold for $\mathbf{U}_i = 0$, which proves Lemma 14 for price $\mathbf{P}_n(\mathbf{W})$.

We use Theorem 5.9 of van der Vaart (1998), which requires verifying two conditions.

²Source: <http://terpconnect.umd.edu/~petersd/666/fixedpoint.pdf>

Lemma 15. *Theorem 5.9 of van der Vaart (1998):*

Let \mathbf{M}_n be random vector-valued functions and let $\boldsymbol{\mu}$ be a fixed vector-valued function of ϕ such that for every $\epsilon > 0$:

$$\begin{aligned} \sup_{\phi \in \Phi} \|\mathbf{M}_n(\phi) - \boldsymbol{\mu}(\phi)\| &\rightarrow_p 0, \\ \inf_{\phi: d(\phi, \phi_0) \geq \epsilon} \|\boldsymbol{\mu}(\phi)\| &> 0 = \|\boldsymbol{\mu}(\phi_0)\|. \end{aligned}$$

Then any sequence of estimators $\hat{\phi}_n$ such that $\mathbf{M}_n(\hat{\phi}_n) = o_p(1)$ converges in probability to ϕ_0 .

Let $\mathbf{Z}_n(\mathbf{p}) = \frac{1}{n} \mathbf{Z}_i(W_i, \mathbf{p} + \mathbf{U}_i)$ and $\mathbf{z}(\mathbf{p}) = \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p})]$. We first show the uniform convergence of the estimating equation for the finite sample price $\mathbf{Z}_n(\mathbf{p})$ to the estimating equation for the mean-field price $\mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p})]$.

Lemma 16. *Under Assumptions 2–4, the finite sample average excess demand converges uniformly to the expected excess demand,*

$$\sup_{\mathbf{p} \in \mathcal{S}} \|\mathbf{Z}_n(\mathbf{p}) - \mathbf{z}(\mathbf{p})\| \rightarrow_p 0.$$

Proof. We have that

$$\begin{aligned} \sup_{\mathbf{p} \in \mathcal{S}} \|\mathbf{Z}_n(\mathbf{p}) - \mathbf{z}(\mathbf{p})\| &\leq \sup_{\mathbf{p} \in \mathcal{S}} \|\mathbf{Z}_n(\mathbf{p}) - \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p} + \mathbf{U}_i)]\| + \sup_{\mathbf{p} \in \mathcal{S}} \|\mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p} + \mathbf{U}_i)] - \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p})]\| \\ &\rightarrow_p 0 \end{aligned}$$

The uniform convergence

$$\sup_{\mathbf{p} \in \mathcal{S}} \|\mathbf{Z}_n(\mathbf{p}) - \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p} + \mathbf{U}_i)]\| \rightarrow_p 0$$

follows directly from Lemma 2.4 of Newey & McFadden (1994), since we have from Assumption 3 that $\mathbf{Z}_i(w, \mathbf{p})$ is weakly continuous and bounded and that \mathcal{S} is compact.

Since $\mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p} + \mathbf{U}_i)]$ is differentiable with a bounded first derivative then it is uniformly continuous, and then $\sup_{\mathbf{p} \in \mathcal{S}} \|\mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p} + \mathbf{U}_i)] - \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p})]\| \rightarrow_p 0$ follows from the uniform continuous mapping theorem, since $\mathbf{U}_i \rightarrow_p 0$. □

The following lemma (the proof is Problem 5.27 of van der Vaart (1998)) is useful for verifying the second condition.

Lemma 17. Strict Zero. *If a function $\boldsymbol{\mu} : \Phi \rightarrow \mathbb{R}^k$ is continuous and has a unique zero at ϕ_0 , then if Φ is compact, we have that for every $\epsilon > 0$:*

$$\inf_{\phi: d(\phi, \phi_0) \geq \epsilon} \|\boldsymbol{\mu}(\phi)\| > 0 = \|\boldsymbol{\mu}(\phi_0)\|.$$

We show that the mean-field price is a strict zero in the following lemma:

Lemma 18. *The mean-field price is a strict zero. Under Assumptions 2–4, for all $\epsilon > 0$,*

$$\inf_{\mathbf{p}: d(\mathbf{p}, \mathbf{p}_\pi^*) \geq \epsilon} \|\mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p})]\| > 0 = \|\mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^*)]\|.$$

Proof. The proof follows immediately from Lemma 17 since \mathcal{S} is compact, $\mathbf{z}(\mathbf{p})$ is continuous in \mathbf{p} , and \mathbf{p}_π^* is unique. \square

Since we have from the version of Assumption 3.5 in Proposition 5, that that $\mathbf{Z}_n(\tilde{\mathbf{P}}) = o_p(n^{-0.5})$, where $\mathbf{Z}_n(\mathbf{p}) = \frac{1}{n} \sum \mathbf{Z}_i(W_i, \mathbf{p} + \mathbf{U}_i)$, then we have met all the conditions of Lemma 15 and we have proved that

$$\tilde{\mathbf{P}} \rightarrow_p \mathbf{p}_\pi^*.$$

A.4 Rate of Market Price Convergence (Theorem 2 and Proposition 5)

We first prove a more general version of Theorem 2 for the perturbed system with market prices $\mathbf{P}_n(\mathbf{W}, \mathbf{U})$.

Theorem 19. *Under Assumptions 2-4, the finite sample equilibrium price in the price experiment converges in distribution to the mean-field price and is asymptotically normal:*

$$\sqrt{n}(\tilde{\mathbf{P}}_n - \mathbf{p}_\pi^*) \rightarrow_d N(0, \boldsymbol{\xi}_z^{-1} \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^*) \mathbf{Z}_i(W_i, \mathbf{p}_\pi^*)^\top] \boldsymbol{\xi}_z^{-1}).$$

Proof. The following result from theory of Z-estimators is used in this proof. We replace the Lipschitz condition from the original Theorem 5.21 with the Donsker condition discussed in Lemma 19.24 of van der Vaart (1998).

Lemma 20. Modified Theorem 5.21 of van der Vaart (1998) : *For each ϕ in an open subset of Euclidean space, let $\mathbf{x} \mapsto \mathbf{m}_\phi(\mathbf{x})$ be a measurable vector-valued function, such that the class $\{\mathbf{m}_\phi : \|\phi - \phi_0\| < \delta\}$ is Donsker for some $\delta > 0$ and $\mathbb{E}[(\mathbf{m}_{\hat{\phi}_n}(\mathbf{X}_i) - \mathbf{m}_{\phi_0}(\mathbf{X}_i))^2] \rightarrow_p 0$. Assume that $\mathbb{E}[\|\mathbf{m}_{\phi_0}(\mathbf{X}_i)\|^2] < \infty$ and that the map $\phi \mapsto \mathbb{E}[\mathbf{m}_\phi(\mathbf{X}_i)]$ is differentiable at a zero ϕ_0 , with nonsingular derivative matrix \mathbf{V}_{ϕ_0} . If $\frac{1}{n} \sum_{i=1}^n \mathbf{m}_{\hat{\phi}_n}(\mathbf{X}_i) = o_p(n^{-1/2})$ and $\hat{\phi}_n \rightarrow_p \phi_0$, then:*

$$\sqrt{n}(\hat{\phi}_n - \phi_0) = -\mathbf{V}_{\phi_0}^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n m_{\phi_0}(\mathbf{X}_i) + o_p(1).$$

We verify each of the conditions of Lemma 20.

- We can show the required quadratic mean convergence using Lemma 30, where $F_i(\mathbf{p}) = \mathbf{Z}_i(W_i, \mathbf{p} + \mathbf{U}_i)$, where the required assumptions are in Assumption 3 - 4.
- By Assumption 3, the class of functions $\mathbf{Z}_i(w, \mathbf{p})$ indexed by \mathbf{p} is Donsker.
- Since $\mathbf{Z}_i(w, \mathbf{p})$ is bounded, then $\mathbb{E}[\|\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i)\|^2] < \infty$.
- By Lemma 14, we have that $\tilde{\mathbf{P}}_n \rightarrow_p \mathbf{p}_\pi^*$.
- By the differentiability of $\mathbf{z}(w, \mathbf{p})$, $\nabla_p \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i)]$ exists. Furthermore, by the continuous mapping theorem, since $\mathbf{U}_i \rightarrow_p 0$, $\nabla_p \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i)] \rightarrow_p \boldsymbol{\xi}_z$ which is non-singular by Assumption 4.

Then, we can apply Theorem 5.21 of [van der Vaart \(1998\)](#):

$$\sqrt{n}(\tilde{\mathbf{P}} - \mathbf{p}_\pi^*) = -[\boldsymbol{\xi}_z]^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i) + o_p(1) \quad (13)$$

$$\begin{aligned} &= -[\boldsymbol{\xi}_z]^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n \left[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i) - \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i)] \right] \\ &+ \sqrt{n} \boldsymbol{\xi}_z^{-1} (\mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^*)] - \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i)]) + o_p(1) \\ &= -[\boldsymbol{\xi}_z]^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n \left[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i) - \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i)] \right] + o_p(1) \end{aligned} \quad (14)$$

Where the last equality comes from Lemma 29, since $\mathbb{E}[\mathbf{z}(w, \mathbf{p} + \mathbf{U}_i)]$ is continuous in \mathbf{p} . We can now apply the CLT to the sum in (14) after scaling by the variance. Let $\mathbf{1}_J$ be the $J \times J$ identity matrix:

$$\left[\frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i) - \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i)] \right] [\text{Var}(\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i))]^{-1} \rightarrow_d N(0, \mathbf{1}_J)$$

Assumption 4 indicates that each element of the variance of $\mathbf{Z}_i(w, \mathbf{p})$ is continuous in \mathbf{p} . We can use the CMT to show that $\text{Var}(\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i)) = \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^*)^2] + o_p(1)$ since $\mathbf{U}_i \rightarrow_p 0$. We have now shown that

$$\sqrt{n}(\tilde{\mathbf{P}} - \mathbf{p}_\pi^*) \rightarrow_d N(0, \boldsymbol{\xi}_z^{-1} \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^*) \mathbf{Z}_i(W_i, \mathbf{p}_\pi^*)^\top] \boldsymbol{\xi}_z^{-1}).$$

□

Proof of Theorem 2 and Proposition 5

Theorem 2 follows directly from Theorem 19. The results hold for any $\mathbf{U}_i = c n^{-\alpha}$, with $\alpha > \frac{1}{4}$ and finite c . Thus, the conclusions hold for $\mathbf{P}_n(\mathbf{W}) = \mathbf{P}_n(\mathbf{W}, \mathbf{0})$.

For the Proposition, the expansion in Equation 13 indicates that

$$\sqrt{n}(\mathbf{P}_n - \tilde{\mathbf{P}}_n) = [\boldsymbol{\xi}_z]^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n [\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i) - \mathbf{Z}_i(W_i, \mathbf{p}_\pi^*)] + o_p(1).$$

Since we have that $\mathbf{p}_\pi^* + \mathbf{U}_i \rightarrow_p \mathbf{p}_\pi^*$, and $\mathbf{Z}_i(W_i, \mathbf{p})$ is bounded and continuous with probability 1 in \mathbf{p} , then we have the quadratic mean convergence $\mathbb{E}[(\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i) - \mathbf{Z}_i(W_i, \mathbf{p}_\pi^*))^2] \rightarrow 0$ (see Lemma 30 for a more detailed argument). Then, from the Donsker assumption of Assumption 3, we can use Lemma 19.24 of [van der Vaart \(1998\)](#) to show that the expansion above is equivalent to:

$$\begin{aligned} \sqrt{n}(\mathbf{P}_n - \tilde{\mathbf{P}}_n) &= [\boldsymbol{\xi}_z]^{-1} \sqrt{n} \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^* + \mathbf{U}_i)] - \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_\pi^*)] + o_p(1) \\ &= o_p(1) \end{aligned}$$

where the second equality is from Lemma 29 since $\alpha > \frac{1}{4}$. The proof of the second part of Proposition 5 follows from Lemma 26, the proof of which is later in the Appendix.

