

Maximum Likelihood Estimation of Differentiated Products Demand Systems

Greg Lewis
Microsoft Research

Bora Ozaltun*
Microsoft Research

Georgios Zervas
Boston University

November 23, 2021

Abstract

We discuss estimation of the differentiated products demand system of Berry et al. (1995) (BLP) by maximum likelihood estimation (MLE). We derive the maximum likelihood estimator in the case where prices are endogenously generated by firms that set prices in Bertrand-Nash equilibrium. In Monte Carlo simulations the MLE estimator outperforms the best-practice GMM estimator on both bias and mean squared error when the model is correctly specified. This remains true under some forms of misspecification. In our simulations, the coverage of the ML estimator is close to its nominal level, whereas the GMM estimator tends to under-cover. We conclude the paper by estimating BLP on the car data used in the original Berry et al. (1995) paper, obtaining similar estimates with considerably tighter standard errors.

*Corresponding author. Email: ozaltun@mit.edu. We are grateful to Ariel Pakes for useful conversations early on in this project, and to Surya Ierokomos for excellent research assistance.

1 Introduction

Demand estimation is a fundamental tool in applied economics. One widely used demand model for differentiated goods markets is Berry et al. (1995) (BLP), which builds on the earlier work of McFadden (1986) concerning the mixed logit. The BLP model is traditionally estimated by generalized method of moments (GMM), but the estimates are often imprecise. It has long been known that the estimation performance is dramatically improved by adding a model of the supply side, typically one in which firms engage in Bertrand-Nash pricing. A small but influential literature has documented other ways of improving estimation performance: replacing a step that inverts from market shares to mean utilities with a set of constraints on the difference between actual and predicted market shares (the MPEC approach of Dubé et al. (2012)), using empirical likelihood instead of GMM (Conlon, 2013)), choosing the right instruments (Reynaert and Verboven (2014); Gandhi and Houde (2019)) and using high-quality software (Conlon and Gortmaker (2020)).

This paper adds to this literature by revisiting an old idea: estimating BLP by maximum likelihood. The theoretical advantages of doing so are clear. First, maximum likelihood is statistically efficient, which translates into more precise estimates. Second, maximum likelihood does not require the researcher to choose their instruments. In fact, provided the model is correctly specified, one can think of MLE as a magical procedure that automatically performs GMM with the optimal instruments. This may be important in practice, for two reasons. First, choosing good instruments can be hard, as attested to by the papers on this topic cited above. Second, BLP is often used for demand analysis in merger cases, and in such cases any degrees of freedom, such as instrument choice, can be exploited by competing expert witnesses in their analysis of the case.

There are, however, some drawbacks. Both price and quantity are endogenous, and so a well-specified likelihood requires modeling how both price and quantity arise, necessitating models of both demand and supply. This requires committing to some model of pricing. Moreover, the researcher needs to specify a distribution for the demand and cost shocks. We investigate these trade offs here.

The paper proceeds in three parts. First, we outline the BLP model and derive the maximum likelihood estimator (MLE) under Bertrand-Nash pricing. Somewhat surprisingly, this appears to be new, as the prior literature we review below has assumed that the demand error enters the pricing equation linearly, which is inconsistent with Bertrand-Nash.

The second part of the paper investigates the performance of the MLE relative to the best practice in implementing GMM (for best practice, we follow Conlon and Gortmaker (2020)). We start with a variety of well-specified scenarios, in which we would expect MLE

to outperform GMM. We find that it does, and quite handily - bias and mean squared error are smaller, standard errors are tighter, and coverage is more precise. We then consider three scenarios in which the model is wrong in some way. The first tests the robustness of MLE to a different error structure, namely Laplace errors. We find that the MLE continues to outperform the GMM benchmark. The second considers a misspecified supply side, where estimation assumes costs are log-linear rather than linear in characteristics. Here we find that both GMM and MLE perform poorly in estimating the demand parameters, as expected, but MLE performs worse. Finally, we consider a case where the ownership matrix is misspecified, which we intend as reduced form shorthand for a mis-specification of the game being played on the supply side. The results of GMM and MLE are comparable in this case.

In the final part of the paper, we replicate BLP, using their original data, and building on the replication exercise done in Andrews et al. (2017) (AGS). We compare the MLE estimates to a series of benchmarks: the original estimates in BLP, the replication results of AGS, and the best practice estimates of Conlon and Gortmaker (2020) (CG). The MLE estimates are very similar to the GMM estimates both in terms of the parameters and implied own price elasticities. Notably, the price coefficient estimate is close to the best practice of CG, unlike BLP and AGS. In addition the standard errors are substantially smaller; often less than 25% on a typical parameter estimate. This performance is “out of the box”; unlike the GMM estimators, we didn’t have to choose our set of instruments.

Much of the work in this paper lies in the details. The likelihood requires computation of a Jacobian term, a large matrix which we analytically derive in the appendix using the implicit function theorem. The Jacobian of the likelihood in turn includes terms from the Hessian of the share function. Because numerical errors can easily accumulate we take extra care to implement state of the art computational methods. We do this by utilizing the PyTorch framework developed in Paszke et al. (2019). One particular feature of PyTorch that helps with optimizing is Automatic Differentiation (AD). For optimizing the objective, finite difference methods can lack speed and reliability compared to analytic derivatives. AD exploits the chain rule and the fact that computer code is made up of elementary mathematical operations to compute gradients that are as precise as analytically derived gradients.

Taken together, these simulation and replication exercises suggest that MLE is a useful addition to the set of tools available to estimate BLP, allowing researchers to avoid choosing moment conditions, and having a substantial advantage in statistical precision, at least where the researcher is willing to commit to a model of the supply side.

Related literature. The idea of estimating differentiated products demand systems by maximum likelihood is by no means new. With individual choice data, it is in fact common (see e.g. Honka et al. (2017) and Abaluck and Adams-Prassl (2021) for more recent examples). The marketing literature has also considered using maximum likelihood with aggregate data. Jiang et al. (2009) suggests a Bayesian Markov Chain Monte Carlo (MCMC) approach for demand estimation in the absence of endogenous prices. They find in simulations that their estimates have lower mean squared error. In an extension of their main model, they add a linear supply side pricing equation and suggest how instrumental variables may be used in the presence of endogeneity. Park and Gupta (2009) perform a similar analysis, using maximum likelihood rather than a Bayesian approach, and again insisting on a linear pricing equation and bivariate normal errors. Finally, in an appendix to their paper, Dubé et al. (2012) once again employ a linear pricing equation when considering the performance of a maximum likelihood estimator of BLP. Our paper differs from these prior contributions in that we set up the supply-side pricing equation to be consistent with Bertrand-Nash pricing, leading to an equation in which the demand error enters non-linearly.

2 Model

Our model closely follows Berry et al. (1995). Our exposition is deliberately concise, and we refer the reader to Berry et al. (1995) for additional details. We use i to index individuals, j for products, and t for markets. We present our full notation in Table 1.

Utility. An individual i buying product j in market t obtains utility equal to:

$$u_{ijt} = \mathbf{x}'_{jt}\boldsymbol{\beta}_i - \alpha p_{jt} + \xi_{jt} + \epsilon_{ijt}, \quad (1)$$

where \mathbf{x}_{jt} represents the observable characteristics of a product, p_{jt} is the associated price of a product, ξ_{jt} is a scalar unobservable product characteristic, δ_{jt} is the product's mean utility, and ϵ_{ijt} is an idiosyncratic shock that is i.i.d. and follows a type-1 extreme value distribution. Product characteristics are allowed to have heterogeneous effects on individuals' utilities through $\boldsymbol{\beta}_i$. Each observable characteristic coefficient k follows $\beta_{ik} = \beta_k + \tilde{\beta}_{ik}$, where $\tilde{\beta}_{ik}$ is an i.i.d. random variable drawn from $N(0, \sigma_\beta)$. Mean (across consumers) utility takes the form:

$$\delta_{jt} = \mathbf{x}'_{jt}\boldsymbol{\beta} - \alpha p_{jt} + \xi_{jt}, \quad (2)$$

and thus individual utility can be re-written as:

$$u_{ijt} = \delta_{jt} + \mathbf{x}'_{jt} \tilde{\boldsymbol{\beta}}_i + \epsilon_{ijt}. \quad (3)$$

Demand. Individual consumers make product choices that maximize their utility. The observed market-level shares of each product are the result of aggregating these individual-level decisions. The assumption on the idiosyncratic shock ϵ_{ijt} and the existence of an outside option where $u_{0t} = 0$ yields market shares:

$$s_{jt} = \int \frac{e^{\delta_{jt} + \mathbf{x}'_{jt} \tilde{\boldsymbol{\beta}}_i}}{1 + \sum_{j' \in \mathcal{J}_t} e^{\delta_{j't} + \mathbf{x}'_{j't} \tilde{\boldsymbol{\beta}}_i}} dF(\tilde{\boldsymbol{\beta}}_i | \sigma_\beta). \quad (4)$$

