Faculty of Business and Law

Department of Economics

Working Paper Series

The Upward Pricing Pressure Test and the Sensitivity of the Diversion Ratio

Lydia Cheung

2014/08

Antitrust Market Definition and the Sensitivity of the Diversion Ratio *

Working Paper

Lydia Cheung

Auckland University of Technology †

December 22, 2014

Abstract

The diversion ratio is a key ingredient to the calculation of the Upward Pricing Pressure (UPP) test, which is a new shortcut for screening mergers. It measures the degree of substitutability between the merging goods, which affects the potential for price increase post-merger. There is currently little existing research on how the diversion ratio is to be estimated (unlike its cousin, the cross-price elasticity). This paper explores one of the methods to estimate diversion ratios, which is through the estimation of a demand system. Specifically, this paper shows that the estimated value of the diversion ratio is, in fact, little affected by one of the most contentious decisions in merger analysis: the definition of the market boundary.

*I thank seminar participants at the second ATE (Applied and Theoretical Economics) Symposium in Sydney, Victoria University in Wellington, Auckland University of Technology, University of Auckland, University of Waikato, and New Zealand Association of Economists (NZAE) 2014 Annual Conference for their constructive comments on this paper. All errors are my own.

†Department of Economics, Faculty of Business, Auckland University of Technology. `lydia.p.cheung@aut.ac.nz`.

1 Introduction

The Upward Pricing Pressure (UPP) test popularized by Farrell and Shapiro (2010) signifies a new emphasis in merger analysis on price effects. While there has always been a need for easy-to-compute shortcuts for screening merger proposals, the old shortcuts that rely market definition and market shares are theoretically proven to be problematic and, in practice, are often subjective. The new UPP test is designed to overcome these flaws. In terms of theory, a focus on post-merger price change is much more appealing than a focus on market share change, because the former has a much more predictable relationship with consumer surplus than the latter. In terms of practice, UPP no longer requires a definition of the market boundary. In addition, UPP addresses a salient feature of the modern retail industry: that most consumer products are highly differentiated, and consumers are keenly aware of products' differences. In other words, two mergers with identical pre-merger market shares, but different characteristics on the two pairs of merging goods, are likely to have sufficiently different outcomes. The UPP is designed to detect these differences, where the old shortcuts fail.

The UPP is an approximation of the post-merger price change for each product. This approximation can be viewed as the change in first order condition for each product, in a Bertrand oligopoly market, from pre-merger to post-merger. Intuitively, before the two firms merge, competition between them imposes a pressure on price increase (hence "upward prising pressure") in terms of diverted profits: should one of the firms increases its price, some of its customers will divert to the competitor, increasing the competitor's profit. However, after merger, this pressure or opportunity cost on price increase disappears, because profits are no longer "diverted", but jointly maximized. In summary, the larger the pre-merger pressure on price increase, the larger the opportunity for actual post-merger price increase.

One of the key ingredients in computing the UPP is the diversion ratio, which measures the degree of substitutability between the two merging goods. Intuitively, this is a crucial input to merger analysis because a merger between two close substitutes causes a larger reduction in competition than a merger between two distant substitutes. The larger the diversion ratio, the larger the proportion of customers that would be diverted to the competitor should the firm increases its price. Thus, two products with high diversion ratios between them will have a large pre-merger pressure on price increase, and subsequently, a large opportunity for actual post-merger price increase. The actual calculation of the UPP requires few inputs: the diversion ratio, the competitor's (pre-merger) price-cost margin, and optionally, the product's own reduction in marginal cost from the merger. Farrell and Shapiro (2010) are rather open-ended on where or how to obtain the diversion ratio, giving suggestions like company internal documents ("to whom do we lose our business?") or customer surveys.

Stemming from the widespread discussion on the UPP, a number of researchers have explored different ways to estimate the diversion ratio, especially on methods that are deemed to have a stronger theoretical basis or higher replicability. There are two main approaches. The first is an econometric approach, where the diversion ratio is derived from various demand parameters. The current paper takes this approach. The second is an experimental approach, where the diversion ratio is derived from observed substitution patterns when some products are taken away, either by natural or designed experiments. Conlon and Mortimer (2013) exogenously remove snack foods from vending machines and record how customers switch. In terms of natural experiments, industries like hospitals are a natural candidate, because of its extensive records on where patients switch to when a hospital shuts down.

