Institute for New Economic Thinking

Working Papers
ineteconomics.org/workingpapers

The Perils of Antitrust Econometrics: Unrealistic Engel Curves, Inadequate Data, and Aggregation Bias

Gabriel A. Lozada*

Working Paper No. 203

March 29th, 2023

ABSTRACT

Some economists argue antitrust policy should be based on empirical methods used by the Industrial Organization subdiscipline of economics, but non-economists must understand that those methods contain certain highly restrictive assumptions. Those assumptions involve econometric "identification," and treating aggregate demand as if it were generated by a representative consumer (Muellbauer's "generalized linear" preferences). We derive new results illustrating how restrictive the representative consumer assumption is; we explain aggregation bias in Almost Ideal Demand System models; and we show that data limitations make it even harder to justify economists' restricting aggregate demands as one would the demand of a single individual.

https://doi.org/10.36687/inetwp203

JEL Codes: D12, C43, L4, C54

Keywords: Antitrust econometrics; Almost Ideal Demand System (AIDS); New Brandeis School

* Department of Economics, University of Utah, Salt Lake City, UT 84112, USA; lozada@economics.utah.edu; 801-581-7650; www.economics.utah.edu/lozada.

This work has been partially supported by the Institute for New Economic Thinking. I would like to thank Mark Glick for suggesting this topic, and thank him and Thomas Ferguson for commenting on earlier drafts.

The cornerstone of the econometric approach to predicting merger effects is the econometrics of consumer demand, given the emphasis (one might say overemphasis) over the last few decades on the effect of mergers on consumer welfare. The econometric approaches used are freighted with assumptions which are not widely and deeply understood, even by many economists, and which are unlikely to hold in practice. Explaining those assumptions is the purpose of this paper.

The prior paper most resembling this one is Baker and Rubinfeld's (1999) survey article on empirical methods in antitrust litigation. Some of the concerns of that paper will be revisited and explained in Section 1. However, Baker and Rubinfeld make no mention of the strong assumptions used in the empirical investigations. Some of those assumptions are simple to understand, have been discussed for decades, and seem to be accepted unquestioningly by antitrust econometricians. These are that consumers are price takers, and that consumers behave as if they were fully rational, have access to full information, and have preferences which never change. None of the last three assumptions are particularly compelling, as much work in the field of behavioral economics, for example the Nobel Prize-winning work of Kahneman and Tversky, has shown. But since those assumptions are so simple and widely-discussed, the purpose of this paper is to analyze another key assumption which is much less widely talked about: the assumption that restrictions from the microeconomic theory of a *single* consumer are applicable to the estimation of the demand behavior of *many* consumers. This assumption allows the "aggregation" of individual consumer demands into a total market demand curve that has the same properties as individual consumer demands.

Section 2 gives an overview of our results, including, in Figure 2, a flowchart capturing the results of Sections 3–5. It turns out that whether restrictions from the microeconomic theory of a single consumer are applicable to the estimation of the demand behavior of many consumers depends, in part, on whether each consumer's "Engel curves" can be assumed to be linear. "Engel curves" are the relationship between a household's income and its consumption of a good. When Engel curves are depicted in graphs, income is on the horizontal axis and demand (or "consumption") is on the vertical axis. On such a graph, when the Engel curve is upward-sloping, the good is "normal": when income rises, its consumption rises. On such a graph, when the Engel curve is downward-sloping, the good is "inferior": when income rises, its consumption falls. If income changes have no effect on a consumer's demand for a product, the Engel curve for the product will be a flat line. Linear Engel curves are, as we explain below, widely believed to be unrealistic. However, most economists think that in the case of Section 4 ("exact nonlinear aggregation"), the assumption of linear Engel curves is unnecessary. We show that contrary to this conventional

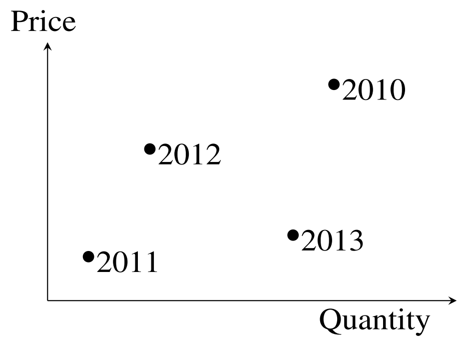

Figure 1. A plot of price and quantity sold for a hypothetical commodity in the years 2010–2013.

wisdom, even in Section 4, the assumption of linear Engel curves is often required. Once it is understood that econometric analyses of consumer demand are built on a framework of assumptions which may be invalid, it becomes possible to raise questions about how accurate those analyses really are. Yet another commonly-used assumption is that income but not wealth determines demand. That makes little sense, but adopting the more realistic assumption would require economists to have data on household wealth, which is often unavailable.

The last section of this paper, Section 6, discusses problems with using consumer surplus as a measure of consumer welfare. Those problems turn out to involve many of the same issues as those involved in estimating consumer demand curves.

1. The Subjectivity of Econometrics

Computation of demand curves is probably the most common empirical analysis in antitrust litigation. For example, it is needed in order to estimate how the consumption of one commodity will change when the price of another commodity increases. However, computation of demand (and supply) curves cannot be done in an a-theoretical, purely empirical way, because those curves are not obvious from market data. Typical observations of market data for price and quantity look like Figure 1, which presents itself as a meaningless scatter of points. This is why the idea of demand curves was unknown to such important and insightful early economists as Adam Smith and David Ricardo, with the first use of demand curves being by Antoine-Augustin Cournot in 1838 (see Humphrey (1992 p. 3)).

Economists make sense of this situation by imposing *non*-empirically-grounded *assumptions* from economic *theory*. However, these assumptions may be inappropriate, or simply wrong, in a particular application. Non-economists in antitrust should inquire about these assumptions and how justifiable they are.

As an example of how these non-empirical, theoretical assumptions work, an economist looking at Figure 1 might require that the data points are intersections between the demand and supply curves of price-taking agents; that demand curves are downward-sloping; and that supply curves are upward-sloping. Those assumptions would enable the economist to rule out one demand curve passing through the data points 2011 and 2012, or 2011 and 2010, or 2011 and 2013, or 2012 and 2010, or 2013 and 2010. If some important non-price determinant of demand, such as income, was similar in 2012 and in 2013, but different in 2011 and in 2010, support increases for one demand curve passing through the 2012 and 2013 points, but a different one for 2011 and for 2010. Assuming that supply curves are upward-sloping narrows down possible locations of supply curves, but still leaves several possibilities: 2011, 2012, and 2010 could all share the same supply curve, or 2011, 2013, and 2010 could all share the same supply curve, or only 2011 and 2010 could share a supply curve and 2012 and 2013 each have their own supply curves; or, each year could have its own supply curve. If this is the market for food and if good or bad weather is known to affect supply, then knowing which years had similar weather would help decide which of these possibilities for supply curves are more likely.

This explains the need for econometricians to impose non-empirically-grounded, theoretical assumptions in order to solve what is called the "identification" problem, which refers to this impossibility of using data *alone* to locate the curves which neoclassical economists believe generated the data.

Baker and Rubinfeld (1999 p. 408) discussed the identification problem in the context of antitrust, and Ulrick (2014 pp. 144–5) considers identification to remain a very serious difficulty in antitrust econometrics. Ulrick sharply criticizes the way Hausman, Leonard, and Zona (1994) solved their identification problem, and Ulrick writes, "Perhaps the preceding discussion makes demand estimation seem a hopeless cause. Indeed, it is highly doubtful that the conditions that make possible consistent estimation of demand are met." The point is that the seemingly most basic question "Is it a demand curve, or is it a supply curve?" (the title of a paper Ulrick cites in his footnote 31) cannot be answered just by collecting data. Measurement in economics is not like measurement in natural science, where tools such as scales and rulers enable objective measurements to be made in a straightforward fashion. Measurement in economics is much more difficult, and making even such basic measurements as demand and supply can be controversial and fraught with difficulties.

Econometricians also need to tell their statistical algorithms what sort of general shape to assume for demand curves. For example, should they be assumed to be linear? Parabolic? Or perhaps they should be assumed to be described by a second-degree polynomial, which is "flexible" enough to include the linear and parabolic shapes as special cases. There are many more possible assumptions, from which econometricians are forced to choose.

In the same way, the statistical algorithms demand the econometrician assume a particular shape for the supply curves. Typically, supply curves are assumed to be upward sloping. This is fine for competitive firms, whose supply curves are their marginal cost curves (above the bottom of average variable cost). Monopolists, however, have no supply curves (they have a supply point), and their marginal cost curves are likely downward-sloping—after all, economies of scale are a common explanation for why a firm was able to become a monopolist in the first place. Oligopolists have no simple supply curve either. Therefore, it is not obvious what the best *a priori* assumption is to impose on the supply side of a possibly non-competitive market.

