Quantitative Marketing and Economics, 1, 293–298, 2003.

Comment

JEAN-PIERRE DUBÉ AND PRADEEP K. CHINTAGUNTA
Graduate School of Business, University of Chicago

The authors develop a Bayesian method to recover the structural parameters from an industry characterized by a heterogeneous logit demand system, on the consumer side, and the Bertrand-Nash equilibrium concept, on the supply-side. The logit demand system is augmented by including error components accounting for unobserved consumer heterogeneity, θ_i, as well as error components accounting for unobserved (to the researcher) product characteristics, ξ_t. By jointly modeling supply and demand, they devise a full-information approach that resolves the potential endogeneity bias in approaches that treat prices as exogenous. Calibrating their model using a household-level scanner panel data set, they draw two empirical conclusions. First, they find that observed prices are correlated with demand shocks, consistent with previous research. Interestingly, this correlation does not appear to lead to biases in the mean price sensitivity across consumers once unobserved heterogeneity is accounted for. This is in variance with previous research that has found such a bias. Second, the results suggest that the additional structure imposed by the supply-side improves the in-sample fit of the model.

Methodologically, the authors make an important contribution to the literature. The Bayesian approach resolves a difficult computational problem that arises in classical likelihood-based approaches. The classical approach assumes the random effects are drawn from a known parametric distribution. Estimation is carried out using the unconditional likelihood, which is integrated over the distribution of the random effects. Integrating over both the random effects as well as unobserved product characteristics ("demand shocks") complicates the evaluation of the unconditional likelihood function as the former vary only across consumers whereas the latter vary over time. The Bayesian approach obviates the need to evaluate this unconditional likelihood thereby making the problem amenable to empirical estimation. This methodological contribution should appeal to researchers modeling the supply-side to control for the endogeneity of marketing variables, and/or to conduct policy experiments using the estimated model parameters. Use of a likelihood-based method also enables researchers to conduct structural tests on the supply-side.

The practical discussion of using full information versus limited information approaches to parameter estimation boils down to trade-offs. A balanced discussion of the relative strengths of alternative approaches to this same empirical problem that rely on fewer modeling assumptions will be a useful addition.

The empirical implementation of the proposed methodology also raises a few issues. Some of these can be resolved through more detail. It is unclear however, whether the data are suitable to showcase the proposed methodology appropriately. Overall, despite the strength of the methodological contribution, it is unclear whether we can make strong conclusions based on the empirical findings.

The remainder of our discussion will focus on two main areas.

1. The trade-offs involved in using full versus limited information approaches as well as maximum likelihood versus instrumental variable (IV) approaches.
2. Empirical implementation.

1. Trade-offs between full-information versus limited information and IV procedures

The authors position existing limited information and IV approaches as a compromise to avoid the technical problems discussed above. In practice however, these IV approaches are typically adopted as they generate consistent demand estimates under far more parsimonious model assumptions.

The full information approach requires taking a stand on the exact model generating the observed prices. In the paper, several static pricing models are compared. This additional structure constrains the estimated demand parameters, as they need to satisfy both the supply and demand equations. In principle, this additional structure leads to more efficient parameter estimates. However, model mis-specification on the supply-side can also lead to biased estimates. In contrast, a limited information approach (Villas-Boas and Winer 1999) remains agnostic about the underlying model generating observed prices. While the demand estimates are estimated with less precision, an inappropriate supply-side model does not contaminate them. Similarly, IV procedures (Chintagunta, Dubé and Goh 2003) do not require assumptions about the precise form of pricing. For policy analysis, one can estimate the supply-side parameters in a second stage that conditions on the demand estimates (Nevo 2000, Berto Villas-Boas 2003). In this way, the demand estimates are not contaminated by an inappropriate supply-side model.

