

Roll on demand estimation: the EC's empirical analysis in Unilever/Sara Lee

1. Some notable prior cases in which simulation approaches were employed include VOLVO/SCANIA (M.1672, March 2000), FRIESLAND/CAMPINA (M.5046, December 2008) and KRAFT FOODS/CADBURY (M.5644, January 2010).

2. Case No COMP/M.5658 – UNILEVER/SARA LEE Article 8 (2) Date: 17/11/2010. In January 2012, the Commission finally published the non-confidential version of its decision. The Technical Annex to the Decision helpfully contains an account of the modelling approach adopted by the Commission.

3. The assumption that firms engage in Bertrand competition, while not discussed in detail in this Brief, is critical and may not be appropriate to all markets.

4. The term "market share" in this context refers not to the share of supply of a properly defined market, but the share of supply within each nest.

The European Commission's decision on Unilever/Sara Lee represents an important step in the use of merger simulations in assessing mergers, placing greater prominence on such analysis than in previous cases where this approach has been used.¹ Estimating price effects from a merger sounds like a panacea for merger control: in most merger cases, this is the central question that merger analysis seeks to answer. However, as we discuss in this Brief, several factors limit the ability of such modelling approaches to provide precise estimates of actual likely price effects, while the bias associated with any simulated price effects will depend on case specific factors. As a result, it is crucial that merger simulation results are not viewed in isolation from other forms of evidence, and that the thresholds for intervention as applied in Unilever/Sara Lee do not become an automatic benchmark for future mergers where these techniques are applied.

Unilever/Sara Lee – case background

The Unilever/Sara Lee merger concerned the sale of branded deodorants in a range of EU countries, with Unilever brands including Axe (Lynx in the UK), Rexona (Sure in the UK), Dove, Vaseline and Impulse, while Sara Lee operated primarily under the Sanex brand.² In determining whether the merger of the parties' branded products was sufficient to justify an SIEC finding, it was necessary to consider the degree of competitive constraint between male deodorants (where the overlap between the parties was limited but Unilever had a strong presence) and non-male (female and unisex) deodorants, where the overlap was more substantial but the merging parties faced a number of alternative competitors (including private label brands).

In this setting, as in many markets where product differentiation plays an important role, market share considerations alone will fail to provide a definitive framework for assessing competitive effects, as they do not capture adequately the closeness of competition between rival brands, products and formats. To address these issues, the Commission chose a merger simulation approach based on a Bertrand-Nash model of competition, relying for its key inputs on a nested logit model of demand estimation.³

Nested logit demand estimation – back to market shares?

Demand estimation is no more than a structured approach for assessing patterns of customer substitution across different products in response to their relative changes in price. As with all economic models, in order to make the analysis tractable, the "nested logit" model applied in Unilever/Sara Lee makes a number of simplifying assumptions. A crucial feature of such "nested logit" models is their underlying reliance on market shares.

In particular, a nested logit model requires that products of interest are grouped together into "nests" on the basis of particular product characteristics. The main role of the model is to estimate the degree of customer substitution between one nest and another, while within each nest, it is assumed that switching between individual brands is proportional to brand market share.⁴

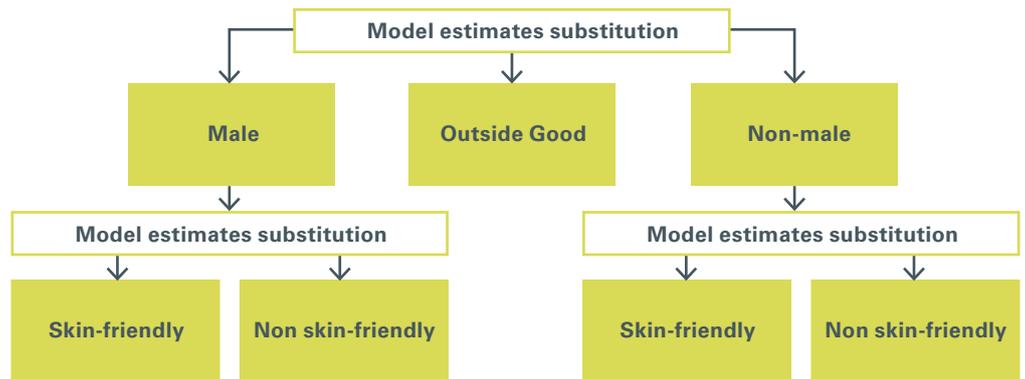
5. The model also included the potential for consumers to switch to an “outside good” – which includes the choice not to buy deodorants at all. For the purposes of this discussion, the outside good option can be ignored.