Supply. For the supply side of the market, we assume a marginal cost structure of:

$$c_{jt} = \mathbf{w}'_{jt} \boldsymbol{\gamma} + u_{jt}, \quad (5)$$

where c_{jt} is the marginal cost of producing a product, \mathbf{w}_{jt} is a vector of characteristics and u_{jt} is a latent scalar supply side cost shock. Firm profits are given by:

$$\pi_{ft} = \max_{p_{ft}} \sum_{j \in \mathcal{J}_{ft}} (p_{jt} - c_{jt}) s_{jt}. \quad (6)$$

As in BLP, we assume that firms have the objective of maximizing their profits in a given market across all of their products. This leads to the following first order conditions (FOCs):

$$\frac{\partial \pi_{ft}}{\partial p_{jt}} = s_{jt} + \sum_{j' \in \mathcal{J}_{ft}} (p_{j't} - c_{j't}) \frac{\partial s_{j't}}{\partial p_{jt}}, \quad (7)$$

where π_{ft} is the profit of firm f in market t . This can be written in matrix form as:

$$\mathbf{c}_t = \mathbf{p}_j - (\mathbf{O}_t \mathbf{J}_t^{s,p})^{-1} \mathbf{s}_t, \quad (8)$$

where $J_t^{s,p}[j, j'] = \frac{\partial s_{j't}}{\partial p_j}$ and O_t is the ownership matrix:

$$O_t = \begin{cases} 1 & \text{if } j \wedge j' \in \mathcal{J}_{ft}, \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

Given $\{\mathbf{p}_t, \mathbf{s}_{jt}\}$ and a set of parameters, Eq. 7 can be used to estimate marginal costs.

These marginal costs can be used to estimate the parameters of Eq. 5:

$$\mathbf{c}_t = \mathbf{p}_t - (\mathbf{O}_t \mathbf{J}_t^{s,p})^{-1} \mathbf{s}_t = \mathbf{w}_t \boldsymbol{\gamma} + \mathbf{u}_t. \quad (10)$$

3 Estimation

This section presents the two main estimation procedures for this model: first, we provide a brief discussion of estimation using Generalized Method of Moments (GMM), and then we provide a more elaborate description of the MLE estimator. We will use notation similar to prior literature: $\Theta = \{\theta, \beta, \gamma\}$ represents all linear and non-linear parameters in the model. The parameters (β, γ) enter the mean utility and cost equations linearly, and we refer to them as the linear parameters. We refer to the remaining parameters $\theta = \{\alpha, \sigma_\beta\}$ as the non-linear parameters.

3.1 Generalized Method of Moments

For GMM estimation, the supply and demand residuals are used to form the moment conditions using valid instruments. Identification relies on finding such valid instruments, which typically arise from exclusion restrictions, i.e., observables that enter only one of the demand and supply equations. For a given set of parameters Θ , Equations 2 and 5 imply the following residuals:

$$\mathbf{u}_t(\Theta) = \mathbf{c}_t(\theta) - \mathbf{w}_t \boldsymbol{\gamma}, \quad (11)$$

$$\boldsymbol{\xi}_t(\Theta) = \boldsymbol{\delta}_t(\theta) - \mathbf{x}_t \boldsymbol{\beta} + \alpha \mathbf{p}_t. \quad (12)$$

As Berry (1994) shows, the share equation 4 is invertible in the sense that given the data $\{\mathbf{x}_t, \mathbf{p}_t, \mathbf{s}_t\}$ and non-linear parameters θ , there is a unique value of the mean utilities δ_{jt} that rationalize the observed market shares. We denote this by $\boldsymbol{\delta}_t(\theta)$. Then as seen in the first equality of Equation 10, the implied costs can be recovered from the prices, ownership matrix and the Jacobian of the shares with respect to prices. Notice that these also depend only on the non-linear parameters, and we write the implied costs as a function $\mathbf{c}_t(\theta)$.

The GMM estimator requires demand-side instruments Z_{jt}^d and supply-side instruments Z_{jt}^s that satisfy the moment conditions $E[\boldsymbol{\xi}_{jt} Z_{jt}^d] = 0$ and $E[u_{jt} Z_{jt}^s] = 0$. Prior work suggests various ways to construct these instruments (Berry et al., 1995; Gandhi and Houde, 2019; Conlon and Gortmaker, 2020). Given instruments, one can form the empirical analogs of

the moment conditions:

$$\hat{g}(\theta) = \frac{1}{N} \sum_t \sum_{j \in \mathcal{J}_t} \begin{bmatrix} u_{jt}(\theta) Z_{jt}^s \\ \xi_{jt}(\theta) Z_{jt}^d \end{bmatrix} = 0, \quad (13)$$

and proceed to minimize the GMM objective:

$$\hat{\theta} = \arg \min_{\theta} \hat{g}(\theta) \hat{\Psi} \hat{g}(\theta), \quad (14)$$

where Ψ is a positive definite weighting matrix. In most cases, GMM estimation proceeds in two steps. In the first step, Ψ is typically set to either the identity matrix, or two-stage least squares (2SLS) weighting matrix.¹ In the second step Ψ is replaced with a heteroscedasticity robust weighting matrix that uses the first-step estimates. More details on this estimation routine can be found in BLP, AGS, and CG.²

Optimization We follow common practice and structure the optimization problem as an outer loop that searches over non-linear parameters, and an inner loop that recovers the optimal linear parameters given the non-linear parameters from the outer loop. As noted above, for any guess of the non-linear parameters θ we can generate the implied mean utilities $\delta_t(\theta)$ and marginal costs $\mathbf{c}_t(\theta)$. We may then concentrate out the linear parameters (β, γ) by setting them to the values that minimize the GMM objective for that θ , a process referred to as the “inner loop”. The inner loop has an analytic solution, using an IV-GMM estimator, as described in CG among other references.

Standard Errors We use the standard formula for heteroskedasticity robust standard errors in GMM. In practice the estimated standard errors depend on how one estimates the gradient of the GMM objective: either by finite differences, automatic differentiation, or using an analytic expression for the gradient. We primarily rely on automatic differentiation.

3.2 Maximum Likelihood Estimation

In this section, we outline our maximum likelihood estimation (MLE) procedure. In this context MLE requires a parametric specification of the joint distribution of the error terms, ξ and u . If this assumption holds, MLE is the most statistically efficient estimation procedure available. The identification requirements for MLE are otherwise the same as for GMM: the researcher should have access to a cost shifter to identify the mean price coefficient, and

¹In our simulations and replication exercise, we initialize Ψ using the 2SLS weighting matrix as in Conlon and Gortmaker (2020).

²CG succinctly outlines the estimation procedure in their Algorithm 1.

choice set variation to identify the distribution of random coefficients.³ Thus MLE requires demand and supply-side instruments just as GMM does, but doesn't require the explicit construction of moment conditions. Since MLE can be thought of as GMM with optimal instruments (the first order conditions of the likelihood in the parameters), one can think of MLE as automatically picking the best instruments. This is distinct from the optimal GMM instruments of Chamberlain (1987), which are only semi-parametrically optimal, since that procedure makes no distributional assumptions on the distribution of the residuals.⁴

The likelihood will depend on the joint density of the observed and endogenous variables in the model: shares (\mathbf{s}_t) and prices (\mathbf{p}_t). As mentioned previously, to estimate the likelihood, we make distributional assumptions. Particularly, we assume that the residuals of the demand and supply unobserved characteristics are i.i.d. (homoskedasticity of the residuals) and follow a joint normal distribution:

$$\begin{bmatrix} \xi_{jt} \\ u_{jt} \end{bmatrix} \sim N(0, \Sigma), \text{ where } \Sigma = \begin{bmatrix} \sigma_\xi^2 & \sigma_{\xi,u} \\ \sigma_{\xi,u} & \sigma_u^2 \end{bmatrix}. \quad (15)$$

The conditional density of interest is $f(\mathbf{s}_t, \mathbf{p}_t | \mathbf{x}_t, \mathbf{w}_t; \Theta)$, where $f(\cdot)$ is probability density function. We can re-write this density by using multi-variate change of variables and the Inverse Function Theorem:

$$f(\mathbf{s}_t, \mathbf{p}_t | \mathbf{x}_t, \mathbf{w}_t; \Theta) = f(\boldsymbol{\delta}_t, \mathbf{c}_t | \mathbf{x}_t, \mathbf{w}_t; \Theta) \times |\mathbf{J}(\mathbf{s}_t, \mathbf{p}_t, \mathbf{x}_t, \mathbf{w}_t; \boldsymbol{\theta})^{-1}| \quad (16)$$

where $|\cdot|$ is the determinant and $\mathbf{J}(\mathbf{s}_t, \mathbf{p}_t | \mathbf{x}_t, \mathbf{w}_t; \boldsymbol{\theta})$ are the full derivatives of shares and prices with respect to mean utility and costs, a $2N_t \times 2N_t$ matrix:

$$\mathbf{J}(\mathbf{s}_t, \mathbf{p}_t | \mathbf{x}_t, \mathbf{w}_t; \boldsymbol{\theta}) = \begin{bmatrix} \frac{d\mathbf{s}_t}{d\boldsymbol{\delta}_t} & \frac{d\mathbf{s}_t}{d\mathbf{c}_t} \\ \frac{d\mathbf{p}_t}{d\boldsymbol{\delta}_t} & \frac{d\mathbf{p}_t}{d\mathbf{c}_t} \end{bmatrix}. \quad (17)$$