The upshot is that econometric analysis of markets requires non-empirically-grounded assumptions to be made. In the next several Sections we look at the most commonly-assumed theoretical structure for consumer demand.

2. Econometric Difficulties in Estimating Consumer Demand

For the rest of this paper we will be concerned with the market's demand side. There, the problem of identification has several facets; Muellbauer (1975 p. 525) describes one:

> Suppose that n groups of commodities have been defined. The minimum number of parameters necessary to define a demand system, even if the equations have been defined in terms of the $n-1$ independent value shares and after the homogeneity restriction has been imposed, is $(n+2)(n-1)$.

The problem is that in most settings, the number of commodities (or commodity groups) n is so large that $(n+2)(n-1)$ is a large number. For example, Hausman, Leonard, and Zona (1994 p. 160) studied 135 different types of breakfast cereals; the number of parameters needed to describe how a change in the price of each cereal affected the demand for each one of the other cereals is, according to Muellbauer's formula: 18,358. It is not possible to reliably estimate that many parameters given the size of economic data sets. The only way for econometricians to solve this problem is to make assumptions that the parameters are related to each other in some ways. Econometricians often look to economic theory to supply assumptions which, once made, will greatly reduce the number of parameters which need to be statistically estimated. But are these assumptions valid?

If the demand-side data comes from surveys of all the expenditures of a household—as is available from, for example, the U.S. government's Consumer Expenditure Survey[1]—then, assuming price-taking fully-informed fully-rational utility-maximizing households, there are four groups of theoretical relationships that have to hold for each household (separately). These can be used to give structure to the data, which reduces the number of unknown parameters that have to be estimated, which ameliorates or even completely solves the problem that $(n + 2)(n - 1)$ is a large number. Therefore, these theoretical relationships are extremely valuable to the econometrician—assuming, of course, that one actually believes that households are price-taking fully-informed fully-rational utility maximizers, which economists in this area typically assume without question. We will call these theoretically-derived properties of demand curves "(T)" (for "theory") and postpone until Section 5 a detailed description of them and the groups they fall into, groups which are called 'adding up,' 'homogeneity,' 'symmetry,' and 'negativity.'

On the other hand, if the demand-side data is aggregate rather than individual, and/or partial—dealing with only one or a few commodities rather than all commodities—then there is no reason to suppose demand curves obey (T). In fact, in this case surprisingly little can be said about what properties the demand curve should have. This is known as the Sonnenschein-Mantel-Debreu theorem: the market demand curve for a market populated with utility-maximizing rational agents can take the shape of any function that is continuous, unchanged if all prices and incomes were to be redenominated (for example, expressed in pennies rather than in dollars), and in accordance with a balanced budget.[2] In particular, market demand curves cannot even be guaranteed to be downward-sloping, and "the problem that $(n+2)(n-1)$ is a large number" looms.

The upshot is that there is a rich theoretical structure available for an econometrician estimating an individual consumer's demand curve, but almost no theoretical support at all for estimating a market demand curve. If one

[1] See Wikipedia, *Consumer Expenditure Survey*, available at `https://en.wikipedia.org/wiki/Consumer_Expenditure_Survey`.

[2] The second criterion is called being "homogeneous of degree zero." The budget balance criterion is called Walras' Law, and comes from consumers' obeying their budget constraints. The original theorem dealt with excess demand curves; it was later extended to market demand curves. See Wikipedia, *Sonnenschein-Mantel-Debreu theorem*, available at `https://en.wikipedia.org/wiki/Sonnenschein%E2%80%93Mantel%E2%80%93Debreu_theorem`.

has comprehensive household-level expenditure data, this is not a problem: estimate each household's demand function, then add them all up to get the market demand function. But if one does not have comprehensive household-level expenditure data—which is the typical situation in antitrust—then there are many obstacles in the way of estimating the market demand curve.

Economists have responded to the difficulty of estimating the market demand curve by investigating the circumstances under which the market demand curve can be thought of as being derived from "the average consumer," called "exact linear aggregation," or from "a representative (but not mathematically average) consumer," called "exact nonlinear aggregation." We will use "exact aggregation" as an overall term to describe these two types of aggregation. If the assumptions permitting exact aggregation are plausible, then (T), the theoretical structure of the individual demand curve analysis, can be applied to the aggregate demand curve analysis.

So two things are needed in order to justify using (T) to help carry out an econometric demand analysis: adequate data, and plausible assumptions. Sections 3–5 reveal the details of what data and assumptions are needed, and that is all summarized in Figure 2. In that figure, the word "comprehensive" means "including all commodities and income," and the opposite of "household-level" is "aggregate." Also, the term "demand system" means a set of equations with the following characteristics. Each equation gives the demand for one commodity, as a function of its own price and of the prices of all other commodities, and as a function of income. Second, there are as many such equations as there are commodities. And last, the equations are not completely general, because that would be impossible to econometrically estimate, but the special functional forms imposed upon them to give them more mathematical structure are fully consistent with the standard neoclassical theory of the *single* utility-maximizing price-taking fully-informed rational consumer.

In this flowchart, there are two types of questions: questions about *data*, in Boxes 1, 4, and 6; and questions about *assumptions* on the behavior of individuals (not aggregates of individuals), in Box 3.

As we discuss at the beginning of Section 5, in antitrust cases, the answer to Box 1's "data" question is usually "no." Therefore, a key question becomes Box 3's "assumptions" question.

Box 3 refers to "Engel curves," which are the relationship between a household's income and its consumption of a good. When Engel curves are depicted in graphs, income is on the horizontal axis and demand (or "consumption") is on the vertical axis. When the Engel curve is upward-sloping, the good is "normal": when income rises, its consumption rises. On such a graph, when the Engel curve is downward-sloping, the good is "inferior":

