

Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Experimental Design in Two-Sided Platforms: An Analysis of Bias

Ramesh Johari, Hannah Li, Inessa Liskovich, Gabriel Y. Weintraub

To cite this article:

Ramesh Johari, Hannah Li, Inessa Liskovich, Gabriel Y. Weintraub (2022) Experimental Design in Two-Sided Platforms: An Analysis of Bias. *Management Science* 68(10):7069-7089. <https://doi.org/10.1287/mnsc.2021.4247>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Experimental Design in Two-Sided Platforms: An Analysis of Bias

Ramesh Johari,^a Hannah Li,^a Inessa Liskovich,^b Gabriel Y. Weintraub^c

^a Management Science and Engineering, Stanford University, Stanford, California 94305; ^b Airbnb Inc., San Francisco, California 94117;

^c Graduate School of Business, Stanford University, Stanford, California 94305

Contact: rjohari@stanford.edu (RJ); hannahli@stanford.edu,  <https://orcid.org/0000-0002-0801-4015> (HL); inessa.liskovich@airbnb.com (IL); gweintra@stanford.edu,  <https://orcid.org/0000-0001-6111-507X> (GYW)

Received: August 14, 2020

Revised: July 2, 2021

Accepted: August 22, 2021

Published Online in Articles in Advance:
January 25, 2022

<https://doi.org/10.1287/mnsc.2021.4247>

Copyright: © 2022 INFORMS

Abstract. We develop an analytical framework to study experimental design in two-sided marketplaces. Many of these experiments exhibit interference, where an intervention applied to one market participant influences the behavior of another participant. This interference leads to biased estimates of the treatment effect of the intervention. We develop a stochastic market model and associated mean field limit to capture dynamics in such experiments and use our model to investigate how the performance of different designs and estimators is affected by marketplace interference effects. Platforms typically use two common experimental designs: demand-side “customer” randomization (CR) and supply-side “listing” randomization (LR), along with their associated estimators. We show that good experimental design depends on market balance; in highly demand-constrained markets, CR is unbiased, whereas LR is biased; conversely, in highly supply-constrained markets, LR is unbiased, whereas CR is biased. We also introduce and study a novel experimental design based on two-sided randomization (TSR) where both customers and listings are randomized to treatment and control. We show that appropriate choices of TSR designs can be unbiased in both extremes of market balance while yielding relatively low bias in intermediate regimes of market balance.

History: Accepted by David Simchi-Levi, revenue management and market analytics.

Funding: This work was supported by National Science Foundation under [Grants 1839229 and 1931696] and Stanford University (Dantzig-Lieberman Operations Research Fellowship).

Supplemental Material: The e-companion and data are available at <https://doi.org/10.1287/mnsc.2021.4247>.

Keywords: statistics: design of experiments • probability: stochastic model applications • two-sided markets

1. Introduction

We develop a framework to study experiments (also known as A/B tests) that two-sided platform operators routinely employ to improve the platform. Platforms use experiments to test many types of interventions that affect interactions between participants in the market; examples include features that change the process by which buyers search for sellers or interventions that alter the information the platform shares with buyers. The goal of the experiment is to introduce the intervention to some fraction of the market and use the resulting outcomes to estimate the effect if the intervention was introduced to the entire market. Platforms rely on these estimated effect sizes to make decisions about whether to launch the intervention to the entire market.

However, in marketplace experiments, these estimates are often biased because of interference between market participants. Market participants interact and compete with each other, and as a result, the treatment

assigned to one individual may influence the behavior of another individual. These interactions violate the Stable Unit Treatment Value Assumption (SUTVA) (Imbens and Rubin 2015) that guarantees unbiased estimates of the treatment effect. Previous work has shown that the resulting bias can be quite large and at times as large as the treatment effect itself (Blake and Coey 2014, Fradkin 2015, Holtz et al. 2020). In this work, we model the platform competition dynamics, investigate how they influence the performance of different canonical experimental designs, and also introduce novel designs that can yield improved performance.

We are particularly motivated by marketplaces where customers do not purchase goods but rather book them for some amount of time. This covers a broad array of platforms, including freelancing (e.g., Upwork), lodging (e.g., Airbnb and Booking.com), and many services (tutoring, dogwalking, child care, etc.). Although we explicitly model such a platform, the model we describe also captures features of a

platform where goods are bought and supply must be replenished for future demand.

Our model consists of a fixed number of listings; customers arrive sequentially over (continuous) time. For example, in an online labor platform, a freelancer offering work is a listing, and a client looking to hire a freelancer is a customer. On a lodging site, listings include hotel rooms or private rooms, and customers are travelers wanting to book. Naturally, an arriving customer can only book available listings (i.e., those not currently booked). The customer forms a consideration set from the set of available listings and then, according to a choice model, chooses which listing to book from this set (including an outside option). Once a listing is booked, it is occupied and becomes unavailable until the end of its occupancy time.

We focus on interventions by the platform that change the parameters governing the choice probability of the customer, such as those described above; we refer to the new choice parameters as the treatment model and the baseline as the control model.¹ We assume that the platform wants to use an experiment to assess the difference between the rate at which bookings would occur if all choices were made according to the treatment parameters and the corresponding rate if all choices were made according to the control parameters. This difference is the global treatment effect, or GTE. In particular, we assume that the quantity of interest is the steady-state (or long-run) GTE, that is, the long-run average difference in rental rates.²

Most platforms employ one of two simple designs for testing such an intervention: either customer-side randomization (what we call the CR design) or listing-side randomization (what we call the LR design). In the CR design, customers are randomized to treatment or control. All customers in treatment make choices according to the treatment choice model, and all customers in control make choices according to the control choice model. In the LR design, listings are randomized to treatment or control, and the utility of a listing is then determined by its treatment condition. As a result, in the LR design, in general each arriving customer will consider some listings in the treatment condition and some listings in the control condition. As an example, suppose the platform decides to test an intervention that shows badges for certain listings. In the CR design, all treatment customers see the badges and no control customers see the badges. In the LR design, all customers see the badges on treated listings and do not see them on control listings.

Each of these designs are associated with natural estimators. In the CR design, the platform measures the difference in the rate of bookings between treatment and control customers; this is what we call the naive CR estimator. In the LR design, the platform measures the difference in the rate at which treatment and

control listings are booked; this is what we call the naive LR estimator.

To develop some intuition for the potential biases, first consider an idealized static setting where listings are instantly replenished upon being booked; in other words, every arriving customer sees the full set of original listings as available. As a result, in the CR design there is no interference between treatment and control customers, and consequently the CR estimator is unbiased for the true GTE. On the other hand, in the LR design, every arriving customer considers both treatment and control listings when choosing whether to book, creating a linkage across listings through customer choice. In other words, in the LR design there is interference between treatment and control, and in general the LR estimator will be biased for the true GTE.

Now return to a dynamic model where the inventory of listings is limited, and listings remain unavailable for some time after being booked. In this case, observe that, on top of the preceding discussion, there is a dynamic linkage between customers; the set of listings available for consideration by a customer is dependent on the listings considered and booked by previously arriving customers. This dynamic effect introduces a new form of bias into estimation and is distinctly unique to our work. In particular, because of this dynamic bias, in general the naive CR estimator will be biased as well.

Our paper develops a dynamic model of two-sided markets with inventory dynamics and uses this framework to compare and contrast both the designs and estimators above. We also introduce and study a novel class of more general designs based on two-sided randomization (of which the two examples above are special cases). In more detail, our contributions and the organization of the paper are as follows.

1.1. Benchmark Model and Formal Mean Field Limit

Our first main contribution is to develop a general, flexible, theoretical model to capture the dynamics described above. In Section 3, we present a model that yields a continuous-time Markov chain in which the state at any given time is the number of currently available listings of each type. In Section 4, we propose a formal mean field analog of this continuous-time Markov chain, by considering a limit where both the number of listings in the system and the arrival rate of customers grow to infinity. Scaling by the number of listings yields a continuum mass of listings in the limit. In the mean field model, the state at a given time is the mass of available listings, and this mass evolves via a system of ODEs. Using a Lyapunov argument, we show that this system is globally asymptotically stable and give a succinct characterization of the resulting

asymptotic steady state of the system as the solution to an optimization problem. We formally establish that the mean field limit arises as the fluid limit of the corresponding finite market model, as the market size grows; in other words, the mean field model is a good approximation to large markets. The mean field model allows us to tractably analyze different estimators as well as to study their biases in the large market regime.

1.2. Designs and Estimators: Two-Sided, Customer-Side, and Listing-Side Randomization

In Section 5, we introduce a more general form of experimental design called two-sided randomization (TSR); an analogous idea was independently proposed recently by Bajari et al. (2021) (see also Section 2). In a TSR design, both customers and listings are randomized to treatment and control. However, the intervention is applied only when a treatment customer considers a treatment listing; otherwise, if the customer is in control or the listing is in control, the intervention is not seen by the customer. (In the example above, a customer would see the badge on a listing only if the customer was treated and the listing was treated.) Notably, the CR and LR designs are special cases of TSR. We also define natural naive estimators for each design.

1.3. Analysis of Bias: The Role of Market Balance

In Section 6, we study the bias of the different designs and estimators proposed. Our main theoretical results characterize how the bias depends on the relative volumes of supply and demand in the market. In particular, in the highly demand-constrained regime (where customers arrive slowly and/or listings replenish quickly), the naive CR estimator becomes unbiased, whereas the naive LR estimator is biased. On the other hand, in the highly supply-constrained regime (where customers arrive rapidly and/or listings replenish slowly), we find that in fact the naive LR estimator becomes unbiased, whereas the naive CR estimator is biased. These findings suggest that platforms can potentially reduce bias by choosing the type of experiment based on knowledge of market balance.

Given the findings that CR and LR experiments offer benefits in different extremes, it is natural to ask whether good performance can be achieved in moderately balanced markets by “interpolating” between the naive CR and LR estimators. We define a naive TSR estimator that achieves this interpolation and has low bias in both market extremes but still has large bias for moderate market balance. We then define more sophisticated TSR estimators that explicitly aim to correct for interference in regimes of moderate market balance. These latter estimators exhibit substantially improved performance in simulations. Online

Appendix C shows that these estimators perform well across a wide range of market parameters.

1.4. Insights from Simulations

In Section 7, we turn to simulations in the finite system to study the variance of the estimators. The simulations corroborate the theoretical findings that TSR offers benefits with respect to bias, albeit at the cost of moderate increases in variance. Among the TSR estimators that we study, we find that those with smaller reductions in bias have smaller increases in variance, whereas those with larger reductions in bias have larger increases in variance, thus revealing a tradeoff between bias and variance for the TSR estimators.

In Section 8, we compare the TSR approach with cluster-randomized experiments, an existing approach that platforms utilize to reduce bias. The simulations suggest that although both approaches can reduce bias when the market is tightly clustered, TSR estimators can reduce bias in highly interconnected markets where cluster randomized experiments cannot.

Taken together, our work sheds light on what experimental designs and associated estimators should be used by two-sided platforms, depending on market conditions, to alleviate the biases from interference that arise in such contexts. We view our work as a starting point toward a comprehensive framework for experimental design in two-sided platforms; we discuss some directions for future work in Section 9.

2. Related Work

2.1. Stable Unit Treatment Value Assumption

The types of interference described in these experiments are violations of the Stable Unit Treatment Value Assumption (SUTVA) in causal inference (Imbens and Rubin 2015). SUTVA requires that the (potential outcome) observation on one unit should be unaffected by the particular assignment of treatments to the other units. A large number of recent works have investigated experiment design in the presence of interference, particularly in the context of markets and social networks.

2.2. Interference in Marketplaces

Biases from interference can be large. Blake and Coey (2014) empirically showed in an auction experiment that the presence of interference among bidders caused the estimate of the treatment effect to be wrong by a factor of two. Fradkin (2019) found through simulations that a marketplace experiment changing search and recommendation algorithms can overestimate the true effect by 50%. More recent work by Holtz et al. (2020) randomized clusters of similar listings to treatment or control and found that the bias due to interference can be almost one-third of the treatment effect.

Interestingly, Holtz et al. (2020) also found weak empirical evidence that the extent of interference depends on market balance; our paper provides strong theoretical grounding for such a claim.

Inspired by the goal of reducing such bias, other work has developed approaches to bias characterization and reduction both theoretically (see, e.g., Basse et al. 2016 in the context of auctions with budgets) as well as via simulation (see, e.g., Holtz 2018, who explored the performance of LR designs). Our work complements this line by developing a mathematical framework for the study of estimation bias in dynamic platforms. Key to our analysis is the use of a mean field model to model both transient and steady-state behavior of experiments. A related approach is taken in Wager and Xu (2019), where a mean field analysis is used to study equilibrium effects of an experimental intervention where treatment is incrementally applied in a marketplace (e.g., through small pricing changes).

