

Romeo, Charles

Working Paper

Incorporating Prior Information into a GMM Objective for Mixed Logit Demand Systems

Economic Analysis Group Discussion Paper, No. EAG 12-1

Provided in Cooperation with:

Expert Analysis Group (EAG), Antitrust Division, United States Department of Justice

Suggested Citation: Romeo, Charles (2012) : Incorporating Prior Information into a GMM Objective for Mixed Logit Demand Systems, Economic Analysis Group Discussion Paper, No. EAG 12-1, U.S. Department of Justice, Antitrust Division, Economic Analysis Group (EAG), Washington, DC

This Version is available at:

<https://hdl.handle.net/10419/202394>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Incorporating Prior Information into a GMM Objective for

Mixed Logit Demand Systems*

Charles J. Romeo
US Department of Justice
450 5th Street, NW–Suite 9300
Washington, DC 20530

charles.romeo@usdoj.gov
(202) 307-5860

First Draft: February 2012

*The Antitrust Division encourages independent research by its economists. The views expressed herein are entirely those of the author and are not purported to reflect those of the US Department of Justice.

Incorporating Prior Information into a GMM Objective for Mixed Logit Demand Systems

Charles J. Romeo
US Department of Justice
600 E Street, NW–Suite 10000
Washington, DC 20530

charles.romeo@usdoj.gov
(202) 307-5860

Abstract

It is well known that random parameters specifications can generate upward sloping demands for a subset of products in the data. Nevo (2001), for example, found 0.7 percent of demands to be upward sloping. Possibly less well known is that demand system estimates can imply margins outside of the theoretical bounds for profit maximization. If such violations are numerous enough, they can confound merger simulation exercises. Using Lerner indices for multiproduct firms playing static Bertrand games, we find that up to 35 percent of implied margins for beer are outside the bounds. We characterize downward sloping demand and the theoretical bounds for profit maximization as prior information and extend the GMM objective function, incorporating inequality moments for product-level own-elasticities and brand-level or product-level Lerner indices. These moments impose a cost when the inequality is violated, and equal zero otherwise. Very few violations remain when an inequality constrained estimator is used. Importantly, the unconstrained GMM objective has multiple minima, while the constrained objective has only one minimum when the product-level constraints are used in our illustration. This is valuable for policy purposes as it enables one to limit attention to a single theoretically consistent model. Inputs to merger simulations are likewise consistent with economic theory, and, as a result, confidence in the output is increased.

In a second innovation, this paper introduces merger simulation for static Stackelberg price competition games. Our illustration uses beer data, a perfect vehicle for introducing Stackelberg games as the economics literature and industry trade press have long considered Anheuser-Busch to be the industry price leader. We find evidence of positive pre-merger price conjectures consistent with beer brands being strategic complements. Allowing the leader to update their conjectures in response to a merger provides dramatically different post-merger price and share changes relative to Bertrand. The Stackelberg conjectures are used as a strategic tool that allows post-merger product repositioning unavailable under Bertrand.

Keywords: Random Coefficients Logit; Inequality Constrained Optimization; Merger Simulation; Stackelberg Pricing

I. Introduction.

We introduce mixed logit estimation under inequality constraints in order to place a cost on parameters falling outside of the region that is consistent with economic theory. The inequalities we propose appending to the GMM objective function are functions of interest at the level of individual data records or at the level of brands. Two sets of inequalities are the focus of this paper: own-elasticities and margins implied by the demand system and pricing game.

The idea of using implicit or explicit constraints in demand estimation is not new. Continuous choice demand systems have long been tested for, or had imposed in estimation, agreement with adding up, homogeneity and symmetry,¹ and a small theoretical coherency literature has developed to incorporate inequality constraints for monotonicity, concavity and other properties implied by theory on estimated systems (see e.g., Chua, Griffiths and O'Donnell (2001), van Soest, Kapteyn and Kooreman (1993), and Kooreman (1990)).

What I offer brings inequality constraints to the mixed logit demand systems while expanding both the scale and the scope of approaches applied to continuous demand systems. To provide a sense of scale, consider that in continuous choice models (AIDS, translog, etc.), where such testing and restrictions have previously been applied, economic theory constraints generally take the form of equality restrictions on functions of the parameters. These restrictions reduce the dimensionality of the parameter space. By contrast, I propose to place inequality restrictions on functions of interest at the level of individual records, or individual brands. These restrictions do not reduce parameter dimensionality, but rather place a cost on parameters lying in regions of the parameter space where the constraints are violated. For example, a mixed logit demand system produces a different $J \times J$ elasticity matrix for each market-time period, yielding as many own-elasticities as there are data records: datasets with 136,000 observations (as in our illustration), yield 136,000 own-elasticities. These models are complex enough that they sometimes produce a small percentage of positive own-elasticities. Nevo (2000) reported 0.7 percent positive own elasticities. I suggest appending all the own-elasticities onto the GMM objective function to place a cost on violations of demand theory. In this way of thinking, that

¹Chapter 3 of Deaton and Muellbauer (1980) provides a discussion and further references on equality constrained estimation for the AIDS and other models. Jorgenson, Lau and Stoker (1982) develop the equality constrained translog model.

demand curves slope downward is prior information. By making it costly for the model to violate this prior we tighten the region where parameters may lie and make it more likely to produce results that are consistent with theory.

To provide a sense of scope, consider the matrix form of the Lerner index for multi-product firms. The key empirical content of the Lerner index is that as long as $\text{price} \geq \text{marginal cost} \geq 0$, the inverse own-elasticity is bounded between 0 and 1. The Lerner index for multi-product firms also provides 0-1 bounds on a function of inverse own- and cross-elasticities for each product. We can estimate a demand model and evaluate how well the margins derived from it conform with Bertrand profit maximization by checking whether the multi-product Lerner index satisfies the 0-1 bounds for all, or a subset, of brands or products in the data. Alternatively we can append the bounds on the GMM objective function to raise the cost of violating Bertrand profit maximization.

Appending the Lerner index bounds on the objective function broadens the scope of previous approaches in that it brings margins directly into the demand objective.² This approach imposes both less and more targeted supply information on demand model estimates than does jointly estimating demand and supply. Appending 0-1 bounds imposes less supply information in that the only component of the supply side that enters are the bounds, and these bounds equal zero unless they are violated. Again, the bounds are evaluated relative to our prior information: a counterfactual Bertrand merger simulation implies a Bertrand prior. If the estimates are consistent with our prior, the bounds shadow price will be zero. In a typical unconstrained optimization with the beer data used in this paper I found that as many as 35 percent of more than 136,000 multiproduct Lerner bounds for beer products are violated with no prior information imposed, while only on the order of one hundred violations remained once the bounds were in place.

Moreover, using bounds targets supply information to the purpose of producing

²Continuous demand approaches incorporated inequality monotonicity constraints on demand that came from a cost function. The cost function, however, was not a structural supply function, but the expenditure form of demand (see e.g., Ryan and Wales (1998) and Moschini (1999)).

consistency with economic theory, and in this sense they may perform better than jointly modeling demand and supply. We include one joint demand-supply model in our illustration. It produces violations of the multiproduct Lerner bounds at about a 15 percent rate. This is much better than when modeling demand alone, but not as good as produced by focusing the supply structure into the bounds.

Introducing constraints is conceptually similar to the micro moments introduced by Berry, Levinsohn, and Pakes (2004) and Petrin (2002), both of which used micro moments on consumers second choices to provide more information on substitutes considered by car buyers. Both concepts bring more information to bear on the demand model to improve the ability of the model to fit moments of interest. A handful of other papers have followed that mix micro and aggregate moments (Hendel and Nevo (2006a,b) and Chintagunta and Dube (2005)), but a lack of readily available micro data has limited the breadth of these applications. By contrast, the prior information we use is available for all applications.

Not surprisingly, improving the theoretical cohesion of demand model estimates to the proposed supply side structure improves the theoretical consistency of counterfactual exercises. We assess Compensating Marginal Cost Reductions (CMCRs), Werden (1996), and simulate mergers for both Miller-Coors and Anheuser Busch-InBev and find that the constraints yield improvements in the margins that we estimate and hence the marginal costs that we back out of the demand model: virtually all margins are in $[0,1]$ and marginal costs are in the range $[0, \text{price}]$. This improves our confidence in the CMCRs and merger simulations.

Other consistency evaluations and constraints are possible. Checking for negative cross-elasticities is straightforward, and imposing that all products be substitutes is possible if the estimation is producing complements among products where it is difficult to formulate a reasonable argument in support of complements. If collusive pricing is a consideration for a particular industry, demand consistency can be evaluated relative to this pricing game. Evaluating consistency with and imposing Stackelberg pricing is possible in principle, but in practice the cost of evaluating the leaders reaction to followers pricing is infeasible for large product level datasets. Given this limitation, we conduct Stackelberg merger simulations using demand estimates obtained under Bertrand margin constraints.