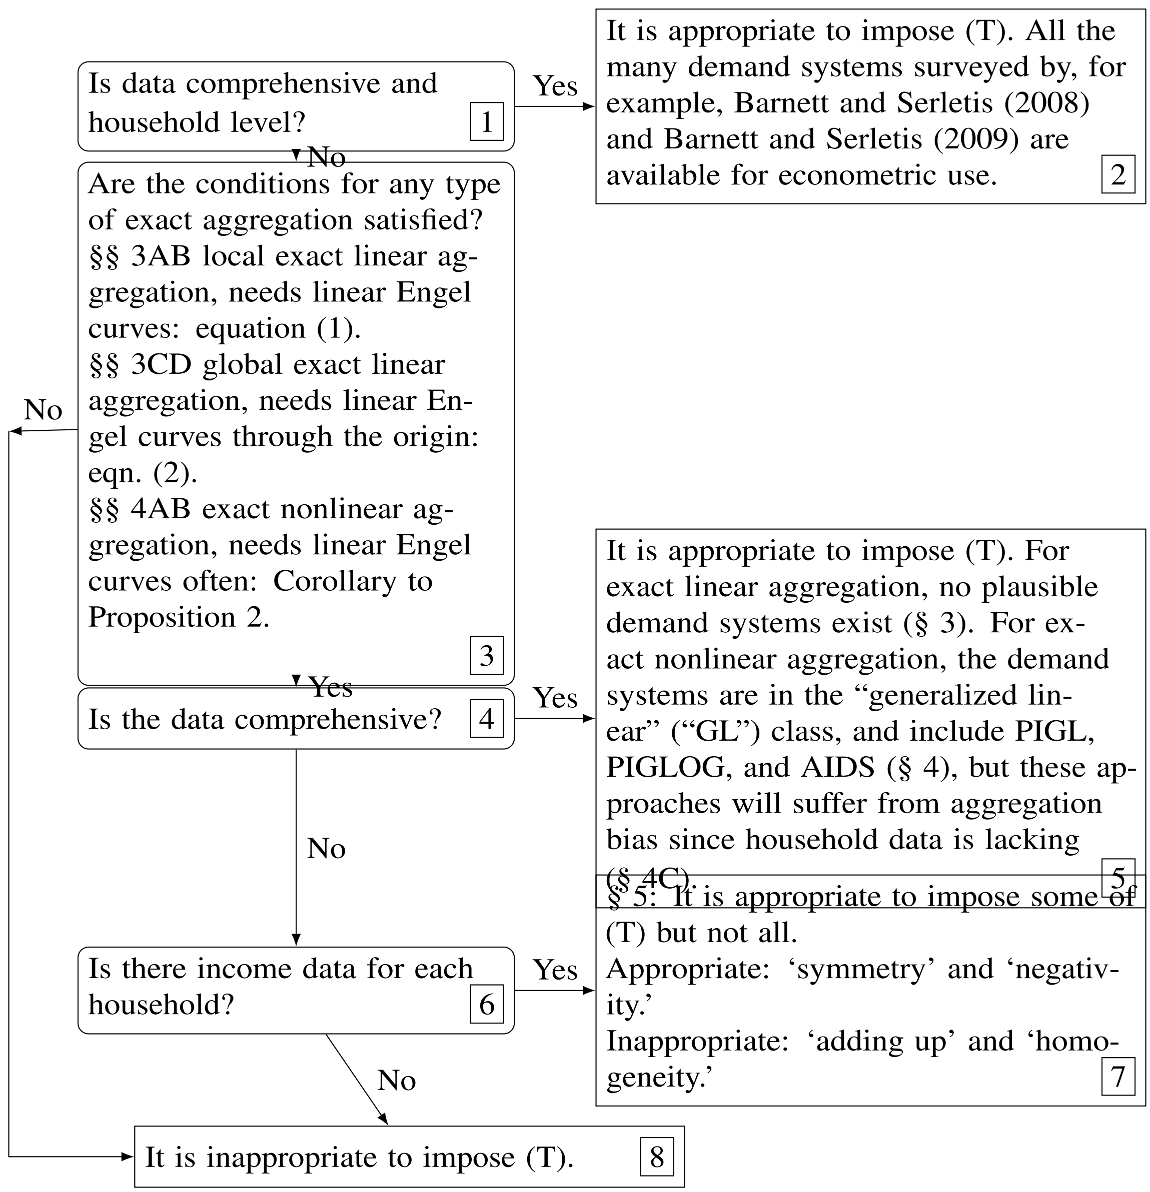

Figure 2. Assuming price-taking fully-informed fully-rational utility-maximizing individual consumers or households, this flowchart describes when it is appropriate to impose the theoretical restrictions (T) on the econometric estimation of demand curves. (Household-level data sets are usually comprehensive, so "no" to Box 1 is assumed to mean the data is not household-level, but it might be comprehensive; this is why Box 5 assumes there is no household-level data.)

when income rises, its consumption falls. If income changes have no effect on a consumer's demand for a product, the Engel curve for the product will be a flat line.

Sections 3 and 4 will show that, in general settings and in particular in antitrust, Engel curves are unlikely to be linear, making the answer to Box 3's question usually "no." That takes the flowchart to Box 8: imposing (T) is inappropriate.

In the implausible situation when Box 3's answer is "yes," one next faces the "data" questions of Boxes 4 and 6. Again as we discuss at the beginning of Section 5, in antitrust cases, the answers to the questions of Boxes 4 and 6 are usually "no," leading again to Box 8, where imposing (T) is inappropriate. Many antitrust econometric consumer demand analyses use one of the demand systems in Box 5, by implicitly, and implausibly, taking the answer to Box 3 to be "yes." Even if this is legitimate, Section 4C explains that there will be aggregation bias. Box 7 is explained in Section 5, which goes into the details of each one of the four (T) categories.

An important conclusion is that the often-hidden (or at least not discussed) assumptions of antitrust econometrics can be implausible. Because the mathematics are complicated, econometric results in antitrust cases may receive a lower level of scrutiny then they deserve. No judge would give a pass to an oral argument by a lawyer with a hidden implausible assumption. Empirical economists should be subject to equal scrutiny in the courtroom.

3. Exact Linear Aggregation: The market demand curve can be thought of as being derived from "the average consumer"

There are two types of exact linear aggregation: "local," for which only small changes in income are contemplated, and "global," for which the aggregation must be able to be carried out for any level of income.[3] In Subsection A we ask: "What assumptions have to be made in order for us to be able to think of the market demand curve as being derived from 'the average consumer,' locally?" In Subsection B we ask how realistic those assumptions are. Subsections C and D repeat Subsections A and B but for the global case. The conclusions we will come to are that the conditions for global exact linear aggregation are too narrow to be at all plausible, while the conditions for local exact linear aggregation are only somewhat more plausible. These dispiriting conclusions will prompt our interest in Section 4, where aggregation turns out to be possible under less restrictive, but still questionable, assumptions.

Throughout this Section we follow Section 6.1 of Deaton and Muellbauer (1980a) so closely that *we will identify text from that source merely by*

[3] Deaton and Muellbauer (1980a) p. 150.

enclosing it in 'single' quote marks. Our notation is slightly different from theirs, however.

A. Needed assumptions for local exact linear aggregation.

We will use the following notation in the rest of this paper. Suppose there are n goods, indexed by $i = \{1, 2, \ldots, n\}$. Also suppose there are H households, indexed by h, each of which has an income of I^h and a demand for good i of $d_i^h(I^h, \mathbf{p})$, where $\mathbf{p}$ is the list of prices $p_1, p_2, \ldots, p_n$ which all households face. *Exact linear aggregation* is possible if and only if for every good i, average demand is some function—call it $\bar{d}$—of average income. This condition is expressed mathematically as:

$$\frac{d_i^1(I^1, \mathbf{p}) + d_i^2(I^2, \mathbf{p}) + \cdots + d_i^H(I^H, \mathbf{p})}{H} = \bar{d}\left(\frac{I^1 + I^2 + \cdots + I^H}{H}, \mathbf{p} \right).$$

The problem for exact linear aggregation is that 'in general, no such function as $\bar{d}$ exists.' Note that $\bar{d}$ 'does not depend on the distribution of [incomes] I^h. Hence, for the equation to hold, a reallocation of a single unit of currency from any one to any other individual must leave market demands unchanged. This can only happen if [...] the marginal propensities to spend are identical for all consumers. Rich consumers must allocate changes in their outlay in exactly the same way as do poor consumers. This observation implies that the functions $d_i^h(I^h, \mathbf{p})$ must be linear in I^h, that is, for some functions α_i^h and β_i of $\mathbf{p}$ alone,

$$d_i^h(I^h, \mathbf{p}) = \alpha_i^h(\mathbf{p}) + \beta_i(\mathbf{p})\, I^h \tag{1}$$

where, although α_i^h is indexed by h, $\beta_i(\mathbf{p})$ is not.' (Equation (1) has the same form as the elementary equation for a straight line, $y = mx + b$, with the left-hand side playing the role of y, $\alpha_i^h(\mathbf{p})$ playing the role of b, $\beta_i(\mathbf{p})$ playing the role of m, and I^h playing the role of x.)

What we want to know is how restrictive assumption (1) is.

Saying, as we did in the paragraph before last, that "$\beta_i(\mathbf{p})$ is not indexed by h" is equivalent to saying that for each good i, *the Engel curve's slope for commodity i is the same for every household*. Furthermore, saying that β is not a function of I^h means that *the Engel curves are straight lines*. 'Suppose now that individuals maximize utility. In this case, (1) will hold if and only if' each consumer has quasi-homothetic preferences.[4]

[4] Blundell and Stoker (2007 p. 4619) attribute this result originally to a 1953 paper by Gorman.

B. How realistic are the needed assumptions for local exact linear aggregation? The consensus is that these assumptions are quite unrealistic.

'Viewed as necessary conditions for aggregation, quasi-homothetic preferences, or equivalently, linear Engel curves, are extremely stringent. For example, any commodity not consumed at low budget levels is immediately excluded. Consequently, if linear aggregation is to work at all, it can only do so for broadly defined composites of goods.'

This is very problematic in the antitrust context, where the goods are often so narrow, such as "different brands of dry cat food," that many households consume zero levels of many of them, even at high and moderate income levels, let alone at low income levels.

Barnett and Serletis (2009 pp. 22–3) bluntly write, "Linearity in expenditure implies marginal budget shares that are independent of the level of expenditure, suggesting that poor and rich households spend the same fraction of an extra dollar on each good. This hypothesis, as well as the hypothesis of expenditure proportionality, are too restrictive for the analysis of household budget data."

We conclude that the assumptions needed for local exact linear aggregation are unrealistic, and economists ought not to assume local exact linear aggregation. They usually do not.

C. Needed assumptions for global exact linear aggregation.

Thinking of (1) as an economic version of $y = mx + b$, we need both of the economic variables, income I^h (the analog of x) and quantity demanded $d_i^h(I^h, \mathbf{p})$ (the analog of y) to be positive or zero, not negative. That will require restrictions on the values of m (which is $\alpha_i^h(\mathbf{p})$) and b (which is $\beta_i(\mathbf{p})$), because unrestricted, $y = mx + b$ can certainly be negative, even if x is positive. This is the economists' next problem: to ensure that none of the d_i^h's (the analog of y) in (1) can be negative. If, as sometimes happens, either $\alpha_i^h(\mathbf{p})$ or $\beta_i(\mathbf{p})$ is negative, 'the permitted range of I^h has to be restricted.' However, if 'we do not wish to place any restriction on the I^h's and' instead demand that 'aggregation be possible for all $I^h \geq 0$, we must delete the intercepts $\alpha_i^h(\mathbf{p})$ since otherwise some demands will be negative.' [Another line of reasoning: if $I^h = 0$ we should have $d_i^h = 0$, and this requires $\alpha_i^h \equiv 0$ if (1) holds.] 'Hence, this "global" aggregation implies that' instead of (1) holding, one has

$$d_i^h(I^h, \mathbf{p}) = \beta_i(\mathbf{p})\, I^h\,. \tag{2}$$

This means that the Engel curves are all straight lines through the origin, and have the same slope for every household; budget shares $p_i d_i^h / I^h = p_i \beta_i(\mathbf{p})$ *are independent of total expenditure (which Barnett and Serletis* *(p. 21) point out contradicts "Engel's Law, according to which the budget share of food is smaller for rich than for poor households"); and expenditure elasticities for all consumers and all goods are unity (meaning that every 1% increase in income induces every consumer to buy 1% more of every good).*[5] If this holds, it implies that

$$\frac{d_i^1(I^1,\mathbf{p}) + d_i^2(I^2,\mathbf{p}) + \cdots + d_i^H(I^H,\mathbf{p})}{H} = \beta_i(\mathbf{p})\frac{I^1 + I^2 + \cdots + I^H}{H},$$

so, indeed, average market demand (the left-hand side) is a function of average income ($(I^1 + I^2 + \cdots + I^H)/H$).

