WORKING PAPERS IN ECONOMICS

No. 07/20

TEIS LUNDE LØMO, FRODE MELAND
AND HÅVARD MORK SANDVIK

DO SLOTTING ALLOWANCES REDUCE PRODUCT VARIETY?

Department of Economics
UNIVERSITY OF BERGEN

Do slotting allowances reduce product variety?[*]

Teis Lunde Lømo[†] Frode Meland[‡] Håvard Mork Sandvik[§]

October 30, 2020

Abstract

Slotting allowances are lump-sum fees paid by manufacturers in return for retail shelf space. We present a novel mechanism by which such upfront payments facilitate vertical foreclosure and thereby reduce product variety. When bidding for the patronage of two retailers, one manufacturer may foreclose a symmetric rival by offering slotting allowances paired with per-unit input prices that offset downstream competition *ex post*. Contrary to the conventional wisdom, slotting allowances can exclude first-rate brands of powerful manufacturers. Our results are in line with recent empirical evidence on slotting allowances but cast doubt on the current policy approach to these payments.

Keywords: vertically related markets, slotting allowances, product variety, vertical foreclosure, exclusion, antitrust policy

JEL classification: L13, L14, L42

[*]We would like to thank Özlem Bedre-Defolie, Øystein Foros, Tommy Staahl Gabrielsen, Germain Gaudin, Bjørn Olav Johansen, Johan Lagerlöf, Eeva Mauring, Dennis Rickert, Greg Shaffer, and Andreas Tveito, as well as audiences at the BECCLE Conference (Bergen, 2019), the CRESSE Conference (Rhodes, 2019), the EARIE Conference (Barcelona, 2019), the MaCCI Conference (Mannheim, 2019), the NORIO Conference (Stockholm, 2019), and the Norwegian School of Economics for helpful comments and discussions. The present paper supersedes an earlier draft circulated under the title "Product variety with competitive retail bottlenecks." The views expressed in this paper do not necessarily reflect those of the Norwegian Competition Authority.

[†]Department of Economics, University of Bergen, and BECCLE. Email: teis.lomo@uib.no.

[‡]Department of Economics, University of Bergen. Email: frode.meland@uib.no.

[§]Norwegian Competition Authority, and BECCLE. Email: hasa@kt.no.

1 Introduction

According to Business Insider, manufacturers of sugary cereal brands make upfront payments to supermarkets who agree to place their colorful boxes exactly 23 in. off the ground, where they wait to catch the eyes of kids eager to influence, if not dictate, their parents' purchase decision.[1] More broadly, lump-sum fees paid in return for access to shelf space are called *slotting allowances*. These payments, which amount to vast amounts of money every year, are pervasive in the grocery industry.[2] They are also widely used for other goods including apparel, toys, electronics, books, and over-the-counter drugs (Klein and Wright, 2007; Raff and Schmitt, 2016).

Slotting allowances have been controversial among antitrust scholars and policy makers for many years.[3] Proponents of the practice claim that it can help a supplier get her product on the market (see Section 2 for a review of the literature). However, this view largely abstracts from the fact that, in a given product category, there are typically many brand manufacturers fighting for retail distribution – a phenomenon which, in the grocery sector, is colloquially referred to as the "shelf space wars."[4] Our contribution to the literature and the policy debate is a novel theory in which manufacturers compete head-to-head for retail shelf space. We find that slotting allowances facilitate vertical foreclosure and thereby reduce product variety.

To see the idea, think of two symmetric manufacturers of differentiated brands, A and B, bidding for the patronage of two competing retailers. Each retailer wants to stock one brand.[5] If A and B sell to one retailer each, *interbrand* competition drives down per-unit input prices, retail prices, and profits. Conversely, if one manufacturer, say A, sells to both retailers, it can set input prices that offset *intrabrand* competition and generate monopoly profits. When brands are not too differentiated and retail competition is sufficiently fierce, A and the retailers can, therefore, earn higher profits than all four firms combined. In that case, A can convince both retailers to reject B's contracts by offering to share the monopoly profits with them

[1] "Why cereal boxes are at eye level with kids," *Business Insider*, January 14, 2019.

[2] For example, using rich panel data from the Chilean grocery market, Elberg and Noton (2019) document slotting allowance payments to *one* mid-sized supermarket chain in the range of 296-514 million USD per year (2010-2012). These authors report also that slotting allowances were paid by more than 95 percent of all suppliers, and close to 98 percent of the retailer's large suppliers. For evidence on slotting allowances in the US grocery industry, see, e.g., Sudhir and Rao (2006).

[3] See, e.g., reports by the US Federal Trade Commission (2001, 2003), UK Competition Commission (2008), and UK Office of Fair Trading (2013).

[4] See, e.g., "The hidden war over grocery shelf space," *Vox.com*, November 11, 2016, and "Buying up the shelves," *The Economist*, June 18, 2015.

[5] Slotting allowances typically occur when shelf space is a scarce resource, see, e.g., Sullivan (1997) and Marx and Shaffer (2010) for anecdotal evidence and Section 3 for further discussion and documentation.

through large slotting allowances. By contrast, if A tried such a move in the absence of slotting allowances, A would have to compensate the retailers by cutting input prices, but this would unleash retail competition. In other words, slotting allowances enable A to disentangle the maximization and redistribution of exclusivity profits. The ability to offer such payments thereby helps A evict B from the market.

We give the above intuitive argument a rigorous foundation by setting up and solving a bidding game where manufacturers offer two-part tariffs, and where slotting allowances (i.e., negative fixed fees) are either feasible or not feasible. The retailers face general consumer demand functions and compete by setting prices. We show that, under mild conditions, a ban on slotting allowances expands the scope for equilibria in which both manufacturers' brands obtain retail distribution. This result is robust to retail quantity competition and contractual arrangements beyond two-part tariffs. Moreover, when considering a representative consumer with a quasilinear utility function, we find that a ban on slotting allowances also lowers retail prices and unambiguously raises consumer surplus.

Conventional wisdom suggests that, one the one hand, a dominant manufacturer may use slotting allowances as a tool to foreclose "weaker" rivals, e.g., credit-constrained suppliers, or potential future entrants. On the other hand, the story goes, exclusionary effects are not a concern if the rivals are on an equal footing and able to make their own counteroffers.[6] This view has been extremely influential on the current policy approach, as illustrated by the following excerpt from the European Commission (2010, pp. 41-42, emphasis added) in their Guidelines on Vertical Restraints:

> "As long as the competitors are sufficiently numerous and strong, no appreciable anti-competitive effects [from slotting allowances or other up-front access payments] can be expected. *Foreclosure of competitors is not very likely where they have similar market positions and can offer similarly attractive products.*"

Our model challenges this view by illustrating that the use of slotting allowances facilitates the exclusion of a manufacturer who, compared with her rival, has exactly the same ability to offer contracts, and makes an equally popular, differentiated

[6] For example, Klein and Wright (2007, p. 422) state that "The primary competitive concern with slotting arrangements is the claim that they may be used by manufacturers to foreclose or otherwise disadvantage rivals [..] It is now well established in both economics and antitrust law that the possibility of this type of anticompetitive effect depends on whether a dominant manufacturer can control a sufficient amount of distribution so that rivals are effectively prevented from reaching minimum efficient scale."

product at the same cost. As explained above, this mechanism is driven by the joint presence of upstream and downstream competition and does not rely on any type of asymmetry between manufacturers. Thus, for courts and antitrust practitioners, the conclusion is that exclusionary effects of slotting allowances cannot be ruled out on the grounds that a given industry has fairly symmetric firms in the upstream sector.[7]

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 sets up the model and discusses its key assumptions. Section 4 characterizes equilibria of the model with and without slotting allowances, and then compares these regimes to state our main results. Section 5 discusses the robustness of these results. Section 6 then concludes.

2 Related literature

This paper relates to the literature on slotting allowances, particularly to the strand that studies how these payments affect product variety and the scope for upstream exclusion.[8]

The paper closest to ours is Shaffer (2005). He studies a model with one dominant manufacturer, a set of small manufacturers ("competitive fringe"), and two homogeneous retailers with room for one product. Shaffer (2005) finds that the dominant firm can sometimes use slotting allowances to effectively raise the price of shelf space and thereby foreclose the fringe firms, who, by assumption, are unable to offer such payments.[9] Instead, our exclusionary mechanism works through a bidding game between two evenly matched manufacturers, both of whom can offer slotting allowances. From a policy perspective, this difference is important because it highlights that slotting allowances facilitate exclusion also of brands from manufacturers with considerable clout vis-à-vis retailers. From a theoretical perspective, the extension to two large manufacturers leads to a contracting game which is sub-

[7] Notably, the US Federal Trade Commission (2001, p. 9), recommended that one should "examine slotting allowances [...] with particular attention to circumstances that could give rise to exclusionary effects."

[8] It should be noted that there also exist other, pro-competitive theories of slotting allowances, e.g., that such payments can help with new product launches (Sullivan, 1997) or promote valuable retail services (Lømo and Ulsaker, forthc.).

[9] Further away from our approach, Asker and Bar-Isaac (2014) find that an incumbent manufacturer can use slotting allowances to foreclose a potential future entrant, who does not have the ability to offer slotting allowances (see also Choi and Stefanadis, 2018). Relatedly, note that our mechanism is also inherently different from the literature on naked exclusion with competing buyers (e.g., Fumagalli and Motta, 2006), which asks whether an incumbent can prevent entry through contracts with explicit exclusionary provisions. In our model, such contract clauses have little bite on equilibrium outcomes and cannot be used to sustain exclusion in the absence of slotting allowances.

stantially richer than the one in Shaffer (2005), and which may be of independent interest.[10] In addition, we allow for differentiation between retailers, the degree of which turns out to be a key determinant of when exclusion occurs.

Our main result stands in contrast to Hamilton and Innes (2017), who argue that slotting allowances *raise* product variety.[11] The difference between their result and ours comes down to how product variety is conceptualized. They consider a Hotelling-model in which two firms (retailers) first choose their product ranges and then compete in prices. In such models, it is well-known that narrower product ranges relax price competition (see Anderson and De Palma, 1992). Hamilton and Innes (2017) point out that slotting allowances (coupled with per-unit input prices above cost à la Shaffer, 1991) provide an alternative way to relax competition, and thus allow product ranges (and thereby variety) to expand. By contrast, in our model, equilibrium product variety is determined by the number of manufacturers who obtain distribution for their brands. On average, slotting allowances reduce this number, and thereby reduce product variety. Notably, this insight does not rely on a certain micro-foundation for retail competition.[12] Compared with the result of Hamilton and Innes (2017), our mechanism should, therefore, be less sensitive to the fine details of a given real-world market. In addition, our prediction is supported by recent empirical evidence (more on this below).

