RBB Brief 47

RBB | Economics

Drink deriving: estimating substitution patterns for hard and soft drinks

1   Demand estimation provides an economic framework for analysing brand level data on prices and volumes to allow inferences to be made regarding patterns of consumer switching.

2   The AIDS framework was originally proposed in "An Almost Ideal Demand System", Angus Deaton, John Muellbauer, The American Economic Review, Vol. 70, No. 3. (June 1980), pp. 312-326. It is one of a number of models that can be employed to analyse demand econometrically. RBB Brief 39 Roll on demand estimation: the EC's empirical analysis in Unilever/ Sara Lee discusses other such frameworks that have been used by the European authorities (see http://www.rbbecon.com/ downloads/2012/11/RBB_B39_ BW.pdf).

3   The diversion ratio from product A to product B is defined as the proportion of sales lost by product A that is won by product B following an increase in the price of Product A.

4   RBB Economics advised the merging parties in both AG Barr/Britvic and Diageo/USL.

5   United Spirits Limited sells whisky (including the Whyte and Mackay blended brand as well as private labels to supermarkets) and vodka (Vladivar and Russian Standard). Diageo sells brands including Bell's blended whisky and Smirnoff vodka.

6   AG Barr sells a number of soft drinks brands in the UK including IRN-BRU, Orangina, Tizer, Rubicon and the Barr range. Britvic is the UK bottler of Pepsi, 7UP and Mountain Dew, and owner of several brands including Tango, R Whites, Robinsons and J20.

7   At the time of writing, the CMA is contemplating remedies proposed by the merging parties to address competition concerns in the blended Scotch whisky segment.

8   The CC focused in particular on analysing the constraint that Britvic's portfolio exerted on two AG Barr brands, IRN-BRU and Orangina.

The UK Office of Fair Trading (OFT) and Competition Commission (CC), since 1 April 2014 combined to form the Competition and Markets Authority (CMA), have both recently analysed high profile mergers involving branded drinks producers using demand estimation.[1] Specifically, both authorities employed the Almost Ideal Demand System (or AIDS) framework to assess quantitatively whether drinks brands supplied by the merging parties are particularly close competitors and so whether these transactions were likely to result in a significant lessening of competition (SLC).[2]

The use of the AIDS model in these cases can be seen as part of an ongoing shift towards a more quantitative approach to analysing mergers in the UK, both at Phase I and Phase II. The AIDS modelling framework is particularly seductive in this context. It appears to be simple to implement (even in Phase I) and requires data on prices and volumes that are often available for fast moving consumer goods (FMCGs). Moreover, on the face of it, the approach produces a clear-cut measure of closeness of competition: estimates of substitution patterns that can be used to derive diversion ratios between brands owned by the merging parties.[3]

However, a simple framework often fails to capture key aspects of a complex economic reality, raising the issue of how much weight can appropriately be placed on estimates of diversion between the merging parties. In this Brief, we discuss how the UK authorities employed the AIDS framework in the AG Barr/Britvic (soft drinks) and Diageo/USL (hard drinks) mergers and explore some of the controversies that arose.[4] In doing so, we highlight the importance of ensuring that econometricians not only ask the right questions, but also (and of no less importance), critically evaluate the answers obtained – particularly if primary reliance is to be placed on the results when analysing mergers.

Background – hard drinks and soft drinks

In both Diageo/USL[5] and AG Barr/Britvic[6], the UK merger authorities employed the AIDS framework to assess whether the merging parties' brands could be considered close competitors from the consumer perspective (as implied by diversion ratios). In the case of hard drinks the OFT concluded that diversion ratios derived using its AIDS model were consistent with there being an incentive to raise prices in relation to each of the parties' main whisky brands resulting in the finding of an SLC for these products, whereas no such incentive was identified in relation to vodka.[7] According to the CC's second phase assessment of the parties' carbonated soft drink (CSD) brands in the soft drinks case, the AIDS model suggested no incentive to increase price for any of the brands considered, a finding that played an apparently decisive role in the CC's decision to clear the merger unconditionally.[8] In each case the econometric analysis was presented as the most authoritative form of evidence on substitution patterns and was given prominence in the decision.