We derive the Jacobian in Appendix A.1.3. Notice that it doesn't depend on the parameters $\{\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\Sigma}\}$. The density of $f(\boldsymbol{\delta}_t, \mathbf{c}_t | \mathbf{x}_t, \mathbf{w}_t; \Theta)$ can be written as a result of the structure of the model (particularly equations 2 and 10) and the distributional assumptions in equation 15:

$$f(\boldsymbol{\delta}_{jt}, c_{jt} | \mathbf{x}_{jt}, \mathbf{w}_{jt}; \Theta, \Sigma) \sim N(\mu(\mathbf{x}_{jt}, \mathbf{w}_{jt}; \Theta), \Sigma), \quad (18)$$

³Instead of a cost shifter, one may instead impose that the cost and demand errors are independent, in which case the cost residual itself is a valid instrument for price (MacKay and Miller, 2021).

⁴In simulations, CG find that using the Chamberlain (1987) instruments generally performs better than other approaches for picking instruments, but the performance gains they document from doing so are small compared to those we find from the use of MLE.

where μ is a vector-valued function (returning a two-vector), and Σ is a 2×2 matrix. For ease of notation, we denote the data in market t as \mathcal{D}_t and the data across all markets as \mathcal{D} . The log-likelihood can be written by expanding out the joint normal distribution for a market t and taking the natural logarithm of all the terms in Equation 16:

$$\begin{aligned} \ell_t(\Theta, \Sigma; \mathcal{D}_t) &= -\frac{N_t}{2} \log(|\Sigma|) \\ &\quad - \frac{1}{2} \sum_{j \in \mathcal{J}_t} \left(([\delta_{jt} \ c_{jt}] - \mu(\mathbf{x}_{jt}, \mathbf{w}_{jt}; \Theta))^T \Sigma^{-1} \begin{bmatrix} \delta_{jt} \\ c_{jt} \end{bmatrix} - \mu(\mathbf{x}_{jt}, \mathbf{w}_{jt}; \Theta) \right) \\ &\quad - \log(|\mathbf{J}(\mathcal{D}_t; \theta)|). \end{aligned}$$

Write $\begin{bmatrix} \xi_{jt}(\Theta) \\ u_{jt}(\Theta) \end{bmatrix} = \begin{bmatrix} \delta_{jt} \\ c_{jt} \end{bmatrix} - \mu(\mathbf{x}_{jt}, \mathbf{w}_{jt}; \Theta)$, substitute into the likelihood and sum across markets to get the aggregate log-likelihood:

$$\begin{aligned} \ell(\Theta, \Sigma; \mathcal{D}) &= \sum_t \ell_t(\Theta, \Sigma; \mathcal{D}_t) \\ &= \sum_t \left(-\frac{N_t}{2} \log(|\Sigma|) - \frac{1}{2} \sum_{j \in \mathcal{J}_t} [\xi_{jt} \ u_{jt}]^T \Sigma^{-1} \begin{bmatrix} \xi_{jt} \\ u_{jt} \end{bmatrix} - \log(|\mathbf{J}(\mathcal{D}_t; \theta)|) \right). \end{aligned} \tag{19}$$

We will refer to the log-likelihood given above as the *unconcentrated* version of the log-likelihood.

Optimization To simplify optimization, we will concentrate out a number of parameters. First consider the variance-covariance matrix Σ . Fixing Θ and therefore the residuals $\{\xi_{jt}(\Theta), u_{jt}(\Theta)\}$, the likelihood maximizing choice of Σ is the implied variance-covariance matrix of the residuals. This is a function of the parameters, which we denote by $\Sigma^*(\Theta)$. As shown in Appendix A.1.1, making this substitution simplifies the middle term of the log-likelihood in eq 19 to a constant, resulting in the following likelihood:

$$\ell(\Theta, \Sigma; \mathcal{D}) = -\frac{N}{2} \log(|\Sigma^*(\Theta)|) - \sum_t \log(|\mathbf{J}(\mathcal{D}_t; \theta)|). \tag{20}$$

Notice that the linear parameters do not enter the Jacobian, only the implied variance-covariance matrix. Thus similar to the GMM procedure, we can concentrate out the linear

parameters by solving an inner loop problem of the form:

$$\boldsymbol{\gamma}^*(\boldsymbol{\theta}), \boldsymbol{\beta}^*(\boldsymbol{\theta}) = \arg \min_{\boldsymbol{\gamma}, \boldsymbol{\beta}} \log(|\boldsymbol{\Sigma}^*(\boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\gamma})|). \quad (21)$$

Now recall that the implied variance covariance-matrix $\boldsymbol{\Sigma}^*(\boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ is a 2×2 matrix, with a determinant equal to $\sigma_{\xi}^2(\boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\gamma})\sigma_u^2(\boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\gamma}) - (\sigma_{\xi,u}(\boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\gamma}))^2$. In the absence of a covariance term, minimizing this objective would be equivalent to minimizing the demand and cost residuals separately. This could be done by separate OLS regressions. But the presence of a covariance term complicates things, leading to a non-linear objective with no closed form solution.⁵ In Appendix A.1.2 we show that $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ can be solved in closed form holding the other parameter fixed and so we estimate them jointly by alternating least squares until convergence.

This yields a concentrated likelihood function that depends only on the non-linear parameters, to be maximized in an outer loop:

$$\ell(\boldsymbol{\theta}; \mathcal{D}) = \underbrace{-\frac{N}{2} \log(|\boldsymbol{\Sigma}^*(\boldsymbol{\theta})|)}_{\text{Part I}} - \underbrace{\sum_t \log(|\mathbf{J}(\boldsymbol{\theta})|)}_{\text{Part II}}, \quad (22)$$

where we will call part I of the likelihood the *covariance term* and part II the *Jacobian term*.

Standard Errors To compute standard errors, we use the inverse of the Fisher information matrix, which is given in Appendix A.3. For accurate estimation of the standard errors, we use the Fisher information matrix of the unconcentrated likelihood, to allow the uncertainty in the inner loop estimation procedures to be accounted for.

Computational Details and Algorithm The optimization of the likelihood initially follows the GMM procedure in inverting from market shares to mean utilities and marginal cost vectors $\{\boldsymbol{\delta}, \mathbf{c}\}$. In order to invert from market shares to mean utilities, we use the SQUAREM method proposed by Varadhan and Roland (2008) and recommended by CG. The marginal cost is calculated from the Bertrand-Nash equations detailed in the Model section.

The rest of the optimization is best understood when thinking about the *covariance* term and the *Jacobian* term of the log-likelihood in equation 22 separately. The covariance term is evaluated by concentrating out the linear terms using the objective in equation 21.

⁵In the GMM estimation procedure, OLS also fails unless there is no covariance in the errors. But there IV-GMM delivers a closed-form solution for $(\boldsymbol{\beta}, \boldsymbol{\gamma})$.

The terms entering the Jacobian terms depend on partial Hessians of the shares with respect to price. They are derived in detail in Appendix A.1.3. Note that these Hessians are the computational bottleneck of the estimation routine, having dimensions of size $(T, N_i, \max(N_t), \max(N_t), \max(F_t))^6$.

Given data, Algorithm 1 outlines the procedure of the estimation objective. The code developed for this application uses the PyTorch package. PyTorch provides a high-performance library for optimizing both the GMM and MLE objectives.⁷

Algorithm 1 MLE estimation of BLP

Require: $\{\theta, data\}$

→ Calculate $\delta(\theta|data)$ using the SQUAREM algorithm.

→ Calculate $c(\theta|data)$ using Equation 7.

→ Calculate $\{\beta^*, \gamma^*\}$ using the ALS process described in appendix section A.1.2.

→ Calculate $\ell_I(\theta, \beta^*, \gamma^*|data)$, part I of the log-likelihood equation 22.

→ Calculate $\ell_{II}(\theta|data)$, part II of the log-likelihood equation 22 as described in appendix sections A.1.3 and A.1.3.

return $\ell = \ell_I + \ell_{II}$

4 Simulations

In this section, we evaluate the performance of GMM and MLE on simulated data, where the ground truth is known. Our simulation procedure is based on Conlon and Gortmaker (2020) and Armstrong (2016).