We introduce Stackelberg merger simulation and find that allowing Anheuser-Busch to be the Stackelberg price leader produces different post-merger predictions relative to Bertrand. This is true whether Anheuser-Busch leads with its flagship Budweiser and Bud Light brands, or with its entire family of brands. Pre-merger, that Anheuser-Busch recognizes beer brands to be strategic complements shows up as positive Stackelberg conjectures. Many conjectures shift in response to the merger; some decrease, reducing some prices and increasing shares of at least some Anheuser-Busch products. It is reasonable to characterize Stackelberg conjectures as a tool that enables Anheuser-Busch to reposition brands post-mergers to a degree unavailable to Bertrand competitors.

This paper proceeds by developing the demand model and the constraints in Section 2. Merger simulation for Stackelberg pricing and implementation issues are discussed in Section 3. In Section 4 we develop the inequality constrained GMM objective function. Section 5 introduces our dataset and issues faced in estimation. Results are in Section 6; conclusions are in Section 7.

2. Demand model and constraints

We represent the indirect utility of a consumer i in market m and time period t with preferences $(a_{im}, v_{im}, \varepsilon_{ijmt})$ from the purchase of one unit of the j th product as

$$\begin{aligned}
 u_{ijmt}(\psi) &= x_{1jmt}\theta + [x_{2jmt}p_{jmt}]\alpha_{im} + \xi_{jmt} + \varepsilon_{ijmt}, \\
 \alpha_{im} &= \bar{\alpha} + \Gamma a_{im} + Yv_{im}, \\
 i &= 1, \dots, I, \quad j = 0, \dots, J, \quad m = 1, \dots, M, \quad t = 1, \dots, T,
 \end{aligned} \tag{1}$$

where we decompose the vector of product characteristics $x_{jmt} = \{x_{1jmt}, x_{2jmt}\}$ into those with fixed coefficients, x_{1jmt} , and those with random coefficients, x_{2jmt} . Decomposing x_{jmt} emphasizes the fact that we are not likely to associate random coefficients with each product characteristic in part because of our inability to identify that many random coefficients, and in part because some

product characteristics are not good candidates for random coefficients.³

ψ is the vector of utility function parameters. ξ_{jmt} represents product characteristics that are observed by market participants, but unobserved by the econometrician. Assuming prices are set strategically, each p_{jmt} will be correlated with all $J+1$ ξ_{jmt} in each market-time period m, t . The ε_{ijmt} are i.i.d. type 1 extreme value errors.

Setting the dimension of α_{im} at $K \times 1$, the second equation in (1) is a $K \times 1$ vector of hierarchical regressions of α_{im} on $D \times 1$ vector of demographic draws a_{im} , and v_{im} a vector of standard normal draws. Γ and Y are matrices of unknown parameters: Γ is $K \times D$ with $r \leq K \times D$ nonzero elements and Y is diagonal $K \times K$.

Normalize u_{0jmt} to zero and assume that each consumer maximizes utility by purchasing one unit of product j at time t if and only if $u_{ijmt} \geq u_{irmt}$, $r = 0, \dots, J$. Then the set of consumers choosing product j in market-time period m, t is represented by

$$A_{jmt} = \{(a_{im}, v_{im}, \varepsilon_{ijmt}) : u_{ijmt} \geq u_{irmt}, r = 0, \dots, J\}.$$

Assuming that there are no ties, that a , v , and ε are mutually independent, and given that we have endowed ε with an i.i.d. type 1 extreme value distribution, yields an expression for market shares of the form

$$\begin{aligned} s_{jmt}(\psi) &= \int \int \frac{\exp(\delta_{jmt} + \mu_{ijmt})}{1 + \sum_r \exp(\delta_{rmt} + \mu_{irmt})} dP_0(a|m) dP_0(v) \\ &= \int \int f_{jmt}(\psi) dP_0(a|m) dP_0(v). \end{aligned} \tag{2}$$

where $\delta_{jmt} = x_{1jmt}\theta + [x_{2jmt}, p_{jmt}]\bar{\alpha} + \xi$, and $\mu_{ijmt} = [x_{2jmt}, p_{jmt}](\Gamma a_{im} + Yv_{im})$. Below we specify a as a discrete distribution of market level demographics from the American Community Survey.

³For example, firms might price discriminate by offering lower per unit prices for larger package sizes, but since there is little variation in available package sizes across markets and time periods the econometrician will not likely be able to explain heterogeneous responses to package size.

a. Constraints on demand

The Lerner index for a Bertrand-Nash equilibrium is straightforward to generalize to the multi-product multi-brand case. Dropping the m and t subscripts, let $P = \text{diag}\{p_1, \dots, p_J\}$, $S = \text{diag}\{s_1, \dots, s_J\}$, and let O be a $J \times J$ matrix of 0s and 1s such that $O(j, r)$ equals one when products j and r are owned by the same firm, and zero otherwise. Letting p , s and mc be $J \times 1$ vectors of prices, shares and marginal costs respectively we represent the first order conditions from profit maximization in terms of optimized margins as

$$\begin{aligned} P^{-1}(p - mc) &= -P^{-1} \left[\left(\frac{\partial S}{\partial p} \right)' * O \right]^{-1} s \\ &= -P^{-1} S^{-1} [E' * O]^{-1} P s, \end{aligned} \quad (3)$$

where “*” represents element by element matrix multiplication, and $E = [e_{jr}]$ is the $J \times J$ matrix of mixed logit elasticities where,

$$e_{jr} = \begin{cases} \frac{p_j}{s_j} \int f_j(\psi) (1 - f_j(\psi)) \partial f_j / \partial p_j P_0(d(a, v)) & r = j \\ -\frac{p_r}{s_j} \int f_j(\psi) f_r(\psi) \partial f_j / \partial p_r P_0(d(a, v)) & r \neq j. \end{cases} \quad (4)$$

Assuming that $p_j \geq mc_j \geq 0$ for each product j , (3) yields J margin conditions of the form

$$0 \leq (p_j - mc_j)/p_j \leq 1 \Rightarrow 0 \leq -(e^{jj} + \sum_{k \in \mathcal{J}} e^{jk} \frac{p_k s_k}{p_j s_j}) \leq 1, \quad (5)$$

where $E^{-1} = [e^{jr}]$, and where we represent by \mathcal{J} the set of products sold by the firm that sells product j . For firms that sell a single product, the right hand side of (5) simplifies to that of the familiar Lerner index, $0 \leq -e^{jj} \leq 1$. In (5) there are as many margin conditions as there are products in the dataset. In some product markets, it may be reasonable to expect that firms focus more on margins at the brand level, possibly using loss leaders, pricing single sizes of a brand

below marginal cost during some sales to build market share. Hence, the above conditions may be too strong. To reset them to the brand level, take the share weighted average of (5) for the set of products within the brand that includes product j , say B , where $B \in \mathcal{F}$

$$0 \leq \sum_{j \in B} \omega_j (p_j - mc_j) / p_j \leq 1 \Rightarrow 0 \geq -\sum_{j \in B} \omega_j (e^{jj} + \sum_{k \in \mathcal{F}} e^{jk} \frac{p_k s_k}{p_j s_j}) \leq 1, \quad (6)$$

where $\omega_j = s_j / \sum_{k \in B} s_k$.

To form elasticity constraints, we use (4), bounding $e_{jj} \leq 0$ for all j . For margin constraints, we bound the right hand sum in either (5) or (6) by $[0,1]$. Both sets of conditions can be stacked as additional moment conditions on the GMM objective. The elasticity condition imposes a cost on upward sloping demand curves, while the margin condition imposes a cost on violations of Bertrand-Nash equilibrium.

There are two further points worth noting about the margin conditions. First, imposing these margin conditions on the objective function is weaker and more targeted than jointly estimating demand and supply. It imposes the supply structure only when profit maximization bounds are violated, and it targets the supply structure on producing theoretical consistency. As we show below, in results obtained without the constraints imposed, on the order of 35 percent of margins violate the bounds. Second, in principle, one is not limited to Bertrand competition when specifying margin constraints. A Stackelberg price leadership profit model is also bounded in $[0,1]$. Likewise, constraints can be imposed assuming a multilevel supply structure as in Villas-Boas (2007) or Yang, Chen, and Allenby (2003). In both the Stackelberg and multilevel case, however, the computational cost of evaluating the reaction functions is quite high, making estimation under either of these constraints infeasible except in small datasets.

3. Merger simulation

Two topics are discussed in this section. Merger simulation for a Stackelberg equilibrium, and some technical details of our application.

a. Stackelberg merger simulation

The economics literature on beer has identified Anheuser-Busch as a price leader. Greer (1998) and Tremblay and Tremblay (2005) provide histories of Anheuser-Busch's price leadership and competitor acquiescence as reported in the industry trade press. Rojas (2008) uses the doubling of beer excise taxes in 1991 as a natural experiment to try and identify whether Stackelberg, Bertrand, or other collusive pricing games better fit brewers' responses to this cost shock.

Limited, as we are, to data that precedes both the Miller/Coors and AB/InBev mergers, we cannot compare predictions of the Bertrand and Stackelberg pricing games to post-merger price increases, but we can present predictions on how Anheuser-Busch's strategic response differs pre- and post-simulated merger for both pricing games.

