Appendix for "Reducing Marketplace Interference Bias Via Shadow Prices"

Authors: Ido Bright, Arthur Delarue, Ilan Lobel

A Reducing Bias in Secondary Metrics

A.1 Secondary Metrics

In the main body of the paper, we assume that the platform's estimand of interest is the global treatment effect on system value, i.e., the difference between the total system value under global treatment $\Phi_\tau(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\boldsymbol{\beta}}, \boldsymbol{S}^\tau)$ and the total system value under global control $\Phi_\tau(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}}, \boldsymbol{S}^\tau)$. In practice, marketplace companies track more than one business metric to assess their performance. For instance, a ride-hailing company might match riders to drivers primarily based on pickup distance, but may also care about fulfilling more rides using electric vehicles. An example of such a setting is presented in Figure 12.

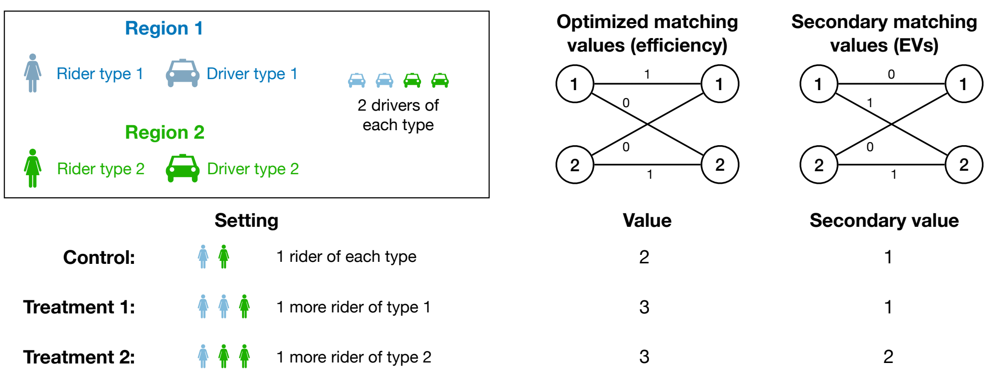

Figure 12: Example of a secondary metric. In a ride-hailing setting, riders and drivers originate from two regions, far enough apart that it is only efficient to match riders and drivers from the same region. Drivers in Region 2 drive electric vehicles (EVs). Therefore, though both treatment 1 and treatment 2 have the same global treatment effect, treatment 2 also improves the secondary metric which counts the number of rides satisfied with electric vehicles.

Even if the total system value is computed as a weighted combination of several metrics of interest, it remains of interest to understand how each component is affected by the treatment. For instance, in ride-hailing, a treatment that increases total welfare by both reducing pickup distance and increasing the number of rides served by electric vehicles may be preferable to a treatment that increases total welfare by the same amount, but by disproportionately increasing the number of rides served by electric vehicles at the expense of pickup distance.

Though it is not the case in the simple example in Figure 12, secondary metrics can also suffer from marketplace interference bias. In our ride-hailing example, a price incentive increasing demand arrival rate could lead to more riders in the treatment group matching with electric vehicles, seemingly suggesting that the price incentive improves sustainability metrics even though the total number of electric vehicle matches remains fixed. Therefore, it is of equal interest to reduce bias in secondary metrics as in the optimized value function.

However, solving the matching LP only provides shadow prices with respect to the original value function, as defined by the edge weights $v_{u,w}$. These shadow prices do not provide information regarding a secondary metric defined by new edge weights $\tilde{v}_{u,w}$. The goal of this section is to develop a method that extends the bias-reducing properties of shadow prices to secondary metrics of interest.

Formally, consider the matching problem defined in (1) with edge weights $v_{u,w}$, and consider a secondary metric defined by edge weights $\tilde{v}_{u,w}$. Let $x^*_{u,w}$ denote the (assumed unique) optimal primal solution of $\Phi_\tau(\boldsymbol{d}, \boldsymbol{s})$. We can define the value of the secondary metric in this setting as

$$\Phi_\tau^{\tilde{\boldsymbol{v}}}(\boldsymbol{d}, \boldsymbol{s}) = \sum_{(u,w)\in\mathcal{A}} \tilde{v}_{u,w} x^*_{u,w} \tag{21}$$

In our experimental setting, the estimand of interest is the *secondary* treatment effect $\Delta^{\tau,\boldsymbol{w}}$, defined as

$$\Delta^{\tau,\tilde{\boldsymbol{v}}} = \frac{1}{\tau}\mathbb{E}\left[\Phi_\tau^{\tilde{\boldsymbol{v}}}(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\boldsymbol{\beta}}, \boldsymbol{S}^\tau) - \Phi_\tau^{\tilde{\boldsymbol{v}}}(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}}, \boldsymbol{S}^\tau)\right]. \tag{22}$$

A.2 Secondary Estimators

The standard RCT estimator for the secondary metric $\tilde{\boldsymbol{v}}$ is given by:

$$\hat{\Delta}^{\tau,\tilde{\boldsymbol{v}}}_{\text{RCT}} = \frac{1}{\tau}\left(\frac{1}{\rho}\sum_{(u,w)\in\mathcal{A}} \tilde{v}_{u,w} X^{\tau,\text{treatment}}_{u,w} - \frac{1}{1-\rho}\sum_{(u,w)\in\mathcal{A}} \tilde{v}_{u,w} X^{\tau,\text{control}}_{u,w}\right), \tag{23}$$

where the optimal matching values $X^{\tau,\text{treatment}}_{u,w}$ and $X^{\tau,\text{control}}_{u,w}$ are computed with respect to the original matching values $v_{u,w}$.

In Section 4, we established that the RCT estimator for the primary metric always suffers from interference bias. The same is not true for the RCT estimator for secondary metrics. Because we assume no particular structure on the values $\tilde{\boldsymbol{v}}$, it is possible to reverse-engineer a secondary metric that does not suffer from interference bias in any marketplace configuration. In practice, metrics are typically not constructed in such an adverse manner, and may indeed suffer from interference bias, justifying our desire to define an analog of the shadow price estimator for secondary metrics. We first propose a way to extend the notion of a shadow price to a secondary metric.

Definition 8. *Let $\Phi_\tau^{\tilde{\boldsymbol{v}}}(\boldsymbol{d}, \boldsymbol{s})$ designate the value function for a secondary metric with edge weights $\tilde{v}_{i,j}$ for some demand vector $\boldsymbol{d}$ and supply vector $\boldsymbol{s}$. Then the secondary shadow price of demand* *type i is given by*

$$a_i^{\tilde{\boldsymbol{v}}} = \Phi_\tau^{\tilde{\boldsymbol{v}}}(\mathbf{d}, \mathbf{s}) - \Phi_\tau^{\tilde{\boldsymbol{v}}}(\mathbf{d} - \mathbf{e}_i, \mathbf{s}).$$

The secondary shadow prices above are defined analogously to the true shadow prices (see (18)). However, unlike the true shadow prices, it is not obvious how to obtain them directly from the matching LP. Applying the definition directly would require solving $n_d + 1$ optimization problems, which creates a much higher computational burden. However, it turns out that we can indeed compute these secondary duals much more efficiently by using the classic complementary slackness conditions.

Assuming the optimal solution of the matching problem is unique and nondegenerate, the complementary slackness conditions form a system of linear equations, which uniquely specifies the optimal shadow prices given the optimal primal solution. It turns out that, by replacing each value term $v_{u,w}$ in this linear system with the corresponding *secondary* value term $\tilde{v}_{u,w}$, we obtain a new linear system which uniquely specifies the *secondary* shadow prices. Therefore, rather than solving n_d additional optimization problems, we can simply solve one linear system. We formalize this result in the following proposition.

Proposition 8. *Assume that the optimal solution $\boldsymbol{x}^*$ to $\Phi_\tau(\boldsymbol{d}, \boldsymbol{s})$ is unique and nondegenerate. Then there exists an invertible matrix $\boldsymbol{M}$ and a vector $\tilde{\boldsymbol{v}}'$ such that*

$$a_i^{\tilde{\boldsymbol{v}}} = \left(\boldsymbol{M}^{-1}\tilde{\boldsymbol{v}}'\right)_i.$$

We provide an example of how to construct the linear system described in Proposition 8 in Figure 13. We first construct the linear system which uniquely determines the optimal dual solution from the optimal primal solution — then replace every occurrence of $v_{u,w}$ with $\tilde{v}_{u,w}$. The proof of Proposition 8 in Section D establishes why this construction is correct.

By solving the matching problem in the experiment state with total demand $\boldsymbol{D}^{\tau,\text{experiment}}$ and applying Proposition 8, we can obtain the secondary shadow prices $\boldsymbol{A}^{\tau,\tilde{\boldsymbol{v}},\text{experiment}}$. Using these marginal values, we can define a secondary version of the shadow price estimator:

$$\hat{\Delta}_{\text{SP}}^{\tau,\tilde{v}} = \frac{1}{\tau}\boldsymbol{A}^{\tau,\tilde{\boldsymbol{v}},\text{experiment}} \cdot \left(\frac{1}{\rho}\boldsymbol{D}^{\tau,\text{treatment}} - \frac{1}{1-\rho}\boldsymbol{D}^{\tau,\text{control}}\right). \tag{24}$$

B Analysis of bias in imbalanced markets

So far, we have analyzed the performance of the SP estimator in a general case. In practice, marketplaces may choose to modulate their behavior depending on macroscopic characteristics such as the balance of supply and demand. It is of interest to understand estimator performance in a range of marketplace balance settings. Intuitively, we expect both estimators to perform well in the limit where supply greatly outnumbers demand; however, the naive RCT estimator is unlikely

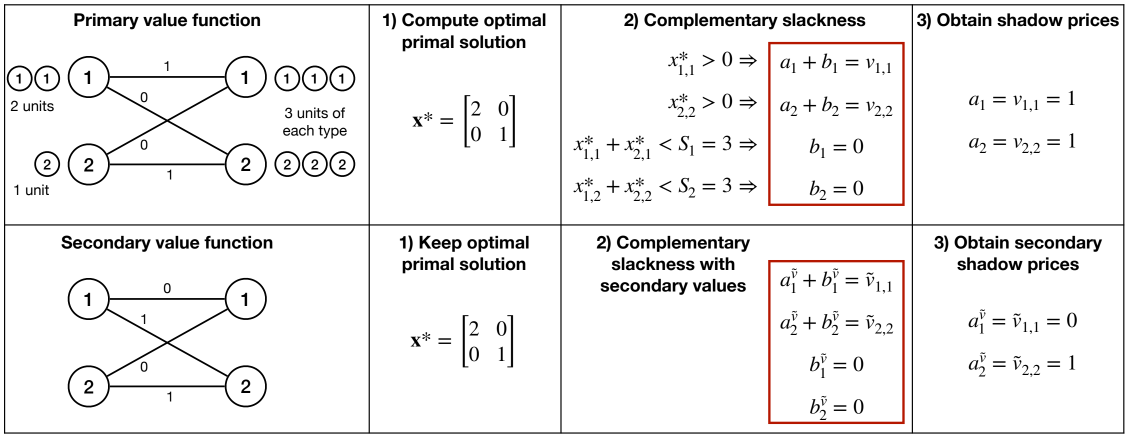

Figure 13: Example of secondary shadow price computation, using the bipartite matching graph from Figure 12. Given a particular supply and demand scenario, we can compute the optimal primal solution, then construct a linear system whose unique solution is the optimal dual solution. Replacing every occurrence of $v_{i,j}$ with $\tilde{v}_{i,j}$ yields a new linear system whose unique solution is the secondary dual solution.

to perform well when demand outperforms supply since it does not capture the notion of contention for scarce supply units.

For simplicity, we restrict our analysis of the undersupply and oversupply limits to the bipartite matching (uncapacitated) setting. We ignore capacities because we are trying to capture the effect of low or high supply relative to demand, and capacity introduces a third dimension that adds complexity but not insight.

We introduce scalar parameters $\bar{\lambda}$ (respectively $\bar{\pi}$) such that $\tilde{\boldsymbol{\lambda}} = \bar{\lambda}\boldsymbol{\alpha}$ with $\|\boldsymbol{\alpha}\|_1 \leq 1$ (respectively $\boldsymbol{\pi} = \bar{\pi}\boldsymbol{\gamma}$ with $\|\boldsymbol{\gamma}\|_1 \leq 1$). In other words, the total demand scales with the single parameter $\bar{\lambda}$ across all types, while the relative proportions of each type remain constant, parametrized by the vector $\boldsymbol{\alpha}$ (and similarly for supply). Recall that $\tilde{\boldsymbol{\lambda}}$ denotes the baseline arrival rate of demand to the platform; when factoring the request probability p_i of demand intent of type i, $\boldsymbol{\lambda} = \tilde{\boldsymbol{\lambda}} \odot \boldsymbol{p} = \bar{\lambda}\boldsymbol{\alpha} \odot \boldsymbol{p}$ denotes the demand arrival rate observed by the platform under global control (where $\odot$ refers to the element-wise vector product); and $\boldsymbol{\beta} = \tilde{\boldsymbol{\lambda}} \odot \boldsymbol{p} = \bar{\lambda}\boldsymbol{\alpha} \odot \boldsymbol{q}$ denotes the "extra" demand arrival rate observed by the platform under global treatment (which could be either negative or positive). As a result, the parameter $\bar{\lambda}$ proportionally scales both the baseline demand arrival rate and the treatment effect.