For each simulation, we set the number of markets to $T = 20$. The number of firms F_t in each market is randomly chosen from the set $\{2, 5, 10\}$. The number of products for each firm \mathcal{J}_{ft} is also chosen randomly from $\{3, 4, 5\}$. The structural error terms are given by:

$$\begin{bmatrix} \xi_j \\ u_j \end{bmatrix} \sim N(0, \Sigma) \text{ where } \Sigma = \begin{bmatrix} 0.2 & 0 \\ 0 & 0.2 \end{bmatrix}.$$

The linear demand characteristics are $[1, x_{jt}, p_{jt}]$ and the linear supply characteristics are $[1, x_{jt}, w_{jt}]$. Both exogenous characteristics (x_{jt}, w_{jt}) are drawn from a standard uniform distribution. We allow for two random coefficients on the demand side. The first is on the

⁶The $\max(N_t)$ is a result of how we programmed markets. Essentially, in order to have all markets concatenated in a matrix or tensor, the max number of products is used. Additionally, the $\max(F_t)$ dimension is a result of the particular hessian objects having a dimension that has to only do with products that are within the same firm.

⁷Two reasons for using PyTorch are: (1) utilizing the efficient tensor and matrix operations, (2) utilizing automatic differentiation for optimization.

demand characteristic x_{jt} , where $\tilde{\beta}_i \sim N(0, \sigma_\beta^2)$ and $\sigma_\beta = 3$. The second one is on price, where $\tilde{\alpha}_i \sim N(0, \sigma_\alpha^2)$ and $\sigma_\alpha = 0.2$. To compute the values of the endogenous variables (p_{jt}, s_{jt}) we use the Morrow and Skerlos (2011) method as described in Conlon and Gortmaker (2020). We use Gauss-Hermite quadrature with a product rule to integrate over the 2-dimensional distribution of random coefficients. We provide details in Appendix A.2. We set the ground-truth values for the demand-side parameters to $[\beta_0, \beta_x, \alpha] = [-7, 6, -1]$, and the supply-side parameters to $[\gamma_0, \gamma_x, \gamma_w] = [2, 1, 0.2]$. To ensure reproducibility, we fix a random seed for each simulation.

We refer to the setup above as “No covariance”, and also consider the following variations:

- **Low covariance:** We set the variance-covariance matrix to $\Sigma = \begin{bmatrix} 0.2 & 0.1 \\ 0.1 & 0.2 \end{bmatrix}$.
- **High covariance:** We set the variance-covariance matrix to $\Sigma = \begin{bmatrix} 0.3 & 0.2 \\ 0.2 & 0.3 \end{bmatrix}$.
- **Laplace, no and low covariance:** The errors are generated so that the marginals are Laplace instead of normal. A Gaussian copula is used to generate the draws in the low covariance case, where Σ takes the form above.
- **Supply Misspecification:** The supply side functional form is misspecified. Marginal costs are linear in characteristics as per Equation 6, but the estimation assumes log-linearity $\ln(c_{jt}) = w_{jt}\gamma + u_{jt}$.
- **Ownership Misspecification:** The ownership matrix is misspecified, in that we set O to be the identity matrix (i.e. each product is assumed to be owned by a single firm). The truth is that each firm produces 3, 4 or 5 products.

For GMM estimation, we construct local differentiation instruments following Gandhi and Houde (2019). We also create a predicted price instrument, by regressing price on valid instruments (Reynaert and Verboven, 2014).

4.1 Results

We compare MLE and GMM for each of the above scenarios on three metrics: parameter estimates, standard errors, and own-price elasticities. For each scenario, we run 1000 simulations, with a different random seed for each simulation⁸. For each simulation, optimization starts at a random initial guess which is drawn uniformly within a 50% range of

⁸Each random seed is also randomly chosen. The random seed is assigned for reproducibility.

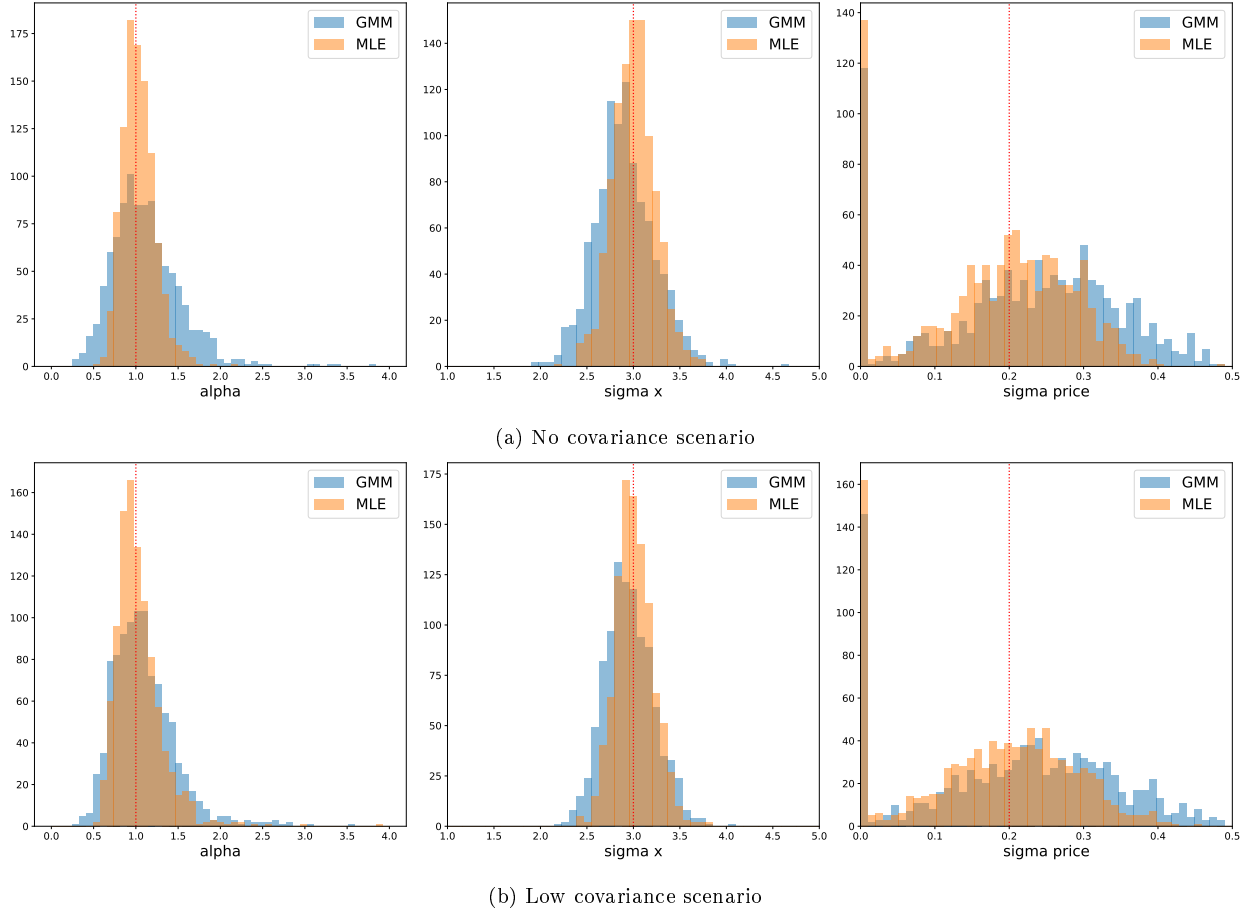


Figure 1: These figures plot the histograms of parameters across the 1000 simulation estimates for the (a) no covariance scenario and the (b) low covariance scenario. The dashed vertical line in each sub-figure is placed at the true parameter value. The results show the overall distribution of parameter estimates.

the true values. Optimization is repeated 3 times with different starting values, after which the parameter estimates are chosen based on the smallest objective value.

To validate that we have correctly coded up the scenarios in CG, as well as our particular implementation of GMM, we compare our estimates to those produced by PyBLP. The parameter estimates are in most cases indistinguishable. In addition, the mean bias and mean absolute bias we report in our GMM results closely match those in Table 5 of CG.

The scenarios that we run simulations for can be viewed in three different categories. First, the *correctly specified* scenarios, where both the MLE and the GMM estimators are correctly specified. Second, the *MLE misspecified* scenarios, which are the Laplace scenarios (GMM is correctly specified in this case). Last, there are the *fully misspecified* scenarios, where both MLE and GMM are misspecified, these are the Supply Misspecification and Ownership Misspecification scenarios.

Mean Bias and RMSE. We begin by summarizing the bias and root mean squared error (RMSE) of the estimators, in two Tables: (a) the RMSE of the estimators can be found in Table 2 and (b) the mean bias can be found in Table 3. The results show that in general the MLE estimates perform somewhat better when it comes to mean bias, and clearly outperform the GMM estimates when looking at the RMSE values. Another way of putting this is for a particular realization of the data, MLE estimates are generally closer to the true values. A visualization of this can be seen in Figure 1, where we see the histogram of the parameter estimates of each of the three parameters for the no covariance and low covariance scenarios.

For RMSE, MLE clearly outperforms the GMM for all the correctly specified and MLE misspecified scenarios. In the fully misspecified scenarios, GMM performs better when there is a supply side misspecification and MLE performs better when the ownership matrix is misspecified.