The econometric approach to diversion ratios comes with a slight paradox: in order to estimate a demand system (albeit a simple one), the researcher must define the market boundary on what products are "in" or "out", which is precisely the subjective decision that plagued the old shortcuts and that UPP is designed to avoid. This paper shows that the choice of market boundary actually has little empirical influence on the estimated diversion ratios. I demonstrate, with common discrete choice demand models in the logit family, that the choice of market boundary has bigger effects on derived own- and cross-price elasticities; but diversion ratios are composed of *ratios* between these elasticities, which do not vary nearly as much. In summary, a demand estimation still requires, in principle, a definition of market boundary; but in practice, different market boundaries give similar diversion ratio values.

2 Demand Model and Diversion Ratio

The simple nested logit model is one of the most commonly used empirical demand models, and the foundation to other more flexible and sophisticated demand models, such as those with random coefficients. Its functional forms are detailed in many papers (such as Berry (1994), p.252–), and are therefore not repeated here. While there are many possible nest designs to categorize products, estimation of multi-nest or multi-layer models become non-trivial very quickly; thus the simple nested model uses only one nest, which contains all the inside goods. This specification implies that when a consumer substitutes away from an inside good, he is more likely to substitute towards another one of the inside goods, rather than the outside good, because the random terms in the utilities of each inside good are correlated with each other, but not with the random term of the outside good. However, no more distinction among the inside goods is made when the model has only one single nest.

The primary advantage of the simple nested logit demand model is that all market shares and elasticities have tractable functional forms. The diversion ratio D_{12} from goods 1 to 2 can be derived explicitly, in terms of estimated model parameters:

$$D_{12} = \frac{\varepsilon_{21}}{|\varepsilon_1|} = \frac{\sigma s_{1|G} + (1-\sigma)s_1}{1 - [\sigma s_{1|G} + (1-\sigma)s_1]}, \text{ where}$$

$$\varepsilon_1 = -\alpha p_1 \left[\frac{1 - s_{1|G}}{1-\sigma} + s_{1|G} - s_1\right] \text{ and } \varepsilon_{21} = -\alpha p_1 s_{1|G}\left[1 - s_G - \frac{1}{1-\sigma}\right]$$

The diversion ratio takes the form $y = \frac{x}{1-x}$, which has asymptotes $x = 1$ and $y = -1$. When we include consumer switching that comes from a change of price only (as opposed to other product characteristics), D_{12} can be expressed as a ratio between price elasticities ε_1 and ε_{21}. This definition of D_{12} will express changes in quantities in terms of percentages, not magnitudes. It is therefore bounded below by zero, but not bounded above; D_{12} is positive when $x = \sigma s_{1|G} + (1-\sigma)s_1$ is between 0 and 1. The (absolute value of the) own-price elasticity $|\varepsilon_1|$ is increasing in σ, while the cross-price elasticity ε_{21} is decreasing in σ. Thus, together, $D_{12} = \frac{\varepsilon_{21}}{|\varepsilon_1|}$ is decreasing in σ. This matches our intuition that, as utilities of inside goods become more correlated, substitution within inside goods become stronger, either in terms of cross-price elasticity or diversion ratio. Because the general function $\frac{x}{1-x}$ is continuous as x varies between $[0, 1]$, the simple nested logit model does not inherently limit the range of values that its implied diversion ratio can take, although it has few parameters. The value of x in tern depends on three variables: the nest parameter σ, which is bounded between 0 and 1, and indicates the degree of correlation between the random terms in utilities of inside goods; the overall market share of the good s_1; and the within-nest (G) market share of the good $s_{1|G}$. The term x can be computed either from observed market shares or fitted values; the former is often used for convenience, although it does not reflect change in model prediction when estimated parameters change. In most consumer goods situations, market size is defined as all prospective consumers (sometimes multiplied by number of servings over the period), and thus the outside good of “not buying” almost always has the largest market share. This means $s_1 = s_{1|G} \cdot s_G$ is usually fairly small, while $s_{1|G}$ is often at least one magnitude larger. In a lot of consumer goods estimations, σ is found to have a moderate magnitude around 0.5; the value of x will then be a rough average between s_1 and $s_{1|G}$, which has the same magnitude as s_1. Thus, although the model does not constrain the values of x, it more often concentrates on the lower end of the range $[0, 1]$, which means that D_{12} will also take a small value in the magnitude of s_1. This is often perceived as "unrealistically small" diversion between substitutes.