Asking the right questions

Econometrics provides an extremely powerful set of tools for asking specific questions of data. But, importantly, the way in which a question is formulated (or, in technical terms, the way a model is specified), has an impact on the answers the data give back.

The AIDS model provides a relatively flexible framework for posing questions to understand how closely merging parties' brands compete. The key steps that should be taken when implementing an AIDS model can be summarised as follows. First, choose a "product set" that contains the merging parties' brands of concern and those brands that may reasonably be viewed as substitutes (e.g. all CSDs).[9] Second, use price and volume data combined with some theoretical assumptions about consumer preferences to estimate how consumers allocate their spending between products within that set as relative prices change over time (controlling for other factors) and use those patterns of substitution to infer diversion ratio estimates. Third, consider and attempt to quantify the extent to which a price rise for brands within the product set may cause expenditure to leak to other brands outside the product set in question (known as diversion to the "outside good"). Where significant leakage may occur, diversion ratios initially calculated within the product set analysed should be down-weighted accordingly.

However, fundamental problems can arise when the specific question asked (i.e. the model's specification) either does not capture an important aspect of the industry or carries an in-built assumption as to the form the answers should take that is not justified by economic reality. Such "misspecification" can in turn mean that seemingly clean measures of closeness of competition derived from this analysis are biased – i.e. they are not expected to reflect true patterns of diversion between brands.[10] We discuss a few of the many issues that may arise using the hard and soft drinks mergers to illustrate their relevance.[11]

Omitted variable bias

Rarely is it possible to take account of all factors that affect consumers' choices between products. However, in any demand modelling exercise, it is critical to control for those factors affecting consumer demand that are closely linked with price changes. For example, in the case of the CC's model in soft drinks it was recognized by both the parties and the CC that in-store "feature and display" advertising, particularly whether a brand is placed on a "gondola end", represents an important determinant of demand that was left out of the CC's model.[12] Leaving out feature and display in this case leads to so called "omitted variable bias": since products advertised on gondola ends also tend to be discounted, demand is estimated to be more sensitive to price reductions than it would have been had the effect of in-store advertising on demand also been taken into account.

The CC recognized that in-store advertising might be an important "omitted variable" but nonetheless appeared to place significant weight on diversion ratios generated by the model.[13] However, if the stated aim of such models is to understand how a price change impacts on volumes it is hard to place much weight on a model that knowingly fails to control for a fundamental driver of volumes that usually comes into effect at the same time as the most significant price reductions (even if it is unclear whether the inferred diversion ratio is likely to be biased upwards or downwards as a result of the omitted variable).[14]

9   The choice of products to be included can have a material impact on the results, and may itself be a source of controversy.

10 Even if they do reflect true patterns of diversion, it is a separate question whether diversion measured at the consumer level appropriately captures closeness of competition at the wholesale level. For a discussion of this point see RBB Brief 39: Roll on demand estimation: the EC's empirical analysis in Unilever/Sara Lee.

11 It is not possible in this Brief to discuss all the issues surrounding AIDS estimation that need to be considered. Other issues include the distortions that can arise from consumer stocking behaviour, whether to carry out the analysis at the SKU level or brand level and, in the latter case, how to calculate the average brand price from the price of individual SKUs. In addition it is important to note that these models are static in nature – they therefore do not take account of dynamic factors (such as brand repositioning and strategic retailer behaviour) that are important when assessing the potential impact of mergers.

12 Specifically, when a product is placed on the gondola end, it is usually sold at a large discount. Many price reductions are therefore associated with large uplifts in volumes not simply because prices were low but also because advertising and product placement was enhanced.

13 The CC did state, however, that it took account of these limitations in its assessment of the model results and that it appropriately reflected them in the weight it put on the model's result in its overall view. AG Barr/Britvic Decision, paragraph 6.50.

