Retail Mergers, Buyer Power and Product Variety*

Roman Inderst† Greg Shaffer‡

August 2004

Abstract

This paper analyses the impact of retail mergers on product variety. We show that, following a merger, a retailer may want to enhance its buyer power *vis a vis* suppliers by committing to a "single-sourcing" purchasing strategy. As we argue, the retailer's gains from single sourcing may be more pronounced in the case of *cross-border mergers.* 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 reduces aggregate profits and, under standard assumptions, also 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.

Keywords: Buyer Power; Retailing; Horizontal Mergers.

---

*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, and Norwich University.

†London School of Economics, INSEAD and CEPR. Address for Correspondence: Department of Economics and Department of Accounting and Finance, London School of Economics, Houghton Street, London WC2A 2AE. Email: r.inderst@lse.ac.uk.

‡University of Rochester. Address for Correspondence: Simon School of Business, University of Rochester, Rochester, NY 14627. Email: shaffer@simon.rochester.edu.

1 Introduction

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 neighborhood, 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 scrutinized by antitrust authorities, who have become increasingly concerned about retailers' growing buyer power *vis-a-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 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 paper 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 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 be content with capturing 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 aggregate profits and, under standard assumptions,

[1]See, for instance, the OECD (1999) report on the buying power of multiproduct retailers as well as various reports on concentration in the U.S. 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 (2000). One issue is whether the retailers' growing buyer power has negative consequences for product quality, innovation, and variety.

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 paper apart from most of the extant literature on buyer power, where delisting goods is only an off-equilibrium threat. (The literature is reviewed below.)

As we discuss in more detail below, 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 regards 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 two predictions that consolidated retailers can obtain more favorable terms of supply and that they may also 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 paper contributes to the growing literature on buyer power. According to the extant literature, larger retailers can obtain more favorable terms and conditions as their orders break collusion between suppliers (Snyder, 1996) or as they can threaten to integrate backwards or sponsor new entry in the upstream industry (Katz, 1987; Fumagalli and Motta, 2000). Also, if buyers compete downstream, a supplier's threat to sell to competing buyers may become less

[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, Davies, Dobson and Waterson (2002).

valuable after a downstream merger (von Ungern-Sternberg, 1996; Dobson and Waterson, 1997; Mazzarotto, 2003).[6] If suppliers have convex production costs, Chipty and Snyder (1999) and Inderst and Wey (2002) show that a larger buyer can negotiate lower prices. This is the case as smaller buyers must negotiate more "on the margin", where average unit costs are higher.[7]

The rest of this paper is organized as follows. Section 2 contains our main results on how a consolidated retailer can use its newly acquired buyer power. Section 3 extends the model by endogenizing product characteristics. Section 4 considers linear contracts. Section 5 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 6 offers concluding remarks.

2 The Main Model

2.1 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 3, 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 the retailer at each outlet to stock only one of the two goods, A or B.

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 can be jointly realized by the retailer and supplier when a good with characteristics θ is sold at outlet r. Note that $\Pi_r(\theta)$ would be realized by an integrated firm that controls both production and final sales.[9] Until Section 3, where θ^s is

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

[7]For experimental results on this, see Normann, Ruffle and Snyder (2003).

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

[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.

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 realized 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". Assumption 1 is tantamount to assuming that consumers in market a have stronger preferences for good A, all else equal, while consumers in market b have stronger preferences for good B.

There are many reasons why Assumption 1 may hold in practice. First, consumers may predictably differ in tastes and preferences as outlets a and b may be located in different regions or even different countries. Still 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. Again, this interpretation seems to be particularly suitable in case the two outlets are located in separate countries. In Section 3, we show how Assumption 1 arises endogenously if suppliers optimally choose their product characteristics. Section 5.2 shows that our main results extend to a model where products are homogeneous but where suppliers are differentiated in how close their factories are located to the different outlets. That is, the resulting differences in (per-unit) transportation costs generate the same outcome as product differentiation.

2.2 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.

After contracts have been determined, retailers set prices at their respective outlets. We place no restrictions on the set of feasible contracts. Optimally, the two contracting parties will thus choose a contract that avoids double-marginalisation and maximizes their joint profit. A simple contract that rules out double marginalization 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, i.e., two sales representatives or managers respectively, who act independently but in the interest of the respective firm. Negotiations proceed simultaneously,

However, to abstract from the standard monopolization 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.

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 realized net surplus. Our assumption of a fixed division of the realized 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 $\beta = 1/2$.[11]

We now proceed as follows. In section 2.3, we present the solutions for the two benchmark cases where $\beta = 1$ and $\beta = 0$. In section 2.4, we consider the case of general values for β.