2.3. Interference in Social Networks

A bulk of the literature in experimental design with interference considers an interference that arises through some underlying social network; for example, Manski (2013) studied the identification of treatment responses under interference, Ugander et al. (2013) introduced a graph cluster-based randomization scheme and analyzed the bias and variance of the design, and many other papers, including Saveski et al. (2017), Athey et al. (2018), and Basse et al. (2019), focused on estimating the spillover effects created by interference. In particular, Zigler and Papadogeorgou (2018) and Pouget-Abadie et al. (2019) considered interference on a bipartite network, which is closer to a two-sided marketplace setting. In general, this line of work considers a fixed interference pattern (social network) over time. Our work is distinct because the interference caused by supply and demand competition is endogenous to the experiment and dynamically evolving over time.

2.4. Other Experimental Designs

In practice, platforms currently mitigate the effects of interference through either clustering techniques that change the unit of observation to reduce spillovers among them (see, e.g., Chamandy 2016), similar to some of the works mentioned above (see, e.g., Ugander et al. 2013 and Holtz 2018), or by switchback testing (Sneider et al. 2019), in which the treatment is turned on and off over time. Both cause a substantial increase in estimation variance due to a reduction in effective sample size, and thus the naive CR and LR designs remain popular workhorses in the platform experimentation toolkit. In addition to these broad classes of experiments, other work has also introduced modified experiment designs for specific types of

interventions, such as Ha-Thuc et al. (2020), for ranking experiments.

2.5. Two-Sided Randomization

Finally, a closely related paper is Bajari et al. (2021). Independent of our own work, here the authors proposed a more general multiple randomized design, of which TSR is a special case. They focused on a static model and provided an elegant and complete statistical analysis under a local interference assumption. By contrast, we focus on a dynamic platform model with market-wide interference patterns and on a mean field analysis of bias.

3. A Markov Chain Model of Platform Dynamics

In this section, we first introduce the basic dynamic platform model that we study in this paper with a finite number N of listings. In the next section, we describe a formal mean field limit of the model inspired by the regime where $N \rightarrow \infty$. This mean field limit model then serves as the framework within which we study the bias of different experimental designs and associated estimators in the remainder of the paper.

We consider a two-sided platform where we refer to the supply side as listings and the demand side as customers. Customers arrive over time, and at the time of arrival the customer forms a consideration set from the set of available listings in the market and then decides whether to book one of them. If the customer books, then the selected listing is occupied for a random length of time, during which it is unavailable to other customers. At the end of this booking, the listing again becomes available for use to other customers.

The formal details of our model are as follows.

3.1. Time

The system evolves in continuous time $t \geq 0$.

3.2. Listings

The system consists of a fixed number N of listings. We refer to “the N ’th system” as the instantiation of our model with N listings present. We use a superscript “ N ” to denote quantities in the N ’th system where appropriate.

We allow for heterogeneity in the listings. Each listing ℓ has a type $\theta_\ell \in \Theta$, where Θ is a finite set (the listing type space). Note that, in general, the type may encode both observable and unobservable covariates; in particular, our analysis does not presume that the platform is completely informed about the type of each listing. For example, in a lodging site, θ_ℓ may encode observed characteristics of a house, such as the number of bedrooms, but also characteristics that are unobserved by the platform because they may be

difficult or impossible to measure. Let $m^{(N)}(\theta)$ denote the total number of listings of type θ in the N 'th system. For each $\theta \in \Theta$, we assume that $\lim_{N \rightarrow \infty} m^{(N)}(\theta) = N \rho(\theta) > 0$. Note that $\rho(\theta) \leq 1$.

3.3. State Description

At each time t , each listing ℓ can be either available or occupied (i.e., occupied by a customer who previously booked it). The system state at time t in the N 'th system is described by $s_t^{(N)}(\sigma_t^{(N)}(\theta))$, where $\sigma_t^{(N)}(\theta)$ denotes the number of listings of type θ available in the system at time t . Let $S_t^{(N)} = \sum_{\theta} \sigma_t^{(N)}(\theta)$ be the total number of listings available at time t . In our subsequent development, we develop a model that makes $s_t^{(N)}$ a continuous-time Markov process.

3.4. Customers

Customers arrive to the platform sequentially and decide whether to book and, if so, which listing to book. Each customer j has a type $\gamma_j \in \Gamma$, where Γ is a finite set (the customer type space) that represents customer heterogeneity. As with listings, the type may encode both observable and unobservable covariates, and again, our analysis does not presume that the platform is completely informed about the type of each customer. Customers of type γ arrive according to a Poisson process of rate $\lambda^{(N)}_\gamma$; these processes are independent across types. Let $\lambda^{(N)} = \sum_\gamma \lambda^{(N)}_\gamma$ be the total arrival rate of customers. Let T_j denote the arrival time of the j 'th customer.

We assume that $\lim_{N \rightarrow \infty} \lambda^{(N)} = N \lambda > 0$; that is, the arrival rate of customers grows proportionally with the number of listings when we take the large market limit. Furthermore, we assume that, for each $\gamma \in \Gamma$, we have $\lim_{N \rightarrow \infty} \lambda^{(N)}_\gamma = \lambda^{(N)}_\gamma \phi_\gamma > 0$. Note that $\phi_\gamma \leq 1$.

3.5. Consideration Sets

In practice, when customers arrive to a platform, they typically form a consideration set of possible listings to book; the initial formation of the consideration set may depend on various aspects of the search and recommendation algorithms employed by the platform. To simplify the model, we capture this process by assuming that, on arrival, each listing of type θ that is available at time t is included in the arriving customer's consideration set independently with probability $\alpha_\gamma(\theta) > 0$ for a customer of type γ . For example, $\alpha_\gamma(\theta)$ can capture the possibility that the platform's search ranking is more likely to highlight available listings of type θ that are more attractive for a customer of type γ , making these listings more likely to be part of the customer's consideration set; this effect is made clear via our choice model presented below. After the consideration set is formed, a

choice model is then applied to the consideration set to determine whether a booking (if any) is made.

Formally, the customer choice process unfolds as follows. Suppose that customer j arrives at time T_j .

For each listing ℓ , let $C_{j\ell} = 0$ if the listing is unavailable at T_j . Otherwise, if listing ℓ is available, then let $C_{j\ell} = 1$ with probability $q_\gamma(\theta_\ell)$, and let $C_{j\ell} = 0$ with probability $1 - q_\gamma(\theta_\ell)$, independently of all other randomness.

Then the consideration set of customer j is $\{\ell : C_{j\ell} = 1\}$.

Our theoretical results in this paper are developed with this model of consideration set formation. However, other models of consideration set formation are also reasonable. As one example, customers might sample a consideration set of a fixed size, regardless of the total number of listings available. We explore such a consideration set model through simulations in Online Appendix C and show that similar insights hold.

3.6. Customer Choice

Customers choose at most one listing to book and can choose not to book at all. We assume that customers have a utility for each listing that depends on both customer and listing types; a type γ customer has utility $v_\gamma(\theta) > 0$ for a type θ listing. Let $q_{j\ell}$ denote the probability that arriving customer j of type γ_j books listing ℓ of type θ_ℓ .

In this paper, we assume that customers make choices according to the well-known multinomial logit choice model. In particular, given the realization of C_j , we have

$$q_{j\ell} = \frac{C_{j\ell} v_{\gamma_j}(\theta_\ell)}{\sum_{\ell' \in C_j} C_{j\ell'} v_{\gamma_j}(\theta_{\ell'})} : \quad (1)$$

Here, $v_\gamma > 0$ is the value of the outside option for type γ customers in the N 'th system. The probability that customer j does not book any listing at all grows with N . We let the outside option scale with N ; this is motivated by the observation that in practical settings the probability that a customer does not book should remain bounded away from zero even for very large systems. In particular, we assume that $\lim_{N \rightarrow \infty} v_\gamma = N \gamma > 0$.

We note that this specification of choice model, although it relies on the multinomial logit model, can be quite flexible because we allow for arbitrary heterogeneity of listings and customers.

For later reference, we define

$$q_j(\theta) = \mathbb{E}_{\ell: \theta_\ell = \theta} q_{j\ell} , \quad (2)$$

where the expectation is over the randomness in C_j . With this definition, $q_j(\theta)$ is the probability that customer j books an available listing of type θ , where the probability is computed prior to realization of the consideration set.

3.7. Dynamics: A Continuous-Time Markov Chain
The system evolves as follows. Initially, all listings are available.³ Every time a customer arrives, the choice process described above unfolds. An occupied listing remains occupied, independent of all other randomness, for an exponential time that is allowed to depend on the type of the listing.⁴ More formally, let $\tau > 0$, and for each type θ define $v(\theta)$ such that, once booked, a listing of this type will remain occupied for an exponential time with parameter $\tau v(\theta)$. We overload notation and define $\tau(\theta) = \tau v(\theta)$. Once this time expires, the listing returns to being available.

When fixing n ($v(\theta), \theta \in \Theta$) and all system parameters except for τ , increasing τ will make the system less supply constrained, and decreasing τ will make the system more supply constrained while preserving the relative occupancy times of each listing type.

Our preceding specification turns $s_t^{(N)}$ into a continuous-time Markov process on a finite state space $S^{(N)} = \{s : 0 \leq \sigma(\theta) \leq m^{(N)}(\theta), \forall \theta\}$. We now describe the transition rates of this Markov process. For a state $s \in S^{(N)}$, $\sigma(\theta)$ represents the number of available listings of type θ .

There are only two types of transitions possible. Either (i) a listing that is currently occupied becomes available, or (ii) a customer arrives and books a listing that is currently available. (If a customer arrives but does not book anything, the state of the system is unchanged.) Let e_θ denote the unit basis vector in the direction θ , that is, $e_\theta(\theta) = 1$, and $e_\theta(\theta) = 0$ for $\theta \neq \theta$. The rate of the first type of transition is

$$R(s, s + e_\theta) = (m^{(N)}(\theta) - \sigma(\theta))\tau(\theta), \quad (3)$$

since there are $m^{(N)}(\theta) - \sigma(\theta)$ booked listings of type θ , and each remains occupied for an exponential time with mean $1 = \tau(\theta)$ independent of all other randomness.

The second type of transition requires some more steps to formulate. In principle, our choice model suggests that both the identity of the arriving guest and individual listings affect system dynamics; however, our state description tracks only the aggregate number of listings of each type available at each time t . The key here is that our entire specification depends on customers only through their type and depends on listings only through their type.

Formally, suppose a customer j of type γ arrives to find the system in state s . For each θ let $D_\gamma(\theta | s)$ be a binomial $(\sigma(\theta), \alpha_\gamma(\theta))$ random variable independently across θ . Recall that for each available listing ℓ , each $C_{j\ell}$ is a Bernoulli($\alpha_\gamma(\theta_\ell)$) random variable. Recalling $q_j(\theta)$ as defined in (2), it is straightforward to check that

$$q_j(\theta) r_\gamma(\theta | s) \in E_{\gamma} \left[\frac{D_\gamma(\theta | s) \alpha_\gamma(\theta)}{\sum_{\gamma'} D_\gamma(\theta | s) \alpha_\gamma(\theta)} \right]. \quad (4)$$

In other words, the probability that an arriving customer of type γ books a listing of type θ when the state is s is given by $r_\gamma(\theta | s)$, and this probability depends on the past history only through the state s (ensuring that the Markov property holds).

With this definition at hand, for states s with $\sigma(\theta) > 0$, the rate of the second type of transition is

$$R(s, s - e_\theta) = \lambda^{(N)} r_\gamma(\theta | s). \quad (5)$$

Note that the resulting Markov chain is irreducible since customers have positive probability of sampling into and booking from their consideration set, and every listing in the consideration set has a positive probability of being booked.

3.8. Steady State

Because the Markov process defined above is irreducible on a finite state space, there is a unique steady-state distribution $\pi^{(N)}$ on $S^{(N)}$ for the process.

4. A Mean Field Model of Platform Dynamics

The continuous-time Markov process described in the preceding section is challenging to analyze directly because the customers' choices and consideration sets induce complex dynamics. Instead, to make progress, we consider a formal mean field limit motivated by the regime where $N \rightarrow \infty$, in which the evolution of the system becomes deterministic. We first present a formal mean field analog of the Markov process introduced in the previous section and provide intuition for its derivation. We then formally prove that the sequence of Markov processes converges to this mean field model as $N \rightarrow \infty$. The mean field model provides tractable expressions in the large market regime for the different estimators we consider, allowing us to study and compare their bias.