A.5 Proof of Mean-Field Causal Effects (Theorem 3)

Convergence of ADE

Proof. As implied by the technical appendix of Sävje et al. (2021) and in Hu et al. (2022), we have that the Horvitz-Thompson estimator is unbiased for the direct effect.

$$\hat{\tau}_{\text{ADE}} = \frac{1}{n} \sum_{i=1}^n \frac{W_i Y_i(W_i, \mathbf{P}_n(\mathbf{W}))}{\pi_i} - \frac{(1 - W_i) Y_i(W_i, \mathbf{P}_n(\mathbf{W}))}{1 - \pi_i}.$$

Taking expectation, we have that

$$\mathbb{E}_\pi[\hat{\tau}_{\text{ADE}}] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\pi[Y_i(W_i, \mathbf{P}_n(\mathbf{W})) | W_i = 1] - \mathbb{E}_\pi[Y_i(W_i, \mathbf{P}_n(\mathbf{W})) | W_i = 0] = \tau_{\text{ADE}}.$$

Subtracting τ_{ADE}^* from both sides of $\tau_{\text{ADE}} = \mathbb{E}_\pi[\hat{\tau}_{\text{ADE}}]$,

$$\tau_{\text{ADE}} - \tau_{\text{ADE}}^* = \mathbb{E}_\pi[\hat{\tau}_{\text{ADE}} - \tau_{\text{ADE}}^*].$$

Next, we show that $\hat{\tau}_{\text{ADE}} \rightarrow_p \tau_{\text{ADE}}^*$. By the expansion in Lemma 30, and the convergence of $\mathbf{P}_n(\mathbf{W})$ to \mathbf{p}_π^* in probability, we have that

$$\begin{aligned} \hat{\tau}_{\text{ADE}} &= \frac{1}{n} \sum_{i=1}^n \frac{W_i Y_i(W_i, \mathbf{p}_\pi^*)}{\pi_i} - \frac{(1 - W_i) Y_i(W_i, \mathbf{p}_\pi^*)}{1 - \pi_i} + o_p(1) \\ &= y(1, \mathbf{p}_\pi^*) - y(0, \mathbf{p}_\pi^*), \end{aligned}$$

where the second line is from the LLN. $\hat{\tau}_{\text{ADE}}$ is uniformly bounded, since it is the average of uniformly bounded functions and π_i is bounded away from 0 and 1. So, we have that $\hat{\tau}_{\text{ADE}} = \tau_{\text{ADE}}^* + o_p(1)$ implies that $\mathbb{E}_\pi[\hat{\tau}_{\text{ADE}} - \tau_{\text{ADE}}^*] = o_p(1)$.

Since $\hat{\tau}_{\text{ADE}}$ is an average of uniformly bounded functions, it is also uniformly bounded. So, we have, where B is finite, that for any $\epsilon > 0$, $\mathbb{E}_\pi[\hat{\tau}_{\text{ADE}} - \tau_{\text{ADE}}^*] \leq B \Pr_\pi(\hat{\tau}_{\text{ADE}} - \tau_{\text{ADE}}^* > \epsilon/2) + \epsilon/2$. This implies $\Pr(\mathbb{E}_\pi[\hat{\tau}_{\text{ADE}} - \tau_{\text{ADE}}^*] > \epsilon) \leq \Pr(\hat{\tau}_{\text{ADE}} - \tau_{\text{ADE}}^* > \epsilon/2) \rightarrow 0$. We have now shown that $\mathbb{E}_\pi[\hat{\tau}_{\text{ADE}} - \tau_{\text{ADE}}^*] = o_p(1)$, which finishes the proof. \square

Convergence of AIE

Proof. The first step is to prove a probability limit for τ_{INF} , which is

$$\tau_{\text{INF}} = \sum_{k=1}^n \frac{\partial}{\partial \pi_k} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\pi[Y_i(W_i, \mathbf{P}_n)].$$

Recall that $W_i \sim \text{Bernoulli}(\pi_i)$, we can expand the expectation with respect to \mathbf{W} :

$$\tau_{\text{INF}} = \sum_{k=1}^n \frac{\partial}{\partial \pi_k} \sum_{\mathbf{w}^j} \frac{1}{n} \sum_{i=1}^n f(\mathbf{w}^j) Y_i(w_i^j, \mathbf{P}_n(\mathbf{w}^j)),$$

where $f(\mathbf{w}^k) = \prod_{k=1}^n \pi_k^{w_k^j} (1 - \pi_k)^{1-w_k^j}$. We can take the derivative and manipulate to get

$$\frac{\partial}{\partial \pi_k} f(\mathbf{w}^j) = \frac{w_k^j - \pi_k}{\pi_k(1 - \pi_k)} f(\mathbf{w}^j).$$

This allows us to write

$$\begin{aligned} \tau_{\text{INF}} &= \sum_{k=1}^n \frac{\partial}{\partial \pi_k} \sum_{\mathbf{w}^j} \frac{w_k^j - \pi_k}{\pi_k(1 - \pi_k)} f(\mathbf{w}^j) \frac{1}{n} \sum_{i=1}^n Y_i(w_i^k, \mathbf{P}_n(\mathbf{w}^k)) \\ &= \mathbb{E}_\pi \left[\sum_{k=1}^n \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} \frac{1}{n} \sum_{i=1}^n Y_i(W_i, \mathbf{P}_n(\mathbf{W})) \right] \end{aligned} \quad (15)$$

We will use the following Lemma in the next few steps, which is proved in Appendix B.2.

Lemma 21. *Under the Assumptions of Theorem 3, for each of the below remainder terms $m \in \{1, 2, 3\}$, $\mathbb{E}_\pi \left[\sum_{k=1}^n \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} R_{mn} \right] \rightarrow_p 0$*

$$R_{1n} = \frac{1}{n} \sum_{i=1}^n Y_i(W_i, \mathbf{P}_n) - \mathbb{E}[Y_i(W_i, \mathbf{P}_n)] - (Y_i(W_i, \mathbf{p}_\pi^*) - \mathbb{E}[Y_i(W_i, \mathbf{p}_\pi^*)]) \quad (16)$$

$$R_{2n} \leq \|\mathbf{P}_n - \mathbf{p}_\pi^*\|^2 D \quad (17)$$

$$R_{3n} = \boldsymbol{\xi}_y^\top \left[(\mathbf{P}_n - \mathbf{p}_\pi^*) + \boldsymbol{\xi}_z^{-1} \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p}_\pi^*) \right], \quad (18)$$

where D is a finite constant.

We now use the expansion from Lemma 30, where $F_i(\mathbf{p}) = Y_i(W_i, \mathbf{p})$. The expansion is valid since $\mathbf{P}_n(\mathbf{W}) = \mathbf{p}_\pi^* + O_p(n^{-0.5})$ (see Theorem 2) and Assumption 3 - 4 provides the necessary continuity and differentiability properties. The expansion is

$$\frac{1}{n} \sum_{i=1}^n Y_i(W_i, \mathbf{P}_n(\mathbf{W})) = \frac{1}{n} \sum_{i=1}^n Y_i(W_i, \mathbf{p}_\pi^*) + (\mathbf{P}_n(\mathbf{W}) - \mathbf{p}_\pi^*)^\top \boldsymbol{\xi}_y + R_{1n} + R_{2n},$$

where R_{1n} and R_{2n} are each $o_p(n^{-0.5})$ and are defined in Lemma 21. Plugging this expansion into (15),

$$\tau_{\text{INF}} = \mathbb{E}_\pi \left[\sum_{k=1}^n \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} \left[\frac{1}{n} \sum_{i=1}^n Y_i(W_i, \mathbf{p}_\pi^*) + (\mathbf{P}_n(\mathbf{W}) - \mathbf{p}_\pi^*)^\top \boldsymbol{\xi}_y \right] \right] \quad (19)$$

$$+ \mathbb{E}_\pi \left[\sum_{k=1}^n \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} (R_{1n} + R_{2n}) \right]$$

$$= \mathbb{E}_\pi \left[\sum_{k=1}^n \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} \left[\frac{1}{n} \sum_{i=1}^n Y_i(W_i, \mathbf{p}_\pi^*) + (\mathbf{P}_n(\mathbf{W}) - \mathbf{p}_\pi^*)^\top \boldsymbol{\xi}_y \right] \right] + o_p(1). \quad (20)$$

We deal with the remainder term using Lemma 21. We can now work with the first part of the expectation directly. Examine $\sum_{k=1}^n \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)}$ and notice that $Y_i(W_i, \mathbf{p}_\pi^*)$ is independent of all

elements of that sum except for $W_i - \pi_i$.

$$\begin{aligned}\mathbb{E}_\pi \left[\sum_{k=1}^N \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} \frac{1}{n} \sum_{i=1}^n Y_i(W_i, \mathbf{p}^*) \right] &= \frac{1}{n} \sum_{i=1}^n \frac{\mathbb{E}_\pi[(W_i - \pi_i)Y_i(W_i, \mathbf{p}^*)]}{\pi_i(1 - \pi_i)} \\ &= \frac{1}{n} \sum_{i=1}^n Y_i(1, \mathbf{p}^*) - Y_i(0, \mathbf{p}^*)\end{aligned}\quad (21)$$

By a similar argument, we work with the second part of the expectation. We can substitute the expansion from Lemma 20 which indicates

$$(\mathbf{P}_n - \mathbf{p}_\pi^*) = -\boldsymbol{\xi}_z^{-1} \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p}_\pi^*) + R_{3n},$$

where $R_{3n} = o_p(n^{-0.5})$. Again, we have that $\mathbf{Z}_i(W_i, \mathbf{p}_\pi^*)$ is independent of all elements of the sum $\sum_{k=1}^N \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)}$ except for $W_i - \pi_i$.

$$\begin{aligned}\mathbb{E} \left[\sum_{k=1}^N \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} (\mathbf{P}_n(\mathbf{W}) - \mathbf{p}_\pi^*) \boldsymbol{\xi}_y \right] &= \mathbb{E} \left[\sum_{k=1}^N \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} (\boldsymbol{\xi}_z)^{-1} \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p}^*) \right] \\ &\quad + \mathbb{E}_\pi \left[\sum_{k=1}^N \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} (R_{3n}) \right] \\ &= \frac{1}{n} \boldsymbol{\xi}_z^{-1} \sum_{i=1}^n \frac{\mathbb{E}[(W_i - \pi_i) \mathbf{Z}_i(W_i, \mathbf{p}^*)]}{\pi_i(1 - \pi_i)} + o_p(1) \\ &= \frac{1}{n} \boldsymbol{\xi}_z^{-1} \sum_{i=1}^n \mathbf{Z}_i(1, \mathbf{p}^*) - \mathbf{Z}_i(0, \mathbf{p}^*) + o_p(1),\end{aligned}\quad (22)$$

where the remainder term converges in probability to zero by Lemma 21. Substituting (21) and (23) into (20), we have now shown that.

$$\tau_{\text{INF}} = \frac{1}{n} \sum_{i=1}^n Y_i(1, \mathbf{p}^*) - Y_i(0, \mathbf{p}^*) - \boldsymbol{\xi}_y^\top [\boldsymbol{\xi}_z]^{-1} \frac{1}{n} \sum_{i=1}^n (\mathbf{Z}_i(1, \mathbf{p}^*) - \mathbf{Z}_i(0, \mathbf{p}^*)) + o_p(1)$$

We can now apply the LLN to show that

$$\tau_{\text{INF}} = y(1, \mathbf{p}_\pi^*) - y(0, \mathbf{p}_\pi^*) - \boldsymbol{\xi}_y [\boldsymbol{\xi}_z]^{-1} [z(1, \mathbf{p}_\pi^*) - z(0, \mathbf{p}_\pi^*)] + o_p(1) = \tau_{\text{INF}}^*$$

We now use this probability limit to derive the limit of τ_{AIE} . Theorem 1 of Hu et al. (2022) indicates that

$$\tau_{\text{INF}} = \tau_{\text{ADE}} + \tau_{\text{AIE}}.$$

Since we have already proved in the previous section that $\tau_{\text{ADE}} \rightarrow_p \tau_{\text{ADE}}^*$, then $\tau_{\text{AIE}} \rightarrow_p \tau_{\text{AIE}}^*$, where

$$\tau_{\text{AIE}}^* = \tau_{\text{INF}}^* - \tau_{\text{ADE}}^*.$$

We have now proved that

$$\tau_{\text{AIE}}^* = -\boldsymbol{\xi}_y [\boldsymbol{\xi}_z]^{-1} [z(1, \mathbf{p}_\pi^*) - z(0, \mathbf{p}_\pi^*)].$$

□

A.6 Proof of Estimator Consistency

This theorem does not appear in the main text but is useful in the proof of other results.