2.3 Analysis for Two Benchmarks

We suppose first that suppliers have all the bargaining power and thus 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, each supply contract maximizes the respective joint surplus. That is, the maximum feasible profit $\Pi_a(\theta^A)$ is realized at outlet a, while $\Pi_b(\theta^B)$ is realized at outlet b. Finally, as all the bargaining power lies with the supplier, each retailer can only extract the value of its respective outside option. Formally, retailer a realizes the profit $\Pi_a(\theta^B)$ and retailer b realizes 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

[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]For a non-cooperative foundation of the asymmetric Nash bargaining solution, see Binmore, Rubinstein, and Wolinksy (1986). We combine both non-cooperative and cooperative concepts, which is common in the literature and allows us to obtain a parsimonious model of negotiations.

[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, $\Pi_a(\theta^A) - \Pi_a(\theta^B)$.

the case of supplier A and $\Pi_b(\theta^B) - \Pi_b(\theta^A)$ in the case of supplier 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 are equal to 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 profit if good A is chosen and total industry profit 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, realizes 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 summarize our results.

Proposition 1. *If suppliers have all the bargaining power ($\beta = 1$), retailers' profits strictly increase after a merger if the consolidated retailer chooses a single-sourcing strategy. In case 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 be surprising given that single sourcing reduces total industry profit (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 profit is equal to $\Pi_b(\theta^B) - \Pi_b(\theta^A)$. Note that this equals supplier B's loss under single sourcing (which is intuitive because in the absence of single sourcing, supplier B earns its good's full added value). 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 exactly the increase realized 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 retailer to extract a larger share of a smaller total industry profit as it makes suppliers less differentiated. The resulting averaging of each good's added value causes competition for the retailer's patronage to be more vigorous, which in turns facilitates the transfer of surplus from the winning supplier to the retailer. The idea that a retailer can increase its share of industry profit by smoothing out the value added of each supplier's good is analogous to the reasons why bundling can be optimal for a monopolist (e.g., Adams and Yellen (1976), Palfrey (1983), and McAfee, McMillan, and Whinston (1989)).[15] However, what is missing from the bundling literature is the analogue to the loss in industry profit that is induced when a good is excluded from distribution. For the benchmark case of $\beta = 1$, this loss in industry profit is fully offset by the decrease in the losing supplier's profit.

For the merger and the subsequent single sourcing to be profitable, it is essential that the

[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 3, where we specify how θ affects demand and industry profits.

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

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 profits.

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. This makes the theory particularly applicable to cross-border retail mergers. The retailers' gains under the new purchasing strategy, e.g., $\Pi_b(\theta^B) - \Pi_b(\theta^A)$ if supplier A is subsequently chosen, are also higher the more differentiated are the suppliers' goods. While the loss from delisting one of the goods increases in their differentiation, the retailers' gains also increase - 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: $\beta = 0$. Here, retailers have all the bargaining power and can extract all of the total industry profit.[16] It is thus immediate that a merger can not increase the retailers' joint profits. What is more, we know that a consolidated retailer would now strictly prefer *not* to choose a single-sourcing strategy. 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. *If retailers have all the bargaining power ($\beta = 0$), there is no gain from merging, as the consolidated retailer would be strictly worse off under a single-sourcing strategy.*

2.4 The General Case

With $0 < \beta < 1$ each of the two parties to a deal can extract a strictly positive share of its added value. Hence, if retailer a negotiates with the two suppliers, its profits are now the sum of its outside option $\Pi_a(\theta^B)$ *plus* the fraction $1 - \beta$ of the added value realized with supplier A, i.e., of the difference $\Pi_a(\theta^A) - \Pi_a(\theta^B)$. Summing up, the two retailers' joint profits are now

$$\Pi_a(\theta^B) + (1-\beta)\Pi_a(\theta^A) - \Pi_a(\theta^B) + \Pi_b(\theta^A) + (1-\beta)\Pi_b(\theta^B) - \Pi_b(\theta^A)$$

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

while suppliers' joint profits equal

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

[16]Incidentally, given that their marginal costs of production are constant, it does not matter whether each supplier also sells to other outlets in markets that do not have overlap with those where retailers a and b are located. The suppliers' gains from these contracts are unaffected by their dealings with a and b.

In the case of single sourcing, if (3) holds, the consolidated retailer's profit equals its outside option $\Pi_a(\theta^B)+\Pi_b(\theta^B)$ plus the fraction $1-\beta$ of the added value of good 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 retailer's profit is now

$$\Pi_a(\theta^B)+\Pi_b(\theta^B)+(1-\beta)\left[\left(\Pi_a(\theta^A)+\Pi_b(\theta^A)\right)-\left(\Pi_a(\theta^B)+\Pi_b(\theta^B)\right)\right]$$
$$=(1-\beta)\left[\Pi_a(\theta^A)+\Pi_b(\theta^A)\right]+\beta\left[\Pi_a(\theta^B)+\Pi_b(\theta^B)\right], \tag{8}$$

while the winning supplier, A, earns a profit equal to

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

What are now the gains and losses from a single-sourcing strategy? The reduction in total industry profit 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

$$(2\beta-1)\left[\Pi_b(\theta^B)-\Pi_b(\theta^A)\right]. \tag{10}$$

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

Proposition 3. *For the case of general β, the consolidated retailer is strictly better off under the single-sourcing strategy if and only if $\beta>1/2$, i.e., if and only if suppliers have sufficient bargaining power. Only in this case does a merger make retailers strictly better off.*

Proposition 3 is intuitive given the results in Propositions 1 and 2. The choice of a single-sourcing strategy creates the following trade-off. Single sourcing makes the total "pie" smaller as it reduces industry profit. The reduction in industry profit 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 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 retailer prefers a larger share of a smaller overall pie) and thus a single sourcing strategy 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 endogenize

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, i.e., potentially 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 endogenize 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 reorganization 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 vigorously "prune" the supplier base of the two merging retailers.