D. How realistic are the needed assumptions for global exact linear aggregation?

The restrictions needed for global exact linear aggregation are so wildly unrealistic that most economists shun (2) in favor of (1) and simply hope that their estimates of $\alpha_i^h(\mathbf{p})$ and $\beta_i(\mathbf{p})$, together with observed or realistic levels of I^h, never lead to a negative d_i^h.

4. Exact Nonlinear Aggregation: The market demand curve can be thought of as being derived from "a representative consumer"

Because Section 3's conditions for exact linear aggregation are so unrealistic, in this Section we investigate the less restrictive assumptions needed for exact nonlinear aggregation. This is an important topic because economists often impose theoretical restrictions from individual demand curves, restrictions we called (T) in Section 2, onto aggregation demand curves, and that is only justified if the conditions for exact linear or nonlinear aggregation hold; and given that we just concluded that the conditions for exact linear aggregation are very unlikely to hold, the conditions for exact nonlinear aggregation are quite important to know.

In Subsection A we ask what assumptions have to be made in order for us to be able to think of the market demand curve as being derived from "a representative consumer" (not by Section 3's "average consumer"). In Subsection B we ask how plausible those assumptions are. In Subsection C we ask what data is required to carry out econometrics based on those assumptions.

Overall, the conclusions are somewhat disappointing. While the conditions for exact nonlinear aggregation are certainly less restrictive than for exact linear aggregation, they are not trivial, and probably require commodities to be considered in large categories rather than analyzed individually—which is troubling for antitrust analyses, which usually concern particular

[5] From Deaton and Muellbauer (1980a p. 144), economists describe such preferences by saying that all the consumers' preferences are "homothetic."

commodities, not broad categories of commodities. In the language of Figure 2, even allowing the option of exact nonlinear aggregation may not lead to an answer of "yes" for Box 3. If it does, and we get to Box 5, it turns out that we will require household-level data in order to carry out exact nonlinear aggregation (that is, without aggregation bias)—but if we had household-level data, we could have been in Box 2, with no need to use the more constrained, aggregate methods of Box 5. Therefore, the most likely situation will be one in which we wish to aggregate because we do not have household-level data (we cannot be in Box 2), but lacking household-level data, we cannot carry out the aggregation exactly. Therefore, an aggregation bias will be present, meaning that the promise of exact aggregation (that there is justification for imposing restrictions (T) without any bias) will not be achievable in practice.

A. Needed assumptions for exact nonlinear aggregation.

Throughout Section 4A we follow Section 6.2 of Deaton and Muellbauer (1980a) so closely that *we will identify text from that source merely by enclosing it in 'single' quote marks*.

Exact nonlinear aggregation restricts average aggregate budget share for the *i*th good, $p_i(d_i^1+d_i^2+\cdots+d_i^H)/(I^1+I^2+\cdots+I^h)$, to depend on prices and, not on *average* income (because that leads to exact linear aggregation), but on a "representative" level of income which we will denote by I_0, 'which itself can be a function of the distribution of expenditures and of prices. If this holds, the market pattern of demand can be thought of as deriving from the behavior of a single representative individual endowed with income I_0 and facing prices **p**.'

In order for such a representative consumer to exist, one must place

> '...strong restrictions upon Engel curves; note, for example, that for each household the slopes of the different Engel curves will vary linearly with one another as total expenditure changes at constant prices. This does not, of course, imply that the Engel curves themselves are linear.
>
> 'Since these linear relations occur...the name given to the conditions for consistent nonlinear aggregation is generalized linearity (GL).... A particularly interesting special case occurs when the representative expenditure level is independent of prices and depends only on the distribution of expenditures. This case, known as price independent generalized linearity (PIGL) occurs when the microcost functions take the form [here follows an equation involving "α"...] When α tends towards zero, [that equation becomes the] form known as PIGLOG.'

Although Deaton and Muellbauer write that in order for a representative consumer to exist, strong restrictions have to be placed on Engel curves, they do not explain those restrictions beyond the first paragraph in the quotation displayed above. We now further explain those restrictions.

To begin, we have to translate into mathematical terms Deaton and Muellbauer's contention that 'for each household the slopes of the different Engel curves will vary linearly with one another as total expenditure changes at constant prices.' We also want to show that that condition is, as Deaton and Muellbauer claim, necessary and sufficient for existence of a representative consumer, because neither Deaton and Muellbauer nor anyone else seems to have published its proof. (All of the proofs for this paper are in the Appendix.)

Proposition 1. *There is a representative consumer if and only if there exist some numbers A_{hij} and B_{hij} such that*

$$\frac{\partial\, d_i^h(I^h, \mathbf{p})}{\partial I^h} = A_{hij} \frac{\partial\, d_j^h(I^h, \mathbf{p})}{\partial I^h} + B_{hij} \quad \text{for each } h \text{ and for all goods } i \neq j. \tag{3}$$

In words: for each household, the slope of the Engel curve for good i is equal to some constant "A_{hij}" times the slope of the Engel curve for good j, plus another constant "B_{hij}."

To help with interpretation, it is going to be useful to have two more results:

Proposition 2. *Suppose a representative household exists. Then the ratio of the second derivative of household h's Engel curve for good i to the second derivative of household h's Engel curve for good j is a constant.*

Furthermore, if a representative household exists and if, for some household h, there exists an interval of incomes $\hat{I}$ on which d_j^h is a linear function of I (for fixed $\mathbf{p}$), then for that household and for all other goods $i \neq j$, d_i^h is a linear function of I on $\hat{I}$.

Corollary. *If for some household h, there exists an interval of incomes $\hat{I}$ on which d_j^h is a constant function of I (meaning that on that interval, household h's consumption of good j stays the same as household h's income changes), then for that household and for all other goods $i \neq j$, d_i^h is a linear function of I on $\hat{I}$.*

Proposition 2's second sentence implies that *if both of a pair of Engel curves are concave (or convex) for one value of I^h, they are concave (or convex) for all values of I^h.* Similarly, *if one of a pair of Engel curves is concave and the other is convex for one value of I^h, then they will have opposite convexity for all values of I^h.*

B. How plausible are the restrictions needed for exact nonlinear aggregation?

The convexity or concavity restrictions just described are arbitrary and artificial. There is absolutely no reason to assume consumers behave in that way, nor do economists claim that they do.

In the passage quoted above, Deaton and Muellbauer admitted that Proposition 1 represents "strong restrictions" that are needed in order for a representative consumer to exist. However, they added, "This does not, of course, imply that the Engel curves themselves are linear." The reason Deaton and Muellbauer were at pains to point out that (3) technically does not imply linear Engel curves is because linear Engel curves are completely implausible (except in the case of zero consumption of a commodity or in the case of a commodity whose consumption has completely flattened out). That was why we rightly rejected, as being unrealistic, the "Engel curves are linear in I for all values of I^h" assumption when discussing exact linear aggregation. Linear Engel curves are so implausible that econometricians usually show no interest in even considering them: instead they study the completely different question of whether or not Engel curves are log-linear (that is, whether the logarithm of consumption is a linear function of the logarithm of income).

The above Corollary is key to understanding that Deaton and Muellbauer's position, that linear Engel curves are not necessary for a representative consumer, while technically true, is very much weaker than they let on. Consider the following situations.

1. *There is a commodity whose consumption by one household completely flattens out for large incomes.* (For example, if the demand for Iams cat food does not change when household h's income I^h is greater than \$150,000.) Then by the Corollary, in order for a representative consumer to exist, *all* the Engel curves of that household *do have to be linear* from that income level on up.

2. *There is a commodity whose consumption flattens out for large incomes for all households.* (For example, if the demand for Iams cat food does not change when any household's income is greater than \$150,000.) Then by the Corollary, in order for a representative consumer to exist, *all* the Engel curves *of every household do have to be linear* from that income level on up.

3. *There is a commodity which a household never consumes.* Then by the Corollary, in order for a representative consumer to exist, *all* the Engel curves of that household *do have to be linear* for *all* values of income.

