The Economic Journal, **117** (*January*), 45–67. Published by Blackwell Publishing, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

RETAIL MERGERS, BUYER POWER AND PRODUCT VARIETY*

Roman Inderst and Greg Shaffer

This article analyses the impact of retail mergers on product variety. We show that, following a merger, a retailer may want to enhance its buyer power by committing to a 'single-sourcing' purchasing strategy. Anticipating further concentration in the retail industry, suppliers will strategically choose to produce less differentiated products, which further reduces product variety. If negotiations are efficient, the overall loss in product variety may reduce consumer surplus and total welfare. With linear tariffs, however, there may be a countervailing effect as the more powerful retailer passes on lower prices to final consumers.

In many OECD countries, retail markets have become increasingly concentrated.[1] Particularly in Europe, the consolidation process does not stop at national borders but involves an increasing number of cross-border mergers. As reported in Dobson (2002), the top ten retailers in the EU account now for more than 30% of all sales of food and dairy products.[2]

As consumers typically choose only among at most a handful of outlets in their neighbourhood, cross-border retail mergers are less likely than other types of retail mergers to have horizontal effects.[3] Over the last years, however, such mergers have been scrutinised by antitrust authorities, who have become increasingly concerned about retailers' growing buyer power *vis-à-vis* suppliers.[4] It is possibly in the UK where competition authorities have started to look most seriously into retailer buyer power (as documented by the UK's Competition Commission's study on grocery retailers in 2000), and, as a result of this, the Code of Practice, which is supposed to govern the relationship of the UK's top five retailers with their suppliers, was formulated.

This article presents a theory to explain why retail mergers may increase buyer power and why they may lead to a socially inefficient reduction in product variety.

* We thank two referees and the associate editor, Ben Polak, for helpful comments. We also thank Paul Dobson, Chiara Fumagalli, Hans Normann, and seminar participants at the Office of Fair Trading, the UK's Competition Commission, the NEI Annual Conference in Lancaster (2004), the London School of Economics, Norwich University and the 6th Annual CEPR Conference on Applied Industrial Organisation in Munich (2005). Inderst thanks the ESRC for supporting this research under its grant 'Buyer Power in Retailing', and Shaffer thanks the ESRC and SSRC for supporting this research under its grant 'Buyer Power in Merger Control'.

[1] See, for instance, the OECD (1999) report on the buying power of multiproduct retailers as well as various reports on concentration in the US and European food distribution sectors, including Dobson Consulting (1999), Competition Commission (2000), Federal Trade Commission (2001), Clarke *et al.* (2002) and Dobson (2002).

[2] Amongst the retailers that are now increasingly active across the EU are Germany's Rewe and Metro, the UK's Tesco and France's Intermarché. Also, Wal-Mart operates now in several European countries after a string of acquisitions, including that of Asda (UK) and Wertkauf (Germany).

[3] The key concern from a horizontal standpoint is whether the merging firms compete in the same markets. If they do, structural remedies can sometimes be applied by prescribing the divestiture of the concerned outlets.

[4] Some of the major policy issues are discussed in Dobson and Waterson (1999) and Rey (2003). One issue is whether the retailers' growing buyer power has negative consequences for product quality, innovation, and variety.

We argue that the consolidated retailer may find it profitable to no longer carry the products of all previous suppliers. By delisting some suppliers, the retailer can make suppliers compete more aggressively for its patronage. The drawback is that, by delisting suppliers whose products provide a better fit to local preferences in some outlets, total industry profits are reduced. The trade-off for the retailer, then, is whether to adopt a single-sourcing policy and capture a larger share of lower industry profits or stay the course and capture a smaller share of higher industry profits. The former is sometimes more profitable. Moreover, we show that the loss of product variety due to single sourcing may be further aggravated as suppliers, in anticipation of further consolidation among their buyers, optimally (re-)position their products and, thereby, reduce product differentiation. Although this makes suppliers better positioned to serve all outlets of a consolidated retailer, the overall reduction in product variety reduces industry profits and, under standard assumptions, leads to lower consumer surplus and welfare when contract negotiations are efficient.

An important counter effect arises if retailers and suppliers negotiate over linear tariffs, i.e., if negotiations are not efficient. Increased competition for the consolidated retailer's account and less product differentiation tend to reduce purchase prices. As some of these savings are passed on to consumers, this reduces double-marginalisation and increases consumer surplus.

According to our theory, a consolidated retailer can obtain better deals from suppliers not only because it threatens to no longer carry their products but because it actually *does* delist some of the previously stocked goods. This has immediate, and potentially adverse, welfare implications, which sets our article apart from most of the extant literature on buyer power, where delisting goods is only an off-equilibrium threat. (The literature is reviewed below.)

It is essential in our model for the exercise of buyer power that the various outlets of the consolidated retailer previously did not all stock the same goods, e.g., due to regional or national differences in consumers' preferences and habits. Consequently, our theory applies particularly to cross-border retail mergers, where existing theories of buyer power have had little to say. In fact, our theory suggests that with regard to buyer power, antitrust authorities may have to be more concerned with mergers that involve *non-overlapping* markets than with mergers where firms' markets overlap considerably. This stands in conflict to standard merger theory which focuses almost exclusively on horizontal pricing effects. There, a merger has – broadly speaking – more serious welfare implications the greater the degree of overlap between the merging parties.

The predictions that consolidated retailers can both obtain more favourable terms of supply and reduce their supplier base seem to accord well with casual observations. The UK's Competition Commission refers in their 2003 report, for instance, to Asda's benefits from the global procurement strategy of Wal-Mart. Data collected for this report and an earlier report (Competition Commission, 2000) also document that the further consolidation of the UK grocery retail industry may have weakened suppliers' negotiating power and led to higher concentration in the retailers' base of suppliers. Perhaps it is not surprising, therefore, that preventing consolidated retailers from delisting (in particular, small and dependent suppliers) seems to be a key objective of antitrust authorities and law makers in several European countries.[5]

The article contributes to the growing literature on buyer power. According to this literature, larger retailers can obtain more favourable terms and conditions as their orders break collusion between suppliers (Snyder, 1996), or as they can threaten to integrate backwards (Katz, 1987; Fumagalli and Motta, 2000). Also, a supplier's threat to sell to competing retailers may become less valuable after a downstream merger (von Ungern-Sternberg, 1996; Dobson and Waterson, 1997; Mazzarotto, 2001).[6] If suppliers have convex costs, Chipty and Snyder (1999) and Inderst and Wey (2002) show that a larger retailer can negotiate lower prices. This is the case as the smaller retailers must negotiate more 'on the margin', where average unit costs are higher.[7]

The rest of this article is organised as follows. Section 1 contains our main results on how a consolidated retailer can use its newly acquired buyer power. Section 2 extends the model by endogenising product characteristics. Section 3 considers linear contracts. Section 4 discusses how some of the model's assumptions about suppliers' costs can be relaxed and how our theory may also relate to the formation of buyer groups. Section 5 offers concluding remarks.

1. The Main Model

1.1. *The Economy*

There are two suppliers $s \in S = \{A, B\}$, each of which produces a single good. Goods can be sold in two retail outlets $r \in R = \{a, b\}$. We assume that the two outlets operate in independent markets, in which the respective retailers act as monopolists.[8]

The characteristics of a supplier's good are captured by a real-valued parameter θ^s. (We denote parameters and functions relating to retailers by subscripts and those relating to suppliers by superscripts.) For the moment, the characteristics θ^s are taken as exogenously given. In Section 2, we let suppliers choose the characteristics of their goods. As shelf space at each outlet is scarce, each retailer must limit its range of product offerings. Thus, we assume that it is optimal for each retailer to stock at most one of the two goods, either A or B but not both.

If a good with characteristics θ is sold at price p, the demand at outlet r equals $D_r(\theta, p)$. The respective inverse demand function is $P_r(\theta, x)$, where x denotes the sold quantity. Suppliers have symmetric and constant marginal costs of production c. Denote $\Pi_r(\theta) := \max_x x[P_r(\theta, x) - c]$, which equals the maximum feasible profit that

[5] For instance, one of the main remedies in the Carrefour/Promodés merger is that contracts with 'economically dependent' suppliers must not be changed to their disadvantage over three years following the merger (Carrefour/Promodes EC/DGIV, 2000, Case No. COMP/M.1648). In France, 'economically dependent' suppliers can sue if they are delisted. For more details on such laws in the EU, see Clarke *et al.* (2002).

[6] See also the seminal work by Horn and Wolinsky (1988). Inderst and Wey (2002) show that, even without downstream interaction among the retailers, the supplier's loss incurred when negotiations break down increases over-proportionally with the retailer's size, which in turn allows a larger retailer to obtain more favourable terms.

[7] For experimental results on this, see Normann *et al.* (2003).

[8] This assumption allows us to abstract from standard monopolisation effects.

can be jointly realised by the retailer and its supplier when a good with characteristics θ is sold at outlet r. Note that $\Pi_r(\theta)$ would be realised by an integrated firm that controls both production and final sales.[9] Until Section 2, where θ^s is endogenously determined, we work with the following assumption.

ASSUMPTION 1. $\Pi_a(\theta^A) > \Pi_a(\theta^B)$ *and* $\Pi_b(\theta^B) > \Pi_b(\theta^A)$.

That is, the maximum feasible profit that can be jointly realised at outlet a is higher if good A is stocked, while the opposite is true at outlet b , where good B provides a better 'fit'.

There are many reasons to explain why Assumption 1 may hold in practice. One possibility is that consumers may predictably differ in tastes and preferences as outlets a and b may be located in different regions or even different countries. Another possibility is that good A may only be well established in the market where outlet a operates, while good B may have brand recognition only for customers of outlet b. Once again, this interpretation seems to be particularly suitable in case the two outlets are located in separate countries. Section 2 shows how Assumption 1 arises endogenously if suppliers choose their product characteristics optimally. Section 4.2 shows that our main results extend to a model where products are homogeneous but where suppliers differ in how close their factories are located to the different outlets. Thus, the resulting differences in (per-unit) transportation costs generate the same outcome as product differentiation.