The interplay between slotting allowances and retailers' product lines features also in Marx and Shaffer (2010) and Chambolle and Molina (2019). Marx and Shaffer (2010) show that the opportunity to receive slotting allowances may lead retailers to reduce their stocking capacity and thereby narrow their product ranges. Chambolle and Molina (2019) find that retailers can elicit slotting allowances by credibly threatening to replace one manufacturer's product with a rival brand (in the spirit of Ho and Lee, 2019), and that prohibiting slotting allowances can expand product variety by disrupting this strategy. However, our result is distinct from both these findings. In particular, Marx and Shaffer (2010) and Chambolle and Molina (2019) consider settings with a single, monopolistic retailer (or, equivalently, multi-

[10] In fact, our game shares many features with settings in which several retailers make offers to a common supplier. This literature was initiated by Marx and Shaffer (2007), and then developed by Miklós-Thal et al. (2011), Rey and Whinston (2013), and Gabrielsen and Johansen (2015). See Section 3 for further discussion.

[11] In the marketing literature, several authors argue that, if a manufacturer offers a slotting allowance on a new product, this may signal to retailers that the product will be in high demand, and therefore deserving of a place on the shelves (Kelly, 1991; Lariviere and Padmanabhan, 1997; Desai, 2000). However, these papers abstract from exclusionary effects by restricting attention to models with an upstream monopolist (and no retail competition).

[12] While our main model has retail price competition, the key insight holds also with retail quantity competition (see Section 5.2).

ple retailers serving separate markets). Thus, in their models, the industry-profit maximizing retail prices can always be induced with two-part tariffs (see Bernheim and Whinston, 1985). By contrast, retail competition is key in our model, as it prevents the maximization of industry profits by two-part tariffs when all four firms are active, which in turn makes exclusion attractive (and sometimes feasible, when slotting allowances are permitted). Another (and related) difference is that a ban on slotting allowances directly impacts retail prices in our model.

On the empirical side, Hristakeva (2019) studies the impact of lump-sum payments such as slotting allowances on product variety by estimating a structural model of the US yogurt market. Specifically, she compares retailers' brand selections across two regimes; one in which manufacturers are free to offer lump-sum payments, and a counterfactual where they can set only per-unit input prices. The main finding of Hristakeva (2019) is that supermarkets stock more products in the counterfactual, i.e., the use of lump-sum payments reduce product variety. Notably, in her data sample period (2001-2010), General Mills and Groupe Danone were the leading suppliers in this industry, with average market shares of 39 and 31 percent, respectively. Thus, the empirical results in Hristakeva (2019) are in line with our point, namely that slotting allowances tend to restrict product variety even in markets with a few large manufacturers.

3 Model

We wish to study the effect of slotting allowances on product variety in an environment with two key features: 1) manufacturers bid for each retailer's patronage, and 2) retailers have a scarcity of shelf space. The first feature is motivated by the "shelf space wars" in the grocery industry, mentioned in Section 1. The second feature captures the fact that, in many product categories, manufacturers in total offer many more varieties than a single retailer demands. This may again be explained by product proliferation[13] or the increasing degree to which retailers sell imported goods (Raff and Schmitt, 2016) and in-house private labels,[14] which limits the share

[13]As an example, in 2018, the average US supermarket carried 306 varieties of yogurt, while over 1,400 varieties were available from manufacturers, see "Yogurt Sales Sour as Options Proliferate," *Wall Street Journal*, April 9, 2019, and "As yogurt options multiply, rivals strive to innovate," *Seattle Times*, November 5, 2018. In a similar vein, Marx and Shaffer (2010, p. 575) state (and provide evidence to support) that "The typical supermarket carries less than 30,000 products, and yet, at any given time, there may be over 100,000 products from which to choose."

[14]According to AC Nielsen, private label market shares lie between 20 and 40 percent in most of the European countries (see `https://www.nielsen.com/ssa/en/insights/report/2018/the-rise-and-rise-again-of-private-label/`). In the US, the private label market share has

of shelf space that is contestable for national brand manufacturers.

To this end, we consider a model with two manufacturers of differentiated brands and two differentiated retailers. We denote the manufacturers and their brands by A and B and the retailers by 1 and 2. Each manufacturer has the capacity to supply both retailers. By contrast, each retailer wants to stock at most one brand. We refer to the situations in which retailers stock the same brand and different brands as "one-brand" and "two-brand," respectively. There are four such configurations: The two-brand cases $\{A1, B2\}$ and $\{B1, A2\}$, and the one-brand cases $\{A1, A2\}$ and $\{B1, B2\}$. We assume that the maximal industry profit is the same in $\{A1, B2\}$ and $\{B1, A2\}$, and in $\{A1, A2\}$ and $\{B1, B2\}$. That is, each manufacturer matches equally well with both retailers and neither manufacturer has a special advantage as an exclusive supplier.

The timing of events is as follows:

1. Each manufacturer makes publicly observable contract offers to both retailers.
2. Each retailer chooses which offer to accept.[15] This decision is public information.
3. Accepted contracts are implemented, and retailers compete in prices downstream.

The solution concept is subgame perfect Nash equilibrium. We focus on symmetric equilibria in which all firms earn weakly positive profits.

A key feature of this game is that manufacturers can make contract offers to *both* retailers at the first stage. By contrast, models of "competing vertical structures" (e.g., Bonanno and Vickers, 1988; Rey and Stiglitz, 1995) typically assume that 1) there is an exogenously determined relationship between one manufacturer and one retailer and 2) manufacturers are unable to make offers to other retailers. By removing these *ad hoc* restrictions, we introduce head-to-head competition for shelf space, where a one-brand structure may emerge in equilibrium.

The firms sign two-part tariffs,[16] which we allow to be contingent on the number of brands sold. Formally, the contract offered by manufacturer $i \in \{A, B\}$ to retailer $j \in \{1, 2\}$ is $\Gamma_{ij}^Y(q) = F_{ij}^Y + w_{ij}^Y q$, where w_{ij}^Y is the (per-unit) input price, $q \geq 0$ is

been estimated to lie at around 23 percent (see https://www.statista.com/statistics/1057671/private-label-unit-sales-share-us/).

[15] Note that we here rule out the possibility that a retailer can reject *both* offers and thereby exit the market. This eases the exposition but does not affect our main results, see Section 5.1 for further details.

[16] Empirical evidence suggests that two-part tariffs are widely used in vertically related industries (e.g., Bonnet and Dubois, 2010; Lafontaine and Slade, 2010; Bonnet and Réquillart, 2013).

the traded quantity, and superscript $Y \in \{T, O\}$ refers to two-brand (T) and one-brand (O) configurations. We call the fixed fee, F_{ij}^{Y}, a slotting allowance whenever $F_{ij}^{Y} < 0$, in which case it is paid by the manufacturer to the retailer. The contract $\Gamma_{ij}^{Y}(q)$ can be interpreted as a menu where the terms (F_{ij}^{O}, w_{ij}^{O}) apply when i is the only active manufacturer, and (F_{ij}^{T}, w_{ij}^{T}) apply when both manufacturers are active. These contracts are similar in spirit to those in Miklós-Thal et al. (2011) and Rey and Whinston (2013), where retailers offer menus contingent on the number of active retailers. In our model, contingent contracts can be justified on the grounds that a drastic change in market structure, i.e., the exit of one manufacturer, cannot be dictated by the rival *ex ante*, and contracts should therefore take account of both market structures. Also, without contingent contracts, the above game admits no pure strategy Nash equilibria in the regime where slotting allowances are feasible, see Schutz (2013) for a rigorous argument.

Moreover, the assumption of publicly observable contracts is not innocuous. In our model, the role of this assumption is to enable manufacturers to offset, at least partially, brand and retail competition through their input prices, w_{ij}^{Y}. By contrast, if contracts were unobservable, one retailer's optimal pricing decision would not depend on the rival's input price and the manufacturers would not be able to commit to this strategy.[17] Thus, before moving on, it is important to assess the likelihood of contract observability.[18] In the grocery industry, large manufacturers are sometimes required by law to post general non-discriminatory terms of sale. In addition, input price discrimination may be prohibited. Under either of these circumstances, contracts are *de facto* observable. More broadly, the assumption can be seen as a shorthand way of capturing the compelling idea that long-run interaction with retailers enables manufacturers to build a reputation for credibility and thereby commit to a set of supply terms even if there is a short-run gain from changing terms behind the retailers' backs.

Let $D_{ij}(p_{ij}, p_{hk})$ be the direct demand for brand $i \neq h \in \{A, B\}$ at retailer $j \neq k \in \{1, 2\}$. We assume that demand functions are symmetric. Furthermore, the function D_{ij} is smooth, with a negative and finite own-price effect, $\partial D_{ij}/\partial p_{ij} < 0$, and a positive cross-price effect, $\partial D_{ij}/\partial p_{hk} > 0$, such that $\partial D_{ij}/\partial p_{ij} + \partial D_{ij}/\partial p_{hk} < 0$.

[17]With unobservable contracts, equilibrium input prices are determined not by the intensity of competition, but by the retailers' out-of-equilibrium beliefs about rivals' contracts. See Pagnozzi and Piccolo (2012) for the case of (exogenously given) competing vertical structures and, e.g., Rey and Vergé (2004) for the case of upstream monopoly.

[18]Note, however, that observable contracts are standard in the literature on slotting allowances (e.g., Shaffer, 1991, 2005) and bidding games (e.g., Miklós-Thal et al., 2011; Rey and Whinston, 2013).

In addition, second order derivatives satisfy $-\partial^2 D_{ij}/\partial p_{ij}^2 > \partial^2 D_{ij}/\partial p_{ij}\partial p_{hk} > 0$. These assumptions are standard (Vives, 1999, p. 150), and ensure that 1) there exists a unique and stable equilibrium in retail prices (for given input prices) and 2) retail prices are strategic complements.

The manufacturers have constant and symmetric marginal costs, $c_A = c_B = c \geq 0$. We normalize all other production costs for the firms to zero.

The maximal industry profit in the two-brand structure $\{A1, B2\}$ (and, by symmetry, in $\{B1, A2\}$) is

$$\Pi_M^T \equiv \max_{p_{A1}, p_{B2}} \{(p_{A1} - c)D_{A1}(p_{A1}, p_{B2}) + (p_{B2} - c)D_{B2}(p_{B2}, p_{A1})\}. \tag{1}$$

The maximal industry profit in the one-brand structure $\{A1, A2\}$ (and in, by symmetry, $\{B1, B2\}$) is

$$\Pi_M^O \equiv \max_{p_{A1}, p_{A2}} \{(p_{A1} - c)D_{A1}(p_{A1}, p_{A2}) + (p_{A2} - c)D_{A2}(p_{A2}, p_{A1})\}. \tag{2}$$

Let p_M^T and p_M^O be the corresponding, industry-profit maximizing (symmetric) retail prices. We make the following assumption.

Assumption 1. $\Pi_M^O < \Pi_M^T < 2\Pi_M^O$.

Assumption 1 implies that consumers see the manufacturers' brands as differentiated but not independent, i.e., imperfect substitutes. It is in this sense that two-brand structures entail more product variety than one-brand structures.

Finally, let $\Pi^Y(w, w)$ be the industry profit in market structure $Y \in \{T, O\}$ for an arbitrary pair of input prices.

Assumption 2. The function Π^Y is quasi-concave and there exist unique (w_M^Y, w_M^Y) such that $\Pi^Y(w_M^Y, w_M^Y) = \Pi_M^Y$, for $Y \in \{T, O\}$.

