ECONOMIC ANALYSIS GROUP
DISCUSSION PAPER

Calibrating the AIDS and Multinomial Logit Models with Observed Product Margins

By
Gloria Sheu and Charles Taragin
EAG 12-7 October 2012

---

* The views expressed in this paper are those of the authors and are not purported to reflect the views of the U.S. Department of Justice.

Calibrating the AIDS and Multinomial Logit Models with Observed Product Margins

Gloria Sheu* Charles Taragin*

August 2012

Abstract

We show how observed product margins may be used in lieu of an observed market elasticity to calibrate parameters for two commonly used demand forms: the Almost Ideal Demand System (AIDS) and the multinomial logit. This technique is useful for antitrust practitioners interested in simulating the effects of a merger, since estimates of product margins are often easier to obtain than estimates of market elasticities.

Keywords: demand calibration, multinomial logit, almost ideal demand system, AIDS
JEL classification: L40, K21

*The views expressed herein are entirely those of the authors and should not be purported to reflect those of the U.S. Department of Justice. Address: Economic Analysis Group, Antitrust Division, U.S. Department of Justice, 450 5th St. NW, Washington DC 20530. E-mail: gloria.sheu@usdoj.gov and charles.taragin@usdoj.gov.

1 Introduction

Counterfactual exercises in industrial organization typically begin by parameterizing a demand system. In order to evaluate the effect of a change in policy or of a shift in the competitive environment, one needs to know how demand will respond and how consumers will be affected. An important example of this process arises in the context of antitrust, with the use of merger simulation. Merger simulation is a tool used to predict the competitive effects of a proposed acquisition. A counterfactual equilibrium is simulated by imposing joint ownership on the products sold by the prospective merging parties and solving for the prices that would result. This technique is commonly used by economists at both of the United States antitrust authorities, the Antitrust Division of the Department of Justice and the Federal Trade Commission. Merger simulation has also been featured in recent antitrust litigation (e.g. United States v. H&R Block, Inc., *et al.* (2011)).[1]

In order to use merger simulation, the antitrust practitioner chooses a functional form for demand and values for its parameters. One method for selecting these parameters is "calibration," where parameters are fitted so as to rationalize certain pieces of observed market data. Calibration is similar to econometric estimation except that i) calibration typically uses significantly less data than estimation and ii) calibration usually assumes that the error terms are exogenous (i.e. the error term is not correlated with observed product characteristics, such as price).[2]

An important question when using merger simulation is how one should allow for demand substitution to products besides those in the market at issue (often termed the "outside good"). If such substitution is likely, it will tend to mitigate the possibility of a post-merger price increase. It is common for antitrust practitioners to assume a baseline outside good sensitivity (as governed by an "aggregate" or "market" elasticity) in their calibrations and to then vary it to check robustness of results.[3] In other cases additional data on consumer behavior are used to validate a certain choice for the level of outside good substitution.[4]

In this paper we show how two commonly used demand forms, the Almost Ideal Demand System (AIDS) and the multinomial logit, in conjunction with the assumption of

[1]See for example the discussion of merger simulation on pages 74 to 78 of the United States v. H&R Block, Inc., *et al.* Memorandum Opinion. Available at http://www.justice.gov/atr/cases/handrblock.html.

[2]As a result, it is often not possible to engage in neo-classical hypothesis testing with calibrated demand parameters.

[3]Epstein and Rubinfeld (2002), for example, suggest setting the aggregate elasticity in their model to -1 as a starting point.

[4]These data could be in the form of consumer survey evidence, for instance.

Bertrand competition, allow one to recover the aggregate elasticity from the calibration routine. We demonstrate how product margin information may be used in lieu of setting the market elasticity a priori. This substitution is important because in many antitrust investigations, estimates of product margins are easier to obtain than estimates of market elasticities.

This paper proceeds as follows. In the next section we briefly introduce the Bertrand framework that is the starting point for most merger simulations. Section 3 explains how calibration can proceed with the AIDS model. Section 4 applies the same principals to the multinomial logit. Section 5 concludes.