1.2. *Contracts and Negotiations*

We consider two different scenarios. In the first scenario, outlets a and b are operated by different retailers and, therefore, each retailer chooses separately which good to stock. In the second scenario, the two outlets are operated by a single consolidated retailer.

Retailers set prices at their respective outlets after their contract terms have been determined. We place no restrictions on the set of feasible contracts. Optimally, the retailer and its supplier will thus choose a contract that avoids double-marginalisation and maximises their joint profit. A simple contract that rules out double-marginalisation is a 'forcing contract', which stipulates that the retailer can purchase a pre-specified quantity – and only this quantity – at a lump-sum price. Alternatively, it suffices for the two sides to agree on a two-part tariff contract which lets the retailer buy at marginal cost and allows the supplier to earn its profits via a 'flat fee'.

To pin down the structure of negotiations, we suppose for the case with two separate retailers that each firm has two agents, that is, two sales representatives or two managers respectively, who act independently but in the interest of the respective firm.

[9] Incidentally, nothing in our model relies on the assumption that retailers are monopolists in their local market. In fact, we could think of $\Pi_r(\theta)$ as retailer r's equilibrium profits given the optimal choices of all competing outlets. However, to abstract from the standard monopolisation effects of a downstream merger, we need that retailer a and b do not compete in the same markets, i.e., retailer a has no outlets in retailer b's markets and *vice versa*.

Negotiations proceed simultaneously, and agents form rational expectations about the outcomes of all other negotiations.[10]

At a given outlet only one supplier is chosen. We specify that the winning supplier can extract the fraction $\beta \in [0, 1]$ of the realised net surplus.[11] Our assumption of a fixed division of the realised net surplus admits several interpretations. If suppliers can make take-it-or-leave-it offers to retailers, we have that $\beta = 1$. In fact, as we argue below in more detail, the outcome of negotiations is then the same as that of an auction conducted by each retailer. If retailers can make take-it-or-leave-it offers, we have that $\beta = 0$. And if the winning supplier and retailer divide the gains from trade equally, as in symmetric Nash bargaining, then $\beta = 1/2$.

We now proceed as follows. In Section 1.3, we present the solutions for the two benchmark cases where $\beta = 1$ and $\beta = 0$. In Section 1.4, we consider the case of general values for β.

1.3. *Analysis for Two Benchmarks*

We begin by supposing that the suppliers have all the bargaining power and thus that they can make take-it-or-leave-it offers to the retailers: $\beta = 1$. As suppliers must still compete with each other to win over a particular retailer, it is as if retailers were to auction off their shelf space.[12]

Suppose first that there are two separate retailers. From Assumption 1, we have that, in equilibrium, good A will be stocked at outlet a and good B at outlet b. Moreover, as we allowed for efficient contracting, we know that each supply contract maximises the respective joint surplus. That is, the maximum feasible profit $\Pi_a(\theta^A)$ is realised at outlet a, while $\Pi_b(\theta^B)$ is realised at outlet b. Finally, as all the bargaining power lies with the suppliers, each retailer can only extract the value of its respective outside option. Formally, retailer a realises the profit $\Pi_a(\theta^B)$ and retailer b realises the profit $\Pi_b(\theta^A)$.[13] By virtue of their bargaining power, each supplier can pocket the full added value (or incremental surplus) generated by its product, that is, $\Pi_a(\theta^A) - \Pi_a(\theta^B)$ in the case of supplier A and $\Pi_b(\theta^B) - \Pi_b(\theta^A)$ in the case of supplier B.

[10] It is known that such a bargaining protocol could give rise to co-ordination failure among the different agents of each retailer. For instance, though A provides a better fit for a the agent that negotiates with B may conclude a forcing contract, which – given that only one good is stocked – makes it in turn optimal for the other agent of *a not* to conclude a contract with A. There are many ways to rule out these less plausible outcomes. A very simple way is to prescribe that any contract allows the retailer to still opt out, i.e., to procure zero units at zero costs.

[11] We combine both non-cooperative and cooperative solution concepts in our model. This is commonly done in the literature, and it allows us to obtain a parsimonious model of negotiations. For a non-cooperative foundation of the asymmetric Nash bargaining solution, see Binmore *et al.* (1986).

[12] Such auctions are also considered in O'Brien and Shaffer (1997), Dana (2004), and Marx and Shaffer (2004).

[13] This specification entails a standard bargaining assumption. Take the pairing of A with a. That $\Pi_a(\theta^B)$ is a's outside option requires that supplier B makes a 'best effort' to obtain the respective retailer's account, even though he knows he will not win it (note that in the bargaining model, suppliers do not compete by 'open outcry', which would immediately imply that B 'bids up' to (at least) $\Pi_a(\theta^B)$). With two-part tariff contracts, each supplier would set the per-unit price equal to its constant marginal cost c, while the fixed fee in the contract with B would be zero and the fixed fee in the contract with A would be equal to A's profit, $\Pi_a(\theta^A) - \Pi_a(\theta^B)$.

For future reference, it is worthwhile to state the two retailers' joint profits, which equal

$$\Pi_a(\theta^B) + \Pi_b(\theta^A), \tag{1}$$

and the two suppliers' joint profits, which equal the sum of the added values,

$$\left[\Pi_a(\theta^A) - \Pi_a(\theta^B)\right] + \left[\Pi_b(\theta^B) - \Pi_b(\theta^A)\right]. \tag{2}$$

Suppose next that the two retailers merge and form a single (consolidated) retailer. If the consolidated retailer still negotiates separately over which good to stock at outlets a and b, it is immediate that the preceding analysis does not change. In particular, the most suitable product will still be stocked at each outlet and the consolidated retailer's overall profit will equal (1).

As the retailer controls the stocking decision in both outlets, he can also adopt a different purchasing strategy, namely to stock only the same good at both outlets. We refer to this as a single-sourcing strategy. Which good will be stocked at both outlets under the single-sourcing strategy? The answer depends on which supplier can promise higher total profit. Hence, if

$$\Pi_a(\theta^A) + \Pi_b(\theta^A) > \Pi_a(\theta^B) + \Pi_b(\theta^B) \tag{3}$$

holds, then supplier A's good provides the better overall fit and A will win, while B will win if the opposite of (3) holds. If the terms on both sides of (3) are equal, both goods provide an equally good 'average fit'. In this case, it does not matter for profits which good is chosen.

Suppose (3) holds. Given that $\beta = 1$, the retailer's profit is again equal to its outside option, which is now equal to what it could earn being supplied under a competitive offer from B:

$$\Pi_a(\theta^B) + \Pi_b(\theta^B). \tag{4}$$

The winning supplier, A, extracts its good's added value, which is the difference between total industry profits if good A is chosen and total industry profits if good B is chosen:

$$\left[\Pi_a(\theta^A) + \Pi_b(\theta^A)\right] - \left[\Pi_a(\theta^B) + \Pi_b(\theta^B)\right]. \tag{5}$$

The losing supplier, B, realises zero profits as it is the one being excluded under single-sourcing.

How does the consolidated retailer's profit compare to the case with separate retailers? Subtracting (1) from (4) yields $\Pi_b(\theta^B) - \Pi_b(\theta^A)$, which is strictly positive by Assumption 1. That is, the retailers' profits increase by $\Pi_b(\theta^B) - \Pi_b(\theta^A)$ after a merger and the implementation of the single-sourcing strategy. For the case where the opposite of (3) holds and, consequently, supplier B is the winner, we get expressions that are symmetric to those in (4) and (5). The gain to the retailers from merging is then equal to $\Pi_a(\theta^A) - \Pi_a(\theta^B)$. We summarise our results.

PROPOSITION 1. *When suppliers have all the bargaining power* ($\beta = 1$), *retailers' profits strictly increase after a merger if the consolidated retailer chooses a single-sourcing policy. In the case* *where good A will subsequently be stocked at both outlets, which holds for* (3)*, the retailers' gain from merging equals* $\Pi_b(\theta^B) - \Pi_b(\theta^A)$*, while if the opposite of* (3) *holds, it equals* $\Pi_a(\theta^A) - \Pi_a(\theta^B)$.

Proposition 1 may at first seem surprising given that single sourcing reduces total industry profits (this follows because the best fit is no longer chosen at each outlet).[14] Formally, if only good *A* is stocked at both outlets, the loss in total industry profits is equal to $\Pi_b(\theta^B) - \Pi_b(\theta^A)$. But note that this loss is exactly offset by the decrease in supplier *B*'s profit (which is intuitive because in the absence of single sourcing, supplier *B* earns its good's full added value), which implies that single sourcing has no effect on the joint profit of the retailer and supplier *A*. Why then does single sourcing make the consolidated retailer strictly better off? The answer is that there is a redistribution of profit from the winning supplier, *A*, to the consolidated retailer. To see this, note that supplier *A*'s loss, which we obtain by subtracting the first pair of bracketed terms in (2) from (5), is $\Pi_b(\theta^B) - \Pi_b(\theta^A)$, which is the increase realised by the retailer. The reasoning is analogous if the opposite of (3) holds, so that only good *B* is stocked at both outlets.

Intuitively, single sourcing allows the consolidated retailer to extract a proportionately larger share of a smaller total industry profit because it makes suppliers less differentiated. The resulting averaging of each good's added values (e.g., good *A*'s added value is positive at outlet *a* but negative at outlet *b*) causes competition for the retailer's patronage to be more vigorous, which in turn facilitates the transfer of surplus from the winning supplier to the retailer. The idea that a retailer can increase its share of industry profits by smoothing out the added values of each supplier's good has some similarities to the reasons why bundling can be optimal for a monopolist (Adams and Yellen, 1976; Palfrey, 1983; McAfee *et al.*, 1989).[15] However, what is missing from the bundling literature is the analogue to the loss in industry profit that is necessarily induced when a good is excluded from distribution. Under pure bundling, for example, both goods continue to be sold and the total realised surplus may actually be higher.