The mean field model we study consists of a continuum unit mass of listings. The total mass of listings of type θ in the system is $\rho(\theta) > 0$ (recall that $\sum_\theta \rho(\theta) = 1$). We represent the state at time t by $s_t = (s(\theta), \theta \in \Theta)$; $s_t(\theta)$ represents the mass of listings of type θ available at time t . The state space for this model is

$$S = \{s : 0 \leq s(\theta) \leq \rho(\theta)\}. \quad (6)$$

We first present the intuition behind our mean field model. Consider a state $s \in S$ with $s(\theta) > 0$ for all θ . We view this state as analogous to a state $s \approx Ns$ in the N 'th system. We consider the system dynamics defined by (3)–(5). Note that the rate at which occupied listings of type θ become available is $(m^{(N)}(\theta) - \sigma(\theta))\tau(\theta)$, from (3). If we divide by N , then this rate becomes $(\rho(\theta) - s(\theta))\tau(\theta)$, as $N \rightarrow \infty$. On the other hand, note that for large N , if $D_\gamma(\theta | s)$ is binomial

$(\sigma(\theta), \alpha_y(\theta))$, then $D_y(\theta | s) = N$ concentrates on $\alpha_y(\theta)s(\theta)$. Thus the choice probability $r_y(\theta | s)$ is approximately

$$p_y(\theta | s) \in \frac{\alpha_y(\theta)v_y(\theta)s(\theta)}{\gamma + \theta\alpha_y(\theta)v_y(\theta)s(\theta)}: \quad (7)$$

(Here we use the fact that $\gamma = N \rightarrow \gamma$ as $N \rightarrow \infty$.) This is the mean field multinomial logit choice model for our system. In the finite model, the rate at which listings of type θ become occupied is $\lambda_y^{(N)}r_y(\theta | s)$, from (5). If we divide by N , this rate becomes $\lambda_y\phi_y p_y(\theta | s)$, as $N \rightarrow \infty$.

Inspired by the preceding observations, we define the following system of differential equations for the evolution of s_t :

$$\frac{ds_t(\theta)}{dt} = (\rho(\theta) - s_t(\theta))\tau(\theta) - \lambda\phi_y p_y(\theta | s_t), \quad \theta \in \Theta: \quad (8)$$

This is our formal mean field model. In the remainder of this section, we first show that this system has a unique solution for any initial condition. Then, we characterize the behavior of the system. By constructing an appropriate Lyapunov function, we show that the mean field model has a unique limit point to which all trajectories converge (regardless of initial condition). This limit point is the unique steady state of the mean field limit. Finally, we prove that the sequence of Markov processes indeed converges to this mean field model (in an appropriate sense). Hence, the mean field model provides a close approximation to the evolution of large finite markets.

4.1. Existence and Uniqueness of Mean Field Trajectory

First, we show the straightforward result that the system of ODEs defined in (8) possesses a unique solution. This follows by an elementary application of the Picard-Lindelöf theorem from the theory of differential equations. The proof is in Online Appendix A.

Proposition 1. Fix an initial state $s \in S$. The system (8) has a unique solution $\{s_t : t \geq 0\}$ satisfying $0 \leq s_t(\theta) \leq \rho(\theta)$ and for all t and θ , and $s_0 = s$.

4.2. Existence and Uniqueness of Mean Field Steady State

Now, we characterize the behavior of the mean field limit. We show that the system of ODEs in (8) has a unique limit point, to which all trajectories converge regardless of the initial condition. We refer to this as the steady state of the mean field system. We prove the result via the use of a convex optimization problem; the objective function of this problem is a Lyapunov

function for the mean field dynamics that guarantees global asymptotic stability of the steady state.

Formally, we have the following result. The proof is in Online Appendix A.

Theorem 1. There exists a unique steady state $s^* \in S$ for (8), that is, a unique vector $s^* \in S$ solving the following system of equations:

$$(\rho(\theta) - s^*(\theta))\tau(\theta) \lambda_y \phi_y p_y(\theta | s^*), \quad \theta \in \Theta: \quad (9)$$

This limit point has the property that $0 < s^*(\theta) < \rho(\theta)$ for all θ ; that is, it is in the interior of S . Furthermore, this limit point is globally asymptotically stable; that is, all trajectories of (8) converge to s^* , as $t \rightarrow \infty$, for any initial condition $s_0 \in S$.

The limit point s^* is the unique solution to the following optimization problem:

$$\begin{aligned} \text{minimize } W(s) \in & \gamma \log \gamma + \theta \alpha_y(\theta)v_y(\theta)s(\theta) \\ & - \tau(\theta)\rho(\theta)\log s(\theta) + \tau(\theta)s(\theta) \end{aligned} \quad (10)$$

$$\text{subject to } 0 \leq s(\theta) \leq \rho(\theta), \quad \theta \in \Theta: \quad (11)$$

The function W appearing in the proposition statement is not convex; our proof proceeds by first noting that it suffices to restrict attention to s such that $s(\theta) > 0$ for all θ , then making the transformation $y(\theta) = \log(s(\theta))$. The objective function redefined in terms of these transformed variables is strictly convex, and this allows us to establish the desired result.

4.3. Convergence to the Mean Field Limit

Finally, we formally describe the sense in which our system converges to the system in (8). We first move from analyzing the number of listings available in the N 'th system to analyzing the proportion of listings available. To this end, define the normalized process $y_t^{(N)}$, where

$$Y_t^{(N)}(\theta) \quad \sigma_t^{(N)}(\theta) = N, \quad \theta \in \Theta:$$

Note that, under this definition, $y_t^{(N)}$ is also a continuous-time Markov process with dynamics induced by the dynamics of $s_t^{(N)}$; in particular, the chain $y_t^{(N)}$ has the same transition rates as $s_t^{(N)}$, but increments are of size $1=N$. The following theorem establishes the convergence of $y_t^{(N)}$ to the solution of the ODE described in (8) as $N \rightarrow \infty$.

Theorem 2. Assume that $\gamma = N \rightarrow \gamma$ for all γ and $\lambda^{(N)} = N \rightarrow \lambda_y$ for all γ as $N \rightarrow \infty$. Fix $s \in S$ and assume that $y_0^{(N)}$ is deterministic, with $y_0^{(N)}(\theta) \rightarrow s(\theta)$ for all θ . Let s_t denote the unique solution to the system defined in (8), with initial condition $s_0 = s$.

Then for all $\delta > 0$ and for all times $u > 0$,

$$P \sup_{0 \leq t \leq u} ||Y_t^{(N)} - s_t|| > \delta \leq \frac{1}{N} : \quad (12)$$

The proof for this result relies on an application of Kurtz's Theorem for the convergence of pure jump Markov processes; full details are in Online Appendix A. We note that this result holds for any sequence of initial conditions $Y_0^{(N)}$, as long as the proportion of available listings at time $t=0$ converges to a constant vector s , as $N \rightarrow \infty$; furthermore, the vector s can be any (feasible) initial state in the mean field model.

We now utilize the mean field model to study experimental designs and interference.

5. Experiments: Designs and Estimators

In this section, we leverage the framework developed in the previous section to undertake a study of experimental designs a platform might employ to test interventions in the marketplace. For simplicity, we focus on interventions that change the choice probability of one or more types of customers for one or more types of listings, and we assume that the platform is interested in estimating the resulting rate at which bookings take place. However, we believe that the same approach we employ here can be applied to study other types of interventions and platform objectives as well.

Formally, the platform's goal is to design experiments with associated estimators to assess the performance of the intervention (the treatment) relative to the status quo (the control). In particular, the platform is interested in determining the steady-state rate of booking when the entire market is in the treatment condition (i.e., global treatment) compared with the steady-state rate of booking, when the entire market is in the control condition (i.e., global control). We refer to the difference of these two rates as the global treatment effect (GTE). We focus on these steady-state quantities since a platform is typically interested in the long-run effect of an intervention.

Two types of canonical experimental designs are typically employed in practice: listing-side randomization (denoted LR) and customer-side randomization (denoted CR). In the former design, listings are randomized to treatment or control; in the latter design, customers are randomized to treatment or control. Each design also has an associated natural "naive" estimator of booking rates, that is, a (scaled) difference in means estimators for the two groups. As we discuss, these estimators will typically be biased due to interference effects.

The LR and CR designs are special cases of a novel, more general two-sided randomization (TSR) design that we introduce in this work, where both listings

and customers are randomized to treatment and control simultaneously. As we discuss, this type of experiment can be combined with design and analysis techniques to reduce bias. On the design side, TSR designs allow us to construct experiments that interpolate between LR and CR designs in such a way that bias is reduced. On the analysis side, TSR designs allow us to observe different competition effects that we can use to heuristically debias our estimators. (TSR designs were also independently introduced and studied in recent work by Bajari et al. (2021); see Section 2 for discussion.) In the next section, we develop the relevant formalism for these designs; we then subsequently define natural "naive" estimators that are commonly used for the LR and CR designs as well as an estimator for a TSR design. In the remainder of the paper, we study the bias of these different designs and estimators under different market conditions.

5.1. Experimental Design

Because CR and LR are special cases of a TSR design, we first describe how to embed TSR experimental designs into our model and then subsequently describe CR and LR designs in our model.

5.1.1. Treatment Condition. We consider a binary treatment; every customer and listing in the market will either be in treatment or control. (Generalization of our model to more than two treatment conditions is relatively straightforward.) We model the treatment condition by expanding the set of customer and listing types. For every customer type $\gamma \in \Gamma$ we create two new customer types $(\gamma, 0), (\gamma, 1)$, and for every listing type $\theta \in \Theta$ we create a two new listing types $(\theta, 0), (\theta, 1)$. The types $(\gamma, 0), (\theta, 0)$ are control types; the types $(\gamma, 1), (\theta, 1)$ are treatment types.

5.1.2. Two-Sided Randomization. In this design, randomization takes place on both sides of the market simultaneously. We assume that a fraction a_c of customers are randomized to treatment and a fraction $1 - a_c$ to control independently, and we assume that a fraction a_l of listings are randomized to treatment and a fraction $1 - a_l$ to control independently.

5.1.3. Treatment as a Choice Probability Shift. Examples of interventions that platforms may wish to test include the introduction of higher-quality photos for a hotel listing on a lodging site, showing previous job completion rates of a freelancer on an online labor market, or reducing the friction for an item in the checkout flow. Such interventions change the choice probability of listings by customers through either the consideration probabilities or perceived utility for a listing. In particular, we continue to assume the multinomial logit choice model, and we assume that for a

type γ customer and a type θ listing that have been given the intervention, the utility becomes $v_\gamma(\theta) > 0$ and the probability of inclusion in the consideration set becomes $\alpha_\gamma(\theta) > 0$. Because we focus on changes in choice probabilities, we assume that the holding time parameter of a listing of type θ is $v(\theta)$, regardless of whether it is assigned to treatment or control.⁵

In the TSR designs that we consider, a key feature is that the intervention is applied only when a treated customer interacts with a treated listing. For example, when an online labor marketplace decides to show previous job completion rates of a freelancer as an intervention, only treated customers can see these rates, and they see them only when they consider treated freelancers. We model this by redefining quantities in the experiment as follows:

$$\begin{aligned} v_{\gamma,0}(\theta, 0) & v_{\gamma,1}(\theta, 0) & v_{\gamma,0}(\theta, 1) & v_\gamma(\theta); \\ v_{\gamma,1}(\theta, 1) & v_\gamma(\theta); \end{aligned} \quad (13)$$

$$\begin{aligned} \alpha_{\gamma,0}(\theta, 0) & \alpha_{\gamma,1}(\theta, 0) & \alpha_{\gamma,0}(\theta, 1) & \alpha_\gamma(\theta); \\ \alpha_{\gamma,1}(\theta, 1) & \alpha_\gamma(\theta); \end{aligned} \quad (14)$$

$$\gamma, 0 \quad \gamma, 1 \quad \gamma; \quad (15)$$

$$v(\theta, 0) \quad v(\theta, 1) \quad v(\theta); \quad (16)$$

This definition is a natural way to incorporate randomization on each side of the market. However, we remark here that it is not necessarily canonical; for example, an alternate design would be one where the intervention is applied when either the customer has been treated or the listing has been treated. Even more generally, the design might randomize whether the intervention is applied, based on the treatment condition of the customer and the listing. In all likelihood, the relative advantages of these designs would depend not only on the bias they yield in any resulting estimators but also in the variance characteristics of those estimators. We leave further study and comparison of these designs to future work.

5.1.4. Customer-Side and Listing-Side Randomization. Two canonical special cases of the TSR design are as follows. When $a_L 1$, all listings are in the treatment condition; in this case, randomization takes place only on the customer side of the market. This is the customer-side randomization (CR) design. When $a_C 1$, all customers are in the treatment condition, and randomization takes place only on the listing side of the market. This is the listing-side randomization (LR) design.

5.1.5. System Dynamics. With the specification above, it is straightforward to adapt our mean field system of

ODEs, cf. (8), and the associated choice model (7), to this setting. The key changes are as follows:

1. The mass of control (resp., treatment) listings of type $(\theta, 0)$ (resp., $(\theta, 1)$) becomes $(1 - a_L)\rho(\theta)$ (resp., $a_L\rho(\theta)$). In other words, abusing notation, we define $\rho(\theta, 0)$ $(1 - a_L)\rho(\theta)$ and $\rho(\theta, 1)$ $a_L\rho(\theta)$.