2 Bertrand Price Competition

The Bertrand model of differentiated price competition forms the basis for most merger simulation routines. Suppose that there are K firms in a market, each selling n_k products for a total of $n = \sum\limits_{k \in K} n_k$ products. Each product is produced using its own distinct constant marginal cost technology, c_i for all $i \in n$.

Firm $k \in K$ chooses the prices $\{p_i\}_{i=1}^{n_k}$ of its products so as to maximize profits. Mathematically, firm k solves

$$\max_{\{p_i\}_{i=1}^{n_k}} \sum_{i=1}^{n_k} (p_i - c_i) q_i,$$

where q_i, the quantity sold of product i, is assumed to be a twice differentiable function of all product prices. Differentiating profits with respect to each p_i yields the following first order conditions (FOCs):

$$q_i + \sum_{j=1}^{n_k} (p_j - c_j) \frac{\partial q_j}{\partial p_i} = 0 \quad \text{for all } i \in n_k,$$

which may be rewritten as

$$r_i + \sum_{j=1}^{n_k} r_i m_j \epsilon_{ji} = 0 \quad \text{for all } i \in n_k, \tag{1}$$

where $r_i \equiv (p_i q_i)/(\sum_{j=1}^{n} p_j q_j)$ is product i's revenue share, $m_i \equiv (p_i - c_i)/p_i$ is product i's gross margin, and $\epsilon_{ij} \equiv (\partial q_i/\partial p_j)(p_j/q_i)$ is the elasticity of product i with respect to the price of product j.

Antitrust practitioners interested in simulating the effects of a merger typically assume that consumer demand is characterized by a particular function, which implies that the price elasticities ϵ_{ij} have a certain functional form. Under this demand assumption, the above system of FOCs is used to first calibrate demand parameters for the chosen demand function, and then to solve for pre- and post-merger equilibrium prices, conditional on the calibrated demand parameters.[5] In the Bertrand model, a merger is modeled as placing the merging parties' products under the control of a single entity, enabling that entity to set prices across *all* of these products so as to maximize profits.

3 The AIDS Model

In Epstein and Rubinfeld (2002), the authors show how the AIDS model may be calibrated for use in merger simulation. They use an AIDS specification based on the original model of Deaton and Muellbauer (1980), but modify it to remove income effects and to have a linear price index. Epstein and Rubinfeld demonstrate that if

1. firms are assumed to play a Bertrand differentiated products pricing game,
2. the *revenue* diversion d_{ij} between any two products, defined as the percentage of product i's revenue that moves from or to product j due to a price change in i, is known,[6]
3. all revenue shares r_i (as calculated amongst the goods inside the market) for all products i are known, and
4. a single product's own-price elasticity ϵ_{ii} for some i as well as the market (aggregate) elasticity ϵ are known,

then the slopes of their modified AIDS demand system can be calibrated and used to simulate equilibrium price changes resulting from a merger. If, additionally, the prices p_i for all products i are known, then the demand intercepts can be recovered as well.[7] In what

[5]Note that the system of FOCs are *necessary* conditions for a Nash-Bertrand equilibrium in prices. Throughout, we assume that pre-merger, a unique price equilibrium exists. We are not aware of any theoretical result demonstrating that a unique equilibrium exists in prices for these demand systems when i) firms produce multiple products and ii) marginal costs are constant.

[6]Specifically, Epstein and Rubinfeld assume that diversion is proportional to revenue share, which is why they refer to their model as the "proportionally calibrated Almost Ideal Demand System" (PCAIDS). This assumption is not strictly necessary in order to identify the demand parameters.

[7]Epstein and Rubinfeld (2002) shows that AIDS intercepts are not needed to predict price *changes* from a merger. AIDS intercepts, however, are necessary to predict pre- and post-merger price *levels*, as well as welfare measures like compensating variation.

follows, we show that knowledge of profit margins can substitute for knowledge of ϵ_{ii} and ϵ when calibrating AIDS parameters.

