Estimating Random Utility Models over Bundles with Large-Scale Linear Programming*

Enzo Di Pasquale†

September 15, 2026

Abstract

This paper develops a simulation-based estimator for random utility models over combinatorial choice sets. When utility is linear in parameters, minimizing the average regret—the difference in utility between the observed choice and the optimal alternative under simulated shocks—yields a convex objective function. Convexity overcomes the computational hurdles of the standard simulated method of moments, enabling the estimation of high-dimensional parameter vectors such as item fixed effects and complementarities. I recast the sample objective as a linear program and solve it via constraint generation. In the numerical experiments, constraint generation terminates in far fewer iterations than worst-case convex optimization benchmarks. I establish formal consistency and asymptotic normality in the standard many-observation limit ($N \to \infty$), with the number of choices and parameters fixed. I then use numerical experiments to demonstrate the method's scalability to single-profile ($N = 1$) large combinatorial objects, such as supermodular potential games. Finally, I apply the framework to the 1996 FCC C-Block spectrum auction. After estimating geographic complementarities and license-specific fixed effects from the realized allocation, I evaluate counterfactual bundling policies where licenses are aggregated into a coarser geographical partition, finding that it would have lowered both seller revenue and bidder surplus.

Keywords: discrete choice, combinatorial optimization, revealed preference, convex estimation

JEL Codes: C14, C25, C51

*I am grateful to Alfred Galichon, Daniel Waldinger, Quang Vuong, and Thomas J. Sargent for their guidance. I also thank Alberto Bisin, Christopher Conlon, Piergiorgio Di Pasquale, Jeff Gortmaker, Antoine Jacquet, Boyan Jovanovic, Konrad Menzel, Paul T. Scott, and Maxime Sylvestre for helpful discussions. I thank participants at the NYU IO Student Seminar, the math+econ+code Seminar Series, and Empirical Models of Matching Markets at IASC Cargèse. Comments welcome.

†New York University. Email: ed2189@nyu.edu.

1 Introduction

Many counterfactuals in industrial organization and trade depend on how strongly goods substitute or complement one another. Gentzkow (2007) develops a model of demand for bundles that allows for both substitution and complementarity, and uses it to estimate the welfare effects of introducing an online newspaper. Head et al. (2026) estimate a model of firms' joint location choices for battery production and vehicle assembly to study how consumer subsidies and local-content requirements affect domestic production and electric-vehicle adoption. Fox and Bajari (2013) estimate geographic complementarities among spectrum licenses to assess the efficiency of the C-Block allocation and quantify the gains from alternative regional assignments.

Consider bundle demand with quasi-linear utility and à la carte prices: consumers pay the sum of the prices of the goods they choose. With few goods, one can estimate a separate mean utility for each bundle, as in Gentzkow (2007). Each additional good doubles the number of possible bundles. I specify one mean utility per item and a limited set of complementarity terms, with item mean utilities absorbing unobserved quality as in Berry et al. (1995). Copurchase frequencies alone do not distinguish substitution or complementarity between goods from correlation in unobserved tastes (Gentzkow, 2007). I allow a flexible joint distribution of item-specific taste shocks, including correlation across goods.

Estimating these parameters by simulated moments requires repeatedly solving consumers' bundle-choice problems (McFadden, 1989). With fixed simulation draws, predicted bundles can switch as parameters change, producing jumps in simulated purchase moments. Inspired by empirical matching models (Galichon and Salanié, 2022), I construct an estimation criterion from the utility attained in these choice problems. Average regret measures the utility gain from replacing the observed bundle with the best feasible bundle under simulated tastes. When utility is linear in the parameters, the value of the best bundle is a maximum of linear functions. Subtracting the utility of the observed bundle gives a convex criterion.

I cast the sample criterion as a linear program and solve it by constraint generation. This avoids nonconvex optimization over the linear parameters, while the computational cost still depends on solving the underlying choice problems. I provide sufficient conditions for point identification and establish consistency of a level-set estimator of the identified set. Under point identification and additional regularity conditions, an exact sample minimizer is asymptotically normal. These results hold as the number of independent observations grows, with the number of elementary decisions, parameters, and simulation draws per observation fixed.

I apply the framework to the same FCC C-Block spectrum auction studied by Fox and Bajari (2013), using it to ask a product-design question: how finely should the FCC divide geographic rights before selling them? Before the auction, the regulator chooses the geographic scope of each license. A fine partition allows different bidders to win different local markets within a region. A coarse partition sells the region as one object, giving the winner the whole footprint but preventing separate assignment of its local components.

The FCC sold 480 continental Basic Trading Area (BTA) licenses. These BTAs can be aggregated into 46 Major Trading Areas (MTAs), each a larger regional unit. I ask how revenue and surplus would change if the same C-Block spectrum had been sold as MTA licenses rather

than BTA licenses.

This counterfactual separates two effects emphasized by the auction-theory literature. The first is an allocative-efficiency effect. Selling a package forces its components to be allocated together, even when different bidders have the highest valuations for different components. Relative to the BTA allocation problem, the MTA allocation problem imposes an additional constraint: all BTAs in an MTA must be assigned to the same bidder. Holding bidder primitives fixed, welfare is therefore weakly lower under the coarser partition. The second is a competition effect. Packaging pools each bidder's valuations across the components of the package. A bidder with a high valuation for one component may be drawn into competition for the larger object, potentially raising revenue. Thus packaging can reduce allocative efficiency and still raise seller revenue (Palfrey, 1983; Jehiel et al., 2007).

The C-Block setting adds two empirical features to this tradeoff. First, licenses may be complementary, so the value of an MTA license need not equal the sum of the values of its BTAs. Second, the auction imposed population-denominated eligibility constraints. A bidder with enough eligibility to bid on selected BTAs within an MTA need not have enough eligibility to bid on the entire MTA. Aggregation can therefore remove some bidders from competition for the regional object even if they were relevant competitors for its constituent local licenses.

Fox and Bajari (2013) estimate geographic complementarities using comparisons with hypothetical license swaps between bidders. Common license values and prices cancel in these comparisons, so their estimation exercise does not allow for a structural counterfactual that recomputes equilibrium prices after licenses are repackaged. I estimate the mean utilities and complementarities jointly, assuming that observed packages maximize utility at observed prices and specifying a distribution for bidder shocks. I then use observed prices and a price instrument to separate price sensitivity from license-specific unobserved quality.

I use the estimates to compute the allocation and supporting prices under the MTA partition, holding the bidder population and eligibility commitments fixed. Seller revenue and bidder surplus both fall relative to the BTA partition.

The framework also applies beyond individual bundle demand. A second class of examples is discrete games with strategic complementarities, including peer effects, network formation, and multi-market entry. In these settings the observed object is an action profile rather than a bundle chosen by one agent. Potential games fit the framework because equilibrium selection can be written as an optimization problem. In a potential game, players' unilateral payoff differences are represented by changes in a common potential function (Monderer and Shapley, 1996). I assume that the observed profile is the equilibrium selected by maximizing the realized potential. Every potential maximizer is a Nash equilibrium, although the game may have other equilibria.

In the supermodular games I study, computing the selected equilibrium amounts to maximizing a supermodular function over binary action profiles. Supermodular maximization over binary profiles is polynomial time (Schrijver, 2000; Iwata et al., 2001). I use this case for the large-game numerical experiments. With strategic substitutes, the corresponding problem is submodular maximization, which is NP-hard in general. The estimator does not change this computational fact.

Strategic complementarities are not the only force that can produce correlated actions. In a peer-effects model, neighbors may choose the same action because one player's action raises another player's payoff, or because neighbors share unobserved payoff shocks. In a network-formation model, links within a group may be correlated because links are complementary, or because players in the group share unobserved tastes for forming links. This is the social-interactions identification problem emphasized by Manski (1993): correlated effects can be mistaken for endogenous social effects.

The numerical experiments use the peer-effects and network-formation examples to ask whether the estimator can work with a single realized profile. In the peer-effects design, the profile is the vector of actions of 500 players on one network. In the network-formation design with 100 players, the profile is a directed adjacency matrix with 9,900 possible links. The Monte Carlo experiments recover the interaction coefficients and fixed effects in these single-profile designs. These single-profile experiments fall outside the asymptotic regime covered by the formal results.

The peer-effects experiment then adds correlated payoff shocks across neighboring players. The correlation parameter is estimated from multiple independent action profiles generated on the same fixed network, using variation across profiles to separate correlated unobserved payoffs from strategic complementarity. The network-formation experiment instead keeps the shocks independent and estimates sender and receiver fixed effects for each player.

The numerical section also reports a small multinomial-probit benchmark. A simulated-likelihood benchmark is feasible in that model, so the comparison measures the efficiency cost of the convex regret criterion and how that cost changes with the number of simulation draws.

2 Related literature

This paper contributes first to the empirical literature on discrete choice over multiple goods. Hendel (1999) estimates multiple-discrete choice in computer adoption, and Gentzkow (2007) estimates demand for print and online newspapers when the two goods may be complements or substitutes. Iaria and Wang (2019) study bundle demand with endogenous prices and bundle-level market shares. Berry et al. (2014) and Fox and Lazzati (2017) discuss the restrictions needed to identify complementarity in discrete-choice models.

The approach here is designed for cases where the number of possible bundles is too large to assign an unrestricted utility to each bundle. I specify package values at the item level: each item has a mean utility, and complementarities enter through a lower-dimensional set of interaction features. This retains the role of product mean utilities in Berry et al. (1995), but estimates those item utilities jointly with the interaction parameters inside a convex program. In the C-Block application, the items are BTA licenses and the counterfactual objects are MTA packages. The estimator therefore recovers primitives at the BTA level and then aggregates them to value packages that were not sold in the observed auction. It does not estimate a separate quality term for each counterfactual MTA license.

The mean-utility step is closest to the matching interpretation of demand inversion in Bonnet et al. (2022). They show that recovering product mean utilities in random-utility demand

can be read as finding equilibrium payoffs in a two-sided matching market. Conditional on the nonlinear shock parameters and the interaction coefficients, there must be weights on simulated bidders' optimal packages such that average item demand clears observed item supply. The dual variables of this package-assignment problem recover license-level net utility shifters. In the full first stage, the same row-generated linear program also estimates the interaction coefficients by matching the observed allocation's feature totals. This connects the estimator to empirical matching models (Galichon and Salanié, 2022) and to competitive-equilibrium formulations for indivisible goods and package assignments (Bikhchandani and Mamer, 1997; Bikhchandani and Ostroy, 2002).

The estimator is also related to simulation-based estimation of discrete-choice models. Like McFadden (1989) and Pakes and Pollard (1989), I simulate unobserved payoff shocks and repeatedly solve the agent's discrete problem. Unlike a conventional simulated method of moments criterion, the criterion here is average regret: it compares the value of observed choice with the indirect utility under simulated shocks. Under the linear-in-parameters payoff restriction, this maximum over alternatives is convex in the parameters. The closest computational comparison is Head et al. (2026), where firms solve large combinatorial production-location problems inside an SMM estimator. Their structural model is nonlinear in parameters. Here the linear part is estimated by constraint generation, while nonlinear shock parameters are handled in an outer loop.

Section 10 returns to the FCC C-Block auction analyzed by Fox and Bajari (2013). They estimate geographic complementarities from pairwise-stability inequalities among winning bidders. The basic comparison swaps one license held by bidder i with one license held by bidder i'. Since the pair of bidders holds the same two licenses before and after the swap, license-level terms enter both sides of the inequality and cancel. This is useful for estimating complementarities without specifying license fixed effects or the distribution of bidder shocks. It also means that the estimator does not recover the license mean utilities or the price coefficient. Those objects are needed for the counterfactual in this paper, because replacing the BTA supply with MTA supply requires solving for new supporting prices. I therefore impose ex-post optimality at observed prices and compare each bidder's observed package with the full feasible set.

Beyond individual bundle demand, the paper relates to econometric models of potential games and social interactions. Potential games are the class in which unilateral payoff changes are represented by changes in a common potential (Monderer and Shapley, 1996). Blume (1993) and Mele (2017) use related potential structures to obtain Gibbs likelihoods for social interaction and network-formation models. This paper instead simulates player-action shocks and selects the action profile that globally maximizes the realized potential. Graham and Pelican (2023) is the closest comparison for large supermodular binary games, but they complete the model by selecting the minimal Nash equilibrium. The correlated-shock peer-effects experiment speaks to the reflection problem in Manski (1993): correlated unobserved payoffs can generate correlated actions that resemble strategic complementarity. The exercise separates these forces only under the maintained graph, shock process, and repeated-profile design of the Monte Carlo.

Moment inequalities provide a different revealed-preference benchmark. In Pakes (2010);

Pakes et al. (2015), choices need only be optimal in expectation conditional on the decision maker's information. A realized choice may fail an ex-post inequality after shocks are realized, while the conditional moment inequality remains valid. The criterion here imposes the stronger restriction that the observed choice is ex-post optimal for the realized payoff shocks. This loses robustness to information frictions, but with linear payoffs the full set of revealed-preference inequalities enters a convex criterion, allowing the estimation of high-dimensional parameters, including fixed effects and complementarities, from exponentially many such inequalities.

3 Setup

Letting $[M] := \{1, \dots, M\}$ denote a finite set of items, I will denote with $d \in \{0,1\}^M$ a generic *bundle* of items. The value

$$V(d; \theta_0, x_i, \nu_i)$$

of a bundle depends on observables x_i, unobservables ν_i and a K-dimensional parameter vector θ_0. The observed bundle d_i is the solution of a combinatorial optimization problem

$$d_i \in \arg \max_{d \in C_i \subseteq \{0,1\}^M} V(d; \theta_0, x_i, \nu_i) \tag{1}$$

where $C_i = C(x_i)$ is an observable choice set. Let $\mathbb{P}$ denote the distribution of (x_i, ν_i, d_i).

Assumption 1 (Hedonic utility)**.** The utility is additively separable

$$V(d; \theta, x_i, \nu_i) \;=\; \theta^\top \phi(d; x_i) + \varepsilon(d; \nu_i) \;=\; \sum_{k \in [K]} \theta_k \, \phi_k(d; x_i) + \varepsilon(d; \nu_i),$$

with $\int \; |\phi_k(d; x_i)| \; d\mathbb{P}(x_i) < \infty$, $\int \; |\varepsilon(d; \nu_i)| \; d\mathbb{P}(\nu_i) < \infty$ for all d and k.

This is the usual hedonic formulation in discrete choice: the value is a sum of characteristics weighted by the parameters (McFadden, 1977).

The focus of this paper is on estimating the *linear* parameters θ, with the component ν_i simulated from a known distribution. I assume that the conditional distribution $\mathbb{P}_{\nu|x}(d\nu_i \mid x_i)$ is known to the econometrician. In practice $\mathbb{P}_{\nu|x}$ might be known up to *non-linear* parameters. I will show in numerical experiments and application how to recover these additional parameters in an outer loop.

Assumption 2 (Simulability)**.** The conditional law $\mathbb{P}_{\nu|x}(\cdot \mid x_i)$ is known, and the econometrician can draw simulation shocks i.i.d. $\nu_{si} \sim \mathbb{P}_{\nu|x}(\cdot \mid x_i)$ for $s \in [S]$, independent across i and conditionally independent of d_i given x_i.

The functions ϕ and ε are specified by the econometrician as a function of observed data and simulated shocks. In most settings illustrated below, the covariates are modular $\phi_k(d; x_i) = \sum_{j \in [M]} d_j r_{ij}$ or quadratic $\phi_k(d; x_i) = \sum_{j,j' \in [M]} d_j d_{j'} w_{ijj'}$, $x_i = (r_i, w_i)$, and the error is modular $\varepsilon(d; \nu_i) = \sum_{j \in [M]} d_j \nu_{ij}$ (see Proposition 2).[1] More generally, the functions ϕ and ε are known to

[1] In this case

$$V(d; \theta, x_i, \nu_i) = (\theta^{mod} r_i + \nu_i)^\top d + \theta^{quad} d^\top W d,$$

the econometrician.

Assumption 3 (Covariate oracle)**.** The econometrician has access to an oracle which given (x, ν, d), returns $(\phi(d; x),\ \varepsilon(d; \nu))$.

The key computational assumption to estimate random utility models by simulation is that given observables, simulated errors and a candidate parameter, the econometrician can solve the combinatorial optimization. I formalize this in the following.

Assumption 4 (Demand oracle)**.** The econometrician has access to an oracle which given (x, ν, θ), returns $d(x, \nu, \theta) \in \arg\max_{d' \in C(x)} \{\phi(d'; x)^\top \theta + \varepsilon(d'; \nu)\}$.

4 Population criterion

Inspired by the empirical analysis of matching markets (Galichon and Salanié, 2022), the population criterion is built from the *regret* at a candidate $\theta \in \mathbb{R}^K$,

$$\rho(\theta, x_i, \nu_i, d_i) \ := \ \max_{d \in C_i} V(d; \theta, x_i, \nu_i) \ - \ V(d_i; \theta, x_i, \nu_i), \tag{2}$$

that is the maximum violation across all revealed preference inequalities at θ.

Regret is non-negative and by (1) it equals zero at the population parameter

$$\rho(\theta, x_i, \nu_i, d_i) \geq 0, \quad \rho(\theta_0, x_i, \nu_i, d_i) = 0,$$

hence,

$$\mathbb{E}\rho(\theta, x_i, \nu_i, d_i) \geq 0, \quad \mathbb{E}\rho(\theta_0, x_i, \nu_i, d_i) = 0.$$

Proposition 1. *The population parameter θ_0 solves the problem*

$$\min_{\theta \in \mathbb{R}^K} \mathbb{E}\ \rho(\theta, x_i, \nu_i, d_i),$$

and the set of solutions is

$$\Theta_0 \ = \ \Big\{\theta \in \mathbb{R}^K : \rho(\theta, x_i, \nu_i, d_i) = 0,\ \mathbb{P}\text{-}a.s.\Big\}.$$

Under Assumption 1, the objective equals

$$Q(\theta) \ = \ \mathbb{E}\Big[\max_{d \in C_i} \{\theta^\top \phi(d; x_i) + \varepsilon(d; \nu_i)\} - \theta^\top \phi(d_i; x_i)\Big] \tag{3}$$

which differs from the expected regret only by the constant $\mathbb{E}\varepsilon(d_i; \nu_i)$. The function Q is convex, and Θ_0 is convex and closed.

Assumption 1 plays two roles. First, additive separable errors allow me to formulate a criterion independent of the unobserved component $\varepsilon(d_i; \nu_i)$ in the observed choice. Second, the linear-in-parameters formulation of the observable utility gives convexity.

with $\theta = (\theta^{mod}, \theta^{quad}) \in \mathbb{R}^{K_1 + K_2}$.

Θ_0 is the set of identifiable parameters given the model $(\mathbb{P}, \theta_0)$. If $\Theta_0 = \{\theta_0\}$ the model is *point identified*. If $\Theta_0 = \{\theta_0\}$ and there are no ties,[2] the first order conditions are

$$\mathbb{E}\Big[\phi(d(x_i, \nu_i, \theta); x_i)\Big] = \mathbb{E}\Big[\phi(d_i; x_i)\Big] \tag{4}$$

where $d(x_i, \nu_i, \theta)$ is the uniquely defined demand given (x_i, ν_i, θ). These are K equations matching the hedonic decomposition of the model to the hedonic decomposition of the observed choice.

Next I provide sufficient conditions for point identification.

Proposition 2. *Suppose*

- $\varepsilon(d; \nu_i) = \sum_{j \in [M]} d_j \nu_{ij}$ *with* $\mathbb{P}_{\nu|x}(\cdot|x_i)$ *supported on all of* $\mathbb{R}^M$,
- *for every* $\theta \neq 0$,

$$\mathbb{P}_x\Big\{\theta^\top \big(\phi(d; x_i) - \phi(d'; x_i)\big) \neq 0 \text{ for some } d, d' \in C_i\Big\} > 0.$$

Then $\Theta_0 = \{\theta_0\}$.

The first condition implies that every bundle in the choice set has positive choice probability. The second is the standard full-rank condition on the covariates, stated without a second-moment assumption.[3]

5 Leading examples

I study two leading classes of models. The first class is demand for bundles under quasi-linear prices. It recovers high-dimensional item mean utilities jointly with complementarities and is the basis of the empirical application in Section 10. The second class is supermodular potential games, which will provide the basis for some numerical experiments in Section 9. Appendix A discusses further applications and limitations of Assumptions 1 and 4.

5.1 Bundle demand with quasi-linear prices

An agent i chooses a bundle $d \in C_i \subseteq \{0,1\}^M$ in a market with quasi-linear prices. The item mean utility is

$$\delta_j \;=\; -\alpha\, p_j \;+\; \xi_j,$$

with price sensitivity α, price p_j, and quality ξ_j unobserved to the econometrician. The value of a bundle d is

$$V(d; \theta, x_i, \nu_i) \;=\; \delta^\top d \;+\; \gamma^\top \psi(d; x_i) \;+\; \varepsilon(d; \nu_i),$$

where ψ collects K_1 interactions between items. This is Assumption 1 with feature map $\phi(d; x_i) = (d,\ \psi(d; x_i))$ and linear parameter $\theta = (\delta, \gamma) \in \mathbb{R}^{M+K_1}$.

[2] That is, $\mathbb{P}\{\theta^\top \phi(d; x_i) + \varepsilon(d; \nu_i) = \theta^\top \phi(d'; x_i) + \varepsilon(d'; \nu_i)\} = 0$ for all distinct $d, d' \in C_i$ and every θ. Q is then differentiable (see appendix).

[3] Under the second-moment condition in Lemma 4, this condition is equivalent to the positive-definite matrix formulation given there.

The parameter $\delta \in \mathbb{R}^M$ carries one mean utility per item. The convex criterion recovers it jointly with the interaction parameter γ, so the dimension $K = M + K_1$ grows with the number of items. These mean utilities absorb item-level unobserved quality, as product mean utilities do in Berry et al. (1995). Here they are recovered inside a single convex program rather than by an inner contraction. Gentzkow (2007) estimates bundle utilities non-parametrically at $M = 3$. The specification above instead restricts interactions to ψ and so scales to large M. Given $\hat{\delta}$, the decomposition $\delta_j = -\alpha p_j + \xi_j$ identifies the price coefficient α in a second stage that instruments for price, with ξ_j the structural residual. The empirical application in Section 10 runs these two stages.

The linear parameters solve the population problem

$$\min_{\delta,\gamma} \; \mathbb{E}\Big[\max_{d \in C_i} \{\delta^\top d + \gamma^\top \psi(d; x_i) + \varepsilon(d; \nu_i)\}\Big] \; - \; \delta^\top \mathbb{E}\big[d_i\big] \; - \; \gamma^\top \mathbb{E}\big[\psi(d_i; x_i)\big].$$

Fixing γ, the residual minimization in δ is the dual of a one-to-many matching with quasi-linear pricing on a continuum of agents (Azevedo et al., 2013b). As such, bundle demand with quasi-linear prices is the one-to-many counterpart of Bonnet et al. (2022), who read BLP inversion as a one-to-one matching and estimate δ with matching algorithms.