The above simple nested logit model is easy to estimate, but researchers often want even richer substitution patterns. The usual next step in further enriching the model is by random coefficients: each demand parameter (representing consumer taste for a product attribute or product group) is not limited being fitted a single value, but instead is given a distribution of values, whose distribution parameters are estimated. This distribution can be discrete or continuous: discrete consumer types are often used for airline products (for leisure vs. business travellers), while continuous taste distribution has been used for many consumer goods, from breakfast cereal to automobiles. Conceptually, a continuous distribution can be thought of as the "limit" to a discrete distribution with very many consumer types, thus in the following discussion I will illustrate with the simpler discrete distribution. In the discrete random coefficient model, each consumer type has its own set of taste parameters, together with its proportion in the population, which is also estimated. The model produces a diversion ratio for each discrete consumer type, and the overall diversion ratio is the weighted average of these individual ratios. Each individual diversion ratio takes the same $\frac{x}{1-x}$ functional form as that of the simple nested logit above, with the same range of possible values. However, the benefit of the random coefficients is that the mixing of diversion ratios gives the overall ratio a different distribution, because the functional form of the diversion ratio is not linear in taste parameters. A future version of this paper will include empirical results with random coefficients.

3 Data and Results

I use a supermarket scanner panel dataset to estimate demand in a complete category of consumer products, and then derive the implied diversion ratios from the estimated parameters. The data used is the IRI Academic dataset, which spans 11 years (from 2001 to 2011). It has weekly prices and sales in 30 product categories, with complete UPC (Universal Product Code)

Table 1: Sugar Substitute Ingredient and Percentage in Observations

Saccharin & Dextrose	20%
Nutra sweet	18%
Aspartame	18%
Saccharin	16%
Sucralose	15%
Fructose	6%
Glucose	3%

coverage under each category. This is especially important for the analysis of product competition and consumer substitution patterns. In terms of types of stores, this data mainly consists of supermarkets, but also includes convenience stores and warehouses. Stores that belong to the same chain are identified. In terms of geography, these stores provide good coverage on 47 major metropolitan areas in the U.S. In addition, there are individual panel shopping data, covering the same product categories, in two small cities, and some limited advertising data.

Supermarket consumer products are a good example of subjective market boundaries due to product proliferation. In this supermarket data, even the smallest product category has more than 25 distinct UPC's; the large categories contain hundreds. Within each category, it is often possible to further divide UPC's into sub-categories. This creates many possible market definitions: should the antitrust market be the entire product category, or just selected sub-categories? Or even particular combinations of individual UPC's? For my preliminary study, I have chosen a medium-sized product category with distinct sub-categories: sugar substitutes. In 2006 data, there are a total of 150 unique UPC's, with 1433 supermarkets carrying at least 10 unique UPC's. The average store carries 20 UPC's (with a minimum of 4 and a maximum of 37). One logical way to sub-categorize these products is via ingredients, shown in table 1. Another possibility is via kinds of packaging: e.g. packets, tablets, or pouring containers.

To test how the estimated diversion ratios vary with market definition, I will use the entire category as the baseline market definition, and experiment with dropping each of the sub-categories in table 1 from the market. In addition, I also experiment with dropping up to eight large UPC's from the baseline market. To compare diversion ratios across the different market definitions, I identify the two largest selling UPC's in the product category ("Sweet'n Low 4.5 oz" and "Equal 3.5 oz", accounting for 3.86% and 3.75% of the observations, respectively), which are always kept in the market. Of course, any other product pair could be used; but a pair of less popular products mean fewer markets contain both of them, and a merger between them is less likely to be problematic, and therefore their diversion ratios are less relevant. I then estimate the simple nested logit demand model in each market and compute the diversion ratios between these two products. By comparing outcomes in these different market definitions, I show that estimated demand elasticities vary more widely than estimated diversion ratios.

Changing the list of products in the market does not only change the particular observations used in the demand estimation, but it changes a few of the derived variables that are used in the demand estimation as well. Total market size M, which is a multiple of population and approximate average consumption, is set to remain constant throughout. This means that as the number of products is changed, the market share of the outside good s_0 and shares within the nest $s_{j|N}$ will change too. Lastly, the derived "BLP instruments" (from Berry, Levinsohn, and Pakes (1995)) will also change, since the instrument for each product is based on product characteristics of all its rivals.