4. *There is a commodity which is an inferior good for a household.* Then by the Corollary, in order for a representative consumer to exist *without* all the household's other Engel curves being linear, the consumption of this good must fall with income but never become constant. For example, if instant ramen noodles are an inferior good, to have a representative consumer exist without all the household's other Engel curves being linear, as income rises, the noodle consumption could never actually reach zero. An asymptotic approach to zero would be allowed, but that is meaningless: even a theorist should reject the claim that there is any actual economic difference between "$\{1/100, 1/200, 1/300, 1/400, \ldots\}$ cups of ramen noodles demanded as income rises" and "$\{0, 0, 0, 0, \ldots\}$ cups of ramen noodles demanded as income rises."

So while Deaton and Muellbauer are technically correct that (3) does not imply linear Engel curves, the Corollary shows that in many if not most real-world situations, (3) is going to imply linear Engel curves.

According to the Corollary, the only way to avoid "a representative consumer implies all Engel curves are linear over a range of income levels," is to avoid having any commodity's consumption be flat (or linear) on that range of income levels. That almost certainly requires defining "goods" i, j, ..., very broadly, in categories such as "clothing," "food," and "shelter," steering clear of finely-detailed categories such as "ramen noodles" or "Iams cat food." However, antitrust cases do not concern monopoly power over broad categories like "clothing," "food," and "shelter"; they concern finely-detailed categories such as ramen noodles or cat food. With finely-detailed categories characterizing antitrust cases, there is grave doubt that the assumptions needed to assume existence of a representative agent will be satisfied in practice.

Moreover, even if those assumptions are satisfied in practice, one still has Proposition 2's weirdly linked concavity and convexity conditions standing in the way of accepting exact nonlinear aggregation as being plausible.

C. Aggregation Bias

What data is needed to implement exact nonlinear aggregation? As mentioned above, in the Deaton and Muellbauer displayed quotation a paragraph before Proposition 1, a representative household exists if and (dismissing the highly implausible exact linear case) only if preferences belong to the "generalized linearity" (GL) class, and special cases within the GL class are the PIGL class and the PIGLOG class. (See also Box 5 of Figure 2.) Within the PIGLOG class is the "Almost Ideal Demand System" (AIDS), which is the most commonly used econometric approach to market demand estimation in antitrust; see Deaton and Muellbauer (1980b, p. 313).

Since the AIDS class of preferences is a subset of the GL class, if one assumes AIDS preferences, then aggregate preferences can be generated by a representative consumer, and there will be no aggregation bias. This makes the title of a 1996 paper by Mittelhammer, Shi, and Wahl initially puzzling: "Accounting for Aggregation Bias in Almost Ideal Demand Systems." The problem turns out to be that in order to calculate the AIDS system, you need to know not only that a representative consumer exists, you also have to be able to calculate exactly who the representative consumer is, and that requires household-level data which, if you had it, would put you in Figure 2's Box 2, where there is no need to use AIDS nor to have a representative consumer. Here is the way this is explained by the authors of the above 1996 paper in a 1994 working paper with almost the same title which lists the authors in a different order (Wahl, Mittelhammer, and Shi (1994 p. 194)):

> Clearly the calculation [for the AIDS model] of the weighted geometric mean of expenditures, x^*, in the aggregate share equation *requires detailed information on the distribution of total expenditures over consumers*. Unfortunately, most empirical data available for applied demand studies are measured at an aggregated level, and the information necessary for computing x^* is often not available in practice. *In demand studies utilizing aggregate time series data, researchers often use the simple average of individual expenditures (i.e., per capita expenditure) to replace the geometric mean.* Deaton and Muellbauer (1980a and 1980[b]) have shown that if the *average* aggregate budget share is to be specified as a function of prices and per capita expenditure, this requires the restrictive conditions of exact linear aggregation. In the case of exact (price-independent) nonlinear aggregation, such as AIDS, it is required that the aggregate budget share, $\bar{w}_i$ depend on prices and a representative level of total expenditure x_0 which itself depends on the distribution of expenditures. In this case, "the market pattern of demand can be thought of as deriving from the behavior of a single representative individual endowed with total expenditure x_0 and facing prices p." (Deaton and Muellbauer 1980a p. 154). In the case of PIGLOG preferences, it is clear from (6) that the appropriate level of representative expenditure is given by $x_0 = \prod x_h^{r_h}$. It follows that *using $\bar{x}$ [the average value of expenditures x] in place of x_0 constitutes a misspecification of the AIDS model.*[6]

[6] Compare the "x" in (16) of Deaton and Muellbauer (1980b) with the "$\bar{x}$" in (20) of that

This confirms that the AIDS model itself has no aggregation bias, but if one misspecifies it, then the misspecified AIDS model has an aggregation bias. The obvious solution—do not misspecify the AIDS model—fails whenever one *lacks the household-level data required to correctly specify the AIDS model*. The authors' 1996 paper (Mittelhammer, Shi, and Wahl (1996)) suggests an approximate solution: supplement one's aggregate data set with a different, household-level data set.

> ...it is evident that, to calculate the expenditure aggregation bias term, time-series information on the number of households and on individual households' shares of aggregate expenditure are needed. Information on the shares of aggregate expenditure across households is generally unavailable or inaccessible. However, time-series information on the number of households in different income categories is readily available for most developed economies and can provide valuable information for closely approximating the income distribution and aggregation bias term in the aggregate AIDS model. (p. 250)

To summarize: to use the AIDS model exactly, one needs to have household-level data, in order to calculate the "representative expenditure" (otherwise you are misspecifying the AIDS model and there will be aggregation bias). But referring back to Figure 2, if one had household-level data, one could be in Box 2, in which one would not need to use AIDS, nor be bothered about existence of a representative consumer, nor about aggregation at all. Many economists want to use AIDS, and assume a representative consumer exists, when they do not have household-level data (the answer to Figure 2's Box 1 is "no"); but we now learn that one cannot use AIDS (at least not exactly) in Box 5 because one cannot calculate the "representative expenditure." The three co-authors suggest an approximation to get out of this Catch-22: substitute "representative expenditure" calculated from economy-wide income distribution data, which is widely available, in the place of the "representative expenditure" for the households that generated one's own data. Essentially, this is a work-around which enables one to proceed, albeit approximately, in Box 5. More recent literature uses "aggregation factors" to describe "the degree of bias in recovering (individual) price and income elasticities from aggregate data alone," as Blundell and Stoker (2007 p. 4621) put it.

If lack of household-level data causes one to be unable to exactly implement AIDS estimation, the other response is to abandon the attempt to impose (T). Stoker (1993 p. 1829) calls this abandoning "descriptions that

paper (it on p. 314 says that "$\bar{x}$ is the average level of total expenditure x_h").

are straitjacketed by the capricious enforcement of restrictions of optimizing behavior by a single individual." Freed from the straitjacket imposed by wanting to use (T), one can then use any of the demand systems in Figure 2's Box 2 *without (T)*, and so one can use a demand system which, unlike AIDS, "seem[s] to do a good job of fitting the data, such as the QUAIDS system of Banks, Blundell and Lewbel (1997)" (Blundell and Stoker 2007 p. 4622). On the other hand, the straitjacket was extremely helpful in shrinking down from $(n + 2)(n - 1)$ the number of parameters that needed to be estimated, so the straitjacket will be missed in some situations: standard errors will be higher, fewer parameters will be significantly different from zero, and the overall fit of the model will be worse.

Nevertheless, Blundell and Stoker's outlook for econometricians who do not have any household-level data (Box 1's answer is "no") is grim (op. cit. p. 4658):

> While we have advanced the idea of using aggregation factors (derived from time-series of individual data) to summarize the impacts of aggregation, the specific method one uses is less important than the ability to use all available types of information to study economic relationships. That is, it is important to study any relationship among economic aggregates with individual data as well as aggregate data, to get as complete a picture as possible of the underlying structure. Even though modeling assumptions will always be necessary to develop explicit formulations of aggregate relationships, testing those assumptions is extremely important, and is not possible without extensive individual data over sequential time periods.

