Slotting Allowances and Scarce Shelf Space

Leslie M. Marx

Fuqua School of Business
Duke University
Durham, NC 27708
marx@duke.edu

Greg Shaffer

Simon School of Business
University of Rochester
Rochester, NY 14627
shaffer@simon.rochester.edu

Slotting allowances are payments made by manufacturers to obtain retail shelf space. They are widespread in the grocery industry and a concern to antitrust authorities. A popular view is that slotting allowances arise because there are more products than retailers can profitably carry given their shelf space. We show that the causality can also go the other way: the scarcity of shelf space may in part be due to the feasibility of slotting allowances. It follows that slotting allowances can be anticompetitive even if they have no effect on retail prices.

1. Introduction

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.[1] To help supermarket retailers decide which products to carry, it has become common in recent years for them to put at least some of their shelf space up for bid and let manufacturers compete for their patronage. Those who offer the best deals obtain shelf space. Those who do not risk being excluded altogether. Typically, the manufacturers' deals include potentially large, upfront payments that are independent

We thank Yaron Yehezkel, two anonymous referees, and seminar participants at Duke, Texas A&M, Washington University in St. Louis, and the Fourth Annual International Industrial Organization Conference for helpful comments.

1. According to the Food Marketing Institute, the typical supermarket carries approximately 15,000 products. A typical superstore (defined as a larger supermarket with at least 40,000 square feet of total selling area) carries approximately 25,000 products. See *http://www.fmi.org/facts-figs/?fuseaction=superfact* (18 June, 2009).

Journal of Economics & Management Strategy, Volume 19, Number 3, Fall 2010, 575–603

of the retailers' subsequent quantity purchases. These upfront payments are referred to in the trade press as slotting allowances.

Slotting allowances have received a lot of attention from policy makers,[2] and there is a growing academic literature that seeks to identify their competitive effects. One concern is that large, dominant firms will buy up scarce shelf space in order to exclude their smaller rivals from the market.[3] According to this view, large firms may obtain shelf space not necessarily because they make better products, but because they are willing to pay more to protect their monopoly rents than small firms are willing to pay to be competitive (Shaffer, 2005).[4] An alternative view is that slotting allowances equate supply and demand and thus represent the market price of a scarce asset (shelf space) which, according to price theory, has optimal allocative properties (Sullivan, 1997). Under this view, the firms that offer the best deals are the ones whose products generate the largest private and social benefit, and thus only the most desirable products will be allocated shelf space.

Both viewpoints take as given the scarcity of shelf space and implicitly assume that slotting allowances arise in response to this scarcity. By contrast, in this paper, we show that the scarcity of shelf space itself may depend on the feasibility of slotting allowances in the sense that the retailer is more likely to limit its shelf space in an environment with slotting allowances than without. Although this departs from conventional thinking on slotting allowances, the idea we wish to explore is straightforward and comes from the theory of rent-shifting.[5] The choice of how much shelf space to build can be viewed as a strategic decision by the retailer. By limiting its shelf space, a retailer can often force manufacturers to compete more vigorously for its patronage, which then allows it to obtain better terms of trade.[6] Slotting allowances can help in this regard by making the transfer

2. See "Slotting: Fair for Small Business and Consumers? Hearing Before the Senate Committee on Small Business," 106th Congress, 1st Session 386 (1999), and the Federal Trade Commission reports, "Report on the Federal Trade Commission's Workshop on Slotting Allowances and Other Marketing Practices in the Grocery Industry," (2001), and "Slotting Allowances in the Retail Grocery Industry: Selected Case Studies in Five Product Categories," (2003).

3. See Canadian Competition Bureau (2002), FTC (2001), and FTC (2003).

4. A related concern is that large firms may have easier access to capital markets and/or be able to obtain better financial terms, giving them a tangible advantage over their smaller rivals when funding upfront payments.

5. The use of contracts to shift rents among vertically related firms was first studied by Aghion and Bolton (1987).

6. The result that the retailer might be able to increase its profit by committing to purchase from fewer suppliers can be found in O'Brien and Shaffer (1997). See also the papers by Dana (2004) and Inderst and Shaffer (2007). The mechanism at work here, however, is fundamentally different. In these other models, the contracting takes place simultaneously. In our model, the contracting occurs sequentially, which allows for the possibility of rent shifting.

of surplus more efficient. In contrast, when slotting allowances are infeasible, the main drawback of carrying fewer products—a smaller overall joint profit—then becomes more salient.

Our model has implications for antitrust policy. Typically the concern of antitrust authorities is whether a dominant firm can plausibly induce a smaller rival to exit by buying up scarce shelf space. Reasoning that the dominant firm will not buy up extra shelf space unless it can be assured that its rival's product is the one being excluded and not some unrelated product, the focus has been on the ancillary provisions that sometimes accompany contracts with slotting allowances. For example, a contract in which a firm promises to pay a retailer a lump-sum amount with no mention of its rivals' products is looked upon more favorably than a contract in which a firm offers a retailer a lump-sum payment in exchange for contingencies on the retailer's sales of rival products.[7] In our model, a prohibition of ancillary provisions in contracts that mention rivals would be ineffective, and it would not address the fundamental problem. This is because the source of the problem is the scarcity of the shelf space itself and not how the shelf space is allocated among the various products. By limiting its shelf space, the retailer links all products, whether or not they are otherwise independent in demand. Thus, irrespective of whether there are ancillary provisions in a particular manufacturer's contract that identify which firms are or are not competitors, some products will necessarily be excluded.

Slotting allowances do not affect consumer prices in our model, nor do they lead retailers to make inefficient choices of which products to carry.[8] Prices are not affected because manufacturers obtain their shelf space at the expense of unrelated products, leaving market conditions unchanged, and product choices are not distorted because, conditional on the available shelf space, the retailer will carry the same products whether or not slotting allowances are feasible. Slotting allowances are nevertheless anticompetitive in our model because they may contribute to a shortage of retail shelf space, leading to a reduction in the number and variety of products sold to consumers.

7. To give an example, consider a contract in which a firm offers to pay the retailer some amount of money upfront if the retailer agrees to purchase at least 80% of its requirements from the firm. Such a contract features both a slotting allowance and a market-share discount. In its 2001 report on slotting allowances, the Federal Trade Commission expressed concern that such contracts might lead to inefficient exclusion and ultimately higher prices for consumers.

8. This contrasts with the view that slotting allowances may lower consumer prices because they are an efficient mechanism for the retailer to equate the supply and demand of shelf space (cf. Sullivan, 1997), and the view that slotting allowances may raise consumer prices because in their absence retailers would use their bargaining powers to negotiate lower wholesale prices, which would then be passed through to the benefit of consumers (cf. Shaffer, 1991).