14 The model would lead to an upward bias in own price elasticity and an upward bias in cross price elasticity implying that the bias in diversion ratios (which are a function of the ratio of these elasticities) would be unclear.

15 The CC noted that it was not able in this case to implement an alternative more complicated approach which would enable one to account for switching from and to drinks outside CSDs. See AG Barr/Britvic, Appendix C, paragraph 44.

16 AG Barr/Britvic, paragraph 43. To see this imagine that a reduction in the price of Pepsi causes a large increase in the purchases of Pepsi but no reduction in the purchases of any other brand such that total expenditure on soft drinks increases. If the price reduction increases Pepsi's share of CSD expenditure, it must by definition reduce the share of other brands. If the model does not take account of the overall increase in expenditure it may mistake the reduction in the share of the other brands and increase in the share of Pepsi for substitution between these brands, when in fact no consumer has substituted away from another CSD brand.

17 By contrast, the OFT did appear to take some account of diversion to the outside good. The OFT's Decision does not provide details of how it sought to quantify the strength of this potentially important constraint however, noting only that *"[The OFT] has also taken greater account, albeit to a limited extent, [of] the constraint from outside goods."*

18 Importantly, the CC did acknowledge the relevance of the outside good in the context of discussing survey evidence, where it stated that *"We agreed with the parties that a correction for outside goods should be made. The 'outside good' option represents the extent to which the market for the products expands and contracts in response to price changes, and this is a real effect and an important one in this market."* See AG Barr/Britvic, Annex 5, paragraph 7e).

Restrictions on the model

Practitioners may impose simplifying restrictions on demand systems in order to make their models easier to apply in practice. Such restrictions may have very significant effects on estimated diversion ratios. This makes it critical either to test if those restrictions are appropriate (i.e. ensure they are not contradicted by other evidence), or to ensure that results are not sensitive to the relaxation of these (potentially artificial) restrictions.

In hard drinks the OFT imposed a restriction known as "symmetry" on its demand system. In essence, it assumed that substitution from brand A to brand B can be inferred from substitution from brand B to brand A. Under this specification the merging parties' brands were found to be one another's closest competitors.

However, there is no reason to expect the symmetry assumption to be valid as a matter of theory. Competition authorities regularly acknowledge the existence of asymmetric constraints between brands and firms in merger assessments. The validity of a modelling assumption that substitution is symmetric in a given case should therefore be tested using the available evidence. In responding to the OFT's findings the merging parties presented empirical analysis suggesting that the OFT's symmetry assumption did not hold. They also showed that allowing for asymmetric switching patterns not only reduced the estimated diversion between the merging parties' brands but also reversed the OFT's finding that the merging parties' brands were one another's closest competitors.

Given that there is no theoretical reason to impose symmetry in the first place and that the evidence pointed to the existence of asymmetries that would result in materially different diversion ratio results, the OFT did not therefore provide a convincing reason to adopt its more restrictive model.

The "outside good"

In order to produce reliable estimates of diversion ratios between a given set of brands, correct account must be taken of the extent to which expenditure leaks out of that product set when prices rise (i.e. the extent to which customers substitute to an "outside good"). Where consideration is not given to this phenomenon (e.g. where it is assumed that customers do not switch their expenditure to other goods when prices rise) diversion ratios between products within the set analysed may be substantially overstated.

In soft drinks, in order to simplify the modelling exercise the CC's estimated model assumed that consumers do not change their total expenditure on CSDs in response to changes in prices of CSDs.[15] The CC recognised that if, contrary to this assumption, consumers significantly decrease their total CSD expenditure when prices increase (and vice versa), then its model would overstate diversion ratios amongst CSD brands.[16] [17]

It may be that the CC felt that having, in effect, estimated an upper bound for a diversion ratio and found that measure to be low, there was no need to address further the issue of the outside good with its econometric model.[18] This approach would, however, be inappropriate where such an upper bound estimate cannot be used to rule out concerns regarding the closeness of competition between brands operated by the parties.