3 Endogenous Variety

3.1 Extending the Model

We now endogenize the suppliers' choice 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, $t = 2$, retailers decide whether or not to merge. The rest of the game is then as described above. That is, in period three, $t = 3$, retailers choose their purchasing strategy, i.e., whether or not to commit to a single-sourcing strategy. (This is only a non-trivial choice for a consolidated retailer.) In the fourth period, $t = 4$, retailers and suppliers negotiate under the chosen purchasing strategy. In the final period, $t = 5$, retailers set prices for final consumers, goods are supplied, and payoffs are realized.

[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 newly introduced stages $t = 1, 2$ deserve some comments. Consider first the choice of product characteristics at $t = 1$.[18] Recall that $\Pi_r(\theta)$ denotes the maximum feasible profits that can be realized 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 that $\Pi_r(\theta) > 0$ for some θ. We also assume that $d\Pi_b(\theta)/d\theta > d\Pi_a(\theta)/d\theta$ holds whenever both $\Pi_r(\theta)$ are differentiable and at least one is strictly positive. Define $\widehat{\theta}_r := \arg\max_\theta \Pi_r(\theta)$ and note that $\widehat{\theta}_a < \widehat{\theta}_b$. It also proves useful to assume that both $\Pi_a(\theta)$ and $\Pi_b(\theta)$ are strictly positive over the "relevant" range $\theta \in [\widehat{\theta}_a, \widehat{\theta}_b]$.

We specify that a retail merger at $t = 2$ happens with some exogenous probability μ. The choices $\mu = 0$ and $\mu = 1$ correspond to cases where the merger occurs never or always. Note that in our model a merger is always weakly profitable for retailers. But a merger may not always be possible. For instance, the owners or the management of a retailer may not be prepared to relinquish control. Likewise, 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]

3.2 Analysis

It is helpful to consider first the case where suppliers anticipate that no merger will occur, $\mu = 0$. In any pure-strategy equilibrium in which $\mu = 0$, each supplier maximizes its profit by focusing on only one outlet; one supplier optimally chooses $\widehat{\theta}_a$ and the other supplier chooses $\widehat{\theta}_b$. Thus, when $\mu = 0$, suppliers choose the product characteristics that maximize 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 re-position its product to balance more adequately the different consumer preferences at the two outlets. Define[20]

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

where the expression inside the brackets is a weighted average of the maximum feasible profit

4 below, as μ increases, $\widetilde{\theta}_r$ increases from a lower bound of $\widehat{\theta}_a$ when $r = a$ and decreases from an upper bound of $\widehat{\theta}_b$ when $r = b$. An immediate implication of (11) is that if $\theta^A = \widetilde{\theta}_a$ and $\theta^B = \widehat{\theta}_b$, then the construction of $\widetilde{\theta}_a$ ensures that total *expected* industry profit is maximized, provided good A is chosen by the consolidated retailer under single sourcing following a merger. The interpretation of $\widetilde{\theta}_b$ as maximizing the total *expected* industry profit is analogous if $\theta^B = \widetilde{\theta}_b$ and $\theta^A = \widehat{\theta}_a$, provided good B is chosen by the consolidated retailer 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 $\widetilde{\theta}_r$ 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 $(\widetilde{\theta}_a, \widehat{\theta}_b)$ or $(\widetilde{\theta}_b, \widehat{\theta}_a)$. To see this, we argue first that some supplier must choose $\theta^s \in \left\{\widehat{\theta}_a, \widehat{\theta}_b\right\}$. Suppose this were not the case. Then if one of the suppliers, say A, wins the single-sourcing contract with probability one following a merger, the other supplier, B, can, by $\mu < 1$, profitably deviate to some $\theta^B \in \left\{\widehat{\theta}_a, \widehat{\theta}_b\right\}$, which maximizes its profit in the case where no merger occurs. Hence, it must be that both suppliers are chosen with positive probability under single sourcing. As this implies that $\Pi_a(\theta^A) + \Pi_b(\theta^A)$ equals $\Pi_a(\theta^B) + \Pi_b(\theta^B)$, both suppliers realize zero profit if a merger takes place. But, it is then again strictly profitable for at least one supplier to deviate and choose $\theta^s \in \left\{\widehat{\theta}_a, \widehat{\theta}_b\right\}$, which maximizes its profit in the case where no merger occurs. This proves that some supplier must choose $\theta^s \in \left\{\widehat{\theta}_a, \widehat{\theta}_b\right\}$.