Our model may also shed light on why some retailers seemingly ask for slotting allowances on all products, while other retailers ask for slotting allowances on some products but not on others, and only for some product categories, while still other retailers, such as Wal-Mart, allegedly do not ask for any slotting allowances.[9] Previous literature has posited that these observed regularities may be attributed to the presence or absence of retail competition (Kuksov and Pazgal, 2007), which may vary from retailer to retailer and by product category, and the presence or absence of noncontractible manufacturer sales efforts (Foros et al., 2009), which may vary by product category. In Kuksov and Pazgal (2007), for example, slotting allowances arise if and only if downstream competition is sufficiently strong. In contrast, we posit that slotting allowances may be related to the scarcity of shelf space, and indeed may be contributing to the scarcity of shelf space by serving as a rent-extraction device. As such, we find that there is an inverse relationship between the observance of slotting allowances and a retailer's bargaining power. Conditional on having at least some market power, the more bargaining power a retailer has over its suppliers, the less likely it will feel a need to limit its shelf space as a means of extracting rent. Slotting allowances do not arise in our model, for example, when the manufacturers have little bargaining power (or conversely, when the retailer has most of the bargaining power). Thus, in contrast to Kuksov and Pazgal, who posit that Wal-Mart does not ask for slotting allowances because it faces less competition in the product market than do other more traditional retailers. Our model suggests that Wal-Mart does not ask for slotting allowances simply because it has more bargaining power than these other retailers have.

Our model applies to any product, new or established, that requires shelf space. Other explanations of slotting allowances either tend to be specific to new products, or are unrelated to whether shelf space is scarce.[10] For example, in the case of new products, one explanation is that slotting allowances serve as a screening device to enable a downstream firm to distinguish high quality products from low quality products, and another explanation is that slotting allowances

9. See FTC (2003; 58) "One supplier reported that all retailers (except Wal-Mart) charge and are paid slotting . . . ." This stylized fact has also been reported numerous times in the trade press literature. See, for example, the article by Mike Duff in *DSN Retailing Today*, October 13, 2003, "Supermarkets charge them, Wal-Mart does not," and the article by Kevin Coupe in *Natural Grocery Buyer*, Spring 2005, "Mention to them that Wal-Mart does not take slotting, but simply drives for the lowest possible cost of goods in its desire to offer the lowest possible prices."

10. There is little empirical literature on the subject, mostly because of the lack of good data, and what there is focuses on new products (Gundlach and Bloom, 1998; Rao and Mahi, 2003; Israilevich, 2004; Sudhir and Rao, 2006).

allow for the efficient sharing of the risk of new product failure. The literature in this vein includes Kelly (1991), Chu (1992), Sullivan (1997), Lariviere and Padmanabhan (1997), Desai (2000), and Bloom et al. (2000). Slotting allowances have also been posited to arise for strategic reasons when retailers have all the bargaining power vis-a-vis their upstream suppliers. Shaffer (1991) and Foros and Kind (2008) suggest that retailers will use their bargaining power to obtain slotting allowances rather than wholesale price concessions in order to avoid dissipating their gains when competing against other downstream firms; and in a model with one upstream firm and competing downstream firms that can make take-it-or-leave-it offers, Marx and Shaffer (2007b) show that slotting allowances arise in equilibrium whenever there is asymmetry downstream (either on the demand or cost side) and lead to exclusion in the sense that the manufacturer does not obtain distribution at all retail outlets.

The rest of the paper proceeds as follows. We introduce the model and notation in Section 2. In Section 3, we solve for the equilibrium payoffs of each firm for the benchmark case in which slotting allowances are infeasible. We then show in Section 4 that the retailer may have an incentive to limit its shelf space in order to obtain better terms of trade, and that slotting allowances weakly increase these incentives. We discuss the model's implications and offer concluding remarks in Section 5.

2. Model

The simplest model in which to capture the idea that slotting allowances can facilitate rent extraction, and that this effect is enhanced when shelf space is scarce, is one in which a retailer must decide which of two products to carry. We assume the two products, X and Y, are produced by manufacturers X and Y at costs $c_X(x)$ and $c_Y(y)$, respectively, where x is the quantity the retailer purchases from manufacturer X and y is the quantity the retailer purchases from manufacturer Y. We also assume that $c_i(\cdot)$ is increasing, continuous, and unbounded, with $c_i(0) = 0$ for all $i = X, Y$.

We assume that X and Y are unrelated and thus independent in demand. This simplifies the model's notation and captures the idea that shelf-space decisions are often "nonlocalized," which means that a retailer that decides to carry a new product in a particular category, say, a new laundry detergent, need not drop one of its old laundry detergents to create space, but can instead obtain the necessary shelf space by dropping a product in some other unrelated category, such as canned soup. Alternatively, one can think of X and Y as being imperfect substitutes, with no change in our qualitative results. Thus, the model also extends to settings in which the retailer's shelf-space decisions are localized (e.g., which products to put in the freezer section).[11]

We say that shelf space is scarce if the retailer can only carry one product and plentiful if the retailer can carry both products. Our focus in this paper is on when the retailer would want to make its shelf space scarce, and how the feasibility of slotting allowances affects this decision.

To isolate the effects of slotting allowances being feasible, we consider two models, one in which slotting allowances are feasible and one in which they are not. The basic model is a four-stage game. As an overview, in stage one the retailer chooses the amount of shelf space; in stage two manufacturers offer slotting allowances if they are feasible; in stage three contracts are negotiated; and in stage four the retailer makes its purchases from the manufacturers. In Section 3, we consider the model in which slotting allowances are not feasible, so the stage-two offers do not occur, and then in Section 4 we consider the full model including slotting allowances.

We now discuss the four stages of the game in greater detail. In stage one, the retailer decides whether to make its shelf space plentiful or scarce. We conceptualize the retailer's shelf space as consisting of slots of fixed width and assume, for simplicity, that the retailer can build either one or two slots. We assume that adequate display and promotion of a product requires exactly one slot of shelf space. Thus, if a manufacturer obtains a slot, it can satisfy any amount of consumer demand for its product, whereas if it does not obtain a slot, it is effectively excluded from making any sales. Plentiful shelf space thus corresponds to building two slots and scarce shelf space corresponds to building one slot. It follows that the retailer can adequately display and sell both products if and only if shelf space is plentiful. By contrast, the retailer can display and sell at most one product if shelf space is scarce.

In stage two, after observing whether the retailer's shelf space is plentiful or scarce, if slotting allowances are feasible, each manufacturer can bid to secure an available slot by offering a slotting allowance.[12] We assume these bids are made simultaneously. The retailer then chooses

11. Although not an issue here, more generally, whether a retailer's shelf space decisions come at the expense of competing products in the same product category or at the expense of unrelated products in unrelated categories is of great concern to policy makers. Under the former, a manufacturer may be accused of attempting to foreclose its rivals if it offers slotting allowances, whereas it would not be accused of trying to foreclose its rivals under the latter.

12. In practice, slotting allowances are typically offered in exchange for a fixed amount of shelf space and a pre-specified length of time for which the firm's product must be

which offer or offers to accept. However, the retailer can accept only as many offers as there are slots available. If shelf space is scarce, then the retailer can accept only one manufacturer's offer. In contrast, if shelf space is plentiful, then the retailer can accept both manufacturers' offers. If an offer is accepted, the retailer is paid and the manufacturer whose slotting allowance was accepted secures the slot. If the retailer rejects both offers, then no payments are made and control over the slot or slots remains in the retailer's hands.