Theorem 22. *Under the assumptions of Theorem 3 and Design 2, we have that the differences in mean estimate converges to the mean-field direct treatment effect:*

$$\hat{\tau}_{ADE} \rightarrow_p \tau_{ADE}^*.$$

And, the indirect effect estimate converges to the mean-field indirect treatment effect:

$$\hat{\tau}_{AIE} \rightarrow_p \tau_{AIE}^*.$$

This section requires the following two lemmas:

Lemma 23. Horvitz-Thompson Consistency

Let $V_i = V_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i)$ for $V_i(\cdot) \in \{Y_i(\cdot), \mathbf{Z}_i(\cdot)\}$ be one of observed outcomes, or excess demand. Let $v(k, \mathbf{p}) = \mathbb{E}[V_i(k, \mathbf{p})]$ and $v(\mathbf{p}) = \mathbb{E}[V_i(W_i, \mathbf{p})]$. Under the assumptions of Theorem 3 and Design 2,

$$\frac{1}{n} \sum_{i=1}^n \frac{W_i V_i}{q} - \frac{(1 - W_i) V_i}{1 - q} \rightarrow_p v(1, \mathbf{p}_q^*) - v(0, \mathbf{p}_q^*).$$

Proof. The Horvitz Thompson Estimator is:

$$\begin{aligned} \hat{\tau}_v &= \frac{1}{n} \sum_{i=1}^n \frac{W_i V_i(1, \tilde{\mathbf{P}}_n + \mathbf{U}_i)}{q} - \frac{(1 - W_i) V_i(0, \tilde{\mathbf{P}}_n + \mathbf{U}_i)}{1 - q} \\ &\stackrel{(1)}{=} \frac{1}{n} \sum_{i=1}^n \frac{W_i V_i(1, \mathbf{p}_q^* + \mathbf{U}_i)}{q} - \frac{(1 - W_i) V_i(0, \mathbf{p}_q^* + \mathbf{U}_i)}{1 - q} + (\tilde{\mathbf{P}}_n - \mathbf{p}_q^*) [\nabla_p v(1, \mathbf{p}_q^*) - \nabla_p v(0, \mathbf{p}_q^*)] + o_p(1) \\ &\stackrel{(2)}{=} \frac{1}{n} \sum_{i=1}^n \frac{W_i V_i(1, \mathbf{p}_q^* + \mathbf{U}_i)}{q} - \frac{(1 - W_i) V_i(0, \mathbf{p}_q^* + \mathbf{U}_i)}{1 - q} + o_p(1) \\ &\rightarrow_p v(1, \mathbf{p}_q^*) - v(0, \mathbf{p}_q^*) \end{aligned}$$

The expansion in (1) comes from Assumptions 3 - 4 and Lemma 30, since we have that $\tilde{\mathbf{P}}_n \rightarrow_p \mathbf{p}_q^*$ and that the class of $G_i V_i(W_i, \mathbf{p})$ indexed by p is Donsker for any bounded random variable G_i by the composition properties of Donsker classes. Then (2) comes from $\tilde{\mathbf{P}}_n + \mathbf{U}_i - \mathbf{p}_q^* = o_p(1)$. The convergence comes from the LLN and the CMT, noting that $\mathbb{E}[W_i] = q$ and that $\mathbf{U}_i \rightarrow 0$. □

Lemma 24. Marginal Effect of Prices Consistency

Let $V_i = V_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i)$ be the observed outcomes, or single good excess demand and $v \in \{y, z_j\}$ for $j \in [J]$ be the corresponding expected value function. Let $v(\mathbf{p}) = \mathbb{E}[V_i(W_i, \mathbf{p})]$. Let \mathbf{V} be the n -dimensional vector of outcomes and \mathbf{U}_i be the $n \times J$ matrix of price perturbations. Let

$$\hat{\xi} = (\mathbf{U}^\top \mathbf{U})^{-1} (\mathbf{U}^\top \mathbf{V}).$$

Under the assumptions of Theorem 3 and Design 2, the regression of the observed data on price perturbations converges in probability to the partial derivative of expected supply or demand on prices:

$$\hat{\xi} \rightarrow_p \nabla_p v(\mathbf{p}_q^*).$$

Proof. By the Law of Large Numbers, as long as $\alpha < \frac{1}{2}$, we have that

$$\frac{\mathbf{U}^\top \mathbf{U}}{h^2 n} \rightarrow_p \mathbf{1}_J,$$

where $\mathbf{1}_J$ is the $J \times J$ identity matrix.

We next use the following Lemma, which we prove in Appendix B.

Lemma 25. *Let $V_i = V_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i)$ for some random function $V_i(w, \mathbf{p})$ that meets the conditions of Lemma 30. Let \mathbf{S}_i be a J -dimensional binary vector with $S_{ij} = 1$ if $U_{ij} = +h_n$ and $S_{ij} = 0$ if $U_{ij} = -h_n$. Under the assumptions of Theorem 3 and Design 2,*

$$\frac{\mathbf{S}^\top \mathbf{V}}{h_n n} = \frac{1}{h_n n} \sum_{i=1}^n \mathbf{S}_i V_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i) + O_p\left(\frac{1}{\sqrt{n}}\right)$$

and

$$\frac{\mathbf{S}^\top \mathbf{V}}{h_n n} \rightarrow_p \nabla_p v(\mathbf{p}_q^*).$$

Now, applying Lemma 25, we have that:

$$\frac{\mathbf{U}^\top \mathbf{V}}{nh_n^2} = \frac{\mathbf{S}^\top \mathbf{V}}{nh_n} \rightarrow_p \nabla_p v(\mathbf{p}_q^*).$$

Using Slutsky's, we have shown that

$$\hat{\xi} \rightarrow_p \nabla_p v(\mathbf{p}_q^*).$$

□

Consistency of Direct Effect Estimate

$$\hat{\tau}_{\text{ADE}} = \frac{1}{n} \sum_{i=1}^n \frac{W_i Y_i}{q} - \frac{(1 - W_i) Y_i}{1 - q}$$

is a Horvitz-Thompson estimate. Lemma 23 applies directly, so we have that

$$\hat{\tau}_{\text{ADE}} \rightarrow_p y(1, \mathbf{p}_q^*) - y(0, \mathbf{p}_q^*) = \tau_{\text{ADE}}^*.$$

Consistency of Indirect Effect Estimate

Proof. Notice that

$$\hat{\tau}_{\text{AIE}} = -\hat{\gamma}^\top \hat{\tau}_z^{\text{HT}}$$

where $\hat{\gamma} = [\hat{\xi}_z]^{-1} \hat{\xi}_y$.

Lemma 24 implies that:

$$\begin{aligned}\hat{\boldsymbol{\xi}}_y &\rightarrow_p \boldsymbol{\xi}_y, \\ \hat{\boldsymbol{\xi}}_z &\rightarrow_p \boldsymbol{\xi}_z\end{aligned}$$

Lemma 23 indicates that

$$\hat{\tau}_z^{HT} \rightarrow_p \mathbf{z}(1, \mathbf{p}_q^*) - \mathbf{z}(0, \mathbf{p}_q^*)$$

Combining these convergence results using Slutsky's theorem as well as Theorem 3, we have that:

$$\hat{\tau}_{\text{AIE}} \rightarrow_p -\boldsymbol{\xi}_y^\top \boldsymbol{\xi}_z^{-1} [\mathbf{z}(1, \mathbf{p}_q^*) - \mathbf{z}(0, \mathbf{p}_q^*)] = \tau_{\text{AIE}}^*.$$

□

A.7 Asymptotic Normality (Theorem 4 and 6)

We begin this section by showing that each of the differences in means and marginal response estimates are asymptotically normal.

Lemma 26. *Asymptotic Distribution of Horvitz-Thompson Estimator*

Let $V_i \in \{Y_i, Z_{ij}\}$ be one of outcomes, or excess demand for a single good $j \in [J]$, and let $v(\mathbf{p}) = \mathbb{E}[V_i(W_i, \mathbf{p})]$, $v(k, \mathbf{p}) = \mathbb{E}[V_i(k, \mathbf{p})]$. Under the assumptions of Theorem 3 and Design 2, we have that the HT estimator is asymptotically normal:

$$\sqrt{n}(\hat{\tau}^{HT} - v(1, \mathbf{p}_q^*) + v(0, \mathbf{p}_q^*)) \rightarrow_d N(0, \sigma_{HT}^2),$$

where

$$\begin{aligned}\sigma_D^2 &= \text{Var} \left[\left(\frac{W_i}{q} - \frac{1 - W_i}{1 - q} \right) V_i(W_i, \mathbf{p}_q^*) + Q_i \right] \\ Q_i &= -\nabla_p [v(1, \mathbf{p}_q^*) - v(0, \mathbf{p}_q^*)]^\top \boldsymbol{\xi}_z^{-1} \mathbf{Z}_i(W_i, \mathbf{p}_q^*)\end{aligned}\tag{24}$$

Proof. We can write the Horvitz-Thompson Estimator as :

$$\hat{\tau} = \frac{1}{n} \sum_{i=1}^n \frac{W_i V_i(1, \tilde{\mathbf{P}}_n + \mathbf{U}_i)}{q} - \frac{(1 - W_i) V_i(0, \tilde{\mathbf{P}}_n + \mathbf{U}_i)}{1 - q}$$

We start by working with the first component of the sum.

$$\begin{aligned}& \frac{1}{n} \sum_{i=1}^n \frac{W_i}{q} V_i(1, \tilde{\mathbf{P}}_n + \mathbf{U}_i) \\ & \stackrel{(1)}{=} \frac{1}{n} \sum_{i=1}^n \frac{W_i}{q} V_i(1, \mathbf{p}_q^* + \mathbf{U}_i) + (\tilde{\mathbf{P}}_n - \mathbf{p}_q^*)^\top \nabla_p v(1, \mathbf{p}_q^*) + o_p(n^{-0.5}) \\ & \stackrel{(2)}{=} \frac{1}{n} \sum_{i=1}^n \frac{W_i}{q} V_i(1, \mathbf{p}_q^* + \mathbf{U}_i) - [\boldsymbol{\xi}_z]^{-1} \nabla_p^\top v(1, \mathbf{p}_q^*) \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i) + o_p(n^{-0.5})\end{aligned}$$

(1) comes from the expansion in Lemma 30. (2) comes from asymptotically linear representation of $\tilde{\mathbf{P}}_n - \mathbf{p}_q^*$ from the proof of Theorem 2. By the same argument,

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \frac{1-W_i}{1-q} V_i(0, \tilde{\mathbf{P}}_n + \mathbf{U}_i) &= \frac{1}{n} \sum_{i=1}^n \frac{1-W_i}{1-q} V_i(0, \mathbf{p}_q^* + \mathbf{U}_i) \\ &\quad - [\boldsymbol{\xi}_z]^{-1} \nabla_p^\top v(0, \mathbf{p}_q^*) \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p}^* + \mathbf{U}_i) + o_p(n^{-0.5}) \end{aligned}$$

Combining these two and centering, letting $\mathbf{z}(\mathbf{p}) = \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p})]$, we can write:

$$\begin{aligned} \sqrt{n} \left(\hat{\tau}^{HT} - v(1, \mathbf{p}_q^*) + v(0, \mathbf{p}_q^*) \right) &= \frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{W_i}{q} V_i(1, \mathbf{p}_q^* + \mathbf{U}_i) - \mathbb{E}[V_i(1, \mathbf{p}_q^* + \mathbf{U}_i)] \\ &\quad - \frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{1-W_i}{1-q} V_i(0, \mathbf{p}_q^* + \mathbf{U}_i) + \mathbb{E}[V_i(0, \mathbf{p}_q^* + \mathbf{U}_i)] \\ &\quad - \mathbf{a}^\top [\boldsymbol{\xi}_z]^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n (\mathbf{Z}_i(W_i, \mathbf{p}^* + \mathbf{U}_i) - \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}^* + \mathbf{U}_i)]) \\ &\quad - \sqrt{n} \mathbf{a}^\top [\boldsymbol{\xi}_z]^{-1} (\mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i)] - \mathbf{z}(\mathbf{p}_q^*)) + \sqrt{n} (\mathbb{E}[V_i(1, \mathbf{p}_q^* + \mathbf{U}_i)] - v(1, \mathbf{p}_q^*)) \\ &\quad + \sqrt{n} (-\mathbb{E}[V_i(0, \mathbf{p}_q^* + \mathbf{U}_i)] + v(0, \mathbf{p}_q^*)) + o_p(1), \end{aligned}$$

where $\mathbf{a} = \nabla_p [v(1, \mathbf{p}_q^*) - v(0, \mathbf{p}_q^*)]$. Using Lemma 29, we have that each of $\sqrt{n} \mathbf{a}^\top [\boldsymbol{\xi}_z]^{-1} (\mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i)] - \mathbf{z}(\mathbf{p}_q^*))$, $\sqrt{n} (\mathbb{E}[V_i(1, \mathbf{p}_q^* + \mathbf{U}_i)] - v(1, \mathbf{p}_q^*))$ and $\sqrt{n} (-\mathbb{E}[V_i(0, \mathbf{p}_q^* + \mathbf{U}_i)] + v(0, \mathbf{p}_q^*))$ are $o_p(1)$. Also note that $\mathbf{z}(\mathbf{p}_q^*) = 0$.