We now argue that some supplier must choose $\theta^s \in \left\{\widetilde{\theta}_a, \widetilde{\theta}_b\right\}$. Once again we offer a proof by contradiction. Suppose neither supplier chooses $\theta^s \in \left\{\widetilde{\theta}_a, \widetilde{\theta}_b\right\}$, and consider the two possibilities. Either one of the suppliers wins the contract with probability one under single sourcing, or both suppliers are chosen with positive probability. If the latter is the case, then the supplier choosing $\theta^s \in \left\{\widehat{\theta}_a, \widehat{\theta}_b\right\}$ can profitably deviate by adjusting θ^s marginally such that it is now the strictly preferred supplier under single sourcing (this allows it to earn higher profit in the event a merger occurs while essentially sacrificing no profit in the event no merger takes place). If the former is the case, and the winning supplier sells to outlet r if no merger takes place, then the winning supplier can profitably deviate by moving θ^s in the direction of $\widetilde{\theta}_r$, which maximizes the expected payoff of the winning supplier.[21] It follows that in both cases there is a profitable deviation by

[21] For example, suppose that supplier B chooses $\theta^s \in \left\{\widehat{\theta}_a, \widehat{\theta}_b\right\}$, and without loss of generality, let $\theta^B = \widehat{\theta}_b$, so that good B will be sold at outlet b if no merger takes place. Then supplier A's expected profit if its good will be sold at outlet a if no merger takes place and at both outlets (given single sourcing) following a merger is

$$\mu\beta\left[\left(\Pi_a(\theta^A) + \Pi_b(\theta^A)\right) - \left(\Pi_a(\widehat{\theta}_b) + \Pi_b(\widehat{\theta}_b)\right)\right] + (1-\mu)\beta\left[\Pi_a(\theta^A) - \Pi_a(\widehat{\theta}_b)\right] \quad (12)$$

one of the suppliers. This proves that some supplier must choose $\theta^s \in \left\{\widetilde{\theta}_a, \widetilde{\theta}_b\right\}$.

Our arguments imply that in any pure-strategy equilibrium, one supplier will maximize 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 $\theta^s = \widetilde{\theta}_r$ (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 $\widetilde{\theta}_a$ while the other chooses $\widehat{\theta}_b$, or one supplier chooses $\widetilde{\theta}_b$ while the other chooses $\widehat{\theta}_a$. Moreover, as μ increases, goods become continuously less differentiated as $\widetilde{\theta}_a \geq \widehat{\theta}_a$ is strictly increasing in μ and as $\widetilde{\theta}_b \leq \widehat{\theta}_b$ is strictly decreasing in μ.*

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

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

Implicit differentiation of (13), while using $d\Pi_a(\theta)/d\theta < d\Pi_b(\theta)/d\theta$ and the strict concavity of $\Pi_r(\theta)$, shows that $d\widetilde{\theta}_a/d\mu > 0$. The proof for $\widetilde{\theta}_b$ is analogous. **Q.E.D.**

Proposition 4 warrants several comments. We consider first the role of μ. For $\mu = 0$, where $\widetilde{\theta}_a = \widehat{\theta}_a$ and $\widetilde{\theta}_b = \widehat{\theta}_b$, the two chosen characteristics provide the best fits for the respective outlets. For $0 < \mu < 1$, one supplier focuses on one outlet, e.g., supplier B chooses $\theta^B = \widehat{\theta}_b$, while the other supplier optimally chooses more "average" characteristics. Proposition 4 does not cover the case with $\mu = 1$, where a merger occurs with probability one. For $\mu = 1$, it is, however, straightforward to establish that in any pure-strategy equilibrium at $t = 1$ a supplier who subsequently wins the single-sourcing contract must choose the unique characteristics that maximize total industry profits $\Pi_a(\theta) + \Pi_b(\theta)$. (Incidentally, for $\mu = 1$, this is equal to $\widetilde{\theta}_a = \widetilde{\theta}_b$.)

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

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

i.e., if $\widetilde{\theta}_a$ and $\widetilde{\theta}_b$ were to be chosen, it would be the supplier with characteristics $\widetilde{\theta}_a$ who would win the single-sourcing contract. Then, it is easy to show that $\theta^A = \widetilde{\theta}_a$ is a best response to $\theta^B = \widehat{\theta}_b$ $\left(\text{because } \Pi_a(\widetilde{\theta}_a) + \Pi_b(\widetilde{\theta}_a) \geq \Pi_a(\widehat{\theta}_b) + \Pi_b(\widehat{\theta}_b)\right)$. We now consider whether $\theta^B = \widehat{\theta}_b$ is a best response to $\theta^A = \widetilde{\theta}_a$. There are two potentially profitable deviation strategies for B.

$$= \beta\left[\Pi_a(\theta^A) - \Pi_a(\widehat{\theta}_b)\right] + \mu\beta\left[\Pi_b(\theta^A) - \Pi_b(\widehat{\theta}_b)\right],$$

which is maximized at $\theta^A = \widetilde{\theta}_a$. The text follows from the assumption that $\Pi_r(\theta)$ is concave.