The paper's view of IV approaches is that they are equivalent to imposing a supply-side model in which prices are a linear function of instruments. This point is well taken. However, to the extent that one can more readily identify exogenous factors that shift prices than one can describe the true nature of the price setting behavior of manufacturers and retailers, our ability to obtain unbiased estimates is not compromised. Further, one can use a flexible functional form instead of a linear regression. Indeed, developing a structural model from which one derives a pricing policy predicting occasional strategically-timed price cuts—a pattern observed in most retail scanner data—is extremely difficult. Existing theories that predict similar pricing involve models with a time-varying price elasticity of demand, on the consumer side, and a dynamic pricing game, on the supply-side. Solving the dynamic programming problem of retailers is beyond the scope of the current paper. However, it is not clear whether the static models considered provide a satisfactory explanation of the observed price variation over time in the data. The key question that one needs to answer is—are we really worse off with less efficient estimates that do not impose inappropriate assumptions on the process generating prices?

2. Empirical implementation

In the empirical section, the authors use scanner data to calibrate the model and demonstrate the proposed methodology. Several surprising results are documented. First, the authors find evidence of correlation between the unobserved product attributes, ξ_{jt}, and prices; however they do not find any evidence of endogeneity bias. Second, they find support for a pricing model that is based on an information-set different from that assumed in the literature. Third, the paper finds substantial improvement in model fit in moving from the limited to the full information approach when heterogeneity is accounted for. In the following, we first discuss the data used in the paper and then the above findings.

A concern that arises is the unusually small data set used to capture features of both supply and demand in this industry. For instance, with 92 weeks, 631 purchase occasions and 3 brands, there are only 2.3 observed choices for a given brand in a given week (on average). The authors estimate weekly brand-specific parameters, ξ_{jt}, to control for unobserved product characteristics. The estimation of these parameters relies either on the information in the data (the 2.3 observed choices for a given brand in a given week) or the information contained in the prior. Given how few choices we have per brand in a typical week, it would help to have a discussion on how well we can inform ourselves of the distribution of ξ_{jt} with the data at hand. Since classical estimation of these parameters would be hopeless with so few observations, the Bayesian approach is indeed better suited to handle this type of small sample problem. But, it seems crucial at this point to discuss the sensitivity of the prior chosen. It would have been helpful to report whether the posterior distribution of ξ_{jt} deviated far from the prior. Similarly, it would have been helpful to see how sensitive the posterior distribution would be to a different choice of prior. Given how little information the data contain about specific realizations of ξ_{jt}, it is unclear how much we can learn about the aggregate correlation between prices and ξ_{jt}. Only in weeks with unusually low prices would one expect to observe enough choices of a brand to learn about that brand's unobserved attribute. In fact, referring to the scatter plots at the end of the paper, we see that the documented correlation between prices and ξ_{jt} is most pronounced at the lowest price levels.

A surprising empirical finding is the stated lack of endogeneity bias. This result obtains despite the positive correlation between prices and ξ_{jt} (between 0.164 and 0.39), and the preponderance of evidence to the contrary in related work. Typically, research on this topic has focused on biases in the mean price sensitivity. The authors find that controlling for the supply side does not impact the estimated mean price sensitivity. Does this mean the endogeneity bias is unimportant? In the case of a linear regression, it is straightforward to characterize the endogeneity bias. Characterizing the bias is not straightforward in the context of the non-linear model used. First, it is unclear how strong the correlation between prices and ξ_{jt} must be to generate statistical bias. Perhaps another useful statistic to report would be the marginal impact of the estimated ξ_{jt} on demand. It is also unclear how the endogeneity bias will manifest itself in the estimates. For instance, Chintagunta et al. (2003) also document that controlling for price endogeneity leads to different results for heterogeneity. The authors report two moments of the parameter distribution; but they do not compare the entire posterior distribution of tastes. It might have been helpful to look at a histogram of the posterior distribution of parameters to see if it is comparable. Perhaps a simulation exercise would have been more instructive than the calibration to scanner data to try and characterize the estimation bias. Also, the authors could have compared the marginal effect of prices on choices under the 10 models (e.g. own price elasticity) as a more informative metric to assess the effect of endogeneity.