In the Stackelberg pricing game, the first order conditions for profit maximization for followers have a standard Bertrand form, while the leaders first order conditions contain its conjecture as to the followers response. Leaders and followers first order conditions can be written,

$$\begin{aligned}\frac{\partial \pi_{\mathcal{J}^L}}{\partial p_r^L} &= s_r^L + \sum_{j \in \mathcal{J}^L} (p_j^L - mc_j^L) \left[\frac{\partial s_j^L}{\partial p_r^L} + \sum_{\ell \in \mathcal{J}^F} \frac{\partial s_j^L}{\partial p_\ell^F} \cdot \frac{\partial p_\ell^F}{\partial p_r^L} \right] \\ \frac{\partial \pi_{\mathcal{J}^F}}{\partial p_r^F} &= s_r^F + \sum_{j \in \mathcal{J}^F} (p_j^F - mc_j^F) \frac{\partial s_j^F}{\partial p_r^F},\end{aligned}\tag{7}$$

where superscript L and F are used to identify the leader and followers prices and shares respectively, with \mathcal{J}^L and \mathcal{J}^F identifying leading and following firms. We assume that the leader leads with J^L products. Though we do not extend our notation to allow for the possibility that the J^L is a subset of the lead firm's products, we allow for this in our analysis.

The two sets of first order conditions can be stacked, written in matrix form as

$$p - mc = - \left[\frac{\partial s}{\partial p} * O + R \right]^{-1} s,\tag{8}$$

where

$$R = \begin{bmatrix} \left(\frac{\partial s^L}{\partial p^F} \cdot \frac{\partial p^F}{\partial p^L} \right)' & 0 \\ 0 & 0 \end{bmatrix}.$$

R is a $J \times J$ matrix with the leaders strategic response in the first $J^L \times J^L$ block and zeros everywhere else. Equation (8) is similar to the familiar Bertrand solution, however, the R term introduces substantial additional computational complexity. $\partial s^L / \partial p^F \geq 0$, while $\partial p^F / \partial p^L \leq 0$, determines whether leaders and followers prices are strategic complements, $\partial p^F / \partial p^L > 0$, or strategic substitutes, $\partial p^F / \partial p^L < 0$. $\partial p^F / \partial p^L$ is formulated by totally differentiating the followers first order conditions

$$\begin{aligned} & \frac{\partial^2 \pi_{\mathcal{F}^F}}{\partial p_r^F \partial p_m^F} dp_m^F + \frac{\partial^2 \pi_{\mathcal{F}^F}}{\partial p_r^F \partial p_n^L} dp_n^L \\ &= \left[\frac{\partial s_r^F}{\partial p_m^F} + \frac{\partial s_m^F}{\partial p_r^F} + \sum_{k \in \mathcal{F}^F} (p_k^F - mc_k^F) \frac{\partial^2 s_k^F}{\partial p_r^F \partial p_m^F} \right] dp_m^F + \\ & \quad \left[\frac{\partial s_r^F}{\partial p_n^L} + \sum_{k \in \mathcal{F}^F} (p_k^F - mc_k^F) \frac{\partial^2 s_k^F}{\partial p_r^F \partial p_n^L} \right] dp_n^L \\ &= 0 \end{aligned} \tag{9}$$

Equation (9) is then solved for $\partial p^F / \partial p^L$.⁴ As (9) indicates, evaluating $\partial p^F / \partial p^L$ is computationally costly and it introduces a complication into merger simulation in that $\partial p^F / \partial p^L$ is an implicit function of marginal cost. Hence, marginal cost estimates can no longer be backed out from observed prices and estimated markups as they can when pricing is Bertrand.

We estimate marginal costs by iterating through a fixed point algorithm. Beginning at $mc^{(0)}$ equal to a vector of Bertrand marginal costs, we update it using

⁴The Appendix in Rojas (2008), provides solution details.

$$mc^{(k+1)} = (1 - \omega)mc^{(k)} + \omega(p - markup^{(k+1)}), \quad (10)$$

where *markup* is the current value of the right hand side in (8). In our experience, setting $\omega = 0.25$ produces a steadily converging sequence of marginal cost estimates.⁵

b. implementation details

We follow the basic approach to merger simulation outlined in Nevo (2000), Peters (2006) and Werden and Froeb (2008); for Bertrand games, we use (3) to back out marginal costs, for Stackelberg games, we solve for marginal costs using (8) and (10). In both games, we then update ownership matrix, O , to its post-merger structure and solve for post-merger equilibrium prices. Beyond this basic approach we discuss two implementation details. First, we discuss our preparation of price and share data for input into merger simulation. Second, we discuss additional structure that we found it necessary to impose on the merger simulation given the large number of products in our product set.

Beer prices and sales, and grocery prices and sales more generally, vary from week-to-week, which raises the question, what are the appropriate vectors of input prices and shares for a merger simulation? We answer this question by forming prices that are 6-month weighted averages, \bar{p}_j , for each product j , using weekly quantities as weights. \bar{p}_j , as such, reflects what the typical beer consumer has paid for a given beer product over the most recent 6-month period. For market shares, we use simple 6-month averages, \bar{s}_j . Given pairs of $J \times 1$ vectors (\bar{p}, \bar{s}) we employ the Berry (1994)/Berry, Levinsohn and Pakes (1995) (henceforth BLP) contraction mapping, solving for δ by equating model shares to \bar{s} .

We found that merger simulations which allowed the entire vector of prices, \bar{p} , to update in response to an ownership change were poorly determined. Different starting points to merger simulation often produced wildly different results, and the simulations had a tendency to crash. While theory does not constrain Bertrand optimization with random coefficients logit demand to

⁵We have tested larger values of ω and found that the algorithm diverged in some tests. Using $\omega = 0.25$ typically produces convergence in less than 20 iterations under the criteria $\max\{|mc^{(k+1)} - mc^{(k)}|\} \leq 0.001$.

have a single optimum,⁶ the wild behavior of the merger simulations suggested the need for additional structure. We imposed a restriction that required price changes for all products within a brand to be the same. We defined $\bar{p}_B = \sum_{j \in B} s_j \bar{p}_j$, and conducted merger simulation on brand level prices, \bar{p}_B , thereby dramatically reducing the dimensionality of the problem to be solved. More specifically, we defined prices pre-merger to be equal to brand level prices plus a difference, η ,

$$\bar{p}_j^{pre} = \bar{p}_B^{pre} + \eta_j^{pre}, j \in B.$$

The simulator updates \bar{p}_B and yields post-merger prices,

$$\bar{p}_j^{post} = \bar{p}_B^{post} + \eta_j^{pre}, j \in B.$$

η remains fixed at its pre-merger values restricting post-merger prices for individual products, \bar{p}_j^{post} , to retain the same absolute relationship to each other that they had pre-merger. These restrictions reduced the number of prices to be solved for from more than 150 to 30 and greatly improved merger simulation performance.

4. Estimation and Inference

We estimate demand using the Nested Fixed Point (NFP) algorithm developed in Berry (1994) and BLP.⁷ Without the constraints in place, the objective function is the standard BLP concentrated GMM function of the nonlinear parameters $\lambda = (\Gamma, Y)$, with $\lambda \in \Lambda = \Re^r \times \Re_+^K$.

$$Q(\lambda) = \xi(\lambda)' Z' W Z \xi(\lambda), \quad (11)$$

where W is any symmetric positive definite weighting matrix. We formulate the constraints to

⁶See Caplin and Nalebuff (1991).

⁷The MPEC algorithm of Dube, Fox, and Su (2011) is an alternative approach.

have smooth first derivatives so as to satisfy standard regularity conditions for optimization and asymptotic normality in Hansen (1982). The elasticity and margin constraints are specified as

$$\|E(\boldsymbol{\varepsilon};\boldsymbol{\eta},\boldsymbol{\lambda})\| = I(\boldsymbol{\varepsilon};\boldsymbol{\eta},\boldsymbol{\lambda})'(\boldsymbol{\varepsilon}(\boldsymbol{\lambda})^2 + 1)^{-5},$$

and

$$\|A(\boldsymbol{m};\boldsymbol{\eta},\boldsymbol{\lambda})\| = \{(1 - I(\boldsymbol{m};\boldsymbol{\eta},\boldsymbol{\lambda})^2)'[\#(\boldsymbol{k}\boldsymbol{\varepsilon}\boldsymbol{B}) * (\boldsymbol{m}(\boldsymbol{\lambda})^2 + 1)] + I(\boldsymbol{m}-1;\boldsymbol{\eta},\boldsymbol{\lambda})^2'[\#(\boldsymbol{k}\boldsymbol{\varepsilon}\boldsymbol{B}) * ((\boldsymbol{m}(\boldsymbol{\lambda})-1)^2 + 1)]\}^{-5},$$