Assumption 3. $\Pi^T(w, w) > \Pi^O(w, w)$ for any $w \in [c, \min\{w_M^T, w_M^O\})$.

Assumption 2 ensures that, in each market structure, there exist a pair of input prices that fully offset competition and induce the maximal industry profit. Assumption 3 states that, for any symmetric input price below these "first-best" levels, the realized industry profit is larger when both brands obtain distribution. For instance, this may be because retailers compete less vigorously for a given input price when offering differentiated brands. Alternatively, Assumption 3 can hold because consumers display "love for variety" such that aggregate demand is larger when both brands are sold. Assumptions 1-3 hold with linear demands, which we return to in Section 4.5.

4 Analysis

The key question to be analyzed is how the use of slotting allowances impacts the scope for a two-brand structure to arise as an equilibrium in our model. To support and reinforce this analysis, we also examine equilibria with one-brand structures.

More specifically, the analysis proceeds as follows. First, in Section 4.1, we derive a key condition for the existence of a two-brand equilibrium. Then, we characterize the (two-brand and one-brand) equilibria of our model, first when slotting allowances are feasible in Section 4.2 and then when they are not in Section 4.3. In Section 4.4, we identify the effect of slotting allowances (or, equivalently, a ban on such fees) on product variety by comparing the set of equilibria obtained in the two regimes. Finally, in Section 4.5, we use a representative consumer with quasi-linear utility to illustrate our results and examine the consumer welfare effect of prohibiting slotting allowances.

4.1 Preliminaries

We use the following notation. The profit of manufacturer $i \in \{A, B\}$ is π_i, and the profit of retailer $j \in \{1, 2\}$ is π_j. Throughout the paper, we denote equilibrium values by upper bars. As such, $\overline{\Pi}^T$ and $\overline{\Pi}^O$ are the industry profits in (candidate) two-brand and one-brand equilibria, respectively. Finally, we denote by Π_D^O the *maximal* industry profit that one manufacturer can generate by deviating from a two-brand structure to a one-brand structure.

Now, to begin examining when a symmetric two-brand equilibrium may exist in our model, consider the following inequality:

$$\frac{\overline{\Pi}^T}{2} - \pi_j \geq \Pi_D^O - 2\pi_j. \tag{3}$$

The left-hand side of (3) is the net profit of one manufacturer in the two-brand candidate. The right-hand side of (3) is the net profit of one manufacturer following the most profitable deviation to a one-brand structure. Note that, as each retailer earns π_j in the candidate, the deviating manufacturer must offer *both* retailers π_j, or marginally more, when seeking to induce the one-brand structure. Thus, whenever

(3) holds, no manufacturer has a profitable deviation from the candidate two-brand equilibrium.

From (3), it follows immediately that whenever $\overline{\Pi}^T/2 \geq \Pi_D^O$, no deviation to a one-brand structure can ever be profitable (for $\pi_j \geq 0$). In those cases, two-brand structures generate enough profits to make one-brand structures obsolete. Next, for $\overline{\Pi}^T/2 < \Pi_D^O$, any candidate two-brand equilibrium has $\pi_j \geq \Pi_D^O - \overline{\Pi}^T/2 > 0$. This is the *minimum* retail profit that deters one manufacturer from deviating to a one-brand structure (i.e., A to $\{A1, A2\}$ or B to $\{B1, B2\}$). Intuitively, high enough retail profits in the two-brand candidate deter deviations to one-brand structures because, in the latter, the deviating manufacturer must compensate both retailers, instead of just one retailer. Following this logic, manufacturer profits in the candidate two-brand equilibrium can be at most $\pi_i = \overline{\Pi}^T/2 - (\Pi_D^O - \overline{\Pi}^T/2) = \overline{\Pi}^T - \Pi_D^O$. Whenever this expression is negative, manufacturers are unable to earn non-negative profits in a two-brand equilibrium. Thus, we have the following necessary condition for the existence of a two-brand equilibrium:

$$\overline{\Pi}^T \geq \Pi_D^O. \tag{4}$$

In the following, we characterize the equilibria of our model with and without slotting allowances and thereby determine if and when (4) can be satisfied.

Note also that, in any two-brand equilibrium, both retailers are indifferent, or as close as possible to indifferent, between all contracts offered.[19] The reason is two-fold. First, retailers can always deviate and induce a one-brand structure, which means that they need to be at least as well off in the candidate two-brand equilibrium. Second, if retailers were strictly better off, manufacturers would want to increase two-brand fixed fees. Thus, any two-brand equilibrium is supported by out-of-equilibrium one-brand contracts that make deviations just unprofitable for the retailers.[20]

4.2 Slotting allowances are feasible

We start with the regime in which fixed fees are unrestricted.

[19] When both retailers are strictly indifferent, there are four Nash equilibria in the Stage 2 subgame (a pair of the two-brand type, and a pair of the one-brand type).

[20] Calzolari et al. (2020) call such out-of-equilibrium terms “barrage” tariffs, as opposed to the “actual” tariffs at which trade takes place when both manufacturers are active.

4.2.1 Two-brand equilibria

Consider a candidate two-brand equilibrium with market structure $\{A1, B2\}$. (Of course, all results apply equally to the structure $\{A2, B1\}$.) From Section 4.1, we know already that such an equilibrium can exist only when $\overline{\Pi}^T \geq \Pi_D^O$, and then only if $\pi_j \geq \Pi_D^O - \overline{\Pi}^T/2$. Then, what remains is to determine a set of equilibrium input prices and corresponding profits.

With two-part tariffs, each manufacturer has an incentive to maximize its bilateral profit with the retailer. Therefore, in the candidate two-brand equilibrium, input prices must be mutual best responses in the following sense: No manufacturer-retailer pair can raise their bilateral profit by adjusting their input price, given the input price of the other manufacturer-retailer pair. Let $\rho(w, w)$ be the *flow* profit of any firm for arbitrary input prices. Then, for $i \in \{A, B\}$ and $j \in \{1, 2\}$, the bilateral flow profit of one channel in the candidate equilibrium is $\rho_i(w_{A1}, w_{B2}) + \rho_j(w_{A1}, w_{B2}) \equiv \rho_{ij}(w_{A1}, w_{B2})$. The best-response function of one channel is then

$$\omega_{ij}^{BR}(w_{hk}) \equiv \arg\max_{w_{ij}} \{\rho_{ij}(w_{ij}, w_{hk})\} \tag{5}$$

and the equilibrium input prices are defined through

$$\overline{w}_{ij}^T = \omega_{ij}^{BR}(\overline{w}_{hk}^T), \tag{6}$$

for ($i \neq h \in \{A, B\}$ and $j \neq k \in \{1, 2\}$). We now make the following assumption.

Assumption 4. 1) The function ω_{ij}^{BR} given by (5) is quasi-concave in w_{ij} and 2) the system given by (6) has at least one symmetric solution, denoted by $\overline{w}^T$.[21]

Several key properties of these equilibrium input prices are known from the literature on competing vertical structures (e.g., Bonanno and Vickers, 1988). On the one hand, each manufacturer has an incentive to lower its input price to give its retailer a competitive advantage in the final market. Hence, starting from the point at which manufacturers choose the industry profit maximizing input prices, downward deviations are profitable (i.e., $\overline{w}^T < w_M^T$). On the other hand, as retail prices are strategic complements, each manufacturer can induce the rival's retail price to go up by raising its own input (and retail) price. Starting from marginal cost pricing, upwards deviations are therefore profitable (i.e., $\overline{w}^T > c$). Thus, when

[21]Note that the solution will be unique in many situations, e.g., the quasilinear model in Secion 4.5. Note also that our formulation of input prices here is reminiscent of the one in Miklós-Thal et al. (2011, p. 14).

slotting allowances are feasible, we have $\overline{w}^T \in (c, w_M^T)$ and a candidate two-brand equilibrium industry profit below the maximal level, $\overline{\Pi}^T < \Pi_M^T$.

We can now return to the two-brand equilibrium existence condition, given by (4). Recall that Π_D^O is the industry profit following the most profitable deviation that successfully induces a one-brand structure. As we discuss in more detail in Section 4.2.2 below, when slotting allowances are feasible, a deviating manufacturer can always induce $\Pi_D^O = \Pi_M^O$ by offering both retailers to buy at w_M^O and using fixed fees to satisfy their participation constraints. Consequently, a two-brand equilibrium exists only if $\overline{\Pi}^T \geq \Pi_M^O$. Importantly, with $\overline{\Pi}^T < \Pi_M^T$ as argued above, this condition need not hold even though $\Pi_M^T \geq \Pi_M^O$. We summarize as follows.

Lemma 1 *When slotting allowances are feasible, a two-brand equilibrium exists whenever $\overline{\Pi}^T \geq \Pi_M^O$ and does not exist if the converse is true.*

Whether or not $\overline{\Pi}^T \geq \Pi_M^O$ generally depends on industry characteristics, e.g., consumers' willingness to substitute between brands and retail outlets. We analyze the circumstances under which the condition is more or less likely to hold in Section 4.4 and Section 4.5, after stating our main result. Note also that, whenever $\overline{\Pi}^T \geq \Pi_M^O$, there may exist a range of two-brand equilibria. The issue of equilibrium multiplicity is examined more closely in Section 4.4.1.

4.2.2 One-brand equilibria

Consider now the following candidate one-brand equilibrium with market structure $\{A1, A2\}$, where A obtains distribution while B is excluded. (Again, all results below apply equally to the case $\{B1, B2\}$.) Both manufacturers set $\overline{w}_{ij} = w_M^O$ for $i \in \{A, B\}$ and $j \in \{1, 2\}$, and use slotting allowances to transfer all profits to the retailers. In addition, out-of-equilibrium two-brand contracts are such that, if realized, retailers earn the same as, or marginally below, what they earn in the one-brand case. For this candidate to constitute a one-brand equilibrium, it must survive the following three types of deviations.

First, manufacturer A must not be able to raise profits within the one-brand structure. Indeed, if this was possible, A could change at least one of the input prices while adjusting fixed fees to make the retailers equally well off as before the change, thereby strictly increasing her own profit. However, such a deviation is not feasible as the candidate has $\overline{w}_{Aj}^O = w_M^O$, for $j = \{1, 2\}$. Intuitively, this is the input price that aligns the pricing incentives of retailer j with the pricing incentives of a fully (vertically and horizontally) integrated firm. The Stage 3 first order condition of retailer j is

$$D_{Aj} + (p_{Aj} - w_{Aj})\frac{\partial D_{Aj}}{\partial p_{Aj}} = 0, \tag{7}$$

and, from (2), we have that the industry-wide pricing incentives are given by

$$D_{Aj} + (p_{Aj} - c)\frac{\partial D_{Aj}}{\partial p_{Aj}} + (p_{Ak} - c)\frac{\partial D_{Ak}}{\partial p_{Aj}} = 0. \tag{8}$$

Combining (7) and (8), and evaluating at (p_M^O, p_M^O) we see that the input price in question is

$$\overline{w}_{Aj}^O = c + (p_M^O - c)\frac{\partial D_{Ak}(p_M^O, p_M^O)/\partial p_{Aj}}{-\partial D_{Aj}(p_M^O, p_M^O)/\partial p_{Aj}}, \tag{9}$$

where the last term is the diversion ratio from retailer j to $k \neq j$, which lies strictly between zero and one, when they both sell brand A.[22] In any one-brand equilibrium, the input prices must be as specified by (9).