5.2 Supermodular potential games

Some discrete games are observed at an action profile that maximizes a common function. For T players and L binary decisions per player, let $d_{t\ell} \in \{0,1\}$ indicate whether player t takes action ℓ. The full action profile $d \in \{0,1\}^M$ then has $M = TL$ coordinates, with player t controlling the block of decisions $d^t = (d_{t\ell})_{\ell \in [L]}$. With payoffs $\pi^t(d^t, d^{-t}; \theta, x, \nu)$, the game is a potential game in the sense of Monderer and Shapley (1996) if there is a function V with

$$\pi^t(d^t, d^{-t}) - \pi^t(\tilde{d}^t, d^{-t}) \; = \; V(d^t, d^{-t}) - V(\tilde{d}^t, d^{-t})$$

for every player t and pair of own-actions $d^t, \tilde{d}^t$. If $d^\star$ maximizes V, then $d^\star$ is a pure-strategy Nash equilibrium (Monderer and Shapley, 1996), so the data-generating process (1) models the observed profile as the potential-maximizing equilibrium, one selection among possibly several. These games are observed as a single large profile, so the framework applies with $N = 1$ and large M.

The leading parametric form is the quadratic potential

$$V(d; \theta, x, \nu) \; = \; \sum_{t \in [T]} \sum_{\ell \in [L]} d_{t\ell} \big(\theta^{\mathrm{mod}\top} r_{t\ell} + \nu_{t\ell}\big) \; + \; \theta^{\mathrm{quad}}\, d^\top W d, \tag{5}$$

where $W = W^\top \in \mathbb{R}^{M \times M}$ has zero diagonal and nonnegative off-diagonal entries, $\theta^{\mathrm{quad}} \geq 0$, and the error $\varepsilon(d; \nu) = \sum_{t,\ell} d_{t\ell}\, \nu_{t\ell}$ is modular. This fits Assumption 1, with modular and quadratic features. The nonnegative interactions make V supermodular, so the Nash set is a complete lattice with a largest and a smallest element (Topkis, 1998), and maximizing V over $\{0,1\}^M$ is a minimum-cut problem solved in $\mathcal{O}(M^3)$, so the demand oracle of Assumption 4 runs in polynomial time. Two examples illustrate this form.

Peer effects give a first example. The same specification can describe binary decisions such as youth smoking, technology adoption, program take-up, or participation in a local activity. In these settings, the payoff from taking the action depends on own covariates and shocks, but also on the choices of neighboring agents. Krauth (2006) estimates this finite-player binary peer-effects model by simulation, building on Brock and Durlauf (2001).

Example 1 (Peer effects, Krauth (2006))**.** There are T players on an undirected network with symmetric adjacency matrix $W \in \mathbb{R}_+^{T\times T}$ and zero diagonal. Player t chooses $d_t \in \{0,1\}$. Payoffs are

$$\pi^t(d_t, d^{-t}; \theta, x, \nu) \;=\; d_t\Big(\beta^\top x_t + \nu_t + \delta \sum_{t'} w_{tt'} d_{t'}\Big), \qquad \delta \geq 0.$$

The associated potential is

$$V(d; \theta, x, \nu) \;=\; \sum_t d_t(\beta^\top x_t + \nu_t) + \frac{\delta}{2} d^\top W d.$$

Positive peer effects are strategic complementarities, so V is supermodular. The observed action profile can therefore be modeled as the equilibrium selected by maximizing this potential. The peer-effects numerical experiment in Section 9.3 instantiates this example on a large fixed network.

The next example is a strategic multi-market entry game in the sense of Jia (2008).

Example 2 (Entry game, Jia (2008))**.** Two chains decide in which of $J = 2065$ locations to open stores. With $y^t \in \{0,1\}^J$ the entry decision of chain $t \in \{A, B\}$, payoffs are

$$\pi^t(y^t, y^{-t}; \theta, x, \nu) \;=\; \sum_j y_j^t(\theta^1 r_j^t + \nu_j^t - \theta^3 y_j^{-t}) \;+\; \theta^2 \sum_{jj'} y_j^t y_{j'}^t w_{jj'},$$

with $w_{jj'} = 1/\text{distance}_{jj'} \geq 0$, $\theta^2 \geq 0$, $\theta^3 \geq 0$, and noise $\nu_j^t = \sigma\eta_{1j} + \sqrt{1-\sigma^2}\,\eta_{2j}^t$, with $\eta_{1j}, \eta_{2j}^A, \eta_{2j}^B$ i.i.d. standard normals across j. I impose the symmetric-congestion restriction $\theta^{3,A} = \theta^{3,B} = \theta^3$ compared to the original paper.[4]

Within-chain proximity is complementary, but cross-chain congestion is not, so the game is not supermodular in (y^A, y^B). Under the change of variables $\bar{y}_j^B := 1 - y_j^B$ the game on $d := (y^A, \bar{y}^B) \in \{0,1\}^{2J}$ ($T = 2$ chains, $L = J$ locations) is supermodular (Milgrom and Roberts, 1990) and is of the form (5), with potential

$$\begin{aligned} V(d; \theta, x, \nu) \;=\; & \sum_j y_j^A(\theta^1 r_j^A + \nu_j^A) \;+\; \sum_j (1 - \bar{y}_j^B)(\theta^1 r_j^B + \nu_j^B) \\ & + \theta^2 \Big[\sum_{jj'} y_j^A y_{j'}^A\, w_{jj'} \;+\; \sum_{jj'} (1 - \bar{y}_j^B)(1 - \bar{y}_{j'}^B)\, w_{jj'} \Big] \\ & - \theta^3 \sum_j y_j^A (1 - \bar{y}_j^B). \end{aligned}$$

[4] In Appendix A.1 I consider asymmetric congestion. Fitting Assumption 1 there requires enlarging the noise specification to a 2×2 covariance with three free non-linear parameters. The estimator then identifies asymmetric congestion and a noise covariance, and Jia's unit-variance assumption becomes a testable restriction rather than an enforced one.

A maximizer $d^\star$ corresponds to the Nash equilibrium $(y^{A,\star}, 1 - \bar{y}^{B,\star})$, computed by minimum cut.

6 Estimation

Let $(x_i, d_i)_{i\in[N]}$ be a sample of i.i.d. observations. By Assumption 2 draw, for each i, i.i.d. shocks $(\nu_{si})_{s\in[S]}$ from $\mathbb{P}_{\nu|x}(\cdot \mid x_i)$. The sample criterion is

$$Q^{N,S}(\theta) \;:=\; \frac{1}{NS} \sum_{i\in[N], s\in[S]} \left[\max_{d\in C_i} \{\theta^\top \phi(d; x_i) + \varepsilon(d; \nu_{si})\} - \theta^\top \phi(d_i; x_i) \right].$$

Lemma 1 (Polyhedrality). *$\theta \mapsto Q^{N,S}(\theta)$ is convex and polyhedral.*

Equivalently, $Q^{N,S}$ is a finite maximum of affine functions, hence continuous and piecewise linear. This polyhedral structure is the basis for the linear programming reformulation and computational methods in Section 7.

For a sequence $c_N \geq 0$ with $c_N \to 0$, the level-set estimator is

$$\hat{\Theta}^{N,S}_{c_N} \;:=\; \{\theta \in \mathbb{R}^K : Q^{N,S}(\theta) \leq \inf_{\vartheta\in\mathbb{R}^K} Q^{N,S}(\vartheta) + c_N\},$$

so that $\hat{\Theta}^{N,S}_0 = \arg\min Q^{N,S}$.

The next proposition establishes consistency of the set of minimizers of the sample criterion. Here d_H denotes the Hausdorff distance between sets.

Proposition 3 (Set-consistency). *Suppose Θ_0 is compact. Then, almost surely,*

$$\sup_{\theta\in\hat{\Theta}^{N,S}_{c_N}} \operatorname{dist}(\theta, \Theta_0) \;\xrightarrow[N\to\infty]{}\; 0.$$

Suppose in addition that $\mathbb{E}\|\phi(d; x_i)\|^2 < \infty$ and $\mathbb{E}|\varepsilon(d; \nu_i)|^2 < \infty$ for every $d \in \{0,1\}^M$. Let (c_N) be a non-negative sequence with $c_N \to 0$ and $c_N\sqrt{N/\log\log N} \to \infty$. Then, almost surely,

$$d_H\big(\hat{\Theta}^{N,S}_{c_N}, \Theta_0\big) \;\xrightarrow[N\to\infty]{}\; 0.$$

Chernozhukov et al. (2007) prove Hausdorff consistency in probability for lower-contour estimators of general criteria on a compact parameter space. Proposition 3 gives an almost-sure result for the convex regret criterion on $\mathbb{R}^K$. Convexity matters because compactness of Θ_0 then implies coercivity, so the Γ-convergence argument does not require the parameter space itself to be compact. The positive tolerance is needed for the inner limit. An exact sample argmin may approach a proper subset of a nonsingleton Θ_0, whereas c_N shrinking at the stated rate recovers all of Θ_0.

When the conditions of Proposition 2 hold, the identified set is a singleton and outer set consistency reduces to consistency of the exact sample minimizer. The next result adds regularity of the conditional shock distribution and second moments of the feature vectors.

Proposition 4 (Asymptotic normality)**.** *Under the conditions of Proposition 2, suppose in addition that $\mathbb{P}_{\nu|x}(\cdot \mid x)$ admits a jointly measurable density with respect to Lebesgue measure that is continuous and strictly positive in ν. Suppose there is a constant $A > 0$ such that, for $\mathbb{P}_x$-almost every x, this density is bounded above by $A(1 + \|\nu\|)^{-M}$ for every $\nu \in \mathbb{R}^M$, and suppose that $\mathbb{E}[\max_{d \in C_i} \|\phi(d; x_i)\|^2] < \infty$. For fixed S, any measurable exact minimizer $\hat{\theta}^{N,S}$ satisfies*

$$\sqrt{N}(\hat{\theta}^{N,S} - \theta_0) \Longrightarrow \mathcal{N}\left(0, H^{-1}\Sigma_S H^{-1}\right),$$

where H is the symmetric positive-definite local curvature matrix of Q and Σ_S is the score covariance.[5]

The proof applies Hjort and Pollard (1993, Theorem 2.1) after verifying its local expansion and remainder conditions for the simulated regret criterion. Their convex argmin result applies even though the sample criterion is polyhedral and generally nonsmooth. As in Pakes and Pollard (1989), Σ_{data} reflects sampling variation across observations, while $S^{-1}\Sigma_{\text{sim}}$ is the additional variation from using finitely many simulated shocks.

Bayesian bootstrap. I compute standard errors using Bayesian bootstrap (Rubin, 1981). Independently of the data, draw $(w_i)_{i \in [N]} \sim N \cdot \text{Dirichlet}(\mathbf{1}_N)$ and re-solve

$$\hat{\theta}^{N,S,*} \in \arg\min_{\theta} \frac{1}{NS} \sum_{i \in [N], s \in [S]} w_i \left[\max_{d \in C_i} \{\theta^\top \phi(d; x_i) + \varepsilon(d; \nu_{si})\} - \theta^\top \phi(d_i; x_i)\right].$$

This avoids estimating the Hessian of the population criterion by finite differences, which is costly when the number of parameters is large. Bayesian bootstrap also avoids bootstrap samples in which the estimator does not exist. If only a few observations choose an item, an ordinary bootstrap sample may omit all of them, in which case the corresponding item fixed effect has no finite estimate. Bayesian-bootstrap weights are almost surely positive, so this problem does not arise.

7 Computation

In this section, I present the computational methods to solve the sample problem.

7.1 Problem and LP formulation

The sample problem from Section 6 is

$$\min_{\theta \in \mathbb{R}^K} \ \frac{1}{SN} \sum_{s \in [S], i \in [N]} \Big[\max_{d \in C(x_i)} \{\phi(d; x_i)^\top \theta + \varepsilon(d; \nu_{si})\} \ - \ \phi(d_i; x_i)^\top \theta\Big]. \tag{6}$$

Hence, the estimation problem can be stated as follows.

[5] It decomposes as $\Sigma_S = \Sigma_{\text{data}} + S^{-1}\Sigma_{\text{sim}}$, where $\Sigma_{\text{data}} := \text{Var}(\mathbb{E}[\nabla_\theta \rho(\theta_0, x_i, \nu, d_i) \mid x_i, d_i])$ and $\Sigma_{\text{sim}} := \mathbb{E}[\text{Var}(\nabla_\theta \rho(\theta_0, x_i, \nu, d_i) \mid x_i, d_i)]$.

Econometrician's problem

Input: sample $(x_i, d_i)_{i\in[N]}$, simulated draws $(\nu_{si})_{s\in[S],\, i\in[N]}$, covariate and demand oracles.
Output: a minimizer $\hat{\theta}$ of (6).

Problem (6) admits a linear-programming formulation. Writing $\hat{m} := \frac{1}{N}\sum_{i\in[N]} \phi(d_i; x_i)$ for the empirical mean of the observed features, the sample problem can be rewritten as

$$\begin{aligned} \min_{\theta\in\mathbb{R}^K,\, u\in\mathbb{R}^{SN}} \quad & \tfrac{1}{SN}\sum_{s,i} u_{si} \;-\; \hat{m}^\top\theta \\ \text{s.t.} \quad & u_{si} \;\geq\; \phi(d; x_i)^\top\theta + \varepsilon(d; \nu_{si}), \quad \forall s, i, \;\; d \in C(x_i), \end{aligned} \tag{7}$$

with dual problem

$$\begin{aligned} \max_{\mu\geq 0} \quad & \tfrac{1}{SN}\sum_{s,i}\sum_{d\in C(x_i)} \mu_{si}(d)\, \varepsilon(d; \nu_{si}) \\ \text{s.t.} \quad & \sum_{d\in C(x_i)} \mu_{si}(d) \;=\; 1, \quad \forall s, i, \\ & \tfrac{1}{SN}\sum_{s,i}\sum_{d\in C(x_i)} \mu_{si}(d)\, \phi(d; x_i) \;=\; \hat{m}. \end{aligned} \tag{8}$$

The primal and dual solutions are related by complementary slackness theorem.

Proposition 5 (LP optimality). *$(\hat{u}, \hat{\theta}, \hat{\mu})$ is a primal-dual solution of (7)–(8) if and only if, for every $s \in [S], i \in [N]$ and $d \in C(x_i)$,*

$$\begin{aligned} & \hat{u}_{si} \;\geq\; \phi(d; x_i)^\top\hat{\theta} + \varepsilon(d; \nu_{si}), \quad \textit{with equality if } \hat{\mu}_{si}(d) > 0, \\ & \hat{\mu}_{si}(d) \;\geq\; 0, \qquad \sum_{d'\in C(x_i)} \hat{\mu}_{si}(d') \;=\; 1, \\ & \tfrac{1}{SN}\sum_{s,i}\sum_{d'\in C(x_i)} \hat{\mu}_{si}(d')\, \phi(d'; x_i) \;=\; \hat{m}. \end{aligned}$$

The primal-dual formulation links estimation to competitive equilibrium. The dual constraints associated with item fixed effects require simulated demand for each item to equal its observed allocation. These are market-clearing conditions, so (8) is an assignment LP for an indivisible-goods market, with additional constraints matching the remaining features (Bikhchandani and Ostroy, 2002).

The dual allows μ_{si} to place positive weight on more than one bundle, which corresponds to a fractional allocation in the assignment interpretation. In the competitive-equilibrium literature, a pure equilibrium corresponds to an integral optimum of the assignment LP (Bikhchandani and Mamer, 1997), while replication makes the convexified and indivisible economies asymptotically equivalent (Starr, 1969). Proposition 6 gives a direct bound for the estimator: there is an optimal dual solution in which at most K of the SN simulated bidders are split across bundles. The fraction split therefore converges to zero as S increases.

7.2 Complexity

The demand and feature oracles allow evaluation of the objective and a subgradient at a candidate parameter.

Lemma 2 (Subgradient). *Let*

$$g \;=\; \tfrac{1}{SN}\sum_{s,i}\phi(d_{si};x_i) \;-\; \hat{m},$$

where $d_{si} \in \arg\max_{d\in C(x_i)}\{\phi(d;x_i)^\top\theta + \varepsilon(d;\nu_{si})\}$. *Then* $g \in \partial Q^{N,S}(\theta)$.

Let $\bar{\Theta} \subset \mathbb{R}^K$ be a known compact set containing a minimizer, and call $\hat{\theta} \in \bar{\Theta}$ τ-optimal if

$$Q^{N,S}(\hat{\theta}) - \inf_{\theta\in\bar{\Theta}} Q^{N,S}(\theta) \le \tau.$$

The ellipsoid method finds a τ-optimal point in $\mathcal{O}(K^2\log(1/\tau))$ full-sample subgradient evaluations (Nesterov, 2013), where the bounded constant depends on the objective's Lipschitz constant and the diameter of $\bar{\Theta}$. Each evaluation solves SN individual demand problems and evaluates the selected features, giving $\mathcal{O}(K^2SN\log(1/\tau))$ individual oracle calls.

Theorem 1 in Appendix B.2 states the bound with its explicit dependence on the geometry of $\bar{\Theta}$ and the Lipschitz constant of the sample objective.

These complexity bounds set a worst-case benchmark. In practice, the constraint-generation algorithms presented below outperform the ellipsoid method in numerical experiments, consistently finding a solution in less than K^2 iterations.

While this bounds the number of outer oracle queries, the overall time complexity depends on the demand problem. When demand reduces to a polynomial-time problem such as minimum cut (Section 5), the entire estimation problem is polynomial-time. For NP-hard problems such as the capacity-constrained quadratic knapsack in the FCC application, each oracle call can be formulated as a mixed-integer program and solved with a standard MIP solver. Constraint generation still terminates after finitely many iterations, but worst-case runtime remains exponential.

7.3 Constraint-generation

The LP (7) has $K + SN$ variables and exponentially many constraints. Constraint-generation avoids forming the full constraint set. It solves a restricted master problem and uses the demand oracle to find violated constraints at the current parameter. Since each $C(x_i)$ is finite, every non-terminating iteration adds at least one new row, so the algorithm terminates after finitely many row additions. In practice, the generated active set is much smaller than the full constraint set.

Constraint-generation begins with a working set $\mathcal{C}_{si} \subseteq C(x_i)$ for each $s \in [S]$, $i \in [N]$, solves the restricted master problem

$$\begin{aligned}\min_{\theta\in\mathbb{R}^K,\,u\in\mathbb{R}^{SN}} \quad & \tfrac{1}{SN}\sum_{s,i} u_{si} \;-\; \hat{m}^\top\theta \\ \text{s.t.} \quad & u_{si} \;\ge\; \phi(d;x_i)^\top\theta + \varepsilon(d;\nu_{si}), \quad \forall s,i,\;\; d\in\mathcal{C}_{si},\end{aligned} \tag{9}$$

for (u, θ) and grows $\mathcal{C}_{si}$ by querying the oracles at the current θ. The pseudocode is shown in Algorithm 1.

Algorithm 1 Constraint-generation

Initialize $\mathcal{C}_{si} \leftarrow \{d_i\}$ for each s, i.
repeat
  Solve (9) for (θ, u).
  for each s, i **do**
    Demand oracle at (x_i, ν_{si}, θ) returns d^*_{si}.
    Covariate oracle at $(x_i, \nu_{si}, d^*_{si})$ returns $(\phi(d^*_{si}; x_i),\ \varepsilon(d^*_{si}; \nu_{si}))$.
    $\bar{u}_{si} \leftarrow \phi(d^*_{si}; x_i)^\top \theta + \varepsilon(d^*_{si}; \nu_{si})$.
  end for
  $\mathcal{C}_{si} \leftarrow \mathcal{C}_{si} \cup \{d^*_{si}\}$ for every s, i with $\bar{u}_{si} > u_{si}$.
until $\bar{u}_{si} \leq u_{si}$ for all s, i.
return (θ, u).

Each non-terminating iteration strictly grows some $\mathcal{C}_{si} \subseteq C(x_i)$, so the algorithm terminates in finitely many iterations. At termination (θ, u), together with the restricted-master dual variables, satisfies the conditions of Proposition 5.

1-slack formulation. The same criterion also has a 1-slack formulation that aggregates the constraints across simulated observations:

$$\begin{aligned} \min_{\theta \in \mathbb{R}^K,\, \bar{u} \in \mathbb{R}} \quad & \bar{u} \;-\; \hat{m}^\top \theta \\ \text{s.t.} \quad & \bar{u} \;\geq\; \tfrac{1}{SN} \sum_{s,i} \Big[\phi(d_{si}; x_i)^\top \theta + \varepsilon(d_{si}; \nu_{si}) \Big], \quad \forall (d_{si})_{s,i} \in \prod_{s \in [S],\, i \in [N]} C(x_i). \end{aligned} \tag{10}$$

Lemma 3 (Equivalence of the slack formulations)**.** *Let $S, N \geq 1$. Problems (7) and (10) have the same optimal value. Moreover, a parameter θ is part of an optimal solution to (7) if and only if it is part of an optimal solution to (10).*

In row generation for (10), the working set contains tuples $(d_{si})_{s,i}$ rather than separate constraints for each simulated observation, hence has more constraints than the formulation above. Given a candidate θ, the new constraints are obtained by solving the same demand problem as in Algorithm 1 for each (s, i). The restricted master in (10) has $K + 1$ variables $(\theta, \bar{u})$ compared to $K + SN$ in the formulation above. In general, the 1-slack formulation requires more oracle queries compared to the baseline version presented in the previous paragraph. Therefore, 1-slack formulation is beneficial for problems with a large number of observations, and a fast demand oracle.

Consider the unit-demand case with $|C(x_i)| = M$ for all i. The complete 1-slack LP has M^{SN} rows, one for each tuple of choices across all simulated observations. The constraint set is therefore exponential even though each component problem is a maximum over only M items. Yet, the restricted 1-slack master has $K + 1$ variables, while the restricted version of (7) has

$K + SN$ variables. Increasing N or S increases the cost of generating a row, because separation requires SN demand calls, but it does not add master variables.

Remark 1. The formulations (7) and (10) presented in this section are two extreme case of the multiple LP reformulations of the sample problem (6). One can easily envision intermediate formulations where the simulation-observation indexes are aggregated at different levels. For instance, if the agents $i \in [N]$ are distributed across T markets, it could be convenient to add a slack variable u_t for each market, rather than for each observation-simulation.

8 Comparison with other methodologies

The comparisons below locate this regret formulation relative to moment inequalities, score estimators, SMM, and likelihood methods for potential games.