In stage three, contract negotiations take place. With one caveat, we assume that these negotiations take place sequentially, with first the retailer and manufacturer X negotiating contract T_X for the purchase of manufacturer X's product and then the retailer and manufacturer Y negotiating contract T_Y for the purchase of manufacturer Y's product.[13] The caveat is that if shelf space is scarce, slotting allowances are feasible, and the only slot has already been secured by manufacturer i, then there are no contract negotiations between the retailer and manufacturer j in stage three. We place no restrictions on the form of the contracts other than to assume that each contract specifies the retailer's payment as a function of how much the retailer buys of that manufacturer's product only. Thus, we do not allow contracts to depend on both manufacturers' quantities.[14] This restriction, which is critical for our results, can be justified if the retailers' purchases cannot easily be observed by the rival manufacturer or verified in court, or if such contracts would invite unwelcome scrutiny from antitrust policy makers (as per the discussion in Section 1 of this paper).

In stage four, the retailer makes its quantity choices and pays the manufacturers according to its contracts with them. Manufacturer X's payoff is thus $T_X(x) - c_X(x)$ gross of any slotting allowance it may have paid, and analogously, manufacturer Y's payoff is $T_Y(y) - c_Y(y)$. In the absence of a contract, a manufacturer's payoff gross of any slotting allowance it may have paid is zero.

The sequential structure of the game is critical for our results. We assume slotting allowances are set prior to the manufacturer–retailer

---

displayed. We simplify here by assuming that slots are indivisible, and that the length of time for which the product must be on display is for the duration of the game.

13. The assumption that the retailer contracts sequentially with the manufacturers as opposed to simultaneously turns out to be inconsequential if shelf space is plentiful. However, if shelf space is scarce, then contracting sequentially turns out to benefit manufacturer X and the retailer, who gain from being able to shift rent from manufacturer Y.

14. Without this restriction, the retailer and manufacturer X could extract all of manufacturer Y's surplus through the use of penalty clauses (a la Aghion and Bolton, 1987). See also the model in Marx and Shaffer (2004), which focuses on the efficiency properties and extent of surplus extraction when contracts can depend on both sellers' quantities.

negotiations. It is our understanding that this structure is typical of much of the supermarket retailing industry and that slotting allowances are generally fixed for a longer period of time than are other terms, such as wholesale prices (see footnote 17 and the surrounding discussion). In addition, we assume the retailer negotiates with manufacturers sequentially. Within the context of our model, this timing is justified because the retailer may gain from being able to shift rents from the second manufacturer. Moreover, it is often what one observes in reality given that not all manufacturer–retailer contracts are in effect for the same length of time.

As a matter of notation, let $R(x, y)$ denote the maximum revenue that can be earned by the retailer if it purchases quantity x from manufacturer X and quantity y from manufacturer Y. If shelf space is plentiful, then our assumption of independent products implies that the retailer's revenue is separable,

$$R(x, y) = R_X(x) + R_Y(y), \tag{1}$$

where $R_X(x) \equiv R(x, 0)$ and $R_Y(y) \equiv R(0, y)$. In contrast, if shelf space is scarce, so that the retailer can only carry one product, then the retailer's maximized revenue is $R_X(x)$ if it carries product X and $R_Y(y)$ if it carries product Y. Thus, in the case of scarce shelf space, it follows that

$$R(x, y) = \max\{R_X(x), R_Y(y)\}. \tag{2}$$

If the retailer has contracts in place with both manufacturers, then its payoff gross of any slotting allowances it may have received is $R(x, y) - T_X(x) - T_Y(y)$. If negotiations with only manufacturer Y fail, the retailer's gross payoff is $R_X(x) - T_X(x)$. If negotiations with only manufacturer X fail, the retailer's gross payoff is $R_Y(y) - T_Y(y)$. If negotiations with both manufacturers fail, then the retailer's gross payoff is zero. We assume that $R(\cdot, \cdot)$ is continuous and bounded, with $R(0, 0) = 0$.

In negotiations, we assume that the retailer and manufacturer i choose T_i to maximize their joint payoff, and that surplus is divided such that each player receives its disagreement payoff (which we define below) plus a share of the incremental gains from trade (the joint payoff of the retailer and manufacturer i if they trade minus their joint payoff if they do not trade), with proportion $\lambda_i \in [0, 1]$ going to manufacturer i. Our assumption of a fixed division of the incremental gains from trade admits several interpretations. For example, if manufacturer i makes a take-it-or-leave-it offer to the retailer, then $\lambda_i = 1$. If the retailer makes a take-it-or-leave-it offer to manufacturer i, then $\lambda_i = 0$. Finally, if the retailer and manufacturer i split the gains from trade equally, then $\lambda_i = \frac{1}{2}$.

We solve for the equilibrium strategies (we restrict attention to pure strategies and contracts for which optimal quantity choices in stage four exist) of the three players by working backwards. The equilibrium we identify corresponds to the subgame-perfect equilibrium of the related four-stage game in which the assumed bargaining solution is embedded in the players' payoff functions.

3. Solving the Model (The Benchmark Case)

In order to isolate the effects of slotting allowances on the retailer's decision of how much shelf space to build, we must solve the model twice, first for the case in which slotting allowances are infeasible (Section 3), and then for the case in which they are feasible (Section 4). Thus, in this section, we solve the model for the case in which slotting allowances are infeasible. To calculate the equilibrium payoffs of each firm, let $\Pi(x, y) \equiv R(x, y) - c_X(x) - c_Y(y)$ denote the overall joint payoff of the manufacturers and the retailer, and let $\Pi_{XY} \equiv \max_{x,y\geq 0} \Pi(x, y)$ denote its maximized value. It is also useful to define $\Pi_X \equiv \max_{x\geq 0} \Pi(x, 0)$ and $\Pi_Y \equiv \max_{y\geq 0} \Pi(0, y)$. Note that if shelf space is plentiful then $\Pi_{XY} = \Pi_X + \Pi_Y$. By contrast, if shelf space is scarce then $\Pi_{XY} = \max\{\Pi_X, \Pi_Y\}$.

3.1 Stage Four—Retailer's Quantity Choices

Let $(x^{**}(T_X, T_Y), y^{**}(T_X, T_Y))$ denote the retailer's profit-maximizing quantity choices if it has contract T_X in place with manufacturer X and contract T_Y in place with manufacturer Y, that is,

$$(x^{**}(T_X, T_Y), y^{**}(T_X, T_Y)) \in \arg\max_{x,y\geq 0} R(x, y) - T_X(x) - T_Y(y). \tag{3}$$

Note that if shelf space is plentiful, then our assumption that the products are independent implies that the retailer's profit-maximizing choice of x will depend only on T_X, and analogously for the retailer's profit-maximizing choice of y. In contrast, if shelf space is scarce, then the two products are necessarily linked in the sense that the retailer can only carry one product, and thus the retailer's profit-maximizing choices of x and y will in general depend on both manufacturer's contracts.

If instead negotiations with manufacturer Y have failed, and the retailer only has a contract in place with manufacturer X, then we denote the retailer's optimal quantity choice by $x^*(T_X)$, where

$$x^*(T_X) \in \arg\max_{x\geq 0} R_X(x) - T_X(x). \tag{4}$$

Define $y^*(T_Y)$ analogously for the case in which negotiations with only manufacturer X have failed.

The quantities x^{**} and x^{*} (and analogously for y^{**} and y^{*}) are related given the simple structure of the model. If shelf space is plentiful, then x^{**} and x^{*} are the same. Otherwise, if shelf space is scarce, then $x^{**} = x^{*}$ if the retailer carries manufacturer X's product, and $x^{**} \neq x^{*}$ otherwise.