Second, manufacturer B must not be able to profitably deviate to the one-brand structure $\{B1, B2\}$. Recall that B cannot generate strictly larger profits than A as an exclusive supplier. Thus, there is no way for B to sign up both retailers and at the same time earn strictly positive profits. However, what the threat of deviations to $\{B1, B2\}$ effectively does, is to drive A's profits in the candidate one-brand equilibrium down to zero, in the usual Bertrand fashion.

Finally, A or B must not be able to profitably deviate to one of the two-brand structures. Consider a deviation by A, whereby A retains the patronage of retailer 1 but chooses not to supply retailer 2, who in turn is picked up by B. Suppose further that the (out-of-equilibrium) input price of B is so low that 2 sets $p_{B2} = 0$ for any w_{A1}^T that is individually rational for A. Denote this input price by $\widetilde{w}_{B2}^T$.[23] This is the worst-case scenario for A's deviation, as competition is at its fiercest in the two-brand case. Then, given B's offers, the key question is whether A can still find a w_{A1}^T that induces p_{A1} such that her channel profit with retailer 1 is at least $\Pi_M^O/2$, which is what they jointly obtain in the one-brand structure. If this is feasible even in the

[22]The point that an upstream monopolist who offers public two-part tariffs can fully offset retail competition has been known at least since Mathewson and Winter (1984). More recently, Miklós-Thal and Shaffer (2019) show that, in symmetric settings, the required input price (equivalent to (9)) can be written as a weighted average $w = \theta c + (1-\theta)p_M^O$, where θ is the retail market conduct parameter, defined as one minus the diversion ratio, evaluated at the industry-profit maximizing retail prices. That is, in their formulation, the input price ranges from p_M^O down to c, depending on the degree of monopolization in the retail sector, as measured by θ (see also Weyl and Fabinger, 2013). In Section 4.3.2, we show that the manufacturer may have to deviate significantly from such a pricing strategy when she faces competition for shelf space from an upstream rival and slotting allowances are not feasible.

[23]Note that B is free to increase the fixed fee in order to curb the retailer's profit.

worst-case scenario, then one-brand equilibria cannot exist. Conversely, if this is not feasible, one-brand equilibria exist and are supported by sufficiently degraded out-of-equilibrium two-brand contracts. We summarize as follows.

Lemma 2 *When slotting allowances are feasible, a one-brand equilibrium exists whenever*

$$\max_{w_{A1}}\{\rho_{A1}(w_{A1}, \widetilde{w}^T_{B2})\} \leq \frac{\Pi^O_M}{2}, \tag{10}$$

and does not exist if the converse is true.

It is clear that a one-brand equilibrium exists when $\overline{\Pi}^T < \Pi^O_M$. In this case, even if manufacturers offer their mutual best-responses in the two-brand case (see (6)), it is never profitable for one manufacturer to induce such a two-brand structure. Beyond this, Lemma 2 states that a one-brand equilibrium may exist even when $\overline{\Pi}^T > \Pi^O_M$, because the inactive manufacturer (e.g., B, if the market structure is $\{A1, A2\}$) can "insist" on making one-brand structures viable. Note that the existence of one-brand equilibria following (10) does not preclude the simultaneous existence of two-brand equilibria, see Section 4.4.2 for further discussion.

4.3 Slotting allowances are not feasible

We now characterize the equilibria of our model under the restriction that all fixed fees be non-negative, that is $\overline{F}^Y_{ij} \geq 0$ must hold for $i \in \{A, B\}$, $j \in \{1, 2\}$, and $Y \in \{T, O\}$.

4.3.1 Two-brand equilibria

Consider a candidate two-brand equilibrium where both manufacturers $i \in \{A, B\}$ offer (out-of-equilibrium) one-brand contracts $(w^O_{ij}, F^O_{ij}) = (c, 0)$, and two-brand contracts that give both retailers a profit of $\Pi^O(c, c)/2$, or marginally more, within the two-brand structure.

Note first that, starting from this candidate equilibrium, manufacturers can earn at most zero by deviating to a one-brand structure. The intuition is simple. Because the retailers must be indifferent between one-brand and two-brand structures, they would have to earn at least $\Pi^O(c, c)/2$ in any one-brand structure. As $w^O_{ij} = c$, this requirement implies zero profits to a deviating manufacturer.

Furthermore, note that because $\Pi^T(w, w) > \Pi^O(w, w)$ for $w \in [c, w^T_M)$ (by Assumption 3), deviating to a one-brand structure can *never* be profitable whenever input prices in the candidate two-brand equilibrium lie between c and $\overline{w}^T (< w_M^T)$. Thus, to check for profitable deviations, we can look at the two-brand equilibrium contracts.

There are two cases to consider. One option is that the non-negativity constraint on $\overline{F}_{ij}^T$ does not bind, and the symmetric equilibrium input price is $\overline{w}^T > c$, just as when slotting allowances were feasible (see (6) in Section 4.2.1). This happens if both retailers earn flow profits of at least $\Pi^O(c,c)/2$ at these input prices; i.e., $\rho_j(\overline{w}^T, \overline{w}^T) \geq \Pi^O(c,c)/2$, for $j \in \{1,2\}$. In this case, the equilibrium fixed fee is strictly positive, and ensures that $\pi_j = \Pi^O(c,c)/2$.

The second case arises if $\rho_j(\overline{w}^T, \overline{w}^T) < \Pi^O(c,c)/2$. In this scenario, without the opportunity to offer slotting allowances, input prices below $\overline{w}^T$ are necessary to retain retailers in the two-brand candidate. Note, however, that under Assumption 2 (quasi-concavity of industry profits in (w,w)), manufacturers always have incentives to keep input prices as high as possible to soften the competitive pressure in the industry. The optimal way of making the retailers indifferent is, therefore, to set $\overline{F}_{ij}^T = 0$ and reduce $\overline{w}_{ij}^T$ until $\pi_j = \Pi^O(c,c)/2$. The manufacturers would ideally want to offer higher input prices and/or fixed fees, but that would lead retailers to induce a one-brand structure. Yet, as $\overline{w}_{ij}^T$ must be weakly above c (otherwise $\pi_i < 0$ given that $\overline{F}_{ij}^T = 0$), we again have that $\overline{w}_{ij}^T \in [c, \overline{w}^T]$ in the two-brand candidate.

To sum up, we see that none of the cases offer a profitable way to deviate from the candidate two-brand equilibrium to a one-brand structure. But then:

Lemma 3 *When slotting allowances are not feasible, at least one two-brand equilibrium always exists.*

When slotting allowances were feasible (in Section 4.2.1), we found that a two-brand equilibrium could exist only subject to conditions on industry profit levels. By contrast, when slotting allowances are not feasible, we see here that a two-brand equilibrium *always* exists. Intuitively, the key difference between the two regimes is that deviations to one-brand structures become less lucrative for manufacturers in the absence of slotting allowances, and this in turn increases the scope for two-brand equilibria to exist. We will discuss this insight in more detail in Section 4.4.

4.3.2 One-brand equilibria

Consider again a candidate one-brand equilibrium where A obtains distribution. (As in the case with slotting allowances, all arguments below apply also if B is the exclusive supplier.) In the candidate, A offers one-brand contracts $(\overline{w}^O_{Aj}, \overline{F}^O_{Aj}) = (c, 0)$ for $j \in \{1, 2\}$, and (out-of-equilibrium) two-brand contracts that give both retailers a profit of $\Pi^O(c, c)/2$, or marginally less. To determine whether this, in fact, constitutes a one-brand equilibrium, we follow the same line of reasoning as in Section 4.2.2.

First, manufacturer A must not be able to increase industry profits within the one-brand structure. Because $\overline{w}^O_{Aj} = c < w^O_M$ (see (9)), industry profits in the candidate are less than Π^O_M, and A would ideally like to raise the input prices. However, this is not possible when slotting allowances are not feasible and each retailer could earn $\Pi^O(c, c)/2$ by inducing a two-brand structure.

Second, B must not be able to profitably deviate to $\{B1, B2\}$. As A and B are equally profitable on their own, an argument parallel to the one in the previous paragraph implies that B would not be able to raise industry profits above $\Pi^O(c, c)$. To sum up, there are no profitable deviations from the candidate one-brand equilibrium to other one-brand structures.

Finally, there must be no profitable deviations to two-brand structures. To consider these incentives, take again the worst-case scenario for a deviation by A to $\{A1, B2\}$, where B sets an (out-of-equilibrium) two-brand input price $\widetilde{w}^T_{B2}$, that induces $p_{B2} = 0$ for any individually rational w_{A1}. The candidate equilibrium survives deviations to $\{A1, B2\}$ if the maximal bilateral profit of A and 1 given $p_{B2} = 0$ is lower than what A has to leave 1 in the one-brand structure, which now is $\Pi^O(c, c)/2$. We therefore have the following result:

Lemma 4 *When slotting allowances are not feasible, a one-brand equilibrium exists whenever*

$$\max_{w_{A1}} \{\rho_{A1}(w_{A1}, \widetilde{w}^T_{B2})\} \leq \frac{\Pi^O(c, c)}{2}, \tag{11}$$

and does not exist if the converse is true.

For one-brand equilibria, the key difference between the regimes with and without slotting allowances is the amount of retail profits needed to deter deviations to two-brand structures. We will discuss the implications this has for product variety in Section 4.4, and also illustrate the point further in Section 4.5.

4.4 The impact of slotting allowances on product variety

We can now identify the impact of slotting allowances on product variety by comparing the equilibrium outcomes in the preceding sections. From Lemma 1 and Lemma 3, we get our first main result.

Proposition 1 *Slotting allowances (weakly) reduce product variety: When $\overline{\Pi}^T \leq \Pi_M^O$, a single manufacturer obtains distribution if slotting allowances are feasible, whereas both A and B can always obtain distribution if they are not feasible.*

Starting from a situation in which both brands are sold, the manufacturers enter a bidding war on their one-brand contracts to sign up both retailers and thereby foreclose the rival brand from the market. Whether slotting allowances are feasible or not dictates the manner in which this bidding war takes place. It is, therefore, the key determinant of whether vertical foreclosure may arise in equilibrium.

The intuition is as follows. If slotting allowances are feasible, manufacturers can set their one-brand input prices high to offset retail competition, and then bid for the retailers' patronage by offering negative fixed fees. In other words, slotting allowances enable a single manufacturer to disentangle the maximization and redistribution of industry profits. Conversely, these two objectives conflict when slotting allowances are not feasible. In this case, manufacturers can only sway the retailers by cutting input prices. Compared with receiving large fixed fees, however, this is less attractive for the retailers, as lower input prices intensify downstream competition. It follows that the equilibrium industry profit in a one-brand structure may be larger than in a two-brand structure when slotting allowances are feasible – this is so whenever $\overline{\Pi}^T < \Pi_M^O$ – whereas it is strictly smaller when they are not.[24] Consequently, it is only when slotting allowances can be used that retailers may be persuaded to stock the same brand rather than different brands.