For the merger and the subsequent single sourcing to be profitable, it is essential that the supplier base of the two outlets be different before the merger. If Assumption 1 did not hold and both retailers had originally the same supplier, i.e., either *A* or *B*, the merger would not lead to a change in the good carried at either outlet nor would it generate additional profit. Since, as noted above, outlets may carry different goods before the merger due to differences in consumer tastes or in brand recognition across regions or countries, it follows that we would expect our theory to be especially applicable to cross-border retail mergers, where these differences are likely to be more pronounced. We discuss the policy implications of this in Section 2.3 below.

[14] Under some additional and quite standard assumptions, the loss in product variety also leads to a reduction in welfare. We explore this issue in Section 2, where we specify how θ affects demand and industry profits.

[15] In the procurement literature, the strand of literature closest to ours is the one on split-award contracts, which asks the opposite question of when choosing multiple suppliers can be optimal. However, the focus there is different as the literature deals with reducing (suppliers') information rents (Riordan and Sappington, 1989), attracting more competition (Perry and Sakovics, 2003), or generating production efficiencies (Anton and Yao, 1989). In the vertical-contracting literature, the strand of literature closest to ours is the one on exclusive-dealing provisions in menu auctions (Bernheim and Whinston, 1998; O'Brien and Shaffer, 1997).

Our theory also suggests that, conditional on the original outlets having different supplier bases, the retailers' gain from merging and then implementing a single-sourcing strategy (e.g., $\Pi_b(\theta^B) - \Pi_b(\theta^A)$ if supplier A is subsequently chosen), is higher the more differentiated are the suppliers' goods (equivalently the stronger are the consumers' preferences for good B over good A at outlet b). While the loss from delisting one of the goods increases in their differentiation, so does the redistribution of profit from the winning supplier to the consolidated retailer, and thus the retailers' gain from merging also increases – for $\beta = 1$ by exactly the same amount.

Before proceeding to the case with general β, it is instructive to deal first with the other extreme: $\beta = 0$. In this case, the retailers have all the bargaining power and can extract all the industry profits.[16] It is thus immediate that a merger cannot increase the retailers' joint profit. What is more, we know that a consolidated retailer would now strictly prefer *not* to choose a single-sourcing policy. For $\beta = 0$ we have in fact the opposite result from that in Proposition 1: a decrease in total industry profit by £1 also reduces the consolidated retailer's profit by £1.

PROPOSITION 2. *When retailers have all the bargaining power ($\beta = 0$), retailers' profits strictly decrease after a merger if the consolidated retailer chooses a single-sourcing policy. Hence, the consolidated retailer would buy from both suppliers, and thus there is no gain from merging.*

1.4. *The General Case*

With general β, each party to a deal can extract a strictly positive share of its added value. Hence, if retailer a negotiates with suppliers A and B, its profit will be the sum of its outside option $\Pi_a(\theta^B)$ *plus* the fraction $1 - \beta$ of the added value realised with supplier A, that is, of the difference $\Pi_a(\theta^A) - \Pi_a(\theta^B)$. Summing up, the two retailers' joint profits will equal

$$\Pi_a(\theta^B) + (1-\beta)\left[\Pi_a(\theta^A) - \Pi_a(\theta^B)\right] + \Pi_b(\theta^A) + (1-\beta)\left[\Pi_b(\theta^B) - \Pi_b(\theta^A)\right]$$

$$= (1-\beta)\left[\Pi_a(\theta^A) + \Pi_b(\theta^B)\right] + \beta\left[\Pi_a(\theta^B) + \Pi_b(\theta^A)\right], \tag{6}$$

while the two suppliers' joint profits will equal β *times* the sum of the added values,

$$\beta\left[\Pi_a(\theta^A) - \Pi_a(\theta^B)\right] + \beta\left[\Pi_b(\theta^B) - \Pi_b(\theta^A)\right]. \tag{7}$$

In the case of a merger and subsequent single sourcing, if (3) holds and supplier A is chosen, the consolidated retailer's profit will equal its outside option $\Pi_a(\theta^B) + \Pi_b(\theta^B)$ plus the fraction $1 - \beta$ of the added value realised with supplier A, i.e., of the difference between $\Pi_a(\theta^A) + \Pi_b(\theta^A)$ and $\Pi_a(\theta^B) + \Pi_b(\theta^B)$. Hence, the consolidated retailer's profit under single sourcing will be

[16] It does not matter whether each supplier also sells to other outlets in markets that do not overlap those where retailers a and b are located as long as the suppliers' gains from these contracts are unaffected by their dealings with a and b. The latter holds in our model because the suppliers' marginal costs are constant.

$$\Pi_a(\theta^B) + \Pi_b(\theta^B) + (1-\beta)\{[\Pi_a(\theta^A) + \Pi_b(\theta^A)] - [\Pi_a(\theta^B) + \Pi_b(\theta^B)]\}$$

$$= (1-\beta)[\Pi_a(\theta^A) + \Pi_b(\theta^A)] + \beta[\Pi_a(\theta^B) + \Pi_b(\theta^B)], \quad (8)$$

while the winning supplier, A, will earn a profit equal to

$$\beta\{[\Pi_a(\theta^A) + \Pi_b(\theta^A)] - [\Pi_a(\theta^B) + \Pi_b(\theta^B)]\}. \quad (9)$$

What are now the gains and losses from a single-sourcing strategy? The reduction in total industry profits is, of course, the same as in the benchmark cases, i.e., $\Pi_b(\theta^B) - \Pi_b(\theta^A)$. Subtracting next (6) from (8), we see that the retailers' profits under single sourcing change by

$$(2\beta - 1)[\Pi_b(\theta^B) - \Pi_b(\theta^A)]. \quad (10)$$

For the case where supplier B is chosen, the expressions are symmetric and we obtain – again in symmetry to expression (10) – that the retailers' profits change by $(2\beta - 1)[\Pi_a(\theta^A) - \Pi_a(\theta^B)]$. As the difference in parentheses in each case is positive, we thus have the following result.

PROPOSITION 3. *The consolidated retailer is strictly better off implementing a single-sourcing policy if and only if* $\beta > 1/2$, *that is, if and only if the suppliers have sufficient bargaining power. Thus, a merger makes the retailers strictly better off only in the case where* $\beta > 1/2$.

Proposition 3, which generalises the results in Propositions 1 and 2 to any $\beta \in [0, 1]$, contains the main result of the article.

In the absence of single sourcing, the consolidated retailer must bargain to make a deal twice. In each case it is bargaining with the same two other parties. However, the bargains differ in a critical way. In the first bargaining situation, joint profit is maximised if the retailer makes a deal with supplier A, whereas in the second bargaining situation, joint profit is maximised if the retailer makes a deal with supplier B. In the first bargain, the retailer gets its outside option plus a share of the extra surplus created by dealing with supplier A. The converse happens in the second bargain. With single sourcing, however, the two formerly separate bargains are treated as one. Although this does not alter the 'bargaining power' of any player, it does alter the 'outside option' facing the consolidated retailer. This is due to the inverse correlation between surplus generated by each of A and B in each of the two deals. Indeed, if the net gain from dealing with supplier A in the first bargain is just equal to the net gain from dealing with supplier B in the second bargain, then by committing to only one negotiation rather than two, the consolidated retailer gets all the surplus from the bargaining.

Will the consolidated retailer gain? This depends on whether the increase in surplus seized by the retailer more than offsets the loss associated with choosing the inefficient party for one of the bargains. Intuitively, single sourcing reduces total industry profits. The reduction in industry profits equals the value added of the losing supplier's good, a portion $1 - \beta$ of which would have been captured by the retailer. On the other hand, single sourcing increases the retailer's proportionate share of the remaining industry profit by making suppliers more homogenous. Proposition 3 shows that if suppliers can extract more than half of the added value of their products, the latter effect dominates the former effect (the retailer prefers a proportionately larger share of a smaller overall pie) and thus a single sourcing policy becomes strictly optimal.

We conclude this Section with several comments regarding Proposition 3. Note first that we apply the same value of β both for separate retailers and for the consolidated retailer. Instead of assuming, for instance, that β increases after a merger, we believe it is important to endogenise the mechanism by which buyer power is created by a merger. Note second that, in our setting, the exogenous factor β is a catch-all measure of bargaining power that arises from other sources. Proposition 3 may thus entail some further, potentially testable, implications of our theory. For instance, it suggests that the level of β may be an important predictor of the profitability of mergers. By Proposition 3, we have that increasing buyer power via single sourcing is a profitable strategy when dealing with strong suppliers, that is, suppliers owning strong brands and 'must-stock' items. In order to explore the implication of this in more generality, we would, however, need to introduce other sources of retailer and supplier power so as to endogenise β.

It is important to note that single sourcing – or, more generally, reducing the number of competing products in the market – serves as an instrument to generate more power and profits for the *retailer*. That is, in our theory it is *not* the strong-brand manufacturer that imposes on the retailer the exclusion of competing brands, but it is the retailer for which this is optimal.[17]

Finally, adopting a single-sourcing strategy necessarily requires some amount of commitment on the part of the retailer. As is immediate from our calculations, *both* suppliers would be strictly better off if they could commit not to participate in negotiations for the combined shelf space and if they could, instead, force the retailer to negotiate separately over both outlets. In practice, retailers should, however, enjoy considerable scope in determining how to allocate their shelf space. Moreover, in the wake of a general reorganisation following a merger, a single-sourcing strategy may be made credible by implementing changes in the distribution system or by top management's directive to 'prune' the supplier base of the two merging retailers vigorously.