Using this framework, we can study the limits $\bar{\lambda}/\bar{\pi} \to 0$ (oversupply) and $\bar{\lambda}/\bar{\pi} \to \infty$ (undersupply). We formalize our results in the following theorem.

Theorem 11. *Consider the bipartite matching setting where $\Phi(\cdot) = \Phi^b(\cdot)$, and assume $\boldsymbol{\beta} \geq 0$.*

(a) In the oversupply limit, both the RCT and SP estimators are unbiased:

$$\lim_{\bar{\lambda}/\bar{\pi}\to 0}\left|\hat{\Delta}_{\mathrm{SP}}-\Delta\right| = 0 = \lim_{\bar{\lambda}/\bar{\pi}\to 0}\left|\hat{\Delta}_{\mathrm{RCT}}-\Delta\right|. \tag{25}$$

(b) In the undersupply limit, only the SP estimator is unbiased:

$$\lim_{\bar{\lambda}/\bar{\pi}\to \infty}\left|\hat{\Delta}_{\mathrm{SP}}-\Delta\right| = 0 \neq \lim_{\bar{\lambda}/\bar{\pi}\to \infty}\left|\hat{\Delta}_{\mathrm{RCT}}-\Delta\right|. \tag{26}$$

Theorem 11 is a direct consequence of the particular form of the partial-treatment value function. In the oversupply limit ($\bar{\lambda}/\bar{\pi} \to 0$), $\Psi(\cdot)$ is linear, and coincides exactly with $\hat{\Psi}_{\mathrm{RCT}}(\eta) = \bar{\boldsymbol{v}}^{*}\cdot(\boldsymbol{\lambda}+\eta\boldsymbol{\beta})$, so the RCT estimator is unbiased. However, in the undersupply limit, $\Psi(\cdot)$ is locally constant, and is not well approximated by $\hat{\Psi}_{\mathrm{RCT}}(\cdot)$, which passes through the origin. In contrast, the SP estimator is unbiased in both limits because the partial-treatment value function is linear and the SP estimator implicitly constructs the best linear approximation in a Taylor series sense.

C Analysis of estimators: proofs

C.1 Convergence to the fluid limit

Denoting ν_τ as the density-averaged total value with density τ, we seek to show that the average total value ν_τ converges to the fluid limit $\Phi(\boldsymbol{\lambda}, \boldsymbol{\pi})$ as τ tends to infinity. We begin with a simple scaling lemma.

Lemma 1. *For any demand vector $\boldsymbol{d} \in \mathbb{Z}_+^{n_d}$, supply vector $\boldsymbol{s} \in \mathbb{Z}_+^{n_s}$, and density τ, we can rewrite the total value as*

$$\frac{1}{\tau}\Phi_\tau(\boldsymbol{d}, \boldsymbol{s}) = \Phi\left(\frac{1}{\tau}\boldsymbol{d}, \frac{1}{\tau}\boldsymbol{s}\right). \tag{27}$$

Moreover, if $x^*_{u,w}$ and $z^*_{u,w}$ are the respective optimal primal solutions of the right-hand and left-hand sides, then $z^*_{u,w} = \tau x^*_{u,w}$.

Proof of Lemma 1. Following Eq. (2), we can write

$$\begin{aligned}
\Phi\left(\frac{1}{\tau}\boldsymbol{d}, \frac{1}{\tau}\boldsymbol{s}\right) = \ & \max \sum_{(u,w)\in\mathcal{A}} v_{u,w} x_{u,w}, \\
& \text{s.t.} \sum_{w:(w,u)\in\mathcal{A}} x_{w,u} \le \frac{1}{\tau} d_i && \forall u = u_i^d \in \mathcal{U}_d, \\
& \sum_{w:(u,w)\in\mathcal{A}} x_{u,w} \le \frac{1}{\tau} s_j && \forall u = u_j^s \in \mathcal{U}_s, \\
& \sum_{w:(u,w)\in\mathcal{A}} x_{u,w} - \sum_{w:(w,u)\in\mathcal{A}} x_{w,u} = 0 && \forall u \in \mathcal{U}_*, \\
& 0 \le x_{u,w} \le k_{u,w} && \forall (u,w) \in \mathcal{A}.
\end{aligned}$$

Applying a change of variables $z_{u,w} = \tau x_{u,w}$ yields

$$\begin{aligned}
\Phi\left(\frac{1}{\tau}\boldsymbol{d}, \frac{1}{\tau}\boldsymbol{s}\right) = \ & \frac{1}{\tau} \max \sum_{(u,w)\in\mathcal{A}} v_{u,w} z_{u,w}, \\
& \text{s.t.} \sum_{w:(w,u)\in\mathcal{A}} z_{w,u} \le d_i && \forall u = u_i^d \in \mathcal{U}_d, \\
& \sum_{w:(u,w)\in\mathcal{A}} z_{u,w} \le s_j && \forall u = u_j^s \in \mathcal{U}_s, \\
& \sum_{w:(u,w)\in\mathcal{A}} z_{u,w} - \sum_{w:(w,u)\in\mathcal{A}} z_{w,u} = 0 && \forall u \in \mathcal{U}_*, \\
& 0 \le z_{u,w} \le \tau k_{u,w} && \forall (u,w) \in \mathcal{A}, \\
& = \frac{1}{\tau} \Phi_\tau(\boldsymbol{d}, \boldsymbol{s}).
\end{aligned}$$

□

Lemma 1 provides an equivalent definition of the total value in the fluid limit as

$$\nu = \lim_{\tau\to\infty} \mathbb{E}\left[\frac{1}{\tau}\Phi_\tau(\mathbf{D}^\tau, \mathbf{S}^\tau)\right] = \lim_{\tau\to\infty} \mathbb{E}\left[\Phi\left(\frac{1}{\tau}\mathbf{D}^\tau, \frac{1}{\tau}\mathbf{S}^\tau\right)\right].$$

Before proving Theorem 1, we introduce a second useful lemma.

Lemma 2. *For $n \in \mathbb{N}$, let $X_n \sim P(\lambda n)$, and $Y_n = \frac{1}{n} X_n$ then $Z = \max_{n=1}^{\infty} Y_n$ has a finite mean.*

Proof of Lemma 2. For any Poisson variable $V \sim P(\lambda_0)$ we can bound the tail distribution as follows [Janson et al., 2000, Corollary 2.4]:

$$P(V \ge t) \le \exp(-t) \quad \forall t \ge 7\lambda_0.$$

Thus for all $t > t_0 = \max(7\lambda, \log 2)$

$$\begin{aligned} P(Z \geq t) &= P(\exists n \text{ s.t. } Y_n \geq t) \\ &\leq \sum_{n=1}^{\infty} P(Y_n \geq t) \\ &= \sum_{n=1}^{\infty} P(X_n \geq nt) \\ &\leq \sum_{n=1}^{\infty} \exp(-nt) \\ &= \exp(-t) / (1 - \exp(-t)) \\ &\leq 2\exp(-t) \end{aligned}$$

Applying this to $\mathbb{E}[z] = \int_0^\infty P(z \geq t)\, dt$ we see that

$$\begin{aligned} \mathbb{E}[z] &= \int_0^{t_0} P(z > t)\, dt + \int_{t_0}^{\infty} P(z > t)\, dt \\ &\leq t_0 + 2\int_{t_0}^{\infty} \exp(-t)\, dt < \infty \end{aligned}$$

□

We are now ready to prove Theorem 1.

Proof of Theorem 1. For simplicity, in this proof we denote $\boldsymbol{D}^{\tau,\boldsymbol{\lambda}}$ by $\boldsymbol{D}^{\tau}$ since we only consider a single arbitrary $\boldsymbol{\lambda}$.

Part 1: convergence of expected total value. We first use the strong law of large numbers, which implies that $\frac{1}{\tau}\mathbf{D}^{\tau} \to \boldsymbol{\lambda}$ and $\frac{1}{\tau}\mathbf{S}^{\tau} \to \boldsymbol{\pi}$ almost surely. We then define

$$Z = \max_{\tau=1}^{\infty} \frac{1}{\tau} Y_{\tau}, \text{ with } Y_{\tau} = \sum_{i=1}^{n_d} D_i^{\tau},$$

and note that Y_τ is a sum of Poisson random variables and is therefore Poisson-distributed with parameter $\tau \sum_{i=1}^{n_d} \lambda_i$. By Lemma 2 we know Z has bounded mean, and furthermore we notice that

$$\Phi\left(\frac{1}{\tau}\mathbf{D}^{\tau}, \frac{1}{\tau}\mathbf{S}^{\tau}\right) < Z \max_{i,j,p} \nu_{i,j,p}.$$

We finally apply the dominated convergence theorem to imply convergence of the expectation which is the desired result.

Part 2: almost sure convergence of optimal decision variables. For simplicity, we assume the optimal primal and dual solutions of $\Phi(\boldsymbol{\lambda}, \boldsymbol{\pi})$ are unique. If there are multiple primal (or dual) solutions, our proof below implies that every converging sub-sequence will tend to *an* optimal primal (or dual) solution of the fluid limit.

Step 1: convergence of dual variables. We first note that for every τ, the optimal dual variables $(\boldsymbol{A}^\tau, \boldsymbol{B}^\tau, \boldsymbol{M}^\tau, \Xi^\tau)$, as well as $(\boldsymbol{a}^*, \boldsymbol{b}^*, \boldsymbol{m}^*, \boldsymbol{\xi}^*)$, must be extreme points of the dual polyhedron (5), of which there are finitely many.

Assume we can find an infinite subsequence $\tau_1, \tau_2, \dots$ such that $(\boldsymbol{A}^{\tau_i}, \boldsymbol{B}^{\tau_i}, \boldsymbol{M}^{\tau_i}, \Xi^{\tau_i}) = (\boldsymbol{a}', \boldsymbol{b}', \boldsymbol{m}', \boldsymbol{\xi}') \neq (\boldsymbol{a}^*, \boldsymbol{b}^*, \boldsymbol{m}^*, \boldsymbol{\xi}^*)$, a distinct extreme point of the dual polyhedron, implying that

$$\boldsymbol{a}' \cdot \frac{\boldsymbol{D}^{\tau_i}}{\tau_i} + \boldsymbol{b}' \cdot \frac{\boldsymbol{S}^{\tau_i}}{\tau_i} + \boldsymbol{\xi}' \cdot \boldsymbol{k} < \boldsymbol{a}^* \cdot \frac{\boldsymbol{D}^{\tau_i}}{\tau_i} + \boldsymbol{b}^* \cdot \frac{\boldsymbol{S}^{\tau_i}}{\tau_i} + \boldsymbol{\xi}^* \cdot \boldsymbol{k}$$

By the strong law of large numbers, the dual objective values of the subsequence converge almost surely to $\boldsymbol{a}' \cdot \boldsymbol{\lambda} + \boldsymbol{b}' \cdot \boldsymbol{\pi} + \boldsymbol{\xi}' \cdot \boldsymbol{k} < \boldsymbol{a}^* \cdot \boldsymbol{\lambda} + \boldsymbol{b}^* \cdot \boldsymbol{\pi} + \boldsymbol{\xi}^* \cdot \boldsymbol{k}$, which contradicts the optimality of $(\boldsymbol{a}^*, \boldsymbol{b}^*, \boldsymbol{m}^*, \boldsymbol{\xi}^*)$ in $\Phi(\boldsymbol{\lambda}, \boldsymbol{\pi})$. Thus there exists $\tau_0 > 0$ such that $(\boldsymbol{a}^*, \boldsymbol{b}^*, \boldsymbol{m}^*, \boldsymbol{\xi}^*)$ is the optimal dual solution of $\Phi\left(\frac{1}{\tau}\mathbf{D}^\tau, \frac{1}{\tau}\mathbf{S}^\tau\right)$ for every $\tau > \tau_0$.

Step 2: convergence of primal variables. For $\tau > \tau_0$, the optimal primal solution is defined by a system of linear equations obtained from the complementary slackness and primal feasibility conditions:

$$\begin{aligned}
X_{u_j^s, u_i^d} &= 0 && \forall (u_j^s, u_i^d) \in \mathcal{A}_{s\to d},\ b_j^* + a_i^* > v_{u_j^s, u_i^d} \\
X_{u_j^s, w} &= 0 && \forall (u_j^s, w) \in \mathcal{A}_{s\to *},\ b_j^* - m_w^* > v_{u_j^s, w} \\
X_{w, u_i^d} &= 0 && \forall (w, u_i^d) \in \mathcal{A}_{*\to d},\ m_w^* + a_i^* > v_{w, u_i^d} \\
X_{u,w} &= 0 && \forall (u,w) \in \mathcal{A}_{*},\ m_u^* - m_w^* > v_{u,w} \\
X_{u,w} &= k_{u,w} && \forall (u,w) \in \mathcal{A},\ \xi_{u,w}^* > 0 \\
\sum_{w:(w,u)\in\mathcal{A}} X_{w,u} &= \frac{1}{\tau} D_i^\tau && \forall u = u_i^d \in \mathcal{U}_d,\ a_i^* > 0 \\
\sum_{w:(u,w)\in\mathcal{A}} X_{u,w} &= \frac{1}{\tau} S_j^\tau && \forall u = u_j^s \in \mathcal{U}_s,\ b_j^* > 0 \\
\sum_{w:(u,w)\in\mathcal{A}} X_{u,w} - \sum_{w:(w,u)\in\mathcal{A}} X_{w,u} &= 0 && \forall u \in \mathcal{U}_*
\end{aligned}$$

The solution of these linear equations is continuous on the inputs $\frac{1}{\tau}\boldsymbol{D}^\tau$ and $\frac{1}{\tau}\boldsymbol{S}^\tau$, both of which converge strongly to $\boldsymbol{\lambda}, \boldsymbol{\pi}$. By continuity $X_{u,w}^\tau$ also converges strongly to $x_{u,w}^*$. □

Proof of Proposition 1. This result follows directly from Chapter 5 of Bertsimas and Tsitsiklis [1997]. Theorem 5.1 proves concavity, and piece-wise linearity with finitely many pieces is established in the paragraphs following Theorem 5.1. □

C.2 Analysis of the RCT estimator

Proof of Proposition 2. We assume without loss of generality that $D_i^\tau > 0$ for each demand type i, otherwise we can simply remove this demand type. In the proof, we first assume for clarity that $D_i^{\tau,\text{treatment}} > 0$ and $D_i^{\tau,\text{control}} > 0$.