3.2 Stage Three—Negotiations with Manufacturers X and Y

We begin by considering the retailer's negotiation with manufacturer Y, which occurs after its negotiation with manufacturer X. Taking contract T_X as given, contract T_Y will be chosen to solve

$$\max_{T_Y} R(x^{**}, y^{**}) - T_X(x^{**}) - c_Y(y^{**}), \tag{5}$$

such that manufacturer Y earns zero plus its share of the incremental gains from trade

$$\pi_Y = \lambda_Y(R(x^{**}, y^{**}) - T_X(x^{**}) - c_Y(y^{**}) - (R_X(x^{*}) - T_X(x^{*}))), \tag{6}$$

where zero is manufacturer Y's disagreement payoff if negotiations with the retailer break down. Because T_X is fixed when T_Y is chosen, it follows from (5) and (6) that T_Y will be chosen such that[15]

$$(x^{**}(T_X, T_Y), y^{**}(T_X, T_Y)) \in \arg\max_{x,y \geq 0} R(x, y) - T_X(x) - c_Y(y). \tag{7}$$

If there is no contract between the retailer and manufacturer X, then T_Y will be chosen to maximize $\Pi(0, y^{*})$ such that manufacturer Y earns $\lambda_Y \Pi(0, y^{*})$ and the retailer earns $(1 - \lambda_Y)\Pi(0, y^{*})$. Thus, manufacturer Y's payoff in this case will be $\lambda_Y \Pi_Y$ and the retailer's payoff will be $(1 - \lambda_Y)\Pi_Y$.

We now consider the retailer's negotiation with manufacturer X. Because this occurs before the retailer's negotiation with manufacturer Y, contract T_X will be chosen to solve

$$\max_{T_X} \Pi(x^{**}, y^{**}) - \pi_Y, \tag{8}$$

subject to

$$(x^{**}, y^{**}) \in \arg\max_{x,y \geq 0} R(x, y) - T_X(x) - c_Y(y) \tag{9}$$

15. Given T_X, there is no reason to distort the retailer and manufacturer Y's joint-profit-maximizing quantities.

and manufacturer X receiving its share of the incremental gains from trade with the retailer

$$\pi_X = \lambda_X(\Pi(x^{**}, y^{**}) - \pi_Y - (1 - \lambda_Y)\Pi_Y). \tag{10}$$

The program in (8) and (10) implies that T_X will be chosen to maximize the difference between $\Pi(x^{**}, y^{**})$ and manufacturer Y's profit, such that (x^{**}, y^{**}) maximizes the retailer's joint payoff with manufacturer Y, and manufacturer X earns its share of the gains from trade with the retailer. In solving it, note that x^*, x^{**}, and y^{**} are functions of T_X only, and that (9) implies that

$$R(x^{**}, y^{**}) - T_X(x^{**}) - c_Y(y^{**}) \geq \max_{y \geq 0} R(x^*, y) - T_X(x^*) - c_Y(y). \tag{11}$$

In choosing their optimal contract, the retailer and manufacturer X may be able to exploit their first-mover advantage by shifting rents from manufacturer Y. However, their ability to do so is constrained by the independence of the products, whether shelf space is scarce or plentiful, and by the inequality in (11). If manufacturer Y's surplus is fully extracted before (11) binds, then $\pi_Y = 0$ and T_X is chosen such that $(x^{**}, y^{**}) \in \arg\max_{x,y \geq 0} \Pi(x, y)$. However, if the constraint in (11) binds before manufacturer Y's surplus is fully extracted, then manufacturer Y's payoff satisfies

$$\begin{aligned}\pi_Y &= \lambda_Y(R(x^{**}, y^{**}) - T_X(x^{**}) - c_Y(y^{**}) - (R_X(x^*) - T_X(x^*)))\\ &= \lambda_Y\left(\max_{y \geq 0} R(x^*, y) - c_Y(y) - R_X(x^*)\right),\end{aligned}$$

where the first line comes from (6), and the second line is obtained under the assumption that the constraint in (11) binds. In this case, the joint payoff of the retailer and manufacturer X is

$$\Pi(x^{**}, y^{**}) - \lambda_Y\left(\max_{y \geq 0} R(x^*, y) - c_Y(y) - R_X(x^*)\right),$$

which is maximized by choosing T_X such that

$$(x^{**}, y^{**}) \in \arg\max_{x,y \geq 0} \Pi(x, y)$$

and

$$x^* \in \arg\min_{x \geq 0} \lambda_Y\left(\max_{y \geq 0} R(x, y) - c_Y(y) - R_X(x)\right).$$

It follows that overall joint payoff will be maximized whether or not (11) binds, and that manufacturer Y's payoff will be zero or $\lambda_Y \min_{x \geq 0} \max_{y \geq 0}(R(x, y) - c_Y(y) - R_X(x))$, whichever is larger.

Proposition 1: *For any $T_X(\cdot)$ that solves the program in (8)–(10), payoffs are given by*

$$\tilde{\pi}_R = \Pi_{XY} - \tilde{\pi}_X - \tilde{\pi}_Y$$
$$\tilde{\pi}_X = \lambda_X(\Pi_{XY} - \tilde{\pi}_Y - (1 - \lambda_Y)\Pi_Y)$$
$$\tilde{\pi}_Y = \max\left\{0,\ \lambda_Y \min_{x\geq 0}\max_{y\geq 0}\,(R(x, y) - c_Y(y) - R_X(x))\right\},$$

where $\tilde{\pi}_R$, $\tilde{\pi}_X$, and $\tilde{\pi}_Y$ are the payoffs of the retailer and manufacturers X and Y, respectively.

For the proof of Proposition 1, see the Appendix.

Proposition 1 implies that the retailer will carry both products and sell the monopoly quantity of each product if shelf space is plentiful. And if shelf space is scarce, then the retailer will only carry the product that offers the higher monopoly profit, and it will sell the monopoly quantity of that product. Thus, if shelf space is scarce and $\Pi_Y > \Pi_X$, the retailer will sell the monopoly quantity of product Y. If $\Pi_X > \Pi_Y$, the retailer will sell the monopoly quantity of product X.

3.3 Stages One and Two—Retailer's Shelf-Space Decision

Proceeding back to the first two stages of the game, we now consider the retailer's decision of how much shelf space to build given that slotting allowances are assumed to be infeasible. Clearly, overall joint payoff will be higher if the retailer builds two slots rather than just one slot of shelf space (with two slots of shelf space, the maximized overall joint payoff is $\Pi_X + \Pi_Y$, whereas with one slot of shelf space, the maximized overall joint payoff is $\max\{\Pi_X, \Pi_Y\}$). But, as we will show, the retailer would then lose the ability to use its contract with manufacturer X to shift rent from manufacturer Y. Thus, the retailer faces a tradeoff; it can earn a larger share of a smaller overall payoff, or it can earn a smaller share of a larger overall payoff. Using Proposition 1, this tradeoff can be seen by substituting $\tilde{\pi}_X$ into $\tilde{\pi}_R$, simplifying, and then rewriting the retailer's payoff as

$$\tilde{\pi}_R = (1 - \lambda_X)\Pi_{XY} + \lambda_X(1 - \lambda_Y)\Pi_Y - (1 - \lambda_X)\tilde{\pi}_Y. \tag{12}$$

This is increasing in the overall joint payoff, Π_{XY}, but decreasing in manufacturer Y's payoff, $\tilde{\pi}_Y$.