2. Endogenous Variety

2.1. *Extending the Model*

We now endogenise the suppliers' choices of product characteristics θ^s. In doing so, we consider the following sequence of events. In the first period, $t = 1$, suppliers choose their real-valued characteristics θ^s non-cooperatively. In the second period, $t = 2$, a retail merger may or may not occur. The rest of the game is then as described above. In the third period, $t = 3$, retailers choose their purchasing strategy, that is, whether or not to commit to a single-sourcing policy. (This is only a non-trivial choice for a consolidated retailer.) In the fourth period, $t = 4$, retailers and suppliers negotiate under

[17] A similar observation has been made in O'Brien and Shaffer (1997) for the case of efficient contracting with exclusive dealing provisions and in Gabrielsen and Sorgard (1999) when contracts are restricted to be linear tariffs.

the retailers' chosen purchasing strategy. In the final period, $t = 5$, retailers set prices for final consumers, goods are supplied, and payoffs are realised.

The newly introduced stages $t = 1, 2$ deserve some comments. Consider first the suppliers' choices of product characteristics at $t = 1$.[18] Recall that $\Pi_r(\theta)$ denotes the maximum feasible profit that can be realised when supplying a good with characteristics θ at outlet r. We assume that $\Pi_r(\theta)$ is strictly concave in θ (where $\Pi_r(\theta) > 0$) and $\Pi_r(\theta) > 0$ for some θ. We also assume that $\mathrm{d}\Pi_b(\theta)/\mathrm{d}\theta > \mathrm{d}\Pi_a(\theta)/\mathrm{d}\theta$ holds whenever both $\Pi_r(\theta)$ are differentiable and at least one is strictly positive. Define $\hat{\theta}_r := \arg\max_\theta \Pi_r(\theta)$ and note that $\hat{\theta}_a < \hat{\theta}_b$. It also proves useful to assume that both $\Pi_a(\theta)$ and $\Pi_b(\theta)$ are strictly positive over the 'relevant' range $\theta \in [\hat{\theta}_a, \hat{\theta}_b]$.

Consider next the possibility of a retail merger at $t = 2$. We specify that this happens with some exogenous probability μ. The choices $\mu = 0$ and $\mu = 1$ correspond to cases where the merger occurs never or always. Although a merger is always weakly profitable for retailers, it may not always be possible. For instance, the owners or the management of a retailer may not be prepared to relinquish control, or a merger may come at prohibitively high transactions costs. What is more relevant for our discussion, however, is that the competition authority may adopt a more or less lenient merger policy for retailers, which is captured by μ in a short-cut way.[19]

2.2 *Analysis*

It is helpful to consider first the case where suppliers anticipate that no merger will occur, $\mu = 0$. In this case, in any pure-strategy equilibrium, each supplier maximises its profit by focusing on only one outlet; one supplier optimally chooses $\hat{\theta}_a$ and the other supplier optimally chooses $\hat{\theta}_b$ (this is analogous to firms choosing to differentiate maximally *à la* Hotelling line, where instead of consumers being uniformly distributed, they are located only at the endpoints). Thus, when $\mu = 0$, suppliers choose the product characteristics that maximise total industry profits.

These strategies are, however, no longer an equilibrium if $\mu > 0$ (and $\beta > 1/2$). Taking into account the possibility that a merger will occur, at least one supplier will reposition its product to balance the different consumer preferences more adequately at the two outlets. Define[20]

$$\tilde{\theta}_r := \arg\max_\theta \{\mu[\Pi_r(\theta) + \Pi_{r'}(\theta)] + (1-\mu)\Pi_r(\theta)\} \text{ where } r' \neq r, \tag{11}$$

where the expression to be maximised is a weighted average of the sum of the maximum feasible profits that can be realised at the two outlets and the maximum feasible profit that can be realised at outlet r when a good with characteristics θ

[18] That suppliers choose θ^s non-cooperatively is a standard assumption that basically rules out negotiations with retailers over the jointly optimal choice of characteristics. However, in our present setting with non-linear tariffs, it turns out that the suppliers' non-cooperative choices also maximise total expected industry profits.

[19] Capturing merger policy in this way seems more appropriate in particular in the case of cross-border mergers where buyer power is a consideration. In these cases, it is fair to say that competition authorities have not yet developed a systematic framework, unlike for, say, horizontal mergers between suppliers in oligopolistic industries.

[20] Uniqueness of $\tilde{\theta}_r$ follows from strict concavity of both $\Pi_a(\theta)$ and $\Pi_b(\theta)$ over $\theta \in [\hat{\theta}_a, \hat{\theta}_b]$.

is supplied. It follows that as μ increases, $\tilde{\theta}_r$ increases from a lower bound of $\hat{\theta}_a$ when $r = a$ and decreases from an upper bound of $\hat{\theta}_b$ when $r = b$. It also follows that if $\theta^A = \tilde{\theta}_a$ and $\theta^B = \hat{\theta}_b$, then the construction of $\tilde{\theta}_a$ ensures that total expected industry profits are maximised, provided that good A is chosen by the consolidated retailer under single sourcing. The interpretation of $\tilde{\theta}_b$ as maximising total expected industry profits is analogous if $\theta^B = \tilde{\theta}_b$ and $\theta^A = \hat{\theta}_a$, provided that good B is chosen under single sourcing.

One might think that both suppliers will want to re-position their products in anticipation of the possibility of a merger, or, at the very least, that the winning supplier under single sourcing will want to distort its product characteristics away from $\tilde{\theta}_r$ in order to deter its rival from moving to challenge on the consolidated market. However, as we now show, this is not the case.

Suppose that $\mu < 1$ and $\beta > 1/2$. Then, in any pure-strategy equilibrium, the chosen pair of characteristics must be either $(\tilde{\theta}_a, \hat{\theta}_b)$ or $(\tilde{\theta}_b, \hat{\theta}_a)$. To see this, we argue first that some supplier must choose $\theta^s \in \left\{\hat{\theta}_a, \hat{\theta}_b\right\}$. Suppose this were not the case. Now if one of the suppliers, say A, wins the single-sourcing contract with probability one following a merger, the other supplier, B, can profitably deviate to some $\theta^B \in \left\{\hat{\theta}_a, \hat{\theta}_b\right\}$, which maximises its profit in the case where no merger occurs (analogously if the names of the winning and losing suppliers are reversed). If both suppliers are chosen with positive probability under single sourcing, then $\Pi_a(\theta^A) + \Pi_b(\theta^A)$ equals $\Pi_a(\theta^B) + \Pi_b(\theta^B)$, which implies that both will realise zero profit if a merger takes place. But this also cannot be an equilibrium as it would again be profitable for at least one supplier to deviate and choose some $\theta^s \in \left\{\hat{\theta}_a, \hat{\theta}_b\right\}$. This proves that some supplier must choose $\theta^s \in \left\{\hat{\theta}_a, \hat{\theta}_b\right\}$.

Suppose now that there is a pure-strategy equilibrium in which one of the suppliers, say A, wins the single-sourcing contract with probability one following a merger.[21] Then, the preceding discussion implies that, in this case, $\theta^B \in \left\{\hat{\theta}_a, \hat{\theta}_b\right\}$ (otherwise, supplier B would have a profitable deviation). If $\theta^B = \hat{\theta}_b$, then supplier A's expected profit given that its good will be sold at outlet a if no merger takes place and at both outlets (given single sourcing) following a merger is

$$\begin{aligned}\mu\beta\Big\{\left[\Pi_a(\theta^A) + \Pi_b(\theta^A)\right] - \Big[\Pi_a(\hat{\theta}_b) + \Pi_b(\hat{\theta}_b)\Big]\Big\} + (1-\mu)\beta\Big[\Pi_a(\theta^A) - \Pi_a(\hat{\theta}_b)\Big] \qquad (12)\\ = \beta\Big[\Pi_a(\theta^A) - \Pi_a(\hat{\theta}_b)\Big] + \mu\beta\Big[\Pi_b(\theta^A) - \Pi_b(\hat{\theta}_b)\Big],\end{aligned}$$

which is maximised at $\theta^A = \tilde{\theta}_a$. If $\theta^B = \hat{\theta}_a$, then $\theta^A = \tilde{\theta}_b$ maximises supplier A's expected profit. The converse holds if supplier B wins the single-sourcing contract with probability one.

Our arguments imply that, in any pure-strategy equilibrium, one supplier will maximise its expected payoff at one outlet (and thus will not re-position its product in anticipation of the possibility of a merger) while the other supplier will choose $\theta^s = \tilde{\theta}_r$

[21] There are no pure-strategy equilibria in which both suppliers are chosen with positive probability under single sourcing. To see this, note that, in any such equilibrium, the preceding discussion implies $\theta^A \in \left\{\hat{\theta}_a, \hat{\theta}_b\right\}$ and $\theta^B \in \left\{\hat{\theta}_a, \hat{\theta}_b\right\}$. But then it is easy to show that one of the suppliers would always want to deviate to $\theta^s \in \left\{\tilde{\theta}_a, \tilde{\theta}_b\right\}$.

(and thus will not distort its product characteristics to preempt a challenge by its rival). We thus have the following result.

Proposition 4. *Suppose that $\mu < 1$ and $\beta > 1/2$. Then in any pure-strategy equilibrium of the game at $t = 1$, either one supplier chooses $\tilde{\theta}_a$ while the other chooses $\hat{\theta}_b$, or one supplier chooses $\tilde{\theta}_b$ while the other chooses $\hat{\theta}_a$. Moreover, as μ increases, goods become continuously less differentiated as $\tilde{\theta}_a \geq \hat{\theta}_a$ is strictly increasing in μ and as $\tilde{\theta}_b \leq \hat{\theta}_b$ is strictly decreasing in μ.*

Proof. It remains to prove the strict monotonicity of $\tilde{\theta}_a$ and $\tilde{\theta}_b$. Consider first the proof for $\tilde{\theta}_a$. Differentiating the terms inside the brackets in (11) yields, for $\tilde{\theta}_a$, the first-order condition

$$\frac{d\Pi_a(\theta)}{d\theta} + \mu\frac{d\Pi_b(\theta)}{d\theta} = 0 \,\text{at}\, \theta = \tilde{\theta}_a. \tag{13}$$

Since $d\Pi_b(\theta)/d\theta > d\Pi_a(\theta)/d\theta$ implies $d\Pi_b(\tilde{\theta}_a)/d\theta > 0$, and since $\Pi_r(\theta)$ is strictly concave, implicit differentiation of (13) shows that $d\tilde{\theta}_a/d\mu > 0$. The proof for $\tilde{\theta}_b$ is analogous. Q.E.D.