8.1 Moment inequalities

To isolate the comparison, set $\varepsilon \equiv 0$ and impose the scale normalization $\theta_1 = \theta_{0,1} = 1$. Utility is $V(d; \theta_0, x_i) = \theta_0^\top \phi(d; x_i)$. The full-information data-generating process (1) implies the pointwise revealed-preference restrictions

$$\theta_0^\top \big(\phi(d_i; x_i) - \phi(d; x_i)\big) \;\geq\; 0 \qquad \forall d \in C_i,\ \mathbb{P}\text{-a.s.} \tag{11}$$

Moment-inequality models (Pakes, 2010; Pakes et al., 2015) distinguish the agent's information $\mathcal{W}_i$ from the econometrician's observables. If the agent maximizes expected utility given $\mathcal{W}_i$, the realized ex-post inequality (11) can fail. Conditional optimality instead implies

$$\mathbb{E}\big[\theta_0^\top (\phi(d_i; x_i) - \phi(d; x_i)) \,\big|\, \mathcal{W}_i\big] \;\geq\; 0 \qquad \forall d \in C_i,\ \mathbb{P}\text{-a.s.}$$

For each $d \in C_i$, let $h_i(d)$ be a nonnegative bounded $\mathcal{W}_i$-measurable weight observed by the econometrician. Then

$$\theta_0^\top\, \mathbb{E}\Big[\sum_{d \in C_i} h_i(d)\big(\phi(d_i; x_i) - \phi(d; x_i)\big)\Big] \;\geq\; 0.$$

Given a class $\mathcal{H}$ of such weights, the associated identified set is

$$\Theta^{MI}(\mathcal{H}) = \Big\{\theta :\ \theta^\top \mathbb{E}\Big[\sum_{d \in C_i} h_i(d)\big(\phi(d_i; x_i) - \phi(d; x_i)\big)\Big] \geq 0 \ \ \forall h_i \in \mathcal{H}\Big\}.$$

When the full-information restrictions (11) hold, $\Theta_0 \subseteq \Theta^{MI}(\mathcal{H})$ for any admissible $\mathcal{H}$, with strict inclusion in general. When the agent's information is coarser, (11) need not hold, $Q(\theta_0)$ need not equal zero, and the regret estimator targets the wrong object. Moment inequalities remain valid by conditional optimality. The tradeoff is robustness to expectational errors in exchange for weaker identifying restrictions.

8.2 Score and rank-order estimators

The maximum-score literature (Manski, 1975; Fox, 2007) does not impose realized ex-post optimality. As above, normalize $\theta_1 = 1$. Identification instead uses the rank-order restriction

$$\theta_0^\top \phi(d; x_i) > \theta_0^\top \phi(d'; x_i) \quad \iff \quad \mathbb{P}(d_i = d \mid x_i, C_i) > \mathbb{P}(d_i = d' \mid x_i, C_i).$$

Thus alternatives with higher deterministic utility are chosen more often. This restriction does not require a parametric distribution for the choice-specific errors.

Fox (2007) shows that the rank ordering is preserved when estimation uses a subset $\bar{C}_i \subseteq C_i$. Using a subset avoids constructing or comparing every alternative, but discards the remaining comparisons. The regret criterion instead considers all inequalities, sharpening identification, and imposes a parametric assumption on the shocks.

Maximum-score criteria are discontinuous, which makes them challenging to optimize in high dimensions. The regret criterion is convex under Assumption 1, and its sample problem is the LP (7).

In the FCC application, Fox and Bajari (2013) use maximum rank correlation under pairwise stability in matches, comparing license swaps between winning bidders. Section 10.6 gives the application-specific comparison with ex-post bundle optimality.

8.3 Simulated method of moments

Under Assumption 1, point identification, and no ties, the first-order conditions of Q are the moment equalities

$$\mathbb{E}\left[\phi(d(x_i, \nu_i, \theta); x_i) - \phi(d_i; x_i)\right] = 0$$

Using standard simulated method of moments (McFadden, 1989; Pakes and Pollard, 1989), the moment conditions would be used in a sample criterion

$$\min_\theta \ \bar{g}(\theta)^\top W \bar{g}(\theta),$$

where $\bar{g}(\theta)$ may include the feature discrepancy above and other moments, and W is positive definite. Evaluating the criterion requires the same oracle assumptions as the regret estimator, in particular the ability to solve the combinatorial optimization problem for each simulated observation. With discrete simulated choices, the SMM criterion is discontinuous and need not be convex in θ. The regret criterion instead aggregates revealed-preference violations through a maximum, and is convex under Assumption 1.

The simulated method of moments can include moments for non-linear parameters. The methods in this paper focus on linear parameters and handle non-linear parameters in a (non-linear) outer loop (Sections 9.3 and 10.2).

8.4 Potential games and likelihood

The entry game of Example 2 and the network game of Section 9.3 are potential games. The data-generating process (1) assumes that the observed action profile d_i globally maximizes the

potential V at the realized ν_i. By Monderer and Shapley (1996), such a maximizer is a Nash equilibrium, although the game may have other Nash equilibria.

A different literature (Blume, 1993; Mele, 2017) models potential-game data with the Gibbs distribution

$$\mathbb{P}_\theta(d \mid x) \propto \exp\big(V(d;\theta,x)\big),$$

micro-founded as the stationary distribution of stochastic logit best-response dynamics. The same choice probabilities can be represented by a static random-utility model (1) with i.i.d. Type-I extreme-value shocks at the bundle level. Likelihood-based estimation then requires the log-partition function $\log \sum_d \exp(V(d;\theta,x))$, or an approximation to it.

For large M, this paper departs from bundle-level i.i.d. Gumbel shocks and works mostly with modular item-level errors. The two approaches therefore correspond to different data-generating processes: likelihood methods exploit the Gibbs structure, while the regret estimator exploits the ability to solve the realized optimization problem.

9 Numerical Experiments

This section studies four simulation designs. The first is a small multinomial probit in which a simulated-likelihood benchmark is feasible, so it measures the efficiency cost of replacing a likelihood-based criterion with the convex regret criterion. The second is bundle demand with quasi-linear endogenous prices. It estimates market-item mean utilities and complementarity parameters in the first stage, then uses the estimated mean utilities in a second-stage price regression. The final two designs use supermodular potential games.

9.1 Small-scale multinomial probit benchmark

Each agent chooses among M items. Item j has covariates $x_{ij} \in \mathbb{R}^K$, drawn independently across agents and items from $\mathcal{N}(0, \Sigma_x)$ with equicorrelation 0.5. The shocks ν_{ij} are iid standard normal. Agent i chooses the item that solves

$$\max_{j \in [M]} \{\theta_0^\top x_{ij} + \nu_{ij}\},$$

I set $K = 3$ and $\theta_0 = (1, -0.5, 0.5)$.

I estimate θ_0 with the simulation-based estimator $\hat{\theta}^{N,S}$ of Section 6, using S simulated shock draws per agent. The benchmark is simulated maximum likelihood for the multinomial probit model, with choice probabilities evaluated by GHK simulation using 1000 draws (Hajivassiliou et al., 1996). Because the likelihood benchmark uses many more simulation draws than the regret estimator in the reported designs, I treat it as a close numerical approximation to the exact multinomial-probit MLE.

Across $R = 500$ Monte Carlo replications, I compute the coordinatewise standard deviation of each estimator. Table 1 reports the ratio of the standard deviation of $\hat{\theta}^{N,S}$ to the standard deviation of the simulated-likelihood benchmark, averaged across coordinates. I vary the number of items $M \in \{2, 10\}$, the sample size $N \in \{100, 500, 1000\}$, and the simulation count

$S \in \{1, 10, 50, 100\}$.

	$S = 1$	$S = 10$	$S = 50$	$S = 100$
Binary choice ($M = 2$)				
$N = 100$	1.35	1.05	1.00	0.99
$N = 500$	1.40	1.08	1.05	1.04
$N = 1000$	1.44	1.10	1.03	1.02
Multinomial ($M = 10$)				
$N = 100$	1.48	1.08	1.02	1.02
$N = 500$	1.49	1.09	1.04	1.03
$N = 1000$	1.40	1.08	1.04	1.03

Table 1: Standard error of $\hat{\theta}^{N,S}$ relative to simulated maximum likelihood.

At $S = 1$, the standard error is inflated by 35–50%. By $S = 10$, the inflation falls to 5–10%. By $S \geq 50$, it is roughly 2–5%. The ratio is nearly flat in N at a fixed S, as predicted by the variance decomposition $\Sigma_S = \Sigma_{\text{data}} + S^{-1}\Sigma_{\text{sim}}$ in Section 6.

Figure 1 shows the same comparison over S, with $M = 10$ and $N = 1000$. The plotted ratio, $\sqrt{V_{CB}/V_{MLE}}$, compares the Monte Carlo variance of $\hat{\theta}^{N,S}$ with that of the simulated-likelihood benchmark. In this small design, the convex criterion is nearly as efficient as simulated likelihood once the regret estimator's simulation error is under control.

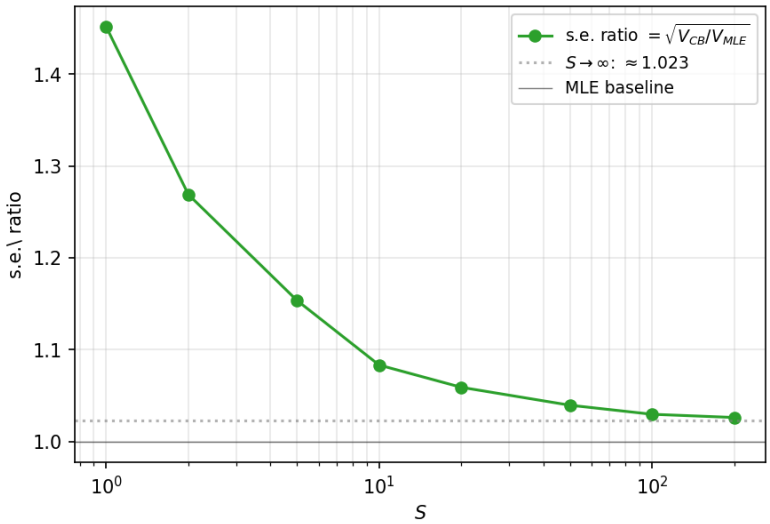


Figure 1: Standard-error ratio of $\hat{\theta}^{N,S}$ against simulated maximum likelihood at $N = 1000$, $M = 10$, across $R = 500$ replications.

9.2 Bundle demand with quasi-linear endogenous prices

There are $T = 15$ markets indexed by t, with $N_t = 100$ agents per market. Each agent i belongs to market $t(i)$ and chooses a bundle $d_i \in \{0,1\}^M$ from $M = 50$ items subject to a capacity

constraint. The value of a bundle is

$$V(d_i;\theta,x_i,\nu_i)=\sum_{j=1}^{M} d_{ij}\left(x_{ij}^{\top}\beta+\delta_{t(i)j}+\nu_{ij}\right)+\sum_{k=1}^{2}\lambda_k\, d_i^{\top}Q^k d_i.$$

The covariates $x_{ij}\in\mathbb{R}^2$ are drawn i.i.d. standard normal. The interaction matrices Q^k are sparse symmetric matrices drawn once and held fixed across Monte Carlo replications. The feasible set is

$$C_i=\{d\in\{0,1\}^M : w^{\top}d\le c_i\}$$

where $w_j=1$ for all j and $c_i=10$ for all i. The linear parameter is $\theta=(\beta,\delta,\lambda)$, where δ contains one mean utility for each market–item pair. The demand oracle solves the resulting capacity-constrained quadratic knapsack problem.

Prices are endogenous through a market-item unobservable. I generate

$$p_{tj}=\gamma_z z_{tj}+\xi_{tj}+\zeta^{p}_{tj},$$
$$\delta_{tj}=-\alpha p_{tj}+\xi_{tj}+\zeta^{\delta}_{tj},$$

where z_{tj} is an excluded cost shifter drawn from a normal distribution. The shocks ξ_{tj}, ζ^{p}_{tj}, ζ^{δ}_{tj}, and ν_{ij} are normal random variables. Across $R=500$ Monte Carlo replications, I hold the market design, instruments, interaction matrices, and structural coefficients fixed, and redraw ξ, ζ^{p}, ζ^{δ}, the demand shocks, and the simulation shocks. The first stage estimates (β,δ,λ) with $S=1$ simulation draw. The second stage regresses $\hat{\delta}_{tj}$ on price with a constant, instrumenting price with z_{tj}.

Parameter	True	Mean	Bias	Std
β_1	+0.6000	+0.6065	+0.0065	+0.0200
β_2	-0.3000	-0.3023	-0.0023	+0.0209
λ_1	+0.6000	+0.6087	+0.0087	+0.0114
λ_2	+0.3000	+0.3041	+0.0041	+0.0098
α (OLS)	+1.0000	+0.3949	-0.6051	+0.1403
α (2SLS)	+1.0000	+1.0415	+0.0415	+0.1893

Table 2: Monte Carlo distribution in the bundle-demand design.

On average, row generation terminated after 30 iterations, with a runtime of 354 seconds per replication.

9.3 Peer effects on a large network

There are $T=500$ players, each choosing $d_t\in\{0,1\}$. The graph $W\in\{0,1\}^{T\times T}$ is the random geometric graph of Graham (2016), and player t has covariates $x_t\in\mathbb{R}^4$. Player t's payoff at a profile is

$$\pi^t(d^t,d^{-t};\theta,x_i,\nu_i)\;=\;d_t\Big(\beta^{\top}x_t+\nu_{it}+\delta\textstyle\sum_{t'}w_{tt'}d_{t'}\Big),\qquad \delta\ge 0,$$

and the associated potential function is

$$V(d;\theta,x_i,\nu_i) \;=\; \sum_t d_t\big(\beta^\top x_t + \nu_{it}\big) \;+\; \frac{\delta}{2}\, d^\top W d.$$

This instantiates Example 1. The linear parameter is $\theta = (\beta,\delta) \in \mathbb{R}^5$. I draw (x, W) once and hold them fixed across replications.

9.3.1 i.i.d. payoff shocks

I first use iid payoff shocks. I set $N = 1$, so that the observation consists of a unique large action profile.

The population parameters are $\theta_0 = (-1, -0.5, -1, 0.5, 0.2)$, and draw $\nu_{it} \overset{iid}{\sim} \mathcal{N}(0,1)$. In each replication, I draw the shocks, generate the observed profile as the potential maximizer at θ_0, and re-estimate θ_0 with $S = 50$ simulation draws. I run $R = 250$ Monte Carlo replications.

Parameter	True	Bias	Std
β_1	-1.00	-0.019	0.140
β_2	-0.50	$+0.003$	0.126
β_3	-1.00	$+0.000$	0.202
β_4	$+0.50$	$+0.010$	0.198
δ	$+0.20$	$+0.000$	0.021

Table 3: Monte Carlo distribution of $\hat{\theta}^{N,S}$, i.i.d. errors.

Table 3 shows small bias for all five parameters. The strategic-complementarity parameter δ is especially precise, with standard deviation 0.021. The exercise shows that the estimator can recover a peer-effect parameter from a single 500-player profile when the shock distribution is correctly specified.

9.3.2 Correlated errors

The next design allows payoff shocks to be correlated across neighboring players and estimates the correlation parameter via an outer loop.

The payoff shock is

$$\nu_i \;=\; (I + \sigma\widetilde{W})\,\eta_i, \qquad \eta_i \sim \mathcal{N}(0, I_T),$$

so $\nu_i \sim \mathcal{N}(0,\, (I+\sigma\widetilde{W})^2)$. Here $\widetilde{W} = \widetilde{D}^{-1/2} W \widetilde{D}^{-1/2}$ is the symmetric degree normalization of W, with $\widetilde{D}_{tt} = \max(\deg_t, 1)$, and σ is a non-linear parameter.

I assume that $N = 100$, so the econometrician observes independent action profiles on the same fixed network. For profile i and player t, I add the observed modular payoff shifter αz_{it}, where $z_{it} \sim$ Bernoulli$(1/2)$ independently across i and t. These two changes are crucial to identify the non-linear parameter. Indeed, variation in modular observable covariates *across* observations is one of the sources to separately identify δ and σ, as discussed in Gentzkow (2007).

Therefore, the linear parameters are $\theta = (\alpha, \beta, \delta) \in \mathbb{R}^6$. The true values are $\sigma_0 = 0.50$ and $\theta_0 = (0.5, -1, -0.5, -1, 0.5, 0.2)$.

For a candidate σ, I run constraint generation with simulated draws $\nu_{si}^{\sigma} = (I + \sigma \widetilde{W})\, \eta_{si}$ and $\eta_{si} \overset{iid}{\sim} \mathcal{N}(0, I_T)$, obtaining the primal–dual solution $(\hat{u}^{\sigma}, \hat{\theta}^{\sigma}, \hat{\mu}^{\sigma})$ of (7)–(8). I then choose σ to match the squared number of active players, $g(d) = (\sum_t d_t)^2$:

$$\hat{\sigma} \;\in\; \arg\min_{\sigma} \; \left(\frac{\bar{g}^{\sigma} - \hat{g}}{|\hat{g}|} \right)^2,$$

where $\bar{g}^{\sigma} := \frac{1}{SN} \sum_{s,i,d} \hat{\mu}_{si}^{\sigma}(d)\, g(d)$ and $\hat{g} := \frac{1}{N} \sum_i g(d_i)$. I use $S = 10$ simulation draws and run $R = 200$ Monte Carlo replications.

Parameter	True	Bias	Std
σ	$+0.50$	$+0.001$	0.121
α	$+0.50$	$+0.000$	0.030
β_1	-1.00	-0.001	0.019
β_2	-0.50	-0.001	0.015
β_3	-1.00	-0.000	0.026
β_4	$+0.50$	$+0.000$	0.031
δ	$+0.20$	$+0.000$	0.004

Table 4: Monte Carlo distribution of $(\hat{\theta}^{\hat{\sigma}}, \hat{\sigma})$, correlated errors, $R = 200$ replications, $N = 100$, $S = 10$.

Table 4 separates correlated unobservables from the strategic interaction. The bias in $\hat{\sigma}$ is 0.001, and the estimates of the six linear parameters have negligible bias. The standard deviation of $\hat{\delta}$ falls to 0.004 because this design observes 100 independent profiles rather than one.

9.4 Network formation with two-sided heterogeneity

The final design lets the number of linear parameters grow with the size of the network. There are T agents. Items $j \in [M]$, with $M = T(T-1)$, index ordered pairs (t, t') with $t \neq t'$. The bundle $d \in \{0,1\}^M$ is the directed adjacency matrix, with $d_{tt'} = 1$ when t links to t'. Player t controls the out-links $d^t = (d_{tt'})_{t' \neq t}$ and receives payoff

$$\pi^t(d^t, d^{-t}; \theta, x_i, \nu_i) \;=\; \sum_{t' \neq t} d_{tt'} \big(\alpha_t^{\mathrm{snd}} + \alpha_{t'}^{\mathrm{rcv}} + \beta^{\top} r_{tt'} + \nu_{i,tt'} + \gamma\, d_{t't} \big), \qquad \gamma \geq 0,$$

which combines a sender effect α_t^{snd}, a receiver effect $\alpha_{t'}^{\mathrm{rcv}}$, dyad covariates $\beta^{\top} r_{tt'}$ with $r_{tt'} \in \mathbb{R}^5$, an idiosyncratic shock $\nu_{i,tt'}$, and a reciprocity bonus γ on each returned link. The potential is

$$V(d; \theta, x_i, \nu_i) \;=\; \sum_{t \neq t'} d_{tt'} \big(\alpha_t^{\mathrm{snd}} + \alpha_{t'}^{\mathrm{rcv}} + \beta^{\top} r_{tt'} + \nu_{i,tt'} \big) \;+\; \gamma \sum_{t < t'} d_{tt'}\, d_{t't}.$$

The linear parameter $\theta = (\alpha^{\text{snd}}, \alpha^{\text{rcv}}, \beta, \gamma) \in \mathbb{R}^{2T+5}$, normalized by $\alpha_1^{\text{snd}} = 0$, contains a sender and receiver fixed effect for every agent. Its dimension is therefore $K = 105$ at $T = 50$ and $K = 205$ at $T = 100$.

The dyad covariates are built from per-agent attributes,

$$r_{tt'} = \big(\mathbb{1}\{g_t^1 = g_{t'}^1\},\ \mathbb{1}\{g_t^2 = g_{t'}^2\},\ -|a_t - a_{t'}|,\ -|b_t - b_{t'}|,\ z_t z_{t'}\big),$$

with $g^1, g^2 \sim \text{Bernoulli}(0.5)$, $a \sim U[0,1]$, and $b, z \sim \mathcal{N}(0,1)$. I draw the attributes and fixed effects $\alpha_t^{\text{snd}}, \alpha_t^{\text{rcv}} \overset{iid}{\sim} \mathcal{N}(-0.8, 0.5^2)$ once. In each Monte Carlo replication, I redraw $\nu_i \overset{iid}{\sim} \mathcal{N}(0, I_M)$, generate the potential-maximizing network at the true parameter, and re-estimate all $2T + 5$ linear parameters.

I set $N = 1$, $\beta_0 = (0.8, 0.5, 1.0, 0.3, 0.4)$, $\gamma_0 = 0.5$, and $S = 100$. I run $R = 250$ replications at each $T \in \{50, 100\}$. Table 5 reports the Monte Carlo distribution of $\hat{\beta}$ and $\hat{\gamma}$. The sender and receiver fixed effects are included in the estimation but omitted from the table.

		$T = 50$		$T = 100$	
Parameter	True	Bias	Std	Bias	Std
β_1	$+0.80$	$+0.083$	0.096	$+0.030$	0.047
β_2	$+0.50$	$+0.044$	0.095	$+0.016$	0.040
β_3	$+1.00$	$+0.065$	0.280	$+0.042$	0.119
β_4	$+0.30$	$+0.033$	0.069	$+0.009$	0.032
β_5	$+0.40$	$+0.039$	0.082	$+0.017$	0.022
γ	$+0.50$	$+0.040$	0.106	$+0.021$	0.047

Table 5: Monte Carlo distribution of $(\hat{\beta}, \hat{\gamma})$, $R = 250$ replications, $S = 100$. The parameter count is $K = 105$ at $T = 50$ and $K = 205$ at $T = 100$. Realized network density is 13% and 15%.

All 250 replications converge at both network sizes. Bias is small at $T = 50$ and smaller at $T = 100$. For the reciprocity parameter, bias falls from 0.040 to 0.021 and the standard deviation falls from 0.106 to 0.047. The covariate coefficients show the same pattern. At $T = 100$, the estimator recovers the main structural parameters while also estimating 200 sender and receiver fixed effects from a single directed network with 9,900 possible links.

10 Empirical Application: the 1996 FCC C-Block

The 1996 FCC C-Block sold 480 BTA-level licenses of 30 MHz each across the continental United States, in a simultaneous-ascending auction held between December 1995 and May 1996 with no package bidding.