Proposition 1 can also be understood as a result of the inability of rival manufacturers to internalize interbrand competition. When the retailers sell differentiated brands rather than the same brand, they could potentially benefit both from being more distant rivals and facing consumers with a higher aggregate willingness to pay. Yet, a two-brand equilibrium may be less profitable than the best one-brand equilibrium because each manufacturer succumbs to the temptation of giving her retailer a discounted input price in order to steal business from the rival channel.

[24] Note that, because exclusion raises industry profits in our model, the mechanism is robust to the so-called "Chicago-school critique," which Inderst and Shaffer (2010, p. 710) summarize as "if the exclusion of competitors reduces industry profits, why can the firms not do better?"

The exclusionary effect of slotting allowances that we highlight in this paper applies as long as manufacturers engage in this type of business stealing. As we discuss in Section 5.2 and Section 5.3, the business stealing incentive persists if retailers compete by choosing quantities, and cannot easily be corrected by more elaborate ways of vertical contracting.

Furthermore, the inequality in Proposition 1 is more likely to be violated when retailers compete fiercely. Intuitively, this is when the exclusive manufacturer's ability to fully offset downstream competition is most important. However, one has to be careful with the practical interpretation of this result. In our model, competition is fierce when consumers see the two retail outlets as close substitutes. In practice, competition can also be fierce in the sense that there are many retailers in the market, for a given (symmetric) degree of substitution. In this case, however, foreclosure could be less of an issue when competition is fierce, simply because manufacturers have a higher number of potential outlets.[25] When it comes to slotting allowances and product variety, therefore, one may get two very different pictures from measuring downstream competition by market concentration (e.g., with HHI) or by closeness of competition (e.g., substitution in utility).

Finally, the result that slotting allowances tend to reduce product variety is reinforced by our analysis of one-brand equilibria. Specifically, by comparing (10) to (11) (see Lemma 2 and Lemma 4), we see that the one-brand equilibrium existence threshold is less stringent when slotting allowances are feasible than when they are not, because $\Pi_M^O > \Pi^O(c,c)$. In other words, we should expect the use of slotting allowances to expand the range of one-brand equilibria. Intuitively, it is more costly to persuade retailers to join a two-brand structure when their profits from remaining in the one-brand structure are larger, which happens when slotting allowances are feasible.

4.4.1 Division of profits

As noted above, there may be a multiplicity of equilibria in our model. In particular, while equilibrium input prices (and, by extension, retail prices, industry profits, and consumer surplus) are uniquely defined, various profit distributions supported by different fixed transfers (both slotting allowances and positive franchise fees) may be part of an equilibrium. Below, we provide some further remarks on this issue.

[25] For example, Inderst and Shaffer (2007) consider a model with two rival manufacturers but only one retailer (who can stock one product) per final market. In this setting, there is always exclusion. Adding a second retailer at least opens for the possibility that both manufacturers get distribution.

Consider first two-brand equilibria. When slotting allowances are feasible and $\overline{\Pi}^T \geq \Pi_M^O$, $\pi_j = \Pi_D^O - \overline{\Pi}^T/2 = \Pi_M^O - \overline{\Pi}^T/2$ are the *minimal* retail profits that can support such an equilibrium. However, there also exists a range of equilibria where retailers are strictly better off. Take, for example, a candidate two-brand equilibrium with market structure $\{A1, B2\}$ where retailers have been offered strictly more than π_j, and where out-of-equilibrium one-brand tariffs make them indifferent between $\{A1, A2\}$ and $\{A1, B2\}$. No firm has a profitable deviation from such a candidate, as long as it entails non-negative profits for manufacturers. In fact, *any* profit distribution where $\pi_j \in [\Pi_M^O - \overline{\Pi}^T/2, \overline{\Pi}^T/2]$ and corresponding manufacturer profits $\pi_i \in [0, \overline{\Pi}^T - \Pi_M^O]$ can be part of a two-brand equilibrium when slotting allowances are feasible. Furthermore, when slotting allowances are not feasible, it is possible that $\overline{\Pi}^T/2 > \Pi_D^O$ because deviations to one-brand structures are less attractive. If this inequality holds, there may even exist two-brand equilibria where manufacturers can extract the entire two-brand industry profit and leave zero profit to the retailers.

When it comes to one-brand equilibria, the firms have somewhat less leeway in dividing their profits. In particular, as noted in Section 4.2.2 and Section 4.3.2, any one-brand equilibrium with or without slotting allowances must involve zero upstream profits. This follows because, in such an equilibrium, the "excluded" manufacturer (B) must not be able to profitably offer each retailer a marginally better deal (than A) and thereby induce her own one-brand equilibrium. The manufacturers can, thus, be seen as offering not specific products or contract terms but rather levels of net profit. In other words, the manufacturers effectively compete in "profit space," where they are homogeneous despite the level of differentiation between their products.[26]

4.4.2 Equilibrium refinement and a ban on slotting allowances

Continuing the theme of equilibrium multiplicity, note that subgame perfect Nash equilibria of the one-brand and two-brand type can coexist in our model, both when slotting allowances are feasible and when they are not. By following Bernheim et al. (1987) in considering perfectly coalition-proof Nash equilibria (PCPNE), we can refine the set of equilibria and get a sharper prediction. As explained by Bernheim

[26] This logic is familiar from other models of exclusion in vertically related markets (e.g., Calzolari et al., 2020). The underlying principle is similar to competition in "utility space," as advanced by Armstrong and Vickers (2001). Alternatively, the property of zero upstream profits can be understood as a consequence of the fact that each manufacturer's one-brand structure offers no *incremental* contribution (relative to the rival's one-brand structure) to the total industry profits (see O'Brien and Shaffer, 1997).

and Whinston (1998), the set of PCPNE contains those outcomes that are Pareto-undominated for the manufacturers. Focus on PCPNE can, therefore, be justified because when manufacturers move first, they should be able to "steer" the market in their preferred direction. We can then state the following result.

Proposition 2 *If two-brand equilibria exist, then any PCPNE is a two-brand equilibrium. When focusing on PCPNE, a ban on slotting allowances, therefore,* guarantees *that both A and B obtain distribution.*

The first part of Proposition 2 follows from the fact that the joint profit of the manufacturers is zero in any one-brand market structure. The second part follows from Lemma 3, which says that there always exists at least one two-brand equilibrium when slotting allowances are prohibited. Proposition 2 clarifies and reinforces the message implicit in Proposition 1, that a ban on slotting allowances tends to raise product variety in the market.

Considering a ban on slotting allowances is particularly relevant in the context of grocery markets, where this policy is often a key part in regulations of so-called "unfair trading practices," which have been widely implemented in recent years. For example, the European Commission's Directive 2019/633, in effect since April 2019, concerns unfair trading practices in vertical relationships in the food supply chain across all EU member states.[27] These regulations aim to protect small and medium-sized suppliers in industries where powerful retailers act as gatekeepers to final consumers. Our analysis – and Proposition 2 in particular – suggests that also large manufacturers, who in principle have significant bargaining power vis-à-vis retailers, can benefit from these laws. However, it should also be noted that prohibitions on slotting allowances is, and has been, highly controversial, in large part due to uncertainty over the welfare effects of such policies. This is a topic that we analyze in the next section.

4.5 Quasilinear utility

To illustrate the results above, and to get a sense of the welfare implications of a ban on slotting allowances, we now consider a version of the model with a representative consumer. Specifically, we suppose that this consumer has quasilinear preferences,

[27] Some country-specific regulations go even further. For example, the French Code of Commerce (article L 442-6) restricts the use of lump-sum payments for shelf space in all business-to-business relations.

$V = y + U$, in which y is a composite good with $p_y = 1$, and

$$U = \sum_{ij} q_{ij} - \frac{1}{2} q_{ij}^2 - \sum_{i \in \{A,B\}} d q_{i1} q_{i2} - \sum_{j \in \{1,2\}} b q_{Aj} q_{Bj} - bd(q_{A1} q_{B2} + q_{A2} q_{B1}) \tag{12}$$

is a quasilinear quadratic utility function.[28] In (12), the q's are quantities consumed, while $b \in (0,1)$ and $d \in (0,1)$ measure the consumer's willingness to substitute between brands and retail outlets, respectively. Given (12), maximization of V with respect to quantities, and subject to budget $y + \sum_{ij}(p_{ij} q_{ij}) = I$, gives a system of inverse demands, $p_{ij} = 1 - q_{ij} - b q_{hj} - d q_{ik} - bd q_{hk}$, for $i \neq h \in \{A, B\}$ and $j \neq k \in \{1, 2\}$. Focusing on the two-brand case $\{A1, B2\}$, we set $q_{A2} = q_{B1} = 0$ and invert this system to get

$$D_{ij}(p_{ij}, p_{hk}) = \frac{1 - bd - p_{ij} + bd p_{hk}}{(1 - bd)(1 + bd)}. \tag{13}$$

Similarly, for the one-brand case $\{A1, A2\}$, we set $q_{B1} = q_{B2} = 0$ and get

$$D_{Aj}(p_{Aj}, p_{Ak}) = \frac{1 - d - p_{Aj} + d p_{Ak}}{(1 - d)(1 + d)}. \tag{14}$$

For simplicity, we here also set $c = 0$.

Under these circumstances, it is straightforward to verify that $p_M^T = p_M^O = p_M = 1/2$ maximizes the industry profit in either market structure. The corresponding levels of industry profits are

$$\Pi_M^T = \frac{1}{2(1 + bd)} \quad \text{and} \quad \Pi_M^O = \frac{1}{2(1 + d)}. \tag{15}$$

In line with Assumption 1, we have $\Pi_M^T \in (\Pi_M^O, 2\Pi_M^O)$ because $b, d < 1$. It is also easy to show that the input prices that induce $p_M = 1/2$ and thereby Π_M^T and Π_M^O are

$$w_M^T = \frac{bd}{2} \quad \text{and} \quad w_M^O = \frac{d}{2}, \tag{16}$$

where $w_M^T < w_M^O$ because $b < 1$. Intuitively, and in line with Assumption 3, there is less competition that needs to be dampened when the retailers sell differentiated brands. In the one-brand structure, when slotting allowances are feasible, the manufacturer can induce the industry profit $1/(2(1 + d))$ as given by (15) with the input price $\overline{w}^O = d/2$ from (16). By contrast, in a two-brand structure, we find that the

[28] Such utility functions are extensively used in IO-models, see Choné and Linnemer (2020) for a comprehensive treatment.