One potentially profitable deviation strategy is for B to choose some $\widetilde{\theta}_a < \theta^B < \widetilde{\theta}_b$ 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 strategy for B is to choose some $\theta^B < \widetilde{\theta}_a$ 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 latter deviation strategy. There are, however, two cases for which it is immediate that the latter deviation strategy 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 $\widetilde{\theta}_a$ increases continuously with μ while $\widetilde{\theta}_a = \widehat{\theta}_a$ holds at $\mu = 0$. Second, B can not profitably deviate to $\theta^B = \widehat{\theta}_a$ in the case where the profit functions $\Pi_r(\theta)$ are symmetric over the relevant range of values θ, i.e., if for all $\theta \in [\widehat{\theta}_a, \widehat{\theta}_b]$ it holds that $\Pi_a(\theta) = \Pi_b(\widehat{\theta}_a + \widehat{\theta}_b - \theta)$ and likewise that $\Pi_b(\theta) = \Pi_a(\widehat{\theta}_a + \widehat{\theta}_b - \theta)$.[23]

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 $\widetilde{\theta}_b$ and A chooses $\widehat{\theta}_a$. When will A find it profitable to deviate and challenge B for the single-sourcing contract, which is optimally done by choosing $\theta^A = \widetilde{\theta}_a$? Substituting into A's expected profit in (12), where we have to replace $\widehat{\theta}_b$ by $\widetilde{\theta}_b$, we find that A's deviation is only profitable if

$$\left[\Pi_a(\widetilde{\theta}_a) + \Pi_b(\widetilde{\theta}_a)\right] - \left[\Pi_a(\widetilde{\theta}_b) + \Pi_b(\widetilde{\theta}_b)\right] > \frac{1-\mu}{\mu}\left[\Pi_a(\widehat{\theta}_a) - \Pi_a(\widetilde{\theta}_a)\right]. \tag{15}$$

That is, the profits when winning the single-sourcing contract (the left-hand side of (15)) must be sufficiently large to compensate the loss at outlet A from choosing $\widetilde{\theta}_a$ instead of $\widehat{\theta}_a$. As the right-hand side of (15) is strictly positive unless $\mu = 1$, condition (15) is weaker than (14).

3.3 Implications for Industry Profit and Welfare

Proposition 4 has the nice implication that total *expected* industry profit is maximized in any pure strategy equilibrium of the game (if the firm choosing $\theta^s \in \left\{\widetilde{\theta}_a, \widetilde{\theta}_a\right\}$ would not be given the contract under single sourcing following a merger, it would be profitable to deviate to $\theta^s \in \left\{\widehat{\theta}_a, \widehat{\theta}_b\right\}$). Intuitively, since the losing supplier under single sourcing earns zero profit if a

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

[23] Formally, note that the value $\theta^* = (\widehat{\theta}_a + \widehat{\theta}_b)/2$ maximizes $\Pi_a(\theta) + \Pi_b(\theta)$ and that $\widetilde{\theta}_a < \theta^*$ holds for all $\mu < 1$. By symmetry it then holds that $\Pi_b(\widehat{\theta}_b) - \Pi_b(\widetilde{\theta}_a) > \Pi_a(\widehat{\theta}_a) - \Pi_a(\widetilde{\theta}_a)$ for all $\mu < 1$.

merger occurs, its incentive is to choose product characteristics that maximize its profit in the non-merger case. And, given this, the winning supplier is free to choose product characteristics that maximize its expected payoff over the merger and non-merger states of the world because it does not have to deter its rival from challenging on the consolidated market. Nevertheless, this is not to say that suppliers would welcome the possibility of retail mergers when $\beta > 1/2$.

By Proposition 4, the supplier that expects to win the single-sourcing contract after a merger, say A, chooses more "average" product characteristics. While this repositioning of A's product increases industry profits (and A's profit) *conditional* on there being a merger, an increase in the ex-ante likelihood of a merger, μ, reduces total expected industry profits (as well as each supplier's expected profit). If no merger occurs, supplier A will have chosen a suboptimal variety for outlet a. And if a merger occurs, delisting good B will further reduce product variety.[24]

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

Proof. We take the case where $\theta^A = \widetilde{\theta}_a$ and $\theta^B = \widehat{\theta}_b$. The other case is analogous. Expected total industry profits are equal to

$$\mu \left[\Pi_a(\widetilde{\theta}_a) + \Pi_b(\widetilde{\theta}_a)\right] + (1-\mu)\left[\Pi_a(\widetilde{\theta}_a) + \Pi_b(\widehat{\theta}_b)\right]. \tag{16}$$

Differentiating (16) with respect to μ and using that $\widetilde{\theta}_a$ satisfies the first-order condition in (13), we have the derivative $\Pi_b(\widetilde{\theta}_a) - \Pi_b(\widehat{\theta}_b)$, which is negative as $\widehat{\theta}_b$ is the unique maximizer of $\Pi_b(\theta)$. **Q.E.D.**

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

$$P_r(\theta, x) = \max\{p_r(x) + \psi_r(\theta), 0\}, \tag{17}$$

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 (17) 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 (17) and captures the preferences of a representative consumer, (expected) welfare is also strictly decreasing in μ.*

[24] 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.