2. The arrival rate of control (resp., treatment) customers of type $(\gamma, 0)$ (resp., $(\gamma, 1)$) becomes $(1 - a_C)\lambda\phi_\gamma$ (resp., $a_C\lambda\phi_\gamma$). Thus we define $\phi_{\gamma,0} (1 - a_C)\phi_\gamma$ and $\phi_{\gamma,1} a_C\phi_\gamma$.

3. The choice probabilities are defined as in (7), with the relevant quantities defined in (13)–(15).

Using Proposition 1 and Theorem 1, we know that there exists a unique solution to the resulting system of ODEs and that there exists a unique limit point to which all trajectories converge, regardless of initial condition. This limit point is the steady state for a given experimental design. For a TSR experiment with treatment customer fraction a_C and treatment listing fraction a_L , we use the notation $s_t(a_C, a_L)$ ($s_t(\theta, j) | a_C, a_L, \theta \in \Theta, j \in \{0, 1\}$) to denote the ODE trajectory, and we use $s^\square(a_C, a_L)$ ($s^\square(\theta, j) | a_C, a_L, \theta \in \Theta, j \in \{0, 1\}$) to denote the steady state.

5.1.6. Rate of Booking. In our subsequent development, it will be useful to have a shorthand notation for the rate at which bookings of listings of treatment condition $j \in \{0, 1\}$ are made by customers of treatment condition $i \in \{0, 1\}$ in the interval $[0, T]$. In particular, we define

$$Q_{ij}(T | a_C, a_L) = \frac{1}{T} \int_0^T \phi_{\gamma,i} \rho_{\gamma,j}(\theta, j | s_t(a_C, a_L)) dt; \quad (17)$$

Because $s_t(a_C, a_L)$ is globally asymptotically stable, bounded, and converges to $s^\square(a_C, a_L)$ as $t \rightarrow \infty$, we have

$$\begin{aligned} Q_{ij}(\infty | a_C, a_L) & \in \lim_{T \rightarrow \infty} Q_{ij}(T | a_C, a_L) \\ & \lambda \phi_{\gamma,i} \rho_{\gamma,j}(\theta, j | s^\square(a_C, a_L)); \end{aligned} \quad (18)$$

5.1.7. Global Treatment Effect. Recall that we assume that the steady-state rate of booking is the quantity of interest to the platform. In particular, the platform is interested in the change in this rate from the global control condition ($a_C = 0, a_L = 0$) to the global treatment condition ($a_C = 1, a_L = 1$).

In the global control condition, the steady state rate at which customers book is $Q^{GC} = Q_{00}(\infty | 0, 0)$, and in the global treatment condition, the steady state rate at which customers book is $Q^{GT} = Q_{11}(\infty | 1, 1)$. Under these definitions, the global treatment effect is $GTE = Q^{GT} - Q^{GC}$.

We remark that the rate of booking decisions made by arriving customers will change over time, even if

the market parameters are constant over time (including the arrival rates of different customer types as well as the utilities that customers have for each listing type). This transient change in booking rates is driven by changes in the state s_t ; in general, such fluctuations will lead the transient rate of booking to differ from the steady-state rate for all values of a_c and a_L (including global treatment and global control). In this work, we focus on the steady-state quantities to capture, informally, the long-run change in behavior due to an intervention.⁶

5.2. Estimators: Transient and Steady State

The goal of the platform is to use the experiment to estimate GTE. In this section, we consider estimators the platform might use to estimate this quantity. We first consider the CR and LR designs, and we define “naive” estimators that the platform might use to estimate the global treatment effect. These designs and estimators are those most commonly used in practice. We define these estimators during the transient phase of the experiment and then define the associated steady-state versions of these estimators. Finally, we combine these estimation approaches in a natural heuristic that can be employed for any general TSR design.

5.2.1. Estimators for the CR Design. We start by considering the CR design, that is, where $a_L \in [0, 1]$ and $a_C \in (0, 1)$. A simple naive estimate of the rate of booking is to measure the rate at which bookings are made in a given interval of time by control customers and compare this to the analogous rate for treatment customers. Formally, suppose the platform runs the experiment for the interval $t \in [0, T]$, with a fraction a_C of customers in treatment. The rate at which customers of treatment condition $i \in \{0, 1\}$ book in this period is $Q_{i1}(T | a_C, 1)$. The naive CR estimator is the difference between treatment and control rates, where we correct for differences in the size of the control and treatment groups by scaling with the respective masses:

$$GTE^{CR}(T | a_C) = \frac{Q_{11}(T | a_C, 1)}{a_C} - \frac{Q_{01}(T | a_C, 1)}{1 - a_C}. \quad (19)$$

We let $GTE^{CR}(\infty | a_C) = Q_{11}(\infty | a_C, 1) - Q_{01}(\infty | a_C, 1) = (1 - a_C)$ denote the steady-state naive CR estimator.

5.2.2. Estimators for the LR Design. Analogously, we can define a naive estimator for the LR design, that is, where $a_C \in [0, 1]$ and $a_L \in (0, 1)$. Formally, suppose the platform runs the experiment for the interval $t \in [0, T]$, with fraction a_L of listings in treatment. The rate at which listings with treatment condition $j \in \{0, 1\}$ are booked in this period is $Q_{1j}(T | 1, a_L)$. The naive LR estimator is the difference between treatment and control rates scaled by the mass of listings in each group:

$$GTE^{LR}(T | a_L) = \frac{Q_{11}(T | 1, a_L)}{a_L} - \frac{Q_{10}(T | 1, a_L)}{1 - a_L}. \quad (20)$$

We let $GTE^{LR}(\infty | a_L) = Q_{11}(\infty | 1, a_L) - Q_{10}(\infty | 1, a_L) = (1 - a_L)$ denote the corresponding steady-state naive LR estimator.

5.2.3. Estimators for the TSR Design. As with the LR and CR designs, it is possible to design a natural naive estimator for the TSR design as well. In particular, we have the following definition of the naive TSR estimator:

$$GTE^{TSRN}(T | a_C, a_L) = \frac{Q_{11}(T | a_C, a_L) c}{a_L a} - \frac{Q_{01}(T | a_C, a_L) + Q_{10}(T | a_C, a_L) + Q_{00}(T | a_C, a_L)}{1 - a_C a_L}. \quad (21)$$

To interpret this estimator, observe that the first term is the normalized rate at which treatment customers booked treatment listings in the experiment; we normalize this by $a_C a_L$, since a mass a_C of customers are in treatment and a mass a_L of listings are in treatment. This first term estimates the global treatment rate of booking. The sum $Q_{01}(T | a_C, a_L) + Q_{10}(T | a_C, a_L) + Q_{00}(T | a_C, a_L)$ is the total rate at which control bookings took place either because the customer was in the control group, the listing was in the control group, or both. (Recall that in the TSR design, the intervention is only seen when treatment customers interact with treatment listings.) This is normalized by the complementary mass, $1 - a_C a_L$. This second term estimates the global control rate of booking. As before, we can define a steady-state version of this estimator as

$GTE^{TSRN}(\infty | a_C, a_L)$, with the steady-state versions of the respective quantities on the right hand side of (21).

It is straightforward to check that as $a_L \rightarrow 1$, we have $GTE^{TSRN}(T | a_C, a_L) \rightarrow GTE^{CR}(T | a_C)$, the naive CR estimator. Similarly, because $a_C \rightarrow 1$, we have $GTE^{TSRN}(T | a_C, a_L) \rightarrow GTE^{LR}(T | a_L)$, the naive LR estimator. In this sense, the naive TSR estimator naturally “interpolates” between the naive LR and CR estimators. In the next section, we exploit this interpolation to choose a_C and a_L as a function of market conditions in such a way as to reduce bias.

More generally, the TSR design also contains much more information about competition in the marketplace, and the resulting interference effects, than either the CR or LR designs. Inspired by this observation, together with the idea of interpolating between the naive CR estimator and the naive LR estimator, in Section 6.3 we explore alternative, more sophisticated TSR estimators that heuristically debias interference due to

competition effects. As we show, these estimators can offer substantial bias reduction over the naive TSRN estimator above.

6. Analysis of Bias

We now utilize the framework defined to analyze the bias of two common experiment types, LR experiments and CR experiments. Recall from Section 1 that, in a setting where listings are immediately replenished, all customers see the full set of original listings as available. There is no competition between customers, but there is still competition between listings, and so intuitively, we expect CR to be unbiased and LR to be biased. Meanwhile, in a setting where listings remain unavailable for some amount of time, the resulting dynamic linkage across customers creates a bias in CR as well. Now, consider the extreme where the market is highly supply-constrained; most customers who arrive see no available listings, but some customers arrive just as a booked listing becomes available and see a single available listing. Such customers compare the listing against the outside option but, because no other listings are available, do not compare listings against each other. In this regime, there is no competition across listings, but there is competition between customers, and so we expect LR to be unbiased and CR to be biased.

In this section, we formalize this intuition about the behavior of the estimators in the extremes of market balance. We establish two key theoretical results; in the limit of a highly supply-constrained market (where $\lambda=\tau \rightarrow \infty$), the naive LR estimator becomes an unbiased estimator of the GTE, whereas the naive CR estimator is biased. On the other hand, in the limit of a highly demand-constrained market (where $\lambda=\tau \rightarrow 0$), the naive CR estimator becomes an unbiased estimator of the GTE, whereas the naive LR estimator is biased. In other words, each of the two naive designs is respectively unbiased in the limits of extreme market imbalance. At the same time, we find empirically that neither estimator performs well in the region of moderate market balance.

Inspired by this finding, we consider TSR and associated estimators that naturally interpolate between the two naive designs, depending on market balance. Given the findings above, we show that a simple approach to adjusting a_c and a_L as a function of market balance yields performance that balances between the naive LR estimator and the naive CR estimator. Nevertheless, we show that there is significant room for improvement by adjusting for the types of interference that arise, using observations from the TSR experiment. In particular, we propose a heuristic for a novel interpolating estimator for the TSR design that

aims to correct these biases and yields surprisingly good numerical performance.

6.1. Theory: Steady-State Bias of CR and LR in Unbalanced Markets

In this section, we study theoretically the bias of the steady-state naive CR and LR estimators in the limits where the market is extremely unbalanced (either demand-constrained or supply-constrained). The key tool we employ is a characterization of the asymptotic behavior of $Q_{ij}(\infty | a_c, a_L)$, as defined in (18) in the limits where $\lambda=\tau \rightarrow 0$ and $\lambda=\tau \rightarrow \infty$. We use this characterization in turn to quantify the asymptotic bias of the naive estimators relative to the GTE. We derive these results in the next two sections and provide a simple example in Section 6.1.3 to illustrate the effects.

6.1.1. Highly Demand-Constrained Markets. We start by considering the behavior of naive estimators in the limit where $\lambda=\tau \rightarrow 0$. We start with the following proposition that characterizes behavior of $Q_{ij}(\infty | a_c, a_L)$ as $\lambda=\tau \rightarrow 0$. The proof is in Online Appendix A.

Proposition 2. Fix all system parameters except λ and τ , and consider a sequence of systems in which $\lambda=\tau \rightarrow 0$. Then along this sequence,

$$\frac{1}{\lambda} Q_{ij}(\infty | a_c, a_L) \xrightarrow{\theta \rightarrow 0} \phi_{y,i} p_{y,i}(\theta, j | r): \quad (22)$$

The expression on the right hand side depends on both a_c and a_L through $\phi_{y,i}$ and r , respectively. In particular, we recall that $\phi_{y,i} = a_c \phi_{y,i}(\theta, j | r)$ and $r = p(\theta, 1) / a_L p(\theta)$. In our subsequent discussion in this regime, to emphasize the dependence of r on a_L below, we will write $r(a_L) = p(\theta, j | a_L) / p(\theta, 0 | a_L)$. With this definition, we have $p(\theta, 1 | a_L) = a_L p(\theta) / (1 - a_L p(\theta))$.

This proposition allows us to characterize the bias of both CR and LR estimators in the demand-constrained limit. Note that Proposition 2 shows that, in this limit, the (scaled) rate of booking behaves as if the available listings of type (θ, j) were exactly $p(\theta, j | a_L)$ for every θ and treatment condition $j \in \{0, 1\}$. That is, it is as if every arriving customer sees the entire mass of listings as being available, and so bookings are immediately replenished. This observation drives our first main result, that in the demand-constrained limit the CR estimator is unbiased and LR estimator is biased.

Theorem 3. Consider a sequence of systems where $\lambda=\tau \rightarrow 0$. Then for all a_c such that $0 < a_c < 1$, $\text{GTE}^{\text{CR}}(\infty | a_c) = \lambda - \text{GTE} = \lambda \rightarrow 0$. However, for $0 < a_L < 1$, generically over parameter values,⁷ we have $\lim \text{GTE}^{\text{LR}}(\infty | a_c) = \lambda - \text{GTE} = \lambda \neq 0$.