We can rewrite each term in Definition 2 as follows:

$$\begin{aligned}\sum_{i=1}^{n_d}\sum_{j=1}^{n_s}\sum_{p=1}^{n_{i,j}} \nu_{i,j,p} Y_{i,j,p}^{\tau,\text{treatment}} &= \sum_{i=1}^{n_d} \frac{\sum_{j=1}^{n_s}\sum_{p=1}^{n_{i,j}} Y_{i,j,p}^{\tau,\text{treatment}} \nu_{i,j,p}}{D_i^{\tau,\text{treatment}}} D_i^{\tau,\text{treatment}} \\ &= \bar{\boldsymbol{V}}^{\tau,\text{treatment}} \cdot \boldsymbol{D}^{\tau,\text{treatment}}, \end{aligned} \tag{28}$$

and similarly

$$\sum_{i=1}^{n_d}\sum_{j=1}^{n_s}\sum_{p=1}^{n_{i,j}} \nu_{i,j,p} Y_{i,j,p}^{\tau,\text{control}} = \bar{\boldsymbol{V}}^{\tau,\text{control}} \cdot \boldsymbol{D}^{\tau,\text{control}}, \tag{29}$$

where $\bar{V}_i^{\tau,\text{treatment}}$ (resp. $\bar{V}_i^{\tau,\text{control}}$) designates the average value obtained from demand type i from the treatment (resp. control) group.

Taking the expectation over treatment-control assignments preserves the total number of units in the treatment and control groups. Therefore,

$$\begin{aligned}\mathbb{E}[\bar{V}_i^{\tau,\text{control}} | \boldsymbol{D}^{\tau,\text{control}}, \boldsymbol{D}^{\tau,\text{treatment}}, \boldsymbol{S}^\tau] &= \mathbb{E}\left[\frac{\sum\limits_{j=1}^{n_s}\sum\limits_{p=1}^{n_{i,j}} \nu_{i,j,p} Y_{i,j,p}^{\tau,\text{control}}}{D_i^{\tau,\text{control}}} | \boldsymbol{D}^{\tau,\text{control}}, \boldsymbol{D}^{\tau,\text{treatment}}, \boldsymbol{S}^\tau\right] \\ &= \sum_{j=1}^{n_s}\sum_{p=1}^{n_{i,j}} \frac{\nu_{i,j,p}}{D_i^{\tau,\text{control}}} \mathbb{E}\left[Y_{i,j,p}^{\tau,\text{control}} | \boldsymbol{D}^{\tau,\text{control}}, \boldsymbol{D}^{\tau,\text{treatment}}, \boldsymbol{S}^\tau\right].\end{aligned}$$

By Assumption 1, among the $Y_{i,j,p}^\tau$ demand units of type i matched with supply type j along path p, each one is assigned to the control group with probability $D_i^{\tau,\text{control}}/D_i^\tau$. Therefore we conclude

$$\begin{aligned}\mathbb{E}[\bar{V}_i^{\tau,\text{control}} | \boldsymbol{D}^{\tau,\text{control}}, \boldsymbol{D}^{\tau,\text{treatment}}, \boldsymbol{S}^\tau] &= \sum_{j=1}^{n_s}\sum_{p=1}^{n_{i,j}} \frac{\nu_{i,j,p}}{D_i^{\tau,\text{control}}} Y_{i,j,p}^\tau \frac{D_i^{\tau,\text{control}}}{D_i^\tau} \\ &= \sum_{j=1}^{n_s}\sum_{p=1}^{n_{i,j}} \frac{Y_{i,j,p}^\tau \nu_{i,j,p}}{D_i^\tau} := \bar{V}_i^\tau,\end{aligned}$$

with the same proof applying to the treatment group. So far, we have assumed that both $\boldsymbol{D}^{\tau,\text{treatment}} > 0$ and $\boldsymbol{D}^{\tau,\text{control}} > 0$. If we assume that one of them is 0, then the right-hand-side in either (28) or

(29) is 0, and the rest of the proof still holds. Therefore, we can write

$$\mathbb{E}\left[\hat{\Delta}^{\tau}_{\mathrm{RCT}} \mid \boldsymbol{D}^{\tau,\mathrm{control}}, \boldsymbol{D}^{\tau,\mathrm{treatment}}, \boldsymbol{S}^{\tau}\right] = \bar{\boldsymbol{V}}^{\tau} \cdot \left(\frac{1}{\rho}\frac{1}{\tau}\boldsymbol{D}^{\tau,\mathrm{treatment}} - \frac{1}{1-\rho}\frac{1}{\tau}\boldsymbol{D}^{\tau,\mathrm{control}}\right).$$

As $\tau \to \infty$, we already know from the strong law of large numbers that $\boldsymbol{D}^{\tau,\mathrm{treatment}}/\tau$ converges to $\boldsymbol{\lambda}^{\mathrm{treatment}}$ (and $\boldsymbol{D}^{\tau,\mathrm{control}}/\tau$ to $\boldsymbol{\lambda}^{\mathrm{control}}$). Furthermore, Theorem 1 implies that $\bar{\boldsymbol{V}}^{\tau}$ converges to $\bar{\boldsymbol{v}}^{*}$. Once again applying Lemma 2 and the dominated convergence theorem, we obtain

$$\lim_{\tau\to\infty} \mathbb{E}\left[\hat{\Delta}^{\tau}_{\mathrm{RCT}} \mid \boldsymbol{D}^{\tau,\mathrm{control}}, \boldsymbol{D}^{\tau,\mathrm{treatment}}, \boldsymbol{S}^{\tau}\right] = \bar{\boldsymbol{v}}^{*} \cdot \left(\frac{1}{\rho}\rho(\boldsymbol{\lambda}+\boldsymbol{\beta}) - \frac{1}{1-\rho}(1-\rho)\boldsymbol{\lambda}\right) = \bar{\boldsymbol{v}}^{*} \cdot \boldsymbol{\beta}.$$

□

Notice that the proof of Proposition 2 analyzes the estimator $\mathbb{E}\left[\hat{\Delta}^{\tau}_{\mathrm{RCT}} \mid \boldsymbol{D}^{\tau,\mathrm{control}}, \boldsymbol{D}^{\tau,\mathrm{treatment}}, \boldsymbol{S}^{\tau}\right]$ in lieu of $\hat{\Delta}^{\tau}_{\mathrm{RCT}}$. Under Assumption 1, we can always compute the former estimator as a lower-variance version of the latter estimator. In the main text, we discuss the latter estimator since it is more intuitive. However, in our results, we analyze the version with variance reduction so that we are always comparing our estimator to to the best possible standard estimator (in terms of variance).

Proof of Theorem 2. For some demand vector $\boldsymbol{d}$ and supply vector $\boldsymbol{s}$, let $\mathcal{Y}(\boldsymbol{d}, \boldsymbol{s})$ denote the set of feasible solutions in path-aggregated form to Eq. (2) with $\tau = 1$ (fluid limit).

Assume without loss of generality that $\boldsymbol{\beta} \geq \boldsymbol{0}$. We first note that the optimal matching value for the experiment state verifies $\Psi(\rho) = \Phi(\boldsymbol{\lambda} + \rho\boldsymbol{\beta}, \boldsymbol{\pi}) = \bar{\boldsymbol{v}}^{*} \cdot (\boldsymbol{\lambda} + \rho\boldsymbol{\beta})$, which follows from the definition of $\bar{\boldsymbol{v}}^{*}$.

Step 1: Starting from the optimal solution to the matching problem in the experiment state $x^{*}_{i,j}$, we build a "scaled" solution $y_{i,j}$ to the matching problem under global control, such that

$$z_{i,j,p} = \frac{\lambda_i}{\lambda_i + \rho\beta_i} y^{*}_{i,j,p}.$$

We can easily verify that $z_{i,j,p} \in \mathcal{Y}(\boldsymbol{\lambda}, \boldsymbol{\pi})$ (feasibility), since:

$$\sum_{j=1}^{n_s}\sum_{p=1}^{n_{i,j}} z_{i,j,p} = \frac{\lambda_i}{\lambda_i + \rho\beta_i} \sum_{j=1}^{n_s}\sum_{p=1}^{n_{i,j}} y^{*}_{i,j,p} \leq \frac{\lambda_i}{\lambda_i + \rho\beta_i}\left(\lambda_i + \rho\beta_i\right) = \lambda_i,$$

and

$$\sum_{i=1}^{n_d}\sum_{p=1}^{n_{i,j}} z_{i,j,p} = \sum_{i=1}^{n_d}\sum_{p=1}^{n_{i,j}} \frac{\lambda_i}{\lambda_i + \rho\beta_i} y^{*}_{i,j,p} \leq \sum_{i=1}^{n_d}\sum_{p=1}^{n_{i,j}} y^{*}_{i,j,p} \leq \pi_j.$$

The flow constraints are satisfied by construction since scaling applies to all variables equally. The capacity constraints are also satisfied because we are only decreasing flow values.

The scaled solution $z_{i,j,p}$ has objective

$$\sum_{i=1}^{n_d}\sum_{j=1}^{n_s}\sum_{p=1}^{n_{i,j}} \nu_{i,j,p} z_{i,j,p} = \sum_{i=1}^{n_d}\sum_{j=1}^{n_s}\sum_{p=1}^{n_{i,j}} \nu_{i,j,p} y^*_{i,j,p} \frac{\lambda_i}{\lambda_i + \rho\beta_i} = \sum_{i=1}^{n_d} \lambda_i \sum_{j=1}^{n_s}\sum_{p=1}^{n_{i,j}} \frac{\nu_{i,j,p} y^*_{i,j,p}}{\lambda_i + \rho\beta_i} = \bar{\boldsymbol{v}}^* \cdot \boldsymbol{\lambda},$$

and it is feasible in $\mathcal{X}(\boldsymbol{\lambda}, \boldsymbol{\pi})$, therefore $\bar{\boldsymbol{v}}^* \cdot \boldsymbol{\lambda} \leq \Phi(\boldsymbol{\lambda}, \boldsymbol{\pi}) = \Psi(0)$. In other words, the linear approximation of the partial-treatment value function constructed implicitly by the RCT estimator underestimates the value of global control.

Step 2: From the concavity of $\Psi(\cdot)$, we can write:

$$\begin{aligned} \Psi(\rho) &\geq \rho\Psi(1) + (1-\rho)\Psi(0) \\ \bar{\boldsymbol{v}}^* \cdot (\boldsymbol{\lambda} + \rho\boldsymbol{\beta}) &\geq \rho\Psi(1) + (1-\rho)\bar{\boldsymbol{v}}^* \cdot \boldsymbol{\lambda} \\ \bar{\boldsymbol{v}}^* \cdot (\boldsymbol{\lambda} + \boldsymbol{\beta}) &\geq \Psi(1). \end{aligned}$$

In other words, the RCT estimator's implicit linear approximation also overestimates the value of global treatment.