3.3.1 The Case of Plentiful Shelf Space

Consider first the case in which the retailer builds two slots. In this case, the retailer's maximized revenue is $R(x, y) = R_X(x) + R_Y(y)$, and thus, from Proposition 1, manufacturer Y's payoff is

$$\tilde{\pi}_Y = \max\left\{0, \lambda_Y \min_{x\geq 0} \max_{y\geq 0} \left(R_Y(y) - c_Y(y)\right)\right\} = \lambda_Y \Pi_Y.$$

Substituting this into the retailer's payoff in (12) and simplifying yields

$$\tilde{\pi}_R = (1-\lambda_X)\Pi_X + (1-\lambda_Y)\Pi_Y. \tag{13}$$

Notice that when shelf space is plentiful, contract T_X has no effect on the negotiations between the retailer and manufacturer Y (nor would contract T_Y affect negotiations between the retailer and manufacturer X if manufacturer Y negotiated first). This follows because, as noted previously, $x^{**} = x^*$ when shelf space is plentiful (and the products are independent), and similarly $y^{**} = y^*$.

3.3.2 The Case of Scarce Shelf Space

Now suppose the retailer builds only one slot. In this case, the retailer's maximized revenue is $R(x, y) = \max\{R_X(x), R_Y(y)\}$, and thus, from Proposition 1, manufacturer Y's payoff is

$$\begin{aligned}\tilde{\pi}_Y &= \max\left\{0,\ \lambda_Y \min_{x\geq 0} \max_{y\geq 0} \left(\max\{R_X(x), R_Y(y)\} - c_Y(y) - R_X(x)\right)\right\} \\ &= \max\left\{0,\ \lambda_Y \left(\Pi_Y - \max_{x\geq 0} R_X(x)\right)\right\}.\end{aligned}$$

Notice that manufacturer Y's payoff is strictly less in this case than when shelf space is plentiful. The reason is that manufacturer X and the retailer can engage in rent shifting when shelf space is scarce, something they could not do when shelf space was plentiful. Manufacturer X offers generous terms that allow the retailer to earn $\max_{x\geq 0} R_X(x)$ if it carries manufacturer X's product instead of Y's product, and this limits the retailer's gains from trade with manufacturer Y. Indeed, when shelf space is scarce, manufacturer Y earns positive payoff if and only if $\Pi_Y > \max_{x\geq 0} R_X(x)$.

If manufacturer Y's payoff is zero, which is the case if the monopoly value of manufacturer X's product is higher, or if $\max_{x\geq 0} R_X(x) > \Pi_Y > \Pi_X$, then the retailer's payoff in (12) simplifies to

$$\tilde{\pi}_R = (1-\lambda_X)\Pi_{XY} + \lambda_X(1-\lambda_Y)\Pi_Y. \tag{14}$$

If instead manufacturer Y's payoff is larger than zero, then the retailer's payoff in (12) simplifies to

$$\tilde{\pi}_R = (1-\lambda_Y)\Pi_Y + \lambda_Y(1-\lambda_X)\max_{x\geq 0} R_X(x). \tag{15}$$

Because (12) is decreasing in $\tilde{\pi}_y$, it follows that the retailer's payoff is smaller in (15) than in (14).

3.3.3 Comparing Payoffs

Comparing the retailer's payoff in (14) with its payoff in (13) yields an immediate result. If $\Pi_X \geq \Pi_Y$, so that $\Pi_{XY} = \Pi_X$ when shelf space is scarce, then the retailer's payoff is weakly higher in (13), which applies when shelf space is plentiful (it is strictly higher in this case if $\lambda_X < 1$). It follows that the retailer will not want to limit its shelf space if it negotiates first with the manufacturer whose monopoly payoff is higher. The decrease in overall joint payoff in this case always exceeds what the retailer could gain by extracting greater surplus from the manufacturer that negotiates second.[16]

More generally, comparing the retailer's payoff in (14) with its payoff in (13), it follows that the retailer earns higher payoff in (14), which applies when shelf space is scarce, if and only if $\lambda_X < 1$ and $\lambda_Y \Pi_Y > \Pi_X$. The analogous condition for when the retailer's payoff in (15) exceeds its payoff in (13) is $\lambda_X < 1$ and $\lambda_Y \max_{x \geq 0} R_X(x) > \Pi_X$. This result is summarized in the following proposition.

Proposition 2: *There exist equilibria in which the retailer chooses to have scarce shelf space when slotting allowances are infeasible if and only if* $\lambda_Y \min\{\max_{x \geq 0} R_X(x), \Pi_Y\} \geq \Pi_X$. *If this inequality is strict and* $\lambda_X < 1$, *then the retailer chooses to have scarce shelf space in all equilibria.*

Proposition 2 implies that the retailer can sometimes earn higher payoff by limiting its shelf space, thereby destroying overall surplus. It follows that the vertically integrated outcome will not always be implemented even in this simple vertical chain with one retailer and two manufacturers. This contrasts with much of the literature on common agency, for example, Bernheim and Whinston (1985), which suggests that the vertically integrated outcome will be obtained when there is complete information. The difference here is that in addition to choosing prices and quantities, the retailer must also decide on its shelf space. As it turns out, for some environments, even though the retailer does not distort its prices under common agency, it may distort the availability

16. This raises the question as to why the retailer would want to negotiate with manufacturer X first if X's product is better than Y's product, when it could earn higher payoff by reversing the order. If the retailer could choose with whom to negotiate first, then clearly in the case of independent products and scarce shelf space, the retailer would prefer to negotiate first with the manufacturer whose product has the smaller monopoly payoff, as this benefits it more in extracting surplus from the stronger manufacturer. For an analysis of the retailer's optimal order of negotiations in a more general rent-shifting setting where products can be substitutes or complements, see Marx and Shaffer (2007a).

of its shelf space (for a similar result when contracts are negotiated simultaneously, see O'Brien and Shaffer, 1997).

The tradeoff for the retailer is whether to build one slot in order to capture a larger share of a smaller overall profit, or two slots in order to capture a smaller share of a larger overall profit. Because the difference in overall profit is Π_X, it follows that the larger is Π_X, (the right-hand side of the condition in Proposition 2), the less likely the retailer will opt to build only one slot, all else being equal. Similarly, the retailer will be less likely to build just one slot the smaller is λ_Y. This is because for a given Π_X, the smaller is λ_Y, the larger is the retailer's share of the profit when shelf space is plentiful. Although the retailer's profit share is also decreasing in λ_Y when shelf space is scarce, the decrease in this case is not as large given the possibilities for rent shifting. Thus, a smaller value of λ_Y makes it less likely that the retailer will build only one slot all else being equal.

4. Slotting Allowances and Scarce Shelf Space

It is sometimes argued that the alleged scarcity of shelf space found in practice is a disequilibrium phenomenon that will be self-correcting in the long run. Proposition 2 suggests otherwise. It implies that a shortage of shelf space can arise in equilibrium as a consequence of profit-maximizing behavior by retailers with market power. It therefore follows that as long as there exist barriers to entry that protect retailers' long-run profits, the alleged shortage of shelf-space need not be self-correcting.