Stoker (1993) explained the basic problem a few years earlier:

> Models that account for individual heterogeneity will typically not be estimable using data on economy-wide averages alone; additional data on distributional composition. . . , or micro data on individual behavior, will need to be incorporated. This should come as no surprise; to study relations that involve heterogeneous individual responses without distributional information is analogous to studying dynamic relations without using data over time. [p. 1836] [. . .]
>
> Whether a representative agent model fits the data or not, there is no realistic paradigm where the parameters of such a model reflect only behavioral effects, uncontaminated by compositional considerations. The application of restrictions appropriate for individual behavior directly to aggregate data [that is, applying (T)] is a practice without any foundation, and leads to

> biases that are impossible to trace or measure with aggregate data alone. [p. 1870]
>
> Hand-in-hand with the necessity of using all relevant data is the necessity of checking or testing all relevant assumptions underlying a model. Aside from a platitude of good empirical work, it is important to stress the testing aspect here because altogether too little attention has been paid to checking or testing assumptions required for aggregation, relative to assumptions on the form of individual behavioral models. [p. 1870] [. . .] Approaches that neglect individual heterogeneity, such as pure representative agent modeling, should be abandoned. [p. 1871]

One cannot test the assumptions required for a representative consumer using only aggregate data. When one tests them using household-level data, what does one find? Barbett and Serletis say, rather mildly, "most of the commonly used PIGLOG specifications are of rank two, and thus do not have enough flexibility in modeling the curvature of Engel curves with large variations in income" (2008 p. 218). Banks, Blundell, and Lewbel (1997 pp. 527–8), writing after extensive use of household-level data from the U.K. Family Expenditure Survey, say that their "quadratic logarithmic class nests both the Almost Ideal (AI) model of Deaton and Muellbauer and [. . .]. Unlike these demand models, however, the quadratic logarithmic model permits goods to be luxuries at some income levels and necessities at others. The empirical analysis we report suggests that this is an important feature. . . . The specific form we propose—the Quadratic Almost Ideal Demand System (QUAIDS)—is constructed so as to nest the AI model and have leading terms that are linear in log income while including the *empirically necessary* [emphasis added] rank 3 quadratic term" (which AIDS lacks). Stoker (1993 p. 1855) is the most emphatic; he says about AIDS (emphasis mine):

> In particular, (4.12) [an equation "used in Deaton and Muellbauer's (1980[b]) estimation"] rests on the assumption that a) (4.7) [the AIDS demand system] is valid, with no individual heterogeneity in demands aside from income effects and b) [. . .]. Each of these assumptions is testable *with micro data* [emphasis added], and patently unrealistic. . . .

Overall, then, it is not only the representative agent approaches in Figure 2's Box 5 which have sustained heavy criticism, but *any* attempt to proceed when the answer to Box 1 is "no." Yet in empirical antitrust work the answer to Box 1 typically is "no."

The next section deals with the situation in Box 7.

5. Demand Estimation when Lacking Data on Some Goods

Ulrick (2014 pp. 138–9) describes the sort of data typically available in antitrust analysis:

> The AIDS model above is often estimated with Nielsen and IRI scanner data to generate brand-level demand elasticities. Scanner data are particularly suited for this type of analysis. The data are almost always available weekly by SKU and city. The data will include total retail dollars, equivalised units, units, and marketing variables. The data generally cover food stores (the grocery channel), but in some cases mass merchandiser data is also available (ie Target, Kmart).

This is similar to the data available to Hausman, Leonard, and Zona (1994). In Figure 2, this implies that the answer to the "data" questions of Boxes 1, 4, and 6 are "no." In this section, we show that if somehow one gets to Box 6 (which requires the "assumptions" Box 3 answer to be "yes," given that the "data" Box 1 answer is "no"), then if the answer to the income "data" question of Box 6 is "no," one gets to Box 8, (T) is inappropriate, and if the answer to the income "data" question of Box 6 is "yes," one gets to the somewhat nuanced conclusions of Box 7.

The question raised if one arrives at Box 6 is to what extent one can impose the restrictions (T)—remember the whole reason why many econometric antitrust analyses work is they have have imposed (T)—if one lacks data on many goods the household purchases. To answer this, we need to go one-by-one through the list of restrictions which constitute (T). Up to this point, we have avoided saying much about what the restriction list (T) actually contains, but now, not only do we have to list the restrictions, but we have to mathematically check each one to make sure they still apply if one only has data on some commodities, not all commodities. What we will find is that two of the restrictions in (T) fail to be applicable to the case when one lacks data on some goods; the other two restrictions in (T) do apply to that case, but only if income data is available, and if income data is not available, those restrictions should not be imposed either. This calls into question the validity of many econometric studies in which (T) is imposed.

We follow the order of restrictions in Ulrick's treatment (2014, Section C).

Adding up: This is the restriction that the sum of expenditures, $p_1 d_1^h + p_2 d_2^h + \cdots + p_n d_n^h$, is equal to household h's income; or the corresponding aggregate restriction that the sum of expenditures of all the households is equal to the sum of all the household's incomes. If one only has data on, for example, commodities 1 and 2, there is no restriction to impose, because expenditures on commodities 1 and 2 are not constrained to be any particular number. There is one potential work-around. Suppose commodities 1 and 2 are two types of pet food, suppose none of the other commodities are pet foods, and suppose that regardless of how much prices and income change, this household always spends a fixed amount of money on pet food. Then $p_1 d_1^h + p_2 d_2^h$ would always equal this fixed amount of money, and this would function like the more general adding up restriction. The problem is that the assumption "regardless of how much prices and income change, this household always spends a fixed amount of money on pet food" is probably incorrect.

It is true that in the context of multi-stage budgeting (as in Hausman, Leonard, and Zona (1994)), assuming a fixed expenditure on, say, pet food may not be problematic. Multi-stage budgeting implies that preferences are not of the AIDS or GL form (but instead have some separability properties—see Deaton and Muellbauer (1980a, Chapter 5)). Hausman, Leonard, and Zona (p. 162) say "our econometric specification at the lowest level is the 'almost ideal demand system' of Deaton and Muellbauer," but to be clear, this only means that it uses the AIDS form for one of their three budgeting stages; their consumers do not, overall, have AIDS preferences, so no representative consumer exists and they should not impose (T). (Unfortunately they do impose one component of (T), symmetry (see pages 163 and notes on Tables 2–4).) On the other hand, the true adding up condition itself (ignoring the other elements of (T)) just comes from the budget constraint, and it applies in the most general cases: it is one of the few conditions which the market demand curve has even in the Sonnenschein-Mantel-Debreu theorem.

Note that even the lowest-level model of Hausman, Leonard, and Zona is not a true AIDS model because its dependent variable is, for example, the amount of money spent on Budweiser premium beer as a fraction of the amount of money spent on all premium beers, whereas in a true AIDS model, the corresponding dependent variable would be the amount of money spent on Budweiser premium beer as a fraction of the amount of money spent on all commodities. With data like Hausman, Leonard, and Zona's, and no additional data on incomes, there is no way to know how much money the consumers are spending on all commodities, so there is no way to construct a true AIDS model, so there is no justification for (T).

Homogeneity: This restriction requires that quantity demanded be unchanged if one multiplies all prices and income by a constant. If household consumes more than (for example) two commodities but one only has data on the first two commodities, one cannot impose this restriction. After all, if the household consumes more than two commodities then it is false that "if only the prices p_1 and p_2 are multiplied by a constant, d_1^h and d_2^h are unchanged."

Symmetry: The Symmetry condition is one of the most commonly-imposed parts of (T), but unfortunately, there is no intuitive, non-mathematical interpretation of it. The Symmetry condition is that, for all goods i and $j \neq i$,[7]

$$\frac{\partial d_i^h(I,\mathbf{p})}{\partial p_j} + \frac{\partial d_i^h(I,\mathbf{p})}{\partial I} d_j^h(I,\mathbf{p}) = \frac{\partial d_j^h(I,\mathbf{p})}{\partial p_i} + \frac{\partial d_j^h(I,\mathbf{p})}{\partial I} d_i^h(j,\mathbf{p}) \qquad (4)$$

or, at the aggregate level, (4) dropping the household indexes h. Notice that the only difference between the two sides of (4) is that the i and the j are interchanged, which gives the condition its name, symmetry. Notice also that if Engel curves were flat, the second terms on each side of (4) (called the "income effect" terms) would be zero, and (4) would collapse to $\partial d_i^h/\partial p_j = \partial d_j^h/\partial p_i$, that is, "the increase in demand of good i when the price of good j changes is equal to the increase in demand of good j when the price of good i changes," which is as close to an intuitive explanation of (4) as one is likely to get. (The Appendix shows how to express the symmetry condition in more conventional but more indirect way; see its Proposition 3.)

The Symmetry condition should hold even if one has data only on a few of the commodities which the consumer purchases. However, the income effect ($\partial d/\partial I$) terms in (4) could only be calculated if one had data on income (at the household or, assuming a representative agent, the aggregate level). If data on income is lacking, as is often the case, then even though the restriction (4) should hold, there would be no way to do the calculations necessary to impose it.

If the Symmetry condition ought to hold, it is a considerable help: Muellbauer (1975 p. 525) writes, "Of the restrictions implied by utility theory, by far the most important saving in parameters results in the $\frac{1}{2}n\,(n-1)$ restrictions implied by the symmetry of compensated cross-price effects," which is this condition. This is why it is imposed so often, even though its imposition is apparently not always theoretically justified.

Negativity: Negativity has two aspects. The first is that for all goods i,

$$\frac{\partial d_i^h(I,\mathbf{p})}{\partial p_i} + \frac{\partial d_i^h(I,\mathbf{p})}{\partial I} d_i^h(I,\mathbf{p}) \leq 0\,. \qquad (5)$$

[7] Ulrick and most other authors express this using "Hicksian," or "compensated," demand functions, instead of the "Marshallian" demand functions used in this paper. The expression using Hicksian demand functions is much more compact.