Proposition 4 warrants several comments. We consider first the role of μ . For $\mu = 0$, where $\tilde{\theta}_a = \hat{\theta}_a$ and $\tilde{\theta}_b = \hat{\theta}_b$, the two chosen characteristics provide the best fits for the respective outlets. For $0 < \mu < 1$, one supplier focuses on one outlet, while the other supplier chooses more 'average' characteristics. For the remaining case, $\mu = 1$, it is straightforward to establish that in any pure-strategy equilibrium a supplier who subsequently wins the single-sourcing contract must choose the unique characteristics that maximise total industry profits, $\Pi_a(\theta) + \Pi_b(\theta)$.

Next, consider the existence of equilibrium, on which Proposition 4 is silent. Suppose that

$$\Pi_a(\tilde{\theta}_a) + \Pi_b(\tilde{\theta}_a) \geq \Pi_a(\tilde{\theta}_b) + \Pi_b(\tilde{\theta}_b), \tag{14}$$

so that if one supplier were to choose $\tilde{\theta}_a$ and the other were to choose $\tilde{\theta}_b$, it would be the supplier with characteristics $\tilde{\theta}_a$ who would win the single-sourcing contract. Then, it is easy to show that $\theta^A = \tilde{\theta}_a$ is a best response to $\theta^B = \hat{\theta}_b$ (because (14) implies $\Pi_a(\tilde{\theta}_a) + \Pi_b(\tilde{\theta}_a) \geq \Pi_a(\hat{\theta}_b) + \Pi_b(\hat{\theta}_b)$).

We now consider whether $\theta^B = \hat{\theta}_b$ is a best response to $\theta^A = \tilde{\theta}_a$. Our previous arguments leave open two potentially profitable deviations for supplier B. One potentially profitable deviation is to choose some $\tilde{\theta}_a < \theta^B < \tilde{\theta}_b$ such that it would supply outlet b if no merger takes place and also win the single-sourcing contract after a merger. A second potentially profitable deviation is to choose some $\theta^B < \tilde{\theta}_a$ such that it is now B who would supply outlet a if no merger occurs. Of these two deviation strategies, we can show, using (14), that the former does not constitute a profitable deviation.[22] Without imposing additional assumptions on $\Pi_r(\theta)$, however, it is not possible to rule out the

[22] The proof is straightforward. As $\tilde{\theta}_b$ maximises $\Pi_b(\theta) + \mu\Pi_a(\theta)$, it follows from $\mu < 1$ and the strict concavity of $\Pi_r(\theta)$ over $\theta \in [\hat{\theta}_a, \hat{\theta}_b]$ that B can only win after a merger if $\theta^B < \tilde{\theta}_b$. Using continuity, denote the highest value $\theta^B < \tilde{\theta}_b$ where $\Pi_a(\theta^B) + \Pi_b(\theta^B) = \Pi_a(\tilde{\theta}_a) + \Pi_b(\tilde{\theta}_a)$ by $\bar{\theta}_b$. At $\theta^B = \bar{\theta}_b$, B can win the single-sourcing contract, but its expected profits would only be $(1-\mu)\beta[\Pi_b(\theta^B) - \Pi_b(\tilde{\theta}_a)]$, i.e., strictly smaller than when choosing $\theta^B = \hat{\theta}_b$. By concavity and $\theta^B < \tilde{\theta}_b$, B's expected profits from the deviation are even lower for all $\tilde{\theta}_a \leq \theta^B < \bar{\theta}_b$.

latter deviation strategy. There are, however, two cases for which it is immediate that this is not optimal. First, for a given $\Pi_r(\theta)$, this is always the case if μ is sufficiently low. This follows from the continuity of $\Pi_r(\theta)$ and because $\tilde{\theta}_a$ increases continuously with μ while $\tilde{\theta}_a = \hat{\theta}_a$ holds at $\mu = 0$. Second, B cannot profitably deviate to $\theta^B < \tilde{\theta}_a$ in the case where the profit functions $\Pi_r(\theta)$ are symmetric over the relevant range of values θ, i.e., if for all $\theta \in [\hat{\theta}_a, \hat{\theta}_b]$ it holds that $\Pi_a(\theta) = \Pi_b(\hat{\theta}_a + \hat{\theta}_b - \theta)$ and likewise that $\Pi_b(\theta) = \Pi_a(\hat{\theta}_a + \hat{\theta}_b - \theta)$. Formally, note that the value $\bar{\theta} = (\hat{\theta}_a + \hat{\theta}_b)/2$ maximises $\Pi_a(\theta) + \Pi_b(\theta)$ and that $\tilde{\theta}_a < \bar{\theta}$ holds for all $\mu < 1$. By symmetry, it then holds that $\Pi_b(\hat{\theta}_b) - \Pi_b(\tilde{\theta}_a) > \Pi_a(\hat{\theta}_a) - \Pi_a(\tilde{\theta}_a)$ for all $\mu < 1$.

There is also scope for the existence of multiple pure-strategy equilibria. To see this, suppose that (14) holds and consider a second pure-strategy candidate equilibrium where B chooses $\tilde{\theta}_b$ and A chooses $\hat{\theta}_a$. When will A find it profitable to deviate and challenge B for the single-sourcing contract, which is optimally done by choosing $\theta^A = \tilde{\theta}_a$? Substituting into A's expected profit in (12), where we have to replace $\hat{\theta}_b$ by $\tilde{\theta}_b$, we find that A's profit under the deviation is

$$\beta\left[\Pi_a(\tilde{\theta}_a) - \Pi_a(\tilde{\theta}_b)\right] + \mu\beta\left[\Pi_b(\tilde{\theta}_a) - \Pi_b(\tilde{\theta}_b)\right], \tag{15}$$

whereas A's expected profit if it does not deviate is $(1-\mu)\beta\left[\Pi_a(\hat{\theta}_a) - \Pi_a(\tilde{\theta}_b)\right]$. Comparing these two profits and rearranging terms, we have that A's deviation is profitable if and only if

$$\left[\Pi_a(\tilde{\theta}_a) + \Pi_b(\tilde{\theta}_a)\right] - \left[\Pi_a(\tilde{\theta}_b) + \Pi_b(\tilde{\theta}_b)\right] > \frac{1-\mu}{\mu}\left[\Pi_a(\hat{\theta}_a) - \Pi_a(\tilde{\theta}_a)\right]. \tag{16}$$

That is, the profit from winning the single-sourcing contract (the left-hand side of (16)) must be sufficiently large to compensate the loss at outlet A from choosing $\tilde{\theta}_a$ instead of $\hat{\theta}_a$. As the right-hand side of (16) is strictly positive unless $\mu = 1$, condition (16) is weaker than (14).

2.3. *Implications for Industry Profits and Welfare*

By Proposition 4, the supplier that expects to win the single-sourcing contract after a merger, say A, chooses more 'average' product characteristics. This repositioning of A's product has both positive and negative effects. On the one hand, it increases industry profits (and A's profit) *conditional* on there being a merger, and it increases expected industry profits conditional on the *ex ante* likelihood of a merger. On the other hand, the decrease in product variety reduces expected industry profits relative to the case of $\mu = 0$ (if no merger occurs, supplier A will have chosen a suboptimal variety for outlet a, and if a merger does occur, supplier A will have chosen a suboptimal variety for outlet a and there will be a reduction in realised profit at outlet b as good B will not be sold). Moreover, as the next result shows, expected industry profits are strictly decreasing in μ, which implies that expected industry profits are lowest when $\mu = 1$.[23]

[23] For this comparative exercise, we assume the existence of a pure-strategy equilibrium. Moreover, if two pure-strategy equilibria exist, we consider the case where following a marginal adjustment of μ we do not switch.

COROLLARY 1. *If* $\beta > 1/2$, *then expected industry profits are strictly decreasing in* μ.

Proof. Suppose $\theta^A = \tilde{\theta}_a$ and $\theta^B = \hat{\theta}_b$. Then expected industry profits are equal to

$$\mu\left[\Pi_a(\tilde{\theta}_a) + \Pi_b(\tilde{\theta}_a)\right] + (1-\mu)\left[\Pi_a(\tilde{\theta}_a) + \Pi_b(\hat{\theta}_b)\right]. \tag{17}$$

Differentiating (17) with respect to μ and using that $\tilde{\theta}_a$ satisfies the first-order condition in (13), we have the derivative $\Pi_b(\tilde{\theta}_a) - \Pi_b(\hat{\theta}_b)$, which is negative because $\hat{\theta}_b$ is the unique maximiser of $\Pi_b(\theta)$. The other case, where $\theta^A = \hat{\theta}_a$ and $\theta^B = \tilde{\theta}_b$, is analogous. Q.E.D.

Without further assumptions on consumer preferences and local demand, we cannot make any claims on how expected consumer surplus and welfare change in μ. However, we can obtain results for the relatively standard case in which the inverse demand takes on the additive form

$$P_r(\theta, x) = \max[p_r(x) + \psi_r(\theta), 0], \tag{18}$$

where $P_r(\theta, x)$ is generated by the preferences of a representative consumer and where $xP_r(\theta, x)$ is strictly concave (recall that x denotes the quantity sold by retailer r). One case where (18) is satisfied is that of linear demand, which is studied below. We have the following result.

COROLLARY 2. *If* $\beta > 1/2$, *and the inverse demand is of the additive form in* (18) *and captures the preferences of a representative consumer, expected welfare is also strictly decreasing in* μ.