Putting the two halves of the proof together, we obtain:

$$\hat{\Delta}_{RCT} = \bar{\boldsymbol{v}}^* \cdot \boldsymbol{\beta} = \bar{\boldsymbol{v}}^* \cdot (\boldsymbol{\lambda} + \boldsymbol{\beta}) - \bar{\boldsymbol{v}}^* \cdot \boldsymbol{\lambda} \geq \Psi(1) - \Psi(0) = \Delta.$$

□

C.3 Analysis of the two-LP estimator

Proof of Proposition 3. The two-LP estimator is defined as:

$$\begin{aligned} \hat{\Delta}^{\tau}_{\text{2LP}} = \Phi\left(\hat{\boldsymbol{\lambda}}(\boldsymbol{D}^{\tau,\text{control}}, \boldsymbol{D}^{\tau,\text{treatment}}) + \hat{\boldsymbol{\beta}}(\boldsymbol{D}^{\tau,\text{control}}, \boldsymbol{D}^{\tau,\text{treatment}}), \hat{\boldsymbol{\pi}}(\boldsymbol{S}^{\tau,\boldsymbol{\pi}})\right) - \\ \Phi\left(\hat{\boldsymbol{\lambda}}(\boldsymbol{D}^{\tau,\text{control}}, \boldsymbol{D}^{\tau,\text{treatment}}), \hat{\boldsymbol{\pi}}(\boldsymbol{S}^{\tau,\boldsymbol{\pi}})\right), \end{aligned}$$

where $\hat{\boldsymbol{\lambda}}(\cdot)$, $\hat{\boldsymbol{\beta}}(\cdot)$ and $\hat{\boldsymbol{\pi}}(\cdot)$ designate maximum-likelihood estimators for the arrival rates $\boldsymbol{\lambda}$, $\boldsymbol{\beta}$ and $\boldsymbol{\pi}$ obtained from the observed demand and supply counts. In particular,

$$\begin{aligned} D_i^{\tau,\text{control}} &\sim \text{Poisson}((1-\rho)\lambda_i) \Rightarrow \hat{\lambda}_i = \frac{D_i^{\tau,\text{control}}}{(1-\rho)\tau} \\ D_i^{\tau,\text{treatment}} &\sim \text{Poisson}(\rho(\lambda_i + \beta_i)) \Rightarrow \hat{\lambda}_i + \hat{\beta}_i = \frac{D_i^{\tau,\text{treatment}}}{\rho\tau} \\ S_j^{\tau,\pi_j} &\sim \text{Poisson}(\pi_j) \Rightarrow \hat{\pi}_j = \frac{S_j^{\tau,\pi_j}}{\tau}. \end{aligned}$$

Using the equations above, we can write:

$$\hat{\Delta}^{\tau}_{\text{2LP}} = \Phi\left(\frac{1}{\rho\tau}\boldsymbol{D}^{\tau,\text{treatment}}, \frac{1}{\tau}\boldsymbol{S}^{\tau,\boldsymbol{\pi}}\right) - \Phi\left(\frac{1}{(1-\rho)\tau}\boldsymbol{D}^{\tau,\text{control}}, \frac{1}{\tau}\boldsymbol{S}^{\tau,\boldsymbol{\pi}}\right),$$

then apply Lemma 1 to complete the proof. □

Proof of Theorem 3. The result follows directly from Proposition 3 and Theorem 1. □

Proof of Proposition 4. Let $n_d = n_s = n \gg 1$. We assume the total demand arrival rate is 1, evenly divided among the demand types, i.e. $\lambda_i = \frac{1}{n}$ for all i. We also assume the treatment effect on demand is $\beta_i = \lambda_i = \frac{1}{n}$. The arrival rate of each supply type i is $\pi_i = m \gg 1$. In particular, we choose $n \geq n_0$ and $m \geq m_n$ such that we can use the following approximations:

$$\begin{aligned}
1 - e^{-m} &= 1 - o\left(\frac{1}{n^2}\right), \\
1 - e^{-\frac{1}{n}} &= \frac{1}{n} - \frac{1}{2n^2} \pm o\left(\frac{1}{n^2}\right), \\
1 - e^{-\frac{2}{n}} &= \frac{2}{n} - \frac{2}{n^2} \pm o\left(\frac{1}{n^2}\right), \\
1 - e^{-\frac{1-\rho}{n}} &= \frac{1-\rho}{n} - \frac{(1-\rho)^2}{2n^2} \pm o\left(\frac{1}{n^2}\right), \\
1 - e^{-\frac{2\rho}{n}} &= \frac{2\rho}{n} - \frac{2\rho^2}{n^2} \pm o\left(\frac{1}{n^2}\right),
\end{aligned}$$

where the last four equations come from the power series expansion of the exponential function.

We assume that the matching network is bipartite (no intermediate nodes), and all edges have capacity 1. Let $\tau = 1$ to simplify notation. The value function is such that:

$$v_{i,j} = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

This value function can be decomposed into the sum of the value obtained from each type. Under global control, the expected value obtained from type i is equal to 1 if at least one demand unit and at least one supply unit of type i arrive, and 0 otherwise. In other words,

$$\begin{aligned}
\mathbb{E}\left[\Phi\left(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}},\boldsymbol{S}^{\tau,\boldsymbol{\pi}}\right)\right] &= \sum_{i=1}^{n} \mathbb{P}\left[D_i^{\tau,\lambda_i} \geq 1\right] \mathbb{P}\left[S_i^{\tau,\pi_i} \geq 1\right] \\
&= \sum_{i=1}^{n} (1 - \mathbb{P}\left[D_i^{\tau,\lambda_i} = 0\right])(1 - \mathbb{P}\left[S_i^{\tau,\pi_i} = 0\right]) \\
&= \sum_{i=1}^{n} (1 - e^{-\frac{1}{n}})(1 - e^{-m}) \\
&= \left(1 - o\left(\frac{1}{n^2}\right)\right) \sum_{i=1}^{n} \left(\frac{1}{n} - \frac{1}{2n^2} \pm o\left(\frac{1}{n^2}\right)\right) \\
&= 1 - \frac{1}{2n} \pm o\left(\frac{1}{n}\right),
\end{aligned}$$

Similarly, for global treatment,

$$\begin{aligned}
\mathbb{E}\left[\Phi\left(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\boldsymbol{\beta}},\boldsymbol{S}^{\tau,\boldsymbol{\pi}}\right)\right] &= \sum_{i=1}^{n} \mathbb{P}\left[D_i^{\tau,\lambda_i+\beta_i} \geq 1\right] \mathbb{P}\left[S_i^{\tau,\pi_i} \geq 1\right] \\
&= \sum_{i=1}^{n} (1 - e^{-\frac{2}{n}})(1 - e^{-m}) \\
&= \left(1 - o\left(\frac{1}{n^2}\right)\right) \sum_{i=1}^{n} \left(\frac{2}{n} - \frac{2}{n^2} \pm o\left(\frac{1}{n^2}\right)\right) \\
&= 2 - \frac{2}{n} \pm o\left(\frac{1}{n}\right).
\end{aligned}$$

Therefore the true global treatment effect is given by

$$\Delta^{\tau} = 1 - \frac{3}{2n} \pm o\left(\frac{1}{n}\right).$$

Now consider an experiment with treatment fraction ρ. For each demand type, we observe $D_i^{\tau,\text{control}} \sim \text{Poisson}\left(\frac{1-\rho}{n}\right)$, and $D_i^{\tau,\text{treatment}} \sim \text{Poisson}\left(\frac{2\rho}{n}\right)$. On average, the two-LP estimate for global control is therefore given by:

$$\begin{aligned}
\mathbb{E}\left[\Phi\left(\frac{1}{1-\rho}\boldsymbol{D}^{\tau,\text{control}},\boldsymbol{S}^{\tau,\boldsymbol{\pi}}\right)\right] &= \sum_{i=1}^{n}\mathbb{P}\left[\frac{1}{1-\rho}D_i^{\tau,\text{control}}\geq 1\right]\mathbb{P}\left[S_i^{\tau,\pi_i}\geq 1\right] \\
&= \sum_{i=1}^{n}\left(1-\mathbb{P}\left[D_i^{\tau,\text{control}}=0\right]\right)\left(1-\mathbb{P}\left[S_i^{\tau,\pi_i}=0\right]\right) \\
&= \sum_{i=1}^{n}(1-e^{-\frac{1-\rho}{n}})(1-e^{-m}) \\
&= \left(1-o\left(\frac{1}{n^2}\right)\right)\sum_{i=1}^{n}\left(\frac{1-\rho}{n}-\frac{(1-\rho)^2}{2n^2}\pm o\left(\frac{1}{n^2}\right)\right) \\
&= 1-\rho-\frac{(1-\rho)^2}{2n}\pm o\left(\frac{1}{n}\right).
\end{aligned}$$

Similarly, the two-LP estimate for global treatment is given by:

$$\begin{aligned}
\mathbb{E}\left[\Phi\left(\frac{1}{\rho}\boldsymbol{D}^{\tau,\text{treatment}},\boldsymbol{S}^{\tau,\boldsymbol{\pi}}\right)\right] &= \sum_{i=1}^{n}\mathbb{P}\left[\frac{1}{\rho}D_i^{\tau,\text{treatment}}\geq 1\right]\mathbb{P}\left[S_i^{\tau,\pi_i}\geq 1\right] \\
&= \sum_{i=1}^{n}\left(1-\mathbb{P}\left[D_i^{\tau,\text{treatment}}=0\right]\right)\left(1-\mathbb{P}\left[S_i^{\tau,\pi_i}=0\right]\right) \\
&= \sum_{i=1}^{n}(1-e^{-\frac{2\rho}{n}})(1-e^{-m}) \\
&= \left(1-o\left(\frac{1}{n^2}\right)\right)\sum_{i=1}^{n}\left(\frac{2\rho}{n}-\frac{2\rho^2}{n^2}\pm o\left(\frac{1}{n^2}\right)\right) \\
&= 2\rho-\frac{2\rho^2}{n}\pm o\left(\frac{1}{n}\right).
\end{aligned}$$

Putting together the above two results yields:

$$\mathbb{E}\left[\hat{\Delta}_{2\text{LP}}\right] = 2\rho-\frac{4\rho^2}{2n}-1+\rho+\frac{(1-\rho)^2}{2n}\pm o\left(\frac{1}{n}\right) = (3\rho-1)\left(1-\frac{1+\rho}{2n}\right)\pm o\left(\frac{1}{n}\right).$$

Henceforth ignoring the $o(1/n)$ terms, and noting that

$$\Delta^\tau = 1-\frac{3}{2n}\leq 1-\frac{1+\rho}{2n}\leq 1-\frac{1}{2n}=\Delta^\tau+\frac{1}{n},$$

we can bound the expected value of the two-LP estimator as follows:

$$\rho \geq \frac{1}{3} \Rightarrow (3\rho - 1)\Delta^\tau \leq \mathbb{E}\left[\hat{\Delta}_{\text{2LP}}\right] \leq (3\rho - 1)\left(\Delta^\tau + \frac{1}{n}\right),$$
$$\rho < \frac{1}{3} \Rightarrow (3\rho - 1)\left(\Delta^\tau + \frac{1}{n}\right) \leq \mathbb{E}\left[\hat{\Delta}_{\text{2LP}}\right] \leq (3\rho - 1)\Delta^\tau.$$

□

C.4 Analysis of the SP estimator

C.4.1 Fluid limit

Proof of Proposition 5. Analogously to Proposition 2, the result follows from the strong law of large numbers, Theorem 1, Lemma 2, and the dominated convergence theorem. □

Proof of Theorem 4. We first assume $\boldsymbol{\beta} \geq 0$. Recall from Proposition 1 that $\Psi(\cdot)$ is concave, implying that $\Psi'(\eta) = \boldsymbol{a}^\eta \cdot \boldsymbol{\beta}$ is decreasing in η. We can therefore separate the SP error into a positive and negative term:

$$\hat{\Delta}_{\text{SP}} - \Delta = \int_0^1 (\mathbf{a}^\rho - \mathbf{a}^\eta) \cdot \boldsymbol{\beta} d\eta = \underbrace{\int_0^\rho (\mathbf{a}^\rho - \mathbf{a}^\eta) \cdot \boldsymbol{\beta} d\eta}_{-\delta_- \leq 0} + \underbrace{\int_\rho^1 (\mathbf{a}^\rho - \mathbf{a}^\eta) \cdot \boldsymbol{\beta} d\eta}_{\delta_+ \geq 0}$$

We call the positive and negative parts above $\delta_- \geq 0$ and $\delta_+ \geq 0$, such that $\hat{\Delta}_{\text{SP}} - \Delta = \delta_+ - \delta_-$.

Recalling step 1 of the proof of Theorem 2, we know that

$$\int_0^\rho \mathbf{a}^\eta \cdot \boldsymbol{\beta} d\eta = \Psi(\rho) - \Psi(0) \leq \rho \hat{\Delta}_{\text{RCT}},$$

therefore

$$\rho \hat{\Delta}_{\text{RCT}} \geq \int_0^\rho \mathbf{a}^\eta \cdot \boldsymbol{\beta} d\eta = \int_0^\rho \mathbf{a}^\rho \cdot \boldsymbol{\beta} + \delta_- = \rho \mathbf{a}^\rho \cdot \boldsymbol{\beta} + \delta_-,$$

which we can re-write as

$$\hat{\Delta}_{\text{RCT}} \geq \hat{\Delta}_{\text{SP}} + \frac{1}{\rho}\delta_-.$$

The above result tells us that not only is the RCT estimator always larger than the SP estimator, but in fact the difference exceeds the (scaled) possible negative bias of the SP estimator.