Given that a shortage exists, temporary or not, slotting allowances have mostly been viewed as an efficient mechanism for allocating the scarce shelf space to those products that generate the highest social surplus. However, this view implicitly takes the scarcity of shelf space as exogenous, and moreover, as we have seen, the role of slotting allowances in helping ensure that the most desirable products are carried may be minimal (in our model, the retailer always selects the product with the higher stand-alone monopoly profit regardless of whether slotting allowances are feasible).

4.1 The Role of Slotting Allowances in the Retailer's Shelf-Space Decision

The question we now ask is what impact will slotting allowances have on the scarcity of shelf space. To consider this, we assume slotting allowances are now feasible and can be offered following the retailer's decision whether to make shelf space plentiful or scarce, but before the contract-negotiation stage. This timing reflects the status of these payments as a long run variable (e.g., an annual payment) that is typically adjusted with less frequency in practice than other contract terms. Many manufacturers claim, for example, that little negotiation takes place over the amount of the slotting allowance, and that this amount is "known by vendors in advance of discussions."[17] Slotting allowances get them in the door, and then negotiations over the other contract terms begin.

We model slotting allowances as follows. After observing whether shelf space is plentiful or scarce, each manufacturer has the option of offering a slotting allowance to secure an available slot. One can think of these offers as being made simultaneously. The retailer then chooses which offer or offers to accept subject to the constraint that it can only accept as many offers as there are slots available. Acceptance of a manufacturer's offer guarantees a slot for that manufacturer's product.

We begin the analysis of this revised game by considering first the case where shelf space is plentiful. Two observations make this case trivial to solve. First, by assumption, demand for a given product does not increase if the manufacturer has two slots versus just one slot, and second, demand for a given product does not decrease if a slot is given to the other (unrelated) product. It follows that each manufacturer cares only about securing a single slot for its own product. When shelf space is plentiful, there is by definition a slot available for both products, and thus it follows that slotting allowances will not arise in equilibrium if shelf space is plentiful (it is a dominant strategy for both manufacturers to forego offering slotting allowances when shelf space is plentiful).

Now suppose that the retailer's shelf space is scarce. In this case, the retailer has only one available slot to allocate, and thus one product will necessarily be excluded. As a result, the manufacturers may be forced to compete for the retailer's patronage by offering to pay for shelf space. Let $S_X \geq 0$ denote manufacturer X's offer and $S_Y \geq 0$ denote manufacturer Y's offer. One can think of S_X and S_Y as lump-sum payments that are offered to the retailer in exchange for shelf space of fixed width (a slot) and duration (the rest of the game). Because the retailer has only one slot available to allocate when shelf space is scarce, it can accept at most one manufacturer's offer, where acceptance implies that the retailer's slot is reserved for that manufacturer's product.

17. See FTC (2003; 58), "The suppliers' responses suggest that they perceive much less flexibility in whether and how much slotting is required." The FTC report goes on to say (same paragraph) that "Most suppliers also stated that the amount of the required slotting payment is set by each retailer and known by vendors in advance of discussions, albeit without written communications. Suppliers stated that very little negotiation over the amount takes place."

If no offer is accepted, the slot remains under the retailer's control, no payment is made, and the rest of the game proceeds as above. In contrast, if the retailer accepts an offer, the manufacturer whose offer was accepted pays the retailer and then chooses along with the retailer the remaining contract terms to maximize their joint payoff, with each receiving its bargaining share of the incremental gains from trade. It follows that if the retailer accepts manufacturer X's offer, then the retailer's payoff is $(1-\lambda_X)\Pi_X + S_X$, manufacturer X's payoff is $\lambda_X\Pi_X - S_X$, and manufacturer Y's payoff is zero.[18] If the retailer accepts manufacturer Y's offer, then the retailer's payoff is $(1-\lambda_Y)\Pi_Y + S_Y$, manufacturer Y's payoff is $\lambda_Y\Pi_Y - S_Y$, and manufacturer X's payoff is zero.

4.1.1 Which Offer will be Accepted?

Solving the game with slotting allowances is straightforward. There are two cases to consider. Consider first the case in which $\Pi_Y > \Pi_X$. Then, if the retailer were to accept manufacturer X's offer, the joint payoff of the retailer and manufacturer Y would be at most Π_X (manufacturer X cannot earn negative profit in equilibrium), which is less than what their joint payoff would be if the retailer were to accept manufacturer Y's offer instead. This implies that there is no equilibrium in which manufacturer X's offer is accepted.[19] Similarly, if the retailer were to reject both offers, then the joint payoff of the retailer and manufacturer Y would, from Proposition 1, be $\Pi_Y - \lambda_X\lambda_Y \min\{\Pi_Y, \max_{x\geq 0} R_X(x)\}$, which is weakly less than what their joint payoff would be if the retailer were to accept manufacturer Y's offer (strictly less if $\lambda_X, \lambda_Y > 0$). This implies that if $\lambda_X, \lambda_Y > 0$, then there is no equilibrium in which shelf space is scarce and both offers are rejected.

It follows that in any equilibrium with scarce shelf space, the retailer will accept manufacturer Y's offer and earn at least as much profit as it could earn if instead it accepted manufacturer X's offer or rejected both offers. It must also be the case that manufacturer X cannot profitably deviate in equilibrium and that manufacturer Y will offer no more than is necessary to obtain the retailer's acceptance. Therefore, in

18. We are implicitly assuming that if the retailer accepts manufacturer X's offer, then it is obligated to carry and sell product X. Alternatively, one might imagine that manufacturer X could obtain the shelf space slot but then allow the retailer to carry and sell product Y instead. We can show that our results are robust to this possibility.

19. Formally, there exists a payment S_Y such that $(1-\lambda_Y)\Pi_Y + S_Y > \Pi_X$ and $\lambda_Y\Pi_Y - S_Y > 0$. By offering an S_Y that satisfies these inequalities, manufacturer Y can induce the retailer to accept its offer and make itself better off.

any equilibrium with scarce shelf space, the retailer's payoff is given by

$$\tilde{\pi}_R = \begin{cases} \max\{\Pi_X, (1-\lambda_X)\Pi_Y + \lambda_X(1-\lambda_Y)\Pi_Y\}, & \text{if } \max_{x\geq 0} R_X(x) \geq \Pi_Y \\ \max\{\Pi_X, (1-\lambda_Y)\Pi_Y + \lambda_Y(1-\lambda_X)\max_{x\geq 0} R_X(x)\}, & \text{if } \max_{x\geq 0} R_X(x) < \Pi_Y. \end{cases} \quad (16)$$

The interpretation of (16) is straightforward. In any equilibrium with scarce shelf space, the retailer can earn Π_X if it accepts manufacturer X's offer, and it can earn either the payoff in (14) or the payoff in (15) if it rejects both offers (where its payoff depends on whether manufacturer Y would earn positive surplus—which depends on the relationship between $\max_{x\geq 0} R_X(x)$ and Π_Y). It follows that (16) represents the retailer's opportunity cost of accepting manufacturer Y's offer.