Proof. We denote for given r and θ the unique optimal quantity by $x_r^*(\theta)$. The envelope theorem then implies that $\mathrm{d}\Pi_r(\theta)/\mathrm{d}\theta = x_r^*(\theta)\mathrm{d}\psi_r(\theta)/\mathrm{d}\theta$. Moreover, implicit differentiation of the first-order condition for $x_r^*(\theta)$ shows that the sign of $\mathrm{d}x_r^*(\theta)/\mathrm{d}\theta$ is determined by the sign of $\mathrm{d}\psi_r(\theta)/\mathrm{d}\theta$. Welfare at outlet r is $W_r = \int_0^{x_r^*(\theta)}[p_r(x) + \psi_r(\theta)]\mathrm{d}x - cx^*(\theta)$. Differentiating welfare with respect to θ, we obtain $\mathrm{d}W_r/\mathrm{d}\theta = \partial W_r/\partial\,\theta + \partial W_r/\partial x\,\mathrm{d}x_r^*(\theta)/\mathrm{d}\theta$, where the signs of $\partial W_r/\partial\theta$ and $\mathrm{d}x_r^*(\theta)/\mathrm{d}\theta$ are equal to the signs of $d\psi_r(\theta)/d\theta$. Additionally, we have from standard results that $\partial W_r/\partial x > 0$ at $x = x_r^*(\theta)$.[24] Hence, we have established that welfare realised at outlet r changes in the characteristics of the supplied good in the same way as industry profits. The assertion follows then from Corollary 1. Q.E.D.

By Corollary 2, expected welfare is decreasing in the likelihood that a retail merger will occur. The reason for this is the same as the reason for why expected industry profits decrease in μ: there is a loss in product variety, both due to the delisting of one supplier if a merger occurs and due to suppliers' optimal choices of characteristics in anticipation of a merger.

According to informal arguments (see the reports listed in footnote 4), the main welfare loss due to the exercise of buyer power stems from suppliers' reduced investment incentives, e.g., on product quality, cost reduction, and variety. However,

[24] To be precise, note first that $\partial W_r/\partial x = P_r(\theta, x) - c$, while the first-order condition for profit maximisation gives $\mathrm{d}\Pi_r/\mathrm{d}x = P_r(\theta, x) - c + x\mathrm{d}P(\theta, x)/\mathrm{d}x$. The claim follows as $P_r(\theta, x)$ is strictly decreasing whenever $P_r(\theta, x) > 0$.

as shown in Inderst and Wey (2002), these arguments do not withstand formal scrutiny – at least not in this sweeping generalisation. While lower profits for suppliers may indeed reduce incentives for entry or the introduction of new products, for more incremental changes such as quality improvement or cost cutting it is the *marginal* change in profits that matters for incentives. From this perspective, Corollary 2 is important because it provides a stronger underpinning for why (even cross-border) retail mergers may have direct adverse welfare implications in the form of reduced product variety.

2.4. *Example*

With linear demand $D = 1 - d - p$, constant marginal costs $c < 1 - d$, and $\mu = 0$, the joint profit of the retailer and its supplier is maximised at the retail price $p = (1 + c - d)/2$. This generates sales $x = (1 - d - c)/2$, profits $\Pi = (1 - d - c)^2/4$ and, if the demand is generated by a representative consumer with a quadratic utility function, welfare $W = 3(1 - d - c)^2/8$.

To model diversity in tastes, for outlet a, we set $d = \theta^2/z$ to obtain $D_a(\theta, p) = 1 - p - \theta^2/z$, while for outlet b, we set $d = (1 - \theta)^2/z$ to obtain $D_b(\theta, p) = 1 - p - (1 - \theta)^2/z$ with $z > 0$. Consequently, it follows that $\hat{\theta}_a = 0$ and $\hat{\theta}_b = 1$ maximise joint profits at the respective outlets.

The case where suppliers can choose any value for θ has no closed-form solution if $\mu > 0$. Therefore, we confine ourselves to the case where θ can only be chosen from a finite set $\theta \in \Theta = \{\hat{\theta}_a, \theta^*, \hat{\theta}_b\}$, where $0 < \theta^* < 0.5$. Moreover, we choose the parameters $c = 0$, $z = 5$ and $\beta = 1$.

What product characteristics will suppliers choose? From Proposition 4, we have that one supplier, say B, must choose $\theta^s \in \{\hat{\theta}_a, \hat{\theta}_b\}$. Without loss of generality, let $\theta^B = \hat{\theta}_b$. Then, by substitution into (12), we have that the other supplier, A, strictly prefers θ^* to $\hat{\theta}_a$ if and only if

$$\mu > \theta^* \frac{10 - (\theta^*)^2}{16 - 4\theta^* + (\theta^*)^3 - 4(\theta^*)^2}. \tag{19}$$

The right-hand side of (19) is strictly increasing in θ^*, which is intuitive because the larger is the difference $\theta^* - \hat{\theta}_a$, the more likely must a merger be to make it optimal to choose θ^*.

For future reference, we specify $\theta^* = 0.2$, for which (19) becomes $\mu > 83/627 \approx 13.2\%$. That is, the likelihood of a merger must exceed 13.2% to induce supplier A to choose the less differentiated product variant. Since our linear demand satisfies (18), we finally have the following.

Results for the linear example:

(*i*) *If* $\mu > 13.2\%$, *supplier A chooses the less differentiated product variant* $\theta^A = \theta^* = 0.2$. *Otherwise, supplier A chooses the more differentiated product variant* $\theta^A = \hat{\theta}_a = 0$.

(*ii*) *Expected welfare is strictly decreasing in* μ.[25]

[25] It is straightforward to show that Corollary 2 implies assertion (*ii*) even though we currently only consider a discrete choice of product characteristics.

3. Linear Contracts

3.1. *Analysis*

So far we have assumed that negotiations are efficient. Retail contracts are indeed often complex and may include, for instance, volume discounts, slotting fees (to obtain shelf space), pay-to-stay fees (for continuation of stocking), display fees (for end-aisle caps), and presentation fees (for making a sales presentation). With efficient negotiations, final consumer prices are not affected by how surplus is distributed between the retailer and each supplier. In contrast, when negotiations are inefficient, for example, when retail contracts determine only a uniform purchase price, issues of double marginalisation arise. To explore these issues further, in what follows, we restrict attention to linear contracts, and we restrict consideration to the case where $\beta = 1$.[26]

3.1.1. *Separate retailers*

Suppose that supplier A wins outlet a with a price of m_a^A, and recall that if a good with characteristics θ is sold at price p, the demand at outlet r equals $D_r(\theta, p)$. Then, as B's offer to a is the uniform price $m_a^B = c$, to at least match B's offer, A 's price m_a^A must satisfy

$$\max_p (p - m_a^A) D_a(\theta^A, p) \geq \Pi_a(\theta^B). \tag{20}$$

In words, the maximum profit that retailer a can realise when buying from supplier A must be at least $\Pi_a(\theta^B)$. There are now two possible cases. In the first case, supplier A's offer is not constrained by B's offer and, thus, the constraint (20) is not binding.[27] In what follows, we focus on the more interesting second case where competition from B constrains A's offer. In this case, A optimally chooses m_a^A such that (20) binds. As retailer a's profits are strictly decreasing in m_a^A (as long as $D_a > 0$), this yields a unique offer.[28] Note that, in equilibrium, the supplier whose product offers the highest feasible profit $\Pi_r(\theta^s)$ wins the contract to supply retailer r.

3.1.2. *Consolidated retailer*

Under a single-sourcing policy, each supplier offers a single price.[29] The analysis is then analogous to the case with separate retailers, i.e.,

(i) the supplier s for which $\Pi_a(\theta^s) + \Pi_b(\theta^s)$ is highest wins the account,
(ii) the losing supplier offers $m^s = c$ and

[26] The case with $\beta < 1$ does not yield any insights beyond those obtained already for efficient negotiations. Moreover, we would have to establish that the bargaining set with linear contracts is still concave in order to apply the axiomatic Nash approach. While this holds for our linear example, it may not be satisfied for more general demand functions. (The standard remedy in this case would be to use lotteries over contracts.)

[27] Formally, suppose $p^*(m_a^A) := \arg\max_p[(p - m_a^A)D_a(\theta^A, p)]$ and $m^* := \arg\max_{m_a^A}\{(m_a^A - c)D_a[\theta^A, p^*(m_a^A)]\}$ are unique. Then supplier A's offer is not constrained by B's offer if $\{[p^*(m^*) - m^*]D_a[\theta^A, p^*(m^*)]\} \geq \Pi_a(\theta^B)$.

[28] Formally, A's offer to a is the highest value solving $\max_p[(p - m_a^A)D_a(\theta^A, p)] = \Pi_a(\theta^B)$.

[29] Alternatively, supplier s could offer two different prices for supplying outlets a and b. This would, however, not be feasible as the consolidated retailer would optimally buy all goods at the lower of the two prices.

(iii) the winning supplier offers m^s such that the retailer is indifferent between the two offers.[30]

In the absence of single-sourcing, the results are the same as in the pre-merger case of two separate negotiations.

Comparing the retailer's profit under single sourcing with that under separate negotiations, it is easily seen that single sourcing is strictly better for the retailer. For $\beta = 1$, if supplier A wins the contract, the retailer's profit under single sourcing is equal to $\Pi_a(\theta^B) + \Pi_b(\theta^B)$, which under Assumption 1 strictly exceeds its profit without single sourcing, $\Pi_a(\theta^B) + \Pi_b(\theta^A)$.

The following Proposition now summarises our results for the case with linear contracts.

PROPOSITION 5. *Suppose that suppliers compete in linear contracts and that good A (B) is sufficiently attractive at outlet b (a) to constrain the offer of the other supplier. Then the retailers' gains from a merger are identical to those with efficient contracting (Proposition 1).*

With efficient contracts, the only welfare effect of a merger is its impact on product availability and the choice of product characteristics. With linear contracts, i.e., where contracts determine only a uniform purchase price, we obtain a new effect. Increasing buyer power and shifting profits to the retailer reduces double-marginalisation. This trade-off applies also if we endogenise product characteristics as was done in Section 2: a higher μ makes suppliers more homogenous, which intensifies competition and further reduces double-marginalisation.

The trade-off between the loss of product variety and a reduction in double-marginalisation complicates the welfare analysis with linear contracts. In what follows, we confine ourselves to discussing the implications on welfare in our previously introduced linear example.