Participation was restricted to small-business "designated entities" (annual revenues below $40 million), with large incumbent carriers ineligible. There were 255 registered bidders and 85 winners in the continental U.S. Bidders posted population-denominated eligibility caps before the auction, with caps reduced over the auction under an activity rule and eight winners ex post at the cap. Winning bids totaled $9.87 billion in the 480-license continental-U.S.

analysis sample. Including the 13 non-continental licenses excluded below, C-Block winning bids totaled $10.1 billion. Designated entities received a 25% bidding credit and a 10-year installment-payment plan, equivalent under standard cost-of-capital assumptions to a further effective subsidy. Many winners subsequently defaulted or sold their licenses through bankruptcy or merger (Fox and Bajari, 2013; Cramton, 1997).

An alternative auction would have packaged the same spectrum into 46 MTA-level licenses, a coarser partition compared to the chosen BTA partition, like the previous AB-Block auction. The goal of this empirical exercise is quantifying the counterfactual revenue and welfare in the counterfactual where licenses are sold at the MTA level, keeping the distribution of bidders fixed.

The revenue effect of pre-packaging multiple objects into a single license is studied in the auction-theory literature for private values and second-price (Palfrey, 1983). Two opposing forces shape the comparison. Bundling pools each bidder's valuations across objects, drawing bidders with high draws on any single object into bidding for the whole bundle — a competition effect that increases competition and raises revenue. Bundling also allocates the entire bundle to one buyer even when other bidders value individual objects more, creating an allocative inefficiency. [6]

I evaluate the BTA-vs-MTA comparison empirically at the C-Block. The estimator of Section 6 runs in two stages. The first stage recovers the assortativity parameter λ^1, the geographic complementarity coefficients (λ^2, λ^3), and the license-level mean utilities δ_j from the realized 480-license allocation. The second stage recovers the price coefficient and the license-specific quality ξ_j from $\hat{\delta}_j$ by 2SLS. The same row-generated linear program, evaluated at the estimated primitives and at the 46-MTA supply structure, returns the LP-relaxation prices and bundles that correspond to the replica-economy competitive equilibrium (Bikhchandani and Mamer, 1997; Bikhchandani and Ostroy, 2002).

Section 10.1 describes the auction and data. Sections 10.2 through 10.4 state the model, report the two-stage estimation, and compute the MTA equilibrium. Section 10.6 contrasts the exercise with Fox and Bajari (2013).

10.1 Data

The auction was a simultaneous-ascending bid mechanism conducted over 184 rounds. In each round, bidders observed the current standing high bids on every license and could place new bids on any subset of licenses, with package bids not allowed. Eligibility — the maximum population a bidder's bids could cover at any time — was committed by an upfront deposit and reduced over the auction by an activity rule: bidders failing to be active on a sufficient fraction of their eligibility lost some of it in subsequent rounds. The auction ended when a round passed with no new bids, leaving each license assigned to the bidder holding the final standing high bid.

License populations are concentrated in a handful of large urban markets, with a long tail

[6] With two bidders, bundling raises revenue over separate sale (Palfrey, 1983). Above a threshold bidder count, the inefficiency effect outweighs the competition effect and separate sale raises revenue compared to bundling (Jehiel et al., 2007). Between these regimes, the comparison depends on the value distribution.

of small BTAs. Winning prices are strongly correlated with population (Table 6). Figure 3 shows the BTA partition and the MTA aggregation used in the counterfactual.

	Mean	Median	Min	Max
Population (millions)	0.51	0.19	0.03	18.10
Winning price ($M)	20.6	4.0	0.1	994.0
Price per person ($)	23.10	19.68	1.44	74.02

Table 6: License-level summary statistics, $N = 480$ continental BTAs (13 non-continental units excluded). Correlation between population and winning price: $\rho = 0.99$.

I restrict the analysis to the 480 continental BTAs and 252 of the 255 registered bidders. I exclude three bidders who won only non-continental licenses, together with the 13 non-continental licenses. Table 7 reports the distribution of eligibility caps and realized packages.

	Mean	Median
All bidders ($N = 252$)		
Eligibility (millions pop)	6.76	0.97
Assets (millions $)	40.5	14.7
Revenues (millions $)	12.67	4.26
Package size (# BTAs)	2	0
Package pop (millions)	0.98	0.00
Total winning bid (millions $)	39.2	0.0

Table 7: Bidder summary statistics.

Figure 2 plots the 12 largest winning packages by winner eligibility and the single-license winners. Large winners hold clusters of contiguous BTAs. Many small bidders win a single license.

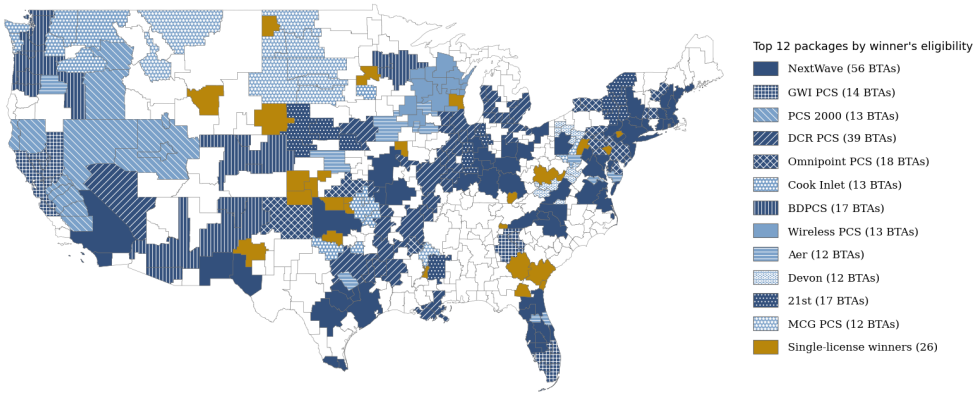


Figure 2: Selected winning packages: the 12 largest packages by winner eligibility and single-license winners.

10.2 Model and estimation

The utility specification follows Fox and Bajari (2013). Bidder i is characterized by an eligibility elig_i and taste shocks $\nu_i = (\nu_{ij})_{j\in[M]}$. For this application, x_i collects bidder eligibility elig_i, license populations $(\text{pop}_j)_{j\in[M]}$, and the pairwise license characteristics used to construct $w_{jj'}$ below. In the first stage I write the license fixed effect as δ_j. The value of license j is

$$\delta_j + \lambda^1 \,\text{elig}_i \,\text{pop}_j + \nu_{ij},$$

where ν_{ij} is a private idiosyncratic shock, whose distribution will be parametrized below. The value of a bundle d is

$$V(d;\delta,\lambda,x_i,\nu_i) \;=\; \sum_j d_j\,(\delta_j + \lambda^1\,\text{elig}_i\,\text{pop}_j + \nu_{ij}) \;+\; (\lambda^2 + \text{elig}_i\,\lambda^3)^\top \sum_{j<j'} d_j d_{j'} w_{jj'}.$$

where $w_{jj'} \in \mathbb{R}^4$ is constructed from four pairwise measures $\tilde{w}^k_{jj'}$.[7] Let $\psi(d,x_i) \in \mathbb{R}^{K_1}$ collect the features corresponding to $\lambda = (\lambda^1,\lambda^2,\lambda^3)$ in this specification.

Feasible bundles satisfy the eligibility constraint,

$$C_i = \{d \in \{0,1\}^M : \textstyle\sum_j d_j\,\text{pop}_j \le \text{elig}_i\}.$$

The errors are a κ-weighted combination of two private shocks:

$$\nu_{ij} \;=\; \kappa\,\text{pop}_j\,\eta^{(1)}_{ij} \;+\; (1-\kappa)\,\overline{\text{pop}}\,\eta^{(2)}_{ij}, \qquad \kappa \in [0,1],$$

with shocks independent across i and j: $\eta^{(1)}_{ij}$ is standard normal, and $\eta^{(2)}_{ij}$ is a Lognormal$(0,1)$ draw centered and normalized to variance one. Finally, $\overline{\text{pop}}$ is the cross-license mean of pop_j.

I do not provide formal asymptotic results for the C-Block exercise. In the first stage, the license mean utilities are estimated from variation across bidders as N grows. The K_1 complementarity coefficients use variation both across bidders and across license pairs within bundles.[8] In the second stage, I use cross-license variation, treating $(\text{pop}_j,\xi_j,p_j)_{j\in[M]}$ as a large-M array in the price regression.[9]

First stage. Fix the non-linear parameter κ. I estimate $(\hat{\lambda},\hat{\delta}) \in \mathbb{R}^{K_1}\times\mathbb{R}^M$ satisfying

1. Bidder optimality: $u_{si} \ge V(d;\hat{\delta},\hat{\lambda},x_i,\nu_{si})$ for every $d \in C_i$, with equality if $\hat{\mu}_{si}(d) > 0$.

[7] The four measures are adjacency $\mathbf{1}\{j \sim j'\}$, gravity $\text{pop}_j\text{pop}_{j'}/\text{dist}^4_{jj'}$, travel-survey trips, and air-passenger flows. Each pair weight sums the two population-scaled, row-normalized directions:

$$w^k_{jj'} \;=\; \frac{\text{pop}_j\,\tilde{w}^k_{jj'}}{\sum_\ell \tilde{w}^k_{j\ell}} + \frac{\text{pop}_{j'}\,\tilde{w}^k_{j'j}}{\sum_\ell \tilde{w}^k_{j'\ell}}, \qquad j \ne j'. \tag{12}$$

Set $w^k_{jj} = 0$ and take a term to be zero when its denominator is zero.

[8] The large-network experiments in Section 9 illustrate how the estimator uses high-dimensional choice profiles to recover common interaction parameters. I do not establish a joint (N,M) asymptotic theory.

[9] See, for instance, Jenish and Prucha (2009). I do not formally establish that primitives $(\text{pop}_j,\xi_j)_{j\in[M]}$ satisfying the assumptions for a law of large numbers imply the same property for the equilibrium vector $(\text{pop}_j,\xi_j,p_j)_{j\in[M]}$.

2. Moment matching:
$$\frac{1}{SN}\sum_{s,i,d}\hat{\mu}_{si}(d)\,\psi(d,x_i) \;=\; \frac{1}{N}\sum_{i}\psi(d_i,x_i).$$

3. License clearing: $\frac{1}{S}\sum_{s,i,d\ni j}\hat{\mu}_{si}(d)=1$ for each $j\in[M]$.

4. $\hat{\mu}_{si}(d)\geq 0$ and $\sum_{d\in C_i}\hat{\mu}_{si}(d)=1$ for every (s,i).

Therefore, conditional on $\hat{\lambda}$, the mean utilities $\hat{\delta}$ make the BTA market clear under preferences $V(d;\hat{\delta},\hat{\lambda},x_i,v_{si})$ in a replica economy with S simulation per observation and the supply scaled by S, which approximates a continuous market (Starr, 1969; Azevedo et al., 2013a).

I select the non-linear parameter κ in an outer loop by matching the model-implied bundle-size variance to the sample variance from the data. The model-side variance is read from the first-stage dual at κ:
$$\frac{1}{NS}\sum_{s,i}\sum_{d\in C_i}\hat{\mu}^{\kappa}_{si}(d)\,|d|^2 \;-\; \Big(\tfrac{1}{NS}\sum_{s,i}\sum_{d\in C_i}\hat{\mu}^{\kappa}_{si}(d)\,|d|\Big)^2,$$
where $\hat{\mu}^{\kappa}$ denotes the first-stage dual solution at κ, and $|d|$ is the bundle size. I take $\hat{\kappa}$ to minimize distance between this and the observed counterpart $\frac{1}{N}\sum_i(|d_i|-\overline{|d|})^2$. In the reported specification this gives $\hat{\kappa}=0.80$. This outer loop is an auxiliary moment selection, and I do not report separate inference for κ.

Second stage. Given $\hat{\delta}_j$ from the first stage, I estimate the price coefficient α from the 2SLS regression
$$\hat{\delta}_j=-\alpha\,p_j+\xi_j,$$
where ξ_j is the price-independent license component, including the common intercept, unobserved primitive quality, and expected resale or merger value.

For each BTA j, let $N_j=\{j'\neq j:1000\leq \text{dist}_{jj'}\leq 2500\,\text{km}\}$ index the BTAs in a distance band 1000–2500 km from j, and let $\text{density}_{j'}$ denote population per square kilometre at j'. The excluded instrument is the unweighted mean of density over N_j,
$$z_j \;=\; \frac{1}{|N_j|}\sum_{j'\in N_j}\text{density}_{j'}. \tag{13}$$

Prices are endogenous because a license with high ξ_j is more valuable and is likely to receive higher bids. The instrument isolates variation in p_j generated by market conditions far from j. Distant population density affects bidding for j through complementarities in bidders' package values and their eligibility constraints. I assume that ξ_j reflects local license attributes and is therefore uncorrelated with population density 1,000–2,500 km away. This assumption fails if unobserved quality or expected resale and merger value is correlated across broad geographic areas.

10.3 Results

First stage. Table 8 reports the nine λ-coefficients from the first stage estimation. Assortativity is positive and significant ($\hat{\lambda}^{1,\text{elig}} > 0$). As for complementarities, gravity and adjacency are positive and significant ($\hat{\lambda}^{2,\text{adj}}, \hat{\lambda}^{2,\text{gravity}} > 0$), with adjacency being particularly important for small bidders ($\hat{\lambda}^{3,\text{adj}} < 0$).

Covariate	Est.	(s.e.)
$\hat{\lambda}^{1,\text{elig}}$	$+1.888$	(0.434)
$\hat{\lambda}^{2,\text{adj}}$	$+0.635$	(0.229)
$\hat{\lambda}^{2,\text{gravity}}$	$+0.693$	(0.168)
$\hat{\lambda}^{2,\text{travel}}$	$+0.674$	(0.655)
$\hat{\lambda}^{2,\text{air}}$	$+0.732$	(0.516)
$\hat{\lambda}^{3,\text{adj}}$	-1.000	(0.382)
$\hat{\lambda}^{3,\text{gravity}}$	-0.398	(0.340)
$\hat{\lambda}^{3,\text{travel}}$	-0.015	(0.878)
$\hat{\lambda}^{3,\text{air}}$	-0.484	(0.861)

Table 8: Estimates of assortativity and complementarity parameters λ. Bayesian-bootstrap standard errors based on 1,000 replications are reported in parentheses.

	OLS	2SLS
c (const)	-0.0000 (0.0002)	-0.0006 (0.0002)
α (price)	-0.0081 (0.0236)	$+0.0752$ (0.0328)
First-stage F	n.a.	10.0 (price)
N	480	480

Table 9: Second-stage estimates: OLS vs. 2SLS regression of $\hat{\delta}_j$ on a constant and price.

Second stage. Second stage regression results are shown in Table 9. The reported standard errors are heteroskedasticity-robust for the second-stage regression, taking the first-stage estimates $\hat{\delta}_j$ as given. The counterfactuals below are evaluated at the 2SLS point estimate.

Baseline surplus decomposition. With $\hat{\alpha}$ in hand, I recover the gross license component as $\hat{\xi}_j = \hat{\delta}_j + \hat{\alpha} p_j$. Dividing utility by $\hat{\alpha}$, total bidder surplus decomposes as

$$\frac{1}{\hat{\alpha} S} \sum_{s,i} \hat{u}_{si} = \underbrace{\frac{1}{\hat{\alpha}} \sum_{i,j} d_{ij} \left(\hat{\lambda}^1 \, \text{elig}_i \, \text{pop}_j + \hat{\xi}_j \right) + \frac{1}{\hat{\alpha} \, S} \sum_{s,i,d} \hat{\mu}_{si}(d) \sum_j d_j \, \nu_{sij}}_{\text{gross standalone license value}} + \underbrace{\frac{1}{\hat{\alpha}} \sum_i \sum_{j<j'} d_{ij} \, d_{ij'} \, w_{jj'}^{\top} (\hat{\lambda}^2 + \text{elig}_i \, \hat{\lambda}^3)}_{\text{complementarities}} - \underbrace{\sum_j p_j}_{\text{prices paid}} , \tag{14}$$

the corresponding estimates of this decomposition are shown in Table 10.

Component	$B
Gross standalone license value	+5.63
Complementarities	+7.65
Adjacency	+1.02
Gravity	+3.23
Travel survey	+2.27
Air travel	+1.14
Price paid	−9.87
Total bidder surplus	**+3.41**

Table 10: Baseline surplus decomposition at estimated parameters. Gross standalone license value includes the price-independent license component, eligibility-population value, and dual-weighted simulated idiosyncratic value. Components defined in eq. (14).

10.4 Counterfactual: MTA-level licensing

I compute the counterfactual auction in which the FCC sells 46 MTA-level licenses. A bidder who wins MTA m acquires all BTAs in m.

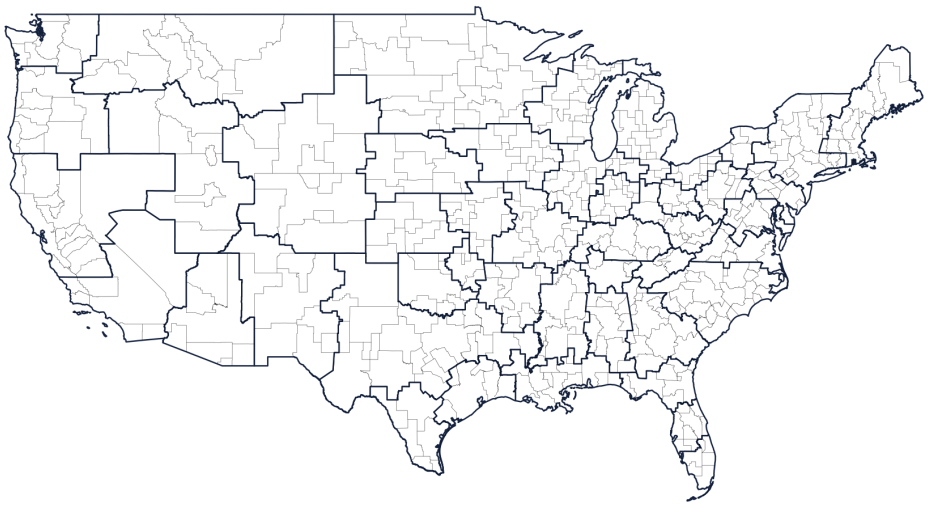


Figure 3: Thin lines: the 480 continental BTAs. Thick lines: 46 MTAs, the counterfactual unit of consolidation.

MTA valuations. Aggregating from BTA to MTA, the standalone value of MTA m for bidder i is

$$\begin{aligned}\gamma_{im}^{\text{MTA}} \;=\; & \sum_{j\in m} \left(\hat{\lambda}^1 \,\text{elig}_i \,\text{pop}_j + \hat{\xi}_j + \nu_{ij}\right) \\ & + \sum_{\substack{j<j' \\ j,j'\in m}} w_{jj'}^{\top}(\hat{\lambda}^2 + \text{elig}_i\,\hat{\lambda}^3),\end{aligned} \tag{15}$$

With $\text{pop}_m^{\text{MTA}} = \sum_{j\in m} \text{pop}_j$, for $m < m'$ across-MTA pairwise weights aggregate via $w_{mm'}^{\text{MTA}} = \sum_{j\in m,\, j'\in m'} w_{jj'}$. Bidder i's MTA payoff is

$$V_i^{\text{MTA}}(d) \;=\; \sum_m d_m\, \gamma_{im}^{\text{MTA}} \;+\; \sum_{m<m'} d_m\, d_{m'}\, (w_{mm'}^{\text{MTA}})^{\top}(\hat{\lambda}^2 + \text{elig}_i\,\hat{\lambda}^3) \;-\; \hat{\alpha}\sum_m d_m\, p_m^{\text{MTA}}, \tag{16}$$

a quadratic 0–1 knapsack in 46 variables, with $\sum_m d_m \,\text{pop}_m^{\text{MTA}} \le \text{elig}_i$.

Equilibrium computation. I compute the MTA counterfactual by solving the same replica-economy assignment LP under the 46-license supply structure. The dual variables on the item constraints are the reported MTA prices. The LP is solved in utility units, so I divide prices and surplus by $\hat{\alpha}$ to report them in dollars. The demand oracle is the 46-variable quadratic knapsack problem described above.

	Baseline	Counterfactual	Δ
Revenue ($B)	9.87	7.34	−25.6%
Bidder surplus ($B)	3.41	1.77	−48.0%
Total welfare ($B)	13.29	9.12	−31.4%

Table 11: Baseline and counterfactual outcomes under full MTA licensing, in billions of dollars.

Table 11 reports the full-MTA counterfactual. Revenue falls from \$9.87 billion to \$7.34 billion, bidder surplus falls from \$3.41 billion to \$1.77 billion, and total welfare falls from \$13.29 billion to \$9.12 billion.

10.5 Partial bundling

The 46-MTA partition is one of many coarsenings of the 480-BTA partition. For a threshold $\tau \in \{0, 1, \ldots, 46\}$, I consider the partition that bundles the τ MTAs with the largest aggregate population, $\text{pop}_m^{\text{MTA}} = \sum_{j \in m} \text{pop}_j$, and sells the remaining $46 - \tau$ MTAs as their constituent BTAs in singleton form. At $\tau = 0$ the partition is the original 480-BTA allocation. At $\tau = 46$ it is the full-MTA allocation. I solve for competitive-equilibrium prices under each such partition by the same row-generation LP, evaluated at $(\hat{\delta}, \hat{\lambda}, \hat{\alpha})$.

I report results for $\tau \in \{15, 20, 25, 30\}$ in Table 12. Total welfare and total revenue are monotone decreasing in τ on this range.

τ	# items	Surplus	Revenue	Welfare
0	480	3.41	9.87	13.29
15	285	2.51	8.93	11.44
20	230	2.30	8.53	10.83
25	187	2.14	8.25	10.39
30	144	1.95	8.02	9.97
46	46	1.77	7.34	9.12

Table 12: Surplus, revenue, and welfare along the partial-bundling path, in billions of dollars.

At $\tau = 15$, price retention $p_m^{\text{CF}} / p_m^{\text{base}}$ on the bundled MTAs ranges from 0.23 to 1.05 across the 15 MTAs. Two MTAs, Houston and Miami, have higher prices than in the baseline.

10.6 Comparison to Fox and Bajari (2013)

Fox and Bajari (2013) estimate the same auction by maximum rank correlation under pairwise stability in matches only. In the notation of Section 10.2, write the structural component of their payoff as

$$V_i^{\text{BF}}(d) \;=\; \text{elig}_i \sum_j d_j \,\text{pop}_j \;+\; \beta^\top \sum_{j<j'} d_j \, d_{j'} \, w_{jj'} \;+\; \sum_j d_j \, \xi_j, \tag{17}$$

with full payoff $V_i^{\text{BF}}(d) + \sum_j d_j \nu_{ij}$, where ν_{ij} is i.i.d. across i and j. Bajari and Fox normalize the assortative coefficient to a sign in $\{-1, +1\}$ and identify $+1$. They estimate β on a three-dimensional $w_{jj'}$ (gravity, travel-survey, air-travel — no adjacency). The rank-correlation estimator counts pairwise-swap inequalities in the structural component

$$V_i^{\text{BF}}(d_i) + V_{i'}^{\text{BF}}(d_{i'}) \;\geq\; V_i^{\text{BF}}\big(d_i - e_j + e_{j'}\big) + V_{i'}^{\text{BF}}\big(d_{i'} - e_{j'} + e_j\big), \tag{18}$$

indexed by pairs of winning bidders (i, i') and licenses j won by i, j' won by i', with $e_j \in \{0,1\}^M$ the j-th standard basis vector. The ξ contribution on each side equals $\sum_\ell (d_{i\ell} + d_{i'\ell})\, \xi_\ell$ — the

swap only reassigns j and j' between the two bidders, so ξ enters both sides identically and cancels. Their estimator returns $\hat{\beta}$ without recovering ξ_j or specifying the distribution of ν_{ij}.