(candidate) equilibrium input price is

$$\overline{w}^T = \frac{(1-bd)b^2d^2}{4-2bd-b^2d^2} < w_M^T \tag{17}$$

and the industry profits are

$$\overline{\Pi}^T = \frac{4(1-bd)(2-b^2d^2)}{(4-2bd-b^2d^2)^2} < \Pi_M^T. \tag{18}$$

This lays the groundwork for Figure 1 below:

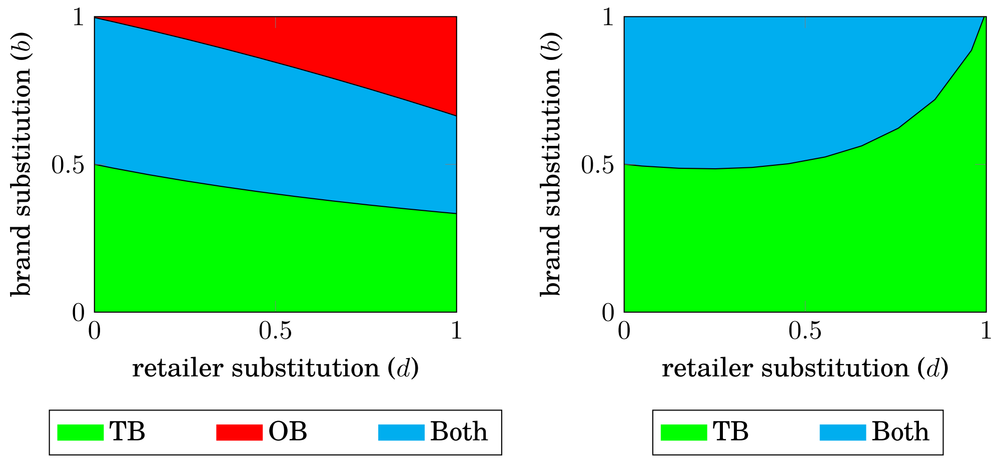

(a) Slotting allowances are feasible.

(b) Slotting allowances are not feasible.

Figure 1: Existence of two-brand (TB) and one-brand (OB) equilibria for all combinations of brand and retailer substitutability, when slotting allowances are feasible (Figure 1(a)) and not feasible (Figure 1(b)). In Figure 1(a), the two-brand existence condition ((4), upper line) uses the profit levels given in (15) and (18). To plot the one-brand existence conditions (i.e., (10) and (11)), we use the fact that when $p_{B2} = 0$, the optimal p_{A1} is $(1-bd)/2$ both with and without slotting allowances, which gives $\rho_{A1} = (1-bd)/(4+4bd)$. When slotting allowances are feasible, we subtract $\Pi_O^M/2 = 1/(4+4d))$ and solve for $b = 1/(2+d)$ (bottom line in Figure 1(a)). When slotting allowances are not feasible, we subtract $\Pi^O(0,0)/2 = (1-d)/((1+d)(2-d)^2)$ and obtain $b = (4-3d+d^2)/(8-4d-3d^2+d^3)$ (line in Figure 1(b)).

In Figure 1(a), where slotting allowances are feasible, two-brand equilibria do not exist when the substitution rates between brands and retailers are sufficiently high. To sustain one-brand equilibria in this region, the exclusive manufacturer pays both retailers a slotting allowance of $\overline{F}^O = -d/(4+4d) < 0$. Further down and left in Figure 1(a), where brands and retailers are more differentiated, two-brand equilibria exist and manufacturers may or may not pay slotting allowances depending on the size of bd and which equilibrium we focus on.[29] By contrast, in Figure 1(b) where slotting allowances are not feasible, two-brand equilibria exist for any combination of (b, d) and the range of parameter values that support one-brand equilibria is smaller than in Figure 1(a).

To illustrate further, we may combine the above to uncover the *change* in sets of equilibria following a ban on slotting allowances:

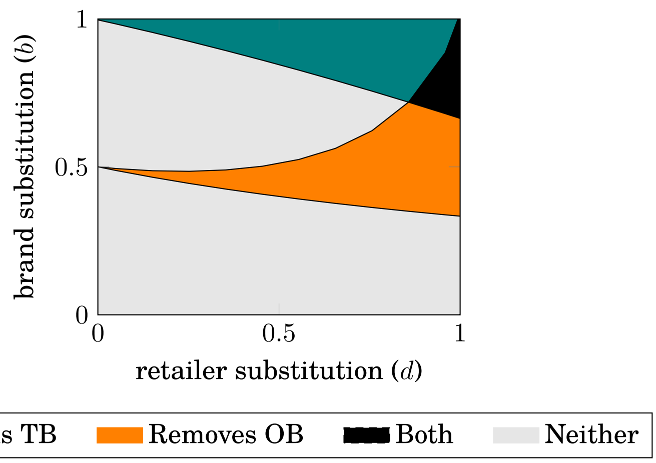

Figure 2: Effect of a ban on slotting allowances on the sets of two-brand (TB) and one-brand (OB) equilibria in the (b, d)-space.

For parameter values in the upper area of Figure 2, a ban on slotting allowances adds two-brand equilibria and thereby opens the possibility that both brands become available to consumers. In the middle area, a ban removes one-brand equilibria and thereby prevents that one brand is excluded when two-brand equilibria exist. Finally, the situation is perhaps at its clearest in the upper right area of Figure 2, where a ban completely flips the outcome: With slotting allowances, there is exclusion in all equilibria, whereas without slotting allowances, there only exist equilibria with full product variety.

By increasing product variety – something consumers tend to appreciate – a ban on slotting allowances can have a positive impact on consumer surplus. On the other hand, it is also possible that such a policy raises retail prices, which has a negative effect on consumer surplus. Thus, in general, in our model the consumer welfare effect of a ban on slotting allowances is ambiguous. However, the effect is easily signed under the assumptions in this section.

[29] For example, if we focus on PCPNE, then slotting allowances arise in two-brand equilibria (i.e., $\overline{F}^T < 0$) whenever $b((4 - bd - b^2d^2)3b^2d^2 - 8(2 - bd)) + 2(1 - bd)(2 + bd)(2 - bd)(2 - b^2d^2) < 0$, which is satisfied if bd is not too small.

Proposition 3 *When consumer preferences are represented by the utility function (12), a ban on slotting allowances that shifts the industry from a one-brand equilibrium to a two-brand equilibrium strictly raises consumer surplus.*

Proposition 3 follows from straightforward calculations using (12), (16), (17), as well as the retailers' best-response functions. Intuitively, the utility function in (12) displays a certain "love for variety," which captures the natural idea that, on average, consumers are better off when they can choose from a broader product range. In addition, retail prices in the "semi-competitive" two-brand equilibria are below the monopoly level, which then ensures the result.

5 Robustness

In this section, we argue that our main insights carry over to the cases in which 1) retailers can exit the market, 2) retailers compete in quantities, and 3) manufacturers can use more sophisticated vertical contracts. Our focus here will lie on the existence of two-brand equilibria.

5.1 Retail exit

In Section 4, we focused on two-brand structures and "triangular" one-brand structures, but did not consider structures with only one active retailer. It is not *a priori* obvious that this restriction is without loss of generality. To this end, we now consider instead a version of the game in which each retailer may also reject *both* received offers and thereby exit the market.

We start by noting that, as long as retailers earn non-negative profits in all candidate two-brand equilibria, they do not have incentives to leave the market voluntarily. However, the retailers may still be forced out of the market following deviations by manufacturers that actively seek to induce this outcome. In the following, therefore, we consider the incentives of one manufacturer to induce the exit of one retailer, and thereby establish a market structure with a single active channel, i.e., a successive monopoly.

To gain intuition, let us consider a candidate two-brand equilibrium with market structure $\{A1, B2\}$, where each retailer earns a profit of τ both within the two-brand structure and in the (out-of-equilibrium) one-brand case. Furthermore, these profits are supported by (out-of-equilibrium) one-brand terms, offered by both manufacturers, denoted by (w_τ^O, F_τ^O). Here, the fixed fee is $F_\tau^O = \rho_j(w_\tau^O, w_\tau^O) - \tau$, where $\rho_j(w_\tau^O, w_\tau^O)$ is retailer j's one-brand flow profit. If slotting allowances are not feasible (i.e., $F_\tau^O \geq 0$), we restrict attention to $\tau \leq \rho_j(w_\tau^O, w_\tau^O)$.

We now ask whether manufacturer A can deviate in a way that, on the one hand, pushes retailer 2 and manufacturer B out of the market, and, on the other hand, retains the patronage of retailer 1. To achieve this, A can give 2 one-brand terms that are bad enough to cause 2 to reject A's offer, and, at the same time, incentivize 1 to set a very low p_{A1} in the two-brand structure, possibly inducing 2 to decline B's offer. The former is easily done by offering 2 a high fixed fee. The latter hinges on whether A can, in fact, find a w_{A1}^T that reduces 2's profits below zero in the two-brand case, thus causing 2 to exit the market. Below, we assume that this is indeed *possible* and then examine whether the deviation to successive monopoly can be *profitable* for A.[30]

If 2 were to leave, the best-response of 1 could be either to go with A's deviation offer, B's original offer, or to reject both offers (and, thus, earn zero profits). Accepting B's offer would give 1 a profit of $\rho_1(w_\tau^O, \infty) - F_\tau^O$. Here, $\rho_1(w_\tau^O, \infty)$ denotes retailer 1's flow profit in the successive monopoly structure with B, with the convention that $w_{A2} = \infty$ means that A and 2 are inactive. To retain the patronage of 1, A must, therefore, offer one-brand terms (w_{A1}^O, F_{A1}^O) for the successive monopoly case ensuring that

$$\rho_1(w_{A1}^O, \infty) - F_{A1}^O \geq \rho_1(w_\tau^O, \infty) - F_\tau^O. \tag{19}$$

Because $F_\tau^O = \rho_j(w_\tau^O, w_\tau^O) - \tau$, (19) means that the fixed fee is bounded above by

$$F_{A1}^O \leq \rho_1(w_{A1}^O, \infty) - \tau - \rho_1(w_\tau^O, \infty) + \rho_1(w_\tau^O, w_\tau^O). \tag{20}$$

From (20), it follows then that A's profit after deviating satisfies

$$\rho_A(w_{A1}^O, \infty) + F_{A1}^O \leq \rho_{A1}(w_{A1}^O, \infty) - \tau - \Delta_1(w_\tau^O), \tag{21}$$

where $\rho_{A1}(w_{A1}^O, \infty) = \rho_A(w_{A1}^O, \infty) + \rho_1(w_{A1}^O, \infty)$ is the channel (and industry) profit in the successive monopoly structure, and $\Delta_1(w_\tau^O) \equiv (\rho_1(w_\tau^O, \infty) - \rho_1(w_\tau^O, w_\tau^O))$.