We can then derive the following bound

$$\begin{aligned}\left|\hat{\Delta}_{RCT} - \Delta\right| = \hat{\Delta}_{\text{RCT}} - \Delta &= \int_0^1 \left(\hat{\Delta}_{\text{RCT}} - \mathbf{a}^\eta \cdot \boldsymbol{\beta}\right) d\eta \\ &\geq 0 + \int_\rho^1 \left(\mathbf{a}^\rho \cdot \boldsymbol{\beta} + \frac{1}{\rho}\delta_- - \mathbf{a}^\eta \cdot \boldsymbol{\beta}\right) d\eta \\ &= \delta_+ + \frac{(1-\rho)}{\rho}\delta_-\end{aligned}$$

When $\rho \leq 0.5$ we obtain

$$\left|\hat{\Delta}_{\text{RCT}} - \Delta\right| \geq \delta_+ + \delta_- \geq |\delta_+ - \delta_-| = \left|\hat{\Delta}_{\text{SP}} - \Delta\right|.$$

We can repeat the proof with $\boldsymbol{\beta} \leq 0$ and obtain the same result with $\rho \geq 0.5$. A symmetric experiment design ($\rho = 0.5$) will therefore exhibit less bias with the SP estimator than the RCT estimator as long as the treatment effect is consistent across types. □

C.4.2 Variance results

Proof of Theorem 5. By the law of total variance, we can decompose the RCT estimator variance into two components, one coming from the demand arrival rates, and the other coming from the randomized treatment assignments.

$$\begin{aligned}\text{Var}\left[\sqrt{\tau}\hat{\Delta}^\tau_{\text{RCT}}\right] &= \mathbb{E}\left[\text{Var}\left[\sqrt{\tau}\hat{\Delta}^\tau_{\text{RCT}}|\boldsymbol{D}^{\tau,\text{control}}, \boldsymbol{D}^{\tau,\text{treatment}}\right]\right] \\ &\quad + \text{Var}\left[\mathbb{E}\left[\sqrt{\tau}\hat{\Delta}^\tau_{\text{RCT}}|\boldsymbol{D}^{\tau,\text{control}}, \boldsymbol{D}^{\tau,\text{treatment}}\right]\right] \\ &\geq \text{Var}\left[\sqrt{\tau}\mathbb{E}\left[\hat{\Delta}^\tau_{\text{RCT}}|\boldsymbol{D}^{\tau,\text{control}}, \boldsymbol{D}^{\tau,\text{treatment}}\right]\right].\end{aligned} \tag{30}$$

From the proof of Proposition 2, we know that we can express the expectation on the right-hand side in Eq. (30) as

$$\mathbb{E}\left[\hat{\Delta}^\tau_{\text{RCT}}|\boldsymbol{D}^{\tau,\text{control}}, \boldsymbol{D}^{\tau,\text{treatment}}\right] = \frac{1}{\tau}\bar{\mathbf{V}}^\tau \cdot \left(\frac{1}{\rho}\boldsymbol{D}^{\tau,\text{control}} - \frac{1}{1-\rho}\boldsymbol{D}^{\tau,\text{treatment}}\right),$$

where $\bar{V}^\tau_i$ denotes the average value obtained from demand type i. Therefore, we can write

$$\lim_{\tau\to\infty} \text{Var}\left[\sqrt{\tau}\hat{\Delta}^\tau_{\text{RCT}}\right] \geq \lim_{\tau\to\infty} \text{Var}\left[\bar{\mathbf{V}}^\tau \cdot \frac{1}{\sqrt{\tau}}\left(\frac{1}{\rho}\boldsymbol{D}^{\tau,\text{control}} - \frac{1}{1-\rho}\boldsymbol{D}^{\tau,\text{treatment}}\right)\right]$$

As in the proof of Proposition 2, we can show that $\bar{\boldsymbol{V}}^\tau$ converges almost surely to $\bar{\boldsymbol{v}}^*$, the vector of average values by demand type in the fluid limit. Therefore, the above is equivalent to

$$\begin{aligned}\lim_{\tau\to\infty} \operatorname{Var}\left[\sqrt{\tau}\hat{\Delta}^{\tau}_{\mathrm{RCT}}\right] &\geq \lim_{\tau\to\infty} \operatorname{Var}\left[\bar{\mathbf{v}}^{*}\cdot\frac{1}{\sqrt{\tau}}\left(\frac{1}{\rho}\boldsymbol{D}^{\tau,\mathrm{control}} - \frac{1}{1-\rho}\boldsymbol{D}^{\tau,\mathrm{treatment}}\right)\right] \\ &= \lim_{\tau\to\infty}\sum_{i=1}^{n_d} \operatorname{Var}\left[\bar{v}^{*}_i\cdot\frac{1}{\sqrt{\tau}}\left(\frac{1}{\rho}D_i^{\tau,\mathrm{control}} - \frac{1}{1-\rho}D_i^{\tau,\mathrm{treatment}}\right)\right] \\ &= \sum_{i=1}^{n_d}(\bar{v}^{*}_i)^2 \lim_{\tau\to\infty} \operatorname{Var}\left[\frac{1}{\sqrt{\tau}}\left(\frac{1}{\rho}D_i^{\tau,\mathrm{control}} - \frac{1}{1-\rho}D_i^{\tau,\mathrm{treatment}}\right)\right] \\ &= \sum_{i=1}^{n_d}(\bar{v}^{*}_i)^2\zeta_i,\end{aligned}$$

where

$$\zeta_i := \lim_{\tau\to\infty} \operatorname{Var}\left[\frac{1}{\sqrt{\tau}}\left(\frac{1}{\rho}D_i^{\tau,\mathrm{control}} - \frac{1}{1-\rho}D_i^{\tau,\mathrm{treatment}}\right)\right] \geq 0$$

is bounded as a result of asymptotic properties of Poisson distributions.

Similarly, we can decompose the asymptotic variance of the SP estimator:

$$\begin{aligned}\operatorname{Var}\left[\sqrt{\tau}\hat{\Delta}^{\tau}_{\mathrm{SP}}\right] &= \mathbb{E}\left[\operatorname{Var}\left[\sqrt{\tau}\hat{\Delta}^{\tau}_{\mathrm{SP}}|\boldsymbol{D}^{\tau,\mathrm{control}},\boldsymbol{D}^{\tau,\mathrm{treatment}}\right]\right] + \operatorname{Var}\left[\mathbb{E}\left[\sqrt{\tau}\hat{\Delta}^{\tau}_{\mathrm{SP}}|\boldsymbol{D}^{\tau,\mathrm{control}},\boldsymbol{D}^{\tau,\mathrm{treatment}}\right]\right] \\ &= \operatorname{Var}\left[\mathbb{E}\left[\sqrt{\tau}\hat{\Delta}^{\tau}_{\mathrm{SP}}|\boldsymbol{D}^{\tau,\mathrm{control}},\boldsymbol{D}^{\tau,\mathrm{treatment}}\right]\right].\end{aligned}$$

We also know that the shadow prices $\boldsymbol{A}^{\tau}$ converge almost surely to the shadow prices in the fluid limit $\boldsymbol{a}^{*}$, so we can write:

$$\begin{aligned}\lim_{\tau\to\infty} \operatorname{Var}\left[\sqrt{\tau}\hat{\Delta}^{\tau}_{\mathrm{SP}}\right] &= \lim_{\tau\to\infty} \operatorname{Var}\left[\mathbf{a}^{*}\cdot\frac{1}{\sqrt{\tau}}\left(\frac{1}{\rho}\boldsymbol{D}^{\tau,\mathrm{control}} - \frac{1}{1-\rho}\boldsymbol{D}^{\tau,\mathrm{treatment}}\right)\right] \\ &= \sum_{i=1}^{n_d}(a^{*}_i)^2 \lim_{\tau\to\infty} \operatorname{Var}\left[\frac{1}{\sqrt{\tau}}\left(\frac{1}{\rho}D_i^{\tau,\mathrm{control}} - \frac{1}{1-\rho}D_i^{\tau,\mathrm{treatment}}\right)\right] \\ &= \sum_{i=1}^{n_d}(a^{*}_i)^2\zeta_i.\end{aligned}$$

From complementary slackness, we know that the average value of demand type i must be no smaller than its marginal value, i.e., $a^{*}_i \leq \bar{v}^{*}_i$. Therefore we conclude $(a^{*}_i)^2 \leq (\bar{v}^{*}_i)^2$ for all i and therefore

$$\lim_{\tau\to\infty} \operatorname{Var}\left[\sqrt{\tau}\hat{\Delta}^{\tau}_{\mathrm{RCT}}\right] \geq \lim_{\tau\to\infty} \operatorname{Var}\left[\sqrt{\tau}\hat{\Delta}^{\tau}_{\mathrm{SP}}\right].$$

□

C.4.3 Imbalanced bipartite matching markets

Proof of Theorem 11. Consider a sequence $\{\bar{\lambda}^k/\bar{\pi}^k\}_{k=1}^{k=\infty}$. Let $\boldsymbol{\lambda}^k$ and $\boldsymbol{\pi}^k$ designate the demand and supply arrival rate vectors corresponding to scale parameters $\bar{\lambda}^k$ and $\bar{\pi}^k$. Similarly, let $\boldsymbol{\beta}^k$ designate the treatment effect on demand corresponding to demand scaling $\bar{\lambda}$.

Oversupply limit: Assume $\{\bar{\lambda}^k/\bar{\pi}^k\}_k$ converges to 0. We can find k_0 such that for all $k \geq k_0$,

$$\min_{j\in[n_s]} \pi_j^k > \sum_{i=1}^{n_d} \lambda_i^k + \beta_i^k.$$

For any $\eta \in [0, 1]$, this result implies that the supply constraint (1c) for each supply type j in the matching problem with demand arrival rate $\boldsymbol{\lambda} + \eta\boldsymbol{\beta}$ verifies

$$\sum_{i=1}^{n_d} x_{i,j} \leq \sum_{i=1}^{n_d} \lambda_i^k + \eta\beta_i^k \leq \sum_{i=1}^{n_d} \lambda_i^k + \beta_i^k < \min_{j'\in[n_s]} \pi_{j'}^k \leq \pi_j^k.$$

From complementary slackness, we obtain $b_j^\eta = 0$ for each supply type j. We then observe that the optimal demand dual variable is given by $a_i^\eta = \max_{j\in[n_s]} v_{i,j}$ for each demand type i, for all $\eta \in [0, 1]$. Furthermore, for the experiment state $(\eta = \rho)$, complementary slackness means the optimal primal variables are given by

$$x_{i,j}^* = \begin{cases} \lambda_i^k + \rho\beta_i^k, & \text{if } j = \arg\max_{j'\in[n_s]} v_{i,j'}, \\ 0, & \text{otherwise.} \end{cases}$$

By proposition 6, we can write the true treatment effect as

$$\Delta = \int_0^1 \boldsymbol{a}^\eta \cdot \boldsymbol{\beta} d\eta = \int_0^1 d\eta \sum_{i=1}^{n_d} \max_{j\in[n_s]} v_{i,j}\beta_i = \sum_{i=1}^{n_d} \max_{j\in[n_s]} v_{i,j}\beta_i.$$

By proposition 5, the SP estimator is given by

$$\hat{\Delta}_{\text{SP}} = \boldsymbol{a}^\rho \cdot \boldsymbol{\beta} = \sum_{i=1}^{n_d} \max_{j\in[n_s]} v_{i,j}\beta_i = \Delta,$$

and by proposition 2, the RCT estimator is given by

$$\hat{\Delta}_{\text{RCT}} = \bar{\boldsymbol{v}}^* \cdot \boldsymbol{\beta} = \sum_{i=1}^{n_d} \frac{\sum_{j=1}^{n_s} x_{i,j}^* v_{i,j}}{\lambda_i^k + \rho\beta_i^k} = \sum_{i=1}^{n_d} \max_{j\in[n_s]} v_{i,j}\beta_i = \Delta.$$

Undersupply limit: Now assume $\{\bar{\lambda}^k/\bar{\pi}^k\}_k \to \infty$. We can find k_0 such that for all $k \geq k_0$,

$$\min_{i\in[n_d]} \lambda_i^k > \sum_{j=1}^{n_s} \pi_j^k.$$

Using a similar complementary slackness argument as in the first part of the proof, we can show that for any $\eta \in [0,1]$, $a_i^\eta = 0$ for each demand type i, and additionally the optimal primal solution for the experiment state ($\eta = \rho$) can be written as

$$\bar{v}_i^* = \frac{\sum_{j=1}^{n_s} \delta_{i,j}\pi_j^k v_{i,j}}{\lambda_i + \rho\beta_i}, \text{ with } \delta_{i,j} = \begin{cases} 1, & \text{if } i = \arg\max_{i'\in[n_d]} v_{i',j}, \\ 0, & \text{otherwise.} \end{cases}$$

Once again applying Propositions 2, 5, and 6, we obtain

$$\Delta = \int_0^1 \boldsymbol{a}^\eta \cdot \boldsymbol{\beta} d\eta = 0 = \boldsymbol{a}^\rho \cdot \boldsymbol{\beta} = \hat{\Delta}_{\mathrm{SP}},$$

and

$$\begin{aligned}
\hat{\Delta}_{\mathrm{RCT}} &= \sum_{i=1}^{n_d} \bar{v}_i^* \beta_i = \sum_{i=1}^{n_d} \frac{\sum_{j=1}^{n_s} \delta_{i,j}\pi_j^k v_{i,j}}{\lambda_i + \rho\beta_i}\beta_i \\
&= \sum_{i=1}^{n_d} \frac{\sum_{j=1}^{n_s} \delta_{i,j}\gamma_j\bar{\pi}^k v_{i,j}}{\bar{\lambda}^k\alpha_i p_i + \rho\bar{\lambda}^k\alpha_i q_i}\bar{\lambda}^k\alpha_i q_i \\
&= \sum_{i=1}^{n_d}\sum_{j=1}^{n_s} \delta_{i,j} v_{i,j}\gamma_j\bar{\pi}^k \frac{q_i}{p_i + \rho q_i}.
\end{aligned}$$

The right-hand side above does not tend to 0 in general as $\bar{\lambda}^k/\bar{\pi}^k \to \infty$. In particular, if $\bar{\pi}^k$ converges to a positive constant (and $\bar{\lambda}^k \to \infty$), then $\hat{\Delta}_{\mathrm{RCT}} > 0$ in the limit. □

C.4.4 Finite-sample analysis

The goal of this section is to build up to the proof of Theorem 6. We first prove Theorem 7, which establishes the existence of a positive-demand and negative-demand dual or subgradient of the matching value function $\Phi_\tau(\cdot,\cdot)$.