Rationality condition. Fox and Bajari (2013) use pairwise stability: at the observed allocation, no two winning bidders can swap one license each and raise their joint valuation. I use ex-post optimality at observed prices: each bidder's bundle is best-response at p. Both conditions hold in a competitive equilibrium with quasilinear utilities. Pairwise stability also holds in the implicit-collusion equilibria of simultaneous ascending auctions, in which observed prices need not equal bidder valuations. Ex-post optimality at p does not hold in those settings.

Parameters recovered. Fox and Bajari (2013)'s criterion differences out ξ and p and does not impose a distribution on ν, so $\hat{\beta}$ is identified semiparametrically. The price coefficient and ξ_j are not separately recovered. My estimator recovers (δ, λ) in the first stage and (α, ξ) in the second stage by 2SLS regression of $\hat{\delta}_j$ on a constant and p_j. Therefore, the stronger rationality conditions deliver identification of the mean-utilities.

Counterfactual. Fox and Bajari (2013)'s counterfactuals evaluate the deterministic component of V^{BF} at some chosen deterministic allocations. They report a 48% deterministic-efficiency gain from consolidating the BTAs into four large regional licenses over the realized C-Block allocation. My counterfactual solves for the new competitive equilibrium under the 46-MTA supply structure at $(\hat{\delta}, \hat{\lambda}, \hat{\alpha})$. Total welfare falls by 31.4%. The two numbers therefore measure different objects under different rationality conditions.

In Table 13 I report the decomposition of the surplus at the deterministic allocations used in Fox and Bajari (2013) for counterfactuals. The table compares the counterfactuals under Fox and Bajari (2013)'s estimates as well as the first stage estimates I computed above.

Allocation	Fox & Bajari		This paper	
	Compl.	Assort + Compl.	Compl.	Assort + Compl.
C-block (85 actual)	0.127	0.521	+0.305	+0.699
480 separate	0.000	0.175	+0.000	+0.175
46 MTAs	0.230	0.428	+0.457	+0.655
4 regional	0.268	0.785	+0.617	+1.134
Nationwide (NextWave)	0.190	0.903	+0.731	+1.444

Table 13: Five candidate allocations of Fox and Bajari (2013) scored under their estimates and under my first-stage estimates.

The deterministic scores agree on the broad ordering that coarser regional packages score highly, but they do not reproduce the same ranking exactly. With complementarities only, my estimates put the nationwide allocation slightly above the four-region and 46-MTA allocations, whereas Fox and Bajari (2013) rank the four-region allocation first. With assortativeness and complementarities, both sets of estimates rank the nationwide allocation first and the realized

C-Block allocation above the 46-MTA allocation. These scores are deterministic allocation comparisons, not equilibrium counterfactuals.

11 Conclusions

This paper introduces a computationally tractable estimator for discrete choice models with large, combinatorial action spaces. By formulating estimation as the minimization of a convex regret criterion, the approach enables the recovery of high-dimensional parameter vectors, including item fixed effects and complementarities, while avoiding the computational bottlenecks of the simulated method of moments. I establish formal consistency and asymptotic normality for this estimator in a fixed-dimension, many-observation limit.

I apply this framework to the FCC C-Block auction to estimate a model with hundreds of license fixed effects and pairwise complementarities from a single realized allocation. The results show that the method scales to empirical environments that have previously been computationally prohibitive. The counterfactuals indicate that aggregating the 480 observed BTA licenses into 46 larger MTA packages reduces total welfare, bidder surplus, and seller revenue compared to the baseline allocation.

There are distinct caveats to this empirical exercise. The formal asymptotic theorems cover a many-independent-observation limit. They do not formally characterize the joint large-market asymptotics of a single auction or propagate first-stage LP uncertainty to the outer-loop parameters. Furthermore, the counterfactual pricing is defined in the simulated replica economy generated by the shock draws. It is not a model of bidding dynamics in the finite simultaneous ascending auction. Despite these limitations, the convex-regret linear-programming approach provides a practical path forward for estimation and counterfactual analysis in high-dimensional matching, network formation, and package auction settings.

References

Andersen, N. T., Giné, E., Ossiander, M., and Zinn, J. (1988). The central limit theorem and the law of iterated logarithm for empirical processes under local conditions. *Probability Theory and Related Fields*, 77(2):271–305.

Antràs, P., Fort, T. C., and Tintelnot, F. (2017). The margins of global sourcing: Theory and evidence from us firms. *American Economic Review*, 107(9):2514–2564.

Azevedo, E. M., Weyl, E. G., and White, A. (2013a). Walrasian equilibrium in large, quasi-linear markets. *Theoretical Economics*, 8(2):281–290.

Azevedo, E. M., Weyl, E. G., and White, A. (2013b). Walrasian equilibrium in large, quasilinear markets. *Theoretical Economics*, 8(2):281–290.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

Berry, S. T., Khwaja, A., Kumar, V., Musalem, A., Wilbur, K. C., Allenby, G. M., Anand, B., Chintagunta, P., Hanemann, W. M., Jeziorski, P., and Mele, A. (2014). Structural models of complementary choices. *Marketing Letters*, 25:245–256.

Bikhchandani, S. and Mamer, J. W. (1997). Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74(2):385–413.

Bikhchandani, S. and Ostroy, J. M. (2002). The package assignment model. *Journal of Economic Theory*, 107(2):377–406.

Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5(3):387–424.

Bonnet, O., Galichon, A., Hsieh, Y.-W., O'Hara, K., and Shum, M. (2022). Yogurts choose consumers? estimation of random-utility models via two-sided matching. *Review of Economic Studies*, 89(6):3085–3114.

Bontemps, C., Gualdani, C., and Remmy, K. (2023). Price competition and endogenous product choice in networks: Evidence from the us airline industry. CRC TR 224 Discussion Paper Series 400, University of Bonn and University of Mannheim, Germany.

Boros, E. and Hammer, P. L. (2002). Pseudo-boolean optimization. *Discrete Applied Mathematics*, 123(1):155–225.

Braides, A. (2006). A handbook of Γ-convergence. In Chipot, M., editor, *Handbook of Differential Equations: Stationary Partial Differential Equations*, volume 3, pages 101–213. Elsevier.

Brock, W. A. and Durlauf, S. N. (2001). Discrete choice with social interactions. *Review of Economic Studies*, 68(2):235–260.

Chernozhukov, V., Hong, H., and Tamer, E. (2007). Estimation and confidence regions for parameter sets in econometric models. *Econometrica*, 75(5):1243–1284.

Cramton, P. (1997). The FCC spectrum auctions: An early assessment. *Journal of Economics & Management Strategy*, 6(3):431–495.

Fox, J. T. (2007). Semiparametric estimation of multinomial discrete-choice models using a subset of choices. *The RAND Journal of Economics*, 38:1002–1019.

Fox, J. T. and Bajari, P. (2013). Measuring the efficiency of an fcc spectrum auction. *American Economic Journal: Microeconomics*, 5(1):100–146.

Fox, J. T. and Lazzati, N. (2017). A note on identification of discrete choice models for bundles and binary games. *Quantitative Economics*, 8(3):1021–1036.

Galichon, A. and Salanié, B. (2022). Cupid's invisible hand: Social surplus and identification in matching models. *Review of Economic Studies*, 89(5):2600–2629.

Gentzkow, M. (2007). Valuing new goods in a model with complementarity: Online newspapers. *American Economic Review*, 97(3):713–744.

Graham, B. S. (2016). Homophily and transitivity in dynamic network formation. NBER Working Paper 22186, National Bureau of Economic Research.

Graham, B. S. and Pelican, A. (2023). Estimating a class of strategic network formation games via maximum likelihood. *Working paper*.

Grötschel, M., Lovász, L., and Schrijver, A. (1981). The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197.

Gul, F. and Stacchetti, E. (2000). The english auction with differentiated commodities. *J. Econ. Theory*, 92:66–95.

Hajivassiliou, V., McFadden, D., and Ruud, P. (1996). Simulation of multivariate normal rectangle probabilities and their derivatives: Theoretical and computational results. *Journal of Econometrics*, 72(1–2):85–134.

Head, K., Mayer, T., Melitz, M., and Yang, C. (2026). Industrial policies for multi-stage production: The battle for battery-powered vehicles. Working Paper 34884, National Bureau of Economic Research.

Hendel, I. (1999). Estimating multiple-discrete choice models: An application to computerization. *Review of Economic Studies*, 66(2):423–446.

Hjort, N. L. and Pollard, D. (1993). Asymptotics for minimisers of convex processes. Unpublished manuscript.

Horst, R. and Thoai, N. V. (1999). DC programming: Overview. *Journal of Optimization Theory and Applications*, 103(1):1–43.

Iaria, A. and Wang, A. (2019). Identification and estimation of demand for bundles. *SSRN Electronic Journal*. Available at SSRN: https://ssrn.com/abstract=3458543.

Iwata, S., Fleischer, L., and Fujishige, S. (2001). A combinatorial strongly polynomial algorithm for minimizing submodular functions. *J. ACM*, 48(4):761–777.

Jehiel, P., Meyer-ter Vehn, M., and Moldovanu, B. (2007). Mixed bundling auctions. *Journal of Economic Theory*, 134(1):494–512.

Jenish, N. and Prucha, I. R. (2009). Central limit theorems and uniform laws of large numbers for arrays of random fields. *Journal of Econometrics*, 150(1):86–98.

Jia, P. (2008). What happens when wal-mart comes to town: An empirical analysis of the discount retail industry. *Econometrica*, 76(6):1263–1316.

Krauth, B. V. (2006). Simulation-based estimation of peer effects. *Journal of Econometrics*, 133(1):243–271.

Lovász, L. M. (1982). Submodular functions and convexity. In *International Symposium on Mathematical Programming*.

Manski, C. F. (1975). Maximum score estimation of the stochastic utility model of choice. *Journal of Econometrics*, 3(3):205–228.

Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3):531–542.

McFadden, D. (1977). Modelling the choice of residential location. *Transportation Research Record*.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995–1026.

Mele, A. (2017). A structural model of dense network formation. *Econometrica*, 85(3):825–850.

Milgrom, P. and Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277.

Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14(1):124–143.

Murota, K. (2022). Discrete convex analysis: A tool for economics and game theory.

Nesterov, Y. (2013). *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87 of *Applied Optimization*. Springer Science & Business Media, New York, NY.

Pakes, A. (2010). Alternative models for moment inequalities. *Econometrica*, 78(6):1783–1822.

Pakes, A. and Pollard, D. (1989). Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027–1057.

Pakes, A., Porter, J., Ho, K., and Ishii, J. (2015). Moment inequalities and their application. *Econometrica*, 83(1):315–334.

Palfrey, T. R. (1983). Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica*, 51(2):463–483.

Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press, Princeton.

Rockafellar, R. T. and Wets, R. J.-B. (1998). *Variational Analysis*, volume 317 of *Grundlehren der mathematischen Wissenschaften*. Springer.

Rubin, D. B. (1981). The bayesian bootstrap. *The Annals of Statistics*, 9(1):130–134.

Schrijver, A. (2000). A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *Journal of Combinatorial Theory, Series B*, 80(2):346–355.

Starr, R. M. (1969). Quasi-equilibria in markets with non-convex preferences. *Econometrica*, 37(1):25–38.

Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press.

A Additional economic applications

A.1 Jia (2008) with asymmetric congestion

I extend Example 2 by allowing the congestion parameter to differ by chain. Payoffs are

$$\pi^t(y^t, y^{-t}; \theta, x, \nu) \;=\; \sum_j y_j^t(\theta^1 r_j^t + \nu_j^t - \theta^{3,t} y_j^{-t}) \;+\; \theta^2 \sum_{jj'} y_j^t y_{j'}^t w_{jj'}, \qquad t \in \{A, B\},$$

with $\theta^{3,A}, \theta^{3,B} > 0$. The shocks satisfy

$$\begin{pmatrix} \nu_j^A \\ \nu_j^B \end{pmatrix} \;\sim\; N\!\left(0,\, \begin{pmatrix} \sigma_A^2 & \rho\,\sigma_A\sigma_B \\ \rho\,\sigma_A\sigma_B & \sigma_B^2 \end{pmatrix}\right) \quad \text{i.i.d. across } j,$$

with $\sigma_A, \sigma_B > 0$ and $\rho \in (-1, 1)$. The main-text example is the symmetric-congestion case $\theta^{3,A} = \theta^{3,B}$.

The change of variables $\bar{y}^B := 1 - y^B$ yields the weighted potential

$$\begin{aligned} V(d; \theta, x, \nu) \;=\; & \tfrac{1}{\theta^{3,A}}\Big[\sum_j y_j^A(\theta^1 r_j^A + \nu_j^A) + \theta^2 \sum_{jj'} y_j^A y_{j'}^A w_{jj'}\Big] \\ & + \tfrac{1}{\theta^{3,B}}\Big[\sum_j (1-\bar{y}_j^B)(\theta^1 r_j^B + \nu_j^B) + \theta^2 \sum_{jj'} (1-\bar{y}_j^B)(1-\bar{y}_{j'}^B) w_{jj'}\Big] \\ & - \sum_j y_j^A (1 - \bar{y}_j^B). \end{aligned}$$

Write the normalized coefficients as

$$\vartheta^{i,t} \;:=\; \theta^i / \theta^{3,t}, \qquad \tilde{\nu}_j^t \;:=\; \nu_j^t / \theta^{3,t}, \qquad \tilde{\sigma}_t \;:=\; \sigma_t / \theta^{3,t}, \qquad i = 1, 2,\; t \in \{A, B\}.$$

$$\mathrm{Var}(\tilde{\nu}_j^t) \;=\; \tilde{\sigma}_t^2, \qquad \mathrm{Cov}(\tilde{\nu}_j^A, \tilde{\nu}_j^B) \;=\; \rho\,\tilde{\sigma}_A \tilde{\sigma}_B.$$

The potential becomes

$$\begin{aligned} V(d; \vartheta, x, \tilde{\nu}) \;=\; & \vartheta^{1,A} \sum_j y_j^A r_j^A \;+\; \vartheta^{1,B} \sum_j (1-\bar{y}_j^B) r_j^B \\ & + \vartheta^{2,A} \sum_{jj'} y_j^A y_{j'}^A w_{jj'} \;+\; \vartheta^{2,B} \sum_{jj'} (1-\bar{y}_j^B)(1-\bar{y}_{j'}^B) w_{jj'} \\ & - \sum_j y_j^A (1 - \bar{y}_j^B) \;+\; \sum_j y_j^A\, \tilde{\nu}_j^A \;+\; \sum_j (1-\bar{y}_j^B)\, \tilde{\nu}_j^B, \end{aligned}$$

with $(\tilde{\nu}_j^A, \tilde{\nu}_j^B) \sim N(0, \tilde{\Sigma})$ parametrized by $(\tilde{\sigma}_A, \tilde{\sigma}_B, \rho)$.

Assume $\vartheta^{i,t} > 0$ for $i = 1, 2$ and $t \in \{A, B\}$. The four normalized linear coefficients cannot be arbitrary. Since the same θ^1 and θ^2 appear in both chains,

$$\frac{\vartheta^{1,A}}{\vartheta^{1,B}} = \frac{\vartheta^{2,A}}{\vartheta^{2,B}}.$$

Conversely, impose the scale normalization $\sigma_A = 1$. Given normalized parameters satisfying

this restriction, recover

$$\theta^{3,A} = \frac{1}{\tilde{\sigma}_A}, \qquad \theta^1 = \frac{\vartheta^{1,A}}{\tilde{\sigma}_A}, \qquad \theta^2 = \frac{\vartheta^{2,A}}{\tilde{\sigma}_A},$$

and

$$\theta^{3,B} = \frac{\vartheta^{1,A}}{\vartheta^{1,B}} \frac{1}{\tilde{\sigma}_A} = \frac{\vartheta^{2,A}}{\vartheta^{2,B}} \frac{1}{\tilde{\sigma}_A}, \qquad \sigma_B = \tilde{\sigma}_B \theta^{3,B}.$$

The ratio restriction makes the two expressions for $\theta^{3,B}$ coincide. With $\sigma_A = 1$, the displayed formulas recover the structural parameters from the normalized ones.

A.2 Other applications

This subsection considers two further applications. The first illustrates estimation with gross-substitutes valuations. The second gives a linear-in-parameters formulation of a multi-stage facility-location problem.

Unlike the supermodular problems in Section 5, general submodular maximization is NP-hard. Nevertheless, a subclass of submodular valuations, namely gross-substitutes valuations (Gul and Stacchetti, 2000), can be maximized in $\mathcal{O}(M^2)$ queries of the valuation and arithmetic operations using a greedy algorithm: starting from the empty set, greedy adds iteratively the item which gives highest marginal utility, if positive, until termination. Therefore, for unconstrained feasible sets ($C_i = \{0,1\}^M$), the estimation problem with gross-substitutes is polynomial time. Nevertheless, unlike supermodular and submodular valuations, gross-substitutes valuations are not generally closed under addition, hence under Assumption 1 the formulation will typically support a single gross-substitute covariate and possibly many modular covariates.[10] The following example illustrates this.

Example 3 (Airline routing network, Bontemps et al. (2023))**.** Airline i has hubs H_i. Let each item $j \in [M]$ be an undirected link $e_j = \{c_j, c_j'\}$ between two cities, at most one of which is a hub in H_i. The airline chooses a routing network $d \in C_i = \{0,1\}^M$. For $h \in H_i$, let

$$q_{ih}(d) = \sum_{j \in [M]} d_j \mathbf{1}\{h \in e_j\}$$

be the number of selected routes incident to hub h. Consider

$$V(d; \theta, x_i, \nu_i) = \sum_{j \in [M]} d_j \left(r_{ij}^\top \theta^{\text{mod}} + \nu_{ij} \right) - \theta^{\text{gs}} \sum_{h \in H_i} q_{ih}(d)^2, \qquad \theta^{\text{gs}} \geq 0.$$

In this example I estimate over the restricted parameter space $\theta^{\text{gs}} \geq 0$. The linear-programming formulation above is unchanged except that the master problem includes this linear inequality. The modular term captures route-level revenue, operating costs, and shocks. The second term captures hub congestion: each additional route using hub h raises delays, gate scarcity, or scheduling costs more when the airline already routes many links through h. Thus routes incident to the same hub are substitutes, while routes assigned to different hubs interact only through the modular terms.

[10] Indeed, the sum of a gross-substitute valuation and a modular valuation is still a gross-substitute valuation.

By construction, a selected route is incident to at most one hub. The congestion term therefore counts routes separately at each hub, with no route appearing in two hub counts. If adding a route sends one more route through hub h, its congestion contribution is

$$-\theta^{gs}\left((q_{ih}(d)+1)^2 - q_{ih}(d)^2\right) = -\theta^{gs}(2q_{ih}(d)+1).$$

Adding a route through hub h lowers the marginal value of other routes through h, but leaves the marginal value of routes through other hubs unchanged. Separable concave congestion of this form is gross substitutes (Murota, 2022). Since the remaining revenue, cost, and shock terms are modular, the full utility is gross substitutes.

The linear-in-parameters specification in Assumption 1 might be restrictive in some applications. For instance, in international trade, firm's utility is usually derived from a CES formulation which features non-linear parameters (see for instance Antràs et al. (2017)). In the next example, I provide a linearization of the setting studied in Head et al. (2026). [11]

Example 4. In Head et al. (2026) a multinational firm decides where to open plants along a multi-stage value chain and which destination markets to serve. With two production stages (cell manufacturing and vehicle assembly), sets of potential locations L_1 and L_2, and destination markets R, the firm chooses cell activations $d^1_{\ell_1} \in \{0,1\}$ ($\ell_1 \in L_1$), assembly activations $d^2_{\ell_2} \in \{0,1\}$ ($\ell_2 \in L_2$), and path indicators $d^3_{r,\ell_1,\ell_2} \in \{0,1\}$. Here $d^3_{r,\ell_1,\ell_2} = 1$ means market r is served from cell location ℓ_1 through assembly location ℓ_2. The bundle is $d = (d^1, d^2, d^3)$. Let $a_r \in \{0,1\}$ be an auxiliary market-entry indicator, subject to

$$\sum_{\ell_1,\ell_2} d^3_{r,\ell_1,\ell_2} = a_r, \qquad d^3_{r,\ell_1,\ell_2} \le d^1_{\ell_1}, \qquad d^3_{r,\ell_1,\ell_2} \le d^2_{\ell_2}.$$

The observables x collect market sizes A_r, inter-stage distances Δ_{ℓ_1,ℓ_2} and $\Delta_{\ell_2,r}$, and headquarters distances $\Delta_{\ell,h}$. The linear parameters $\theta = (\alpha_0, \alpha_1, \alpha_2, \beta_1, \beta_2, \gamma_1, \gamma_2)$ collect a fixed revenue coefficient (α_0), per-stage trade-cost coefficients on path distance (α_1, α_2), per-stage baseline facility costs (β_1, β_2), and per-stage headquarters-distance coefficients (γ_1, γ_2). The firm's payoff is

$$\begin{aligned} V(d;\theta,x,\nu) = &\sum_{r\in R} A_r \sum_{\ell_1,\ell_2} d^3_{r,\ell_1,\ell_2}\left(\alpha_0 - \alpha_1\Delta_{\ell_1,\ell_2} - \alpha_2\Delta_{\ell_2,r}\right) \\ &- \sum_{\ell_1\in L_1} d^1_{\ell_1}\left(\beta_1 + \gamma_1\Delta_{\ell_1,h}\right) - \sum_{\ell_2\in L_2} d^2_{\ell_2}\left(\beta_2 + \gamma_2\Delta_{\ell_2,h}\right) + \varepsilon(d;\nu), \end{aligned}$$

linear in θ, with modular Gaussian noise ε, so Assumption 1 holds.[12]

The maximization over $(d,a) \in C$ is the Multi-level Uncapacitated Facility Location problem with endogenous market entry, formulated and solved as a mixed integer linear program in Head et al. (2026).

[11] The computational methods in this paper can be extended to settings where the utility is convex-in-parameters, as in CES preferences. In this case, the resulting estimation problem becomes a difference-of-convex (DC) optimization problem (see, e.g., Horst and Thoai, 1999). Such problems are nonlinear and do not generally have global convergence guarantees, but the methods in this paper can still be used in DC algorithms.