6. This feature – which follows from the independence of irrelevant alternatives (IIA), itself a corollary of the assumptions underlying the logit demand model – vastly simplifies the modelling requirements of demand estimation, limiting the number of parameters that the model has to estimate. A popular alternative demand estimation framework (the almost ideal demand system – AIDS) relaxes this assumption substantially, by estimating separate parameters for any two product pairs.

For example, in the “one-level” nest structure first considered by the Commission in Unilever/Sara Lee, all deodorants were grouped as being either male or non-male. The model then estimated only the degree of substitution between male and non-male deodorants, while it simply assumed that switching between different brands of male deodorants took place in proportion to each brand’s market share.⁵ In addition, if Unilever was the largest brand in the male segment, then customer switching from non-male to male brands was assumed to be captured to the greatest extent by Unilever, and less so by smaller competitors.⁶

As such, this model had no ability to shed light on the closeness of competition between the parties’ products within male or non-male deodorants respectively over and above what may be indicated by a simple assessment of market shares. Recognising this critical limitation of the “one-level” nest structure, the Commission preferred an alternative “two-level” nest structure (as shown in Figure 1). This more complex model subdivides male and non-male deodorants according to a further product characteristic, namely whether or not the deodorants are branded as being “skin friendly”. It then estimates both the degree of customer substitution between skin friendly and non-skin friendly deodorants (within each of the male and non-male segments), as well as the extent of substitution between male and non-male deodorants.

Figure 1: Two-level nested logit



By adding a second level of characteristics, the “two-level” model reduces the reliance on market shares. The assumption that consumers switch between products in proportion to market shares alone is limited to a narrower product set (such as male skin-friendly deodorants rather than all male deodorants), while the model allows for lower levels of substitution between skin friendly and non-skin friendly deodorants than may be the case if substitution is assumed to take place on the basis of market shares alone.

However, this two-level nesting structure by no means completely overcomes the limitations inherent in a nested logit modelling framework. Even in this more complex structure, brand “share” remains an important determinant of diversion ratios between brands: all else being equal, customers are assumed to switch to a greater degree to larger brands than they are to smaller ones. Of course, in some markets, an assumption that customers are more likely to switch to larger brands may provide a reasonable proxy of actual consumer behaviour. Indeed, in Unilever/Sara Lee, the Commission cited some survey evidence supporting such findings. But there remains a risk that this assumption may misrepresent actual competitive constraints.

7. These results are based on the Commission's preferred two-level model. The Decision relies on average price effects across each of the relevant markets defined as male and non-male deodorants respectively. Interestingly, predicted price increases for individual brands could be significantly higher. An assessment of market-level predicted price effects contrasts with UPP approaches adopted by other authorities, where the focus has been on the price rise of the individual merging brands.

Further, the two-level model still fails to take into account many other potentially important aspects of differentiation across brands, which in the case of deodorants may be fragrance, format (roll on versus aerosol) or efficacy. To the extent that these characteristics substantially affect customers' purchasing decisions, the model will necessarily be limited in estimating actual customer switching between different products, and hence the true effects of the merger. Judgments as to which competitive parameters should be taken into account in the econometric modelling should be guided by other available evidence, as well as by sensitivity tests (for example, by testing a number of alternative models). On the face of it, it seems surprising that the split between roll-on and aerosol deodorants, where the Decision explicitly stated that certain evidence indicated that "different formats are not suitable for many end-customers", was not incorporated in the Commission's modelling framework.

Ultimately, the extent to which the restrictions of the modelling approach compromise the results will depend on the facts of the merger at hand, and on the questions that the model is used to address. In particular, the model as used by the Commission may perform well in addressing the degree of substitution between male and non-male deodorants, but it would

be substantially less informative in analysing the closeness of competition between brands of skin-friendly non-male deodorants because key competitive interactions between brands in this category are entirely assumed by the modelling structure. Since merger simulation relies on estimates of substitution across all brands, the predictions of the merger simulation are only as good as all of the modelling assumptions that underpin the demand estimation exercise.

Relying on the results: what threshold for intervention?

Although the Commission has not explicitly stated its threshold for intervention on the basis of the price effects predicted by its merger simulation (and was at pains to point out that the findings of merger simulation must be put in the context of other qualitative evidence), an SIEC was found to arise – and a structural divestment was required - in all markets where the estimated price increases exceeded 2% (while the highest predicted price increase where the Commission did not find an SIEC was 1.2%).⁷ This critical debate on defining a threshold at which a simulated merger effect is deemed to trigger an SIEC needs to be informed by a number of factors.