The full proof can be found in Online Appendix A. The key insight is that as the market becomes more demand constrained, there is a weakening of the competition between arriving customers, which leads to less interference in a CR experiment. In the limit, the CR estimator becomes unbiased. On the other hand, in an LR experiment, there is a positive mass of control and treatment listings available in steady state, leading to competition between listings and bias in the LR estimator.

6.1.2. Heavily Supply-Constrained Markets. We now characterize the behavior of naive estimators in the limit where $\lambda=\tau \rightarrow \infty$. We start with the next proposition, where we study the behavior of Q_{ij} as $\lambda=\tau \rightarrow \infty$. The proof is in Online Appendix A. To state the proposition, we define

$$g_{y,i}(\theta, j) \frac{\alpha_{y,i}(\theta, j)v_{y,i}(\theta, j)}{\phi_{y,i}(\theta, j)}.$$

Proposition 3. Fix all system parameters except for λ and τ , and consider a sequence of systems in which $\lambda=\tau \rightarrow \infty$. Along this sequence, the following limit holds

$$\frac{1}{\tau} Q_{ij}(\infty | a_C, a_L) \rightarrow \frac{\phi_{y,i}(\theta, j)}{\theta} \frac{g_{y,i}(\theta, j)}{\phi_{y,i}(\theta, j)} \rho(\theta, j)v(\theta); \quad (23)$$

As before, the expression on the right-hand side depends on both a_C and a_L through $\phi_{y,i}$ and r respectively. In particular, we recall that $\phi_{y,1} = a_C \phi_y, \phi_{y,0} = (1 - a_C) \phi_y$ and $\rho(\theta, 1) = a_L \rho(\theta), \rho(\theta, 0) = (1 - a_L) \rho(\theta)$.

A key intermediate result we employ is to demonstrate that in the steady state in this limit, $s^{\mathbb{P}}(\theta, j | a_C, a_L) \rightarrow 0$ for all θ, j . We know that in the steady state of the mean field limit, the rate at which occupied listings become available must match the rate at which available listings become occupied (flow conservation). We use this fact to show that to first order in $\lambda=\tau$, in the limit where $\lambda=\tau \rightarrow \infty$ we have

$$s^{\mathbb{P}}(\theta, j | a_C, a_L) \approx \frac{1}{\lambda=\tau} \frac{\rho(\theta, j)v(\theta)}{\phi_{y,i}(\theta, j)}.$$

The proposition follows by using this limit to characterize the choice probabilities.

The proof of the preceding proposition reveals that in the limit where $\lambda=\tau \rightarrow \infty$, we have

$$p_{y,i}(\theta, j | s^{\mathbb{P}}(a_C, a_L)) \approx g_{y,i}(\theta, j) s^{\mathbb{P}}(\theta, j | a_C, a_L) \frac{\alpha_{y,i}(\theta, j)v_{y,i}(\theta, j)s^{\mathbb{P}}(\theta, j | a_C, a_L)}{g_{y,i}(\theta, j)}.$$

This preceding expression is the formalization of our intuition that, in the limit where the market is heavily supply constrained, it is as if each arriving customer seeing an available listing compares only that listing

to the outside option; there is no longer any competition between listings.

We can use the preceding proposition to understand the behavior of the GTE, the naive LR estimator, and the naive CR estimator in steady state as $\lambda=\tau \rightarrow \infty$. For simplicity, we hold τ constant and take the limit $\lambda \rightarrow \infty$. In this case, the preceding proposition shows that

$$Q_{11}(\infty | 1, 1) \rightarrow \tau \rho(\theta)v(\theta); \quad Q_{00}(\infty | 0, 0) \rightarrow \tau \rho(\theta)v(\theta); \quad \epsilon$$

The global treatment effect is GTE $\rightarrow 0$ in this limit. Bookings occur essentially instantaneously after a listing becomes available, which happens at rate $\tau \rho(\theta)v(\theta)$.

We also note that

$$Q_{11}(\infty | 1, a_L) \rightarrow a_L \tau \rho(\theta)v(\theta); \quad Q_{10}(\infty | 1, a_L) \rightarrow (1 - a_L) \tau \rho(\theta)v(\theta);$$

The preceding two expressions reveal that the steady-state naive LR estimator GTE $(\infty | a_L)$ in this setting approaches zero, matching the GTE; thus it is asymptotically unbiased.

It is also now straightforward to see why the CR design will be biased. Note that

$$Q_{11}(\infty | a_C, 1) \rightarrow a_C \tau \frac{\phi_{y,1}(\theta, 1)}{\phi_{y,1}(\theta, 1) + \phi_{y,0}(\theta, 1)} g_{y,1}(\theta, 1) \rho(\theta)v(\theta);$$

An analogous expression holds for $Q_{01}(\infty | a_C, 1)$. We see that the right-hand side reflects the dynamic interference created between treatment and control customers; just as in our simple example, whether an available listing is seen by, for example, a control customer depends on whether it has previously been booked by a treatment customer. That is, customers compete for listings. As in the example, the naive CR estimator will remain nonzero in general in the limit, even though the GTE approaches zero.

We summarize our discussion in the following theorem.

Theorem 4. Consider a sequence of systems where $\lambda=\tau \rightarrow \infty$. Then, GTE $= \tau \rightarrow 0$, and for all a_L such that $0 < a_L < 1$, there also holds GTE $(\infty | a_L) = \tau \rightarrow 0$. However, for $0 < a_C < 1$, generically over parameter values we have $\lim_{\lambda \rightarrow \infty} \text{GTE}(\infty | a_C) = \tau - \text{GTE} = \tau \neq 0$.

Although the preceding theorem shows that the absolute bias of the naive LR estimator approaches zero, in fact in general the relative bias $(\text{GTE}(\infty | a_L) - \text{GTE})/\text{GTE}$ will not generally approach zero; this is because the GTE is also approaching zero, and so the second-order behavior of the naive LR estimator

matters. This is in contrast to the behavior of the naive CR estimator in the demand-constrained limit; in that limit, the GTE remains nonzero in general, and so the naive CR estimator is both absolutely and relatively unbiased. Nevertheless, note that relative bias of the naive LR estimator will be significantly smaller than the relative bias of the naive CR estimator in the supply-constrained limit since the naive CR estimator has a nonzero absolute bias in this limit, whereas the GTE approaches zero.

We finish this section by noting that Theorems 3 and 4 are driven by fundamental competition dynamics in the respective limiting regimes, and therefore, we believe that they hold under a much more general class of models than the ones considered here. We leave this for future investigation.

6.1.3. An Example: Homogeneous Customers and Listings. To more clearly understand the behavior of the bias, in this section we apply Propositions 2 and 3 to a simpler setting where both listings and customers are homogeneous; that is, there is only one type of customer and one type of listing. This example illustrates the symmetry between the two sides of the market and the resulting implications for bias in marketplace experiments.

Let v denote the control utility and \tilde{v} the treatment utility of a customer for a listing. Let ρ denote the outside option value of both control and treatment customers, $\alpha_0(0) = \alpha_1(1) = 1$, and $v(0) = v(1) = 1$. In this example, we consider two limits: one where λ is fixed and $\tau \rightarrow \infty$ (the demand-constrained regime) and one where τ is fixed and $\lambda \rightarrow \infty$ (the supply-constrained regime).

In the first case, when $\tau \rightarrow \infty$ with λ fixed, if we apply Proposition 2, we obtain

$$\begin{aligned} Q_{00}(\infty | a_C, a_L) &\rightarrow \lambda \cdot \frac{(1 - a_C)(1 - a_L)\rho v}{+ \rho v}; \\ Q_{01}(\infty | a_C, a_L) &\rightarrow \lambda \cdot \frac{(1 - a_C) a_L \rho v}{+ \rho v}; \\ Q_{10}(\infty | a_C, a_L) &\rightarrow \lambda \cdot \frac{(a_C(1 - a_L)\rho v)}{+ (1 - a_L)\rho v + a_C \tilde{v}}; \\ Q_{11}(\infty | a_C, a_L) &\rightarrow \lambda \cdot \frac{a_C a_L \rho \tilde{v}}{+ (1 - a_C)\rho v + a_L \tilde{v}}. \end{aligned} \quad (24)$$

In this limit,

$$GTE \rightarrow \lambda \cdot \frac{\rho \tilde{v}}{+ \rho v} - \frac{\rho v}{+ \rho v}.$$

From these expressions it is clear that the naive CR estimator is unbiased, whereas the naive LR estimator is biased. Furthermore, the expressions reveal that listing-side randomization creates interference across listings.

In the second case, when $\lambda \rightarrow \infty$ with τ is fixed, if we apply Proposition 3, we obtain

$$\begin{aligned} Q_{00}(\infty | a_C, a_L) &\rightarrow \tau(1 - a_C)(1 - a_L)\rho; \\ Q_{01}(\infty | a_C, a_L) &\rightarrow \tau \cdot \frac{(1 - a_C)v}{(1 - a_C)v + a_C \tilde{v}} a_L \rho; \\ Q_{10}(\infty | a_C, a_L) &\rightarrow \tau a_C(1 - a_L)\rho; \\ Q_{11}(\infty | a_C, a_L) &\rightarrow \tau \cdot \frac{a_C \tilde{v}}{(1 - a_C)v + a_C \tilde{v}} a_L \rho. \end{aligned} \quad (25)$$

In this limit, $GTE \rightarrow 0$. From these expressions, it is clear that the naive CR estimator is biased, whereas the naive LR estimator is unbiased. Furthermore, these expressions also reveal that customer-side randomization creates interference across customers.

Interestingly, these expressions highlight a remarkable symmetry. As expected, in the limit of a highly demand-constrained market, customers choose among listings; thus there is competition for customers among listings, and this is the source of potential interference in LR designs. The expressions reveal that in the limit of a highly supply-constrained market, it is as if listings choose among customers; thus there is competition among customers, and this is the source of potential interference in CR designs. Indeed, the expressions for Q_{01} and Q_{11} in (25) take the form of a multinomial logit choice model of listings for customers. We believe this type of symmetry provides important insight into the nature of experimental design in two-sided markets and in particular the roots of the interference typically observed in such settings.

6.1.4. Sign of the Bias in CR and LR Estimates. Theorems 3 and 4 state that the LR estimate is biased in the demand-constrained limit and the CR estimate is biased in the supply-constrained limit but make no claim as to whether the estimators overestimate or underestimate the GTE. In general, we cannot provide guarantees for the sign of the bias because it depends on the distribution of listings, the rates at which listings replenish, and the lift on the individual $\alpha_y(\theta)$ and $v_y(\theta)$ induced by the interventions. However, for a broad class of interventions, we can show that the LR estimate overestimates in the demand-constrained limit and CR overestimates in the supply-constrained limit. In such cases where we know the bias to be positive, CR and LR experiments can be used to bound the size of the GTE.

We call an intervention positive if $\alpha_y(\theta)v_y(\theta) > \alpha_y(\theta)v_y(\theta)$ for all y and θ . Such an intervention can be viewed as an improvement on the platform for all customer and listing type pairs, since for each pair at least one of the customer's consideration probabilities or utilities for the listing type must increase. Note that this class of interventions is broad enough to allow for

heterogeneous treatment effects across different (γ, θ) pairs.⁸

For positive interventions, straightforward applications of Propositions 2 and 3 show that GTE^{LR} overestimates the GTE in the demand-constraint limit and GTE^{CR} overestimates the GTE in the supply-constrained limit. The result follows from the fact that in a customer-side (resp., listing-side) experiment in a supply-constrained (resp., demand-constrained) setting, the individuals in the treatment group face less competition than they would in the global treatment setting, whereas the individuals in the control group face more competition than in the global control setting.

Proposition 4. Suppose that the treatment is positive, that is, $\alpha_\gamma(\theta)v_\gamma(\theta) > \alpha_\gamma(\theta)v_\gamma(\theta)$ for all γ, θ . Then we have the following.

1. LR bias when demand constrained: Consider a sequence of systems where $\lambda=\tau \rightarrow 0$. For any $0 < a_L < 1$, we have $\lim \text{GTE}^{\text{LR}}(\infty | a_C) = \lambda - \text{GTE} = \lambda > 0$.

2. CR bias when supply constrained: Consider a sequence of systems where $\lambda=\tau \rightarrow \infty$. For any $0 < a_C < 1$, we have $\lim \text{GTE}^{\text{CR}}(\infty | a_C) = \tau - \text{GTE} = \tau > 0$.

Furthermore, we find through simulations that CR and LR overestimate the GTE with positive treatments in intermediate ranges of market balance for all parameter regimes that we study in the examples in this section (see Figure 3 as well as Online Appendix C). We do find in some cases that the TSRN estimators underestimate the GTE, and so, because we plot bias on a log scale, we report the absolute value of the bias.