Standard Errors and Coverage. Standard errors and confidence intervals are fundamental for inference. These results are summarized in two tables: (a) the mean standard error for a scenario in Table 4 and (b) the coverage percentile of the true values for the 95% confidence intervals in Table 5.⁹ The coverage percentile represents the percent of simulations where:

$$\theta_{true} \in [\hat{\theta} - 1.96\hat{\sigma}_{\theta}/\sqrt{n}, \hat{\theta} + 1.96\hat{\sigma}_{\theta}/\sqrt{n}]. \quad (23)$$


The mean standard errors for MLE are tighter under all correctly specified scenarios as well as the Laplace scenarios (the one exception to this is the mean standard error of α in the high covariance scenario). Table 5 shows the coverage of each estimator for each parameter, for a nominal 95%. In general GMM tends to badly under-cover, with true coverage of between 80-90% for the mean price parameter, 90-95% for the standard deviation of the random coefficient on x and 70-80% for the standard deviation of the random coefficient on price. By contrast for MLE those coverages are respectively 88-95%, 90-95% and 90-95%, so that the undercoverage problem is less severe. The sole exception to this general pattern is the case where the cost structure is misspecified, where MLE performs terribly.

Own-price elasticities Demand systems are often used to generate own-price elasticities. Given the interest in these statistics, we examine what the mean bias and absolute mean

⁹For the standard errors of the σ_{price} parameter in the GMM procedure, we find that about 10% of the simulations have numerically infeasible values (standard error values ranging from 1000 to 1e13). To validate that this is not an error in our standard error computations, we take a handful of simulations with such issues and run the pyBLP package provided by CG. The outputs were *NaN* for the corresponding standard error values. It therefore appears that this problem is not specific to our code, but to the GMM procedure implemented here and in pyBLP itself. We drop these cases in the relevant columns of Tables 4 and 5; if they were included the performance of GMM would be much worse.



(a) No covariance scenario



(b) Low covariance scenario

Figure 2: This figure has a set of histograms summarizing the bias (left) and absolute bias (right) of product-level own-price elasticities in the (a) no covariance and (b) low covariance scenarios. Within each subfigure, the vertical dashed lines are the average bias and average absolute bias for GMM (blue) and MLE (red).

bias of product-level own-price elasticities are for a given set of estimated parameters. Figure 2 shows this for the no covariance and low covariance scenarios. The results support the conclusion from the previous sections, the MLE produces tighter results in most of our scenarios.

Jacobian. One notable feature of MLE is the presence of a Jacobian in the likelihood. This accounts for the non-linear mapping from marginal costs and mean utilities to the (endogenous) prices and quantities. One might think that minimizing the covariance term in the likelihood would be enough to ensure good performance, since doing so in some sense minimizes the variance of the demand and cost residuals. We investigate the importance of the Jacobian term in the likelihood here.

Figure 3 plots the likelihood decomposition into the covariance and Jacobian terms sep-




Figure 3: This figure plots the likelihood and its components for different values of σ_x . Each subplot has the same x-axis, which is a range of σ_x around the true value of 3, while α and σ_{price} stay fixed at the true value of 1. The figure on the left shows the full likelihood detailed in equation 22, with the dashed green line highlighting the GMM loss. The middle figure is the covariance term of the likelihood, and the figure on the right is the Jacobian term of the likelihood. Finally, the vertical dashed red line is the true value of sigma.

arately, where we vary the σ_x parameter at the true $\alpha = 1$ value (for a random simulation). The covariance term is minimized below the truth, and the Jacobian term is necessary to correct the bias in the covariance term and arrive at the correct solution. The Jacobian thus plays an important role, in some sense regularizing the estimator by penalizing parameters under which the derivatives of the endogenous variables in the exogenous variables are big (i.e. parameters at which the predictions are unstable).

5 Empirical exercise: replicating BLP'95

In this section, we describe our replication of BLP'95 using MLE. We use the same data as Berry et al. (1995), provided by Andrews et al. (2017). The data contains US automobile data from 1971 to 1990. The dataset also includes 5 product characteristics, horse power to weight ratio (hpwt), if a car has air conditioning (air), miles per dollar of gasoline (mpd), width times length (space), and miles per gallon (mpg).

We start by replicating the GMM estimation procedure. We follow the procedure outlined in AGS, which is based off of BLP. This includes dropping certain instruments due to collinearity as well as starting the estimation procedure from the estimated parameters in BLP. As in BLP, we have 2217 model/years with 997 distinct models. The instruments that we use are BLP sums of characteristic instruments. With this setup, the local GMM estimation produces estimates that are almost identical to the AGS parameters.

Next, we estimate the parameters using MLE. For the MLE setup, we run the procedure on the same dataset. Our integration approach follows AGS by using the same random

draws and importance sampling. One change we make is to double the number of draws to ensure symmetry. More specifically, we take the set of draws, multiply the vector associated with air conditioning with negative one, and concatenate these to the original draws. This leaves us with two times the total draws.¹⁰

To get a sense of how reasonable our results are, we compare our estimates to existing benchmarks in the literature that use this dataset. These benchmarks are the estimates reported in BLP, AGS, and CG. These benchmarks differ in two main ways, the instruments selected and the choice of draws (and weights) used for integration. BLP use the sum of rival characteristics (BLP) instruments and integrate using importance sampling. AGS follows BLP closely in order to replicate their estimates. CG uses the Chamberlain (1987) optimal instruments and quasi-Monte Carlo integration (10,000 scrambled Halton draws in each market).

The estimation results can be found in Table 6 with standard error comparisons in parentheses below the parameter estimates. The parameter estimates of MLE are closely aligned with the GMM estimates. In particular, the α parameter is closest to the CG best practice estimates. The sigma estimates tend to align with the BLP results, and generally indicate larger heterogeneity in the data than the CG results. The similarity of the parameter estimates is surprising since the effort put into calibrating the MLE estimator to the automobile data was minimal i.e. these results are “out-of-the-box”. No calibration or instrument construction was necessary. Standard error estimates of the MLE parameters are significantly tighter for MLE, regularly one quarter of the smallest GMM standard error. The standard error on α is 0.769, compared to a range of 5.6-11.7 for the GMM. This is nearly 10-fold reduction compared to GMM. Since we don’t know the ground truth we cannot comment on whether the estimated standard errors imply confidence intervals with correct coverage, but the simulations above give us confidence that if the model is correctly specified the coverage will be approximately correct.

Finally, we look at the model level comparison of the own-price elasticities of the GMM and MLE estimators in Figure 4. We can see that there is a fairly large mass around the diagonal, indicating generally similar estimates. The most significant trend off of the diagonal is a section of the AGS estimates. This trend can be fully attributed to the large difference in the σ_{air} estimates, where AGS estimates this parameter to be 4.2, compared to BLP, CG, and MLE estimating it to between 1.4-2. Since the air conditioning variable is binary, it leads to two separate trend lines, where the one above the diagonal representing

¹⁰When estimating the model by MLE using the AGS draws, the parameter estimates were mostly similar, but σ_{air} value was estimated to be 0. This is a result of an artificial lack of symmetry around zero of the objective function in σ_{air} due to sampling variance in the draws; doubling the number of draws fixes this problem.




Figure 4: This figure plots the own-price elasticity comparisons between MLE and GMM parameters. Each dot represents a model own-price elasticity, with the different colors representing the three main GMM estimates we use as benchmarks in this paper.

models that have air conditioning.

Overall, we were able to use MLE on the BLP automobile data with relative ease and produce estimates that are similar to prior work but with tighter standard errors and without having to construct valid instruments.

6 Conclusion

We develop and evaluate a maximum likelihood estimator for differentiated products demand. Unlike the traditionally used GMM estimator, the ML estimator requires stricter distributional assumptions, while offering the advantages of statistical efficiency and limiting choices required by the researcher to fit the model. While we are not the first to propose an ML estimator, prior work has relied on reduced-form supply specifications, whereas our supply-side pricing is consistent with Bertrand-Nash pricing.

In the simulation section, we show that under correct specification, our method outperforms GMM as well as showing its robustness and sensitivity to different types of misspecification. Finally, we replicate BLP using the new estimation method, finding that the estimates are very similar by comparing the estimated values to other benchmarks. These results demonstrate that MLE can be a useful addition to the tools available to researchers estimating BLP.