respectively, where $I(\boldsymbol{\omega};\boldsymbol{\eta},\boldsymbol{\lambda}) = 1/(1 + \exp(-\boldsymbol{\eta}\boldsymbol{\omega}(\boldsymbol{\lambda})))^{-5}$. With $\boldsymbol{\eta}$ large enough, $I(\cdot;\boldsymbol{\eta},\boldsymbol{\lambda})$ resembles a smoothed indicator function; together $I(\cdot;\boldsymbol{\eta},\boldsymbol{\lambda})$ and $1 - I(\cdot;\boldsymbol{\eta},\boldsymbol{\lambda})$ give $A(\boldsymbol{m};\boldsymbol{\eta},\boldsymbol{\lambda})$ a bathtub shape which turns off the constraint for values of values of $\boldsymbol{m} \in [0,1]$ while $E(\boldsymbol{\varepsilon};\boldsymbol{\eta},\boldsymbol{\lambda})$ is turned off for $\boldsymbol{\varepsilon} \leq 0$.⁸ We add one to the elasticity and margin values, e.g., $\boldsymbol{\varepsilon}(\boldsymbol{\lambda})^2 + 1$, to increase the value of the constraints, giving them more weight in the optimization. When the margin constraint is applied at the brand-level, $\#(\boldsymbol{k}\boldsymbol{\varepsilon}\boldsymbol{B})$ captures the cardinality of the number of products within a brand. This varies across brands, markets and time periods. When used at the product-level, $\#(\boldsymbol{k}\boldsymbol{\varepsilon}\boldsymbol{B})$ is a vector of ones. Satisfying the constraints restricts $\boldsymbol{\lambda}$ to a subset of Λ : $\boldsymbol{\lambda} \in \Lambda^*$, where $\Lambda^* = \{\boldsymbol{\lambda}: E(\boldsymbol{\varepsilon};\boldsymbol{\eta},\boldsymbol{\lambda}) = 0 \text{ and } A(\boldsymbol{m};\boldsymbol{\eta},\boldsymbol{\lambda}) = 0\}$.

The constrained GMM optimization problem takes the form

$$C(\boldsymbol{\lambda}) = \boldsymbol{\xi}(\boldsymbol{\lambda})'Z'WZ\boldsymbol{\xi}(\boldsymbol{\lambda}) + \|E(\boldsymbol{\varepsilon};\boldsymbol{\eta},\boldsymbol{\lambda})\|^2 + \|A(\boldsymbol{m};\boldsymbol{\eta},\boldsymbol{\lambda})\|^2. \quad (12)$$

We formulate the constraints as squared Euclidean norms as they are a natural formulation for constructing the variance estimator for the constrained model.⁹ In both the unconstrained and constrained formulations, for the first stage GMM estimates we set $W = (Z'Z)$ and use Newey and (1987) West's heteroscedastic and autocorrelation consistent estimates for our second stage

⁸We set $\boldsymbol{\eta} = 50$ in all estimation work. At this setting, $I(-0.1;\boldsymbol{\eta},\cdot)^2 = 0.007$, while $I(0.1;\boldsymbol{\eta},\cdot)^2 = 0.993$. As such, it is a slight abuse of language to say that the constraint gets turned off for $\boldsymbol{m} \in [0,1]$ and $\boldsymbol{\varepsilon} \leq 0$. More accurately, the cost of the constraint is rapidly diminishing in these ranges.

⁹The variance function is formulated in the Appendix.

W.

5. The Data

Development of the outside share and discussion of instrumental variables are contained in Appendix B in associate file Romeo (2012).

We use 26 weeks of Nielsen grocery store scanner data, covering the first half of 2008, on beer pricing and sales from 37 Nielsen markets. For product characteristics we include variables for promotions (feature, display, feature and display, and discount), and dummy variables for package size (6, 12, 18, 24, and 30 packs), container type (bottles, longnecks, and cans), beer type (light, regular, and ice), package type (freezer packs, non-freezer packs), and brand (Busch, Busch Light, Bud Select, Budweiser, Bud Light, Coors, Coors Light, Corona, Heineken, Keystone, Keystone Light, Labatt's, Labatt's Light, Michelob, Michelob Light, Michelob Ultra, Miller Genuine Draft, Miller Genuine Draft Light, Miller High Life, Miller High Life Light, Miller Lite, Milwaukee's Best, Milwaukee's Best Light, Natural, Natural Light). The preceding list contains all the major beer brands in terms of sales in our 37 markets. In the time period of these data, all of these brands were owned by one of five major brewers: Anheuser-Busch, Coors, Grupo-Modelo, Heineken, and Miller. Together they accounted for more than 83 percent of the total reported quantity sales. The data also includes hundreds of smaller brands in each market-week. We define five aggregate "brands" (Craft, Import, Premium, Sub-Premium, Super Premium) and create share weighted aggregates of the product characteristics for these brands. With these aggregations in place, the data contain an average of 136 products in each of $37 \times 26 = 962$ market-time periods for a total of 130,634 observations.

The promotional variables—discount, feature, display, and feature and display—record for each brand the percentage of stores, weighted by All Commodity Volume (%ACV),¹⁰ in which these brands are on sale. All of these promotional variables are likely to be endogenous. There are two possible sources of endogeneity. First, ξ likely includes unobserved advertising that is correlated with observed promotions. Second, since these variables are only observed at the

¹⁰All Commodity Volume is total sales of all products in each store.

brand level they are really only proxies for the product level promotion variables. Any measurement error in these proxies gets absorbed into the ξ .

These four promotional variables together with price give us a total of five endogenous variables. We assume that product characteristics are exogenous, and employ these as instruments in addition to the following four sets of variables: own-firm other product instruments of the type described by BLP, lagged share differences, lagged discount effectiveness, and the means of included demographics.

5. Results

To begin, we estimate the unconstrained model 25 times, from random standard normal starting points, in order to assess the likely number of minima of the unconstrained objective. This exercise produced four minima.¹¹ The joint demand-supply and the constrained optimizations are then started at these four unconstrained minima. With the product-level constraints applied, the four minima collapsed to a single constrained minimum. This result, in and of itself, is valuable as it enables us, in this instance, to limit attention to a single model with few remaining bounds violations. Beyond that, the parameter estimates and random coefficient functions for the product-level constrained model all agree with economic intuition, something that cannot be said of the unconstrained demand results. The joint demand-supply and brand-level constraint results ended up somewhere in between. They are more economically defensible than the unconstrained minima, the brand-level constrained model more so than the joint model, but there are still four of them in both cases.

Parameter estimates for our “global” unconstrained, joint demand-supply and brand-level constrained minima and the unique product-level minima are in Table 1. We limit attention to estimates of parameters with random coefficients (price, discount, light beer and Corona) and parameters for other endogenous variables (feature, display, and feature and display). Results

¹¹We identify unique minima using an eyeball comparison to group outputs from each optimization, and then verify that our groupings each identify a single minimum by evaluating the within group standard deviation of each parameter. In general standard deviations for each parameter are on the order of 0.001 for estimates classified as having the same minimum.

for the three local unconstrained minima, the three local joint demand-supply and the three local brand-level constrained minima are in Appendix Tables C1 - C3 respectively, for your perusal.¹²

The constraints have a dramatic affect on the estimates. The price coefficients, for example, increase in magnitude as we move from unconstrained demand system to the unconstrained demand-supply system, to the brand-level constrained system, and finally to the product-level constrained system. For the product-level constrained system, the price coefficient has roughly doubled in magnitude relative to either unconstrained system, thereby making demand more elastic in the constrained models. The parameters, not surprisingly, are more precisely determined in the jointly estimated system and in the constrained models; all the parameters in Table 1 are significant at least at the five percent level with the product-level constraints in place. In addition, the coefficients in the model with product-level constraints, in particular, are more consistent with expectations. For example, the mean coefficients for light beer and discount are both positive and significant indicating that marginal effects on utility of light beer and larger discounts are positive. Discounts are a vertical characteristic with more always preferred to less, and light beer is the only domestic beer category exhibiting positive year-to-year growth since the early 1990s suggesting light beer to be utility enhancing.¹³ The brand-level constrained results also shows a positive, though not statistically significant response to discounts, but finds light beer to be utility decreasing. Both unconstrained results indicate that mean utility decreases for both discounts and light beer.

The demographic coefficients are also substantially affected by the constraints. Focusing on the product-level constrained results, the light beer-age interactions show a larger positive affect for both younger age categories indicating both higher levels of utility and heterogeneity from light beer consumption for these age groups. Similarly, the price-age interactions are both negative and significant for this model indicating that demand for beer is more elastic among younger consumers. Price interactions with income show that demand becomes less elastic with

¹²All Appendix Tables are in the associated file, Romeo (2012).

¹³Light beer trends are not presented here, but are found in Section 5 of The U.S. Beer Market: Impact Databank Review and Forecast, 2008 Edition.

income while the discount-income interactions indicate that demand for discounts becomes more elastic with income. The brand-level constrained results, though not as strong statistically, all have the same signs and are similarly intuitive economically. We point to two key differences between the constrained and the unconstrained results. First, that unconstrained results show much weaker price-income heterogeneity. The income² coefficient in the joint model is much smaller than either constrained model, while the demand only model shows demand becoming more elastic with income, though this effect is not statistically significant. Second, the large negative discount-income² coefficient indicates the value to discounts to decline precipitously with income. The following figures make clear how stark these differences are between the different specifications.