Proof of Theorem 7. Let us first assume that $\boldsymbol{\epsilon}$ is a rational, and that $\epsilon_i = \frac{f_i}{N}$ for some $N \in Z_+$ and $f_i \in [N]$. Using Lemma 1 we rewrite

$$\Phi_\tau(\mathbf{d},\mathbf{s}) = \frac{1}{N}\Phi_{N\tau}(N\mathbf{d}, N\mathbf{s})$$

and

$$\Phi_\tau\left(\mathbf{d}+\boldsymbol{\epsilon},\mathbf{s}\right)=\frac{1}{N}\Phi_{N\tau}\left(N\mathbf{d}+\mathbf{f},N\mathbf{s}\right).$$

In order to get from an optimal solution $x^*_{u,w}$ of $\Phi_\tau\left(\mathbf{d},\mathbf{s}\right)$ to an optimal solution of $\Phi_\tau\left(\mathbf{d}+\mathbf{e}_k,\mathbf{s}\right)$, for some $k\in[n_d]$, we know that we need to find an augmenting path, denoted by $z^k_{u,w}$, such that $x^*_{u,w}+z^k_{u,w}$ is an optimal solution of $\Phi_\tau\left(\mathbf{d}+\mathbf{e}_k,\mathbf{s}\right)$.

From diminishing returns we know that if $f_i>0$ then

$$\Phi_{N\tau}\left(N\mathbf{d}+\mathbf{f},N\mathbf{s}\right)-\Phi_{N\tau}\left(N\mathbf{d}+\mathbf{f}-\mathbf{e}_k,N\mathbf{s}\right)\le\Phi_{N\tau}\left(N\mathbf{d}+\mathbf{e}_k,N\mathbf{s}\right)-\Phi_{N\tau}\left(N\mathbf{d},N\mathbf{s}\right).$$

To compute the right-hand side, we observe that to get from the optimal solution of $\Phi_{N\tau}\left(N\mathbf{d},N\mathbf{s}\right)$ to the optimal solution of $\Phi_{N\tau}\left(N\mathbf{d}+\mathbf{e}_k,N\mathbf{s}\right)$, we follow the same augmenting path $z^k_{u,w}$, thus,

$$\Phi_{N\tau}\left(N\mathbf{d}+\mathbf{f},N\mathbf{s}\right)-\Phi_{N\tau}\left(N\mathbf{d}+\mathbf{f}-\mathbf{e}_k,N\mathbf{s}\right)\le\Phi_\tau\left(\mathbf{d}+\mathbf{e}_k,\mathbf{s}\right)-\Phi_\tau\left(\mathbf{d},\mathbf{s}\right).$$

and by diminishing returns

$$\Phi_{N\tau}\left(N\mathbf{d}+\mathbf{f},N\mathbf{s}\right)-\Phi_{N\tau}\left(N\mathbf{d},N\mathbf{s}\right)\le\sum_{k=1}^{n_d}f_k\left(\Phi_\tau\left(\mathbf{d}+\mathbf{e}_k,\mathbf{s}\right)-\Phi_\tau\left(\mathbf{d},\mathbf{s}\right)\right).$$

Using Lemma 1 again, we see that

$$\Phi_\tau\left(\mathbf{d}+\boldsymbol{\epsilon},\mathbf{s}\right)-\Phi_\tau\left(\mathbf{d},\mathbf{s}\right)\le\sum_{k=1}^{n_d}\epsilon_k\left(\Phi_\tau\left(\mathbf{d}+\mathbf{e}_k,\mathbf{s}\right)-\Phi_\tau\left(\mathbf{d},\mathbf{s}\right)\right),$$

which we can also write as

$$\Phi_\tau\left(\mathbf{d}+\boldsymbol{\epsilon},\mathbf{s}\right)\le\left(1-\sum_{k=1}^{n_d}\epsilon_k\right)\Phi_\tau\left(\mathbf{d},\mathbf{s}\right)+\sum_{k=1}^{n_d}\epsilon_k\Phi_\tau\left(\mathbf{d}+\mathbf{e}_k,\mathbf{s}\right).$$

To show the equality is strict we observe that the concavity of $\Phi_\tau\left(\mathbf{d},\mathbf{s}\right)$ implies that

$$\begin{aligned}\Phi_\tau\left(\mathbf{d}+\boldsymbol{\epsilon},\mathbf{s}\right)&=\Phi_\tau\left(\mathbf{d}+\sum_{k=1}^{n_d}\epsilon_k\mathbf{e}_k,\mathbf{s}\right)\\&=\Phi_\tau\left(\left(1-\sum_{k=1}^{n_d}\epsilon_k\right)\mathbf{d}+\sum_{k=1}^{n_d}\epsilon_k\left(\mathbf{d}+\mathbf{e}_k\right),\mathbf{s}\right)\\&\ge\left(1-\sum_{k=1}^{n_d}\epsilon_k\right)\Phi_\tau\left(\mathbf{d},\mathbf{s}\right)+\sum_{k=1}^{n_d}\epsilon_k\Phi_\tau\left(\mathbf{d}+\mathbf{e}_k,\mathbf{s}\right).\end{aligned}$$

Continuity of $\Phi_\tau\left(\cdot,\cdot\right)$ implies that the result extends to real-valued $\boldsymbol{\epsilon}$. A similar proof holds for $\Phi_\tau(\mathbf{d},\mathbf{s})-\Phi_\tau\left(\mathbf{d}-\boldsymbol{\epsilon},\mathbf{s}\right)$. □

Having proven Theorem 7, we now introduce three lemmas on the form of the derivative of a function of Poisson random variables.

Lemma 3. *Suppose $D \sim$ Poisson(λ), with $0 < \lambda < \infty$, and $\boldsymbol{x} \in \mathbb{R}^m$ be a constant vector. Let $g(\boldsymbol{x}, d)$ be a continuous function satisfying $0 \leq g(\boldsymbol{x}, d) \leq M\left(|d| + \|\boldsymbol{x}\|_1\right)$. Then $\mathbb{E}\left[g(\boldsymbol{x}, D)\right]$ is differentiable with respect to λ, and in particular*

$$\frac{\partial}{\partial \lambda}\mathbb{E}\left[g(\boldsymbol{x}, D)\right] = \mathbb{E}\left[g(\boldsymbol{x}, D+1) - g(\boldsymbol{x}, D)\right].$$

We note that the expectation above is taken only with respect to D, while keeping $\boldsymbol{x}$ constant.

Proof of Lemma 3. We use a constructive argument for the proof. In particular, using the exact form of the Poisson probability mass function, we can write

$$\mathbb{E}\left[g(\boldsymbol{x}, D)\right] = \sum_{d=0}^{\infty} g(\boldsymbol{x}, d)\mathbb{P}(D = d) = \sum_{d=0}^{\infty} g(\boldsymbol{x}, d)\frac{e^{-\lambda}\lambda^d}{d!}.$$

We need to establish that the derivative of $\mathbb{E}\left[g(\boldsymbol{x}, D)\right]$ with respect to λ exists. For any finite $n \geq 0$, we define the sequence of partial sums:

$$f_n(\lambda) = \sum_{d=0}^{n} g(\boldsymbol{x}, d)\frac{e^{-\lambda}\lambda^d}{d!},$$

such that $\mathbb{E}\left[g(\boldsymbol{x}, D)\right] = \lim_{n\to\infty} f_n(\lambda)$. We then compute the derivative:

$$\begin{aligned}
f_n'(\lambda) &= \frac{\partial}{\partial \lambda}\sum_{d=0}^{n} g(\boldsymbol{x}, d)\frac{e^{-\lambda}\lambda^d}{d!} \\
&= \sum_{d=0}^{n} \frac{g(\boldsymbol{x}, d)}{d!}\frac{\partial}{\partial \lambda}\left(e^{-\lambda}\lambda^d\right) \\
&= \sum_{d=0}^{n} \frac{g(\boldsymbol{x}, d)}{d!}\left(d\lambda^{d-1}e^{-\lambda} - \lambda^d e^{-\lambda}\right) \\
&= \sum_{d=1}^{n} g(\boldsymbol{x}, d)\frac{\lambda^{d-1}e^{-\lambda}}{(d-1)!} - \sum_{d=0}^{n} g(\boldsymbol{x}, d)\frac{e^{-\lambda}\lambda^d}{d!} \\
&= \sum_{d=0}^{n} g(\boldsymbol{x}, d+1)\frac{\lambda^{d}e^{-\lambda}}{d!} - \sum_{d=0}^{n} g(\boldsymbol{x}, d)\frac{e^{-\lambda}\lambda^d}{d!} \\
&= g_n(\lambda) - f_n(\lambda),
\end{aligned}$$

where we define

$$g_n(\lambda) = \sum_{d=0}^{n} g(\boldsymbol{x}, d+1)\frac{\lambda^d e^{-\lambda}}{d!}.$$

Because $g(\boldsymbol{x}, d) \le M(\|x\|_1 + |d|)$, we can use the Weierstrass M-test to establish uniform convergence of f_n to $\mathbb{E}[g(\boldsymbol{x}, D)]$, and of g_n to $\mathbb{E}[g(\boldsymbol{x}, D+1)]$. Therefore f_n' also converges uniformly to the derivative of $\mathbb{E}[g(\boldsymbol{x}, D)]$, and we have that

$$\frac{\partial}{\partial \lambda}\mathbb{E}[g(\boldsymbol{x}, D)] = \lim_{n\to\infty}(g_n(\lambda) - f_n(\lambda)) = \mathbb{E}[g(\boldsymbol{x}, D+1) - g(\boldsymbol{x}, D)].$$

□

Lemma 4. *Suppose demand D_i for type i is Poisson-distributed with parameter λ_i and supply S_j for type j is Poisson distributed with parameter π_j. If $\lambda_i > 0$, we can write*

$$\frac{\partial}{\partial \lambda_i}\mathbb{E}[\Phi(\boldsymbol{D}, \boldsymbol{S})] = \mathbb{E}\left[A_i^+\right],$$

where A_i^+ is the positive-demand dual or subgradient associated with the optimal solution $\Phi(\boldsymbol{D}, \boldsymbol{S})$, as defined in Theorem 7.

Proof of Lemma 4. Let $i = 1$ without loss of generality. By the law of total expectation, we have

$$\mathbb{E}[\Phi(\boldsymbol{D}, \boldsymbol{S})] = \mathbb{E}_{D_1}\left[\mathbb{E}_{\boldsymbol{D}_{2:n_d}, \boldsymbol{S}}[\Phi(\boldsymbol{D}, \boldsymbol{S})|D_1]\right].$$

Setting $g(\boldsymbol{x}, d) = \mathbb{E}_{\boldsymbol{D}_{2:n_d}, \boldsymbol{S}}[\Phi(\boldsymbol{D}, \boldsymbol{S})|D_1 = d]$, we can apply Lemma 3 to obtain

$$\begin{aligned}
\frac{\partial}{\partial \lambda_1}\mathbb{E}[\Phi(\boldsymbol{D}, \boldsymbol{S})] &= \mathbb{E}_D\left[\mathbb{E}_{\boldsymbol{D}_{2:n_d}, \boldsymbol{S}}[\Phi(\boldsymbol{D}, \boldsymbol{S})|D_1 = D+1] - \mathbb{E}_{\boldsymbol{D}_{2:n_d}, \boldsymbol{S}}[\Phi(\boldsymbol{D}, \boldsymbol{S})|D_1 = D]\right] \\
&= \mathbb{E}_{D_1}\left[\mathbb{E}_{\boldsymbol{D}_{2:n_d}, \boldsymbol{S}}[\Phi(\boldsymbol{D} + \boldsymbol{e}_1, \boldsymbol{S})|D_1]\right] - \mathbb{E}_{D_1}\left[\mathbb{E}_{\boldsymbol{D}_{2:n_d}, \boldsymbol{S}}[\Phi(\boldsymbol{D}, \boldsymbol{S})|D_1]\right] \\
&= \mathbb{E}[\Phi(\boldsymbol{D} + \boldsymbol{e}_1, \boldsymbol{S}) - \Phi(\boldsymbol{D}, \boldsymbol{S})] \\
&= \mathbb{E}\left[A_1^+\right].
\end{aligned}$$

□

Lemma 5. *Let D be a Poisson-distributed random variable with parameter λ, and let the function $f(\cdot)$ verify $\mathbb{E}[f(D)] < \infty$. Then we can write*

$$\mathbb{E}[Df(D)] = \lambda\mathbb{E}[f(D+1)].$$

Proof of Lemma 5. We can write out the right-hand side explicitly using the Poisson mass function:

$$\begin{aligned}\mathbb{E}\left[Df(D)\right] &= \sum_{d=0}^{\infty} df(d)\frac{e^{-\lambda}\lambda^d}{d!} \\ &= \sum_{d=1}^{\infty} df(d)\frac{e^{-\lambda}\lambda^d}{d!} \\ &= \lambda\sum_{d=1}^{\infty} f(d)\frac{e^{-\lambda}\lambda^{d-1}}{(d-1)!}.\end{aligned}$$

We can set $d' = d - 1$ (equivalently $d = d' + 1$) and obtain

$$\mathbb{E}\left[Df(D)\right] = \lambda\sum_{d'=0}^{\infty} f(d'+1)\frac{e^{-\lambda}\lambda^{d'}}{d'!} = \lambda\mathbb{E}\left[f(D+1)\right].$$

□

We are now ready to prove Theorem 6.