References

- J. Abaluck and A. Adams-Prassl. What do consumers consider before they choose? identification from asymmetric demand responses. *The Quarterly Journal of Economics*, 136(3): 1611–1663, 2021.
- I. Andrews, M. Gentzkow, and J. M. Shapiro. Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics*, 132(4):1553–1592, 2017.
- T. B. Armstrong. Large market asymptotics for differentiated product demand estimators with economic models of supply. *Econometrica*, 84(5):1961–1980, 2016.
- S. Berry, J. Levinsohn, and A. Pakes. Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890, 1995.
- S. T. Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262, 1994.
- G. Chamberlain. Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of econometrics*, 34(3):305–334, 1987.
- C. Conlon and J. Gortmaker. Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics*, 51(4):1108–1161, 2020.
- C. T. Conlon. The empirical likelihood mpec approach to demand estimation. *Available at SSRN 2331548*, 2013.
- J.-P. Dubé, J. T. Fox, and C.-L. Su. Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation. *Econometrica*, 80(5):2231–2267, 2012.
- A. Gandhi and J.-F. Houde. Measuring substitution patterns in differentiated-products industries. Technical report, National Bureau of Economic Research, 2019.
- E. Honka, A. Hortaçsu, and M. A. Vitorino. Advertising, consumer awareness, and choice: Evidence from the us banking industry. *The RAND Journal of Economics*, 48(3):611–646, 2017.
- R. Jiang, P. Manchanda, and P. E. Rossi. Bayesian analysis of random coefficient logit models using aggregate data. *Journal of Econometrics*, 149(2):136–148, 2009.

- A. MacKay and N. Miller. Estimating models of supply and demand: Instruments and covariance restrictions. Technical report, Harvard University, 2021.
- D. McFadden. The choice theory approach to market research. *Marketing science*, 5(4): 275–297, 1986.
- W. R. Morrow and S. J. Skerlos. Fixed-point approaches to computing bertrand-nash equilibrium prices under mixed-logit demand. *Operations research*, 59(2):328–345, 2011.
- S. Park and S. Gupta. Simulated maximum likelihood estimator for the random coefficient logit model using aggregate data. *Journal of Marketing Research*, 46(4):531–542, 2009.
- A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019.
- M. Reynaert and F. Verboven. Improving the performance of random coefficients demand models: the role of optimal instruments. *Journal of Econometrics*, 179(1):83–98, 2014.
- B. S. Skrainka and K. L. Judd. High performance quadrature rules: How numerical integration affects a popular model of product differentiation. *Available at SSRN 1870703*, 2011.
- R. Varadhan and C. Roland. Simple and globally convergent methods for accelerating the convergence of any em algorithm. *Scandinavian Journal of Statistics*, 35(2):335–353, 2008.

Table 1: Notation.

Variable	Definition
t	Markets in range $\{1, 2, 3, \dots, T\}$
F_t	The number of firms in market t .
f	Firm f in range $\{1, \dots, F_t\}$
j	Product j in range $\{1, 2, 3, \dots, N_t\}$
\mathcal{J}_t	Set of products in market t
\mathcal{J}_{ft}	Set of products in firm f and market t
N_t	Number of products in market t
N_{ft}	Number of products in firm f and market t
x_j	Observed characteristics of product j
β_i	Consumer i 's taste parameter
p_j	Price of product j
c_j	Marginal cost of product j
ξ_j	Unobserved characteristics of product j
ϵ_{ij}	Distribution of consumer preferences around the mean
θ	A vector of all parameters.

Table 2: RMSE of simulations.

	GMM			MLE		
	α	σ_x	σ_{price}	α	σ_x	σ_{price}
No covariance	0.42	0.34	0.13	0.20	0.23	0.10
Low covariance	0.41	0.27	0.13	0.28	0.20	0.11
High covariance	0.40	0.29	0.14	0.43	0.20	0.11
Laplace, no covariance	0.46	0.34	0.13	0.22	0.23	0.11
Laplace, low covariance	0.38	0.27	0.13	0.33	0.21	0.12
Supply misspecification	0.42	0.32	0.14	0.99	0.41	0.13
Ownership misspecification	0.39	0.32	0.12	0.34	0.21	0.11

Notes This table presents the root mean squared error of different simulation scenarios. Each column is a combination of an estimator and a parameter. We use 1000 simulations to generate each row in this table. Note that the supply misspecification and the ownership misspecification scenarios also have low covariance.

Table 3: Mean bias of simulations.

	GMM			MLE		
	α	σ_x	σ_{price}	α	σ_x	σ_{price}
No covariance	0.12	-0.09	0.03	0.03	0.01	-0.02
Low covariance	0.10	-0.04	0.01	0.03	0.01	-0.02
High covariance	0.04	-0.08	0.01	0.05	-0.00	-0.03
Laplace, no covariance	0.11	-0.08	0.02	0.03	-0.00	-0.01
Laplace, low covariance	0.08	-0.06	0.02	0.05	0.00	-0.02
Supply misspecification	0.20	0.08	-0.01	0.94	0.30	-0.05
Ownership misspecification	-0.05	-0.20	0.03	-0.00	-0.09	-0.02

Notes This table presents the root mean bias of different simulation scenarios. Each column is a combination of an estimator and a parameter. We use 1000 simulations to generate each row in this table. Note that the supply misspecification and the ownership misspecification scenarios also have low covariance.

Table 4: Mean standard error values.

	GMM			MLE		
	α	σ_x	σ_{price}	α	σ_x	σ_{price}
No covariance	0.30	0.30	0.13	0.20	0.22	0.12
Low covariance	0.28	0.27	0.22	0.23	0.19	0.12
High covariance	0.28	0.29	0.16	0.30	0.20	0.14
Laplace, no covariance	0.30	0.32	0.14	0.20	0.22	0.13
Laplace, low covariance	0.28	0.28	0.15	0.25	0.20	0.13
Supply misspecification	0.29	0.30	0.16	0.20	0.21	0.10
Ownership misspecification	0.26	0.24	0.10	0.30	0.19	0.13

Notes This table presents the standard errors of different simulation scenarios. Each column is a combination of an estimator and a parameter. We use 1000 simulations to generate each row in this table. Note that the supply misspecification and the ownership misspecification scenarios also have low covariance. In the GMM sigma price column, we omit approximately 10% of the simulations because the standard error calculation for these simulations were producing numerically infeasible values.

Table 5: Fraction of 95% CI that include the true values.

	GMM			MLE		
	α	σ_x	σ_{price}	α	σ_x	σ_{price}
No covariance	0.86	0.92	0.79	0.94	0.93	0.94
Low covariance	0.89	0.95	0.77	0.93	0.94	0.94
High covariance	0.84	0.93	0.72	0.88	0.94	0.95
Laplace, no covariance	0.86	0.93	0.79	0.94	0.95	0.92
Laplace, low covariance	0.88	0.95	0.77	0.91	0.94	0.90
Supply misspecification	0.93	0.97	0.70	0.01	0.64	0.86
Ownership misspecification	0.83	0.81	0.80	0.89	0.91	0.93

Notes This table presents the coverage of different estimators. Each column is a combination of an estimator and a parameter, and the entries are the fraction of the simulations in which the truth lay within the 95% confidence intervals produced by each estimator. We use 1000 simulations to generate each row in this table. Note that the supply misspecification and the ownership misspecification scenarios also have low covariance.

Table 6: BLP automobile data parameter estimates with standard errors.

	Variable	BLP	AGS	CG	MLE
Term on Price (α)	$\ln(y - p)$	43.501 (6.43)	42.870 (5.56)	45.898 (11.75)	45.227 (0.77)
Standard deviations	Constant	3.612 (1.49)	2.522 (2.54)	2.962 (1.64)	1.924 (0.16)
	HP/Weight	4.628 (1.89)	3.525 (2.84)	1.388 (2.11)	4.611 (0.28)
	Air	1.818 (1.70)	4.167 (1.41)	1.424 (0.43)	1.991 (0.23)
	MP\$	1.050 (0.27)	0.393 (0.28)	0.072 (1.00)	0.299 (0.04)
	Size	2.056 (0.58)	1.937 (0.60)	0.231 (3.84)	2.168 (0.11)
Means	Constant	-7.061 (0.94)	-7.728 (1.16)	-6.679 (1.30)	-5.439 (0.33)
	HP/Weight	2.883 (2.02)	4.620 (1.13)	2.774 (0.83)	3.175 (0.35)
	Air	1.521 (0.89)	-1.227 (1.38)	0.572 (0.35)	0.385 (0.16)
	MP\$	-0.122 (0.32)	0.293 (0.16)	0.340 (0.10)	0.086 (0.06)
	Size	3.460 (0.61)	3.992 (0.35)	3.920 (0.32)	3.327 (0.17)
Cost Side Parameters	Constant	0.952 (0.19)	2.751 (0.08)	2.780 (0.10)	2.327 (0.04)
	$\ln(\text{HP/Weight})$	0.477 (0.06)	0.812 (0.06)	0.731 (0.07)	0.495 (0.03)
	Air	0.619 (0.04)	0.430 (0.05)	0.528 (0.04)	0.601 (0.02)
	$\ln(y - p)$	-0.415 (0.06)	-0.610 (0.05)	-0.651 (0.07)	-0.377 (0.04)
	$\ln(\text{Size})$	-0.046 (0.08)	-0.352 (0.11)	-0.472 (0.12)	0.028 (0.06)
	Trend	0.019 (0.00)	0.027 (0.00)	0.018 (0.00)	0.015 (0.00)

Notes This table presents parameter estimates of different estimators on the BLP automobile data. Column (1) has the original estimates, column (2) is from the AGS data, (3) is the best practice estimates from the CG paper, (4) is the parameter estimates from the MLE estimator. Standard errors are shown parentheses.