Notice now that the difference $\Delta_1(w_\tau^O)$ is positive: The flow profit of one retailer (here, 1) is increasing in the input price of the rival retailer (here, w_{A2}). Intuitively, a

[30]Note that the deviations discussed above are not "costly" for A in the sense that they only concern out-of-equilibrium terms; specifically, the two-brand input price to retailer 1 (for whom one-brand terms would apply in successive monopoly) and the one-brand fixed fee to retailer 2 (who would be inactive).

higher w_k, $k \neq j$, leads to a higher p_k, which both diverts some customers to retailer j and permits a higher p_j. Furthermore, observe that if $\Delta_1(w_\tau^O)$ is sufficiently large, then the right hand side of (21) is negative (for any τ), which means that A is unable to earn non-negative profits following the deviation. Recalling that retailer flow profits are of the form $(p(w,w) - w)D(p(w,w), p(w,w))$, we can use the Envelope Theorem to obtain

$$\frac{\partial \Delta_1(w_\tau^O)}{\partial w_\tau^O} = -\left[D_1(p_\tau, \infty) - D_1(p_\tau, p_\tau)\right] - \left[(p_\tau - w_\tau^O)\frac{\partial D_1(p_\tau, p_\tau)}{\partial p_2}\frac{dp_\tau}{dw_\tau^O}\right] < 0 \tag{22}$$

where $p_\tau = p_{ij}(w_\tau^O, w_\tau^O)$, and where we have used the symmetry of demand functions to write $D_{A1} = D_{B1} = D_1$. The sign of (22) follows from the positive cross-price effect $\partial D_{ij}/\partial p_{hk} > 0$, which ensures the positivity of both bracketed terms (for $p_\tau > w_\tau^O$).

Importantly, then, (22) implies that $\Delta_1(w_\tau^O)$ can be made very large by setting w_τ^O very low. Also, as a very low input price would generally not be combined with a negative fixed fee, such a reduction of w_τ^O is possible both when slotting allowances are feasible and not feasible. We can, therefore, conclude that the profit A can obtain by deviating from the candidate two-brand equilibrium to a successive monopoly structure can be made negative by sufficiently low out-of-equilibrium one-brand input prices. Thus, if two-brand equilibria exist when retailers cannot exit the market (as in Section 4), two-brand equilibria will also exist when they can.

Intuitively, deviations to successive monopoly are not profitable for manufacturers because the out-of-equilibrium one-brand contract of the non-deviating manufacturer is always an option for the active retailer. This means, for instance, that any such deviation by A must provide 1 with at least what the offer from B yields. Furthermore, as retail outlets are substitutes, a retailer sells more in successive monopoly than in a one-brand structure for a given set of input prices. A small reduction in the input price, therefore, raises the retailer's profit more in the former case. Thus, when the out-of-equilibrium one-brand input prices are sufficiently low, the deviating manufacturer is unable to compensate the retailer without earning negative profits.

5.2 Retail quantity competition

Suppose now that, instead of setting prices, the retailers compete by choosing quantities to sell in the final market. How do slotting allowances affect the scope for two-brand equilibria to exist in this case?

We note first that the two-brand existence condition given by (4) still applies. Slotting allowances reduce product variety if they restrict the scope for this condition to be satisfied (that is, for any profit given to the retailers in a candidate two-brand equilibrium).

Now, whether slotting allowances are feasible or not has no impact on $\overline{\Pi}^T$, which is the industry profit in a two-brand structure. This follows from the fact that, as quantities are strategic substitutes (which holds under mild conditions, see Vives, 1999), two-brand equilibrium input prices lie *below* upstream marginal costs, as each manufacturer has an incentive to subsidize its retailer in order to expand output of its own brand at the expense of the rival's brand (see, e.g., Irmen, 1998). Accordingly, with input prices below cost, slotting allowances are not used in any two-brand structure. On the other hand, and just as under retail price competition, the possibility of paying slotting allowances does affect Π_D^O. When slotting allowances are feasible, an exclusive manufacturer can always induce Π_M^O by setting input prices that incentivize retailers to choose the industry-profit maximizing quantities (Inderst and Shaffer, 2010),[31] and redistribute this profit by the means of lump-sum transfers. When slotting allowances are prohibited, this strategy no longer works and Π_D^O must, therefore, be lower (although not necessarily identical as under price competition). To sum up, a ban on slotting allowances expands the range where two-brand equilibria exist when retailers compete in quantities, just as it did under retail price competition.

While the mechanism for how slotting allowances affect product variety is qualitatively the same in the two cases, a change from price to quantity competition may alter the quantitative effect of a ban on slotting allowances. On the one hand, quantity competition gives manufacturers an extra incentive to cut input prices, as explained above. On the other hand, quantity competition tends to be less fierce than price competition (see Vives, 1999, p. 155). Thus, although quantity competition gives rise to lower input prices, it may not give lower retail prices and two-brand industry profits, because retailers can sustain higher margins for given costs. In other words, going from price to quantity competition in the retail market generally has an ambiguous effect on $\overline{\Pi}^T$, the two-brand existence condition (4), and ultimately

[31] The required one-brand input price under quantity competition – corresponding to (9) under price competition – can be written as

$$w_{ij} = c + \frac{\partial P_{ik}(q_{ij}^*, q_{ik}^*)}{\partial q_{ij}} q_{ik}^*,$$

in which (q_{ij}^*, q_{ik}^*) are the optimal quantities and P_{ik} is the inverse demand function, $k \neq j = \{1, 2\}$.

the degree to which a ban on slotting allowance raises product variety.

To gain further intuition, we can again use the quasilinear utility approach from Section 4.5. Figure 3 below illustrates the existence of two-brand equilibria with and without slotting allowances under retail quantity competition, and compares this to the case of price competition.

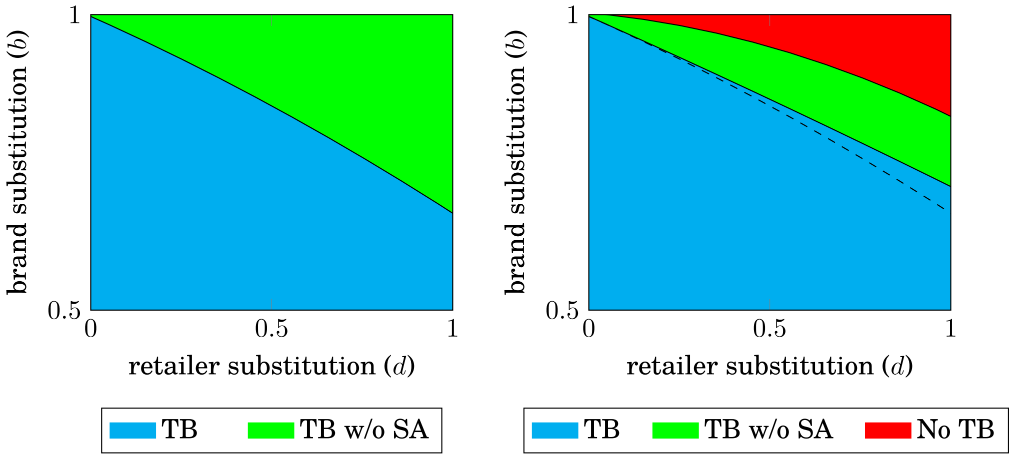

(a) Retail price competition.

(b) Retail quantity competition.

Figure 3: Existence of two-brand equilibria (TB) with and without slotting allowances (SA) with competition in prices (Figure 3(a)) and quantities (Figure 3(b)) for $b \geq 0.5$. (For $b < 0.5$, existence is unaffected by going from price to quantity competition.) In Figure 3(b), the dashed line is the two-brand existence condition (4) under price competition. The lower solid line is (4) under quantity competition, where we use the fact that inverse demand $p_{ij} = 1 - q_{ij} - bdq_{hk}$ yields $\overline{w}^T = -b^2d^2/(4 + 2bd - b^2d^2)$ and an industry profit of $(8 - 4b^2d^2)/(4 + 2bd - b^2d^2)^2$. (Note that, as $\overline{w}^T < c = 0$, we require here that retailers have to sell all they order to rule out the artificial case where they buy infinitely many units to collect the subsidy.) The upper solid line is a second existence condition, specific to quantity competition (see explanation below).

Figure 3, and part (b) in particular, shows that under retail quantity competition, there is a range of fairly high (b, d) in which two-brand equilibria exist if and only if slotting allowances are prohibited, see the middle area in Figure 3(b), marked "TB w/o SA." Moreover, relative to the case of retail price competition, analyzed in Section 4.5 and shown in Figure 3(a), there are two differences.

First, (4) is slightly less stringent under quantity competition, which means that two-brand equilibria exist when slotting allowances are feasible for a slightly larger share of the parameter space. Intuitively, this follows from the fact that $\overline{\Pi}^T$ is marginally larger when retailers compete in quantities, reflecting that the above-mentioned effect of weaker downstream competition for given input prices, just dominates the opposing effect of lower input prices.

Second, under quantity competition and combinations of (b, d) sufficiently close to unity, two-brand equilibria do not exist even when slotting allowances are prohibited, see the upper right area in Figure 3(b), marked "No TB." (By contrast, with price competition, two-brand equilibria always exist when slotting allowances are not feasible (see Lemma 3).) In this area, even with zero fixed fees, there exists a pair of one-brand input prices that give $\pi_i > 0$ for the active manufacturer, and $\pi_j = \overline{\Pi}^T/2$ to both retailers. But then, as two-brand equilibria with high retail profits are generally those that are most likely to exist, and as $\overline{\Pi}^T/2$ is the maximal retail profit in any (candidate) two-brand equilibrium, we can conclude that there always is a profitable deviation to a one-brand structure.[32] Intuitively, one-brand structures without slotting allowances are relatively more attractive under quantity competition than under price competition because the former mode is less intense for given input prices.

5.3 Other vertical contracts

For vertical foreclosure to occur in our model, the active manufacturer must be able to compensate both retailers for the loss they incur by selling the same brand rather than differentiated brands. As such, exclusion by the use of slotting allowances may be feasible only if a single manufacturer is relatively better than competing manufacturers at softening competitive pressure in the industry. Whether this is the case clearly depends on the set of admissible contracts.

In Section 4, we found that, at least when slotting allowances are feasible, two-part tariffs are sufficient for a single manufacturer to induce the fully integrated outcome, but insufficient for competing manufacturers to do the same. Thus, when opening for more sophisticated contracts, it is primarily in the two-brand structures that things could change. Below, we discuss whether the manufacturers may be able to soften more inter- and intrabrand competition by using other vertical contracts.

Note first that manufacturers would actually lessen their ability to soften competition by vertically integrating with one retailer each. Under vertical integration, it would be as though all input prices were at marginal cost, with the manufacturers competing directly in retail prices in the final market. Indeed, one of the main

[32] Specifically, such profitable deviations exist whenever $\sqrt{4-2b^2d^2}/(4+2bd-b^2d^2)(2+d)^{-1} < 1$. The upper solid line in Figure 3(b) plots this expression with an equality sign.

insights in the seminal paper by Bonanno and Vickers (1988) is that vertical integration yields strictly lower profits than vertical separation, where manufacturers can commit to softening retail price competition with input prices above marginal cost. By a similar logic, the manufacturers could do no better than the two-part tariff outcome by taking direct control of the retailers' pricing decisions through resale price maintenance (see Gabrielsen et al., 2018).