[12] Linear-in-θ analogue of Head et al. (2026). Their CES variable profits and lognormal fixed costs together place the full model outside Assumption 1.

B Computational details

B.1 Supermodular maximization

Most of the optimization literature states the problem as submodular minimization. I state the equivalent problem as supermodular maximization because the main text writes choice and equilibrium selection as maximization of V.

Let V be a pseudo-Boolean function on $\{0,1\}^M$. It is supermodular if, for every $d, d' \in \{0,1\}^M$,

$$V(d \vee d') + V(d \wedge d') \;\geq\; V(d) + V(d'),$$

where $\vee$ and $\wedge$ denote coordinatewise maximum and minimum. Equivalently, the marginal value of changing coordinate j from zero to one is weakly increasing in the other coordinates. With $f := -V$, maximizing V is the same problem as minimizing the submodular function f.

The Lovász extension gives a continuous formulation of this sign change. Identify a subset $A \subseteq [M]$ with its incidence vector and write $V(A)$ for the value of that vector. For $z \in [0,1]^M$, order coordinates so that $z_{\pi(1)} \geq \cdots \geq z_{\pi(M)}$. Let $A_k := \{\pi(1), \ldots, \pi(k)\}$ and $A_0 := \varnothing$. The Lovász extension of V is

$$\widehat{V}(z) \;:=\; V(A_0) + \sum_{k=1}^{M} z_{\pi(k)} \big(V(A_k) - V(A_{k-1})\big).$$

It agrees with V at vertices of the cube. Moreover, V is supermodular if and only if $\widehat{V}$ is concave, equivalently $-\widehat{V}$ is the convex Lovász extension of the submodular function $f = -V$ (Lovász, 1982). Hence

$$\max_{d \in \{0,1\}^M} V(d) \;=\; \max_{z \in [0,1]^M} \widehat{V}(z).$$

After solving the continuous problem at a maximizer $z^\star$, a discrete maximizer is obtained by thresholding $z^\star$: for some threshold t, the set $\{j : z^\star_j \geq t\}$ attains the same value. Thus the general supermodular demand oracle can be implemented either through the Lovász-extension formulation or through combinatorial submodular-minimization algorithms (Schrijver, 2000; Iwata et al., 2001).

For the quadratic potentials used in Section 5, there is a simpler graph-cut construction. Let

$$V(d) \;=\; c + \sum_{j \in [M]} a_j d_j + \sum_{j<k} b_{jk} d_j d_k, \qquad b_{jk} \geq 0.$$

The nonnegative interaction coefficients are exactly the pairwise supermodularity restrictions. Since

$$d_j d_k \;=\; \tfrac{1}{2} d_j + \tfrac{1}{2} d_k - \tfrac{1}{2} |d_j - d_k|,$$

minimizing $-V$ is equivalent, up to an additive constant, to minimizing

$$\sum_{j \in [M]} \alpha_j d_j + \sum_{j<k} \frac{b_{jk}}{2} |d_j - d_k|, \qquad \alpha_j := -a_j - \frac{1}{2} \sum_{k \neq j} b_{jk}.$$

Build an s–t cut problem with one node for each $j \in [M]$. Add an undirected edge $\{j,k\}$ with

capacity $b_{jk}/2$. If $\alpha_j \geq 0$, add an edge $\{j, t\}$ with capacity α_j. If $\alpha_j < 0$, add an edge $\{s, j\}$ with capacity $-\alpha_j$ and add the constant α_j to the objective. For a cut (S, T) with $s \in S$ and $t \in T$, set $d_j = 1\{j \in S\}$. The cut value equals $-V(d)$ plus a constant, so a minimum cut gives a maximizer of V (Boros and Hammer, 2002).

This graph-cut statement is for maximization over the full binary lattice. Extra restrictions in the feasible set, such as knapsack constraints, require additional problem-specific structure or a separate oracle.

B.2 Oracle complexity

For the realized sample, define

$$G_{N,S} := \max_{(d_{si}) \in \prod_{s,i} C(x_i)} \left\| \frac{1}{SN} \sum_{s,i} \phi(d_{si}; x_i) - \widehat{m} \right\|.$$

Theorem 1 (Oracle complexity). *Let $S, N \geq 1$. Let $\bar{\Theta} \subset \mathbb{R}^K$ be a full-dimensional compact convex set with a strong separation oracle and known inner and outer radii $0 < r \leq R$. Then $Q^{N,S}$ is $G_{N,S}$-Lipschitz on $\mathbb{R}^K$. For every $\tau > 0$, the ellipsoid method computes a τ-optimal parameter using*

$$\mathcal{O}\left(K^2 \left[\log \frac{R}{r} + \log\left(1 + \frac{R G_{N,S}}{\tau} \right) \right] \right)$$

full-sample subgradient evaluations. Each full-sample evaluation requires SN demand-oracle calls and SN feature evaluations.

If the inputs are rational and the demand oracle returns an exact optimizer in polynomial time, an exact primal-dual solution to (7)–(8) is computable in time polynomial in the input length and the running time of the demand oracle.

Proof. Each tuple $(d_{si})_{s,i}$ indexes an affine row with slope

$$\frac{1}{SN} \sum_{s,i} \phi(d_{si}; x_i) - \widehat{m}.$$

Its norm is at most $G_{N,S}$, so the maximum of these rows, $Q^{N,S}$, is $G_{N,S}$-Lipschitz. At any θ, choose a demand optimizer for each (s, i). The resulting tuple indexes an active row, and Lemma 2 identifies its slope as a subgradient. The ellipsoid bound then follows from the standard subgradient-separation argument for a convex body with inner and outer radii r and R (Nesterov, 2013). Each subgradient evaluation solves one demand problem and evaluates one feature vector for every (s, i).

For rational inputs, the demand oracle is a strong separation oracle for the feasible polyhedron of (7). The equivalence of strong separation and exact optimization gives the final claim (Grötschel et al., 1981). □

B.3 Slack formulations

Proof of Lemma 3. Fix θ and write

$$F_{si}(d;\theta) := \phi(d;x_i)^\top\theta + \varepsilon(d;\nu_{si}).$$

Profiling out the slacks in (7) gives

$$\frac{1}{SN}\sum_{s,i}\max_{d\in C(x_i)} F_{si}(d;\theta) - \widehat{m}^\top\theta.$$

Profiling out $\bar{u}$ in (10) gives

$$\begin{aligned}\max_{(d_{si})\in\prod_{s,i}C(x_i)} \frac{1}{SN}\sum_{s,i} F_{si}(d_{si};\theta) - \widehat{m}^\top\theta \\ = \frac{1}{SN}\sum_{s,i}\max_{d\in C(x_i)} F_{si}(d;\theta) - \widehat{m}^\top\theta,\end{aligned}$$

where the equality follows because the choice set is a Cartesian product across (s,i). The profiled objectives therefore agree at every θ, so the two programs have the same optimal value and the same optimal parameters. □

B.4 Limited fractionality

For a dual-feasible μ, define

$$\operatorname{supp}(\mu_{si}) := \{d \in C(x_i) : \mu_{si}(d) > 0\}, \qquad \mathcal{I}(\mu) := \{(s,i) : |\operatorname{supp}(\mu_{si})| \geq 2\}.$$

Proposition 6 (Limited fractionality under replication)**.** *Suppose $S, N \geq 1$ and $d_i \in C(x_i)$ for every i. Every extreme point μ of the feasible set of (8) satisfies*

$$|\mathcal{I}(\mu)| \leq K.$$

Consequently, (8) has an optimal solution $\widehat{\mu}$ for which

$$\frac{|\mathcal{I}(\widehat{\mu})|}{SN} \leq \frac{K}{SN}.$$

For fixed N and K, the fraction of simulated bidders split across bundles therefore converges to zero as $S \to \infty$.

Proof. Let μ be an extreme point and suppose $|\mathcal{I}(\mu)| > K$. For each $a = (s,i) \in \mathcal{I}(\mu)$, choose two distinct bundles $d_a^0, d_a^1 \in \operatorname{supp}(\mu_{si})$. Since there are more than K vectors

$$\phi(d_a^1;x_i) - \phi(d_a^0;x_i) \in \mathbb{R}^K, \qquad a \in \mathcal{I}(\mu),$$

they are linearly dependent. Hence there are coefficients $(c_a)_a$, not all zero, such that

$$\sum_{a=(s,i)\in\mathcal{I}(\mu)} c_a\big[\phi(d_a^1;x_i)-\phi(d_a^0;x_i)\big]=0.$$

Define a perturbation h by

$$h_{si}(d_a^1)=c_a, \qquad h_{si}(d_a^0)=-c_a,$$

for $a=(s,i)\in\mathcal{I}(\mu)$, and set all other entries to zero. The entries of h_{si} sum to zero for every (s,i), and the displayed linear dependence implies that its aggregate feature vector is zero. Thus $\mu+th$ and $\mu-th$ satisfy all equality constraints of (8). Since the perturbed entries of μ are strictly positive, both vectors remain nonnegative for sufficiently small $t>0$. They are distinct feasible points whose average is μ, contradicting extremality. Therefore $|\mathcal{I}(\mu)|\le K$.

The feasible set of (8) is a nonempty compact polytope. It is nonempty because assigning each replica the observed bundle d_i satisfies the simplex and feature constraints. The linear dual objective therefore has an optimal extreme point, to which the first part applies. The ratio bound and its limit follow immediately. □

C Identification and asymptotic theory

This appendix proves Proposition 2, establishes outer and Hausdorff set-consistency, and derives asymptotic normality in the point-identified case for a fixed number of simulation draws.

C.1 Point identification

Lemma 4 (Matrix form of the rank condition)**.** *Suppose*

$$\mathbb{E}\Big[\max_{d\in C_i}\|\phi(d;x_i)\|^2\Big]<\infty.$$

Then the rank condition in Proposition 2 is equivalent to

$$\mathbb{E}\Big[\sum_{d,d'\in C_i}\big(\phi(d;x_i)-\phi(d';x_i)\big)\big(\phi(d;x_i)-\phi(d';x_i)\big)^\top\Big]\succ 0.$$

Proof. For every $\theta\in\mathbb{R}^K$, the quadratic form associated with the matrix in the lemma is

$$\mathbb{E}\Big[\sum_{d,d'\in C_i}\big(\theta^\top(\phi(d;x_i)-\phi(d';x_i))\big)^2\Big].$$

The integrand is non-negative, and it is positive exactly on the event

$$\big\{\exists\, d,d'\in C_i:\theta^\top(\phi(d;x_i)-\phi(d';x_i))\neq 0\big\}.$$

Its expectation is positive if and only if this event has positive probability. This holds for every $\theta\neq 0$ precisely under the rank condition. □

Proof of Proposition 2. Fix $\theta \in \Theta_0$. By the characterization of Θ_0 in §4, $\rho(\theta, x_i, \nu_i, d_i) = 0$ $\mathbb{P}$-a.s., so d_i maximizes $d \mapsto \theta^\top \phi(d; x_i) + \sum_{j\in[M]} d_j \nu_{ij}$ over C_i, and by (1) it maximizes the same objective at θ_0. By Fubini, for $\mathbb{P}_x$-almost every x, the two objectives have a common maximizer for $\mathbb{P}_{\nu|x}(\cdot \mid x)$-almost every ν. Fix such an x, write $C_x := C(x)$, and define

$$h(\nu) := \max_{d\in C_x}\{\theta^\top \phi(d; x) + \nu^\top d\} - \max_{d\in C_x}\{\theta_0^\top \phi(d; x) + \nu^\top d\}.$$

As a difference of finite maxima of affine functions, h is continuous.

Step 1: h is constant. At any ν admitting a common maximizer,

$$h(\nu) \;=\; \big[\theta^\top \phi(d^\star; x) + \nu^\top d^\star\big] - \big[\theta_0^\top \phi(d^\star; x) + \nu^\top d^\star\big] \;=\; (\theta - \theta_0)^\top \phi(d^\star(\nu); x),$$

because the shock term cancels. Thus h takes at most $|C_x|$ values on a set of full $\mathbb{P}_{\nu|x}$-measure. Full support makes that set dense, so continuity implies that $h(\mathbb{R}^M)$ is finite. It is also connected. Hence $h \equiv c$ for some constant c.

Step 2: $c = (\theta - \theta_0)^\top \phi(d; x)$ for every $d \in C_x$. Fix $d \in C_x$ and put $\nu_t := t(2d - \mathbf{1}_M)$ for $t > 0$, where $\mathbf{1}_M$ is the vector of ones of size M. For any $e \in \{0,1\}^M$, $\sum_{j\in[M]}(2d_j - 1)(d_j - e_j) = \sum_{j\in[M]} |d_j - e_j|$, which is at least 1 when $e \neq d$. Hence for $t > \max_{e\in C_x} |\theta_0^\top(\phi(e; x) - \phi(d; x))|$, d is the unique maximizer of $e \mapsto \theta_0^\top \phi(e; x) + \nu_t^\top e$ on C_x. These are finitely many strict inequalities, so they persist on an open neighbourhood of ν_t, which has positive $\mathbb{P}_{\nu|x}$-measure by full support and therefore meets the full-measure set of Step 1. At such a ν, uniqueness of the θ_0-maximizer forces $d^\star(\nu) = d$, so $c = h(\nu) = (\theta - \theta_0)^\top \phi(d; x)$.

Step 2 gives $(\theta - \theta_0)^\top \big(\phi(d; x) - \phi(d'; x)\big) = 0$ for all $d, d' \in C_x$, for $\mathbb{P}_x$-almost every x. If $\theta \neq \theta_0$, the second condition of the proposition applied to $\theta - \theta_0$ gives a set of x of positive $\mathbb{P}_x$-measure on which some such difference is non-zero, a contradiction. Hence $\Theta_0 \subseteq \{\theta_0\}$, and $\theta_0 \in \Theta_0$ by (1). □

C.2 Regret criterion

Identification. Recall from §4

$$Q(\theta) \;=\; \mathbb{E}\big[\rho(\theta, x_i, \nu_i, d_i) + \varepsilon(d_i; \nu_i)\big], \qquad \Theta_0 \;=\; \arg\min_{\theta\in\mathbb{R}^K} Q(\theta).$$

Set $\omega_i := (x_i, \nu_i, d_i)$ and write the regret as $\rho(\theta, \omega_i)$. Thus Q differs from $\int \rho(\theta, \omega)\, \mathbb{P}(d\omega)$ only by a finite constant.

Lemma 5 (Convex Lipschitz regret)**.** *Under Assumption 1, $\theta \mapsto \rho(\theta, x_i, \nu_i, d_i)$ is convex on $\mathbb{R}^K$, and*

$$|\rho(\theta_1, x_i, \nu_i, d_i) - \rho(\theta_2, x_i, \nu_i, d_i)| \;\le\; L_i \,\|\theta_1 - \theta_2\|, \qquad L_i \;:=\; 2\max_{d\in C_i} \|\phi(d; x_i)\|.$$

Proof. By Assumption 1, $V(\cdot; \theta, x_i, \nu_i)$ is affine in θ, so $\theta \mapsto \max_{d\in C_i} V(d; \theta, x_i, \nu_i)$ is a finite maximum of affine functions and hence convex. Subtracting the affine $\theta \mapsto \theta^\top \phi(d_i; x_i) +$

$\varepsilon(d_i; \nu_i)$ preserves convexity. For the Lipschitz bound,

$$\rho(\theta_1, x_i, \nu_i, d_i) - \rho(\theta_2, x_i, \nu_i, d_i) = \max_{d \in C_i} [\theta_1^\top \phi(d; x_i) + \varepsilon(d; \nu_i)] - \max_{d \in C_i} [\theta_2^\top \phi(d; x_i) + \varepsilon(d; \nu_i)] \\ - (\theta_1 - \theta_2)^\top \phi(d_i; x_i),$$

The difference of maxima is bounded in absolute value by $\|\theta_1 - \theta_2\| \max_{d \in C_i} \|\phi(d; x_i)\|$. The final term obeys the same bound because $d_i \in C_i$. □

Lemma 6 (Moment bound). *Fix $p \in \{1, 2\}$. If $\mathbb{E}|\phi_k(d; x_i)|^p < \infty$ for every $d \in \{0, 1\}^M$ and $k \in [K]$, then $L_i \in \mathcal{L}^p(\mathbb{P})$ and $\rho(\theta_0, \cdot) \in \mathcal{L}^p(\mathbb{P})$.*

Proof. $C_i \subseteq \{0, 1\}^M$ is finite, so $L_i \leq 2 \sum_{d \in \{0,1\}^M} \sum_{k \in [K]} |\phi_k(d; x_i)|$. A finite sum of $\mathcal{L}^p$ random variables is in $\mathcal{L}^p$. For $\rho(\theta_0, \cdot)$, (1) gives $\rho(\theta_0, x_i, \nu_i, d_i) = 0$ $\mathbb{P}$-a.s. □

Lemma 7 (Coercivity and the identified set). *Under Assumption 1, Q is convex and finite on $\mathbb{R}^K$. Then*

$$\Theta_0 \text{ is non-empty and compact} \iff Q(\theta) \to \infty \text{ as } \|\theta\| \to \infty.$$

Proof. Finiteness and convexity of the expected regret follow from Lemmas 5–6 with $p = 1$, since

$$|\rho(\theta, x_i, \nu_i, d_i)| \leq |\rho(\theta_0, x_i, \nu_i, d_i)| + L_i \|\theta - \theta_0\| = L_i \|\theta - \theta_0\|.$$

Assumption 1 makes $\mathbb{E}\varepsilon(d_i; \nu_i)$ finite, so the same properties hold for Q. A finite convex function on $\mathbb{R}^K$ is closed and proper. If Q is coercive, continuity implies that it attains its minimum and that its argmin is compact. Conversely, if Θ_0 is non-empty and compact, the minimum sublevel set is bounded. By Rockafellar (1970, Cor. 8.7.1), every sublevel set is then bounded, which is equivalent to coercivity. □

Simulated sample objective. For $S \geq 1$ let $\tilde{\omega}_i := (x_i, d_i, \nu_{1i}, \ldots, \nu_{Si})$ with $(\nu_{si})_{s \in [S]}$ i.i.d. draws from $\mathbb{P}_{\nu|x}(\cdot \mid x_i)$, independent across i. Define

$$\tilde{\rho}(\theta, \tilde{\omega}_i) := \frac{1}{S} \sum_{s=1}^{S} \left[\max_{d \in C_i} \{\theta^\top \phi(d; x_i) + \varepsilon(d; \nu_{si})\} - \theta^\top \phi(d_i; x_i) \right], \qquad Q^{N,S}(\theta) := \frac{1}{N} \sum_{i=1}^{N} \tilde{\rho}(\theta, \tilde{\omega}_i).$$

Proof of Lemma 1. Since each C_i is finite,

$$Q^{N,S}(\theta) = \max_{(d_{si}) \in \prod_{s,i} C_i} \frac{1}{NS} \sum_{i,s} \left[\theta^\top \phi(d_{si}; x_i) + \varepsilon(d_{si}; \nu_{si}) - \theta^\top \phi(d_i; x_i) \right],$$

which is a finite maximum of affine functions of θ. □

The simulated criterion requires a separate moment bound. Fix $p \in \{1, 2\}$. For each s,

$$\left| \max_{d \in C_i} \{\theta_0^\top \phi(d; x_i) + \varepsilon(d; \nu_{si})\} - \theta_0^\top \phi(d_i; x_i) \right| \leq 2\|\theta_0\| \max_{d \in C_i} \|\phi(d; x_i)\| + \max_{d \in C_i} |\varepsilon(d; \nu_{si})|.$$

Because $C_i \subseteq \{0, 1\}^M$ is finite and ν_{si} has the same conditional law as ν_i given x_i, the right-hand side is in $\mathcal{L}^p$ when the feature and shock moments of order p are finite. Since S is fixed,

$\tilde{\rho}(\theta_0, \tilde{\omega}_i) \in \mathcal{L}^p$. The $(\tilde{\omega}_i)_{i\geq 1}$ are i.i.d., and Lemma 5 applied to each summand shows that $\tilde{\rho}(\cdot, \tilde{\omega}_i)$ is convex with Lipschitz envelope L_i. By Assumption 2,

$$\mathbb{E}\,\tilde{\rho}(\theta, \tilde{\omega}_i) = Q(\theta).$$

The simulated population criterion therefore has identified set Θ_0.

Existence of $\hat{\theta}^{N,S}$. Lemma 1 shows that $Q^{N,S}$ is polyhedral convex. Subtracting the simulated shock at the observed bundle gives the corresponding average simulated regret, which is non-negative. Thus $Q^{N,S}$ is bounded below by a finite random constant. By Rockafellar (1970, Cor. 27.3.2), $\arg\min_{\theta\in\mathbb{R}^K} Q^{N,S}(\theta)$ is non-empty for every $N \geq 1$ and every realization.

Corollary 1 (Outer set-consistency)**.** *Under Assumptions 1–2, suppose that Θ_0 is non-empty and compact. Fix $S \geq 1$. For any sequence $(c_N) \subseteq [0, \infty)$ with $c_N \to 0$, the i.i.d. observations and simulation draws above satisfy, with probability one,*

$$\sup_{\theta\in\hat{\Theta}^{N,S}_{c_N}} \operatorname{dist}(\theta, \Theta_0) \xrightarrow[N\to\infty]{} 0,$$

where $\hat{\Theta}^{N,S}_{c_N} := \{\theta \in \mathbb{R}^K : Q^{N,S}(\theta) \leq \inf Q^{N,S} + c_N\}$.

Proof. Lemmas 5–7 and the moment bound for the simulated criterion establish (A1)–(A4) with $p = 1$. Existence of an exact sample minimizer gives (A5), so Theorem 4 gives the stated convergence. □

Corollary 2 (Hausdorff set-consistency)**.** *Suppose that Assumptions 1–2 hold, Θ_0 is non-empty and compact, and $\mathbb{E}\|\phi(d; x_i)\|^2 < \infty$ and $\mathbb{E}|\varepsilon(d; \nu_i)|^2 < \infty$ for every $d \in \{0, 1\}^M$. Fix $S \geq 1$. If $(c_N) \subseteq (0, \infty)$ satisfies $c_N \to 0$ and $c_N\sqrt{N/\log\log N} \to \infty$, then, with probability one,*

$$d_H(\hat{\Theta}^{N,S}_{c_N}, \Theta_0) \to 0.$$

Proof. Lemmas 5–7, the simulated moment bound, and existence of an exact sample minimizer establish (A1)–(A5) with $p = 2$. Theorem 6 applies. □

C.3 Preliminaries: Γ-convergence

Throughout this subsection $f^N, f : \mathbb{R}^K \to \mathbb{R}$ are finite convex functions, hence continuous, and we call f *coercive* if $f(\theta) \to \infty$ as $\|\theta\| \to \infty$, equivalently if every sublevel set $\{f \leq t\}$ is bounded.