Questioning the answers

Given the inherent risk associated with any misspecification of an econometric model (and not just the AIDS model), it is crucial to question the results that are produced using these methods.

It is often easy to spot when misspecification gives rise to highly implausible results. In particular, if higher prices are predicted to increase demand or to win sales from substitute products these counter-intuitive results will point to a misspecified model. In both soft and hard drinks, however, controversy arose when the same modelling approach gave plainly implausible results when applied to one data set but plausible results when applied to another. When the CC applied its CSD demand model to Scotland-only data the results did not provide clear evidence of misspecification, whereas the same model performed extremely poorly for England and Wales.[19] In hard drinks, when the OFT applied its model to individual supermarket level data the estimated diversion ratios were subject to extremely high levels of variation and often failed a basic economic sense test with diversion ratios over 100% being calculated for some products.

In both cases this gave rise to an important question: was the model well specified when faced with some datasets but not others? Or was the model at all times misspecified but, by chance, produced less obviously implausible results when presented with certain datasets?

Ultimately in soft drinks the CC concluded that its preferred model met basic plausibility tests and it had sufficient confidence in the estimated diversion ratios to use them in an "illustrative price rise" calculation.[20] However, the threshold applied to determine plausibility (when price goes down, demand goes up and demand for most substitutes goes down) is a relatively lax one:[21] satisfaction of these criteria is the very least we should expect from a well specified model of demand.[22] Even a misspecified model that produces estimates of diversion ratios relatively far from their "true" values may satisfy this *"it's not clearly wrong"* criterion. Where the reality checks applied are so easy to pass, this should be reflected in the confidence that is placed in the model's outputs. Authorities should be particularly cautious when seeking to draw inferences from the absolute size of estimated diversion ratios in this context.

Conclusion

The soft and hard drinks cases underline that the measurement of switching patterns using econometric models is now an established feature of the UK (and other major) competition authorities. Such techniques may offer a valuable way to shed light on how customers respond to price changes. However, where it is clear that a model omits features that are fundamental to understanding how demand works in a given industry, this should ideally be addressed or, if data limitations prevent this, the evidential value of the model down-weighted accordingly. It is not appropriate to presume that a knowingly misspecified model's output is meaningful simply on the grounds that the direction and magnitude of the bias is unknown.

Particular care should be taken if (as in the soft and hard drinks cases) demand estimation is to be presented as foremost amongst the various forms of evidence on the magnitude of diversion between the merging parties. Both the parties and competition authorities should by all means strive to produce econometric models that provide insights into closeness of competition; but they should at the same time recognise that reliance on a misspecified demand model can be worse than no model at all.

---

19 The fact that IRN-BRU is consumed and marketed more in Scotland may in theory give rise to different substitution patterns in different regions, but it would not explain why the England and Wales model found diversion ratios not only above 100% but to quite implausible substitutes in the working paper versions of the model shared with the parties. Such extreme results cast doubt on the model's specification.

20 The parties submitted a version of the CC's model based on an alternative dataset that covered convenience stores. The parties' claim was not that the submitted model was better than the CC's but that all models should be treated with caution as they appeared to give rise to quite different looking diversion ratios. The CC nonetheless subsequently used the diversion ratios predicted by the parties' model in an illustrative price rise calculation which suggested no material incentive to increase price.

21 The CC did also carry out a basic sense check of whether the model was identifying Coke and Pepsi as close substitutes.

22 This is particularly problematic in the case of the AIDS framework because the specific question that is being posed implies that when there is very little information in the data the estimates of the own price elasticity will be computed as -1 (rather than zero), making it appear as if the data are revealing something intuitively plausible about the true nature of demand. Put differently, the absence of any meaningful information in the data can give rise to an elasticity estimate that is subsequently used to validate the model, even though the estimate arises purely from assumptions in the setup of the model.