Comparisons of the price-income, discount-income and Corona-%Hispanic functions are graphed in Figure 1. Estimated price-income functions are compared in Figure 1a. The function for the unconstrained demand model trends in the wrong direction, though there is no statistically significant price-income heterogeneity. This result is not an outlier, both Nevo (2000,2001) and Villas-Boas (2007) report negative and significant income² coefficients in their random coefficients specifications. In the joint demand-supply model and with either set of constraints in place, demand becomes less elastic with income as one would expect, though this effect is the strongest for the product-level constrained model.

The discount-income comparisons in Figure 1b show that the constraints moderated the heterogeneity in this relationship. All four models show positive utility for low income levels with the affects trailing off at higher levels, and unreasonably the effects of discounts turn negative at higher income levels. For both constrained models the negative downturn is small, while the unconstrained models predicts substantial decreases in utility with discounting at higher income levels. For this large utility decrease to be economically sound, we would have to postulate the existence of something like a “snob effect:” higher income individuals who turn up their noses at a sale.

The Corona-%Hispanic interaction is consistent across all four models: utility of consuming Corona decreases as the percentage of Hispanics in the population increases, and turns negative between about 18 and 22 percent Hispanics in all four models. I cannot offer any

clear insights as to why this occurs. I can only postulate that it may have to do with increased competition from other Latin American beers in markets with a large proportion of Hispanics.

Table 2 contains statistics for each model in Table 1, while Appendix Table C4 contains the statistics for the remaining local minima for all models. Two things are immediately apparent when examining the table: demand becomes more elastic and the number of Lerner violations decrease as we move from left to right across the table: from the unconstrained demand only model, to the joint demand-supply model, and then to the two constrained models. Lerner bound violations decrease from more than 30 percent and 15 percent for the two unconstrained models to 5.5 and then 0.09 percent with the brand- and product-level constraints in place respectively. Under either set of constraints, more than 99 percent of individual product demands are in the elastic range, and the aggregate elasticity is likewise in the elastic range.¹⁴

Given the mean outside share of 61.5 percent used in estimation, all our aggregate inside elasticity estimates are in the midst of market elasticities for beer reported in the literature. Table 2 shows aggregate elasticities ranging from -0.908 for the unconstrained model to -1.135 with the product-level constraints in place. The range of estimated market elasticities spans from -0.1 estimated by Clements and Johnson (1983), to -1.36 in Hausman, Zona, and Leonard (1994). In between are Johnson, et. al. (1992) at -0.31, Pinske and Slade (2004) at -0.5, Lee and Tremblay (1992) at -0.6, and Hogarty and Elzinga (1972) at -0.9. This wide range of results may be attributable to the different modeling approaches and data sources. All of these authors use continuous choice representative consumer demand systems, Hausman, et. al., and Pinske and Slade use multi-stage budgeting, Clements and Johnson, and Johnson et. al. limit attention to pure time-series. As a result either of data limitations or model specification, or both, all of these papers are conducting inference with data that are more highly aggregated than the data we use.

Turning next to merger simulation, Table 3 contains aggregate results for Bertrand and Stackelberg merger simulations for Cincinnati and Buffalo-Rochester. For Cincinnati, we

¹⁴We formulate the aggregate elasticity as follows. Let E_{mt} be the $J_{mt} \times J_{mt}$ elasticity matrix for inside goods in market-time period m, t , let $\mathbf{1}_{mt}$ be a J_{mt} vector of ones, and let \tilde{s}_{mt} be the associated vector of renormalized inside shares: $\tilde{s}_{mt}' \mathbf{1}_{mt} = 1$. The aggregate elasticity, e , is formulated as, $e = (MT)^{-1} \sum_{m,t} \tilde{s}_{mt}' E_{mt} \mathbf{1}_{mt}$.

simulate the Miller/Coors merger, while for Buffalo-Rochester we simulate the Anheuser-Busch/InBev merger.¹⁵ The Miller/Coors merger occurred in June 2008, while the Anheuser-Busch/InBev merger was consummated in September 2009. For the Bertrand pricing game we simulate mergers for our global minima using parameters from both unconstrained models and both constrained models. For Stackelberg, we limit attention to the product-level constrained model, but we follow Rojas (2008) and evaluate merger effects under two leadership scenarios: Anheuser-Busch leads with its flagship Budweiser and Bud Light brands; Anheuser-Busch leads with all of its products.

Focusing on the Bertrand merger simulations in the first four columns, the first two rows for both Cincinnati and Buffalo-Rochester provide city specific aggregate elasticity and marginal cost range violation statistics (these are equivalent to Lerner bound violations) similar to those in Table 2. These tell the same story of increasing demand elasticity and decreasing bounds violations for the joint model and with brand-level and then product-level constraints imposed. The percent pre-merger margins in row three are related to the pattern of aggregate elasticities: as demand become more elastic margins shrink in both markets. The percent CMCR¹⁶ and price change results are less consistent across models. Cincinnati shows the smallest competitive affects with product-level constraints imposed, while Buffalo-Rochester shows the largest CMCR and second largest price change in this case. The details of why this is are buried in the elasticity matrices, key points of which will be discussed below.

Turning to the Stackelberg results, we note that, at the aggregate level, these do not differ substantially from the product-level Bertrand results. The pre- and post-merger leadership strategy results, however, are specific to Stackelberg. All four of these results show that pre-

¹⁵InBev owns Labatts, a Canadian beer that has roughly a 10 percent share in Buffalo-Rochester. We chose Cincinnati because both Miller and Coors have large shares there, and we chose Buffalo-Rochester because that is the US market where Labatt's has its largest share.

¹⁶Mergers are often driven by the potential for merger specific efficiencies. Marginal cost efficiencies reduce upward pressure on prices that results from the merger. Compensating Marginal Cost Reductions, CMCRs, indicate the size of marginal costs efficiencies that would be consistent with zero upward pricing pressure post-merger. Following Werden (1996), we calculate $\%CMCR = 100 * (mc_{post} - mc_{pre}) / p_{pre}$.

merger Anheuser-Busch's prices are strategic complements to follower prices. In three of four cases, strategic complementarity gets stronger post merger. One case, in Buffalo-Rochester, where Anheuser-Busch leads with Budweiser and Bud Light, shows the leader's prices becoming strategic substitutes for the followers prices post merger. Underlying this seemingly odd result is Anheuser-Busch adjusting its strategy post-merger. Explicitly accounting for the leader's conjecture of follower price changes gives the leader more ability to adjust its pricing strategy to the post-merger environment than Bertrand allows. To expand upon this and other points, we turn to brand-level elasticity and Bertrand and Stackelberg merger simulation tables for Buffalo-Rochester in Tables 4-7. Appendix Tables C5-C8 contain a matching set for Cincinnati with discussion.

Table 4 contains a brand-level elasticity matrix for Buffalo-Rochester for the model estimated under product-level constraints.¹⁷ One highlight of this table is simply the extent of cross-elasticity variations within each column. Variations of an order of magnitude or more are typical indicating that the Independence from Irrelevant Alternatives (IIA) property has been substantially weakened by the combination of random coefficients and constraints. Furthermore, many of the largest cross-elasticity pairings are intuitive. Examples are, Bud Light's share responds most strongly to the price of Labatt's Light and vice-versa, while Labatt's share responds most strongly to the prices of Busch and Budweiser. In addition, Labatt's Light's share also responds most strongly to Busch Light, Michelob Light and Coors Light. This mix of strong responses including Anheuser-Busch and Labatt's brands are behind the shift to Budweiser and Bud Light's prices becoming strategic substitutes for follower prices reported above; what we observe as a shift to strategic substitution is the merged firm shifting share to its highest margin products post-merger.

Tables 5-7 contain the Bertrand and two Stackelberg merger simulations for Buffalo-Rochester aggregated to the brand-level. All three merger simulations predict price increases for

¹⁷Brands with share of < 0.50 percent are excluded, where share measures include an outside share = 61.5 percent in Buffalo-Rochester.

most products.¹⁸ In the Bertrand simulation we list the percent CMCRs while in the Stackelberg simulations we list Anheuser-Busch's optimal response to follower price increases pre- and post-merger. The presence of these conjectures produces very different simulation predictions. They predict that Anheuser-Busch will limit price increases on its higher margin products, while increasing the price of lower margin Labatt's products more than they would under Bertrand. This is especially true when Anheuser-Busch leads with all of its products as in Table 7. In this case, we predict that Anheuser-Busch will reduce the price of Busch Light and Bud Light post-merger, each of these having margins in excess of 60 percent, while raising the price of low margin Labatt's Light by 13.87 percent. In total, Anheuser-Busch's share changes indicate that it actually increases its share post-merger from 11.97 percent to 15.71 percent. When Anheuser-Busch leads with only Budweiser and Bud Light, it increases the prices of these two products less than it would under Bertrand, and it does gain some share in four of its brands, but its tools to produce markedly different results are limited. Moreover, casting the conjectures as tools provides an explanation for the strategic complements finding in Table 3.