Table 2 shows the estimated own-price elasticities and diversion ratios of the two largest UPC's in the baseline market, across 1381 supermarkets. The elasticities all have the correct sign, with very few cases in the inelastic range. Their values are plausible and in line with estimates from many other studies on consumer goods demand. The values for the derived diversion ratios, however, might appear at first glance to be quite small. These diversion ratios are ratios of elasticities and thus measure quantity changes in percentages, not absolute magnitudes. For example, a diversion ratio from Sweet'n Low to Equal of 0.017 implies that, when Sweet'n Low loses 1% of its sales (due to an increase in price), Equal will see a 0.017% increase in its sales. Alternatively, one can translate percentage changes into magnitudes. If we consider the mean total annual sales of these two products across supermarkets in 2006 (449 packs of Sweet'n Low and 319 packs of Equal), the above diversion ratio implies that, when Sweet'n Low loses sales of 4.49 units, Equal will see a 0.054 unit increase in sales. Equal's average diversion ratio to Sweet'n Low is slightly smaller at 0.0087. Unlike demand elasticities, there is virtually no existing research on what the "typical values" of diversion ratios between consumer goods are. One could rationalize their seemingly small values with the fact that the average supermarket carries 20 UPC's in this category, plus the overwhelming substitution towards the outside good ("not buying"), which almost always has the largest market share. It should also be reminded that the diversion ratio includes only consumer switching due to a change in price, not switching observed. Consumers have any number of other reasons to switch products: product attributes,

Table 2: Estimation Results from Baseline Market

UPC	Variable	Mean	Min	Max
Sweet'n Low 4.5 oz.	elasticity	−1.27	−1.89	−0.91
	diversion ratio	0.017	0.00028	0.058
Equal 3.5 oz.	elasticity	−2.96	−4.14	−1.48
	diversion ratio	0.0087	0.000039	0.035

unobserved characteristic, or sheer randomness. These other factors should not be included in the calculation of the UPP, either because they are present both before and after the merger (e.g. random term in utility), or that they are not substitutions that can be induced post-merger by a change of price alone (e.g. product attributes).

Figure 1 shows the distributions on own-price elasticities and diversion ratios in each set of experimental market boundaries. In the diagram on the left, the first set of box plots on top is the baseline market, while each set below is a market resulting from dropping one of the ingredient types, starting from the largest type ("Saccharin & Dextrose"). In the diagram on the right, each set of box plot is a market resulting from dropping an increasing number of individual UPC's, from dropping the largest UPC to dropping the largest eight UPC's. Within each set of box plot are the distributions on two own-price elasticities and two diversion ratios, of Sweet'n Low and Equal. All of these experiments demonstrate that the elasticities have a much wider range than the diversion ratios. Of course, if one wants to test whether the mean of an elasticity of diversion ratio is statistically different between experiments, tests such as ANOVA will take into account the different dispersions in elasticities and diversion ratios. But in terms of the effect on a calculated UPP value, the narrow ranges of the diversion ratios show that the difference between various markets is inconsequential. In other words, excluding various product groups or UPC's in the market definition gives little magnitude difference to the derived diversion ratios.

Now I zoom into the distribution of diversion ratios and how it changes between experiments. Figure 2 shows the kernel densities of diversion ratios in each set of experimental market from the second graph of figure 1, where the largest eight UPC's in the category are dropped one by one. The first graph in figure 2 is diversion ratios from Sweet'n Low to Equal, while the second graph is diversion ratios from Equal to Sweet'n Low. The mass in each set of density is coherent with table 2, where Sweet'n Low has a larger diversion ratio. Both graphs show that, as more UPC's are dropped, the density of the diversion ratio shifts to the right, with an increase in

Figure 1: Box Plots of Elasticities and Diversion Ratios over Various Market Definitions