6.2. Discussion: Violation of SUTVA

Our results on the bias of the naive CR and LR experiments can be interpreted through the lens of the classical potential outcomes model. An important result from this literature is that when the stable unit treatment value assumption (SUTVA) holds, then naive estimators of the sort we consider will be unbiased for the true treatment effect. SUTVA requires that the treatment condition of units other than a given customer or listing should not influence the potential outcomes of that given customer or listing. The discussion above illustrates that in the limit where $\lambda \rightarrow 0$, there is no interference across customers in the CR design; this is why the naive CR estimator is unbiased. Similarly, in the limit where $\lambda \rightarrow \infty$, there is no interference across listings in the LR design; this is why the naive LR estimator is unbiased. On the other hand, the cases where each estimator is biased involve interference across experimental units.

6.3. Estimation With the TSR Design

The preceding sections reveal that each of the naive LR and CR estimators has its virtues, depending on market balance conditions. In this section, we explore whether we can develop TSR designs and estimators in which a_C and a_L are chosen as a function of $\lambda=\tau$ to obtain the beneficial asymptotic performance of the naive CR estimator in the highly demand-constrained regime as well as the LR estimator in the highly supply-constrained regime. We also expect that an appropriate interpolation should yield a bias for TSR that is comparable with if not lower than CR and LR in intermediate regimes of market balance.

Recall the naive TSRN estimator defined in (21) and in particular the steady-state version of this estimator. Suppose the platform observes $\lambda=\tau$; note that this is reasonable from a practical standpoint because this is a measure of market imbalance involving only the overall arrival rate of customers and the average rate at which listings become available. For example, consider the following heuristic choices of a_C and a_L for the TSR design for some fixed values of \bar{a}_C and \bar{a}_L :

$$\begin{aligned} a_C(\lambda=\tau) &= (1 - e^{-\lambda=\tau}) + \bar{a}_C e^{-\lambda=\tau}; \\ a_L(\lambda=\tau) &= \bar{a}_L (1 - e^{-\lambda=\tau}) + e^{-\lambda=\tau}. \end{aligned} \quad (26)$$

Then, as $\lambda=\tau \rightarrow 0$, we have $a_C(\lambda=\tau) \rightarrow \bar{a}_C$ and $a_L(\lambda=\tau) \rightarrow 1$, whereas as $\lambda=\tau \rightarrow \infty$, we have $a_C(\lambda=\tau) \rightarrow 1$ and $a_L(\lambda=\tau) \rightarrow \bar{a}_L$.⁹ With these choices, it follows that in the highly demand-constrained limit ($\lambda=\tau \rightarrow 0$), the TSRN estimator becomes equivalent to the naive CR estimator, whereas in the highly supply-constrained limit ($\lambda=\tau \rightarrow \infty$), the TSRN estimator becomes equivalent to the naive LR estimator. In particular, using Propositions 2 and 3, it is straightforward to show that the steady-state naive TSRN estimator is unbiased in both limits; we state this as the following theorem and omit the proof.

Corollary 1. For each $\lambda=\tau$, consider the TSR design with a_C and a_L defined as in (26). Consider a sequence of systems where either $\lambda=\tau \rightarrow 0$ or $\lambda=\tau \rightarrow \infty$. Then, in either limit,

$$\text{GTE}^{\text{TSRN}}(\infty | a_C(\lambda=\tau), a_L(\lambda=\tau)) - \text{GTE} \rightarrow 0.$$

We are also led to ask whether we can improve upon the naive TSRN estimator when the market is moderately balanced. Note that the TSRN estimator does not explicitly correct for either the fact that there is interference across listings or the fact that there is interference across customers. We now suggest a heuristic for correcting these effects, which we use to define two improved interpolating TSR estimators; these estimators are the fourth and fifth estimators appearing in Figure 2,

which we call “TSR-Improved (1)” and “TSR-Improved (2)”. These effects are visualized in Figure 1.

First, abusing notation, let $\text{GTE}^{\text{CR}}(T | a_C, a_L)$ denote the estimator in (19) using the same terms from a TSR design and dividing through by a_L on both terms as normalization. Similarly abusing notation, let $\text{GTE}^{\text{LR}}(T | a_C, a_L)$ denote the estimator in (20) using the same terms from a TSR design and dividing through by a_C on both terms as normalization. Motivated by these naive estimators, we explicitly consider an interpolation between the LR and CR estimators of the form

$$\begin{aligned} & \beta \text{GTE}^{\text{CR}}(T | a_C, a_L) + (1 - \beta) \text{GTE}^{\text{LR}}(T | a_C, a_L) \\ & \beta \frac{Q_{11}(T | a_C, a_L)}{a_C a_L} - \frac{Q_{01}(T | a_C, a_L)}{(1 - a_C)a_L} \\ & + (1 - \beta) \frac{Q_{11}(T | a_C, a_L)}{a_C a_L} - \frac{Q_{10}(T | a_C, a_L)}{a_C(1 - a_L)} : \end{aligned} \quad (27)$$

Now, consider the quantity $Q_{00}(T | a_C, a_L) = ((1 - a_C)(1 - a_L)) - Q_{10}(T | a_C, a_L) = ((1 - a_C)a_L)$ in a TSR design. This is the (appropriately normalized) difference between the rate at which control customers book control listings and the rate at which treatment customers book control listings. Note that both treatment and control customers have the same utility for control listings, because of the TSR design, but potentially different utilities for treatment listings. Hence, the difference in steady-state rates of booking among control and treatment customers on control listings must be driven by the fact that treatment customers substitute bookings from control listings to treatment listings (or vice versa). This difference captures the “cannibalization” effect (i.e., interference) that was found in LR designs in the demand-constrained regime.

Thus motivated, we can think of this difference as a “correction term” for the LR design from our interpolating TSR estimator in (27). Using a symmetric argument, we can also consider an appropriately weighted correction term associated to interference across customers in a CR design: $Q_{00}(T | a_C, a_L) = ((1 - a_L)(1 - a_C)) - Q_{01}(T | a_C, a_L) = ((1 - a_C)a_L)$. (Similar estimates were also studied in Bajari et al. (2021); see the related work for further details on this work.) See Figure 1 for an illustration of these competition effect estimates.

We can weight these correction terms with different factors $k > 0$ to control their impact. In addition, we can choose these weights in a market balance-dependent fashion based on the direction of market balance in which we have seen that the respective interference grows. Combining these insights, for $\beta \in (0, 1)$ and $k > 0$, we define a class of improved TSR estimators given by

$$\begin{aligned} & \text{GTE}^{\text{TSRI-}k}(T | a_C, a_L) \beta \frac{Q_{11}(T | a_C, a_L) - Q_{01}(T | a_C, a_L) c}{a_C a_L} \frac{c}{(1 - a_C)a_L} \\ & - k(1 - \beta) \frac{Q_{00}(T | a_C, a_L)}{(1 - a_C)c(1 - a_L)} \\ & - \frac{Q_{01}(T | a_C, a_L)}{(1 - a_C)a_L} \\ & + (1 - \beta) \frac{Q_{11}(T | a_C, a_L)}{a_C a_L} - \frac{Q_{10}(T | a_C, a_L)}{a_C(1 - a_L)} \\ & - k\beta \frac{Q_{00}(T | a_C, a_L)}{(1 - a_C)(1 - a_L)} - \frac{Q_{10}(T | a_C, a_L)}{a_C(1 - a_L)} : \end{aligned} \quad (28)$$

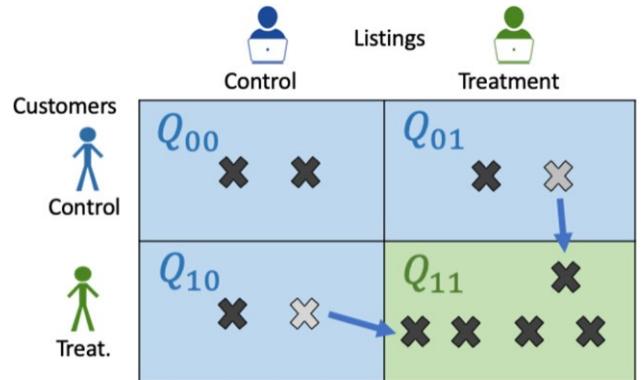
Given market balance $\lambda = \tau$, we set $\beta = e^{-\lambda = \tau}$, and we choose a_C and a_L as in (26).

In the limit where $\lambda = \tau \rightarrow 0$, note that $\text{GTE}^{\text{TSRI-}k}(T | a_C(\lambda = \tau), a_L(\lambda = \tau))$ approaches $\text{GTE}^{\text{CR}}(T | \bar{a}_C)$ as expected. Similarly, in the limit where $\lambda = \tau \rightarrow \infty$, $\text{GTE}^{\text{TSRI-}k}(T | a_C(\lambda = \tau), a_L(\lambda = \tau))$ approaches $\text{GTE}^{\text{LR}}(T | \bar{a}_L)$. It is straightforward to show that $\text{GTE}^{\text{TSRI-}k}$ for any k is unbiased in both the highly demand-constrained and highly supply-constrained regimes, because the correction terms play no role in the limits.

For moderate values of market balance, both the cannibalization correction terms kick in, which leads to improvements over naive TSRN, as seen in Figure 2. To simplify the exposition, we consider only two factors

k_1, k_2 ; we see that TSR-1 has lower bias than the native TSR, but TSR-2, which has a higher weight in front of the correction terms, has a lower bias than both

Figure 1. (Color online) Illustration of TSR Design and Competition Effects



\mathbb{X} denotes bookings made

Notes. Intervention applies only when treatment customers view treatment listings (green cell). Suppose that the intervention makes the listing more attractive.

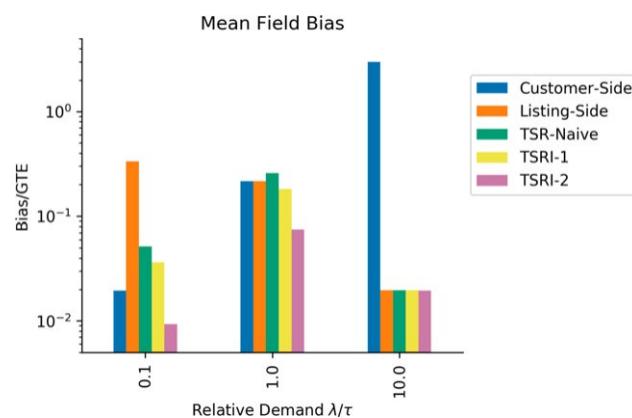
naive TSRN and TSRI-1 as well as the naive CR and LR estimators. In Online Appendix C, we explore the robustness of our results to other model primitives, specifically scenarios with smaller or larger utilities and the introduction of heterogeneity on one or both sides of the market. We find that the bias of CR and LR estimators can increase with the introduction of these factors, but remarkably the bias of the TSRI estimators remains low across the ranges that we study. We emphasize the fact that the three TSR estimators presented here are examples to illustrate the potential for bias reduction using this new design. There is of course a much broader range of both TSR designs and estimators; some of these may offer even better performance.

We conclude this section with two additional observations. First, note that all our analysis in this section has been carried out in the mean field steady state; in particular, Figure 2 shows the bias of the estimators in steady state. For practical implementation, it is also important to consider the relative bias in the candidate estimators in the transient system, since experiments are typically run for relatively short time horizons. For discussion of the transient behavior for a finite time horizon, see Online Appendix D. Second, in the next section, we discuss variance of the estimators we have studied. There we find that the TSR estimators with the lowest bias also have the highest variance; in other words, there is a bias-variance tradeoff.

7. Bias-Variance Tradeoff of Estimators

Our mean field model is deterministic, so it does not allow us to study the variance of the different estimators.

Figure 2. (Color online) Difference Between Estimator and GTE in Steady State



Notes. We consider variation in $\lambda=\tau$ by fixing $\tau=1$ and varying λ ; analogous results are obtained if λ is fixed and τ is varied. We consider a market with homogeneous customers and homogeneous listings pretreatment. We set $\alpha=1$ and $\alpha=0.5$. In the CR design, $a_C=0.5$. In the LR design, $a_L=0.5$. Customers have utility $v=0.315$ for control listings and $v=0.394$ for treatment listings, which corresponds to a steady state booking probability of 20% in global control and 23% global treatment when $\lambda=\tau$.