Fershtman et al. (1991) show that if the contract of one manufacturer-retailer pair can be contingent on the strategies of the rival retailer, then virtually *any* outcome can be supported in equilibrium. In particular, with such a contract clause, manufacturers could achieve $\overline{\Pi}^T = \Pi_M^T$, and two-brand equilibria would always exist irrespective of whether slotting allowances were feasible or not. While effective in theory, we believe that contracts with such a "horizontal flavor" are less relevant in practice because they conflict with current antitrust laws and, therefore, are liable to fines. For example, one contract clause in this category is a closed territory distribution agreement, whereby each retailer gets the exclusive right to serve all consumers within a given territory. However, the European Commission (2010, p. 18) lists as a hardcore restriction the "market partitioning by territory or by customer group [...] such as the obligation not to sell to certain customers or to customers in certain territories or the obligation to refer orders from these customers to other distributors."

Several supply contracts share the property that the quantity-dependent part of a retailer's payment is a non-linear function of the ordered quantity (e.g., retroactive rebates or premiums). While this set of contracts is potentially very large, we can draw on the work of Kühn (1997), who analyzes a model with two manufacturers selling their products through exclusive retailers. When demand is linear and products are differentiated – such as the setting analyzed in Section 4.5 – Kühn (1997) shows that the optimal contracts are quadratic functions of quantity. However, although such contracts outperform regular two-part tariffs, it is easily verified (again, in the context of Section 4.5) that the profit gains are generally small, and that the manufacturers are still unable to induce the industry profit maximizing outcome (i.e., $\overline{\Pi}^T < \Pi_M^T$). Thus, even with quadratic input prices, the two-brand existence condition given by (4) may be violated, and slotting allowances can be used to sustain vertical foreclosure.

6 Conclusion

In this paper, we have examined the effect of slotting allowances on product variety in a vertically related market. Prior work has shown that a dominant manufacturer may use slotting allowances to exclude brands from weaker rivals who lack the ability to offer such payments. On the other hand, when manufacturers are on an equal footing and compete head-to-head for shelf space (as often appears to be the case in practice), scholars and policy makers have gone a long way in dismissing the possibility of exclusionary effects.

This paper has challenged this conventional wisdom by showing that slotting allowances facilitate vertical foreclosure in a setting where two symmetric manufacturers of differentiated brands bid for the right to supply two competing retailers. In our model, exclusion is driven not by asymmetries between manufacturers but by the joint presence of upstream and downstream competition. Furthermore, using a quasilinear utility model, we have shown that a ban on slotting allowances tends to raise product variety and increase consumer surplus.

Slotting allowances remain a subject of academic debate and scrutiny from policy makers. Our paper lends support to a stricter legal treatment of slotting allowances. In particular, by illustrating how slotting allowances restrict product variety even in bidding wars between evenly matched manufacturers, it expands the range of market circumstances under which courts and practitioners should emphasize exclusionary effects from such payments. This anti-competitive effect must then, as usual, be weighed against other, pro-competitive effects on a case-by-case basis.

References

Anderson, S. P. and De Palma, A. (1992), 'Multiproduct firms: A nested logit approach', *The Journal of Industrial Economics* **40**(3), 261–276.

Armstrong, M. and Vickers, J. (2001), 'Competitive price discrimination', *The RAND Journal of economics* **32**(4), 579–605.

Asker, J. and Bar-Isaac, H. (2014), 'Raising retailers' profits: on vertical practices and the exclusion of rivals', *American Economic Review* **104**(2), 672–686.

Bernheim, B. D., Peleg, B. and Whinston, M. D. (1987), 'Coalition-proof nash equilibria i. concepts', *Journal of Economic Theory* **42**(1), 1–12.

Bernheim, B. D. and Whinston, M. D. (1985), 'Common marketing agency as a device for facilitating collusion', *The RAND Journal of Economics* **16**(2), 269–281.

Bernheim, B. D. and Whinston, M. D. (1998), 'Exclusive dealing', *Journal of Political Economy* **106**(1), 64–103.

Bonanno, G. and Vickers, J. (1988), 'Vertical separation', *The Journal of Industrial Economics* **36**(3), 257–265.

Bonnet, C. and Dubois, P. (2010), 'Inference on vertical contracts between manufacturers and retailers allowing for nonlinear pricing and resale price maintenance', *The RAND Journal of Economics* **41**(1), 139–164.

Bonnet, C. and Réquillart, V. (2013), 'Impact of cost shocks on consumer prices in vertically-related markets: the case of the french soft drink market', *American Journal of Agricultural Economics* **95**(5), 1088–1108.

Calzolari, G., Denicolò, V. and Zanchettin, P. (2020), 'The demand-boost theory of exclusive dealing', *The RAND Journal of Economics* **51**(3), 713–738.

Chambolle, C. and Molina, H. (2019), 'Buyer power, upstream bundling, and foreclosure', *Available at SSRN 3511692* .

Choi, J. P. and Stefanadis, C. (2018), 'Sequential innovation, naked exclusion, and upfront lump-sum payments', *Economic Theory* **65**(4), 891–915.

Choné, P. and Linnemer, L. (2020), 'Linear demand systems for differentiated goods: Overview and users guide', *International Journal of Industrial Organization* p. 102663.

Desai, P. S. (2000), 'Multiple messages to retain retailers: Signaling new product demand', *Marketing Science* **19**(4), 381–389.

Elberg, A. and Noton, C. (2019), 'What drives trade allowances? new evidence from actual payments to a big-box retailer', *Available at SSRN 3452974* .

European Commission (2010), 'Guidelines on vertical restraints'.

Fershtman, C., Judd, K. L. and Kalai, E. (1991), 'Observable contracts: Strategic delegation and cooperation', *International Economic Review* **32**(3), 551–559.

Fumagalli, C. and Motta, M. (2006), 'Exclusive dealing and entry, when buyers compete', *American Economic Review* **96**(3), 785–795.

Gabrielsen, T. S. and Johansen, B. O. (2015), 'Buyer power and exclusion in vertically related markets', *International Journal of Industrial Organization* **38**, 1–18.

Gabrielsen, T. S., Johansen, B. O. and Lømo, T. L. (2018), 'Resale price maintenance in two-sided markets', *The Journal of Industrial Economics* **66**(3), 570–609.

Hamilton, S. and Innes, R. (2017), 'Slotting allowances and retail product variety under oligopoly', *Economics Letters* **158**, 34–36.

Ho, K. and Lee, R. S. (2019), 'Equilibrium provider networks: Bargaining and exclusion in health care markets', *American Economic Review* **109**(2), 473–522.

Hristakeva, S. (2019), 'Vertical contracts with endogenous product selection: An empirical analysis of vendor-allowance contracts', *Available at SSRN 3506265* .

Inderst, R. and Shaffer, G. (2007), 'Retail mergers, buyer power and product variety', *The Economic Journal* **117**(516), 45–67.

Inderst, R. and Shaffer, G. (2010), 'Market-share contracts as facilitating practices', *The RAND Journal of Economics* **41**(4), 709–729.

Irmen, A. (1998), 'Precommitment in competing vertical chains', *Journal of Economic Surveys* **12**(4), 333–359.

Kelly, K. (1991), 'The antitrust analysis of grocery slotting allowances: The procompetitive case', *Journal of Public Policy & Marketing* **10**(1), 187–198.

Klein, B. and Wright, J. D. (2007), 'The economics of slotting contracts', *The Journal of Law and Economics* **50**(3), 421–454.

Kühn, K.-U. (1997), 'Nonlinear pricing in vertically related duopolies', *The RAND Journal of Economics* **28**(1), 37–62.

Lafontaine, F. and Slade, M. (2010), Inter-firm contracts: Evidence, *in* 'Handbook of organizational economics', New Jersey: Princeton University Press.

Lariviere, M. A. and Padmanabhan, V. (1997), 'Slotting allowances and new product introductions', *Marketing Science* **16**(2), 112–128.

Lømo, T. L. and Ulsaker, S. A. (forthc.), 'Lump-sum payments and retail services: A relational contracting perspective', *The Journal of Industrial Economics* .

Marx, L. M. and Shaffer, G. (2007), 'Upfront payments and exclusion in downstream markets', *The RAND Journal of Economics* **38**(3), 823–843.

Marx, L. M. and Shaffer, G. (2010), 'Slotting allowances and scarce shelf space', *Journal of Economics & Management Strategy* **19**(3), 575–603.

Mathewson, G. F. and Winter, R. A. (1984), 'An economic theory of vertical restraints', *The RAND Journal of Economics* **15**(1), 27–38.

Miklós-Thal, J., Rey, P. and Vergé, T. (2011), 'Buyer power and intrabrand coordination', *Journal of the European Economic Association* **9**(4), 721–741.

Miklós-Thal, J. and Shaffer, G. (2019), 'Input price discrimination by resale market', *Available at SSRN 3191951* .

O'Brien, D. P. and Shaffer, G. (1997), 'Nonlinear supply contracts, exclusive dealing, and equilibrium market foreclosure', *Journal of Economics & Management Strategy* **6**(4), 755–785.

Pagnozzi, M. and Piccolo, S. (2012), 'Vertical separation with private contracts', *The Economic Journal* **122**(559), 173–207.

Raff, H. and Schmitt, N. (2016), Retailing and international trade, *in* 'Handbook on the Economics of Retailing and Distribution', Edward Elgar Publishing.

Rey, P. and Stiglitz, J. (1995), 'The role of exclusive territories in producers' competition', *The RAND Journal of Economics* **26**(3), 431–451.

Rey, P. and Vergé, T. (2004), 'Bilateral control with vertical contracts', *The RAND Journal of Economics* **35**(4), 728–746.

Rey, P. and Whinston, M. D. (2013), 'Does retailer power lead to exclusion?', *The RAND Journal of Economics* **44**(1), 75–81.

Schutz, N. (2013), 'Competition with exclusive contracts in vertically related markets: An equilibrium non-existence result', *SFB/TR 15 Discussion Paper* .

Shaffer, G. (1991), 'Slotting allowances and resale price maintenance: a comparison of facilitating practices', *The RAND Journal of Economics* **22**(1), 120–135.

Shaffer, G. (2005), 'Slotting allowances and optimal product variety', *The BE Journal of Economic Analysis & Policy* **5**(1).

Sudhir, K. and Rao, V. R. (2006), 'Do slotting allowances enhance efficiency or hinder competition?', *Journal of Marketing Research* **43**(2), 137–155.

Sullivan, M. W. (1997), 'Slotting allowances and the market for new products', *The Journal of Law and Economics* **40**(2), 461–494.

UK Competition Commission (2008), *The supply of groceries in the UK market investigation*.

UK Office of Fair Trading (2013), 'Impact of reverse-fixed-payments on competition'.

US Federal Trade Commission (2001), 'Report on the federal trade commission workshop on slotting allowances and other marketing practices in the grocery industry'.

US Federal Trade Commission (2003), 'Slotting allowances in the retail grocery industry: Selected case studies in five product categories'.

Vives, X. (1999), *Oligopoly pricing: old ideas and new tools*, MIT press.

Weyl, E. G. and Fabinger, M. (2013), 'Pass-through as an economic tool: Principles of incidence under imperfect competition', *Journal of Political Economy* **121**(3), 528–583.

Institutt for økonomi
Universitetet i Bergen
Postboks 7800
5020 Bergen
Besøksadresse: Fosswinckels gate 14
www.uib.no/econ/