Proof. We denote for given r and θ the unique optimal quantity by $x_r^*(\theta)$. The envelope theorem then implies that $\frac{d\Pi_r(\theta)}{d\theta} = x_r^*(\theta)\frac{d\psi_r(\theta)}{d\theta}$. Moreover, implicit differentiation of the first-order condition for $x_r^*(\theta)$ shows that the sign of $\frac{dx_r^*(\theta)}{d\theta}$ is determined by the sign of $\frac{d\psi_r(\theta)}{d\theta}$. Welfare at outlet r is $W_r = \int_0^{x_r^*(\theta)}[p_r(x) + \psi_r(\theta)]dx - cx^*(\theta)$. Differentiating welfare with respect to θ, we obtain $\frac{dW_r}{d\theta} = \frac{\partial W_r}{\partial\theta} + \frac{\partial W_r}{\partial x}\frac{dx_r^*(\theta)}{d\theta}$, where the signs of $\frac{\partial W_r}{\partial\theta}$ and $\frac{dx_r^*(\theta)}{d\theta}$ are equal to the signs of $\frac{d\psi_r(\theta)}{d\theta}$. Additionally, we have from standard results that $\frac{\partial W_r}{\partial x} > 0$ at $x = x_r^*(\theta)$.[25] Hence, we have established that welfare realized 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 total 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 choice 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, as shown in Inderst and Wey (2002), these arguments do not withstand formal scrutiny - at least not in this sweeping generalization. While lower total 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 welfare implications in the form of reduced product variety.

3.4 Example

With linear demand $D = 1 - d - p$ and constant marginal costs $c < 1 - d$, joint profits are maximized 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 welfare $W = 3(1 - d - c)^2/8$.[26] 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, $\widehat{\theta}_a = 0$ and $\widehat{\theta}_b = 1$ maximize industry profits at the respective outlets.

The case where suppliers can choose any value for θ has no closed-form solution if $\mu > 0$.

[25] Precisely, note first that $\frac{\partial W_r}{\partial x} = P_r(\theta, x) - c$, while the first-order condition for profit maximization gives $\frac{d\Pi_r}{dx} = P_r(\theta, x) - c + x\frac{dP(\theta,x)}{dx}$. The claim follows as $P_r(\theta, x)$ is strictly decreasing whenever $P_r(\theta, x) > 0$.

[26] To obtain the last expression we assume that the linear demand is generated by a representative consumer with quadratic utility function.

Therefore, we confine ourselves to the case where θ can only be chosen from a finite set $\theta \in \Theta = \left\{ \widehat{\theta}_a, \theta^*, \widehat{\theta}_b \right\}$, 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 2 we have that $\theta^B = \widehat{\theta}_b$ and thus, by substitution into (12), that A strictly prefers θ^* to $\widehat{\theta}_a$ if and only if

$$\mu > \theta^* \frac{10 - (\theta^*)^2}{16 - 4\theta^* + (\theta^*)^3 - 4(\theta^*)^2}. \tag{18}$$

The right-hand side of (18) is strictly increasing in θ^*, which is intuitive because the larger is the difference $\theta^* - \widehat{\theta}_a$, the more likely must a merger be to make it optimal to choose θ^*.

For future reference, we specify $\theta^* = 0.2$, for which (18) becomes $\mu > \frac{83}{627} \approx 13.2\%$. That is, the likelihood of a merger must exceed 13.2% to induce supplier A to choose the less differentiated product. Since the linear demand satisfies (17), we finally have the following result.

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 $\theta^A = \widehat{\theta}_a = 0$.
ii) Expected welfare is strictly decreasing in μ.[27]

4 Linear Contracts

4.1 Analysis

So far we have assumed that negotiations are efficient. Retail contracts are indeed often complex and 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 what follows, we now consider the opposite extreme where contracts determine only a uniform purchase price, and we restrict consideration to the case where the suppliers have all the bargaining power ($\beta = 1$).[28]

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

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

[28] The case with negotiations and $\beta < 1$ does not yield new 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.)

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

$$\max_p \ (p - m_a^A) D_a(\theta^A, p) \geq \Pi_a(\theta^B). \tag{19}$$

In words, the maximum profit that retailer a can realize 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 (19) is not binding.[29] 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 (19) binds. As retailer a's profits are strictly decreasing in m_a^A (as long as $D_a > 0$), this yields a unique offer.[30] Note that, in equilibrium, the supplier whose product offers the highest feasible profits $\Pi_r(\theta^s)$ wins the contract to supply retailer r.