The AIDS model (without income effects) assumes that the demand for each product $i \in n$ in the market is governed by

$$r_i = \alpha_i + \sum_{j \in n} \beta_{ij} \log(p_j) \quad \text{for all } i \in n, \text{ where } \beta_{ii} < 0 \text{ and } \beta_{ij} > 0.$$

These equations may be written in matrix notation as

$$r = \alpha + B \log(p), \tag{2}$$

where r and p are vectors of product revenue shares and prices, respectively, α is a vector of product-specific demand intercepts, and B is a matrix of slopes. Deaton and Muellbauer (1980) demonstrate that in order for the AIDS model to be consistent with consumer choice theory, B must be symmetric. Epstein and Rubinfeld (2002) further assume that the industry price index P takes the form of Stone's index,

$$\log(P) = \sum_{i \in n} r_i \log(p_i),$$

which is the revenue share-weighted geometric average of prices.[8]

This model yields the following own- and cross-price elasticities:

$$\begin{aligned} \epsilon_{ii} =& -1 + \frac{\beta_{ii}}{r_i} + r_i(1+\epsilon) \text{ and} \\ \epsilon_{ij} =& \frac{\beta_{ij}}{r_i} + r_j(1+\epsilon). \end{aligned} \tag{3}$$

It is typically assumed that $\epsilon \leq -1$ and $|\epsilon| \leq |\epsilon_{ii}|$ for all $i \in n$.

3.1 Calibrating Demand Parameters

Suppose that revenue diversion $d_{ij} \equiv -\frac{\partial r_j}{\partial p_i} / \frac{\partial r_i}{\partial p_i}$ is observed for all products $i, j \in n$. The form of demand in equation (2) then implies that $d_{ij} = -\beta_{ji}/\beta_{ii}$. When combined with symmetry

[8] The version of the AIDS model using Stone's index is often called the "linear approximate AIDS" (LA-AIDS). This form allows the aggregate elasticity ϵ to enter linearly (with a revenue share weight r_i) in the equations for the own and cross-price elasticities of each product. See equation (3).

of B, this yields the following for all products $i, j \in n$:

$$\begin{aligned} \beta_{jj} &= \frac{d_{ij}}{d_{ji}} \beta_{ii} \text{ and} \\ \beta_{ij} &= -d_{ij} \beta_{ii}. \end{aligned} \tag{4}$$

The above equations imply that when the full matrix of diversion ratios is known, knowledge of a single element on the diagonal of B (without loss of generality, β_{ii}) is sufficient to determine all of B. Epstein and Rubinfeld (2002) solve for β_{ii} by assuming that the industry elasticity ϵ and the own-price elasticity ϵ_{ii} are known, and then use the definition of ϵ_{ii} to solve for β_{ii}.

However, data on margins can substitute for knowledge of the elasticities ϵ and ϵ_{ii}. Margins are an important profitability metric and as such are often computed by firms in the normal course of business. Price elasticities, in contrast, are less frequently measured, unless a firm has taken it upon itself to engage in detailed demand estimation or consumer survey work. Estimates of aggregate elasticities are even more rare, as individual firms do not directly need to know the behavior of consumers leaving the market in order to maximize profits. From a single firm's perspective, such industry-wide behavior only matters in so far as it is reflected in sales of its own goods.

Recall that the FOC for one good, as shown in equation (1), provides a relationship between revenue shares, margins, and the own- and cross-price elasticities. These price elasticities, according to equations (3) and (4), in turn can be expressed as a function of revenue shares, the industry elasticity, and a single price coefficient β_{ii}. As a result, each FOC can be written in terms of revenue shares, margins, the coefficient β_{ii}, and ϵ. Therefore, if all the margins appearing in two FOCs are known (that is, if the margins of all the goods sold by one multi-product firm or of two goods sold by separate single-product firms are known), a system of two linear equations for the unknowns of β_{ii} and ϵ is obtained.[9] In the single-product case these FOCs take the form

$$\begin{aligned} -1 + \frac{\beta_{ii}}{r_i} + r_i(1+\epsilon) &= -\frac{1}{m_i} \text{ and} \\ -1 + \frac{d_{ik}\beta_{ii}}{d_{ki} r_k} + r_k(1+\epsilon) &= -\frac{1}{m_k}. \end{aligned} \tag{5}$$

[9]To be precise, we have a system of two linear equations and two inequality constraints, $\epsilon \leq -1$ and $|\epsilon| \leq |\epsilon_{ii}|$. If a particular set of margins and revenue shares causes these inequalities to be violated, they can be incorporated as constraints in the choice of β_{ii} and ϵ. For example, one could use a minimum distance routine that sets the FOCs as close as possible to zero without violating these constraints.