In the absence of income effects, this simply says that "demand curves are downward-sloping." In the presence of income effects, it is difficult to fruitfully express (5) in words.[8] This aspect of the Negativity restriction should hold even if one has data only on a few of the commodities which the consumer purchases, although just like for the almost identical terms in (4), income data is needed to compute the left-hand side.

The second aspect of Negativity involves more mathematics, and the reader unfamiliar with quadratic forms is invited to skip this paragraph. Let the left-hand side of (4) be abbreviated S_{ij}. If one lacks data on how demand varies when income varies, the S_{ij}'s cannot be calculated and "Negativity's second aspect" restrictions cannot be imposed. Otherwise, with n being the number of commodities, denote by **S** (sometimes called the "Slutsky Matrix" or the "Substitution Matrix") the $n \times n$ matrix whose (i, j) element is S_{ij}. The restriction is that **S** be negative semidefinite (and symmetric, which underlies the above symmetry restriction). It can be shown that the submatrix of **S** obtained by retaining only some of its rows and the corresponding columns should also be negative semidefinite, so this restriction does carry through to the case where there is data only on some commodities.[9]

6. Consumer Surplus and Demand-curve Estimation

One purpose for which econometric estimates of demand and supply curves in antitrust are calculated is to make judgments about economic welfare changes. If this is done using Kaldor's Compensating Variation and Hicks's Equivalent Variation, no restrictive assumptions about consumer preferences have to be made. However, if this is done instead using consumer surplus as the welfare measure, it has been known for a long time—in some sense, since Alfred Marshall[10]—that the commodity being studied must have no income effect, that is, that its consumption must not vary

[8] Unless one is familiar with footnote 7's Hicksian demand curves, in which case (5) simply says that Hicksian demand curves are downward-sloping.

[9] Proof: Reorder the commodities so that the commodities which one has data on are the first ones. Suppose there are $n' < n$ such commodities. Since **S** is negative semidefinite, all of its principal minors of order r alternate in sign beginning with ≤ 0 for $r = 1, 2, \ldots, n$. This means that all of **S**'s principal minors of order r alternate in sign beginning with ≤ 0 for $r = 1, 2, \ldots, n' < n$. The latter means that the submatrix of **S** obtained by retaining only n' of its first few rows and the corresponding columns is also negative semidefinite.

[10] Marshall wrote that "In regard to different people allowance may have to be made where necessary for differences of sensibility and for differences of wealth: but it is seldom needed in considering large groups of people" in his margin notes for pages 130 and 131, Book III ("On Wants and their Satisfaction") Chapter VI ("Value and Utility") Section 3, of Marshall (1920).

with income and therefore that its Engel curve must be horizontal:

$$d_i^h(I^h, \mathbf{p}) = \alpha_i^h(\mathbf{p})\,. \tag{6}$$

It is helpful to contrast this with the other assumptions discussed in this paper. For global exact linear aggregation, we needed

$$d_i^h(I^h, \mathbf{p}) = \beta_i(\mathbf{p})\, I^h\,, \tag{2}$$

"Engel curves are all straight lines through the origin, and have the same slope for every household." For local exact linear aggregation, we needed

$$d_i^h(I^h, \mathbf{p}) = \alpha_i^h(\mathbf{p}) + \beta_i(\mathbf{p})\, I^h\,, \tag{1}$$

"the Engel curve's slope for commodity i is the same for every household, and Engel curves are straight lines." And for exact nonlinear aggregation, we needed

$$\frac{\partial}{\partial I^h}\, d_i^h(I^h, \mathbf{p}) = A_{hij}\, \frac{\partial}{\partial I^h}\, d_j^h(I^h, \mathbf{p}) + B_{hij} \tag{3}$$

for each h and for all goods $i \neq j$: for each household, the slopes of the household's Engel curves for different commodities will vary linearly with one another as total expenditure changes at constant prices (and the implications of this in Proposition 2 and its corollary). While (6) clearly contradicts (2), at first glance it seems to be compatible with (1), in the special case where $\beta_i(\mathbf{p}) = 0$, and it seems to be compatible with (3), which it turns into $0 = 0 + B_{hij}$, in the special case where $B_{hij} = 0$. However, (1) has to hold for every commodity i, and (3) has to hold for every pair of commodities i and j, whereas (6) cannot hold for every commodity i, because then income levels would affect no one's consumption for any commodity, which is completely at odds with empirical findings and would lead to violation of consumers' budget constraints. One could have (1) hold for all commodities and have $\beta_i(\mathbf{p})$ be zero for a subset of commodities, so that consumer surplus would be an exact welfare measure on that subset. One could also have (3) hold for all commodities and have B_{hij} be zero for a subset of commodities on which (6) also holds, so that consumer surplus would be an exact welfare measure on that subset. Before taking these steps, though, one should first determine whether (6) is at all a plausible description of the way consumers buy the good i of interest. If i is an important good, as is likely if it is of antitrust interest, then the plausibility of (6), with consumption independent of income, is questionable.

Conclusion

To non-economists, econometric analyses such as consumer behavior predictions stemming from the Almost Ideal Demand System may seem to be unimpeachable mathematical edifices. They are not. As summarized in Figure 2, even accepting the questionable assumptions that consumers are fully informed, fully rational, and maximize a single utility function even if they are part of a multi-person household, econometric results from consumer demand systems are only valid if particular data is available and if further particular assumptions on the shape of Engel curves, backed by no empirical data, are made. In antitrust litigation, disputes concern narrow types of commodities, rather than broad categories such as "food" or "shelter." In that situation, we have shown that the only way to assume existence of a representative consumer—the assumption giving rise to the "(T)" restrictions so commonly reflected in consumer demand systems—is to assume linear Engel curves, which are implausible. Without a representative consumer, however, estimation of demand systems becomes much more challenging, and much less reliable, unless one has comprehensive, household-level data, which antitrust economists usually lack. Moreover, the most commonly-used tool to make welfare judgments in antitrust, which is consumer surplus, is only valid under assumptions which are rejected by economists who specialize in studying consumer behavior. In summary, antitrust econometrics relies on often-implausible assumptions and usually has available only inadequate data, making it an imprecise, limited tool which should be wielded with humility.

Appendix

Proof of Proposition 1. Following Muellbauer (1975), let $w_i = (p_i d_i)/I$ be the value share of good i, where we only consider one household and we suppress the h superscript. From Muellbauer's Theorem 3, using slightly different notation, there exist functions $A'_i(\mathbf{p})$, $B'_i(\mathbf{p})$, and $v(I, \mathbf{p})$ such that[11]

$$w_i(I, \mathbf{p}) = v(I, \mathbf{p}) A'_i(\mathbf{p}) + B'_i(\mathbf{p}) \tag{7}$$

$$w_j(I, \mathbf{p}) = v(I, \mathbf{p}) A'_j(\mathbf{p}) + B'_j(\mathbf{p}) \,. \tag{8}$$

Solving (8) for $v(I, \mathbf{p})$ and substituting it into (7) leads to

$$\begin{aligned} w_i(I, \mathbf{p}) &= \frac{w_j(I, \mathbf{p}) - B'_j(\mathbf{p})}{A'_j(\mathbf{p})} A'_i(\mathbf{p}) + B'_i(\mathbf{p}) \\ &= \frac{A'_i(\mathbf{p})}{A'_j(\mathbf{p})} w_j(I, \mathbf{p}) + B'_i(\mathbf{p}) - \frac{B'_j(\mathbf{p}) A'_i(\mathbf{p})}{A'_j(\mathbf{p})} \end{aligned}$$

which for simplicity we will write as

$$w_i(I, \mathbf{p}) = A''_{ij}(\mathbf{p}) \, w_j(I, \mathbf{p}) + B''_{ij}(\mathbf{p}) \,.$$

Substituting in the definition of value shares leads to

$$\begin{aligned} \frac{p_i \, d_i(I, \mathbf{p})}{I} &= A''_{ij}(\mathbf{p}) \frac{p_j \, d_j(I, \mathbf{p})}{I} + B''_{ij}(\mathbf{p}) \,, \text{ so} \\ d_i(I, \mathbf{p}) &= A''_{ij}(\mathbf{p}) \frac{p_j}{p_i} d_j(I, \mathbf{p}) + \frac{B''_{ij}(\mathbf{p})}{p_i} I \,, \end{aligned}$$

which for simplicity we can write as

$$d_i(I, \mathbf{p}) = A_{ij}(\mathbf{p}) \, d_j(I, \mathbf{p}) + B_{ij}(\mathbf{p}) \, I \,.$$

Take the partial derivative of both sides with respect to I. ▮

Proof of Proposition 2. Differentiating both sides of (3) with respect to I^h,

$$\frac{\partial^2 d_i^h(I^h, \mathbf{p})}{\partial (I^h)^2} = A_{hij} \frac{\partial^2 d_j^h(I^h, \mathbf{p})}{\partial (I^h)^2} \quad \text{for each } h \text{ and for all goods } i \neq j. \tag{9}$$

This proves the proposition's second sentence.