Consolidated retailer

Under a single-sourcing strategy, each supplier offers a single price,[31] and the analysis is 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 (iii) the winning supplier offers m^s such that the retailer is indifferent between the two offers.[32] 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 summarizes 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

[29] 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 $(p^*(m^*) - m^*)D_a(\theta^A, p^*(m^*)) \geq \Pi_a(\theta^B)$.

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

[31] 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.

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

shifting profits to the retailer reduces double marginalization. This trade-off applies also if we endogenize product characteristics as was done in Section 3: a higher μ makes suppliers more homogenous, which intensifies competition and further reduces double marginalization.

The trade-off between the loss of variety and a reduction in double marginalization 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.

4.2 Example

Given the demand $D = 1 - d - p$ and a 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 $c = 0$ and $z = 5$. 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

$$\begin{aligned} \Pi_a &= \frac{1}{4}\left[1 - \frac{1}{5}\theta^2 - m\right], \\ \Pi_b &= \frac{1}{4}\left[1 - \frac{1}{5}(1-\theta)^2 - m\right]. \end{aligned} \tag{20}$$

Finally, recall that we restricted consideration to product characteristics in $\theta \in \Theta = \left\{\widehat{\theta}_a, \theta^*, \widehat{\theta}_b\right\}$, where $\widehat{\theta}_a = 0$, $\widehat{\theta}_b = 1$ and $\theta^* = 0.2$.

Suppose again that A wins under single sourcing, which implies from Proposition 4 that $\theta^B = \widehat{\theta}_b = 1$. 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 (20), we can solve for $m_a^A = (1 - (\theta^A)^2)/5$, 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 and consider how the choice of θ^A affects welfare if retailers stay separate. With efficient contracts, 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 in turn reduces the double-marginalisation problem. Straightforward calculations establish that this countervailing effect dominates and welfare is higher for $\theta^A = \theta^*$.[33]

We next consider the case of a consolidated retailer. For $\theta^A = 0$, symmetry implies that suppliers realize zero profits and make the offers $m^A = m^B = c = 0$. If $\theta^A = \theta^*$, supplier A

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

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 (20) yields $m^A = 0.0285$.

Putting the results together and solving for the optimal choice of θ^A, we obtain the following results. (See the Appendix for the complete calculations.)

Results for the linear example with linear contracts:
i) If $\mu > 11.1\%$, *supplier A chooses the less differentiated product variant* $\theta^A = \theta^*$. *Otherwise, supplier A chooses the more differentiated product variant* $\theta^A = \widehat{\theta}_a$.
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* θ^*, *expected welfare jumps up. This is also the highest feasible value for expected welfare.*

We can now compare the outcomes with efficient and linear contracts. In the former case, 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 $\theta^A = \theta^*$ and no merger took place.[34] However, to induce the supplier to choose a less differentiated product, it is necessary to have $\mu > 0$.

Though our comparison is clearly confined to a very specific example, it highlights an important question for analyzing 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.

5 Discussion

Buyer Groups

A merger enhances buyer power by allowing the consolidated retailer to adopt a single-sourcing strategy. In principle, these benefits could also be achieved if separate retailers agreed to purchase from a single supplier. In fact, buyer groups (or buyer alliances) have gained considerable importance.[35] While our theory is in principle applicable to buyer groups, we feel that the applicability and welfare implications of our analysis will be more pronounced in the

[34] The total welfare order is as follows: Welfare equals 0.655 with no merger and $\theta^A = \theta^*$, 0.641 with a merger and $\theta^A = \theta^*$, 0.640 without a merger and $\theta^A = \widehat{\theta}_a$, and 0.615 with a merger and $\theta^A = \widehat{\theta}_a$.

[35] Dobson (2002) documents this trend for Europe.

case of a 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 consolidate their supplier base.

Suppose the two separate 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 purchase only one good (single sourcing). By Assumption 1, this implies that at one outlet an inferior good is sold. Absent side payments between the two retailers, it may thus be difficult to ensure that the winning supplier's offer is beneficial to both retailers.[36] 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 b may be $B's$ private information. Likewise, b may not know what profits 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.

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.[37]

Suppose thus that demand at both outlets is characterized by the same function $D(p, \theta)$ and that both goods have the same characteristics $\theta^A = \theta^A = \theta$.[38] 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, we have that $t_a^A < t_b^A$. 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, Assumption 1 is still satisfied. But this was all that we needed to derive our main results.[39]

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

[37] In the working paper version 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. Consequently, our results hold *a fortiori* if Assumption 1 holds while costs are strictly convex.

[38] 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.

[39] 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*.

6 Conclusion

This paper analyses the impact of retail mergers on product variety. We conceptualize a merger as extending the range of outlets controlled by a single retailer. The larger retailer may continue to buy from all previous suppliers or it may choose to buy from only a subset of the previous suppliers (single sourcing). We show that a single-sourcing policy may be profitable as it increases competition between suppliers by reducing their differentiation. The resulting benefits may more than outweigh the loss in industry profits due to a reduction in product variety.