Let

$$A_i(\mathbf{p} + \mathbf{U}_i) = \frac{W_i}{q} V_i(1, \mathbf{p} + \mathbf{U}_i) - \frac{1-W_i}{1-q} V_i(0, \mathbf{p} + \mathbf{U}_i) - \mathbf{a}^\top \boldsymbol{\xi}_z^{-1} \mathbf{Z}_i(W_i, \mathbf{p} + \mathbf{U}_i).$$

We have shown that

$$\sqrt{n} \left(\hat{\tau}^{HT} - v(1, \mathbf{p}_q^*) + v(0, \mathbf{p}_q^*) \right) = \frac{1}{\sqrt{n}} \sum_{i=1}^n [A_i(\mathbf{p}_q^* + \mathbf{U}_i) - \mathbb{E}[A_i(\mathbf{p}_q^* + \mathbf{U}_i)]] + o_p(1)$$

Given that A_i is i.i.d. and its variance is finite by the boundedness of Y_i and Z_i and the positive definiteness of $\boldsymbol{\xi}_z$, then the CLT implies that:

$$\sqrt{n} \frac{\left(\hat{\tau}^{HT} - v(1, \mathbf{p}_q^*) + v(0, \mathbf{p}_q^*) \right)}{\text{Var}[A_i(\mathbf{p}_q^* + \mathbf{U}_i)]} \rightarrow_d N(0, 1)$$

We can write the asymptotic variance as $\text{Var}[A_i(\mathbf{p}_q^*)]$ since $\mathbf{U}_i \rightarrow 0$ and, by the continuity almost everywhere in \mathbf{p} and boundedness of $A_i(\mathbf{p})$ and the continuity of $\mathbb{E}[A_i(\mathbf{p})]$, the variance of $A_i(\mathbf{p})$ is continuous in \mathbf{p} .

$$\sqrt{n} \left(\hat{\tau}^{HT} - v(1, \mathbf{p}_q^*) + v(0, \mathbf{p}_q^*) \right) \rightarrow_d N(0, \text{Var}[A_i(\mathbf{p}_q^*)]),$$

$$\text{where } A_i(\mathbf{p}_q^*) = \frac{W_i}{q} V_i(1, \mathbf{p}_q^*) - \frac{1-W_i}{1-q} V_i(0, \mathbf{p}_q^*) - \nabla_p [v(1, \mathbf{p}_q^*) - v(0, \mathbf{p}_q^*)]^\top \boldsymbol{\xi}_z^{-1} \mathbf{Z}_i(W_i, \mathbf{p} + \mathbf{U}_i).$$

□

We can now prove Theorem 4.

Asymptotic Normality of τ_{ADE}

$\hat{\tau}_{\text{ADE}}$ is a differences in means estimate when $\mathbf{U}_i = 0$, so we can conclude directly from Lemma 26 that it is asymptotically normal.

$$\sqrt{n}(\hat{\tau}_{\text{ADE}} - \tau_{\text{ADE}}) \rightarrow_d N(0, \sigma_D^2),$$

where

$$\begin{aligned} \sigma_D^2 &= \text{Var} \left[\left(\frac{W_i}{q} - \frac{1 - W_i}{1 - q} \right) Y_i(W_i, \mathbf{p}_q^*) + A_i \right] \\ A_i &= -\nabla_p [y(1, \mathbf{p}_q^*) - y(0, \mathbf{p}_q^*)]^\top \boldsymbol{\xi}_z^{-1} \mathbf{Z}_i(W_i, \mathbf{p}_q^*) \end{aligned}$$

Asymptotic Normality of τ_{AIE}

$$\sqrt{nh_n}(\hat{\tau}_{\text{AIE}} - \tau_{\text{AIE}}) = \sqrt{nh_n}(\hat{\boldsymbol{\gamma}}^\top \hat{\boldsymbol{\tau}}_z^{HT} - \boldsymbol{\gamma}^\top (\mathbf{z}(1, \mathbf{p}_q^*) - \mathbf{z}(0, \mathbf{p}_q^*))),$$

where we can write as an IV estimator, with $\mathbf{S} = \mathbf{U}/h_n$:

$$\begin{aligned} \hat{\boldsymbol{\gamma}} &= (\mathbf{S}^\top \mathbf{Z})^{-1} (\mathbf{S}^\top \mathbf{Y}) \\ \boldsymbol{\gamma} &= \boldsymbol{\xi}_y^\top \boldsymbol{\xi}_z^{-1}. \end{aligned}$$

First, we characterize the distribution of $\hat{\boldsymbol{\gamma}}$.

Lemma 27. *Under the assumptions of Theorem 3 and Design 2, the weak-IV estimator, appropriately scaled, has an asymptotically normal distribution.*

$$\sqrt{nh_n}(\hat{\boldsymbol{\gamma}} - \boldsymbol{\gamma}) \sim N(0, \boldsymbol{\xi}_z^{-1} v^2 \boldsymbol{\xi}_z^{-1}),$$

where $v^2 = \mathbb{E}[(Y_i(W_i, \mathbf{p}_q^*) - \mathbf{Z}_i(W_i, \mathbf{p}_q^*)^\top \boldsymbol{\gamma})^2]$.

Proof. We can write

$$Y_i = \mathbf{Z}_i \boldsymbol{\gamma} + R_i,$$

where $R_i = R_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i) = Y_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i) - \mathbf{Z}_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i)^\top \boldsymbol{\gamma}$. Plugging this in to the equation for $\hat{\boldsymbol{\gamma}}$, we have that

$$\begin{aligned} \hat{\boldsymbol{\gamma}} &= \left(\frac{\mathbf{S}^\top \mathbf{Z}}{nh_n} \right)^{-1} \left(\frac{\mathbf{S}^\top \mathbf{R}}{nh_n} \right) + \boldsymbol{\gamma}, \\ \sqrt{nh_n}(\hat{\boldsymbol{\gamma}} - \boldsymbol{\gamma}) &= \left(\frac{\mathbf{S}^\top \mathbf{Z}}{nh_n} \right)^{-1} \left(h_n \sqrt{n} \frac{\mathbf{S}^\top \mathbf{R}}{nh_n} \right). \end{aligned}$$

From the second part of Lemma 25, we have that the denominator converges in probability:

$$\left(\frac{\mathbf{S}^\top \mathbf{Z}}{nh_n} \right) \rightarrow_p \boldsymbol{\xi}_z. \quad (25)$$

Examining the numerator more closely, we can use the first part of Lemma 25 to write

$$\begin{aligned}
h_n \sqrt{n} \frac{\mathbf{S}^\top \mathbf{R}}{nh_n} &= \frac{h_n \sqrt{n}}{h_n n} \sum_{i=1}^n \mathbf{S}_i R_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i) + O_p(h_n) \\
&\stackrel{(1)}{=} \frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{S}_i R_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i) - \mathbb{E}[R_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i)] \\
&\quad + \sqrt{n}(\mathbb{E}[\mathbf{S}_i R_i(W_i, \mathbf{p}_q^*)] - \mathbb{E}[\mathbf{S}_i R_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i)]) + o_p(1) \\
&\stackrel{(2)}{=} \frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{S}_i R_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i) - \mathbb{E}[R_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i)] + o_p(1)
\end{aligned}$$

The expansion (1) is valid since $\mathbb{E}[\mathbf{S}_i R_i(W_i, \mathbf{p}_q^*)] = 0$, since $\mathbb{E}[\mathbf{S}_i] = 0$ and \mathbf{S}_i and $R_i(W_i, \mathbf{p}_q^*)$ are uncorrelated. The expansion (2) comes from Lemma 29, which indicates that $\sqrt{n}(\mathbb{E}[R_i(W_i, \mathbf{p}_q^*)] - \mathbb{E}[R_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i)]) = o_p(1)$. The CLT indicates that

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{\mathbf{S}_i R_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i) - \mathbb{E}[R_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i)]}{\text{Var}[(\mathbf{S}_i R_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i))]} \rightarrow_d N(0, 1)$$

$\text{Var}[(\mathbf{S}_i R_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i))] = \mathbb{E}[\mathbf{S}_i \mathbf{S}_i' R_i(W_i, \mathbf{p}_q^*)^2]$, by the CMT, since the variance is continuous in \mathbf{p} by the continuity almost everywhere and boundedness of $R_i(W_i, p)$ in p . Since $\mathbb{E}[\mathbf{S}_i] = 0$ and \mathbf{S}_i and $R_i(W_i, \mathbf{p}_q^*)$ are independent, we have that $\mathbb{E}[\mathbf{S}_i \mathbf{S}_i' R_i(W_i, \mathbf{p}_q^*)^2] = \mathbf{1}_J \mathbb{E}[R_i(W_i, \mathbf{p}_q^*)^2]$. We have now shown that

$$\sqrt{n} h_n \left(\frac{\mathbf{S}^\top \mathbf{R}}{nh_n} \right) \rightarrow_d N(0, \mathbf{1}_J \mathbb{E}[R_i(W_i, \mathbf{p}_q^*)^2])$$

Now applying Slutsky's theorem to combine this with the convergence in probability of the denominator in (25), this implies that the distribution of the IV estimator is:

$$\sqrt{n} h_n (\hat{\gamma} - \gamma) \rightarrow_d N(0, \Sigma),$$

$$\text{where } \Sigma = \xi_z^{-1} \mathbb{E}[(Y_i(W_i, \mathbf{p}_q^*) - \mathbf{Z}_i(W_i, \mathbf{p}_q^*)^\top \gamma)^2] \xi_z^{-1}.$$

□

From Lemma 26, we have that

$$\sqrt{n} (\hat{\tau}_z^{HT} - z(1, \mathbf{p}_q^*) + z(0, \mathbf{p}_q^*)) \rightarrow_d N(0, \Omega),$$

where $\Omega = \text{Var} \left[\left(\frac{W_i}{q} - \frac{1-W_i}{1-q} \right) \mathbf{Z}_i(W_i, \mathbf{p}_q^*) + \mathbf{B}_i \right]$, and

$$\mathbf{B}_i = -\nabla_p [z(1, \mathbf{p}_q^*) - z(0, \mathbf{p}_q^*)]^\top \xi_z^{-1} \mathbf{Z}_i(W_i, \mathbf{p}_q^*).$$

Then we use the delta method:

$$\sqrt{n} h_n (\hat{\tau}_{\text{AIE}} - \tau_{\text{AIE}}) \rightarrow_d N(0, \sigma_I^2),$$

$$\begin{aligned}
\sigma_I^2 &= [z(1, \mathbf{p}_q^*) - z(0, \mathbf{p}_q^*)]^\top \Sigma [z(1, \mathbf{p}_q^*) - z(0, \mathbf{p}_q^*)] + h_n^2 \gamma^\top \Omega \gamma \\
&= [z(1, \mathbf{p}_q^*) - z(0, \mathbf{p}_q^*)]^\top \Sigma [z(1, \mathbf{p}_q^*) - z(0, \mathbf{p}_q^*)] + o_p(1) \\
&= \mathbb{E}[(Y_i(W_i, \mathbf{p}^*) - \mathbf{Z}_i(W_i, \mathbf{p}^*)^\top \gamma)^2] \boldsymbol{\nu}^\top \boldsymbol{\nu},
\end{aligned}$$

where $\boldsymbol{\nu} = \boldsymbol{\xi}_z^{-1}[\mathbf{z}(1, \mathbf{p}_q^*) - \mathbf{z}(0, \mathbf{p}_q^*)]$.