Definition 1 (Γ-convergence)**.** $f^N \xrightarrow{\Gamma} f$ on $\mathbb{R}^K$ if

- (lim inf) for every $\theta_N \to \theta$, $\liminf_N f^N(\theta_N) \geq f(\theta)$,
- (lim sup) for every $\theta \in \mathbb{R}^K$, there is $\theta_N \to \theta$ with $\limsup_N f^N(\theta_N) \leq f(\theta)$.

Definition 2 (Equi-coercivity)**.** $\{f^N\}_N$ is equi-coercive on $\mathbb{R}^K$ if for every $t \in \mathbb{R}$ there is a compact $K_t \subseteq \mathbb{R}^K$ with $\{f^N \leq t\} \subseteq K_t$ for N large enough.

Theorem 2 (Rockafellar and Wets (1998, Thm. 7.17)). *For finite convex $f^N, f : \mathbb{R}^K \to \mathbb{R}$, the following are equivalent:*

(i) $f^N \xrightarrow{\Gamma} f$ *on* $\mathbb{R}^K$;

(ii) $f^N \to f$ *pointwise on a dense subset of* $\mathbb{R}^K$;

(iii) $f^N \to f$ *uniformly on every compact subset of* $\mathbb{R}^K$.

The classical fundamental theorem of Γ-convergence takes equi-coercivity as a hypothesis. Under convexity, coercivity of the limit is already enough:

Lemma 8 (Rockafellar and Wets (1998, Exer. 7.32(c))). *If $f^N \xrightarrow{\Gamma} f$ on $\mathbb{R}^K$ and f is coercive, then $\{f^N\}$ is equi-coercive.*

Theorem 3 (Fundamental theorem of Γ-convergence, Braides (2006, Thm. 2.10)). *Let $f^N, f : \mathbb{R}^K \to \mathbb{R}$, not necessarily convex, with f continuous and attaining its minimum. Suppose $f^N \xrightarrow{\Gamma} f$ on $\mathbb{R}^K$, $\{f^N\}$ is equi-coercive, and each f^N attains its infimum. Then for every sequence $c_N \geq 0$ with $c_N \to 0$,*

$$\sup_{\theta\,:\,f^N(\theta) \leq \inf f^N + c_N} \operatorname{dist}(\theta, \arg\min f) \xrightarrow[N\to\infty]{} 0.$$

C.4 Convex criterion

Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space, $\rho : \mathbb{R}^K \times \Omega \to \mathbb{R}$ jointly measurable, and

$$Q(\theta) \;:=\; \int_\Omega \rho(\theta, \omega)\, \mathbb{P}(d\omega), \qquad \Theta_0 \;:=\; \arg\min_{\theta \in \mathbb{R}^K} Q(\theta).$$

Assume:

(A1) $\theta \mapsto \rho(\theta, \omega)$ is convex on $\mathbb{R}^K$ for every $\omega \in \Omega$.

(A2) there exists a random variable $C : \Omega \to [0, \infty)$ such that

$$|\rho(\theta_1, \omega) - \rho(\theta_2, \omega)| \;\leq\; C(\omega)\, \|\theta_1 - \theta_2\| \qquad \text{for every } \theta_1, \theta_2 \in \mathbb{R}^K \text{ and } \omega \in \Omega.$$

(A3) $C \in \mathcal{L}^p(\mathbb{P})$ and $\rho(\theta_0, \cdot) \in \mathcal{L}^p(\mathbb{P})$ for some $\theta_0 \in \mathbb{R}^K$ and some $p \in [1, \infty)$.

(A4) Θ_0 is non-empty and compact.

(A5) $Q^N(\cdot, \bar{\omega})$ attains its infimum for every N and every $\bar{\omega} \in \Omega^\infty$.

Under (A1)–(A3), the Lipschitz bound $|\rho(\theta, \omega)| \leq |\rho(\theta_0, \omega)| + C(\omega)\|\theta - \theta_0\|$ gives $\rho(\theta, \cdot) \in \mathcal{L}^p(\mathbb{P})$ for every $\theta \in \mathbb{R}^K$, so Q is convex and finite-valued on $\mathbb{R}^K$. For finite convex Q, (A4) is equivalent to Q being coercive.

Define for $\theta \in \mathbb{R}^K$ and $\bar{\omega} = (\omega_i)_{i \geq 1} \in \Omega^\infty$

$$Q^N(\theta, \bar{\omega}) \;:=\; \frac{1}{N} \sum_{i \in [N]} \rho(\theta, \omega_i), \qquad \inf Q^N(\bar{\omega}) \;:=\; \inf_{\theta \in \mathbb{R}^K} Q^N(\theta, \bar{\omega})$$

$$\hat{\Theta}_c^N(\bar{\omega}) \;:=\; \{\theta \in \mathbb{R}^K : Q^N(\theta,\bar{\omega}) \le \inf Q^N(\bar{\omega}) + c\} \quad (c \ge 0),$$

and, when (A5) holds, write $\hat{\Theta}^N := \hat{\Theta}_0^N = \arg\min Q^N$.

For each $N \in \mathbb{N}$, Q^N is $\otimes^\infty \mathcal{A}$-measurable in $\bar{\omega}$ at each θ and convex in θ at each $\bar{\omega}$, i.e. a Carathéodory function on $\Omega^\infty \times \mathbb{R}^K$ (Rockafellar and Wets, 1998, Example 14.29).

Lemma 9. *For each $N \in \mathbb{N}$ and $c \ge 0$ the multifunction*

$$\bar{\omega} \rightrightarrows \hat{\Theta}_c^N(\bar{\omega})$$

is closed-valued and $\otimes^\infty \mathcal{A}$-measurable. Consequently, for every continuous $g : \mathbb{R}^K \to \mathbb{R}$, both

$$\bar{\omega} \;\longmapsto\; \sup_{\theta \in \hat{\Theta}_c^N(\bar{\omega})} g(\theta) \qquad \text{and} \qquad \bar{\omega} \;\longmapsto\; \inf_{\theta \in \hat{\Theta}_c^N(\bar{\omega})} g(\theta)$$

are measurable as extended-real-valued functions on Ω^∞.

Proof. Q^N is a Carathéodory function on $\Omega^\infty \times \mathbb{R}^K$, hence a normal integrand. By Rockafellar and Wets (1998, Theorem 14.37), $\inf Q^N$ is measurable. The function

$$h(\theta,\bar{\omega}) \;:=\; \max\big(Q^N(\theta,\bar{\omega}) - \inf Q^N(\bar{\omega}) - c,\; 0\big)$$

is again Carathéodory, and (A5) implies $\arg\min_{\theta\in\mathbb{R}^K} h(\theta,\bar{\omega}) = \hat{\Theta}_c^N(\bar{\omega})$. A second application of Theorem 14.37 shows that $\hat{\Theta}_c^N$ is closed-valued and measurable.

For continuous g and $\alpha \in \mathbb{R}$, both $g^{-1}((\alpha,\infty))$ and $g^{-1}((-\infty,\alpha))$ are open subsets of $\mathbb{R}^K$, so by the definition of multifunction measurability,

$$\begin{aligned}
\Big\{ \sup_{\theta\in\hat{\Theta}_c^N} g(\theta) > \alpha \Big\} \;&=\; (\hat{\Theta}_c^N)^{-1}\big(g^{-1}((\alpha,\infty))\big),\\
\Big\{ \inf_{\theta\in\hat{\Theta}_c^N} g(\theta) < \alpha \Big\} \;&=\; (\hat{\Theta}_c^N)^{-1}\big(g^{-1}((-\infty,\alpha))\big),
\end{aligned}$$

are measurable. □

Lemma 10. $\{\bar{\omega} \in \Omega^\infty : Q^N(\cdot,\bar{\omega}) \xrightarrow{\Gamma} Q\}$ *is $\otimes^\infty \mathcal{A}$-measurable.*

Proof. By the equivalence (i) ⇔ (iii) of Theorem 2, Γ-convergence of $\{Q^N\}$ to Q is equivalent to uniform convergence on every compact set, equivalently on every closed ball $\overline{B}(0,R)$, $R \in \mathbb{N}$. Hence pathwise

$$\{Q^N \xrightarrow{\Gamma} Q\} \;=\; \bigcap_{R\in\mathbb{N}} \Big\{ \sup_{\theta\in\overline{B}(0,R)} |Q^N(\theta) - Q(\theta)| \to 0 \Big\}.$$

For each $R \in \mathbb{N}$, $\bar{\omega} \mapsto \sup_{\theta\in\overline{B}(0,R)} |Q^N(\theta,\bar{\omega}) - Q(\theta)|$ is a random variable (supremum over a deterministic compact set of a Carathéodory function), so the corresponding convergence event is measurable. □

C.5 Outer set-consistency

Theorem 4 (Outer set-consistency). *Assume (A1)–(A5) with $p = 1$. Then, for each non-negative sequence (c_N) with $c_N \to 0$,*

$$\otimes^{\infty}\mathbb{P}\left\{ \sup_{\theta\in\hat{\Theta}^N_{c_N}} \operatorname{dist}(\theta,\Theta_0) \xrightarrow[N\to\infty]{} 0\right\} = 1.$$

Proof. By (A3), $\rho(\theta,\cdot) \in \mathcal{L}^1(\mathbb{P})$ for every $\theta \in \mathbb{R}^K$. The strong law therefore gives $Q^N(\theta) \to Q(\theta)$ almost surely for each θ. Taking a countable intersection over $\mathbb{Q}^K$ yields pointwise convergence on this dense set with probability one. By Theorem 2, $Q^N \xrightarrow{\Gamma} Q$ on this event.

Assumption (A4) makes Q coercive, and Lemma 8 then makes $\{Q^N\}$ equi-coercive. Theorem 3 gives

$$\sup_{\theta\in\hat{\Theta}^N_{c_N}} \operatorname{dist}(\theta,\Theta_0) \to 0.$$

Lemmas 10 and 9 ensure that the events used above are measurable. □

C.6 Hausdorff set-consistency of the level-set estimator

Theorem 5 (Bounded LIL for Lipschitz parametric classes). *Let Θ be a compact subset of $\mathbb{R}^K$ and let $\rho : \Theta \times \Omega \to \mathbb{R}$ be jointly measurable. Suppose there exists $\alpha \in \mathcal{L}^2(\mathbb{P})$ such that*

$$|\rho(\theta_1,\omega) - \rho(\theta_2,\omega)| \;\le\; \alpha(\omega)\,\|\theta_1-\theta_2\| \qquad \text{for every } \theta_1,\theta_2\in\Theta \text{ and } \omega\in\Omega,$$

and $\rho(\bar\theta,\cdot) \in \mathcal{L}^2(\mathbb{P})$ for some $\bar\theta \in \Theta$. Then for $(\omega_i)_{i\ge1}$ i.i.d. from $\mathbb{P}$,

$$\limsup_{N\to\infty} \sqrt{\frac{N}{2\log\log N}}\; \sup_{\theta\in\Theta} \left| \frac{1}{N}\sum_{i=1}^{N} \rho(\theta,\omega_i) \;-\; \int \rho(\theta,\omega)\,\mathbb{P}(d\omega)\right| \;<\; \infty \quad \otimes^{\infty}\mathbb{P}\text{-a.s.}$$

Proof. Apply Andersen et al. (1988, Theorem 5.2) to the centered process

$$\rho^{\circ}(\theta,\omega) := \rho(\theta,\omega) - \mathbb{E}\rho(\theta,\omega).$$

If

$$F_0(\omega) := |\rho(\bar\theta,\omega)| + \alpha(\omega)\operatorname{diam}(\Theta),$$

then $\sup_{\theta\in\Theta}|\rho(\theta,\omega)| \le F_0(\omega)$ and $\sup_{\theta\in\Theta}|\rho^{\circ}(\theta,\omega)| \le F_0(\omega) + \mathbb{E}F_0$. This envelope is in $\mathcal{L}^2(\mathbb{P})$. Moreover,

$$\|\rho^{\circ}(\theta_1,\cdot) - \rho^{\circ}(\theta_2,\cdot)\|_{\mathcal{L}^2} \le 2\|\alpha\|_{\mathcal{L}^2}\|\theta_1-\theta_2\|,$$

and pointwise increments are bounded by $(\alpha + \mathbb{E}\alpha)\|\theta_1-\theta_2\|$. The Gaussian process $\theta \mapsto a g^{\top}\theta$, with $g \sim \mathcal{N}(0, I_K)$ and $a \ge \max\{2\|\alpha\|_{\mathcal{L}^2}, 1\}$, dominates both the canonical metric of ρ° and the Euclidean metric on Θ. Its paths are bounded and uniformly continuous on Θ, while $\alpha + \mathbb{E}\alpha \in \mathcal{L}^2(\mathbb{P})$. These are the envelope, comparison, and increment conditions of Theorem 5.2. The resulting bounded LIL for the partial sums of ρ° is the displayed conclusion after division by N. □

Theorem 6 (Hausdorff set-consistency of the level-set estimator). *Assume (A1)–(A5) with $p = 2$, and let $(c_N) \subseteq (0, \infty)$ satisfy $c_N \to 0$ and $c_N \sqrt{N / \log\log N} \to \infty$. Then*

$$\otimes^{\infty}\mathbb{P}\left\{d_H(\hat{\Theta}^N_{c_N}, \Theta_0) \to 0\right\} \;=\; 1.$$

Proof. Theorem 4 gives

$$\sup_{\theta \in \hat{\Theta}^N_{c_N}} \operatorname{dist}(\theta, \Theta_0) \to 0 \qquad \otimes^{\infty}\mathbb{P}\text{-a.s.} \tag{19}$$

It remains to prove the reverse inclusion. Set $C_0 := \{\theta \in \mathbb{R}^K : \operatorname{dist}(\theta, \Theta_0) \le 1\}$ and define

$$\Delta^N := \sup_{\theta \in C_0} |Q^N(\theta) - Q(\theta)|, \qquad \sigma^N := \sup_{\theta \in \hat{\Theta}^N} \operatorname{dist}(\theta, \Theta_0).$$

The Carathéodory property and compactness of C_0 make Δ^N measurable, while Lemma 9 gives measurability of σ^N. The same lemma, together with compactness of Θ_0, also makes $d_H(\hat{\Theta}^N_{c_N}, \Theta_0)$ measurable.

Theorem 4, applied with $c_N = 0$, gives $\sigma^N \to 0$ almost surely. Theorem 5, applied on C_0, and the assumed rate for c_N give

$$\frac{2\Delta^N}{c_N} \to 0 \qquad \otimes^{\infty}\mathbb{P}\text{-a.s.}$$

On the event $\{\sigma^N \le 1\}$, every minimizer of Q^N belongs to C_0. Thus, for every $\theta \in \Theta_0$,

$$\begin{aligned} Q^N(\theta) - \inf Q^N &= Q^N(\theta) - \min_{C_0} Q^N \\ &\le Q^N(\theta) - \min_{C_0} Q + \Delta^N \\ &= Q^N(\theta) - Q(\theta) + \Delta^N \\ &\le 2\Delta^N. \end{aligned}$$

Here $|\min_{C_0} Q^N - \min_{C_0} Q| \le \Delta^N$, $\min_{C_0} Q = \min Q$, and $Q(\theta) = \min Q$ on Θ_0. Almost surely, $\sigma^N \le 1$ and $2\Delta^N \le c_N$ for all sufficiently large N. The bound above then implies $\Theta_0 \subseteq \hat{\Theta}^N_{c_N}$ eventually. Combining this inclusion with (19) proves Hausdorff convergence. □

C.7 Asymptotic normality (point-identified case)

Lemma 11 (Adjacency of the choice cells). *Let C be a nonempty subset of $\{0,1\}^M$, let $\phi_d \in \mathbb{R}^K$ for $d \in C$, and fix $\theta_0 \in \mathbb{R}^K$. For $d \in C$, define*

$$\mathcal{A}_d := \left\{\nu \in \mathbb{R}^M : \theta_0^\top \phi_d + \nu^\top d > \theta_0^\top \phi_e + \nu^\top e \text{ for every } e \in C \setminus \{d\}\right\}.$$

Every $\mathcal{A}_d$ is nonempty. Moreover, the graph on C that joins d and e when $\overline{\mathcal{A}_d}$ and $\overline{\mathcal{A}_e}$ share an $(M-1)$-dimensional facet is connected.

Proof. For fixed d, the construction $\nu = \lambda(2d - \mathbf{1}_M)$ used in the proof of Proposition 2 shows that $\nu \in \mathcal{A}_d$ for all sufficiently large λ.

For three distinct bundles d_1, d_2, d_3, the vectors $d_1 - d_2$ and $d_1 - d_3$ are linearly independent. Otherwise one would be a nonzero scalar multiple of the other. Since both vectors have entries in $\{-1, 0, 1\}$, that multiple must be 1 or -1. The first case gives $d_2 = d_3$. The second gives $2d_1 = d_2 + d_3$, which is impossible for distinct binary vectors. Therefore

$$\mathcal{T}(d_1, d_2, d_3) := \left\{\nu : \theta_0^\top \phi_{d_1} + \nu^\top d_1 = \theta_0^\top \phi_{d_2} + \nu^\top d_2 = \theta_0^\top \phi_{d_3} + \nu^\top d_3\right\}$$

is either empty or an affine subspace of codimension two.

Fix distinct $d, e \in C$. The union of the pairwise tie hyperplanes is Lebesgue null, so choose $\nu_d^\circ \in \mathcal{A}_d$ outside this union. For each nonempty $\mathcal{T}(d_1, d_2, d_3)$,

$$\dim \operatorname{span}\big(\mathcal{T}(d_1, d_2, d_3) - \nu_d^\circ\big) \le M - 1.$$

The union of these finitely many proper linear subspaces is Lebesgue null. Since $\mathcal{A}_e - \nu_d^\circ$ is nonempty and open, there is $\nu_e^\circ \in \mathcal{A}_e$ such that

$$\nu_e^\circ - \nu_d^\circ \notin \bigcup_{\substack{d_1,d_2,d_3 \in C \text{ distinct} \\ \mathcal{T}(d_1,d_2,d_3) \neq \varnothing}} \operatorname{span}\big(\mathcal{T}(d_1, d_2, d_3) - \nu_d^\circ\big).$$

If the segment $[\nu_d^\circ, \nu_e^\circ]$ met a nonempty $\mathcal{T}(d_1, d_2, d_3)$ at $\bar{\nu}$, then $\bar{\nu} \neq \nu_d^\circ$ and

$$\nu_e^\circ - \nu_d^\circ \in \operatorname{span}(\bar{\nu} - \nu_d^\circ) \subseteq \operatorname{span}\big(\mathcal{T}(d_1, d_2, d_3) - \nu_d^\circ\big),$$

contrary to the choice of ν_e°. Thus the segment meets no triple-tie set. Its supporting line is contained in no pairwise tie hyperplane because ν_d° lies on none of them, and it therefore meets each such hyperplane at most once. Hence the maximizer changes only finitely many times. At a change point $\bar{\nu}$, exactly two bundles, say d_1 and d_2, attain the maximum. For every $d_3 \notin \{d_1, d_2\}$,

$$\theta_0^\top \phi_{d_1} + \bar{\nu}^\top d_1 = \theta_0^\top \phi_{d_2} + \bar{\nu}^\top d_2 > \theta_0^\top \phi_{d_3} + \bar{\nu}^\top d_3.$$

These strict inequalities persist on a relative neighborhood in the hyperplane where the utilities of d_1 and d_2 are equal. Hence $\overline{\mathcal{A}_{d_1}}$ and $\overline{\mathcal{A}_{d_2}}$ share a facet. The successive maximizers along the segment form a path from d to e. □

Lemma 12 (Conditional curvature)**.** *Let C be a nonempty subset of $\{0,1\}^M$, let $\phi_d \in \mathbb{R}^K$ for $d \in C$, and let ν be an $\mathbb{R}^M$-valued random vector with probability density p and $\mathbb{E}\|\nu\| < \infty$. Define*

$$\mathcal{V}_C(\theta) := \int_{\mathbb{R}^M} \max_{d \in C} \{\theta^\top \phi_d + \nu^\top d\} p(\nu)\, d\nu.$$

Suppose p is continuous and strictly positive and, for some $A > 0$,

$$p(\nu) \le A(1 + \|\nu\|)^{-M}.$$

Then, for every θ_0, there is a symmetric positive-semidefinite matrix H_C such that

$$\mathcal{V}_C(\theta_0 + au) - \mathcal{V}_C(\theta_0) - a\nabla\mathcal{V}_C(\theta_0)^\top u = \frac{a^2}{2} u^\top H_C u + o(a^2)$$

for every $u \in \mathbb{R}^K$. There is a constant B, depending only on (A, M), such that, for every $a > 0$,

$$0 \le \frac{\mathcal{V}_C(\theta_0 + au) - \mathcal{V}_C(\theta_0) - a\nabla\mathcal{V}_C(\theta_0)^\top u}{a^2} \le B\|u\|^2 \sum_{d,e\in C} \|\phi_d - \phi_e\|^2, \tag{20}$$

$$\|H_C\| \le B \sum_{d,e\in C} \|\phi_d - \phi_e\|^2. \tag{21}$$

Moreover, $u^\top H_C u > 0$ whenever $d \mapsto u^\top \phi_d$ is not constant on C.