3.2. *Example*

With linear demand $D = 1 - d - p$, and constant purchase price $m < 1 - d$, a retailer optimally chooses the price $p = (1 + m - d)/2$. This generates profit $(1 - m - d)^2/4$. Recall that we chose the parameters $c = 0$, $z = 5$, and $\beta = 1$. Recall next that we substituted $d = \theta^2/z$ for $r = a$ and $d = (1 - \theta)^2/z$ for $r = b$. Retailers' profits at the two outlets from an offer m are then

$$\Pi_a = \frac{1}{4}\left(1 - \frac{1}{5}\theta^2 - m\right), \tag{21}$$

$$\Pi_b = \frac{1}{4}\left[1 - \frac{1}{5}(1 - \theta)^2 - m\right]. \tag{22}$$

Finally, recall that we confined ourselves to the case where θ can only be chosen from a finite set $\theta \in \Theta = \left\{\hat{\theta}_a, \theta^*, \hat{\theta}_b\right\}$, where $\hat{\theta}_a = 0, \hat{\theta}_b = 1$ and $\theta^* = 0.2$.

Suppose again that A wins under single sourcing, which implies from Proposition 4 that $\theta^B = \hat{\theta}_b = 1$. If no merger takes place, A's offer to a must match B's offer of $m_a^B = c = 0$. Using the expressions in (21), we can solve for $m_a^A = \left[1 - (\theta^A)^2\right]/5$,

[30] That is, if supplier A wins the single-sourcing contract we have that m^A is the highest value solving $\max_p[(p - m^A)D_a(\theta^A, p)] + \max_p[(p - m^A)D_b(\theta^A, p)] = \Pi_a(\theta^B) + \Pi_b(\theta^B)$.

which yields $m_a^A = 0.2$ in case $\theta^A = 0$ and $m_a^A = 0.192$ in case $\theta^A = 0.2$. For B's offer to b we get $m_b^B = 0.2$ if $\theta^A = 0$ and $m_b^B = 0.128$ if $\theta^A = 0.2$. Intuitively, m_b^B is lower if A's product becomes more attractive to b.

It is instructive to pause here and consider how the choice of θ^A affects expected welfare if retailers stay separate. With efficient contracts, expected welfare was lower with the less differentiated variant $\theta^A = \theta^* = 0.2$. With linear contracts, however, there is a countervailing effect in that less differentiation improves the retailers' outside option and thus lowers the purchasing price. This reduces the double-marginalisation problem. Straightforward calculations establish that this countervailing effect dominates and thus expected welfare is higher for $\theta^A = \theta^* = 0.2$.[31]

Now consider the case in which a merger occurs. With single sourcing, if $\theta^A = 0$ and $\theta^B = 1$, then symmetry implies that suppliers realise zero profits and make the offers $m^A = m^B = 0$. But if $\theta^A = 0.2$ and $\theta^B = 1$, then supplier A chooses m^A such that the retailer earns no more than its outside option $\Pi_a(\theta^B) + \Pi_b(\theta^B)$, which, by substitution into (21), yields $m^A = 0.0285$.

Putting these results together and solving for supplier A's profit-maximising choice of θ^A as a function of μ, we obtain the following:

Results for the linear example with linear contracts:

(*i*) *If* $\mu > 11.1\%$, *supplier A chooses the less differentiated product variant* $\theta^A = \theta^* = 0.2$. *Otherwise, supplier A chooses the more differentiated product variant* $\theta^A = \hat{\theta}_a = 0$.

(*ii*) *Expected welfare is strictly decreasing in* μ *over both regimes, i.e., for* $\mu < 11.1\%$ *and* $\mu > 11.1\%$. *At* $\mu = 11.1\%$, *where supplier A switches to* $\theta^A = \theta^* = 0.2$, *expected welfare jumps up. This is also the highest feasible value for expected welfare.*

(See the Appendix for the complete calculations.)

We can now compare the outcomes with efficient and linear contracts. In the case of efficient contracts, a very stringent merger policy ($\mu = 0$) is best. In contrast, with linear contracts, expected welfare is maximal at an interior choice $\mu = 11.1\%$. In fact, we can show that *ex-post* welfare would be maximal if $\theta^A = \theta^* = 0.2$ and no merger took place.[32] However, to induce the supplier to choose a less differentiated product variant, it is necessary to have $\mu > 0$.

Though our comparison is clearly confined to a very specific example, it highlights an important question for analysing the welfare implications of buyer power. Should we reasonably assume that contracts are sufficiently complex to allow for efficient contracting or should we assume that contracts are relatively incomplete and simple, with linear contracts as a good approximation? In the first case, shifting rents to retailers has no direct impact on output and welfare, whereas in the second instance it increases output and welfare. An answer to this question, while being key for the analysis of welfare, may depend on the specific circumstances.

[31] We obtain for outlet b the quantity $1/2 - (1 - \theta^A)^2/10$ and welfare $3/8 - (3/40)(1 - \theta^A)^2$, while we obtain for outlet a the quantity $2/5 - (\theta^A)^2/10$ and welfare $(8/25) - (7/50)(\theta^A)^2 + (3/200)(\theta^A)^4$. Summing up, total welfare equals $(3/25)(\theta^A) - (11/50)(\theta^A)^2 + (1/50)(\theta^A)^3 + (1/100)(\theta^A)^4 + (16/25)$, which yields 0.655 for $\theta^A = 0.2$ and 0.640 for $\theta = 0$.

[32] The total welfare ordering is as follows: welfare equals 0.655 with no merger and $\theta^A = \theta^* = 0.2$, 0.641 with a merger and $\theta^A = \theta^* = 0.2$, 0.640 without a merger and $\theta^A = \hat{\theta}_a = 0$, and 0.615 with a merger and $\theta^A = \hat{\theta}_a = 0$.

4. Discussion

4.1. *Buyer Groups*

A merger enhances buyer power by allowing the consolidated retailer to adopt a single-sourcing policy. In principle, the benefits from single sourcing could also be achieved if separate retailers formed an alliance and agreed to bundle their purchases. In fact, as reported in Dobson (2002), buyer groups (or buyer alliances) have gained considerable importance. Buyer groups are typically alleged to be advantageous for the member retailers because of the additional 'clout' they confer on the retailers when contracts are negotiated and because, by pooling all purchases, they allow quantity discounts to be realised on a scale that would not be possible otherwise.

Our analysis suggests that a policy of actually delisting suppliers (as opposed to merely threatening to delist suppliers) may be one way of operationalising the alleged 'clout' that is conferred on the retailers. However, there are some hurdles to overcome that are not necessarily present in the case of a single large retailer. Thus, while our analysis is in principle applicable to buyer groups, we feel that its applicability and welfare implications will likely be more pronounced in the case of an outright merger. As we argue next in more detail, this follows as buyer groups are likely to have less scope than a single large retailer to prune their supplier base.

Suppose the two retailers in our model could form a buyer group in order to bundle their purchases. Of course, this only makes a difference if they also decide to adopt single sourcing. By Assumption 1, this implies that a suboptimal good is sold at one outlet. Without side payments between the two retailers, it may thus be difficult to ensure that the winning supplier's offer is beneficial to both retailers.[33] What is more, though this is admittedly outside our model, limited information about each others' profits may render even an agreement with side payments difficult. For instance, while it may be known that good B provides a better fit for outlet b, the extent to which good A reduces sales and profits at outlet b may be b's private information. Likewise, retailer b may not know what profits retailer a can make with the two different goods. As is well known from the bargaining and mechanism-design literature, such two-sided private information typically leads to failure of agreement, at least with positive probability.

4.2 *Suppliers' Costs*

Different plant locations and the presence of non-negligible transportation costs provide an alternative source of differentiation for suppliers. We now illustrate that, in our model, transportation costs would fulfill the same role as differences in consumers' preferences.[34]

[33] Repeated interaction could provide an (imperfect) substitute for side payments.

[34] In the working paper version of this article, we showed that our main results are also robust to the introduction of (strictly) convex costs. In fact, we showed that convex costs *alone* can make single sourcing optimal for the consolidated retailer, while obviously leading to production inefficiencies (the proof is available on request). Consequently, our results hold *a fortiori* if Assumption 1 holds while costs are strictly convex.

Suppose that consumer demand at both outlets is characterised by the same function $D(p, \theta)$ and that both goods have the same characteristics $\theta^A = \theta^B = \theta$.[35] Producing and shipping an additional unit of good s to outlet r comes now at the constant marginal costs $c + t_r^s$. If A's factory is closer to outlet a than to outlet b, then we have that $t_a^A < t_b^A$. And if an analogous relation holds for B, i.e., if $t_b^B < t_a^B$, and if no supplier has lower costs in supplying both outlets, then Assumption 1 is still satisfied. But this was all that we needed to derive our main results.[36]

5. Conclusion

This article analyses the impact of retail mergers on product variety. We consider the effects of a merger on the consolidated retailer's choice of which products to carry and on the upstream suppliers' choices of how differentiated to make their products. Thus, in contrast to a more standard focus on the increase in market power that may arise at the retail level, we provide a way of understanding merger effects in terms of an endogenous adjustment in the supply chain.

We find that a merger may create an incentive for the retailer to reduce product variety by consolidating its supplier base. When the products of different suppliers are valued differently by buyers across retail markets (e.g., brands across regions or countries), a supplier whose product is well suited to a particular market is in a stronger bargaining position with a firm who is concentrated in that retail market. This allows the suppliers with the upper hand in each market to capture rents from the individual retailers. To reduce these rents, after a merger, (we conceptualise a merger as extending the range of outlets controlled by a single retailer) a retailer can shift the terms of supplier competition by forcing them to compete to become the exclusive supplier (i.e., the retailer offers only one product version that is common across stores). While it is true that this outcome will sacrifice the additional joint surplus that is generated when product variety is tailored to preferences in each market, this additional surplus was previously captured primarily by the well-matched suppliers and not the retailers. The 'sole source' competition forces suppliers to bid based on the surplus their product generates across all retail markets. This typically increases competitive pressure and works to the advantage of the consolidated retailer whenever the suppliers were previously in a strong bargaining position.