First, the predicted price increases from a merger simulation are only estimates, and the level of confidence over these predictions will inter alia depend on the reliability of the inputs, and in particular on the results of the demand estimation analysis. If patterns of customer switching between products can only be measured imprecisely, the confidence interval around a predicted price increase may be very wide. While the Commission in Unilever/Sara Lee found its predicted price effects to be statistically significant, it is entirely plausible that in other cases, a predicted price increase of 2% may not be statistically different from zero. Without fully considering the statistical significance of the estimates obtained, it is not appropriate to make a judgment on the evidentiary weight that should be given to the predicted price change.

Second, merger simulations generally will, absent an assumption on efficiencies, predict that a merger of substitute products will lead to an increase in price, even if the competitive constraint between them is limited. As a result, there is a very real risk that applying a low threshold for intervention on the basis of merger simulation could lead to a lower SIEC threshold than is typically applied on cases where these techniques are not used.

8. This point is acknowledged by the Commission in the Technical Annex, para 181.

Finally, merger simulation techniques are inherently imperfect in capturing all aspects of competition, and these technical limitations introduce the scope for systematic divergence between actual and predicted effects. Therefore it is vital that the “thresholds” for intervention are informed by an assessment of whether such analyses may be expected to over or understate the true price effects of a merger.

Merger simulation and the “predicted” price increases – the role of dynamic effects

A key technical limitation of merger simulation tests as applied by the Commission in Unilever/Sara Lee is that they are by their nature static, and do not take into account dynamic aspects of competition, such as innovation, product repositioning, supply side substitution or entry. By disregarding these dynamic considerations, static merger simulations tell only a part of the story.

In many cases, not taking into account the potential for these dynamic supply side responses – such as the scope for new entry – will cause modelling predictions to overstate any price increases from the merger. However, in other cases, the competitive effects of the merger may be more severe than indicated by a merger simulation that does not capture this dynamic of competition: for example, if the merging firms competed particularly strongly in developing product innovation. In its Unilever/Sara Lee Decision, the Commission usefully acknowledged this deficiency. But it provides little comment as to how this omission may have biased the results of its merger simulation in this particular case, or as to what approaches might be taken to fill the gap. As a result, this Decision provides few policy indications as to how these dynamic considerations should be taken into account when determining the threshold for intervention.

Retail versus wholesale level analysis

A further important limitation stems from reliance on retail-level data to consider the impact of a merger between branded goods manufacturers who operate and compete at the wholesale level. The models estimate the predicted price increases at the retail level as if manufacturers set the retail product price and sell directly to the end consumer. Since this assumption does not reflect the true relationship between manufacturers and retailers, it introduces a further risk that simulation results will be biased. While the Commission openly recognised this problem in Unilever/Sara Lee, its approach was simply to state that the direction of any bias can never be known with confidence – and, it appears, to treat its results as if such a bias were not present.

It is indeed the case that a large number of factors may affect the relationship between the predicted retail price effects and the price effects that may materialise at the wholesale level – including factors that determine how retailers pass on wholesale price increases to end consumers, and the relative bargaining strengths of manufacturers and retailers in the negotiation of wholesale prices. Many of these factors cannot be fully assessed in the scope of a merger enquiry. However, it is also well recognised in economic theory that, where retailers are able to threaten to significantly reduce a brand’s sales (for example by partly or fully delisting it or by reducing its visibility on the shelf) manufacturers may face further disciplining effects on their ability to raise price. In these circumstances – and if competition between the merging parties is not the main source of such retailer threats being credible – the assessment of a merger simulation exercise based on retail estimates may be biased towards an overstatement of price increases.⁸ At the very least, this intuition should be tested against the facts of the case.

Key lessons for future cases

Demand estimation and merger simulation can be valuable techniques for evaluating mergers in differentiated product markets. As with other approaches that seek to capture the “upward price pressure (UPP)” of a merger, they provide a systematic framework for processing data in differentiated product markets. However, the value that they add is crucially dependent on the ability of the modelling framework chosen to take into account the key parameters of competition, and to answer the crucial questions that arise in the context of a given case.

For these reasons, merger simulation models face some inherent limitations in their ability to predict accurately the price effects of a merger – particularly in industries where dynamic effects are particularly important, or when they are used to assess wholesale level mergers on the basis of retail data. Paradoxically, the ability of these techniques to improve merger decisions can be realised only if a clear perspective is taken on these inherent limitations.

The Decision on Unilever/Sara Lee makes clear that merger simulation analysis must be informed by, and placed in the context of, other forms of evidence. Unless or until some of these broader questions are addressed, however, relying on the implicit 2% price effect threshold that was adopted in this case as a de facto benchmark for determining intervention against a merger would fail to make best use of simulation techniques and risks deterring a number of pro-competitive, or competitively neutral, transactions.