Proof of Theorem 6. **Step 1:** Explicit form of the derivative (right-hand side). Using the chain rule, we can write

$$\Psi'_\tau(\eta) = \frac{1}{\tau}\frac{d}{d\eta}(\boldsymbol{\lambda} + \eta\boldsymbol{\beta}) \cdot \nabla_{\boldsymbol{\lambda}+\eta\boldsymbol{\beta}}\mathbb{E}\left[\Phi\left(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\eta\boldsymbol{\beta}}, \boldsymbol{S}^{\tau,\boldsymbol{\pi}}\right)\right],$$

where both terms on the right-hand side are vectors of length n_d. We can then apply Lemma 4 to obtain

$$\Psi'_\tau(\eta) = \frac{1}{\tau}\boldsymbol{\beta} \cdot \mathbb{E}\left[\boldsymbol{A}^+\right].$$

Step 2: Explicit form of the shadow price estimator in expectation (left-hand side). Recall that the shadow price estimator is given by:

$$\hat{\Delta}^\tau_{\text{SP}} = \frac{1}{\tau}\boldsymbol{A}^{\tau,\text{experiment},-} \cdot \left(\frac{1}{\rho}\boldsymbol{D}^{\tau,\rho(\boldsymbol{\lambda}+\boldsymbol{\beta})} - \frac{1}{1-\rho}\boldsymbol{D}^{\tau,(1-\rho)\boldsymbol{\lambda}}\right).$$

To simplify notation, we denote the conditional expectation on supply of any random variable $\boldsymbol{Z}$ by $\mathbb{E}_{|\boldsymbol{S}}\left[\boldsymbol{Z}\right] = \mathbb{E}[\boldsymbol{Z}|\boldsymbol{S}^{\tau,\boldsymbol{\pi}}]$. We further notice that after conditioning on supply, the shadow price only depends on total demand, and therefore simplify notation to $A_i^{\tau,\text{experiment},-} = A_i^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})$.

Taking the conditional expectation on supply yields

$$
\begin{aligned}
\mathbb{E}_{|\boldsymbol{S}}\left[\hat{\Delta}^{\tau}_{\mathrm{SP}}\right] &= \frac{1}{\tau}\mathbb{E}_{|\boldsymbol{S}}\left[\sum_{i=1}^{n_d} A_i^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\left(\frac{1}{\rho}D_i^{\tau,\rho(\lambda_i+\beta_i)} - \frac{1}{1-\rho}D_i^{\tau,(1-\rho)\lambda_i}\right)\right] \\
&= \frac{1}{\tau}\sum_{i=1}^{n_d}\mathbb{E}_{|\boldsymbol{S}}\left[A_i^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\left(\frac{1}{\rho}D_i^{\tau,\rho(\lambda_i+\beta_i)} - \frac{1}{1-\rho}D_i^{\tau,(1-\rho)\lambda_i}\right)\right] \\
&= \frac{1}{\tau}\sum_{i=1}^{n_d}\mathbb{E}_{|\boldsymbol{S}}\left[\mathbb{E}_{|\boldsymbol{S}}\left[A_i^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\left(\frac{1}{\rho}D_i^{\tau,\rho(\lambda_i+\beta_i)} - \frac{1}{1-\rho}D_i^{\tau,(1-\rho)\lambda_i}\right)\middle|\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}}\right]\right] \\
&= \frac{1}{\tau}\sum_{i=1}^{n_d}\mathbb{E}_{|\boldsymbol{S}}\left[\mathbb{E}_{|\boldsymbol{S}}\left[\left(\frac{1}{\rho}D_i^{\tau,\rho(\lambda_i+\beta_i)} - \frac{1}{1-\rho}D_i^{\tau,(1-\rho)\lambda_i}\right)\middle|\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}}\right]A_i^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\right], \qquad (31)
\end{aligned}
$$

where the last two steps follow from the law of total expectation. We now recall that the total demand is the sum of the control and treatment demand, i.e., $\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}} = \boldsymbol{D}^{\tau,\rho(\boldsymbol{\lambda}+\boldsymbol{\beta})} + \boldsymbol{D}^{\tau,(1-\rho)\boldsymbol{\lambda}}$. As a result, we can write

$$
\begin{aligned}
\mathbb{E}_{|\boldsymbol{S}}\left[D_i^{\tau,\rho(\lambda_i+\beta_i)}\middle|\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}}\right] &= D_i^{\tau,\lambda_i+\rho\beta_i}\frac{\rho(\lambda_i+\beta_i)}{\lambda_i+\rho\beta_i}, \text{ and} \\
\mathbb{E}_{|\boldsymbol{S}}\left[D_i^{\tau,(1-\rho)\lambda_i}\middle|\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}}\right] &= D_i^{\tau,\lambda_i+\rho\beta_i}\frac{(1-\rho)\lambda_i}{\lambda_i+\rho\beta_i}. \qquad (32)
\end{aligned}
$$

Substituting the two expressions above into (31), we obtain:

$$
\begin{aligned}
\mathbb{E}_{|\boldsymbol{S}}\left[\hat{\Delta}^{\tau}_{\mathrm{SP}}\right] &= \frac{1}{\tau}\sum_{i=1}^{n_d}\mathbb{E}_{|\boldsymbol{S}}\left[\left(\frac{\lambda_i+\beta_i}{\lambda_i+\rho\beta_i} - \frac{\lambda_i}{\lambda_i+\rho\beta_i}\right)D_i^{\tau,\lambda_i+\rho\beta_i}A_i^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\right] \\
&= \frac{1}{\tau}\sum_{i=1}^{n_d}\frac{\beta_i}{\lambda_i+\rho\beta_i}\mathbb{E}_{|\boldsymbol{S}}\left[D_i^{\tau,\lambda_i+\rho\beta_i}A_i^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\right].
\end{aligned}
$$

Fix $i = 1$. Using the law of total expectation, we can expand the expectation on the right-hand side to

$$
\begin{aligned}
\mathbb{E}_{|\boldsymbol{S}}\left[D_1^{\tau,\lambda_1+\rho\beta_1}A_1^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\right] &= \mathbb{E}_{|\boldsymbol{S}}\left[\mathbb{E}_{|\boldsymbol{S}}\left[D_1^{\tau,\lambda_1+\rho\beta_1}A_1^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\middle|D_1^{\tau,\lambda_1+\rho\beta_1}\right]\right] \\
&= \mathbb{E}_{|\boldsymbol{S}}\left[D_1^{\tau,\lambda_1+\rho\beta_1}\mathbb{E}_{|\boldsymbol{S}}\left[A_1^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\middle|D_1^{\tau,\lambda_1+\rho\beta_1}\right]\right].
\end{aligned}
$$

Let $f(d) = \mathbb{E}_{|\boldsymbol{S}}\left[A_1^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\middle| D_1^{\tau,\lambda_1+\rho\beta_1} = d\right]$. Now we apply Lemma 5 to obtain

$$\begin{aligned}
\mathbb{E}_{|\boldsymbol{S}}\left[D_1^{\tau,\lambda_1+\rho\beta_1} A_1^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\right] &= \mathbb{E}_{|\boldsymbol{S}}\left[D_1^{\tau,\lambda_1+\rho\beta_1} f\left(D_1^{\tau,\lambda_1+\rho\beta_1}\right)\right] \\
&= (\lambda_1+\rho\beta_1)\tau\, \mathbb{E}_{|\boldsymbol{S}}\left[f(D_1^{\tau,\lambda_1+\rho\beta_1}+1)\right] \\
&= (\lambda_1+\rho\beta_1)\tau\, \mathbb{E}_{|\boldsymbol{S}}\left[A_1^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}}+\boldsymbol{e}_1)\right] \\
&= (\lambda_1+\rho\beta_1)\tau\, \mathbb{E}_{|\boldsymbol{S}}\left[\Phi_\tau(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}}+\boldsymbol{e}_1, \boldsymbol{S}^{\tau,\pi}) - \Phi_\tau(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}}, \boldsymbol{S}^{\tau,\pi})\right] \\
&= (\lambda_1+\rho\beta_1)\tau\, \mathbb{E}_{|\boldsymbol{S}}\left[A_1^+(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\right],
\end{aligned}$$

where we use the definitions of the positive-demand and negative-demand subgradients in the last step. We can apply the derivation above for each $i \geq 1$, and substitute into (32) to obtain

$$\begin{aligned}
\mathbb{E}_{|\boldsymbol{S}}\left[\hat{\Delta}_{\mathrm{SP}}^\tau\right] &= \frac{1}{\tau}\sum_{i=1}^{n_d} \frac{\beta_i}{\lambda_i+\rho\beta_i} \mathbb{E}_{|\boldsymbol{S}}\left[D_i^{\tau,\lambda_i+\rho\beta_i} A_i^-(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\right] \\
&= \frac{1}{\tau}\sum_{i=1}^{n_d} \frac{\beta_i}{\lambda_i+\rho\beta_i} (\lambda_i+\rho\beta_i)\tau\, \mathbb{E}_{|\boldsymbol{S}}\left[A_i^+(\boldsymbol{D}^{\tau,\boldsymbol{\lambda}+\rho\boldsymbol{\beta}})\right].
\end{aligned}$$

Taking the expectation over both demand and supply, we obtain that

$$\mathbb{E}\left[\hat{\Delta}_{\mathrm{SP}}^\tau\right] = \frac{1}{\tau}\boldsymbol{\beta}\cdot\mathbb{E}\left[\boldsymbol{A}^+\right],$$

which, together with the result from step 1, completes the proof. □

C.4.5 Asymmetric experiments

Proof of Theorem 8. We recall the definition of the shadow price estimator in the fluid limit

$$\hat{\Delta}_{\mathrm{SP}} = \boldsymbol{a}^\rho \cdot \boldsymbol{\beta}.$$

Taking the expectation over ρ uniformly random over $[0,1]$ yields

$$\begin{aligned}
\mathbb{E}_\rho\left[\hat{\Delta}_{\mathrm{SP}}\right] &= \int_0^1 \boldsymbol{a}^\rho \cdot \boldsymbol{\beta} d\rho \\
&= \Psi(1) - \Psi(0) \qquad \text{(by Proposition 6)} = \Delta.
\end{aligned}$$

Similarly, in the finite-sample case, we can write:

$$\mathbb{E}_{\boldsymbol{D},\boldsymbol{S}}\left[\mathbb{E}_\rho\left[\hat{\Delta}_{\mathrm{SP}}^\tau\right]\right] = \int_0^1 \Psi_\tau'(\rho)d\rho = \Psi_\tau(1) - \Psi_\tau(0) = \Delta^\tau.$$

Note that we can change the order of integration by Fubini's theorem. $\square$

Proof of Theorem 9. We fix $\boldsymbol{\beta} \leq \mathbf{0}$ and $\rho < 0.5$. We first recall that for any demand type i, $a_i \leq \bar{v}_i^*$, therefore

$$\hat{\Delta}_{\text{SP}} = \sum_{i=1}^{n_d} a_i \beta_i \geq \sum_{i=1}^{n_d} \bar{v}_i^* \beta_i = \hat{\Delta}_{\text{RCT}}. \tag{33}$$

Using the definition of the SP+ estimator, we can therefore write

$$\hat{\Delta}_{\text{SP+}} = (1-2\rho)\hat{\Delta}_{\text{RCT}} + 2\rho\hat{\Delta}_{\text{SP}} \geq (1-2\rho)\hat{\Delta}_{\text{RCT}} + 2\rho\hat{\Delta}_{\text{RCT}} = \hat{\Delta}_{\text{RCT}}.$$

This implies that

$$\Delta - \hat{\Delta}_{\text{RCT}} \geq \Delta - \hat{\Delta}_{\text{SP+}}.$$

From Theorem 2, we know that $\Delta - \hat{\Delta}_{\text{RCT}} \geq 0$. To complete the proof, we simply need to show that $\Delta - \hat{\Delta}_{\text{RCT}} \geq \hat{\Delta}_{\text{SP+}} - \Delta$. We can write their difference as

$$\begin{aligned}
\Delta - \hat{\Delta}_{\text{RCT}} - \hat{\Delta}_{\text{SP+}} + \Delta &= 2\Delta - \hat{\Delta}_{\text{RCT}} - (1-2\rho)\hat{\Delta}_{\text{RCT}} - 2\rho\hat{\Delta}_{\text{SP}} \\
&= 2\Delta - 2(1-\rho)\hat{\Delta}_{\text{RCT}} - 2\rho\hat{\Delta}_{\text{SP}} \\
&= 2\int_0^1 \boldsymbol{a}^\eta \cdot \boldsymbol{\beta} d\eta - 2\int_\rho^1 \bar{\boldsymbol{v}}^* \cdot \boldsymbol{\beta} d\eta - 2\int_0^\rho \boldsymbol{a}^\rho \cdot \boldsymbol{\beta} d\eta \\
&= 2\underbrace{\int_0^\rho (\boldsymbol{a}^\eta - \boldsymbol{a}^\rho) \cdot \boldsymbol{\beta} d\eta}_{T_1} + 2\underbrace{\int_\rho^1 (\boldsymbol{a}^\eta - \bar{\boldsymbol{v}}^*) \cdot \boldsymbol{\beta} d\eta}_{T_2}.
\end{aligned}$$

We can separetely prove each term is non-negative.

- First term: by concavity of the partial-treatment value function, we know that $\boldsymbol{a}^\eta \cdot \boldsymbol{\beta} \geq \boldsymbol{a}^\rho \cdot \boldsymbol{\beta}$ for all $\eta \leq \rho$. Therefore $T_1 \geq 0$.