A.8 Variance Estimators (Proof of Theorem 7)

Estimated Variance of Direct Effect Converges

$$\begin{aligned}\hat{\sigma}_D^2 &= \widehat{\text{Var}}[G_i], \\ G_i &= G_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i, \hat{\mathbf{a}}, \hat{\boldsymbol{\xi}}_z) \\ &= \frac{W_i}{q} Y_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i) - \frac{1 - W_i}{1 - q} Y_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i) - \hat{\mathbf{a}}^\top \hat{\boldsymbol{\xi}}_z^{-1} \mathbf{Z}_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i)\end{aligned}\tag{26}$$

and $\hat{\mathbf{a}} = \hat{\boldsymbol{\xi}}_{y1} - \hat{\boldsymbol{\xi}}_{y0}$. From Lemma 24, we have that:

$$\begin{aligned}\hat{\boldsymbol{\xi}}_{y1} - \hat{\boldsymbol{\xi}}_{y0} &\rightarrow_p \nabla_p y(1, \mathbf{p}_q^*) - \nabla_p y(0, \mathbf{p}_q^*) = \mathbf{a}, \\ \hat{\boldsymbol{\xi}}_z &\rightarrow_p \boldsymbol{\xi}_z.\end{aligned}$$

From Lemma 14, we have that $\tilde{\mathbf{P}}_n + \mathbf{U}_i \rightarrow_p \mathbf{p}_q^*$.

Let $\bar{G} = \frac{1}{n} \sum_{i=1}^n G_i$. From the expansion in Lemma 30 (given the composition rules of Donsker classes) and an application of the LLN and the CMT, we have that

$$\bar{G} \rightarrow_p \mathbb{E}[G_i(W_i, \mathbf{p}_q^*, \mathbf{a}, \boldsymbol{\xi}_z)].$$

Now we can work with $\hat{\sigma}_D^2$.

$$\begin{aligned}\hat{\sigma}_D^2 &= \frac{1}{n} \sum_{i=1}^n (G_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i, \hat{\mathbf{a}}, \hat{\boldsymbol{\xi}}_z) - \bar{G})^2 \\ &\stackrel{(1)}{=} \frac{1}{n} \sum_{i=1}^n (G_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i, \mathbf{a}, \boldsymbol{\xi}_z) - \mathbb{E}[G_i(W_i, \mathbf{p}_q^*, \mathbf{a}, \boldsymbol{\xi}_z)])^2 + o_p(1) \\ &\stackrel{(2)}{=} \frac{1}{n} \sum_{i=1}^n (G_i(W_i, \mathbf{p}_q^*, \mathbf{a}, \boldsymbol{\xi}_z) - \mathbb{E}[G_i(W_i, \mathbf{p}_q^*, \mathbf{a}, \boldsymbol{\xi}_z)])^2 + [\tilde{\mathbf{P}}_n - \mathbf{U}_i - \mathbf{p}_q^*]^\top \mathbf{D} + o_p(1) \\ &\stackrel{(3)}{=} \frac{1}{n} \sum_{i=1}^n (G_i(W_i, \mathbf{p}_q^*, \mathbf{a}, \boldsymbol{\xi}_z) - \mathbb{E}[G_i(W_i, \mathbf{p}_q^*, \mathbf{a}, \boldsymbol{\xi}_z)])^2 + o_p(1) \\ &= \text{Var}[G_i(W_i, \mathbf{p}_q^*, \mathbf{a}, \boldsymbol{\xi}_z)] + o_p(1) \\ &= \text{Var} \left[\frac{W_i}{q} Y_i(W_i, \mathbf{p}_q^*) - \frac{1 - W_i}{1 - q} Y_i(W_i, \mathbf{p}_q^*) - \nabla_p [y(1, \mathbf{p}_q^*) - y(0, \mathbf{p}_q^*)]^\top \boldsymbol{\xi}_z^{-1} \mathbf{Z}_i(W_i, \mathbf{p}_q^*) \right] \\ &= \sigma_D^2\end{aligned}$$

(1) comes from the CMT. (2) comes from the expansion in Lemma 30 given the composition rules of Donsker classes and that $Y_i(w, \mathbf{p})$ and $\mathbf{Z}_i(w, \mathbf{p})$ are continuous with probability 1 in \mathbf{p} and bounded. $\mathbf{D} = \nabla_p \mathbb{E}[(G_i(W_i, \mathbf{p}, \mathbf{a}, \boldsymbol{\xi}_z) - \mathbb{E}[G_i(W_i, \mathbf{p}, \mathbf{a}, \boldsymbol{\xi}_z)])^2]$, the derivatives exist by the differentiability of $\mathbb{E}[Y_i(w, \mathbf{p})]$, $\mathbb{E}[Y_i(w, \mathbf{p}) \mathbf{Z}_i(w, \mathbf{p})]$, and $\mathbb{E}[\mathbf{Z}_i(w, \mathbf{p})]$ in \mathbf{p} . (3) comes from the convergence in probability of $\tilde{\mathbf{P}}_n + \mathbf{U}_i$ to \mathbf{p}_q^* .

Estimated Variance of Indirect Effect Converges

$$\hat{\sigma}_I^2 = [\hat{\boldsymbol{\xi}}_z^{-1} \hat{\boldsymbol{\tau}}_z^{HT}]^\top [\hat{\boldsymbol{\xi}}_z^{-1} \hat{\boldsymbol{\tau}}_z^{HT}] \hat{v}^2,$$

where $\hat{v}^2 = \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{Z}_i^\top \hat{\boldsymbol{\gamma}})^2$ and $\hat{\boldsymbol{\gamma}} = \hat{\boldsymbol{\xi}}_z^{-1} \hat{\boldsymbol{\xi}}_y$.

From Lemma 25 and Lemma 23, we have that

$$\hat{\boldsymbol{\xi}}_z \rightarrow_p \boldsymbol{\xi}_z, \quad \hat{\boldsymbol{\xi}}_y \rightarrow_p \boldsymbol{\xi}_y, \quad \hat{\boldsymbol{\tau}}_z^{HT} \rightarrow_p \mathbf{z}(1, \mathbf{p}^*) - \mathbf{z}(0, \mathbf{p}^*).$$

Combining using Slutsky's, this implies that

$$\hat{\boldsymbol{\gamma}} = \boldsymbol{\gamma} + o_p(1).$$

By the composition rules for Donsker classes we have that $\{(Y_i(w, \mathbf{p} + \mathbf{u}) - \mathbf{Z}_i(w, \mathbf{p} + \mathbf{u})\boldsymbol{\gamma})^2 : \mathbf{p} \in \mathcal{S}\}$ forms a universal Donsker-class, so we can use Lemma 30 to show that

$$\begin{aligned} \hat{v}^2 &= \frac{1}{n} \sum_{i=1}^n (Y_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i) - \mathbf{Z}_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i) \hat{\boldsymbol{\gamma}})^2 \\ &= \frac{1}{n} \sum_{i=1}^n (Y_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i) - \mathbf{Z}_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i) \hat{\boldsymbol{\gamma}})^2 + o_p(1) \\ &= \mathbb{E}[(Y_i(W_i, \mathbf{p}^*) - \mathbf{Z}_i(W_i, \mathbf{p}^*))^\top \boldsymbol{\gamma}]^2 + o_p(1) \end{aligned}$$

where the second step uses the CMT and the LLN. Thus, we have shown that $\hat{v}^2 = v^2 + o_p(1)$. Combining using Slutsky's, we have shown that

$$\hat{\sigma}_I^2 \rightarrow \sigma_I^2.$$

A.9 Heterogeneous Effects (Proof of Proposition 8 and Theorem 9)

Proof of Proposition 8

Proof. The proof of Theorem 1 in Appendix B.1 of Hu et al. (2022) shows that under a general interference pattern that $\frac{\partial}{\partial \pi_k} \mathbb{E}_{\pi'} [Y_i(\mathbf{W})]_{\pi'=\pi} = \mathbb{E}_\pi [Y_i(w_k = 1; W_{-k}) - Y_i(w_k = 0; W_{-k})]$. Plugging this into our conditional expectation of interest,

$$\begin{aligned} n \mathbb{E} \left[\frac{\partial}{\partial \pi_k} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\pi [Y_i(\mathbf{W})] \middle| X_k = x \right] &= \mathbb{E} \left[\sum_{i=1}^n \mathbb{E}_\pi [Y_i(w_k = 1; W_{-k}) - Y_i(w_k = 0; W_{-k})] \middle| X_k = x \right] \\ &= \sum_{i=1}^n \mathbb{E} [Y_i(w_k = 1; W_{-k}) - Y_i(w_k = 0; W_{-k}) | X_k = x] \\ &= \mathbb{E} [Y_i(w_i = 1; W_{-k}) - Y_i(w_i = 0; W_{-i}) | X_i = x] \\ &\quad + \sum_{j \neq i} \mathbb{E} [Y_j(w_i = 1; W_{-i}) - Y_j(w_i = 0; W_{-i}) | X_i = x] \\ &= \tau_{\text{CADE}}^Y(x) + \tau_{\text{CAIE}}(x) \end{aligned}$$

□

Proof of Theorem 9

We use the following lemma, which are proved in the technical Appendix B.

Lemma 28. *Under the assumptions of Theorem 9, when the treatment of a single individual is fixed, the random finite sample price that arises converges in probability to the mean-field price:*

$$\mathbf{P}(W_i = k; \mathbf{W}_{-i}) \rightarrow_p \mathbf{p}_\pi^*$$

Now, we turn to the proof of the Theorem.

Proof. We can use the Continuous Mapping Theorem for the CADE convergence result. We have from Lemma 28 that $\mathbf{P}(W_i = 1; \mathbf{W}_{-i}) \rightarrow_p \mathbf{p}_\pi^*$ and $\mathbf{P}(W_i = 0; \mathbf{W}_{-i}) \rightarrow_p \mathbf{p}_\pi^*$. Furthermore, $y(w, \mathbf{p}, x)$ is continuous in \mathbf{p} . So, from the CMT, we have that:

$$\mathbb{E}[Y_i(W_i = 1; \mathbf{P}(W_i = 1; \mathbf{W}_{-i})) - Y_i(W_i = 0; \mathbf{P}(W_i = 0; \mathbf{W}_{-i})) | X_i = x] = y(1, \mathbf{p}_\pi^*, x) - y(0, \mathbf{p}_\pi^*, x) + \mathbb{E}[o_p(1)].$$

The limit of the expectation of the remainder term that is $o_p(1)$ is zero because the remainder term is uniformly bounded. We have shown that:

$$\lim_{n \rightarrow \infty} \mathbb{E}[Y_i(W_i = 1; \mathbf{P}(W_i = 1; \mathbf{W}_{-i})) - Y_i(W_i = 0; \mathbf{P}(W_i = 0; \mathbf{W}_{-i})) | X_i = x] = y(1, \mathbf{p}_\pi^*, x) - y(0, \mathbf{p}_\pi^*, x).$$

Proposition 8 indicates that

$$n \mathbb{E} \left[\frac{\partial}{\partial \pi_k} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\pi [Y_i(\mathbf{W})] \middle| X_k = x \right] = \tau_{\text{CADE}}^Y(x) + \tau_{\text{CAIE}}(x).$$

Thus, to finish the proof for the CAIE, it is enough to show that

$$\lim_{n \rightarrow \infty} n \mathbb{E} \left[\frac{\partial}{\partial \pi_k} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\pi [Y_i(\mathbf{W})] \middle| X_k = x \right] = \tau_{\text{CADE}}^*(x) + \tau_{\text{CAIE}}^*(x).$$

We can do this using similar techniques as in the unconditional case. From the proof of Theorem 3, we have the expansion

$$\begin{aligned} & \frac{\partial}{\partial \pi_k} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\pi [Y_i(W_i, \mathbf{P}_n(\mathbf{W}))] = \\ & \mathbb{E}_\pi \left[\frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} \left(\frac{1}{n} \sum_{i=1}^n Y_i(W_i, \mathbf{p}_\pi^*) - \boldsymbol{\xi}_y^\top \boldsymbol{\xi}_z^{-1} \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p}_\pi^*) + R_{1n} + R_{2n} + R_{3n} \right) \right]. \end{aligned}$$

where the remainder terms are defined in Lemma 21. Taking the conditional expectation with respect to x , we get

$$\begin{aligned} & = n \mathbb{E} \left[\frac{\partial}{\partial \pi_k} \frac{1}{n} \sum_{i=1}^n \mathbb{E}_\pi [Y_i(\mathbf{W})] \middle| X_k = x \right] \\ & = \mathbb{E} \left[\frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} \sum_{i=1}^n Y_i(W_i, \mathbf{p}_\pi^*) - \boldsymbol{\xi}_y^\top \boldsymbol{\xi}_z^{-1} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p}_\pi^*) \middle| X_k = x \right] + R_n \\ & = y(1, \mathbf{p}_\pi^*, x) - y(0, \mathbf{p}_\pi^*, x) - \boldsymbol{\xi}_y^\top \boldsymbol{\xi}_z^{-1} [z(1, \mathbf{p}_\pi^*, x) - z(0, \mathbf{p}_\pi^*, x)] + R_n \\ & = \tau_{\text{CADE}}^* + \tau_{\text{CAIE}}^* + R_n \end{aligned}$$

To complete the proof, we just need to confirm that $\lim_{n \rightarrow \infty} R_n = 0$, where $R_n = \sum_{m=1}^3 \mathbb{E} \left[n \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} R_{mn} | X_k = x \right]$. From Cauchy-Schwartz, we have that for each $m \in \{1, 2, 3\}$,

$$\begin{aligned} \left| \mathbb{E} \left[n \frac{W_k - \pi_k}{\pi_k(1 - \pi_k)} R_{mn} | X_k = x \right] \right| &\leq \sqrt{\mathbb{E}[nR_{mn}^2 | X_i = x]} \\ &= o(1) \end{aligned}$$

In the proof of Lemma 21, we showed that $\lim_{n \rightarrow \infty} \mathbb{E}[nR_{mn}^2] = 0$. Since $R_{mn}^2 \geq 0$, we this implies that $\lim_{n \rightarrow \infty} \mathbb{E}[nR_{mn}^2 | X_i = x] = 0$ for any x with non-zero density. \square

A.10 Optimal Targeting (Proof of Proposition 10 and Proposition 11)

Proposition 10

We first show that $\mathbf{p}_\pi^* = \mathbf{p}_q^*$ if and only if $\mathbb{E}[\pi(X_i)\mathbf{Z}_i(1, \mathbf{p}_q^*) + (1 - \pi(X_i))\mathbf{Z}_i(0, \mathbf{p}_q^*)] = \mathbb{E}[q\mathbf{Z}_i(1, \mathbf{p}_q^*) + (1 - q)\mathbf{Z}_i(0, \mathbf{p}_q^*)]$.