As suppliers anticipate the single-sourcing strategy 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.

Our model provides a parsimonious theory of the origins and (welfare) consequences of buyer power. It emphasizes 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 retailers' *previous* supplier base 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-a-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 paper. 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 paper.

Appendix: Omitted Proofs

Omitted calculations for Example 2

We first analyze the optimal choice of θ^A, given that $\theta^B = 1$. Suppose $\theta^A = 0$. If a merger takes place, supplier A realizes zero profits. If a merger does not take place, supplier A supplies retailer a at the previously derived price of $m_a^A = (1 - (\theta^A)^2)/5 = 0.2$. As retailer a chooses the output $x_a = (1 - (\theta^A)^2/5 - m_a^A)/2 = 0.4$, supplier A realizes $m_a^A x_a = 0.8$. Thus, if supplier A chooses $\theta^A = 0$, its expected profits are $(1 - \mu)0.8$. Suppose next that $\theta^A = \theta^* = 0.2$. If a merger takes place, we have $m^A = 0.0285$. Moreover, the merged 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 $\theta^A = \theta^*$, we obtained $m^A = (1 - (1/5)^2)/5 = 0.192$ and the supply of $x_a = (1 - (\theta^A)^2/5 - m^A)/2 = 0.400$ to outlet a, yielding the profits $x_a m^A = 0.077$. In total, for $\theta^A = \theta^*$ the expected profits of A are $0.0258\mu + (1 - \mu)0.0768$. Comparing profits for $\theta^A = \widehat{\theta}_a$ and $\theta^A = \theta^*$, we obtain that supplier A prefers $\widehat{\theta}_a$ for $\mu < 0.111$.

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

References

Adams, W. and Yellen, J. (1976), Commodity Bundling and the Burden of Monopoly, Quarterly Journal of Economics 90, 475-498.

Anton, J. and Yao, D. (1989), Split Awards, Procurement and Innovation, Rand Journal of Economics 20, 538-552.

Bernheim, D. and Whinston, M. (1998), Exclusive Dealing, Journal of Political Economy 106, 64-103.

Binmore, K., Rubinstein, A., and Wolinsky, A. (1986), The Nash Bargaining Solution in Economic Modeling, Rand Journal of Economics 17, 176-188.

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 81, 326-340.

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

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.

Dobson, P. and Waterson, M. (1997), Countervailing Power and Consumer Prices, Economic Journal 107, 418-430.

Dobson, P. and Waterson, M. (1999), Retailer Power: Recent Developments and Policy Implications, Economic Policy 28, 133-164.

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.

Dobson, P. (2002), Retailer Buyer Power in European Markets: Lessons from Grocery Supply, Loughborough University Business School, WP 2002-1.

Fumagalli, C. and Motta, M. (2000), Buyers' Coordination, mimeo.

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.

Horn, H. and Wolinsky, A. (1988), Bilateral Monopolies and Incentives for Merger, Rand Journal of Economics 19, 408-419.

Inderst, R. and Wey, C. (2003), Market Structure, Bargaining, and Technology Choice in Bilaterally Oligopolistic Industries, Rand Journal of Economics 34, 1-19.

Inderst, R. and Wey, C. (2002), Buyer Power and Supplier Incentives, CEPR Discussion Paper No. 3547, London.

Katz, M.L. (1987), The Welfare Effects of Third Degree Price Discrimination in Intermediate Goods Markets, American Economic Review 77, 154-167.

Marx, L. and G. Shaffer (2004), Slotting Allowances and Scarce Shelf Space, Duke University working paper.

Mazzarotto, N. (2003), Retail Mergers and Buyer Power, mimeo.

McAfee, R., McMillan, J. and Whinston, M. (1989), Multiproduct Monopoly, Commodity Bundling, and Correlation of Values, Quarterly Journal of Economics 29, 371-388.

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.

O'Brien, D.P. and Shaffer, G. (1997), Nonlinear Supply Contracts, Exclusive Dealing, and Equilibrium Market Foreclosure, Journal of Economics and Management Strategy 6, 755-785.

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

Palfrey T, (1983), Bundling Decisions by a Multiproduct Monopolist with Incomplete Information, Econometrica 51, 463-484.

Perry, M. and Sakovics, J. (2003), Auctions for Split-Award Contracts, The Journal of Industrial Economics, 51, 215-242.

Rey, P., Retailer Buyer Power and Competition Policy, in: Annual Proceedings of the Fordham Corporate Law Institute (Chapter 27), 2000.

Riordan, M. and Sappington, D (1989), Second Sourcing, Rand Journal of Economics 20, 41-57

Snyder, C.M. (1996), A Dynamic Theory of Countervailing Power, Rand Journal of Economics 27, 747-769.

von Ungern-Sternberg, T. (1996), Countervailing Power Revisited, International Journal of Industrial Organization 14, 507-520.