Comparing the retailer's payoff in (16) with the payoffs in (14) and (15), it can be seen that, conditional on shelf space being scarce, the retailer is always weakly (sometimes strictly) better off when slotting allowances are feasible than when they are not. Intuitively, the retailer cannot be worse off with slotting allowances because it can always reject both offers and guarantee itself the minimum of the payoffs in (14) and (15). And it can sometimes be better off (e.g., when $\lambda_X, \lambda_Y = 1$) because the auctioning of scarce shelf space forces the weaker manufacturer to compete away its entire surplus even if this manufacturer would otherwise have been able to earn positive payoff.

It follows that because the retailer can always choose whether to restrict its shelf space, the feasibility of slotting allowances cannot make the retailer worse off. Comparing the retailer's payoff in (16) with its payoff in (13), it also follows that the feasibility of slotting allowances may contribute to the scarcity of shelf space. Because the retailer earns at least Π_X when shelf space is scarce and slotting allowances are permitted, a sufficient condition for the retailer to prefer to have one slot is

$$\Pi_X > (1-\lambda_X)\Pi_X + (1-\lambda_Y)\Pi_Y,$$

which holds if and only if $\lambda_X\Pi_X > (1-\lambda_Y)\Pi_Y$. Combining this condition with the condition $\lambda_Y \min\{\max_{x\geq 0} R_X(x), \Pi_Y\} \geq \Pi_X$ from Proposition 3, yields the first half of Proposition 3 below.

Now consider the alternative case in which $\Pi_X \geq \Pi_Y$. Then the analogue to (16) is

$$\tilde{\pi}_R = \max\{\Pi_Y, (1-\lambda_X)\Pi_X + \lambda_X(1-\lambda_Y)\Pi_Y\}, \quad (17)$$

from which it follows that the retailer will strictly prefer to build one slot if and only if

$$\Pi_Y > (1 - \lambda_X)\Pi_X + (1 - \lambda_Y)\Pi_Y,$$

or in other words if and only if $\lambda_Y \Pi_Y > (1 - \lambda_X)\Pi_X$. This yields the second half of Proposition 3.

Proposition 3: *Assume slotting allowances are feasible. If $\Pi_Y > \Pi_X$, there exist equilibria in which the retailer chooses scarce shelf space if and only if $\lambda_Y \min\{\max_{x\geq 0} R_X(x), \Pi_Y\} \geq \Pi_X$ or $\lambda_X \Pi_X \geq (1 - \lambda_Y)\Pi_Y$; and scarce shelf space arises in all equilibria if $\lambda_X < 1$ and either inequality is strict. If $\Pi_X > \Pi_Y$, the retailer chooses scarce shelf space if and only if $\lambda_Y \Pi_Y \geq (1 - \lambda_X)\Pi_X$; and scarce shelf space arises in all equilibria if the inequality is strict.*

Conventional wisdom suggests that slotting allowances arise in response to scarce shelf space. In this view, the scarcity of shelf space is assumed to be exogenous. Here, we have shown that slotting allowances do not arise when shelf space is plentiful, and in that sense our findings support the view that scarce shelf space causes slotting allowances. However, our findings also imply that causality can go the other way—the scarcity of shelf space may, in part, be caused by slotting allowances. This suggests that slotting allowances may be contributing to the scarcity of retail shelf space.

We can illustrate the differences between Proposition 2 (where slotting allowances are not feasible) and Proposition 3 (where slotting allowances are feasible) in Figure 1 below. Assume $\lambda_X = \lambda_Y$ and $\max_{x\geq 0} R_X(x) > \Pi_Y > \Pi_X$. Then Proposition 2 implies that shelf space

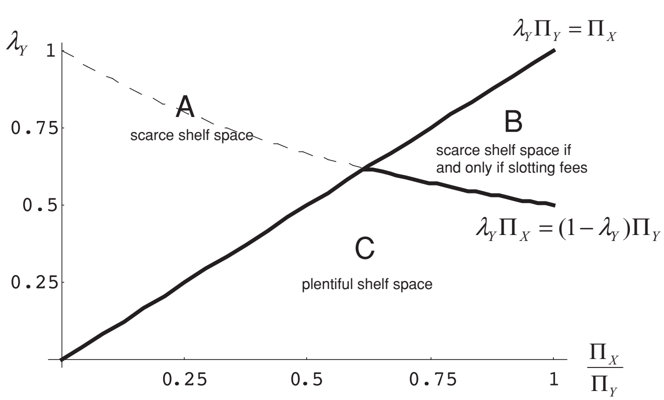

FIGURE 1. SLOTTING ALLOWANCES AND SHELF SPACE (ASSUMING $\lambda_X = \lambda_Y$ AND $\max_{x\geq 0} R_X(x) > \Pi_Y > \Pi_X$)

will be scarce if and only if $\lambda_Y \Pi_Y \geq \Pi_X$ (region A), whereas Proposition 3 implies that shelf space will be scarce if and only if $\lambda_Y \Pi_Y \geq \Pi_X$ or $\lambda_Y \Pi_X \geq (1 - \lambda_Y)\Pi_Y$ (regions A and B). When slotting allowances are feasible, scarce shelf space arises when parameter values are in region A or B. But when they are not feasible, scarce shelf space arises only when parameter values are in region A.

Slotting allowances contribute to the scarcity of shelf space in some environments (region B in Figure 1) because they provide the retailer with an additional means of shifting rent when shelf space is scarce. By auctioning off its shelf space upfront, the retailer can induce the winning manufacturer to give it a payoff that is at least equal to the payoff it could earn if it gave the space to the loser. This guarantees for the retailer a payoff equal to the smaller of the monopoly profits of the two manufacturers' products (in equilibrium, the manufacturer of the weaker product offers to sell its product at cost and earns zero profit). Note that without slotting allowances, the retailer has no such guarantee. Whether shelf space is plentiful or scarce, the retailer's payoff in the absence of slotting allowances is decreasing in the manufacturers' bargaining powers and is zero in the polar case in which both manufacturers can make take-it-or-leave-it offers. In contrast, the retailer always earns positive payoff with slotting allowances when shelf space is scarce. Thus, it follows that slotting allowances may contribute to the scarcity of shelf space and be particularly valuable to retailers that have local monopoly power but otherwise have little bargaining power.

4.2 Bargaining Power and Slotting Allowances

The retailer's decision whether to limit its shelf space, and the extent to which the feasibility of slotting allowances contributes to this decision, depends on the retailer's bargaining power. Surprisingly, we see from Proposition 3 that there is an inverse relationship between the retailer's bargaining power with respect to each manufacturer and whether slotting allowances arise in equilibrium.[20] Slotting allowances are less likely to arise in equilibrium when the retailer has more bargaining

20. In our model, the retailer has some market power even if the manufacturers make take-it-or-leave-it offers, as the manufacturers have no alternative retailer to go to if they want to sell to the local consumers. For this reason, the retailer earns positive profit even when the manufacturers have all the bargaining power. Therefore, a more complete interpretation of our result is the following. First, suppose that the retailer does not have *any* market power, in that there are many potential retailers that each manufacturer can turn to. Clearly, in such a case, the retailer cannot benefit from limiting its shelf space (the manufacturers will go to another retailer), so slotting allowances will not arise in equilibrium. Second, as we have shown, slotting allowances do not arise if the retailer can make take-it-or-leave-it offers. These polar cases suggest that while the retailer must

power (when λ_X and λ_Y are low) because with higher bargaining power the focus of the retailer shifts toward maximizing the overall profit rather than limiting its shelf space to extract rent. Slotting allowances do not arise, for example, if the retailer can make take-it-or-leave-it offers ($\lambda_X = \lambda_Y = 0$). This may explain why a retailer like Wal-Mart, which has high bargaining power vis-a-vis its suppliers, allegedly does not avail itself of these payments. We call this the "Wal-Mart phenomenon." The idea is that the usual tradeoff of whether to opt for a larger share of a smaller overall profit or a smaller share of a larger overall profit does not apply to Wal-Mart because it can already capture most or all of the profit. Hence, Wal-Mart prefers to have plentiful shelf space.