A Appendix

A.1 Deriving and simplifying the log-likelihood

In this section of the appendix, we are concerned with deriving terms and simplifying the unconcentrated log likelihood. We first show how the parameters (Σ, β, γ) may be concentrated out, and then provide an analysis of the Jacobian term in the concentrated log likelihood.

A.1.1 Concentrating out Σ

Define the following equalities for a given dataset ranging with observations (product-markets) ranging from 1 to m :

$$\begin{aligned} \sum_{i=1}^m \xi_i^2 &= \sum_{i=1}^m (\delta_i - x_i \beta)^2 = (\delta - X\beta)^T (\delta - X\beta) = \xi^T \xi \\ \sum_{i=1}^m u_i^2 &= \sum_{i=1}^m (c_i - w_i \gamma)^2 = (C - W\gamma)^T (C - W\gamma) = u^T u \end{aligned} \tag{24}$$

Recall Part II of the log-likelihood is given by:

$$\begin{aligned} &= ([\delta_j \ c_j] - E[\delta_j \ c_j | x_t, w_t]) \Sigma^{-1} \begin{bmatrix} \delta_j \\ c_j \end{bmatrix} - E[\delta_j \ c_j | x_t, w_t]^T \\ &= ([\delta_j \ c_j] - E[\delta_j \ c_j | x_t, w_t]) \begin{bmatrix} \sigma_\xi^2 & \sigma_{\xi,u} \\ \sigma_{\xi,u} & \sigma_u^2 \end{bmatrix}^{-1} \begin{bmatrix} \delta_j \\ c_j \end{bmatrix} - E[\delta_j \ c_j | x_t, w_t]^T \end{aligned} \tag{25}$$

where for simplicity we index observations by j instead of jt . The maximum likelihood estimate of Σ , holding the other parameters fixed, is the sample variance-covariance matrix of the error terms.

$$\begin{aligned} &= \begin{bmatrix} \xi_j & u_j \end{bmatrix} \begin{bmatrix} \sum_{i=1}^m \xi_i^2 & \sum_{i=1}^m \xi_i u_i \\ \sum_{i=1}^m \xi_i u_i & \sum_{i=1}^m u_i^2 \end{bmatrix}^{-1} \begin{bmatrix} \xi_j \\ u_j \end{bmatrix} \\ &= \frac{1}{\sum_{i=1}^m \xi_i^2 \sum_{i=1}^m u_i^2 - (\sum_{i=1}^m \xi_i u_i)^2} \begin{bmatrix} \xi_j & u_j \end{bmatrix} \begin{bmatrix} \sum_{i=1}^m u_i^2 & -\sum_{i=1}^m \xi_i u_i \\ -\sum_{i=1}^m \xi_i u_i & \sum_{i=1}^m \xi_i^2 \end{bmatrix} \begin{bmatrix} \xi_j \\ u_j \end{bmatrix} \\ &= \frac{\xi_j^2 \sum_{i=1}^m u_i^2 - 2u_j \xi_j \sum_{i=1}^m \xi_i u_i + u_j^2 \sum_{i=1}^m \xi_i^2}{\sum_{i=1}^m \xi_i^2 \sum_{i=1}^m u_i^2 - (\sum_{i=1}^m \xi_i u_i)^2} \end{aligned} \tag{26}$$

Next, we can aggregate the likelihood across all data points:

$$\begin{aligned}
&= \sum_{j=1}^m \left(\frac{\xi_j^2 \sum_{i=1}^m u_i^2 - 2u_j \xi_j \sum_{i=1}^m \xi_i u_i + u_j^2 \sum_{i=1}^m \xi_i^2}{\sum_{i=1}^m \xi_i^2 \sum_{i=1}^m u_i^2 - (\sum_{i=1}^m \xi_i u_i)^2} \right) \\
&= \frac{\sum_{j=1}^m \xi_j^2 \sum_{i=1}^m u_i^2 - 2 \sum_{j=1}^m u_j \xi_j \sum_{i=1}^m \xi_i u_i + \sum_{j=1}^m u_j^2 \sum_{i=1}^m \xi_i^2}{\sum_{i=1}^m \xi_i^2 \sum_{i=1}^m u_i^2 - (\sum_{i=1}^m \xi_i u_i)^2} \\
&= \frac{2 \sum_{j=1}^m \xi_j^2 \sum_{i=1}^m u_i^2 - 2(\sum_{j=1}^m u_j \xi_j)^2}{\sum_{i=1}^m \xi_i^2 \sum_{i=1}^m u_i^2 - (\sum_{i=1}^m \xi_i u_i)^2} \\
&= 2
\end{aligned} \tag{27}$$

A.1.2 Estimating the linear parameters

This section is concerned with the estimation of Part I of eq. 22. As mentioned before, u and ξ have a bivariate normal distribution with arbitrary correlation. The local objective function that we are concerned with for this part is is:

$$\begin{aligned}
\min_{\gamma, \beta} \log(|\hat{\Sigma}|) &= \min_{\gamma, \beta} \log \left(\left| \begin{bmatrix} \sum_{i=1}^m \xi_i^2 & \sum_{i=1}^m \xi_i u_i \\ \sum_{i=1}^m \xi_i u_i & \sum_{i=1}^m u_i^2 \end{bmatrix} \right| \right) \\
&= \min_{\gamma, \beta} \log \left(\sum_{i=1}^m \xi_i^2 \sum_{i=1}^m u_i^2 - \left(\sum_{i=1}^m \xi_i u_i \right)^2 \right)
\end{aligned} \tag{28}$$

where $|\cdot|$ is the determinant. Since $\log(\cdot)$ is monotonic, the objective above is equivalent to:

$$\min_{\gamma, \beta} \log(|\Sigma|) = \min_{\gamma, \beta} \left(\sum_{i=1}^m \xi_i^2 \sum_{i=1}^m u_i^2 - \left(\sum_{i=1}^m \xi_i u_i \right)^2 \right) \tag{29}$$

Next, we can write out the objective in matrix form and find the FOCs:

$$O = (\delta - X\beta)^T (\delta - X\beta) (C - W\gamma)^T (C - W\gamma) - [(\delta - X\beta)^T (C - W\gamma)]^2 \tag{30}$$

Taking a derivative in β :

$$\begin{aligned}
\frac{\partial O}{\partial \beta} &= -2X^T (\delta - X\beta) (C - W\gamma)^T (C - W\gamma) - 2X^T (C - W\gamma) (C - W\gamma)^T (\delta - X\beta) \\
0 &= X^T (\delta - X\beta) u^T u - 2X^T u u^T (\delta - X\beta) \\
X^T X \beta u^T u + X^T u u^T X \beta &= X^T \delta u^T u + X^T u u^T \delta \\
(X^T X u^T u + X^T u u^T X) \beta &= X^T \delta u^T u + X^T u u^T \delta \\
(X^T X u^T u + X^T u u^T X) \beta &= X^T \delta u^T u + X^T u u^T \delta \\
\hat{\beta} &= (X^T X u^T u + X^T u u^T X)^{-1} (X^T \delta u^T u + X^T u u^T \delta)
\end{aligned} \tag{31}$$

Similarly, for γ :

$$\begin{aligned}\frac{\partial O}{\partial \gamma} &= -2W^T(C - W\gamma)(\delta - X\beta)^T(\delta - X\beta) - 2W^T(\delta - X\beta)(\delta - X\beta)^T(C - W\gamma) \\ \hat{\gamma} &= (W^TW\xi\xi^T\xi + W^T\xi\xi^TW)^{-1}(W^TC\xi\xi^T\xi + W^T\xi\xi^TC)\end{aligned}\quad (32)$$

Equations 31 and 32 are dependent on each other. To solve for $\hat{\beta}$ and $\hat{\gamma}$, we use Alternating Least Squares (ALS).

A.1.3 The Jacobian in the Likelihood Function

To derive the Jacobian component in the likelihood function, we will use a triangular system to estimate the total derivatives: (1) Use the supply side first order conditions (FOCs) to get derivatives with respect to price, mean utility, and costs, (2) get the price total derivatives using the Implicit Function Theorem (IFT) on the FOCs, (3) get the share total derivatives using the price total derivatives and the share partial derivatives.