As shown in Table 3, the overall Anheuser-Busch price increase in Buffalo-Rochester is muted under Stackelberg relative to Bertrand. Having this additional strategic component available in price setting allows them to better optimize their product repositioning post-merger without raising the prices of their highest margin products. The post-merger finding of strategic substitutes in Table 3 can be explained in this context. Decreasing Bud Light's response to followers prices post-merger was a case of effectively using the conjecture on Bud Light as a tool to limit its price increase, so as to limit its share loss.¹⁹

6. Conclusions

We extend the GMM objective function for mixed logit demand systems to include prior

¹⁸Marginal cost efficiencies were not taken into account in conducting these merger simulations.

¹⁹The discussion associated with the matched set of table for Cincinnati, Tables C5-C8, tells a very different story. In this case, Anheuser-Busch's strategic response to the Miller-Coors merger causes its Stackelberg price increases to be above the Bertrand level.

information. Formally, we impose a cost on violations of downward sloping demand curves and on violating brand-level or product-level Lerner bounds for firms playing a static Bertrand pricing game. Incorporating these priors into the objective raises the cost of the objective producing parameter estimates outside the region that is consistent with economic theory. We show that the resulting estimates accord better with economic intuition than either estimating a demand system alone, or jointly estimating demand and structural supply. Moreover, when product-level Lerner bounds are used, the number of minima for the GMM objective shrinks from four to one in our illustration.

We conduct counterfactual merger simulation exercises using Bertrand and Stackelberg pricing games; for Cincinnati we predict effects from the merger of Miller and Coors, and for Buffalo-Rochester we predict effects from InBev's merger with Anheuser-Busch. We find that Bertrand and Stackelberg models predict similar overall price effects. Both sets of merger simulations produced relatively small overall price increases, though in Buffalo-Rochester the Stackelberg estimates were substantially smaller than the Bertrand predictions. Digging into the elasticity matrices and merger simulation results showed why this occurred. The strategy term in the Stackelberg simulation gives Anheuser-Busch additional flexibility to respond to a merger that Bertrand pricing does not allow.

Imposing constraints makes the mixed logit demand system a better policy tool. Limiting attention to the theoretically consistent subset of the parameter space has a long history in demand estimation for continuous choice models for good reason. Working with theoretically consistent results eliminates caveats regarding economically inconsistent results and provides confidence in a policy setting.

Appendix A

A.1. Formulating the Constrained Variance Estimator.

Assume $E[\xi_{jmt}|z_{mt}] = 0$ and $E[\xi'_{jmt}\xi_{jmt}|z_{mt}] = \Omega(z_{jmt})$. Stack the moment conditions that contribute to the constrained objective function and use a delta method expansion in the parameter vector ψ .

$$\begin{bmatrix} z' \xi(\psi) \\ \|E(\epsilon, \cdot; \psi)\| \\ \|A(m, \cdot; \psi)\| \end{bmatrix} - \begin{bmatrix} z' \xi(\tilde{\psi}) \\ \|E(\epsilon, \cdot; \tilde{\psi})\| \\ \|A(m, \cdot; \tilde{\psi})\| \end{bmatrix} = \begin{bmatrix} z' \partial \xi(\tilde{\psi}) / \partial \psi \\ \partial \|E(\epsilon, \cdot; \tilde{\psi})\| / \partial \psi \\ \partial \|A(m, \cdot; \tilde{\psi})\| / \partial \psi \end{bmatrix} (\psi - \tilde{\psi})$$

where $\tilde{\psi} = \arg \min C(\lambda) \equiv C(\lambda(\alpha, \beta)) = C(\psi)$. Multiplying both side by their transpose and rearranging yields

$$Var(\psi) = (G(\tilde{\psi})' W^{-1} G(\tilde{\psi}))^{-1},$$

where $G(\tilde{\psi}) = [(\partial \xi(\tilde{\psi}) / \partial \psi)' Z \quad \partial \|E(\epsilon, \cdot; \tilde{\psi})\| / \partial \psi \quad \partial \|A(m, \cdot; \tilde{\psi})\| / \partial \psi]$ and $W = \text{block diag} \{ \text{Var}(Z' \xi(\psi)), \text{Var}(\|E(\epsilon, \cdot; \psi)\|), \text{Var}(\|A(m, \cdot; \psi)\|) \}$.

References

- Berry, Steven, James Levinsohn, and Ariel Pakes, (1995), "Automobile Prices in Market Equilibrium," Econometrica, 63, 4, 841-90.
- Berry, Steven, (1994), "Estimating Discrete Choice Models of Product Differentiation," RAND Journal of Economics, 25, 2, 242-62.
- Caplin, Andrew, and Barry Nalebuff, (1991), "Aggregation and Imperfect Competition: on the Existence of Equilibrium," Econometrica, 59, 1, 25-59.
- Chua, C. L., W. E. Griffiths and C. J. O'Donnell, (2001) "Bayesian Model Averaging in Consumer Demand Systems with Inequality Constraints," Canadian Journal of Agricultural Economics, 49, 3, 269-91.
- Clements, Kenneth, and Lester Johnson, (1983), "The Demand for Beer, Wine, and Spirits: a Systemwide Analysis," Journal of Business, 56, 3, 273-304.
- Dube, Jean-Pierre, Jeremy Fox, and Che-lin Su, (2011), "Improving the Numerical Performance of Discrete Choice Random Coefficients Demand Estimation, Working Paper, University of Chicago Booth School of Business.
- Greer, Douglas, (1998), "Beer: Causes of Structural Change," in Duetsch, L. (Ed.), Industry Studies, Prentice Hall: New Jersey, 28-62.
- Hansen, Lars, (1982), "Large Sample Properties of Generalized Method of Moments Estimators," Econometrica, 50, 1029-54.
- Hausman, Jerry, Gregory Leonard, and Douglas Zona, (1994), "Competitive Analysis with Differentiated Products," Annales D'Economie et de Statistique, 34, 159-80.
- Hendel, Igal, and Aviv Nevo, (2006a), "Sales and Consumer Inventory," RAND Journal of Economics, 37, 3, 543-61.
- _____ and _____, (2006b), "Measuring the Implications of Sales and Consumer Inventory Behavior," Econometrica, 74, 6, 1637-73.
- Hogarty, Thomas, and Kenneth Elzinga, (1972), "The Demand for Beer," The Review of Economics and Statistics, 54, 2, 195-98.

- Johnson, James, Ernest Oksanen, Michael Veall, and Deborah Fretz, (1992), "Short-Run and Long-Run Elasticities for Canadian Consumption of Alcoholic Beverages: and Error-Correction Mechanism/Cointegration Approach," The Review of Economics and Statistics, 74, 1, 64-74.
- Jorgenson, Dale, Lawrence Lau and Thomas Stoker, (1982), "the Transcendental Logarithmic Model of Aggregate Consumer Behavior," Advances in Econometrics, V1, JAI Press: Greenwich, CT, 97-238.
- Kooreman, Peter, (1990), "Coherency and Maximum Likelihood Estimation in Demand Systems with Binding Inequality Constraints," Advanced Lectures in Quantitative Economics, 605-20.
- Lee, B., and V. J. Tremblay, (1992), "Advertising and the Market Demand for Beer," Applied Economics, 24, 69-76.
- Moschini, Giancarlo, "Imposing Local Curvature Conditions in Flexible Demand Systems," Journal of Business & Economic Statistics, 17, 4, 487-90.
- Nevo, Aviv., (2001), "Measuring Market Power in the Ready-to-Eat Cereal Industry," Econometrica, 69, 2, 307-42.
- _____, (2000), "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry," Rand Journal of Economics, 31, 3, 395-421.
- Newey, Whitney, and Kenneth West, (1987), "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," Econometrica, 55, 3, 703-08.
- Peters, Craig, (2006), "Evaluating the Performance of Merger Simulation: Evidence from the U.S. Airline Industry," Journal of Law and Economics, 49, 2, 627-49.
- Petrin, Amil, (2002), "Quantifying the Benefits of New Products: the Case of the Minivan," Journal of Political Economy, 110,4, 705-29.
- Pinske, Joris, and Margaret Slade, (2004), "Mergers, Brand Competition, and the Price of a Pint," European Economic Review, 48, 617-43.
- Rojas, Christian, (2008), "Price Competition in Brewing," Journal of Industrial Economics, LVI, 1, 1 - 31.

- Ryan, David and Terrence Wales, (1998), "A Simple Method for Imposing Local Curvature in some Flexible Consumer-Demand Systems," Journal of Business & Economic Statistics, 16, 3, 331-38.
- Romeo, Charles, (2012), "Appendices B and C: Incorporating Prior Information into a GMM Objective for Mixed Logit Demand Systems," Economic Analysis Group, US Department of Justice.
- The U.S. Beer Market: Impact Databank Review and Forecast, M. Shanken Communications, Inc.: New York, NY, (2008).
- Tremblay, Victor, and Carol Tremblay, (2005), The U.S. Brewing Industry: Data and Economic Analysis, MIT Press, Cambridge, Massachusetts.
- van Soest, Arthur, Arie Kapteyn and Peter Kooreman, (1993), "Coherency and Regularity of Demand Systems with Equality and Inequality Constraints," Journal of Econometrics, 57, 161-88.
- Villas-Boas, Sofia, (2007), Vertical Contracts between Manufacturers and Retailers: Inference with Limited Data," Review of Economic Studies, 74, 2, 625-52.
- Werden, Gregory, and Luke Froeb, (2008), "Unilateral Competitive Effects of Horizontal Mergers," Handbook of Antitrust Economics, ed. Paolo Buccirossi, Cambridge, MA: MIT Press, 43-104.
- Werden, Gregory, (1996), "A Robust Test for consumer Welfare Enhancing Mergers among Sellers of Differentiated Products," Journal of Industrial Economics, 44, 4, 409-13.
- Yang, Sha, Yuxin Chen and Mark Allenby, (2003), Bayesian Analysis of Simultaneous Demand and Supply," Quantitative marketing and Economics, 1, 251-75.