Another unusual finding in the paper is the support for a model in which firms are assumed to aggregate demand over the posterior distribution of heterogeneity based on the sample. Typically, researchers integrate over an assumed population distribution of heterogeneity (e.g. the normal distribution using the hyperparameters for the mean and variances).[1] The authors' finding provides the basis for an interesting discussion about the information sets available to firms. While it may be too strong to assume the population distribution of tastes is known, at the same time, do we truly believe that retailers and manufacturers set prices based on this exact panel data set? In general, the finding opens a potentially interesting discussion about the information used by manufacturers and retailers when they set prices. A more thorough investigation of the appropriate information sets facing firms seems like an interesting area for future research.

The authors also observe that the limited and full information approaches result in different price effects in the no heterogeneity case whereas the two models result in similar estimates when heterogeneity is accounted for. They conclude, "mis-specification of heterogeneity can mask the importance of the supply-side model." However, after accounting for heterogeneity, the estimates from the full information model become close to those from the "no endogeneity" specification, whose results are similar to the limited information case under both heterogeneity and no heterogeneity specifications. A plausible re-interpretation of the results would be that accounting for heterogeneity mitigates the bias resulting from imposing an incorrect full information supply-side model. For instance, the IIA assumption in the homogeneous logit demand system restricts a brand's price-cost margin on the supply-side to be proportional to its market share. However, in the data (Table 1),

1 The results from this second model are not reported. However, the authors state that it provided an inferior fit to the data.

Coors Light has the largest share on average; but not the highest price.[2] The inclusion of heterogeneity offsets the IIA property and its consequent effect on the pricing model.

A related finding is the superior fit of the same full-information model versus a limited information approach (model M9 versus M10), although the fit of the limited information approach is superior when consumer heterogeneity is not modeled. The intuition for this improved fit is unclear. Looking at the results in Table 2, we see that accounting for heterogeneity in the limited information model improves the log-marginal density from -3400 to -2464. Since accounting for heterogeneity does not influence the supply-side in this model, all the improvement in fit can be attributed to the demand model. In the full information case, where accounting for heterogeneity does influence the pricing model, adding heterogeneity improves the log-marginal density from -3410 to -2377 which is more than the improvement in the limited information case. This suggests that heterogeneity is also helping the fit of the pricing model in the full information case. As noted previously, it does appear that accounting for heterogeneity resolves the potential mis-specification of the supply-side model in the full information case.

Finally, lagged price is used as an instrument for price in the limited information specification of the model. As noted previously, the no endogeneity finding, comparing the results from the limited information model with those from the specification that ignores endogeneity, is conditional on the chosen instruments for price. One concern is whether these instruments are appropriate. A discussion of this issue, in the light of other types of instruments chosen by researchers (factor prices, wholesale prices, prices in other markets, etc.), is warranted. It would also be helpful to have more information about the quality of the instruments, such as an R^2. To assess the validity of the instruments, the authors could check the correlation between the estimated ξ_{jt} from the full information model (M9) with the chosen instruments, i.e., lagged prices.

In summary, we reiterate that we believe that this paper makes an important methodological contribution to the burgeoning area of understanding market demand and supply behavior. We look forward to further research in this area that can shed some light on the issues raised in this discussion.

References

Berto Villas-Boas, S. (2003). "Vertical Contracts Between Manufacturers and Retailers: An Empirical Analysis", Discussion paper, University of California—Berkeley.

Chintagunta, P.K., J.-P. Dubé, and K.-Y. Goh. (2003). "Beyond the Endogeneity Bias: The Effect of Unmeasured Brand Characteristics on Household-level Brand Choice Models", working paper, University of Chicago.

2 We compare prices as the estimates marginal costs (brand intercepts) are all statistically the same (M3, M4, M8 and M9).

Nevo, A. (2000). "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry", The RAND Journal of Economics 31(3), 395–421.
Villas-Boas, J.M. and R. Winer. (1999). "Endogeneity in Brand Choice Models", Management Science 45, 10, 1324–1338.