Intuitively, the price sensitivity for one good can be recovered from its margin using the usual Lerner relationship. With only a single margin, however, one cannot differentiate between price-driven substitution to other goods versus to outside the market (as governed by the interplay between ϵ_{ii} and ϵ). In order to make that distinction, a second margin and its attendant Lerner condition is required.[10]

Hence, using margins, one can solve these equations for β_{ii} and ϵ. Knowledge of β_{ii} then allows one to recover the other entries in B using equation (4). The intercepts in α can be derived by plugging prices and revenue shares into the demand equations in (2).

4 The Multinomial Logit Model

Werden and Froeb (1994) highlight how the multinomial logit demand model may be used in merger simulations. Werden and Froeb (1994) demonstrate that if

1. firms are, as in the the AIDS simulation framework, assumed to play a Bertrand differentiated products pricing game,
2. all *quantity* shares $s_{i|I}$ (as calculated amongst the goods inside the market) for all products i are known,
3. all prices p_i for all products i are known,
4. the market elasticity ϵ is known, and
5. the margin m_i for some product i is known,

then the parameters of the multinomial logit can be calibrated.[11] Below, we show that knowledge of additional margins can substitute for knowledge of ϵ.

[10] In the multi-product case, all the margins entering into two full FOCs are needed, giving the system of equations,

$$r_i + r_i m_i \left(-1 + \frac{\beta_{ii}}{r_i} + r_i(1+\epsilon) \right) + \sum_{j=1, j\neq i}^{n_k} r_j m_j \left(\frac{-d_{ij}\beta_{ii}}{r_j} + r_i(1+\epsilon) \right) = 0 \text{ and}$$

$$r_k + r_k m_k \left(-1 + \frac{d_{ik}\beta_{ii}}{d_{ki} r_k} + r_k(1+\epsilon) \right) + \sum_{j=1, j\neq k}^{n_k} r_j m_j \left(\frac{-d_{kj} d_{ik}\beta_{ii}}{d_{ki} r_j} + r_k(1+\epsilon) \right) = 0.$$

[11] In the empirical application discussed by Werden and Froeb (1994), they use a pre-existing estimate of the price coefficient instead of a margin, thus avoiding calibration for that parameter. However, conditional on all the other pieces of data enumerated above being available, there is a one to one mapping between a margin and the price coefficient. See equation (14) in Werden and Froeb (1994).

The multinomial logit model assumes that each consumer, here indexed by k, has the following indirect utility function for each product $i \in n$ and for the outside good:

$$u_{ik} = \delta_i - \gamma p_i + e_{ik}.$$

Let $V_i \equiv \delta_i - \gamma p_i$. In order to identify the δ_i parameters, one of them must be normalized. Typically this is achieved by setting $V_0 = 0$ for the outside good.

If the e_{ik} are distributed Type I extreme value, the probability (market share) that a consumer purchases product i can be written as

$$s_i = \frac{\exp(V_i)}{1 + \sum\limits_{j \in I} \exp(V_j)}. \tag{6}$$

Similarly, the probability (conditional market share) that a consumer purchases good i conditional on choosing an inside good is

$$s_{i|I} = \frac{\exp(V_i)}{\sum\limits_{j \in I} \exp(V_j)}.$$

Then $s_{i|I}$ is related to s_i according to

$$s_i = s_{i|I} s_I, \tag{7}$$

where s_I is the probability that a consumer chooses an inside good. If we denote by s_0 the probability that a customer chooses the outside good, then $s_I = 1 - s_0$.

The own- and cross-price elasticities may be expressed as

$$\begin{aligned} \epsilon_{ii} =& -\gamma(1 - s_{i|I}(1 - s_0))p_i \\ \epsilon_{ij} =& \gamma s_{j|I}(1 - s_0)p_j. \end{aligned} \tag{8}$$

Furthermore, the aggregate elasticity of demand is

$$\epsilon = -\gamma s_0 \bar{p}, \tag{9}$$

where $\bar{p} = \sum_{i \in n} s_{i|I} p_i$.