Table 1. Unconstrained and constrained mixed logit estimates for beer demand

variables	unconstrained global minima				constrained global minima			
	demand only		joint demand and supply		brand-level		product-level	
	means (α, θ)	std devs Y	means (α, θ)	std devs Y	means (α, θ)	std devs Y	means (α, θ)	std devs Y
price	-5.502 (0.417) [†]	0.0008 (0.0006)	-6.986 (0.325) [†]	0.029 (0.007) [†]	-7.663 (0.325) [†]	0.012 (0.005)**	-11.407 (0.147) [†]	0.0007 (8.7e-5) [†]
light	-0.252 (0.064) [†]	0.0002 (0.0001)**	-0.208 (0.022) [†]	0.375 (0.098) [†]	-0.235 (0.138)*	0.002 (0.001) [†]	0.297 (0.097) [†]	0.0006 (0.0001) [†]
Corona	0.533 (0.128) [†]	0.0002 (0.0001)	0.683 (0.014) [†]	0.0004 (0.0002) [†]	0.813 (0.116) [†]	0.014 (0.013) [†]	0.705 (0.096) [†]	0.007 (0.003) [†]
discount	-0.655 (0.226) [†]	0.0003 (0.0002)	-1.362 (0.041) [†]	0.075 (0.029) [†]	0.114 (0.162)	0.0001 (0.0001)*	0.340 (0.093) [†]	0.0002 (3.5e-5) [†]
feature	0.410 (0.091) [†]		0.942 (0.067) [†]		0.554 (0.112) [†]		0.550 (0.133) [†]	
display	0.644 (0.076) [†]		0.319 (0.033) [†]		0.636 (0.084) [†]		0.501 (0.082) [†]	
feature and display	0.829 (0.080) [†]		1.034 (0.049) [†]		0.833 (0.087) [†]		0.995 (0.117) [†]	
demographic interactions								
price: income	0.036 (0.920)		-0.780 (0.179) [†]		-0.323 (0.419)		-1.222 (0.110) [†]	
income ²	-0.348 (2.329)		0.994 (0.383) [†]		1.368 (1.339)		2.749 (0.419) [†]	
age 21 - 34	0.291 (0.869)		-1.094 (0.077) [†]		-2.932 (0.298) [†]		-0.651 (0.106) [†]	
age 35 - 54	0.330 (0.697)		3.497 (0.197) [†]		0.121 (0.205)		-4.702 (0.145) [†]	
light: age 21 - 34	0.782 (0.355)**		1.034 (0.085) [†]		3.359 (0.388) [†]		5.983 (0.236) [†]	
age 35 - 54	0.989 (0.501)*		-1.337 (0.197) [†]		5.364 (0.557) [†]		5.601 (0.217) [†]	
Corona: %Hispanic	-2.420 (0.316) [†]		-2.405 (0.320) [†]		-2.720 (0.305) [†]		-4.065 (0.263) [†]	
discount: income	0.852 (0.546)*		1.138 (0.223) [†]		0.253 (0.415)		0.168 (0.098)*	
income ²	-3.410 (1.246) [†]		-3.256 (0.689) [†]		-1.094 (1.051)		-0.806 (0.285) [†]	

(Standard deviations in parentheses)

*Significant at the 10% level; **significant at the 5% level; [†]significant at the 1% level.

Table 2. Model statistics

statistic	unconstrained global minima		constrained global minima	
	demand only	joint demand and supply	brand-level	product-level
own-elasticities: # (%) > 0	0 (0)	0 (0)	0 (0)	0 (0)
# (%) ≤ -1	100,384 (76.8)	128,534 (98.4)	130,424 (99.8)	130,622 (99.99)
cross-elasticities: # (%) < 0	0 (0)	0 (0)	0 (0)	0 (0)
aggregate elasticity	-0.908	-0.990	-1.049	-1.135
Lerner bound violations: # (%) < 0	0 (0)	0 (0)	0 (0)	0 (0)
# (%) > 1	42,863 (32.8)	19,815 (15.2)	2,467 (1.9)	121 (0.09)

Table 3. Aggregate merger simulation results for unconstrained and constrained demand models: Bertrand and Stackelberg games

statistics	Bertrand				Stackelberg: AB leads with	
	unconstrained global minima		constrained global minima		Bud and Bud Light	all products
	demand only	joint demand-supply	brand-level	product-level	product-level constraints	
Miller/Coors merger Cincinnati: pre-merger outside share = 57.57%						
aggregate elasticity	-0.89	-0.990	-1.09	-1.14	-1.14	-1.14
# (%) marginal cost $\notin [0,p]$	56 (32.6)	29 (16.9)	15 (8.7)	0 (0)	0 (0)	0 (0)
%pre-merger margin	72.66	72.37	58.10	49.66	49.82	50.06
%CMCR	3.64	4.04	3.62	1.55	1.55	1.55
%Miller/Coors price change	1.29	1.52	1.20	0.46	0.47	0.51
100*leader strategy:						
pre-merger	—	—	—	—	0.13	0.12
post-merger	—	—	—	—	0.28	0.34
Anheuser-Busch/InBev merger Buffalo-Rochester: pre merger outside share = 62.06%						
aggregate elasticity	-1.03	-0.99	-1.11	-1.21	-1.21	-1.21
# (%) marginal cost $\notin [0,p]$	42 (25.8)	10 (6.1)	1 (0.6)	0 (0)	0 (0)	0 (0)
%pre-merger margin	61.17	62.58	50.95	43.27	42.74	45.74
%CMCR	5.30	6.78	5.70	7.45	7.45	7.45
%AB/InBev price change	1.73	1.81	1.56	1.67	1.14	0.18
100*leader strategy:						
pre-merger	—	—	—	—	0.006	0.44
post-merger	—	—	—	—	-1.15	0.76

The Cincinnati merger simulations contained 172 products, the Buffalo-Rochester ones contained 163 products.

Table 4. Buffalo-Rochester brand-level elasticity matrix for model with product-level constraints

price change →	Anheuser-Busch							InBev		Coors			Miller			Grpo - Modlo	Heinkn
share change	Busch	Busch Light	Bud	Bud Light	Mich Light	Mich Ultra	Natural	Labatt	Labatt Light	Light	Keystn	Keystn Light	High Life	Lite	Milwk Best	Corona	Heinkn
Busch	-2.312	0.004	0.194	0.008	0.002	0.008	0.028	0.243	0.007	0.011	0.025	0.004	0.036	0.074	0.052	0.062	0.037
Busch Light	0.005	-2.671	0.005	0.316	0.033	0.080	0.006	0.002	0.304	0.370	<u>0.002</u>	0.118	0.003	0.001	0.014	0.031	0.026
Budweiser	0.114	0.003	-2.388	0.006	0.002	0.006	0.027	0.310	0.005	0.008	0.025	0.003	0.036	0.095	0.048	0.067	0.037
Bud Light	<u>0.003</u>	0.096	0.003	-2.751	0.034	0.077	0.004	0.001	0.360	0.441	<u>0.002</u>	0.116	<u>0.002</u>	0.000	0.010	0.032	0.025
Michelob Light	0.007	0.081	0.006	0.269	-3.092	0.083	0.008	0.003	0.258	0.324	0.003	0.096	0.004	0.001	0.018	0.026	0.021
Michelob Ultra	0.009	0.072	0.009	0.222	0.031	-3.076	0.011	0.004	0.211	0.270	0.004	0.085	0.005	0.001	0.026	<u>0.022</u>	0.019
Natural	0.088	0.014	0.140	0.031	0.009	0.032	-2.337	0.167	0.028	0.042	0.021	0.014	0.030	0.052	0.057	0.046	0.029
Labatt's	0.119	<u>0.001</u>	0.261	<u>0.002</u>	0.001	<u>0.002</u>	0.027	-2.356	<u>0.002</u>	0.003	0.025	<u>0.001</u>	0.037	0.104	0.047	0.070	0.038
Labatt's Light	<u>0.003</u>	0.098	0.003	0.385	0.035	0.078	<u>0.004</u>	<u>0.001</u>	-2.839	0.465	<u>0.002</u>	0.116	<u>0.002</u>	<u>0.000</u>	<u>0.009</u>	0.033	0.025
Coors Light	0.004	0.093	<u>0.004</u>	0.362	0.034	0.077	0.005	0.002	0.356	-2.651	<u>0.002</u>	0.112	<u>0.002</u>	0.000	0.010	0.032	0.024
Keystone	0.105	0.007	0.179	0.016	0.004	0.016	0.028	0.216	0.014	0.021	-2.459	0.007	0.034	0.066	0.054	0.057	0.035
Keystone Light	0.004	0.099	<u>0.004</u>	0.316	0.033	0.078	0.005	0.002	0.299	0.366	<u>0.002</u>	-2.622	<u>0.002</u>	0.001	0.013	0.031	0.026
Miller High Life	0.106	0.005	0.171	0.012	0.003	0.012	0.028	0.210	0.011	0.016	0.023	0.005	-2.272	0.066	0.056	0.056	0.035
Miller Light	0.118	<u>0.001</u>	0.256	<u>0.002</u>	<u>0.001</u>	<u>0.002</u>	0.027	0.336	<u>0.002</u>	<u>0.002</u>	0.025	<u>0.001</u>	0.037	-2.544	0.047	0.069	0.038
Milwaukee's :Best	0.083	0.015	0.123	0.036	0.010	0.037	0.028	0.144	0.033	0.047	0.020	0.016	0.029	0.045	-2.381	0.043	0.028
Corona	0.085	0.027	0.158	0.094	0.009	0.021	0.020	0.200	0.092	0.113	0.018	0.031	0.026	0.060	0.037	-2.695	0.035
Heineken	0.079	0.036	0.132	0.112	0.012	0.028	0.020	0.165	0.108	0.131	0.017	0.043	0.025	0.050	0.038	0.054	-2.731
Outside	0.030	0.030	0.039	0.075	0.014	0.044	0.013	0.041	0.067	0.093	0.008	0.035	0.011	0.013	0.028	0.023	<u>0.016</u>