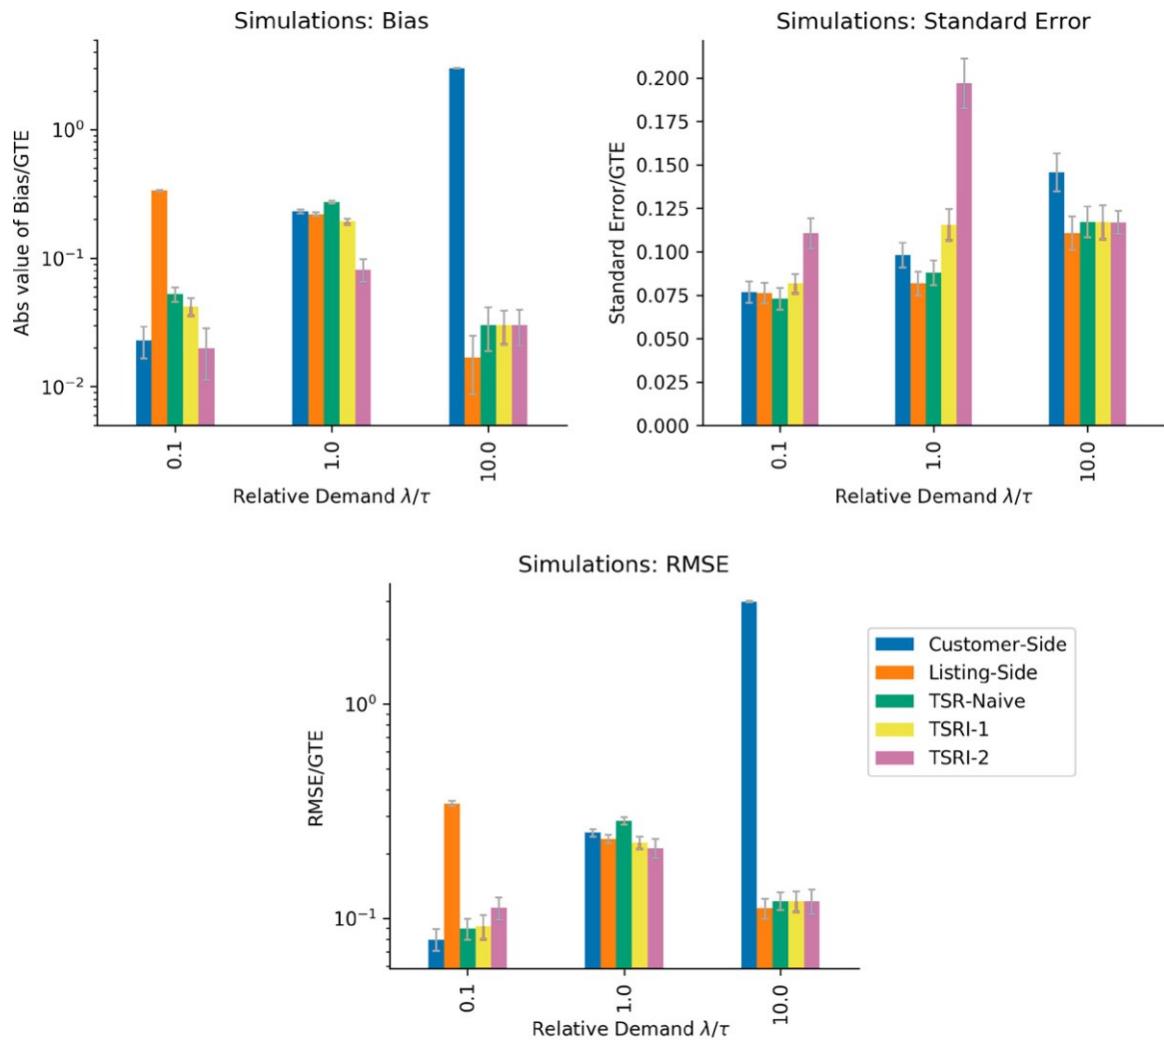
In practice, however, markets consist of finitely many listings, and experiments are run for finite time horizon T , and so the variance of any estimator will be non-zero.¹⁰ In particular, the variance of estimators becomes an important consideration alongside bias, particularly in choosing between multiple estimators with similar bias. The variance of the TSR estimators is especially important, given the earlier discussion that many heuristics that platforms use to minimize bias do so at the cost of increased variance, leading to underpowered experiments (see Section 2).

With this background as motivation, in this section we provide a preliminary yet suggestive simulation study of variance. The simulations highlight two important considerations that a platform must take into account when designing and analyzing an experiment. First, similar to the results from Section 6 on bias, we find that the estimator with the lowest variance depends on market balance. Second, we see a bias-variance tradeoff between the CR, LR, and TSR estimators, with the TSR estimators offering bias improvements at the cost of an increase in variance. We emphasize the point that whether a platform should care more about bias or variance depends on the size of the platform (number of listings N) and the time horizon on which the experiment is run. The bias of the experiment is relatively unaffected by changes in these two factors, but of course variance decreases in the size of the market and the length of the time horizon. Thus these two factors dictate whether bias or variance contributes more to the overall RMSE.

Full details of the simulation environment and parameters are in Online Appendix C, which we briefly summarize here. We simulate marketplace experiments with varying market parameters for a finite system with a number of listings $N=5,000$ and fixed time horizon T . For each run of the simulation, we fix an experiment design (e.g., CR, LR, TSR) and simulate customer arrivals and booking decisions until time T . System evolution is simulated according to the continuous time Markov chain specified in (3)–(5). We calculate the estimator corresponding to the experiment design defined in (19)–(21) and (28) (for $k=1, 2$) for the time interval $[T_0, T]$, where T_0 is chosen to eliminate the transient burn-in period. We then simulate multiple runs and compare the bias and standard error of the estimators across runs. Note that we report the true standard errors calculated across simulation runs. For discussion on the estimation of standard errors, see Section 9.

Figure 3 shows simulations for a homogeneous system with only one customer and one listing type, with the same parameters as the mean field numerics presented in Figure 2. Note that the bias of the estimators in these large-market simulations echoes the qualitative insights about bias obtained from the mean field

Figure 3. (Color online) Homogeneous Listings and Customers



Notes. Top Left: bias of each estimator. Top right: Standard error of estimates. Bottom: RMSE of the estimates. Statistics are normalized by GTE. All statistics are calculated across 500 runs, with bootstrapped 95 percentile confidence intervals provided for each statistic. We consider a setting with homogeneous listings and customers, with the same utilities as defined in Figure 2. In the CR design, $a_C = 0.5$. In the LR design, $a_L = 0.5$. Simulation parameters are defined in Online Appendix C.

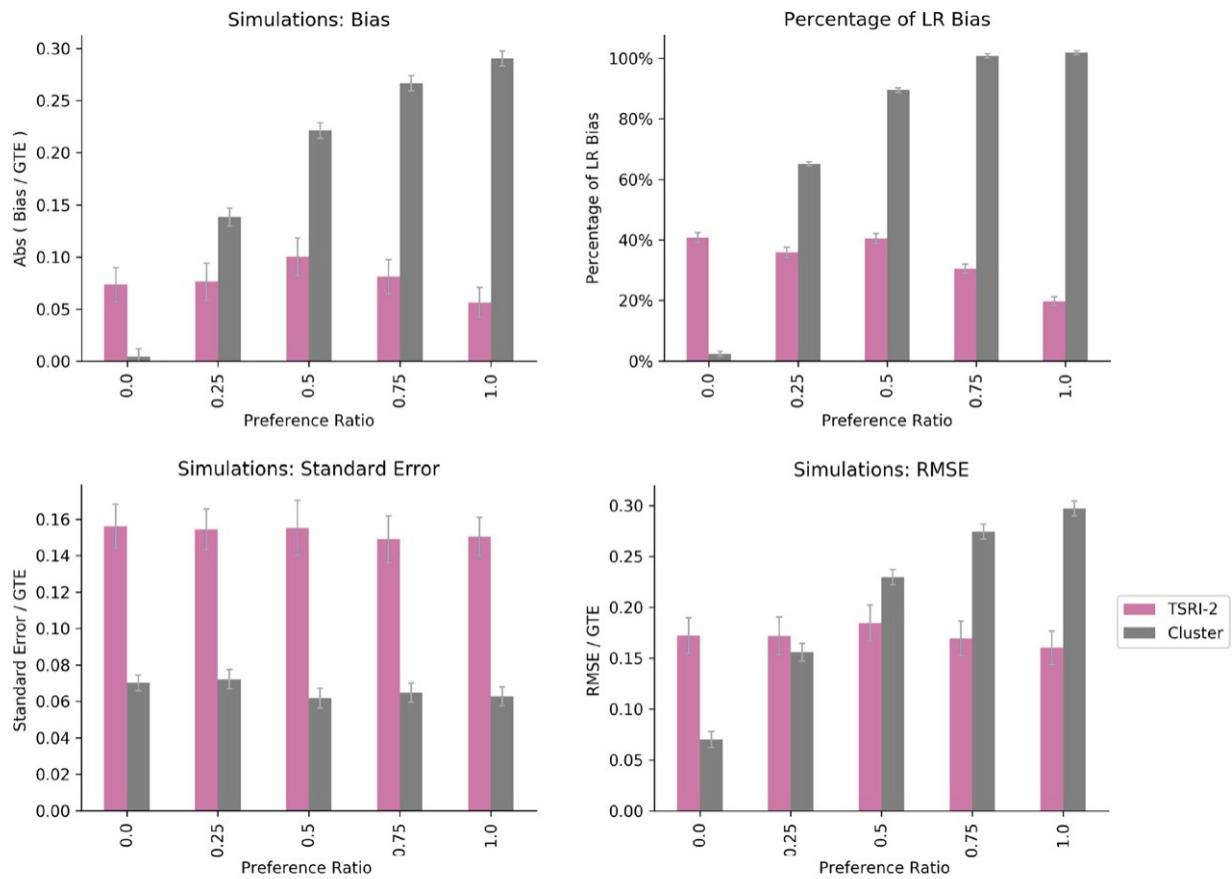
model. Similar findings are obtained in more general scenarios, see Appendix C, where we investigate the effect of heterogeneity in the marketplace.

These simulations point to a bias variance-tradeoff between TSR estimators and the naive CR and LR estimators as well as between the three TSR estimators themselves. The TSR estimators, as discussed earlier, offer benefits over the naive CR and LR estimators with respect to bias, but they do so at the cost of an increase in variance. Moreover, among the three TSR estimators that we explore, those with lower bias also have higher variance. The naive TSRN estimator has similar variance with the lowest of CR and LR, but the bias of this estimator is also similar to the lowest bias of CR and LR. On the other hand, TSR-2 shows a substantial improvement in bias over both CR and

LR for several market conditions, but this estimator also has the largest variance among all five estimators, especially in the regime of intermediate market balance (cf. Online Appendix C).

Furthermore, the minimum variance estimator depends on market conditions. For example, in a demand-constrained market with $\lambda=\tau = 0.1$, GTE^{CR} has the lowest standard error, whereas in a supply-constrained market with $\lambda=\tau = 10$, GTE^{CR} has the highest standard error.

We conclude this section by highlighting the potential for the class of TSR experiments, which open up a large class of both designs and estimators. Among the three estimators we explore, we see that there is a TSR estimator with low bias and another TSR estimator with low variance. It is possible that, with further optimization of

Figure 4. (Color online) Simulations of Cluster-Randomized Estimator and TSR Estimators as Preference Ratio y/x Varies

Notes. Relative demand is fixed at $\lambda=\tau=1$. We fix $x=0.5$ and vary y . Bootstrapped 95 percentile confidence intervals are provided for each statistic.

these designs and estimators, one can devise a new estimator that optimizes this bias-variance tradeoff.

8. Comparison with Cluster-Randomized Experiments

In this section, we compare the TSR approach to existing approaches to reduce bias in marketplace experiments. One such approach is to run a cluster-randomized experiment, which changes the unit of randomization in order to reduce interference effects across units. The typical approach is to divide the marketplace into clusters, such as geographical regions, such that there is less interaction of market participants across different clusters. All participants within a cluster receive the same treatment condition. The platform then estimates the GTE by comparing the outcomes within the treatment clusters versus the outcomes in the control clusters. It is important to note that many markets and social networks are highly connected, and it is not possible to avoid all interference across clusters; see, for example, Holtz et al. (2020) for an example in the context of Airbnb. Thus cluster-randomized experiments will reduce but not fully remove the bias.

To compare the performance of the cluster-randomized and TSR approaches, we use our existing model to define a regime that gives the best-case performance for cluster-randomized experiments, where there are tightly clustered preferences in the marketplace and the platform knows *ex ante* the true clusters (without having to learn them).

The simulations suggest that cluster-randomized estimators offer substantial bias reductions over when the market is tightly clustered, but these improvements diminish if the market becomes more interconnected. The TSR estimators, however, offer bias reductions in both clustered and interconnected markets and, in interconnected markets, are less biased than the cluster-randomized estimator. In our simulations, the variance of the cluster-randomized estimator is lower than that of TSR-2 in our example, although the variance of the cluster-randomized estimator will likely change if we deviate from this best-case scenario with perfect knowledge of the clusters, identical listings within clusters, and uniform treatment effects across different clusters. Hence, although our model for market clusters is stylized, we believe our results suggest that TSR designs

can be a useful alternative to cluster-randomized experiments in interconnected markets.