As suppliers anticipate the single-sourcing policy of a consolidated retailer, the likelihood of a retail merger influences their optimal choice of product characteristics. In particular, suppliers will be induced to produce less differentiated products, which further reduces product variety. If negotiations are efficient, we find that, as retail mergers become more likely, e.g., due to a more lenient merger policy, expected industry profits and potentially also welfare are reduced (because of the reduction in product variety). With linear contracts, however, there may be a countervailing effect as the more powerful retailer passes on lower input prices to final consumers.

[35] To rule out benefits from single sourcing due to differences in demand it is, of course, sufficient that either of the two conditions holds, i.e., that local demand is homogenous or that goods are not differentiated.

[36] If supplier A has both lower transportation costs *and* a product that is more suitable for outlet a than supplier B – and if analogous conditions apply for B – Assumption 1, obviously, continues to hold *a fortiori*.

Our model provides a parsimonious theory of the origins and (welfare) consequences of buyer power. It emphasises the role of single sourcing, both as an (off-the-equilibrium) threat and as an active (on-the-equilibrium) strategy to exert buyer power. The profitability of a retail merger and of a subsequent single-sourcing strategy depends crucially on differences in the retailers' *previous* supplier bases and, thereby, on differences in consumer preferences at their respective outlets. This makes our theory of buyer power and retail mergers particularly applicable to cross-border mergers, where standard explanations based on horizontal merger theory seem to be less appropriate and where competition authorities often see no issues arising.

Looking only at the *downstream* market, mergers between firms operating in 'overlapping' markets should have more serious consequences for price strategies and welfare. In retail mergers, stipulating the divestiture of outlets in overlapping markets is a common way to deal with these concerns. In contrast, looking at the *upstream* market, our analysis suggests that mergers in non-overlapping markets may provide more scope for firms to lever up their position *vis-à-vis* their suppliers. As we show, this may have serious consequences for product variety and welfare.

There are some obvious ways to enrich the simple model studied in this article. First, to obtain a descriptive theory of retail mergers, we would like to have a countervailing force that makes it sometimes unprofitable for retailers to merge. In the current model, a merger between retailers is always weakly profitable. Second, to study overall industry dynamics one should also allow for mergers between suppliers. These extensions are beyond the scope of this article.

Appendix: Omitted Calculations for the Example with Linear Contracts

We first analyse the optimal choice of θ^A, given that $\theta^B = 1$. Suppose $\theta^A = \hat{\theta}_a = 0$. If a merger takes place (and assuming there is single sourcing), supplier A realises zero profits. If no merger takes place, we obtain $m_a^A = \left[1 - (\theta^A)^2\right]/5 = 0.2$. As retailer a chooses the output $x_a = \left[1 - (\theta^A)^2/5 - m_a^A\right]/2 = 0.4$, supplier A realises profit $m_a^A x_a = 0.8$. Thus, if supplier A chooses $\theta^A = 0$, its expected profit is $(1 - \mu)0.8$. Suppose next that $\theta^A = \theta^* = 0.2$. If a merger takes place (and assuming there is single sourcing), we obtain $m^A = 0.0285$. Moreover, the consolidated retailer will choose $x_a = [1 - (\theta^A)^2/5 - m^A]/2 = 0.482$ for outlet a and $x_b = [1 - (1 - \theta^A)^2/5 - m^A]/2 = 0.422$ for outlet b. Hence, supplier A's profits are $(x_a + x_b)\, m^A = 0.0258$. If no merger takes place and $\theta^A = 0.2$, we obtain $m^A = [1 - (1/5)^2]/5 = 0.192$ and the output $x_a = [1 - (\theta^A)^2/5 - m^A]/2 = 0.400$ to outlet a, yielding the profit $x_a m^A = 0.077$. Thus, if supplier A chooses $\theta^A = 0.2$, its expected profit is $0.0258\mu + (1 - \mu)0.0768$. Comparing profits for $\theta^A = \hat{\theta}_a = 0$ and $\theta^A = \theta^* = 0.2$, we obtain that supplier A prefers $\hat{\theta}_a$ if and only if $\mu < 0.111$.

We calculate expected welfare for the two scenarios next. Suppose first $\mu > 0.111$, implying $\theta^A = 0.2$. From previous results we know that total welfare equals 0.655 in the case of no merger and 0.641 in the case of a merger. This yields the *ex ante* welfare $0.655 - 0.0140\mu$. Proceeding likewise for $\mu < 0.111$ and $\theta^A = 0$, we obtain the expected welfare of $0.640 - 0.0250\mu$. Finally, substitution shows that welfare is maximised at the lowest feasible value μ at which $\theta^A = 0.2$.

London School of Economics
University of Rochester

Submitted: 8 July 2004
Accepted: 3 August 2005

References

Adams, W. and Yellen, J. (1976). 'Commodity bundling and the burden of monopoly', *Quarterly Journal of Economics*, vol. 90, pp. 475–98.

Anton, J. and Yao, D. (1989). 'Split awards, procurement and innovation', *Rand Journal of Economics*, vol. 20, pp. 538–52.

Bernheim, D. and Whinston, M. (1998). 'Exclusive dealing', *Journal of Political Economy*, vol. 106, pp. 64–103.

Binmore, K., Rubinstein, A. and Wolinsky, A. (1986). 'The Nash bargaining solution in economic modeling', *Rand Journal of Economics*, vol. 17, pp. 176–88.

Chipty, T. and Snyder, C.M. (1999). 'The role of firm size in bilateral bargaining: a study of the cable television industry', *Review of Economics and Statistics*, vol. 81, pp. 326–40.

Clarke, R., Davies, S., Dobson, P. and Waterson, M. (2002). *Buyer Power and Competition in European Food Retailing*, Cheltenham: Edward Elgar.

Competition Commission (2000). *Supermarkets: a Report on the Supply of Groceries from Multiple Stores in the United Kingdom*, Report Cm-4842.

Competition Commission (2003). *Safeway plc and Asda Group Limited (owned by Wal-Mart Stores Inc); Wm Morrison Supermarkets plc; J Sainsbury's plc; and Tesco plc: A Report on the Mergers in Contemplation.*

Dana, J. (2004). 'Buyer groups as strategic commitments', mimeo, Northwestern University.

Dobson, P. (2002). 'Retailer buyer power in European markets: lessons from grocery supply', Loughborough University Business School, Working Paper No. 2002-1.

Dobson, P. and Waterson, M. (1997). 'Countervailing power and consumer prices', ECONOMIC JOURNAL, vol. 107, pp. 418–30.

Dobson, P. and Waterson, M. (1999). 'Retailer power: recent developments and policy implications', *Economic Policy*, vol. 28, pp. 133–64.

Dobson Consulting (1999). *Buyer Power and its Impact on Competition in the Food Retail Distribution Sector of the European Union*, Report produced for the European Commission, DG IV, Brussels.

Federal Trade Commission (2001). *Report on the Federal Trade Commission Workshop on Slotting Allowances and Other Marketing Practices in the Grocery Industry*, Report by the Federal Trade Commission Staff, Washington, D.C.

Fumagalli, C. and Motta, M. (2000). 'Buyers' coordination', mimeo, University of Bocconi, CEPR Discussion Paper No. 2908.

Gabrielsen, T. and Sorgard, L. (1999). 'Discount chains and brand policy', *Scandinavian Journal of Economics*, vol. 101, pp. 127–42.

Horn, H. and Wolinsky, A. (1988). 'Bilateral monopolies and incentives for merger', *Rand Journal of Economics*, vol. 19, pp. 408–19.

Inderst, R. and Wey, C. (2002). 'Buyer power and supplier incentives', London School of Economics, CEPR Discussion Paper No. 3547.

Katz, M.L. (1987). 'The welfare effects of third degree price discrimination in intermediate goods markets', *American Economic Review*, vol. 77, pp. 154–67.

Marx, L. and Shaffer, G. (2004). 'Slotting allowances and scarce shelf space', mimeo, University of Rochester.

Mazzarotto, N. (2001). 'Competition policy towards retailers: size, seller market power and buyer power', University of East Anglia Centre for Competition & Policy, CCR Discussion Paper No. 01–4.

McAfee, R., McMillan, J. and Whinston, M. (1989). 'Multiproduct monopoly, commodity bundling, and correlation of values', *Quarterly Journal of Economics*, vol. 29, pp. 371–88.

Normann, H., Ruffle, B and Snyder, C. (2003). 'Do buyer-size discounts depend on the curvature of the surplus function? experimental tests of bargaining models', mimeo, George Washington University.

O'Brien, D.P. and Shaffer, G. (1997). 'Nonlinear supply contracts, exclusive dealing, and equilibrium market foreclosure', *Journal of Economics and Management Strategy*, vol. 6, pp. 755–85.

OECD (1999). *Buying Power of Multiproduct Retailers,* Series Roundtables on Competition Policy DAFFE/CLP(99)21, Paris: OECD.

Palfrey T, (1983). 'Bundling decisions by a multiproduct monopolist with incomplete information', *Econometrica*, vol. 51, pp. 463–84.

Perry, M. and Sakovics, J. (2003). 'Auctions for split-award contracts', *Journal of Industrial Economics*, vol. 51, pp. 215–42.

Rey, P. (2003). 'Retailer buyer power and competition policy', in (B. Hawk, ed.) *Annual Proceedings of the Fordham Corporate Law Institute*, pp. 487–511, New York: Juris Publishing.

Riordan, M. and Sappington, D (1989). 'Second sourcing', *Rand Journal of Economics*, vol. 20, pp. 41–57

Snyder, C.M. (1996). 'A dynamic theory of countervailing power', *Rand Journal of Economics*, vol. 27, pp. 747–69.

von Ungern-Sternberg, T. (1996). 'Countervailing power revisited', *International Journal of Industrial Organization*, vol. 14, pp. 507–20.