Outside share = 61.547 percent. Aggregate brands (Craft, Import, Sub-Premium, Premium, and Super-Premium) and brands with share < 0.50 percent are excluded from table.

In each column: largest cross-elasticity is in **bold**; smallest is underlined.

Table 5. Buffalo-Rochester Bertrand merger simulation results for Anheuser-Busch/InBev (Labatt's) merger. Estimated under product-level constraints.

Brewer	brand	%price change	%CMCR	pre-merger share	share change	%pre-merger margin	change in %margin
Anheuser-Busch	Busch	1.74	6.24	1.84	-0.40	47.66	1.867
	Busch Light	1.44	7.40	1.21	-0.24	46.22	1.644
	Bud Select	1.00	6.91	0.16	-0.02	45.70	1.160
	Budweiser	1.36	7.15	2.67	-0.46	44.49	1.686
	Bud Light	1.24	7.62	3.29	-0.59	41.10	1.752
	Michelob	1.02	5.64	0.14	-0.01	51.46	0.952
	Michelob Light	0.73	5.46	0.46	-0.04	40.11	1.097
	Michelob Ultra	0.37	4.23	1.27	-0.01	38.53	0.597
	Natural	0.37	4.78	0.63	0.01	48.19	0.394
	Natural Light	0.48	5.38	0.25	0.00	45.35	0.595
InBev	Labatt's	3.48	9.65	3.09	-1.49	41.76	4.693
	Labatt's Light	3.46	11.71	2.99	-1.62	36.27	5.873
Coors	Coors	-0.12	—	0.26	0.05	37.78	-0.212
	Coors Light	0.23	—	4.03	0.30	41.01	0.317
	Keystone	-0.05	—	0.43	0.05	40.71	-0.074
	Keystone Light	0.56	—	1.44	-0.02	45.26	0.685
Miller	Miller Genuine Draft	-0.17	—	0.17	0.03	41.76	-0.239
	Miller High Life	-0.02	—	0.67	0.07	45.31	-0.022
	Miller Lite	0.01	—	0.98	0.13	40.25	0.000
	Milwaukee's Best	-0.03	—	1.25	0.13	42.85	-0.023
	Milwaukee's Best Light	-0.14	—	0.34	0.04	37.74	-0.212
Grupo-Modelo	Corona	-0.09	—	1.21	0.18	37.01	-0.162
Heineken	Heineken	-0.04	—	0.79	0.11	36.77	-0.082

Outside good: initial share = 62.06; share change = 3.56.

Table 6. Buffalo-Rochester Stackelberg merger simulation results for Anheuser-Busch/InBev (Labatt's) merger. Estimated under product-level constraints: Anheuser-Busch leads with Budweiser and Bud Light.

Brewer	brand	%price change	100*Stackelberg conjecture		pre-merger share	share change	%pre-merger margin	change in %margin
			pre-merger	post-merger				
Anheuser-Busch	Busch	1.53	–	–	1.84	-0.39	47.82	1.65
	Busch Light	0.13	–	–	1.21	0.00	46.37	0.15
	Bud Select	0.45	–	–	0.16	-0.01	45.87	0.52
	Budweiser	1.08	0.00	0.03	2.67	-0.39	44.69	1.32
	Bud Light	0.66	0.01	-2.09	3.29	-0.45	41.29	0.94
	Michelob	1.04	–	–	0.14	-0.02	51.63	0.97
	Michelob Light	-0.30	–	–	0.46	0.05	40.24	-0.45
	Michelob Ultra	-0.16	–	–	1.27	0.09	38.64	-0.23
	Natural	0.02	–	–	0.63	0.03	48.33	0.02
	Natural Light	-0.08	–	–	0.25	0.01	45.49	-0.09
InBev	Labatt's	3.57	–	–	3.09	-1.57	41.92	4.77
	Labatt's Light	2.56	–	–	2.99	-1.44	36.40	4.37
Coors	Coors	-0.42	–	–	0.26	0.06	37.93	-0.69
	Coors Light	-0.81	–	–	4.03	1.27	41.15	-1.17
	Keystone	-0.98	–	–	0.43	0.14	40.84	-1.44
	Keystone Light	0.09	–	–	1.44	0.01	45.41	0.11
Miller	Miller Genuine Draft	-0.35	–	–	0.17	0.03	41.92	-0.48
	Miller High Life	-0.22	–	–	0.67	0.08	45.45	-0.24
	Miller Lite	-0.47	–	–	0.98	0.22	40.40	-0.69
	Milwaukee's Best	-0.32	–	–	1.25	0.16	42.97	-0.42
	Milwaukee's Best Light	-0.17	–	–	0.34	0.02	37.84	-0.26
Grupo-Modelo	Corona	-0.69	–	–	1.21	0.32	37.14	-1.18
Heineken	Heineken	-0.67	–	–	0.79	0.20	36.89	-1.14

Outside good: initial share = 62.06; share change = 1.16.

Table 7. Buffalo-Rochester Stackelberg merger simulation results for Anheuser-Busch/InBev (Labatt's) merger. Estimated under product-level constraints: Anheuser-Busch leads with all products.

Brewer	brand	%price change	100*Stackelberg conjecture		pre-merger share	share change	%pre-merger margin	change in %margin
			pre-merger	post-merger				
Anheuser-Busch	Busch	2.67	0.06	0.05	1.84	-0.44	48.54	2.76
	Busch Light	-4.15	0.61	1.06	1.21	3.08	65.36	-2.29
	Bud Select	2.43	0.01	0.00	0.16	-0.04	46.61	2.72
	Budweiser	2.75	0.10	0.08	2.67	-0.74	45.39	3.22
	Bud Light	-0.10	1.89	1.68	3.29	0.71	60.02	-0.07
	Michelob	2.20	0.00	0.00	0.14	-0.02	52.31	1.95
	Michelob Light	0.82	0.19	0.13	0.46	-0.05	54.82	0.67
	Michelob Ultra	1.28	0.40	0.25	1.27	-0.29	50.40	1.25
	Natural	1.61	0.02	0.02	0.63	-0.08	49.12	1.65
	Natural Light	1.10	0.09	0.06	0.25	-0.04	60.77	0.71
InBev	Labatt's	4.49	–	0.06	3.09	-1.55	41.92	5.94
	Labatt's Light	13.87	–	0.08	2.99	-2.86	36.40	21.29
Coors	Coors	2.37	–	–	0.26	-0.05	37.93	3.80
	Coors Light	1.77	–	–	4.03	-1.03	41.15	2.48
	Keystone	1.50	–	–	0.43	-0.03	40.84	2.13
	Keystone Light	1.33	–	–	1.44	-0.24	45.41	1.59
Miller	Miller Genuine Draft	2.16	–	–	0.17	-0.03	41.92	2.93
	Miller High Life	1.61	–	–	0.67	-0.05	45.45	1.91
	Miller Lite	2.19	–	–	0.98	-0.17	40.40	3.17
	Milwaukee's Best	1.51	–	–	1.25	-0.17	42.97	1.98
	Milwaukee's Best Light	1.61	–	–	0.34	-0.09	37.84	2.62
Groupo-Modelo	Corona	1.93	–	–	1.21	-0.22	37.14	3.20
Heineken	Heineken	1.85	–	–	0.79	-0.14	36.89	3.12

Outside good: initial share = 62.06; share change = 5.47. Aggregate brands (Craft, Import, Sub-Premium, Premium, and Super-Premium) and brands with pre-merger share < 0.10 percent excluded from table.

Figure 1. Price and discount response heterogeneity by income, Corona heterogeneity by percent Hispanic