Proof. *Differentiability.* The moment assumption makes $\mathcal{V}_C(\theta)$ finite for every θ. Let $\mathcal{A}_d$ be the strict choice cell of d at θ_0. The tie sets are finite unions of Lebesgue-null hyperplanes, so the maximizer is unique almost surely. Denote it by $d^\star(\nu)$. Fix $u \in \mathbb{R}^K$. Since C is finite, the strict inequalities that select $d^\star(\nu)$ at θ_0 remain strict at $\theta_0 + au$ for all sufficiently small $|a|$. Hence $d^\star(\nu)$ remains the unique maximizer. For $a \neq 0$,

$$\left| \frac{\max_{d\in C}\{(\theta_0 + au)^\top \phi_d + \nu^\top d\} - \max_{d\in C}\{\theta_0^\top \phi_d + \nu^\top d\}}{a} \right| \le \max_{d\in C} |u^\top \phi_d|.$$

The quotient equals $u^\top \phi_{d^\star(\nu)}$ for all sufficiently small $|a|$. Dominated convergence gives

$$u^\top \nabla\mathcal{V}_C(\theta_0) = \sum_{d\in C} \mathbb{P}(\mathcal{A}_d) u^\top \phi_d, \qquad \nabla\mathcal{V}_C(\theta_0) = \sum_{d\in C} \mathbb{P}(\mathcal{A}_d)\phi_d.$$

One-gap probabilities. For $d, e \in C$, set

$$\mathrm{gap}_{de}(\nu) := \theta_0^\top(\phi_d - \phi_e) + \nu^\top(d - e).$$

For $d \neq e$ and $z \in (d-e)^\perp$, define

$$\nu_{de}(r, z) := \frac{r - \theta_0^\top(\phi_d - \phi_e)}{\|d - e\|^2}(d - e) + z.$$

Then $\mathrm{gap}_{de}(\nu_{de}(r,z)) = r$ and the Jacobian is $\|d-e\|^{-1}$. Hence, for $s > 0$,

$$P_{de}(s) := \mathbb{P}\{\nu \in \mathcal{A}_d,\ 0 < \mathrm{gap}_{de}(\nu) \le s\} = \int_0^s p_{de}(r)\, dr, \tag{22}$$

$$p_{de}(r) := \frac{1}{\|d-e\|} \int_{(d-e)^\perp} \mathbf{1}\Big\{\mathrm{gap}_{df}(\nu_{de}(r,z)) > 0 \text{ for all } f \notin \{d,e\}\Big\}\, p(\nu_{de}(r,z))\, dz.$$

Since $\|\nu_{de}(r,z)\| \ge \|z\|$ and $\|d-e\| \ge 1$,

$$0 \le p_{de}(r) \le B_{\mathrm{slice}} := A \int_{\mathbb{R}^{M-1}} (1 + \|z\|)^{-M}\, dz < \infty. \tag{23}$$

For $f \notin \{d, e\}$, the linear part of the affine map

$$z \longmapsto \mathrm{gap}_{df}(\nu_{de}(0, z))$$

is $z \mapsto z^\top(d - f)$. If it vanished on $(d - e)^\perp$, then $d - f \in \mathrm{span}(d - e)$, contradicting the linear independence proved in Lemma 11. Its zero set therefore has measure zero. For almost every z, the indicator in (22) converges to its value at $r = 0$, and continuity of p gives $p(\nu_{de}(r, z)) \to p(\nu_{de}(0, z))$. Since the integrand is bounded by $A(1 + \|z\|)^{-M}$, dominated convergence gives

$$\lim_{r \to 0} p_{de}(r) = p_{de}(0), \qquad \lim_{s \downarrow 0} \frac{P_{de}(s)}{s} = \lim_{s \downarrow 0} \frac{1}{s} \int_0^s p_{de}(r)\, dr = p_{de}(0).$$

Since $e - d = -(d - e)$ and $\theta_0^\top(\phi_e - \phi_d) = -\theta_0^\top(\phi_d - \phi_e)$,

$$\nu_{ed}(0, z) = \nu_{de}(0, z) \qquad (z \in (d - e)^\perp).$$

Moreover, whenever $\mathrm{gap}_{de}(\nu) = 0$, $\mathrm{gap}_{df}(\nu) = \mathrm{gap}_{ef}(\nu)$ for every $f \notin \{d, e\}$. Hence the integrals defining $p_{de}(0)$ and $p_{ed}(0)$ are equal, and

$$p_{de}(0) = p_{ed}(0), \qquad 0 \le p_{de}(0) \le B_{\mathrm{slice}}. \tag{24}$$

Simultaneous gaps. For distinct $d, e, f \in C$, $M \ge 2$. Let $\mathsf{G}_{d,e,f}$ be the matrix with rows $(d - e)^\top$ and $(d - f)^\top$. It has rank two. For $y \in \mathbb{R}^2$, define

$$\nu^\perp_{d,e,f}(y) := \mathsf{G}^\top_{d,e,f}(\mathsf{G}_{d,e,f}\mathsf{G}^\top_{d,e,f})^{-1} \left[y - \begin{pmatrix} \theta_0^\top(\phi_d - \phi_e) \\ \theta_0^\top(\phi_d - \phi_f) \end{pmatrix} \right].$$

The fiber on which the two gaps equal y is $\nu^\perp_{d,e,f}(y) + \ker \mathsf{G}_{d,e,f}$, and $\nu^\perp_{d,e,f}(y) \perp \ker \mathsf{G}_{d,e,f}$. The coarea formula therefore gives

$$\begin{aligned} &\mathbb{P}\{0 < \mathrm{gap}_{de}(\nu) \le s,\ 0 < \mathrm{gap}_{df}(\nu) \le s\} \\ &\quad = \frac{1}{\sqrt{\det(\mathsf{G}_{d,e,f}\mathsf{G}^\top_{d,e,f})}} \int_{(0,s]^2} \int_{\ker \mathsf{G}_{d,e,f}} p(\nu^\perp_{d,e,f}(y) + z)\, dz\, dy \\ &\quad \le \frac{As^2}{\sqrt{\det(\mathsf{G}_{d,e,f}\mathsf{G}^\top_{d,e,f})}} \int_{\mathbb{R}^{M-2}} (1 + \|z\|)^{-M}\, dz \le B_{\mathrm{joint}} s^2, \end{aligned} \tag{25}$$

The last inequality holds because $\det(\mathsf{G}_{d,e,f}\mathsf{G}^\top_{d,e,f})$ is a positive integer. Thus B_{joint} is finite and depends only on (A, M). If no such triple exists, set $B_{\mathrm{joint}} = 0$.

Quadratic expansion. It suffices to consider $a \downarrow 0$, since negative a are covered by replacing u with $-u$. Fix u and define, for $d, e \in C$,

$$\mathrm{gain}_{d \to e}(u) := u^\top(\phi_e - \phi_d).$$

Since the strict cells cover $\mathbb{R}^M$ up to a null set, for $a > 0$,

$$\begin{aligned}\mathcal{R}_C(a,u) &:= \mathcal{V}_C(\theta_0 + au) - \mathcal{V}_C(\theta_0) - a\nabla\mathcal{V}_C(\theta_0)^\top u \\ &= \sum_{d\in C}\int_{\mathcal{A}_d} \max_{e\in C}\{a\,\mathrm{gain}_{d\to e}(u) - \mathrm{gap}_{de}(\nu)\}_+ p(\nu)\,d\nu.\end{aligned} \tag{26}$$

For an ordered pair (d,e) with $d \neq e$, the layer-cake identity and the substitution $s = ar$ give

$$\begin{aligned}\frac{1}{a^2}\int_{\mathcal{A}_d}\{a\,\mathrm{gain}_{d\to e}(u) - \mathrm{gap}_{de}(\nu)\}_+ p(\nu)\,d\nu &= \frac{1}{a^2}\int_0^{a\{\mathrm{gain}_{d\to e}(u)\}_+} P_{de}(s)\,ds \\ &= \int_0^{\{\mathrm{gain}_{d\to e}(u)\}_+} \frac{P_{de}(ar)}{a}\,dr \\ &\longrightarrow \frac{1}{2}p_{de}(0)\{\mathrm{gain}_{d\to e}(u)\}_+^2.\end{aligned} \tag{27}$$

By (22)–(23), the integrand is bounded by $B_{\mathrm{slice}}r$.

Define

$$\mathrm{gain}_{\max}(u) := \max_{d,e\in C}|\,\mathrm{gain}_{d\to e}(u)|.$$

For every $d \in C$ and $\nu \in \mathcal{A}_d$,

$$\begin{aligned}0 &\le \sum_{e\in C}\{a\,\mathrm{gain}_{d\to e}(u) - \mathrm{gap}_{de}(\nu)\}_+ - \max_{e\in C}\{a\,\mathrm{gain}_{d\to e}(u) - \mathrm{gap}_{de}(\nu)\}_+ \\ &\le \sum_{\substack{\{e,f\}\subseteq C\\ e\neq f}} \min\Big\{\{a\,\mathrm{gain}_{d\to e}(u) - \mathrm{gap}_{de}(\nu)\}_+, \{a\,\mathrm{gain}_{d\to f}(u) - \mathrm{gap}_{df}(\nu)\}_+\Big\}.\end{aligned}$$

Each nonzero minimum on the right is at most $a\,\mathrm{gain}_{\max}(u)$ and can occur only when

$$0 < \mathrm{gap}_{de}(\nu) \le a\,\mathrm{gain}_{\max}(u) \quad \text{and} \quad 0 < \mathrm{gap}_{df}(\nu) \le a\,\mathrm{gain}_{\max}(u).$$

Therefore (25) implies

$$\begin{aligned}0 &\le \sum_{d\in C}\int_{\mathcal{A}_d}\left[\sum_{e\in C}\{a\,\mathrm{gain}_{d\to e}(u) - \mathrm{gap}_{de}(\nu)\}_+ - \max_{e\in C}\{a\,\mathrm{gain}_{d\to e}(u) - \mathrm{gap}_{de}(\nu)\}_+\right] p(\nu)\,d\nu \\ &\le |C|\binom{|C|}{2} B_{\mathrm{joint}}a^3\,\mathrm{gain}_{\max}(u)^3 = o(a^2).\end{aligned}$$

Combining this bound with (24) and (27) gives

$$\begin{aligned}\lim_{a\downarrow 0}\frac{\mathcal{R}_C(a,u)}{a^2} &= \frac{1}{2}\sum_{\substack{d,e\in C\\ d\neq e}} p_{de}(0)\{\mathrm{gain}_{d\to e}(u)\}_+^2 \\ &= \frac{1}{2}\sum_{\substack{\{d,e\}\subseteq C\\ d\neq e}} p_{de}(0)\{u^\top(\phi_d - \phi_e)\}^2 = \frac{1}{2}u^\top H_C u,\end{aligned}$$

where

$$H_C := \sum_{\substack{\{d,e\}\subseteq C \\ d\neq e}} p_{de}(0)(\phi_d - \phi_e)(\phi_d - \phi_e)^\top.$$

Bounds and strict curvature. Since the maximum of nonnegative terms is bounded by their sum,

$$\begin{aligned} 0 \le \mathcal{R}_C(a,u) &\le \sum_{\substack{d,e\in C \\ d\neq e}} \int_{\mathcal{A}_d} \{a\,\mathrm{gain}_{d\to e}(u) - \mathrm{gap}_{de}(\nu)\}_+ p(\nu)\,d\nu \\ &= \sum_{\substack{d,e\in C \\ d\neq e}} \int_0^{a\{\mathrm{gain}_{d\to e}(u)\}_+} P_{de}(s)\,ds \\ &\le \frac{B_{\mathrm{slice}}a^2}{2} \sum_{\substack{d,e\in C \\ d\neq e}} \{\mathrm{gain}_{d\to e}(u)\}_+^2 \\ &\le \frac{B_{\mathrm{slice}}a^2}{2}\|u\|^2 \sum_{d,e\in C} \|\phi_d - \phi_e\|^2. \end{aligned}$$

Also, by (24),

$$\|H_C\| \le \sum_{\substack{\{d,e\}\subseteq C \\ d\neq e}} p_{de}(0)\|\phi_d - \phi_e\|^2 \le \frac{B_{\mathrm{slice}}}{2} \sum_{d,e\in C} \|\phi_d - \phi_e\|^2.$$

These two bounds prove (20)–(21). If the cells of d and e share a facet, its relative interior contains a set of the form $\{\nu_{de}(0,z) : z \in U\}$ for some nonempty open $U \subset (d-e)^\perp$. Hence

$$p_{de}(0) \ge \frac{1}{\|d-e\|} \int_U p(\nu_{de}(0,z))\,dz > 0.$$

Consequently,

$$u^\top H_C u = \sum_{\substack{\{d,e\}\subseteq C \\ d\neq e}} p_{de}(0)\{u^\top(\phi_d - \phi_e)\}^2.$$

If this expression is zero, then

$$u^\top \phi_d = u^\top \phi_e \quad \text{whenever the cells of } d \text{ and } e \text{ share a facet.}$$

The graph in Lemma 11 is connected. Hence $d \mapsto u^\top \phi_d$ is constant on C. □

For the simulated loss $\tilde{\rho}$, write

$$\dot{\tilde{\rho}}_i := \nabla_\theta \tilde{\rho}(\theta_0, \tilde{\omega}_i)$$

whenever the derivative exists.

Theorem 7 (Asymptotic normality). *Under the conditions of Proposition 2, suppose in addition that $\mathbb{P}_{\nu|x}(\cdot \mid x)$ admits a jointly measurable density with respect to Lebesgue measure that is continuous and strictly positive in ν. Suppose there is a constant $A > 0$ such that, for $\mathbb{P}_x$-almost every x, this density is bounded above by $A(1+\|\nu\|)^{-M}$ for every $\nu \in \mathbb{R}^M$, and suppose that $\mathbb{E}[\max_{d\in C_i} \|\phi(d;x_i)\|^2] < \infty$. For fixed S, let $\hat{\theta}^{N,S}$ be any measurable exact minimizer of $Q^{N,S}$. Then $\dot{\tilde{\rho}}_i$ exists almost surely, and there*

is a symmetric positive-definite matrix H such that

$$Q(\theta_0 + t) - Q(\theta_0) = \frac{1}{2} t^\top H t + o(\|t\|^2)$$

and

$$\sqrt{N}(\hat{\theta}^{N,S} - \theta_0) = -H^{-1} \frac{1}{\sqrt{N}} \sum_{i=1}^{N} \dot{\rho}_i + o_p(1).$$

Consequently,

$$\sqrt{N}(\hat{\theta}^{N,S} - \theta_0) \Longrightarrow \mathcal{N}\left(0, H^{-1} \Sigma_S H^{-1}\right),$$

where

$$\Sigma_S := \mathrm{Var}(\dot{\rho}_i) = \Sigma_{\text{data}} + \frac{1}{S} \Sigma_{\text{sim}},$$

with

$$\Sigma_{\text{data}} := \mathrm{Var}(\mathbb{E}[\nabla_\theta \rho(\theta_0, x_i, \nu, d_i) \mid x_i, d_i])$$

and

$$\Sigma_{\text{sim}} := \mathbb{E}[\mathrm{Var}(\nabla_\theta \rho(\theta_0, x_i, \nu, d_i) \mid x_i, d_i)].$$

Proof. *Score.* For fixed x, ties lie on finitely many affine hyperplanes and therefore have conditional probability zero. Differentiating at the almost-surely unique maximizer and averaging over s gives

$$\dot{\rho}_i = \frac{1}{S} \sum_{s=1}^{S} \left[\phi\big(d(x_i, \nu_{si}, \theta_0); x_i\big) - \phi(d_i; x_i)\right].$$

Lemma 5 gives $\|\dot{\rho}_i\| \le L_i$, and the feature moment condition gives $L_i \in \mathcal{L}^2$. Thus $\dot{\rho}_i \in \mathcal{L}^2$. By (1) and Assumption 2, each simulated choice has the same conditional distribution given x_i as the observed choice. Therefore

$$\mathbb{E}[\dot{\rho}_i \mid x_i] = \mathbb{E}\left[\phi\big(d(x_i, \nu, \theta_0); x_i\big) \mid x_i\right] - \mathbb{E}[\phi(d_i; x_i) \mid x_i] = 0,$$

and hence $\mathbb{E}\dot{\rho}_i = 0$.

Population curvature. For almost every x, define

$$Q_x(\theta) := \int_{\mathbb{R}^M} \max_{d \in C(x)} \{\theta^\top \phi(d; x) + \nu^\top d\} \, \mathbb{P}_{\nu|x}(d\nu \mid x) - \theta^\top \mathbb{E}[\phi(d_i; x_i) \mid x_i = x].$$

Then $Q(\theta) = \mathbb{E}[Q_{x_i}(\theta)]$. Since the error of a singleton bundle is one coordinate of ν_i, Assumption 1 implies $\mathbb{E}[\|\nu_i\| \mid x_i] < \infty$ almost surely. Apply Lemma 12 with $C = C(x)$, $\phi_d = \phi(d; x)$, and p equal to the conditional density, and denote the resulting matrix by $H(x)$. The second term in Q_x is affine, so the lemma gives, for every $u \in \mathbb{R}^K$,

$$Q_x(\theta_0 + au) - Q_x(\theta_0) - a \nabla Q_x(\theta_0)^\top u = \frac{a^2}{2} u^\top H(x) u + o(a^2). \tag{28}$$

The bounds (20) and (21) hold with $\mathcal{V}_C$ and H_C replaced by Q_x and $H(x)$, with the same constant B for almost every x. Measurability of $x \mapsto C(x)$ and $x \mapsto \phi(d; x)$, together with joint

measurability of the conditional density, makes $x \mapsto Q_x(\theta)$ measurable for every θ. Applying (28) to u and $-u$ yields

$$u^\top H(x)u = \lim_{n\to\infty} n^2 \left[Q_x\left(\theta_0 + \frac{u}{n}\right) + Q_x\left(\theta_0 - \frac{u}{n}\right) - 2Q_x(\theta_0) \right].$$

Thus $x \mapsto u^\top H(x)u$ is measurable for every u. If $e_1, \ldots, e_K$ are the coordinate vectors, then

$$\begin{aligned} H_{jj}(x) &= e_j^\top H(x) e_j, \\ H_{jk}(x) &= \frac{1}{2}\left[(e_j + e_k)^\top H(x)(e_j + e_k) - e_j^\top H(x) e_j - e_k^\top H(x) e_k \right], \qquad j \neq k. \end{aligned}$$

Hence every entry of $H(x)$ is measurable. The feature moment condition makes these conditional bounds integrable. Differentiating the finite maximum at its almost-surely unique maximizer gives

$$\begin{aligned} \nabla Q_x(\theta_0) = &\int_{\mathbb{R}^M} \phi(d(x, \nu, \theta_0); x)\, \mathbb{P}_{\nu|x}(d\nu \mid x) \\ &- \mathbb{E}[\phi(d_i; x_i) \mid x_i = x] = \mathbb{E}[\dot{\rho}_i \mid x_i = x]. \end{aligned}$$

Therefore

$$\mathbb{E}[\nabla Q_{x_i}(\theta_0)] = \mathbb{E}[\dot{\rho}_i] = 0,$$

dominated convergence in (28) gives, for every $u \in \mathbb{R}^K$,

$$\lim_{a \downarrow 0} \frac{2\{Q(\theta_0 + au) - Q(\theta_0)\}}{a^2} = u^\top \mathbb{E}[H(x_i)]u.$$

Set

$$H := \mathbb{E}[H(x_i)].$$

For $u \neq 0$, Lemma 12 and the rank condition in Proposition 2 give

$$\mathbb{P}_x\{u^\top H(x_i)u > 0\} \geq \mathbb{P}_x\{u^\top \phi(d; x_i) \text{ is not constant over } d \in C_i\} > 0.$$

Since $u^\top H(x_i)u \geq 0$ almost surely,

$$u^\top H u = \mathbb{E}[u^\top H(x_i)u] > 0.$$

Let (a_n) be an arbitrary positive sequence converging to zero. The finite convex functions

$$u \longmapsto \frac{Q(\theta_0 + a_n u) - Q(\theta_0)}{a_n^2}$$

converge pointwise to $u \mapsto u^\top H u/2$. By Theorem 2, they converge uniformly on the closed unit ball. Since (a_n) was arbitrary,

$$\sup_{\|u\| \leq 1} \left| \frac{Q(\theta_0 + au) - Q(\theta_0)}{a^2} - \frac{1}{2} u^\top H u \right| \longrightarrow 0.$$

Taking $a = \|t\|$ and $u = t/\|t\|$ gives

$$Q(\theta_0 + t) - Q(\theta_0) = \frac{1}{2} t^\top H t + o(\|t\|^2).$$

Local remainder. For $t \in \mathbb{R}^K$, write

$$R_i(t) := \tilde{\rho}(\theta_0 + t, \tilde{\omega}_i) - \tilde{\rho}(\theta_0, \tilde{\omega}_i) - \dot{\tilde{\rho}}_i^\top t.$$

Almost surely, all S simulated maximizers at θ_0 are unique. Since S and C_i are finite, they remain unchanged on a neighborhood of θ_0, so $R_i(t) = 0$ for all sufficiently small $\|t\|$. For arbitrary t, for each s,

$$\begin{aligned} 0 &\le \max_{d \in C_i} \{(\theta_0 + t)^\top \phi(d; x_i) + \nu_{si}^\top d\} - \max_{d \in C_i} \{\theta_0^\top \phi(d; x_i) + \nu_{si}^\top d\} \\ &\quad - t^\top \phi(d(x_i, \nu_{si}, \theta_0); x_i) \\ &\le \max_{d \in C_i} t^\top \{\phi(d; x_i) - \phi(d(x_i, \nu_{si}, \theta_0); x_i)\} \le L_i \|t\|. \end{aligned}$$

Averaging gives

$$0 \le R_i(t) \le L_i \|t\|.$$

Since L_i^2 is integrable and $R_i(t)/\|t\| \to 0$ almost surely, dominated convergence gives

$$\mathbb{E} R_i(t)^2 = o(\|t\|^2).$$

Moreover,

$$\mathbb{E} R_i(t) = Q(\theta_0 + t) - Q(\theta_0) - (\mathbb{E}\dot{\tilde{\rho}}_i)^\top t = Q(\theta_0 + t) - Q(\theta_0) = \frac{1}{2} t^\top H t + o(\|t\|^2)$$

and

$$\mathrm{Var}(R_i(t)) \le \mathbb{E} R_i(t)^2 = o(\|t\|^2).$$

Argmin and covariance. The function $\theta \mapsto \tilde{\rho}(\theta, \tilde{\omega}_i)$ is convex, $\dot{\tilde{\rho}}_i \in \mathcal{L}^2$, and $H \succ 0$. The decomposition defining $R_i(t)$, together with $\mathbb{E} R_i(t) = t^\top H t/2 + o(\|t\|^2)$ and $\mathrm{Var}(R_i(t)) = o(\|t\|^2)$, satisfies the hypotheses of Hjort and Pollard (1993, Theorem 2.1). In the notation of that theorem, its quadratic matrix is H and its score covariance is $\mathrm{Var}(\dot{\tilde{\rho}}_i) = \Sigma_S$. That theorem gives

$$\sqrt{N}(\hat{\theta}^{N,S} - \theta_0) = -H^{-1} \frac{1}{\sqrt{N}} \sum_{i=1}^{N} \dot{\tilde{\rho}}_i + o_p(1).$$

The multivariate central limit theorem gives the stated normal limit with $\Sigma_S = \mathrm{Var}(\dot{\tilde{\rho}}_i)$. Conditional on (x_i, d_i), the summands in

$$\dot{\tilde{\rho}}_i = \frac{1}{S} \sum_{s=1}^{S} \nabla_\theta \rho(\theta_0, x_i, \nu_{si}, d_i)$$

are independent and identically distributed. Therefore

$$\mathbb{E}[\dot{\tilde{\rho}}_i \mid x_i, d_i] = \mathbb{E}[\nabla_\theta \rho(\theta_0, x_i, \nu, d_i) \mid x_i, d_i]$$

and

$$\mathrm{Var}(\dot{\tilde{\rho}}_i \mid x_i, d_i) = \frac{1}{S} \mathrm{Var}(\nabla_\theta \rho(\theta_0, x_i, \nu, d_i) \mid x_i, d_i).$$

The law of total variance therefore gives

$$\begin{aligned} \Sigma_S &= \mathrm{Var}(\mathbb{E}[\dot{\tilde{\rho}}_i \mid x_i, d_i]) + \mathbb{E}[\mathrm{Var}(\dot{\tilde{\rho}}_i \mid x_i, d_i)] \\ &= \Sigma_{\text{data}} + \frac{1}{S}\Sigma_{\text{sim}}. \end{aligned}$$

□