- Second term: the proof of Theorem 2 established the following statements for $\boldsymbol{\beta} \leq \mathbf{0}$:

$$\begin{aligned}
\bar{\boldsymbol{v}}^* \cdot (\boldsymbol{\lambda} + \boldsymbol{\beta}) &\leq \Psi(1) \\
\bar{\boldsymbol{v}}^* \cdot (\boldsymbol{\lambda} + \rho\boldsymbol{\beta}) &= \Psi(\rho).
\end{aligned}$$

Subtracting the two above equations yields

$$(1-\rho)\bar{\boldsymbol{v}}^{\boldsymbol{*}} \cdot \boldsymbol{\beta} \leq \Psi(1) - \Psi(\rho) \Leftrightarrow \int_\rho^1 \bar{\boldsymbol{v}}^{\boldsymbol{*}} \cdot \boldsymbol{\beta} d\eta \leq \int_\rho^1 \boldsymbol{a}^\eta \cdot \boldsymbol{\beta} d\eta$$
$$\Leftrightarrow \int_\rho^1 (\boldsymbol{a}^\eta - \bar{\boldsymbol{v}}^{\boldsymbol{*}}) \cdot \boldsymbol{\beta} d\eta \geq 0$$
$$\Leftrightarrow T_2 \geq 0.$$

□

Proof of Theorem 10. Assume without loss of generality that $\boldsymbol{\beta} \geq 0$. By Proposition 1, we recall that for any $0 \leq \eta \leq 1$,

$$\boldsymbol{a}^0 \cdot \boldsymbol{\beta} \geq \boldsymbol{a}^\rho \cdot \boldsymbol{\beta} \geq \boldsymbol{a}^1 \cdot \boldsymbol{\beta}.$$

By definition,

$$\hat{\Delta}_{\text{SP}} - \Delta = \int_0^1 (\boldsymbol{a}^\rho - \boldsymbol{a}^\eta) \cdot \boldsymbol{\beta} d\eta,$$

which we can bound as:

$$(\boldsymbol{a}^\rho - \boldsymbol{a}^0) \cdot \boldsymbol{\beta} \leq \hat{\Delta}_{\text{SP}} - \Delta \leq (\boldsymbol{a}^\rho - \boldsymbol{a}^1) \cdot \boldsymbol{\beta},$$

and in absolute value

$$\text{Bias}_{\text{SP}} = \left|\hat{\Delta}_{\text{SP}} - \Delta\right| \leq (\boldsymbol{a}^0 - \boldsymbol{a}^1) \cdot \boldsymbol{\beta}.$$

This proves the second result.

Similarly, recalling Theorem 2, we can write the RCT bias as

$$\text{Bias}_{\text{RCT}} = \hat{\Delta}_{\text{RCT}} - \Delta = \int_0^1 (\bar{\boldsymbol{v}}^* - \boldsymbol{a}^\eta) d\eta \geq (\bar{\boldsymbol{v}}^* - \boldsymbol{a}^0) \cdot \boldsymbol{\beta}.$$

The last two inequalities complete the proof of the first result. □

D Secondary metrics: proofs

Proof of Proposition 8. **Step 1:** We start with $\boldsymbol{x}^{\boldsymbol{*}}$, the unique and nondegenerate optimal primal solution to $\Phi_\tau(\boldsymbol{d}, \boldsymbol{s})$, and write down the complementary slackness conditions on the optimal dual variables:

$$
\begin{aligned}
b_j^* + a_i^* + \xi^*_{u_j^s,u_i^d} &= v_{u_j^s,u_i^d} && \text{when } x^*_{u_j^s,u_i^d} > 0,\ (u_j^s, u_i^d) \in \mathcal{A}_{s\to d},\\
b_j^* - m_w^* + \xi^*_{u_j^s,w} &= v_{u_j^s,w} && \text{when } x^*_{u_j^s,w} > 0,\ (u_j^s, w) \in \mathcal{A}_{s\to *},\\
m_w^* + a_i^* + \xi^*_{w,u_i^d} &= v_{w,u_i^d} && \text{when } x^*_{w,u_i^d} > 0,\ (w, u_i^d) \in \mathcal{A}_{*\to d},\\
m_u^* - m_w^* + \xi^*_{u,w} &= v_{u,w} && \text{when } x^*_{u,w} > 0, (u,w) \in \mathcal{A}_*,\\
a_i^* &= 0 && \text{when } \sum_{w:(w,u)\in\mathcal{A}} x^*_{w,u} < d_i,\ u = u_i^d \in \mathcal{U}_d,\\
b_j^* &= 0 && \text{when } \sum_{w:(u,w)\in\mathcal{A}} x^*_{u,w} < s_j,\ u = u_j^s \in \mathcal{U}_s,\\
\xi^*_{u,w} &= 0 && \text{when } x^*_{u,w} < \tau k_{u,w},\ (u,w) \in \mathcal{A}.
\end{aligned}
$$

Because the primal solution is nondegenerate, the above system of linear equations consists of $|\mathcal{U}| + |\mathcal{A}|$ equations which uniquely determine the optimal dual solution. We can write this system in matrix form as

$$
\boldsymbol{M} \begin{pmatrix} \boldsymbol{a}^* \\ \boldsymbol{b}^* \\ \boldsymbol{m}^* \\ \boldsymbol{\xi}^* \end{pmatrix} = \boldsymbol{v}'.
$$

Each row of the matrix $\boldsymbol{M}$ is a sparse vector, and all nonzero entries are equal to 1 or -1. The first $|\mathcal{A}|$ rows have exactly three nonzeros, while the last $|\mathcal{U}|$ rows are $(|\mathcal{U}| + |\mathcal{A}|)$-dimensional unit vectors. Meanwhile, the vector $\boldsymbol{v'}$ consists of two parts: first, the sequence of $v_{u,w}$ corresponding to $x^*_{u,w} > 0$, then a sequence of zeros.

Let $\varepsilon > 0$. Consider an alternative matching problem $\Phi^\varepsilon_\tau(\boldsymbol{d}, \boldsymbol{s})$ in which the edge weights are given by $v_{u,w} + \varepsilon \tilde{v}_{u,w}$. This is a small perturbation of the original weights of the matching problem: if we take ε small enough, the optimal solution $\boldsymbol{x^*}$ in $\Phi_\tau(\boldsymbol{d}, \boldsymbol{s})$ remains optimal in $\Phi^\varepsilon_\tau(\boldsymbol{d}, \boldsymbol{s})$. As a result, we can compute the dual variables for $\Phi^\varepsilon(\boldsymbol{d}, \boldsymbol{s})$ directly from $\boldsymbol{x^*}$ using the same complementary slackness system:

$$
\boldsymbol{M} \begin{pmatrix} \boldsymbol{a}^\varepsilon \\ \boldsymbol{b}^\varepsilon \\ \boldsymbol{m}^\varepsilon \\ \boldsymbol{\xi}^\varepsilon \end{pmatrix} = \boldsymbol{v'} + \varepsilon \boldsymbol{\tilde{v}'},
$$

where the matrix $\boldsymbol{M}$ is constructed exactly as before, while the vector $\boldsymbol{\tilde{v}'}$ is constructed from the weights $\tilde{v}_{u,w}$ exactly as the vector $\boldsymbol{v'}$ was constructed from the weights $v_{u,w}$.

Step 2: By definition of $\Phi_\tau^{\tilde{\boldsymbol{v}}}(\cdot,\cdot)$, we can write

$$\Phi_\tau^\varepsilon(\boldsymbol{d},\boldsymbol{s}) = \sum_{(u,w)\in\mathcal{A}} (v_{u,w} + \varepsilon \tilde{v}_{u,w}) x^*_{u,w} = \Phi_\tau(\boldsymbol{d},\boldsymbol{s}) + \varepsilon \Phi_\tau^{\tilde{\boldsymbol{v}}}(\boldsymbol{d},\boldsymbol{s}).$$

We obtain a new expression for the marginal values $a_i^{\tilde{\boldsymbol{v}}}$:

$$\begin{aligned}
\varepsilon a_i^{\tilde{\boldsymbol{v}}} &= \varepsilon \Phi^{\tilde{\boldsymbol{v}}}(\boldsymbol{d}+\boldsymbol{e}_i,\boldsymbol{s}) - \varepsilon \Phi^{\tilde{\boldsymbol{v}}}(\boldsymbol{d},\boldsymbol{s}) \\
&= (\Phi^\varepsilon(\boldsymbol{d}+\boldsymbol{e}_i,\boldsymbol{s}) - \Phi(\boldsymbol{d}+\boldsymbol{e}_i,\boldsymbol{s})) - (\Phi^\varepsilon(\boldsymbol{d},\boldsymbol{s}) - \Phi(\boldsymbol{d},\boldsymbol{s})) \\
&= (\Phi^\varepsilon(\boldsymbol{d}+\boldsymbol{e}_i,\boldsymbol{s}) - \Phi^\varepsilon(\boldsymbol{d},\boldsymbol{s})) - (\Phi(\boldsymbol{d}+\boldsymbol{e}_i,\boldsymbol{s}) - \Phi(\boldsymbol{d},\boldsymbol{s})) \\
&= a_i^\varepsilon - a_i^*,
\end{aligned}$$

where the penultimate equality follows from Eq. (18).

Step 3: We know from step 1 that we can obtain a_i^* and a_i^ε by solving two linear systems. In other words, we can write:

$$\begin{aligned}
a_i^* &= \left(\boldsymbol{M}^{-1}\boldsymbol{v'}\right)_i, \\
a_i^\varepsilon &= \left(\boldsymbol{M}^{-1}(\boldsymbol{v'} + \varepsilon\tilde{\boldsymbol{v}}')\right)_i.
\end{aligned}$$

Applying these identities to the result from step 2, we obtain

$$\begin{aligned}
a_i^{\tilde{\boldsymbol{v}}} &= \frac{1}{\varepsilon}\left[\left(\boldsymbol{M}^{-1}(\boldsymbol{v'} + \varepsilon\tilde{\boldsymbol{v}}')\right)_i - \left(\boldsymbol{M}^{-1}\boldsymbol{v'}\right)_i\right] \\
&= \frac{1}{\varepsilon}\left(\boldsymbol{M}^{-1}\boldsymbol{v'} + \varepsilon\boldsymbol{M}^{-1}\tilde{\boldsymbol{v}}' - \boldsymbol{M}^{-1}\boldsymbol{v'}\right)_i \\
&= \left(\boldsymbol{M}^{-1}\tilde{\boldsymbol{v}}'\right)_i.
\end{aligned}$$

□

References

Ravindra K Ahuja, Thomas L Magnanti, and James B Orlin. *Network flows: theory, algorithms, and applications*, chapter 9, page 299. John Wiley & Sons, Ltd, 1993.

Patrick Bajari, Brian Burdick, Guido W Imbens, Lorenzo Masoero, James McQueen, Thomas Richardson, and Ido M Rosen. Multiple Randomization Designs. *Working Paper, Stanford*, 2021.

Dimitris Bertsimas and John N Tsitsiklis. *Introduction to Linear Optimization*, volume 6. Athena Scientific Belmont, MA, 1997.

Thomas Blake and Dominic Coey. Why Marketplace Experimentation is Harder Than It Seems: The Role of Test-Control Interference. In *Proceedings of the 15th ACM Conference on Economics and Computation*, pages 567–582, 2014.

Iavor Bojinov, David Simchi-Levi, and Jinglong Zhao. Design and Analysis of Switchback Experiments. *Management Science (forthcoming)*, 2022.

Nicholas Chamandy. Experimentation in a Ridesharing Marketplace. *Lyft Engineering Blog, https://eng.lyft.com/experimentation-in-a-ridesharing-marketplace-b39db027a66e*, 2016.

Alex Chin. Regression Adjustments for Estimating the Global Treatment Effect in Experiments with Interference. *Journal of Causal Inference*, 7(2), 2019.

WG Cochran. Long-Term Agricultural Experiments. *Supplement to the Journal of the Royal Statistical Society*, 6(2):104–148, 1939.

Andrey Fradkin. A Simulation Approach to Designing Digital Matching Platforms. *Working Paper, Boston University*, 2019.

David Holtz, Ruben Lobel, Inessa Liskovich, and Sinan Aral. Reducing Interference Bias in Online Marketplace Pricing Experiments. *Working Paper, UC Berkeley*, 2020.

Svante Janson, Tomasz Luczak, and Andrzej Rucinski. *Random Graphs*, chapter 2, pages 25–51. John Wiley & Sons, Ltd, 2000.

Ramesh Johari, Hannah Li, Inessa Liskovich, and Gabriel Y Weintraub. Experimental Design in Two-Sided Platforms: An Analysis of Bias. *Management Science (forthcoming)*, 2022.

Hannah Li, Geng Zhao, Ramesh Johari, and Gabriel Y Weintraub. Interference, bias, and variance in two-sided marketplace experimentation: Guidance for platforms. In *Proceedings of the ACM Web Conference 2022*, pages 182–192, 2022.

Evan Munro, Stefan Wager, and Kuang Xu. Treatment effects in market equilibrium. *arXiv preprint arXiv:2109.11647*, 2021.

New York City Taxi and Limousine Commission. Trip record data. https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page, 2022.

Carla Sneider and Yixin Tang. Experiment Rigor for Switchback Experiment Analysis. *Doordash Engineering Blog, https://doordash.engineering/2019/02/20/experiment-rigor-for-switchback-experiment-analysis/*, 2019.

Johan Ugander, Brian Karrer, Lars Backstrom, and Jon Kleinberg. Graph Cluster Randomization: Network Exposure to Multiple Universes. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 329–337, 2013.

Stefan Wager and Kuang Xu. Experimenting in Equilibrium. *Management Science*, 67(11):6694–6715, 2021.