Getting the supply side FOCs. Using the Bertrand-Nash equilibrium, we assume that firms act using the following objective:

$$\max_{p_j} \pi^{f,t} = \sum_{j \in N_{ft}} (p_{j,t} - c_{j,t})s_{j,t}. \quad (33)$$

For the rest of this derivation, we drop the market notation (all math below is for a single market). The product specific first order condition is:

$$F_j^f = \frac{\partial \pi^f}{\partial p_j} = s_j + \sum_{k \in F(j)} (p_k - c_k) \frac{\partial s_k}{\partial p_j} \quad (34)$$

Next, we take the derivative of the FOCs with respect to price, mean utility, and costs:

$$\frac{\partial F_j^f}{\partial p_l} = \frac{\partial s_j}{\partial p_l} + \underbrace{\frac{\partial s_k}{\partial p_l}}_{\text{if } k \in F(j)} + \underbrace{\sum_{k \in F(j)} (p_k - c_k) \frac{\partial s_k}{\partial p_j \partial p_l}}_{\Xi_{j,l}^{p,p}} \quad (35)$$

$$\frac{\partial F_j^f}{\partial \delta_l} = \frac{\partial s_j}{\partial \delta_l} + \underbrace{\sum_{k \in F(j)} (p_k - c_k) \frac{\partial s_k}{\partial p_j \partial \delta_l}}_{\Xi_{j,l}^{p,\delta}} \quad (36)$$

$$\frac{\partial F_j^f}{\partial c_l} = - \underbrace{\frac{\partial s_l}{\partial p_j}}_{\text{if } l \in F(j)}. \quad (37)$$

We can write this in matrix notation, where the partial derivative of vector y with respect to vector x can be denoted as $J^{y,x}$. Additionally, as noted below the terms above, the Ξ terms are sparse Hessians (sparsity comes from the $k \in F(j)$). With these two notational changes in mind, we can write the above equations in a simpler form:¹¹

$$J^{F,p} = J^{s,p} + O \circ J^{s,p} + \Xi^{p,p} \quad (38)$$

$$J^{F,\delta} = J^{s,\delta} + \Xi^{p,\delta} \quad (39)$$

$$J^{F,c} = -O \circ J^{s,p}, \quad (40)$$

where \circ is the Hadamard (i.e., element-wise) matrix-product operator. This gives us the variables of interest. Formulae for the Hessian tensors in Ξ will be derived later in the last paragraph of this section.

Getting the price derivatives. To get the price derivatives we will use the IFT on the first derivative of the supply side equation as follows:

$$\frac{\partial \pi_{ft}}{\partial p_{jt}} = F_{fjt}(\delta_t, p_t, c_t) \equiv \sum_{j' \in N_{ft}} (p_{j't} - c_{j't}) \frac{\partial s_{j't}}{\partial p_{jt}} + s_{jt} 1_{j \in N_{ft}} = 0. \quad (41)$$

The IFT implies:

$$\left[\frac{dp_t}{d\delta_t} \mid \frac{dp_t}{dc_t} \right] = -(J_t^{F,p})^{-1} \left[J_t^{F,\delta} \mid J_t^{F,c} \right] \quad (42)$$

¹¹Note that we can do this since $J^{s,p}$ is symmetric.

Applying the IFT to the $\mathcal{J}_t \times 2\mathcal{J}_t$ matrix formed by stacking the matrices next to each other gives:

$$\begin{aligned} \left[\frac{dp_t}{d\delta_t} \mid \frac{dp_t}{dc_t} \right] &= - \left(\frac{\partial F_t}{\partial p_t} \right)^{-1} \left[\frac{\partial F_t}{\partial \delta_t} \mid \frac{\partial F_t}{\partial c_t} \right] \\ &= - \left[J_t^{s,p} + O_t \circ J_t^{s,p} + \Xi_t^{p,p} \right]^{-1} \left[J_t^{s,\delta} + \Xi_t^{p,\delta} \mid -O_t \circ J_t^{s,p} \right], \end{aligned} \quad (43)$$

which is a $\mathcal{J}_t \times (2\mathcal{J}_t)$ matrix with derivatives of all the prices (rows) in δ (first set of \mathcal{J}_t columns) and c (last set of \mathcal{J}_t columns) respectively.

Share derivatives. Using the derived values above, we can calculate the last two entries of the Jacobian matrix of interest, the full derivatives of shares with respect to mean utility and costs:

$$\frac{ds_t}{d\delta_t} = \frac{\partial s_t}{\partial \delta_t} + \frac{\partial s_t}{\partial p_t} \frac{dp_t}{d\delta_t} \quad (44)$$

$$\frac{ds_t}{dc_t} = \frac{\partial s_t}{\partial p_t} \frac{dp_t}{dc_t} \quad (45)$$

$$(46)$$

where the derivatives of p w.r.t. to exogenous variables are given by the IFT above. Now, we have a way of getting the Jacobian matrix for the likelihood:

$$J = \begin{bmatrix} \frac{ds}{d\delta} & \frac{ds}{dc} \\ \frac{dp}{d\delta} & \frac{dp}{dc} \end{bmatrix} \quad (47)$$

Deriving the Hessians. From above, we need to calculate the following Hessian terms: $H_{k,j,l}^{s,p,p} = \frac{\partial s_k}{\partial s_j \partial s_l}$ and $H_{k,j,l}^{s,p,\delta} = \frac{\partial s_k}{\partial s_j \partial \delta_l}$ for each market. We can write shares for product k as:

$$s_k = \int \frac{e^{V_{ki}}}{1 + \sum_{w \in \mathcal{J}_t} e^{V_{wi}}} d\epsilon_{ki} \quad (48)$$

$$= \int s_{ki} d\epsilon_{ki}, \quad (49)$$

which we can use to get the derivative of shares with respect to price:

$$\frac{\partial s_k}{\partial p_j} = \int \frac{\partial V_{ji}}{\partial p_j} \left(\underbrace{s_{ki}}_{\text{if } k=j} - s_{ki} * s_{ji} \right) d\epsilon_{ki}. \quad (50)$$

Taking a second derivative with respect to prices:

$$\frac{\partial s_k}{\partial p_j \partial p_l} = \int \frac{\partial V_{ji}}{\partial p_j} \frac{\partial V_{li}}{\partial p_l} \left(\underbrace{s_{ki}}_{\text{if } k=j=l} - \underbrace{s_{ki}s_{li}}_{\text{if } k=j} - \underbrace{s_{ki}s_{ji}}_{\text{if } j=l} - \underbrace{s_{ki}s_{ji}}_{\text{if } k=l} + 2s_{ki}s_{ji}s_{li} \right) d\epsilon_{ki}. \quad (51)$$

$H_{k,j,l}^{s,p,\delta}$ can be derived similarly:

$$\frac{\partial s_k}{\partial p_j \partial \delta_l} = \int \frac{\partial V_{ji}}{\partial p_j} \frac{\partial V_{li}}{\partial \delta_l} \left(\underbrace{s_{ki}}_{\text{if } k=j=l} - \underbrace{s_{ki}s_{li}}_{\text{if } k=j} - \underbrace{s_{ki}s_{ji}}_{\text{if } j=l} - \underbrace{s_{ki}s_{ji}}_{\text{if } k=l} + 2s_{ki}s_{ji}s_{li} \right) d\epsilon_{ki}. \quad (52)$$

A.2 Integrating market shares

Accurate integration is crucial for optimization. One option is Monte Carlo integration, which uses randomly drawn nodes with equal weights of $1/\text{draws}$. A more precise and computationally efficient way is Gauss-Hermite quadrature, which deals with integrals of the form $\int e^{-x^2} f(x) \approx \sum_{i=1}^n w_i f(x_i)$. For the multinomial logit setup of BLP, Skrainka and Judd (2011) propose the quadrature setup:

$$s_{jt} = \int \frac{e^{\delta_{jt} + x_{jt}\sigma_\beta v_{it}}}{1 + \sum_{j' \in \mathcal{J}_t} e^{\delta_{j't} + x_{j't}\sigma_\beta v_{it}}} f(v_{it}) dv_{it} \quad (53)$$

where v_{it} is a normal variable with an identity for diagonal. The Gauss-Hermite estimation of this integral can be calculated once $f(v_{it})$ is replaced with the normal distribution formula and a change of variables setup is used to satisfy the form of $\int e^{-x^2} f(x)$:

$$s_{jt} \approx \pi^{-\frac{K}{2}} \sum_{i=1}^n w_i \frac{e^{\delta_{jt} + x_{jt}\sigma_\beta v_{it}}}{1 + \sum_{j' \in \mathcal{J}_t} e^{\delta_{j't} + x_{j't}\sigma_\beta v_{it}}}. \quad (54)$$

A.3 Standard Error Estimation

We refer the reader to Conlon and Gortmaker (2020) for GMM standard error estimation. For MLE, we have:

$$\sqrt{n}(\hat{\Theta} - \Theta_0) \rightarrow N(0, I^{-1}) \quad (55)$$

for I the Fisher Information matrix:

$$I = -E \left[\frac{\partial^2 \ell(\mathcal{D}; \Theta)}{\partial \Theta \partial \Theta'} \right] \quad (56)$$

We estimate this as follows:

$$\hat{I} = -\frac{\partial^2}{\partial\Theta\partial\Theta'} \left(\frac{1}{T} \sum_t \ell(\mathcal{D}_t; \Theta) \right) \quad (57)$$

where we have interchanged expectation and derivative, and replaced the expectation with its empirical analog. We compute derivatives using automatic differentiation.