8.1. Market Setup for Cluster Randomization

We consider the case where clusters are defined on the listing side so that the cluster-randomized experiment is expected to improve upon a listing-randomized experiment.¹¹ To induce a clustered structure, we model a setting with two customer types and two listing types, where each type of customer prefers a different type of listing. Formally, there are customer types $\{\gamma_1, \gamma_2\}$ and listing types $\{\theta_1, \theta_2\}$. All customers consider all listings ($\alpha_y(\theta) = 1$ for all y, θ), but customers have different utilities for different listings.¹² The global control utilities $v_y(\theta)$ have the following form:

$$\begin{array}{c|cc} & \theta_1 & \theta_2 \\ \hline \gamma_1 & x & y \\ \gamma_2 & y & x \end{array}$$

where $x \geq y \geq 0$. Note that if $y = 0$, then the market can be perfectly decomposed into two submarkets, where customers of type γ_1 (resp., γ_2) book only listings of type θ_1 (resp., θ_2). If $y = x$, then each customer prefers both listings equally. Thus we can interpret the ratio y/x as a measure of how equally a customer prefers both products, where intuitively the market is tightly clustered when y/x is small. We call y/x the preference ratio.

The platform then runs a cluster-randomized experiment where it first assigns listings to clusters and then randomizes entire clusters to either treatment or control. In practice, the platform must learn how to create the clusters, likely through observational data in the global control setting,¹³ but in these simulations, we assume that the platform observes the cluster structure perfectly. The platform assigns all θ_1 listings to one cluster and θ_2 listings to another and runs a completely randomized design on the clusters, assigning one of the clusters to treatment and one to control. For simplicity, assume that the intervention has a multiplicative lift $\delta > 1$ on all customer-listing pairs so that the treatment utilities satisfy $v_y(\theta) = \delta v_y(\theta)$.

The cluster-randomized estimator GTE_{Cluster} , with clusters defined on the listing side, compares the (scaled) rate of bookings of listings in treatment clusters to the rate of bookings of listings in control clusters. Formally, in the mean field setting, once the clusters are randomized, let Z denote the mass of listings assigned to a treatment cluster. Then

$$GTE_{\text{Cluster}}(T | Z) = \frac{Q_{T1}(T | 1, Z)}{Q_{T0}(T | 1, Z)} : \quad (29)$$

We can similarly define an analogous estimator in the finite model.

The full set of parameters is as follows. The market has an equal number of listings of both types so that $m^{(N)}(\theta_1) = m^{(N)}(\theta_2) = 0.5$ and the same arrival rate for both customers types ($\lambda^{(N)} = \lambda_{\gamma_1}^{(N)}$). We set $\alpha = 0.5$ and vary $y \in [0, 1]$. We set $\delta = 1.3$. We fix $\tau = 1$ and consider $\lambda \in \{0.1, 1, 10\}$.

8.2. Results

Figure 4 shows how the performance of the estimators changes when we vary the preference ratio (at a fixed market balance). We choose TSRI-2 as a representative estimator for the TSR approach; see Figure 5 in Online Appendix B for a comparison with all estimators.

We see a clear takeaway that, perhaps unsurprisingly, the cluster-randomized estimator offers substantial bias improvements when the market is tightly clustered (i.e., a customer strongly prefers one type of item over another) but offers little reduction in bias when the market is more interconnected. In particular, the cluster-randomized estimator outperforms the TSR estimators when the market is tightly clustered, whereas the TSR estimators outperform the cluster-randomized estimator when the market is more connected.

The standard error changes little across the preference ratios, with the standard error of the cluster-randomized estimate lower than that of TSRI-2, although the variance of the cluster-randomized estimator will likely change if we deviate from this best-case scenario with perfect knowledge of the clusters, identical listings within clusters, and uniform treatment effects across clusters. In our scenario, the cluster-randomized estimator has lower RMSE for tightly clustered markets and higher otherwise.

9. Conclusion

This paper proposes a general mean field framework to study the dynamics of inventory bookings in two-sided platforms, and we leverage this framework to study the design and analysis of a number of different experimental designs and estimators. We study both commonly used designs and estimators (CR, LR) and also introduce a family of more general two-sided randomization designs and estimators (TSR). Our work sheds light on the market conditions in which each approach to estimation performs best based on the relative supply and demand balance in the marketplace.

For bias minimization, we suggest two directions for future work. The first is further optimization of the TSR design as a standalone experiment design. We have proposed three natural TSR estimators, but the space of both designs and estimators is much richer, and it is worth asking which are optimal with respect to bias and variance as well as how this answer may change with differing market conditions. The second

direction is to develop the TSR design as a method to debias one-sided experiments. The design allows us to measure competition effects between customers and between listings; this observation suggests that these measurements can be used to approximately debias existing CR and LR experiments, providing another route for platforms to utilize two-sided randomization designs.

To make technical progress in this paper, we employed several simplifying assumptions on the choice model and booking behavior. We believe that our core insight of market balance mediating competition effects, and thus affecting the resulting bias in an experiment, extends to other settings as well. We hypothesize, however, that some of our results are more robust to modeling choices than others. For example, the result that the CR estimator is unbiased in the demand-constrained limit depends only on the fact that booked listings are replenished in between customer arrivals and likely extends to any reasonable choice model. On the other hand, the result that the absolute bias of the LR estimator approaches 0 in the supply-constrained limit may be more sensitive to different choice models and in particular how the customer weighs options on the platform compared with the outside option. In this supply-constrained limit, we conjecture that the relative performance of the two experiment types still holds even in modified settings; that is, the LR estimator has lower bias than the CR bias. Furthermore, we believe that the approach of using TSR to observe competition effects and heuristically debias estimators also extends beyond our model.

Another practical consideration for platforms is that not all experiment designs are suitable for all types of interventions. In this paper, we have largely focused on interventions such as user interface changes that change how an individual customer perceives an individual listing. There are also other interventions that operate between subsets of customers or subsets of listings. For example, a modification in the ranking algorithm over the listings changes operates not on an individual customer-listing pair but rather changes how a customer interacts with a subset of listings. This intervention is more conducive to a CR experiment than an LR or TSR experiment. Beyond these feasibility constraints, it remains an open question whether certain types of experiments lead to lower bias in different classes of interventions.

Finally, we emphasize the importance of inference in these settings, which we do not study in this paper. In practice, standard errors are also estimated “naively”; they are typically computed assuming independence of observations. However, because of interference, observations are clearly not independent. In these settings, how biased might the standard error estimates be? How can experimenters derive valid confidence

intervals in these settings? Such questions are critical for any platforms controlling the false-positive and false-negative results arising from their experiments.

Acknowledgments

The authors thank Navin Sivanandam from Airbnb for numerous fruitful conversations. The authors also thank Guillaume Basse, Dominic Coey, Nikhil Garg, David Holtz, Michael Luca, Johan Ugander, and Nicolas Stier for their feedback as well as seminar participants at ACM EC’20, Facebook, Marketplace Algorithms and Design, CODE 2020, INFORMS 2020, Lyft, Vinted, Marketplace Innovations Workshop 2021, MSOM 2021, and Stanford. Finally, the authors thank the ACM EC’20 review team for their thoughtful comments and engagement with their work.

Endnotes

¹ The same modeling framework that we employ in this paper can be used to consider interventions that change other parameters, such as customer arrival rates or the time that listings remain occupied when booked; such application is outside the scope of our current work.

² Our framework can also be used to evaluate other metrics of interest based on experimental outcomes; for simplicity, we focus on rate of booking in this work.

³ Because the system we study is irreducible and we analyze its steady-state behavior, it would not matter if we chose a different initial condition.

⁴ An even more general model might allow the occupancy time to depend on both listing type and the type of the customer who made the booking; such a generalization remains an interesting open direction.

⁵ Our current work allows us to relatively easily incorporate v , depending on treatment condition of the listing, and as such we can extend our results to study LR designs where v varies with treatment condition. In general, however, when customers are also randomized to treatment or control, the holding time parameter of a listing should also depend on the treatment condition of the customer who booked that listing. Adapting our framework to incorporate this possibility remains an interesting direction for future work.

⁶ Note that in two-sided markets, certain types of interventions will also cause long-run economic equilibration due to strategic responses on the part of market participants; for example, if prices are lowered during an experiment, this may affect entry decisions of both buyers and sellers and thus the long-run market equilibrium. Although our model allows the choice probabilities to change due to treatment, a more complete analysis of long-run economic equilibration due to interventions remains a direction for future work.

⁷ Here “generically” means for all parameter values, except possibly for a set of parameter values of Lebesgue measure zero.

⁸ A symmetric analysis can be applied for “negative” interventions, where $\hat{\alpha}_y(\theta)\hat{v}_y(\theta) < \alpha_y(\theta)v_y(\theta)$ for all y and θ , although, of course, interventions known to be negative in advance are less likely to be desirable from the platform’s perspective.

⁹ Our choice of exponent here is somewhat arbitrary; the same analysis follows even if we replace $e^{-\lambda=\tau}$ with $e^{-c\lambda=\tau}$ for any value of $c > 0$.

¹⁰ We note that even if the system consists of only a finite number of listings N , because $T \rightarrow \infty$, the standard error of the various estimators proposed in this paper will converge to zero. However, for finite T , this is not the case; because A/B tests are always run to a finite horizon T , this nonzero variance will impact the accuracy of any estimates obtained.

¹¹ Because we analyze a model where customers are short-lived, whereas listings remain on the platform, it is likely that the platform has more information on the listings and is better able to learn clusters on the listing side.

¹² Alternatively, we can induce a clustered structure by modifying the consideration probabilities $\alpha_v(\theta)$. Both approaches of modifying the $\alpha_v(\theta)$ and modifying the $v_v(\theta)$ are equivalent in the mean field model.

¹³ See Holtz et al. (2020) in the context of marketplaces and Ugander et al. (2013) in the context of social networks.

References

Athey S, Eckles D, Imbens GW (2018) Exact p-values for network interference. *Journal Am Stat Assoc* 113(521):230–240.

Bajari P, Burdick B, Imbens G, Masoero L, McQueen J, Richardson T, Rosen I (2021) Multiple randomization designs. Preprint, submitted December 27, <https://arxiv.org/abs/2112.13495>.

Basse GW, Feller A, Toulis P (2019) Randomization tests of causal effects under interference. *Biometrika* 106(2):487–494.

Basse GW, Soufiani HA, Lambert D (2016) Randomization and the pernicious effects of limited budgets on auction experiments. Gretton A, Robert CC, eds., *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS 2016, Cadiz, Spain, May 9–11, 2016, volume 51 of JMLR Workshop and Conference Proceedings*, 1412–1420.

Blake T, Coey D (2014) Why marketplace experimentation is harder than it seems: The role of test-control interference. *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, 567–582, EC '14 (Association for Computing Machinery, New York).

Chamandy N (2016) Experimentation in a ridesharing marketplace. Accessed June 20, 2020, <https://eng.lyft.com/experimentation-in-a-ridesharing-marketplace-b39db027a66e>.

Fradkin A (2015) Search frictions and the design of online marketplaces. Working paper, NBER.

Fradkin A (2019) A simulation approach to designing digital matching platforms. *SSRN Electronic Journal* ISSN 1556-5068.

Ha-Thuc V, Dutta A, Mao R, Wood M, Liu Y (2020) A counterfactual framework for seller-side a/b testing on marketplaces. *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2288–2296, SIGIR '20 (Association for Computing Machinery, New York).

Holtz D (2018) Limiting bias from test-control interference in online marketplace experiments. Ph.D. thesis, <https://dspace.mit.edu/handle/1721.1/117999>.

Holtz D, Lobel R, Liskovich I, Aral S (2020) Reducing interference bias in online marketplace pricing experiments.

Imbens GW, Rubin DB (2015) *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. (Cambridge University Press, Cambridge, MA).

Manski CF (2013) Identification of treatment response with social interactions. *Econometrics J* 16(1):S1–S23.

Pouget-Abadie J, Aydin K, Schudy W, Brodersen K, Mirrokni V (2019) Variance reduction in bipartite experiments through correlation clustering. *Adv. Neural Inf. Process. Syst.* 32: 13288–13298.

Saveski M, Pouget-Abadie J, Saint-Jacques G, Duan W, Ghosh S, Xu Y, Airoldi EM (2017) Detecting network effects: Randomizing over randomized experiments. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1027–1035, KDD '17 (Association for Computing Machinery, New York).

Sneider C, Tang Y, Tang Y (2019) Experiment rigor for switchback experiment analysis. Accessed June 20, 2020, <https://doordash.engineering/2019/02/20/experiment-rigor-for-switchback-experiment-analysis/>.

Ugander J, Karrer B, Backstrom L, Kleinberg J (2013) Graph cluster randomization: Network exposure to multiple universes. *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 329–337, KDD '13 (Association for Computing Machinery, New York).

Wager S, Xu K (2019) Experimenting in equilibrium. Preprint, submitted March 6, <https://arxiv.org/abs/1903.02124>.

Zigler CM, Papadogeorgou G (2018) Bipartite causal inference with interference. Preprint, submitted July 23, <https://arxiv.org/abs/1807.08660>.