4.1 Calibrating Demand Parameters

Assume that all conditional market shares $s_{i|I}$ and prices p_i are known. Based on the form of the multinomial logit laid out in the previous section, the task for calibration is to find the unknown values of γ, δ_i for all $i \in n$, and s_0. As in the AIDS model, calibration of the multinomial logit relies on using the FOCs that result from Bertrand price competition. In the case of single-product firms, the FOC for a product i takes the form

$$\gamma(1 - s_{i|I}(1 - s_0))p_i = \frac{1}{m_i}. \tag{10}$$

This is the usual Lerner relationship. Following Werden and Froeb (1994), if the margin m_i and the market elasticity ϵ are known, then an estimate of γ may be formed by solving equation (9) for s_0, substituting this expression into equation (10), and then solving for γ. Once an estimate of γ has been recovered, the value of s_0 follows immediately from (9).

However, knowledge of additional margins can substitute for knowledge of ϵ. Again using single-product firms as an example, suppose that two margins, m_i and m_j are observed. Then taking equation (10) for both goods i and j, we have a system of two equations in two unknowns that can be solved for s_0 and γ.[12] This result extends to multi-product firms, so long as all of the margins that appear in two FOCs are known.[13]

All that remains is to calibrate the vector of δ_i parameters. Once s_0 is recovered, the unconditional market shares can be constructed using equation (7). Furthermore, the underlying probability form of the unconditional market shares in (6) implies that

$$\log(s_i) - \log(s_0) = \delta_i - \gamma p_i, \tag{11}$$

which allows one to back out the δ vector.[14] Therefore, we find that the multinomial logit

[12] Specifically, these give that $1 - s_0 = (m_i p_i - m_j p_j)/(m_i p_i s_{i|I} - m_j p_j s_{j|I})$. Then γ can be solved for by plugging back into either FOC.

[13] In this case the system of two equations is given by

$$r_i + \gamma r_i m_i (s_{i|I}(1 - s_0) - 1)p_i + \gamma \sum_{l=1, l\neq i}^{n_k} r_l m_l s_{i|I}(1 - s_0)p_i = 0 \text{ and}$$
$$r_j + \gamma r_j m_j (s_{j|I}(1 - s_0) - 1)p_j + \gamma \sum_{l=1, l\neq j}^{n_k} r_l m_l s_{j|I}(1 - s_0)p_j = 0,$$

where it is assumed that all of the margins for the products made by firm k are known.

[14] Recall that $V_0 = 0$ for the outside good. This means that $s_0 = 1/(1 + \sum_{j\in I} \exp(V_j))$ by equation (6). Taking the log of s_i/s_0 then yields this result.

model can be calibrated without knowing the aggregate elasticity ϵ.

5 Conclusion

In this paper, we have demonstrated how margin information can be substituted for the market elasticity when calibrating the AIDS and multinomial logit parameters. We believe this substitution is useful, as in our experience antitrust practitioners are more likely able to obtain reasonable margin estimates than estimates of the market elasticity.

In the above discussion, we focused on the case when the market elasticity parameter was just-identified. We argued that at a minimum, at least two margins from two single-product firms or all the margins from a multi-product firm are necessary to calibrate demand parameters. If, however, additional margin information is available, then there are more equations than unknowns and the demand parameters are over-identified. In this case, we recommend using a minimum distance procedure that finds the parameters that set the FOCs as close to zero as possible, subject to any theoretical restrictions on the signs of the demand parameters.

References

Deaton, A. and J. Muellbauer (1980). An Almost Ideal Demand System. *The American Economic Review 70*(3), 312–326.

Epstein, R. J. and D. L. Rubinfeld (2002). Merger Simulation: A Simplified Approach with New Applications. *The Antitrust Law Journal 69*(3), 883–919.

Werden, G. J. and L. M. Froeb (1994). The Effects of Mergers of Differentiated Products Industries: Logit Demand and Merger Policy. *The Journal of Law, Economics, and Organization 10*(2), 407–426.