If $\mathbf{p}_\pi^* = \mathbf{p}_q^*$, then by this equality and the market-clearing condition,

$$\begin{aligned} 0 &= \mathbb{E}[\pi(X_i)\mathbf{Z}_i(1, \mathbf{p}_\pi^*) + (1 - \pi(X_i))\mathbf{Z}_i(0, \mathbf{p}_\pi^*)] \\ &= \mathbb{E}[\pi(X_i)\mathbf{Z}_i(1, \mathbf{p}_q^*) + (1 - \pi(X_i))\mathbf{Z}_i(0, \mathbf{p}_q^*)] \\ &= \mathbb{E}[q\mathbf{Z}_i(1, \mathbf{p}_q^*) + (1 - q)\mathbf{Z}_i(0, \mathbf{p}_q^*)]. \end{aligned}$$

To prove the converse direction, if

$$\mathbb{E}[\pi(X_i)\mathbf{Z}_i(1, \mathbf{p}_q^*) + (1 - \pi(X_i))\mathbf{Z}_i(0, \mathbf{p}_q^*)] = \mathbb{E}[q\mathbf{Z}_i(1, \mathbf{p}_q^*) + (1 - q)\mathbf{Z}_i(0, \mathbf{p}_q^*)],$$

and note that $\mathbb{E}[\pi(X_i)\mathbf{Z}_i(1, \mathbf{p}_q^*) + (1 - \pi(X_i))\mathbf{Z}_i(0, \mathbf{p}_q^*)] = 0$ by the market-clearing condition for \mathbf{p}_q^* . Then, by the uniqueness of the market-clearing price, we must have $\mathbf{p}_\pi^* = \mathbf{p}_q^*$ (there cannot be two different prices that both clear the market under treatment allocation $\pi(\cdot)$).

This means the constraint in optimization problem in Equation 11 can be replaced with the constraint

$$\mathbb{E}[\pi(X_i)\mathbf{Z}_i(1, \mathbf{p}_q^*) + (1 - \pi(X_i))\mathbf{Z}_i(0, \mathbf{p}_q^*)] = \mathbb{E}[q\mathbf{Z}_i(1, \mathbf{p}_q^*) + (1 - q)\mathbf{Z}_i(0, \mathbf{p}_q^*)].$$

We can transform the constraint by subtracting $\mathbb{E}[\mathbf{Z}_i(0, \mathbf{p}_q^*)]$ from both sides, and using the law of iterated expectations to get the constraint in the Proposition.

$$\mathbb{E}[\pi(X_i)\mathbb{E}[\mathbf{Z}_i(1, \mathbf{p}_q^*) - \mathbf{Z}_i(0, \mathbf{p}_q^*) | X_i = x]] = \mathbb{E}[q\mathbf{Z}_i(1, \mathbf{p}_q^*) - q\mathbf{Z}_i(0, \mathbf{p}_q^*)].$$

This is now a J -dimensional constraint that is linear in $\pi(x)$. For the value equation, we expand and drop a constant that does not depend on π . Maximizing $\mathbb{E}[Y_i(W_i, \mathbf{p}_\pi^*)] = \mathbb{E}[\pi(X_i)Y_i(1, \mathbf{p}_\pi^*) + (1 - \pi(X_i))Y_i(0, \mathbf{p}_\pi^*)]$ is equivalent to maximizing $\mathbb{E}[\pi(X_i)Y_i(1, \mathbf{p}_q^*) + (1 - \pi(X_i))Y_i(0, \mathbf{p}_q^*)]$ under the price equality constraint. Then, we can drop $\mathbb{E}[Y_i(0, \mathbf{p}_q^*)]$ which does not depend on $\pi(\cdot)$ and use to the law of iterated expectations to get the objective in the Proposition:

$$\mathbb{E}[\pi(X_i)\mathbb{E}[Y_i(1, \mathbf{p}_q^*) - Y_i(0, \mathbf{p}_q^*) | X_i = x]].$$

Proposition 11

First define some shorthand notation to make the exposition easier. For this proof, we write $Z_i(1) = Z_i(1, \mathbf{p}_q^*)$, $Z_i(0) = Z_i(0, \mathbf{p}_q^*)$, and

$$\delta_Z(x) = \mathbb{E}[Z_i(1, \mathbf{p}_q^*) - Z_i(0, \mathbf{p}_q^*) | X_i = x], \quad \delta_Y(x) = \mathbb{E}[Y_i(1, \mathbf{p}_q^*) - Y_i(0, \mathbf{p}_q^*) | X_i = x].$$

$V(c) = \mathbb{E}[\mathbb{1}(\delta_Y(X_i) > c\delta_Z(X_i))\delta_Y(X_i)]$ is maximized at 0 and is concave. $\beta(c) = \mathbb{E}[\mathbb{1}(\delta_Y(X_i) > c\delta_Z(X_i))\delta_Z(X_i)] - q\mathbb{E}[Z_i(1) - Z_i(0)]$ is non-increasing in c . Define $\bar{\beta} = \beta(0)$. $\bar{\beta}$ can be positive or negative. As c increases above 0, then those with $\delta_Z(x) > 0$ who were previously treated are controlled and those with $\delta_Z(x) < 0$ who were previously controlled are treated. The reverse happens as c decreases below 0. We can define the parameter c_q as follows:

$$c_q = \begin{cases} \inf \{c : \beta(c) = 0\} & \text{if } \bar{\beta} \geq 0 \\ \inf \{-c : \beta(c) = 0\} & \text{if } \bar{\beta} < 0 \end{cases}$$

We have that the optimal rule is $\pi^*(x) = \mathbb{1}(\delta_Y(x) \geq c_q\delta_Z(x))$. If $c_q = 0$, then $\pi^*(x)$ is optimal, since any other rule will treat individuals with a negative CATE or control those with a positive CATE, which is suboptimal when $c_q = 0$ satisfies the stability condition. For $c_q \neq 0$, we complete the proof by showing that any other rule $r(X_i)$ that meets the budget, so that $\mathbb{E}[r(X_i)\delta_Z(X_i)] - q\mathbb{E}[Z_i(1) - Z_i(0)] = 0$, has weakly lower expected outcomes. If $r(X_i)$ is different from $\pi^*(X_i)$, then when $\delta_Y(X_i) \geq c_q\delta_Z(X_i)$, then $\pi^*(X_i) \geq r(X_i)$. When $\delta_Y(X_i) < c_q\delta_Z(X_i)$, then $\pi^*(X_i) \leq r(X_i)$. This gives the following inequality

$$\begin{aligned} \mathbb{E}[\delta_Y(X_i)(\pi^*(X_i) - r(X_i))] &\geq c_q\mathbb{E}[\delta_Z(X_i)(\pi^*(X_i) - r(X_i))] \\ &\geq 0 \end{aligned}$$

where the last step is because both $r(X_i)$ and $\pi^*(X_i)$ satisfy the equilibrium stability condition. This completes the proof.

B Proofs of Technical Lemmas

B.1 Statement and Proofs of Lemmas 29 - 30

These two lemmas are used repeatedly throughout the proofs.

Lemma 29. *Let $f(\mathbf{p})$ be a function that is twice continuously differentiable in \mathbf{p} . Let \mathbf{U}_i be a random variable with $\mathbb{E}[\mathbf{U}_i] = 0$. \mathbf{U}_i is bounded by A_n , $|\mathbf{U}_i| \leq A_n$, and $A_n = O(n^{-\alpha})$, $\frac{1}{4} < \alpha < \frac{1}{2}$. We have that*

$$\mathbb{E}[f(\mathbf{p} + \mathbf{U}_i)] - f(\mathbf{p}) = o_p(n^{-0.5})$$

Proof. First, take a Taylor expansion of $f(\mathbf{p} + \mathbf{U}_i)$ around \mathbf{p} .

$$\mathbb{E}[f(\mathbf{p} + \mathbf{U}_i)] = f(\mathbf{p}) + \mathbb{E}[\mathbf{U}_i \nabla_{\mathbf{p}} f(\mathbf{p})] + \mathbb{E}[\mathbf{U}_i^\top \nabla_{\mathbf{p}}^2 f(\mathbf{p}) \mathbf{U}_i] + R_n$$

For the first order term, note that $\mathbb{E}[\mathbf{U}_i \nabla_{\mathbf{p}} f(\mathbf{p})] = 0$ since the gradient is bounded and $\mathbb{E}[\mathbf{U}_i] = 0$. For the second order term, we know that each element of the Hessian is bounded by B and that \mathbf{U}_i

is bounded by A . So, we have that

$$\begin{aligned} \left| \mathbb{E}[\mathbf{U}_i^\top \nabla_p^2 f(\mathbf{p}) \mathbf{U}_i] \right| &\leq A_n^2 B \\ &= O(n^{-2\alpha}) \\ &= o(n^{-0.5}). \end{aligned}$$

The last equality is because we have that $\alpha > 0.25$. This completes the result. \square

Lemma 30. *Let $F_i(\mathbf{p}) = f(\mathbf{p}, \theta_i)$ for random θ_i be a bounded random vector-valued function. The class of $f(\mathbf{p}, \theta)$ indexed by $p \in \mathcal{S}$ is a Donsker function class. $\mathbb{E}[F_i(\mathbf{p})]$ is twice continuously differentiable in \mathbf{p} with bounded derivatives, and $F_i(\mathbf{p})$ is continuous in \mathbf{p} with probability 1. Let \mathbf{P}_n be some random variable such that $\mathbf{P}_n \rightarrow_p \mathbf{p}^*$. Then, we have the following quadratic mean convergence:*

$$\mathbb{E}[(F_i(\mathbf{P}_n) - F_i(\mathbf{p}^*))^2] \rightarrow_p 0$$

The following expansion holds:

$$\frac{1}{n} \sum_{i=1}^n F_i(\mathbf{P}_n) = \frac{1}{n} \sum_{i=1}^n F_i(\mathbf{p}^*) + \mathbb{E}[F_i(\mathbf{P}_n)] - \mathbb{E}[F_i(\mathbf{p}^*)] + o_p(n^{-0.5}).$$

And, if we also have that $\mathbf{P}_n = \mathbf{p}^* + O_p(\sqrt{n})$ then the following expansion holds:

$$\frac{1}{n} \sum_{i=1}^n F_i(\mathbf{P}_n) = \frac{1}{n} \sum_{i=1}^n F_i(\mathbf{p}^*) + (\mathbf{P}_n - \mathbf{p}^*)^\top \nabla_p \mathbb{E}[F_i(\mathbf{p}^*)] + o_p(n^{-0.5}).$$

Proof. First, for the quadratic mean convergence, we show the function $\gamma(\mathbf{p}) = \mathbb{E}[(F_i(\mathbf{p}) - F_i(\mathbf{p}^*))^2]$ for any \mathbf{p} is continuous. Then, the result holds from the continuous mapping theorem. Write $F_i(\mathbf{p}) = f(\mathbf{p}, \theta_i)$ for random θ_i . Let the set of θ such that $f(\mathbf{p}, \theta)$ is discontinuous at \mathbf{p} be DC_p . The set of θ such that the function is continuous is C_p .

$$\begin{aligned} \mathbb{E}[(F_i(\mathbf{p}) - F_i(\mathbf{p}^*))^2] &= \int_{\theta} [f(\mathbf{p}, \theta) - f(\mathbf{p}^*, \theta)]^2 p(\theta) d\theta \\ &= \int_{\theta \in C_p} [f(\mathbf{p}, \theta) - f(\mathbf{p}^*, \theta)]^2 p(\theta) d\theta - \int_{\theta \in DC_p} [f(\mathbf{p}, \theta) - f(\mathbf{p}^*, \theta)]^2 p(\theta) d\theta \\ &= \int_{\theta \in C_p} [f(\mathbf{p}, \theta) - f(\mathbf{p}^*, \theta)]^2 p(\theta) d\theta \\ &= \gamma(\mathbf{p}) \end{aligned}$$

We can drop the integral over the discontinuous functions because it is a sum of bounded terms that happen with zero probability. By the dominated convergence theorem, which lets us exchanged the limit and the integral since $f(\mathbf{p}, \theta)$ is bounded, then we have that $\gamma(\mathbf{p})$ is continuous, since it is equal to the integral of functions each of which are continuous. We have now proved that the desired quadratic mean convergence holds.