A function is linear if and only if its second derivative is zero.[12] Therefore, wherever d_j^h is linear, the right-hand side of (9) is zero, so the left-hand side of (9) must be zero as well, making d_i^h is linear. ▮

[11] Here v is not an indirect utility function; Muellbauer (p. 530) interprets it as being the value share for an arbitrary (say, the first) commodity. Muellbauer uses an upper-case V to denote indirect utility functions, as in his equation (12).

[12] A mathematician would use "affine" wherever we use "linear."

Proof of Corollary. A constant function is linear; use Proposition 2. ∎

Symmetry, re-expressed. Deaton and Muellbauer (1980a p. 45) state our equation (4) as their equation ([2.]4.6), but on p. 76, their equation ([3.]4.11) expresses symmetry in a completely different way, without any explanation, and the latter expression is the one used by most later authors. Ulrick (2014 p. 132) has an explanation; here is a different one, which is more elementary in some respects. We use Deaton and Muellbauer's notation for income and amount consumed instead of ours.

Proposition 3. *If w_i is the value share of consumption of the ith commodity, p_i is the price of the ith commodity, and x is income, the AIDS demand system is*

$$w_i = \alpha_i + \sum_j \gamma_{ij} \log p_j + \beta_i \log\{x/P\}$$

where by definition P satisfies

$$\log P = \alpha_0 + \sum_k \alpha_k \log p_k + \frac{1}{2} \sum_j \sum_k \gamma_{kj} \log p_k \log p_j \,.$$

In the AIDS demand system, the condition that the Slutsky Substitution Matrix be symmetric is equivalent to the condition that $\gamma_{ij} = \gamma_{ji}$.

Proof. This description of the AIDS demand system comes from Deaton and Muellbauer's (1980b) equations (8) and (9). Their equation (6) can be written as

$$\frac{p_i q_i}{x} = \alpha_i + \sum_j \gamma_{ij} \log p_j + \beta_i u \beta_0 \prod_k p_k^{\beta_k}$$

where q_i is quantity demanded and x is income (or expenditures). Therefore

$$q_i = \frac{x}{p_i}\alpha_i + \frac{x}{p_i} \sum_j \gamma_{ij} \log p_j + \frac{x}{p_i}\beta_i u \beta_0 \prod_k p_k^{\beta_k} \,.$$

For example, if $n = 3$ then

$$q_i = \frac{x}{p_i}\alpha_i + \frac{x}{p_i}(\gamma_{i1} \log p_1 + \gamma_{i2} \log p_2 + \gamma_{i3} \log p_3) + \frac{x}{p_i}\beta_i u \beta_0 p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3}$$

and

$$q_j = \frac{x}{p_j}\alpha_j + \frac{x}{p_j}(\gamma_{j1} \log p_1 + \gamma_{j2} \log p_2 + \gamma_{j3} \log p_3) + \frac{x}{p_j}\beta_j u \beta_0 p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3} \,.$$

Note that since

$$\frac{\partial}{\partial p_1} p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3} = \beta_1 p_1^{\beta_1 - 1} p_2^{\beta_2} p_3^{\beta_3} = \frac{\beta_1}{p_1} p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3}$$

we can write by analogy

$$\frac{\partial}{\partial p_i} p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3} = \frac{\beta_i}{p_i} p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3} ,$$

which makes it easy to calculate the following derivatives for $i \neq j$:

$$\begin{aligned}
\frac{\partial q_i}{\partial p_j} &= \frac{x}{p_i}\gamma_{ij}\frac{1}{p_j} + \frac{x}{p_i}\beta_i u \beta_0 \frac{\beta_j}{p_j} p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3} \\
&= \frac{x\gamma_{ij}}{p_i p_j} + \frac{x\beta_i\beta_j\beta_0 u}{p_i p_j} p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3} \qquad (10) \\
\frac{\partial q_j}{\partial p_i} &= \frac{x}{p_j}\gamma_{ji}\frac{1}{p_i} + \frac{x}{p_j}\beta_j u \beta_0 \frac{\beta_i}{p_i} p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3} \\
&= \frac{x\gamma_{ji}}{p_i p_j} + \frac{x\beta_i\beta_j\beta_0 u}{p_i p_j} p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3} . \qquad (11)
\end{aligned}$$

The symmetry condition on the Slutsky Substitution Matrix is that the left-hand side of (10) be equal to the left-hand side of (11) (since these q's are Hicksian (compensated) demand curves, not Marshallian demand curves). Clearly this holds if and only if $\gamma_{ij} = \gamma_{ji}$.

To extend this to Marshallian demand curves, in Deaton and Muellbauer's equation (4), replace the expenditure function $c(u, p)$ with income (or expenditure) x and rearrange to obtain, for $n = 3$,

$$\beta_0 u p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3} = \log x - \left(\alpha_0 + \sum_{k=1}^{3} \alpha_k \log p_k + \frac{1}{2}\sum_{k=1}^{3}\sum_{j=1}^{3} \gamma_{ij}^* \log p_k \log p_j\right) . \qquad (12)$$

Replacing the $\beta_0 u p_1^{\beta_1} p_2^{\beta_2} p_3^{\beta_3}$ portions of (10) and (11) with the right-hand side of (12) drops u and introduces x, thus turning the expressions into their Marshallian form.[13]

The extension to other values of n is straightforward. ▮

As Ulrick (2014 p. 132) also notes, in the first sentence of Proposition 3, the γ_{ij} constants are defined by Deaton and Muellbauer (1980b, Equation (7)) as $\gamma_{ij} = \frac{1}{2}(\gamma_{ij}^* + \gamma_{ji}^*)$, "hence, $\gamma_{ij} = \gamma_{ji}$ and symmetry holds. (When estimating the AIDS model, it is possible to obtain estimated parameters such that $\gamma_{ij} \neq \gamma_{ji} \ldots$.)"

References

Baker, Jonathan B., and Daniel L. Rubinfeld (1999), "Empirical Methods in Antitrust Litigation: Review and Critique," *American Law and Economics Review* 1/1: 386–435.

[13] Usually, Marshallian demand curves do not have the property that $\partial q_i/\partial p_j = \partial q_j/\partial p_i$, but this proof suggests that in the AIDS model, the Marshallian demand curves do have that property.

Banks, James, Richard Blundell, and Arthur Lewbel (1997), “Quadratic Engel Curves And Consumer Demand,” *The Review of Economics and Statistics* 79/4: 527–539.

Barnett, William A., and Apostolos Serletis (2008), “Consumer Preferences and Demand Systems,” *Journal of Econometrics* 147: 210–224.

Barnett, William A., and Apostolos Serletis (2009), “Measuring Consumer Preferences And Estimating Demand Systems.” Chapter 1, pages 1–35 of Daniel J. Slottje, ed., *Quantifying Consumer Preferences*. Bingley, UK: Emerald Group Publishing Limited.

Blundell, Richard, and Thomas M. Stoker (2007), “Models of Aggregate Economic Relationships that Account for Heterogeneity.” Chapter 68 of James J. Heckman and Edward E. Leamer, eds., *Handbook of Econometrics, Volume 6A*. Amsterdam: Elsevier.

Deaton, Angus, and John Muellbauer (1980a), *Economics and Consumer Behavior*. Cambridge: Cambridge University Press.

Deaton, Angus, and John Muellbauer (1980b), “An Almost Ideal Demand System,” *The American Economic Review* 70/3: 312–326.

Hausman, Jerry, Gregory Leonard, and J. Douglas Zona (1994), “Competitive Analysis with Differenciated Products,” *Annales d’Économie et de Statistique* 34: 159–180.

Humphrey, Thomas M. (1992), “Marshallian Cross Diagrams and Their Uses before Alfred Marshall: The Origins of Supply and Demand Geometry,” *Economic Review, Federal Reserve Bank of Richmond* 78/March: 3–23.

Marshall, Alfred (1920), *Principles of Economics, 8th edition*. London: Macmillan.

Muellbauer, John (1975), “Aggregation, Income Distribution and Consumer Demand,” *The Review of Economic Studies* 42/4: 525–543.

Stoker, Thomas M. (1993), “Empirical Approaches to the Problem of Aggregation Over Individuals,” *Journal of Economic Literature* 31/December: 1827–1874.

Ulrick, Shawn (2014), “A Primer on AIDS-Based Models in Antitrust Analysis,” *European Competition Journal* 10/1: 123–153.

Wahl, Thomas I., Ron C. Mittelhammer, and Hongqi Shi (1994), “Accounting for Aggregation Bias in Empirical Demand Models—The Case of Almost Ideal Demand Systems.” Proceedings of the NCR-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management. Chicago, Illinois.