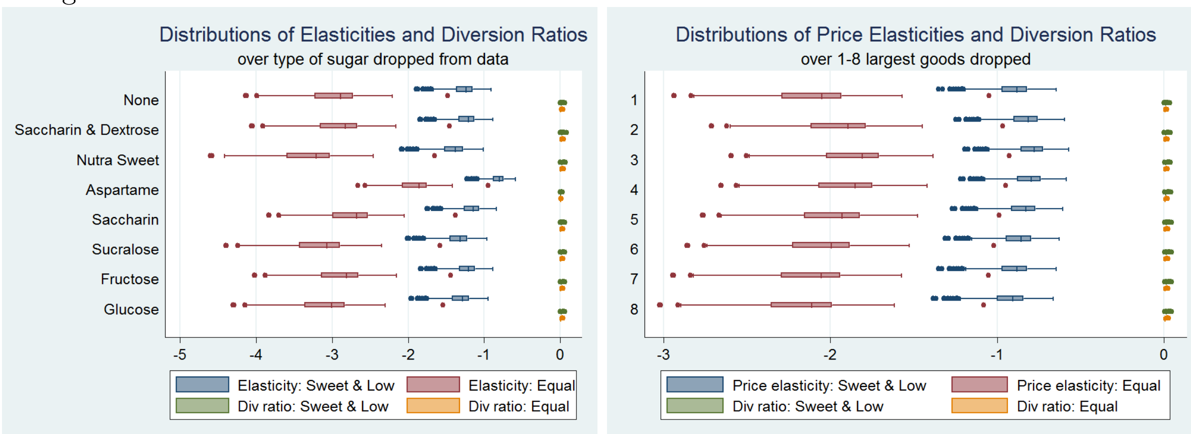

Figure 2: Kernel Densities of Diversion Ratios

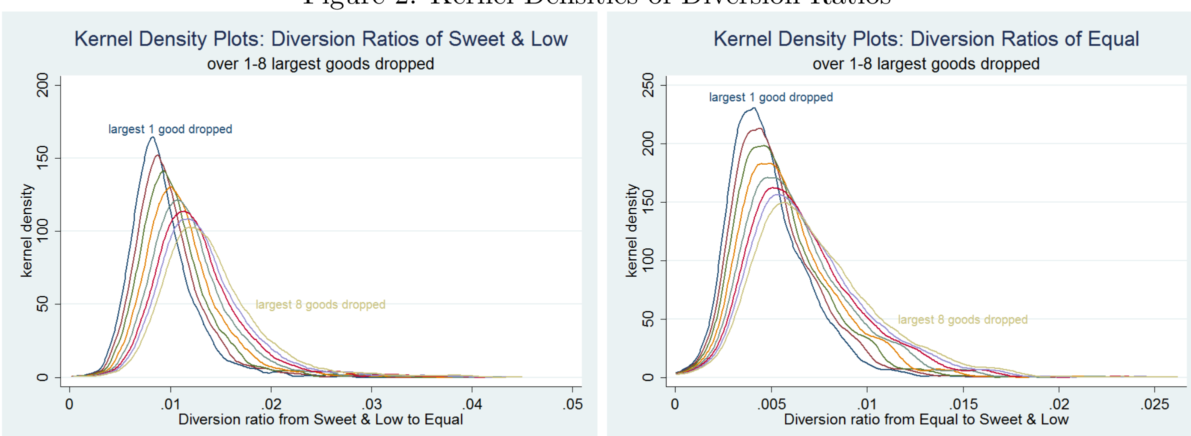

dispersion, albeit both changes are mild. These changes are likely direct results of shrinking choices in the market. On average, each inside good gets a slightly larger share of customers; but this effect is not uniform across all markets, because not all markets carried the deleted products in the first place.

4 Conclusion

The diversion ratio is a key ingredient to the calculation of the Upward Pricing Pressure (UPP) test, which is a new shortcut for screening mergers. However, there is currently little existing research on how the diversion ratio is to be estimated or approximated, unlike its cousin, the cross-price elasticity, which is much more familiar to economists. This paper explores one of the methods to estimate diversion ratios, which is through the estimation of a demand system.

Specifically, this paper shows that the estimated value of the diversion ratio is, in fact, little affected by one of the most contentious decisions in merger analysis: the definition of the market boundary.

However, more empirical research is needed to compare diversion ratios estimated from different methods. A promising alternative approach is an experimental one, where products are exogenously removed from the market and consumers' new purchase patterns are recorded. A similar, albeit less than perfect, experiment can also be done with the supermarket scanner dataset used in this paper. When a product is not observed to be purchased for a sufficiently long time in a store, it can be assumed to be withdrawn; the researcher can then compare the "before" and "after" purchase patterns. A future version of this paper will include these experiments.

References

Berry, S. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 25(2), 242–262.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841–890.

Competition Commission and Office of Fair Trading (2010): "Merger Assessment Guidelines," `http://www.competition-commission.org.uk/about_us/our_organisation/workstreams/analysis/pdf/100916_merger_assessment_guidelines.pdf`.

Conlon, C. T., and J. H. Mortimer (2013): "An Experimental Approach to Merger Evaluation," *NBER Working Paper*.

Department of Justice and Federal Trade Commission (2010): "Horizontal Merger Guidelines," `http://www.justice.gov/atr/public/guidelines/hmg-2010.html`.

Farrell, J., and C. Shapiro (2010): "Upward Pricing Pressure in Horizontal Merger Analysis: Reply to Epstein and Rubinfeld," *The BE Journal of Theoretical Economics*, 10(1), 41.