Next, for the expansion. Given we have shown the quadratic mean convergence and that the function class is Donsker, we can use Lemma 19.24 of [van der Vaart \(1998\)](#) to show the first

expansion directly.

$$\sqrt{n} \left[\left(\frac{1}{n} \sum_{i=1}^n F_i(\mathbf{P}_n) - \mathbb{E}[F_i(\mathbf{P}_n)] \right) - \left(\frac{1}{n} \sum_{i=1}^n F_i(\mathbf{p}^*) - \mathbb{E}[F_i(\mathbf{p}^*)] \right) \right] \rightarrow_p 0,$$

which is equivalent to

$$\frac{1}{n} \sum_{i=1}^n F_i(\mathbf{P}_n) = \frac{1}{n} \sum_{i=1}^n F_i(\mathbf{p}^*) + \mathbb{E}[F_i(\mathbf{p})] - \mathbb{E}[F_i(\mathbf{p}^*)] + o_p(n^{-0.5}).$$

For the final expansion. Since we have that $\mathbb{E}[F_i(\mathbf{p})]$ is twice continuously differentiable in \mathbf{p} , we can use a Taylor expansion to write that

$$\mathbb{E}[F_i(\mathbf{P}_n)] = \mathbb{E}[F_i(\mathbf{p}^*)] + (\mathbf{P}_n - \mathbf{p}^*)^\top \nabla_p \mathbb{E}[F_i(\mathbf{p}^*)] + R_n$$

$R_n = o_p(n^{-0.5})$ since derivatives are bounded and $\mathbf{P}_n - \mathbf{p}^* = O_p(n^{-0.5})$. Plugging this into the first expansion, we have now shown that the second expansion holds:

$$\frac{1}{n} \sum_{i=1}^n F_i(\mathbf{P}_n) = \frac{1}{n} \sum_{i=1}^n F_i(\mathbf{p}^*) + (\mathbf{P}_n - \mathbf{p}^*)^\top \nabla_p \mathbb{E}[F_i(\mathbf{p}^*)] + o_p(n^{-0.5}).$$

□

B.2 Concentration Results: Proof of Lemma 21

We first state Lemma 31 - 33 which are proved at the end of this subsection.

Lemma 31. *If $\lim_{n \rightarrow \infty} \mathbb{E}[nR_n^2] = 0$, then $\mathbb{E}_\pi \left[\sum_{i=1}^N \frac{W_i - \pi_i}{\pi_i(1-\pi_i)} R_n \right] = o_p(1)$.*

Lemma 32. *If R_n meets the following conditions:*

1. $R_n = o_p(n^{-0.5})$,
2. $R_n \leq M_n$, and
3. $\limsup_{n \rightarrow \infty} n^2 \mathbb{E}[M_n^4] \leq \infty$,

then $\mathbb{E}[nR_n^2] \rightarrow 0$.

Lemma 33. *For bounded random function $F_i(\mathbf{p})$ which is an element of a Universal Donsker class indexed by \mathbf{p} , with expected value $\mathbb{E}[F_i(\mathbf{p})] = f(\mathbf{p})$. Then,*

$$\limsup_{n \rightarrow \infty} \mathbb{E} \left[\sup_{\mathbf{p} \in \mathcal{S}} \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n F_i(\mathbf{p}) - f(\mathbf{p}) \right)^4 \right] \leq C$$

for $C < \infty$.

We now use these to prove Lemma 21.

Proof. By Lemma 31, it is sufficient to prove that for each remainder term, $\lim_{n \rightarrow \infty} \mathbb{E}[nR_n^2] = 0$.

For the first, we can use Lemma 32. For R_{1n} , Condition 1 is satisfied by Lemma 19.24 of van der Vaart (1998). Let $y(p) = \mathbb{E}[Y_i(W_i, p)]$. Condition 2 is satisfied since $\sqrt{n}R_{1n} \leq C \sup_{p \in \mathcal{S}} \frac{1}{\sqrt{n}} \sum_{i=1}^n (Y_i(W_i, p) - y(p))$, which is the supremum of a centered and bounded empirical process. Lemma 33 provides the required bound on the fourth moment that satisfies Condition 3 of Lemma 32.

For R_{2n} , we can show directly. Let $z(p) = \mathbb{E}[Z_i(W_i, p)]$. Let C be a finite constant whose value we don't keep track of exactly and $Z_n(p) = \frac{1}{n} \sum_{i=1}^n Z_i(W_i, p)$.

$$\begin{aligned}
nR_{2n}^2 &\leq n \|\mathbf{P}_n - \mathbf{p}_\pi^*\|^4 \\
&\leq n \|\mathbf{z}(\mathbf{P}_n) - \mathbf{z}(\mathbf{p}_\pi^*)\|^4 C \\
&\stackrel{(1)}{=} n \|\mathbf{z}(\mathbf{P}_n)\|^4 C \\
&\leq nJ \max_{j \in \{1, \dots, J\}} \mathbb{E}[(z_j(\mathbf{P}_n))^4] C \\
&\stackrel{(2)}{\leq} Cn \mathbb{E}[(Z_{jn}(\mathbf{P}_n))^4] + Cn \mathbb{E}[(z_j(\mathbf{P}_n) - Z_{jn}(\mathbf{P}_n))^4] \\
&\stackrel{(3)}{\leq} o(1) + \frac{C}{n} \sup_{p \in \mathcal{S}} \max_{j \in \{1, \dots, j\}} \left(\sum_{i=1}^n \frac{1}{\sqrt{n}} (Z_{ij}(p) - z_j(p)) \right)^4 \\
&\stackrel{(4)}{\rightarrow} 0
\end{aligned}$$

The first line comes from a Taylor expansion of $\mathbf{z}(\mathbf{P}_n)$ around \mathbf{p}_π^* . (1) is because $\mathbf{z}(\mathbf{p}_\pi^*) = 0$. (2) is because $\mathbb{E}[(A+B)^4] \leq 16 \max\{\mathbb{E}[A^4], \mathbb{E}[B^4]\}$ with $A = Z_{jn}(\mathbf{P}_n)$ and $B = z_j(\mathbf{P}_n) - Z_{jn}(\mathbf{P}_n)$.

For (3), we have that the norm of $\mathbf{Z}_n(\mathbf{P}_n)$ is bounded by M , since each $\mathbf{Z}_i(p)$ is bounded. This means that for any j , $\mathbb{E}[(Z_{jn}(\mathbf{P}_n))^4] \leq M^2 \mathbb{E}[(Z_{jn}(\mathbf{P}_n))^2] = o(\frac{1}{n})$ by Assumption 3. We finish in (4) since we can show that for any j that $\limsup_{n \rightarrow \infty} \mathbb{E} \left[\sup_{p \in \mathcal{S}} \left(\sum_{i=1}^n \frac{1}{\sqrt{n}} (Z_{ij}(p) - z_j(p)) \right)^4 \right]$ is bounded using Lemma 33.

For

$$R_{3n} = \boldsymbol{\xi}_y^\top \left[(\mathbf{P}_n - \mathbf{p}_\pi^*) + \boldsymbol{\xi}_z^{-1} \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_i(W_i, \mathbf{p}_\pi^*) \right],$$

we again use Lemma 32. Condition 1 is satisfied by the asymptotically linear expansion in the proof of Theorem 2 given by Theorem 5.21 of van der Vaart (1998). We next show Condition 2 of Lemma 32, where we use C to represent a finite constant that we don't keep track of the exact

value. The first step uses the triangle inequality.

$$\begin{aligned}
R_{3n} &\leq C\|\mathbf{P}_n - \mathbf{p}_\pi^*\| + C\|\mathbf{Z}_n(\mathbf{p}_\pi^*)\| \\
&\stackrel{(1)}{\leq} C\|z(\mathbf{P}_n) - z(\mathbf{p}^*)\| + C\|\mathbf{Z}_n(\mathbf{p}_\pi^*)\| \\
&\stackrel{(2)}{\leq} C\|\mathbf{Z}_n(\mathbf{P}_n)\| + C\|z(\mathbf{P}_n) - \mathbf{Z}_n(\mathbf{P}_n)\| + C\|\mathbf{Z}_n(\mathbf{p}^*) - z(\mathbf{p}^*)\| \\
&\leq 2CJ \max_{j \in \{1, \dots, J\}} \sup \frac{1}{n} \sum_{i=1}^n Z_{ij}(W_i, \mathbf{p}) - z_j(\mathbf{p}) + C\|\mathbf{Z}_n(\mathbf{P}_n)\| \\
&= M_n
\end{aligned}$$

(1) is by the contraction property of $z(p)$ and (2) is via the triangle inequality. We can bound $n^2\mathbb{E}[M_n^4]$ by bounding the fourth moment of each of its two components.

$\limsup_{n \rightarrow \infty} n^2\|\mathbf{Z}_n(\mathbf{P}_n)\|^4 = 0$ is bounded by Assumption 3. Last, we can use Lemma 33 to show Condition 3. For each j ,

$$\limsup_{n \rightarrow \infty} \mathbb{E} \left[\left(\sup \frac{1}{\sqrt{n}} \sum_{i=1}^n Z_{ij}(W_i, \mathbf{p}) - z_j(\mathbf{p}) \right)^4 \right] \leq \infty$$

□

Proof of Lemma 33

Proof. Let M be an upper bound on $|F_i(\mathbf{p}) - f(\mathbf{p})|$ and B is an upper bound on $|F_i(\mathbf{p})|$. If we set $Z = \sup_{\mathbf{p} \in \mathcal{S}} \sum_{i=1}^n \frac{1}{\sqrt{n}} \left(\frac{F_i(\mathbf{p}) - f(\mathbf{p})}{M} \right)$, then Z is a centered empirical process where each element of the sum takes values between -1 and 1 . Theorem 12.2 of Boucheron et al. (2013) indicates that, where C_1 and C_2 are finite constants:

$$Pr(Z - E[Z] \geq t) \leq \exp\left(\frac{-t^2}{C_1 + 2t}\right).$$

since we have that $\mathbb{E} \left[\sup_{\mathbf{p} \in \mathcal{S}} \sum_{i=1}^n \left(\frac{F_i(\mathbf{p}) - f(\mathbf{p})}{\sqrt{n}M} \right)^2 \right] \leq 1$.

Then, when k is even, so that $(Z - E[Z])^k$ is always positive, this implies that:

$$\limsup_{n \rightarrow \infty} \mathbb{E}[(Z - E[Z])^k] \leq \int_0^\infty P((Z - E[Z])^k > t) dt \leq \int_0^\infty C_2 \exp(-t^{1/k}) < \infty$$

Then, using $\mathbb{E}[(A + B)^4] \leq 16 \max\{\mathbb{E}[(A)^4], \mathbb{E}[(B)^4]\}$

$$\mathbb{E}[Z^4] \leq C\mathbb{E}[(Z - \mathbb{E}[Z])^4] + C\mathbb{E}[Z]^4.$$

where C is a finite constant. We have already shown that the first term is bounded. From Chapter 19 of van der Vaart (1998) (see Lemma 19.35), we also have that $E[Z]$ is bounded since $F_i(\mathbf{p})$ is an element of a Donsker class.

Thus, we have shown that $\mathbb{E}[Z^4]$ is bounded, which completes our proof.

□

Proof of Lemma 32

Proof. Pick f_n, ϵ_n such that $nf_n \rightarrow 0, \epsilon_n \rightarrow 0$, and $P(R_n^2 > f_n) \leq \epsilon_n$. It is possible to choose these sequences by Condition 1.

$$\begin{aligned} \lim_{n \rightarrow \infty} n\mathbb{E}[R_n^2] &\leq \lim_{n \rightarrow \infty} nf_n + n\mathbb{E}[\mathbb{1}(R_n^2 \geq f_n)M_n^2] \\ &\leq \lim_{n \rightarrow \infty} \sqrt{P(R_n^2 \geq f_n)}\sqrt{n^2\mathbb{E}[M_n^4]} \\ &= 0 \end{aligned}$$

□

Proof of Lemma 31

Proof. Let $B_n = \sum_{i=1}^N \frac{W_i - \pi_i}{\pi_i(1-\pi_i)}$. This is the sum of n standardized Bernoulli random variables where $E_\pi[B_n^2] = n$. Using Cauchy-Schwarz for conditional expectations, we have that

$$\begin{aligned} ||\mathbb{E}_\pi[B_n R_n]|| &\leq \sqrt{\mathbb{E}_\pi[B_n^2]\mathbb{E}_\pi[R_n^2]} \\ &= \sqrt{n}\sqrt{\mathbb{E}_\pi[R_n^2]} \end{aligned}$$

If $n\mathbb{E}_\pi[R_n^2] \rightarrow_p 0$, then we have the desired result that $\mathbb{E}_\pi[B_n R_n] \rightarrow_p 0$.

We note that $n\mathbb{E}[R_n^2] = n\mathbb{E}[E_\pi[R_n^2]]$. The RHS is an expectation of a positive random variable. Then, we can use Markov's inequality and the squeeze theorem to finish the proof. Let $X = E_\pi[R_n^2]$. For any $a > 0$,

$$0 \leq P(X \geq a) \leq \frac{E[X]}{a}$$

Taking limits,

$$0 \leq \lim_{n \rightarrow \infty} P(X \geq a) \leq \lim_{n \rightarrow \infty} \frac{\mathbb{E}[X]}{a} = 0$$

This completes the proof since we have shown that $n\mathbb{E}_\pi[R_n^2] \rightarrow_p 0$.

□

B.3 Proof of Lemma 25

Proof. First, we use the expansion from Lemma 30,

$$\frac{1}{h_n n} \sum_{i=1}^n \mathbf{S}_i V_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i) = \frac{1}{h_n n} \sum_{i=1}^n \mathbf{S}_i V_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i) + (\tilde{\mathbf{P}}_n - \mathbf{p}_q^*) \frac{\mathbb{E}[\mathbf{S}_i V_i(\mathbf{p}_q^* + \mathbf{U}_i)]}{h_n} + o_p(n^{-0.5}h_n^{-1})$$

We can now use the following Lemma twice.

Lemma 34. *Let \mathbf{S}_i be a J -dimensional binary vector with $S_{ij} = 1$ if $\mathbf{U}_{ij} = +h_n$ and $S_{ij} = 0$ if $\mathbf{U}_{ij} = -h_n$. $h_n = cn^{-\alpha}$ with $\frac{1}{4} < \alpha < \frac{1}{2}$. Let $f(\mathbf{p})$ be a twice differentiable and bounded function in p .*

$$\lim_{n \rightarrow \infty} \frac{\mathbb{E}[\mathbf{S}_i f(\mathbf{p} + \mathbf{U}_i)]}{h_n} = \nabla_p f(\mathbf{p})$$

Since we have that $\tilde{\mathbf{P}}_n - \mathbf{p}_q^* = O_p(n^{-0.5})$ and from Lemma 34 that $\frac{\mathbb{E}[\mathbf{S}_i V_i(\mathbf{p}_q^* + \mathbf{U}_i)]}{h_n} = O(1)$, then $(\tilde{\mathbf{P}}_n - \mathbf{p}_q^*) \frac{\mathbb{E}[\mathbf{S}_i V_i(\mathbf{p}_q^* + \mathbf{U}_i)]}{h_n} = O_p(n^{-0.5}) = o_p(n^{-0.5} h_n^{-1})$. We can then write

$$\frac{1}{h_n n} \sum_{i=1}^n \mathbf{S}_i V_i(W_i, \tilde{\mathbf{P}}_n + \mathbf{U}_i) = \frac{1}{h_n n} \sum_{i=1}^n \mathbf{S}_i V_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i) + o_p(n^{-0.5} h_n^{-1})$$

To finish the proof, we have from the LLN

$$\frac{1}{h_n n} \sum_{i=1}^n \mathbf{S}_i V_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i) - \frac{\mathbb{E}[\mathbf{S}_i V_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i)]}{h_n} = o_p(1)$$

Then, we have from Lemma 34, using $F_i = qV_i(1, \mathbf{p}_q^* + \mathbf{U}_i) + (1 - q)V_i(0, \mathbf{p}_q^* + \mathbf{U}_i)$, that $\frac{\mathbb{E}[\mathbf{S}_i V_i(W_i, \mathbf{p}_q^* + \mathbf{U}_i)]}{h_n} = \nabla_{\mathbf{p}} v(\mathbf{p}_q^*) + o_p(1)$.

As a result, we can write.

$$\frac{\mathbf{S}^\top \mathbf{V}}{h_n n} = \nabla_{\mathbf{p}} v(\mathbf{p}_q^*) + o_p(1).$$

We finish the proof by proving Lemma 34. We examine the expectation of each term, where $F_i = f(\mathbf{p} + \mathbf{U}_i)$:

$$\frac{\mathbb{E}[\mathbf{S}_{ij} F_i]}{h_n} = \frac{\frac{1}{2} \mathbb{E}[f(\mathbf{p} + \mathbf{U}_i) | U_{ij} = h_n] - \frac{1}{2} \mathbb{E}[f(\mathbf{p} + \mathbf{U}_i) | U_{ij} = -h_n]}{h_n}.$$

Let

$$f_r(\mathbf{p}) = \frac{\partial^r f(\mathbf{p})}{\partial^r \mathbf{p}}$$

be the r -th partial derivative evaluated at \mathbf{p} . We can take a Taylor expansion of the expected outcome function, where \mathbf{e}_j is the J -length basis vector that is all 0s except for a 1 at position j :

$$\mathbb{E}_u[f(\mathbf{p} + \mathbf{U}_i) | U_{ij} = h_n] = f(\mathbf{p} + \mathbf{e}_j h) + \mathbb{E}_u[\mathbf{f}_1(\mathbf{p} + \mathbf{e}_j h_n)(\mathbf{U}_i - \mathbf{e}_j h_n) | U_{ij} = h_n] + O(h_n^2).$$

Note that the first order term drops out because $E[\mathbf{U}_{-i}] = 0$. We can write a similar expansion for $\mathbb{E}_u[f(\mathbf{p} + \mathbf{U}_i) | U_{ij} = -h_n]$. This implies that

$$E_u[f(\mathbf{p} + \mathbf{U}_i) | U_{ij} = h_n] - E_u[f(\mathbf{p} + \mathbf{U}_i) | U_{ij} = -h_n] = f(\mathbf{p} + \mathbf{e}_j h_n) - f(\mathbf{p} - \mathbf{e}_j h_n) + O(h_n^2).$$

We can now write

$$\frac{\mathbb{E}[\mathbf{S}_{ij} F_i]}{h_n} = \frac{f(\mathbf{p} + \mathbf{e}_j h_n) - f(\mathbf{p} - \mathbf{e}_j h_n)}{2h_n} + O(h_n).$$

Next, we use a Taylor expansion to derive the error of the centered difference approximation to the derivative. For $k \in \{0, 1\}$,

$$\frac{f(\mathbf{p} + \mathbf{e}_j h) - f(\mathbf{p} - \mathbf{e}_j h)}{2h_n} = \frac{\partial f(\mathbf{p})}{\partial p_j} + O(h_n^2).$$

This implies that:

$$\frac{\mathbb{E}[\mathbf{S}_{ij} F_i]}{h_n} = \frac{\partial f(\mathbf{p})}{\partial p_j} + o_p(1).$$

Stacking into a vector, we have that:

$$\frac{\mathbb{E}[\mathbf{S}_i F_i]}{h_n} = \nabla_p f(\mathbf{p}) + o_p(1).$$

□

B.4 Proof of Lemma 28

We have that $\mathbf{P}_n(W_i = k; W_{-i})$ is a Z-estimator that sets the criterion function $\tilde{\mathbf{Z}}_n(p)$ approximately equal to zero. Let $\mathbf{Z}_n(\mathbf{p}) = \frac{1}{n} \sum_{j=1}^n \mathbf{Z}_j(W_j, p)$ be the criterion function that $\mathbf{P}_n(\mathbf{W})$ approximately sets to zero. We can define

$$\tilde{\mathbf{Z}}_n(\mathbf{p}) = \mathbf{Z}_n(\mathbf{p}) - \frac{1}{n} \mathbb{1}(W_i \neq k) \mathbf{Z}_i(k, \mathbf{p}).$$

Let $\mathbf{z}(\mathbf{p}) = \mathbb{E}[\mathbf{Z}_i(W_i, \mathbf{p})]$. We use Theorem 5.9 of [van der Vaart \(1998\)](#) to show that $\mathbf{P}(W_i = k; W_{-i}) \rightarrow_p \mathbf{p}_\pi^*$, which requires verifying three conditions. For any $\epsilon > 0$,

$$\sup_{\mathbf{p} \in \mathcal{S}} \|\tilde{\mathbf{Z}}_n(\mathbf{p}) - \mathbf{z}(\mathbf{p}_\pi^*)\| \rightarrow_p 0 \quad (27)$$

$$\inf_{\phi: d(\mathbf{p}, \mathbf{p}_\pi^*) \geq \epsilon} \|\mathbf{z}(\mathbf{p})\| > 0 = \|\mathbf{z}(\mathbf{p}_\pi^*)\| \quad (28)$$

$$\tilde{\mathbf{Z}}_n(\mathbf{P}(W_i = k; W_{-i})) = o_p(1) \quad (29)$$

The first condition is easily verified since we already have the uniform convergence of $\mathbf{Z}_n(p)$ to $\mathbf{z}(p)$ from the proof of Lemma 14. For any $\epsilon > 0$, noting that excess demand is uniformly bounded by M ,

$$\begin{aligned} \lim_{n \rightarrow \infty} \sup_{\mathbf{p} \in \mathcal{S}} Pr(\|\tilde{\mathbf{Z}}_n(\mathbf{p}) - \mathbf{z}(\mathbf{p}_\pi^*)\| > \epsilon) &\leq \lim_{n \rightarrow \infty} \sup_{\mathbf{p} \in \mathcal{S}} Pr\left(\|\mathbf{Z}_n(\mathbf{p}) - \mathbf{z}(\mathbf{p}_\pi^*)\| + \left\|\frac{1}{n} \mathbb{1}(W_i \neq k) \mathbf{Z}_i(k, \mathbf{p})\right\| > \epsilon\right) \\ &\leq \lim_{n \rightarrow \infty} \sup_{\mathbf{p} \in \mathcal{S}} Pr(\|\mathbf{Z}_n(\mathbf{p}) - \mathbf{z}(\mathbf{p}_\pi^*)\| > \epsilon/2) + \lim_{n \rightarrow \infty} \mathbb{1}\left(\frac{M}{n} > \epsilon/2\right) \\ &\leq 0 + 0 \end{aligned}$$

The second condition has already been verified by the proof of Lemma 14. To verify the third condition is also straightforward, remembering that $\mathbf{Z}_n(\mathbf{P}_n(\mathbf{W})) = o_p(n^{-0.5})$ from Assumption 3,

$$\begin{aligned} \lim_{n \rightarrow \infty} Pr(\|\tilde{\mathbf{Z}}_n(\mathbf{P}(W_i = k; W_{-i}))\| > \epsilon) &\leq \lim_{n \rightarrow \infty} Pr(\|\tilde{\mathbf{Z}}_n(\mathbf{P}_n(\mathbf{W}))\| > \epsilon) \\ &\leq \lim_{n \rightarrow \infty} Pr(\|\mathbf{Z}_n(\mathbf{P}_n(\mathbf{W}))\| > \epsilon/2) + \mathbb{1}(M/n > \epsilon/2) \\ &= 0 \end{aligned}$$