The Wal-Mart phenomenon is unusual in that Wal-Mart's bargaining power is thought to be high whether it is dealing with small manufacturers of produce or potentially large manufacturers of differentiated, consumer packaged goods. More generally, however, one can imagine that a retailer may have high bargaining power with respect to suppliers in some product classes, but low bargaining power with respect to suppliers in other product classes. If the retailer's shelf space decisions are localized (for example, products that require freezer space are competing only against themselves and not against products that do not need refrigeration), then it is possible that a given retailer may want strategically to limit its shelf space and obtain slotting allowances for some product categories, but not limit its shelf space and thus not obtain slotting allowances for other product categories. In this case, our results imply that the incidence of slotting allowances will be higher in those product categories in which manufacturers have relatively more bargaining power.

Conditional on observing slotting allowances, one might think that a retailer with more bargaining power will obtain a larger slotting allowance than a retailer with less bargaining power. This intuition turns out to hold in our model for some environments, but overall the relationship between a retailer's bargaining power and the size of the slotting allowance it obtains is nonmonotonic. To see this, recall that the retailer's payoff when shelf space is scarce and $\Pi_Y > \Pi_X$ is given by[21]

$$(1-\lambda_Y)\Pi_Y + S_Y = \max\left\{\Pi_X, \min\left\{(1-\lambda_X\lambda_Y)\Pi_Y, (1-\lambda_Y)\Pi_Y + \lambda_Y(1-\lambda_X)\max_{x\geq 0} R_X(x)\right\}\right\},$$

have market power (few alternative retailers) for slotting allowances to arise, it cannot have significant bargaining power. We thank Yaron Yehezkel for this observation.

21. The right-hand side of the equality in the displayed equation is just (16) written in a more compact way.

which implicitly defines S_Y. Subtracting $(1 - \lambda_Y)\Pi_Y$ from both sides of this condition gives

$$S_Y = \max\left\{\Pi_X - (1 - \lambda_Y)\Pi_Y, (1 - \lambda_X)\lambda_Y \min\left\{\Pi_Y, \max_{x \geq 0} R_X(x)\right\}\right\},$$

from which it follows that in the case of $\Pi_Y > \Pi_X$, manufacturer Y's slotting offer, S_Y, is increasing in λ_Y but weakly decreasing in λ_X (strictly decreasing if the retailer's payoff exceeds Π_X).[22] Intuitively, the larger is λ_X, the smaller is the retailer's disagreement payoff if it rejects both offers and thus the smaller is the slotting allowance that manufacturer Y will offer in equilibrium. The larger is λ_Y, the less the retailer will earn in the continuation game if it accepts manufacturer Y's offer, and hence the more manufacturer Y will have to offer upfront in the form of a slotting allowance.

As an example of how slotting allowances interact with the retailer's bargaining parameters, consider the case in which the retailer has equal bargaining power with respect to each manufacturer. Then, if the retailer's equilibrium payoff exceeds Π_X, slotting allowances will initially be increasing in the retailer's bargaining power and then decreasing, obtaining a maximum at $\lambda_X = \lambda_Y = 1/2$.

4.3 Slotting Allowances, Welfare, and Profits

Our results have implications for social welfare and the distribution of profits (who gains and who loses). Notice that if shelf space is plentiful, then the retailer sells both products and the overall joint payoff is $\Pi_X + \Pi_Y$. But if shelf space is scarce, then the retailer sells only one product and the overall joint payoff is $\max\{\Pi_X, \Pi_Y\}$. In both cases, consumers face monopoly prices, which is true whether or not slotting allowances are feasible. Thus, conditional on the availability of shelf space, it follows that slotting allowances do not affect product selection or retail prices. However, as we have seen, the feasibility of slotting allowances affects whether shelf space is scarce. In particular, it is seen from Propositions 2 and 3 that when slotting allowances are feasible, the retailer is more likely to limit its shelf space (see region B in Figure 1). When this happens (i.e., when parameter values are in this region), consumers are worse off with slotting allowances than

22. It is easily verified that $S_Y = 0$ if and only if either $\lambda_Y = 0$ or $\lambda_X = 1$ and $(1 - \lambda_Y)\Pi_Y > \Pi_X$. If $\lambda_Y = 0$, then the retailer prefers to build two slots of shelf space. If $\lambda_X = 1$ and $(1 - \lambda_Y)\Pi_Y > \Pi_X$, then the retailer is indifferent between building one slot or two, as its payoff in both cases is $(1 - \lambda_Y)\Pi_Y$. Thus, it follows that slotting allowances will be positive for all parameter values for which the retailer strictly prefers to limit its shelf space to one slot.

without. Without slotting allowances, consumers would be able to select between products X and Y, each priced at the monopoly price. With slotting allowances, consumers would only be able to select product Y, and there would be no change in product Y's price. Because overall profits are also lower in region B with slotting allowances than without, it follows that welfare is unambiguously lower in this region.

Turning to the distribution of firm profits, the retailer is weakly better off with slotting allowances, and when $\Pi_Y > \Pi_X$, manufacturer X is weakly worse off. Slotting allowances add to the retailer's profit when the monopoly profit of the losing manufacturer's product exceeds what the retailer could earn if it rejected both offers or built two slots of shelf space. Given $\Pi_Y > \Pi_X$, it follows that slotting allowances contribute to the retailer's profit if and only if Π_X is larger than the retailer's payoff in (14) and (15), which for $\max_{x \geq 0} R_X(x) \geq \Pi_Y$, implies that $\Pi_X > (1 - \lambda_X \lambda_Y)\Pi_Y$, and larger than its payoff when shelf space is plentiful, which implies that $\lambda_X \Pi_X > (1 - \lambda_Y)\Pi_Y$.

The upper envelope of these two inequalities is shown in Figure 2. For all points in regions A′ and B, the retailer is better off with slotting allowances. For all points in regions A and C, the retailer is indifferent to slotting allowances. Thus, for a given λ_Y, the retailer is more likely to gain from slotting allowances the larger is Π_X relative to Π_Y (the higher is the guaranteed payoff from slotting allowances). And, for a given ratio of Π_X to Π_Y, the retailer is more likely to gain from slotting allowances the larger is λ_Y (the higher is manufacturer Y's bargaining power). Because manufacturer X's profit decreases to zero when slotting allowances are feasible and shelf space is scarce, manufacturer X is worse off in regions A, A′, and B. However, there is no change in its

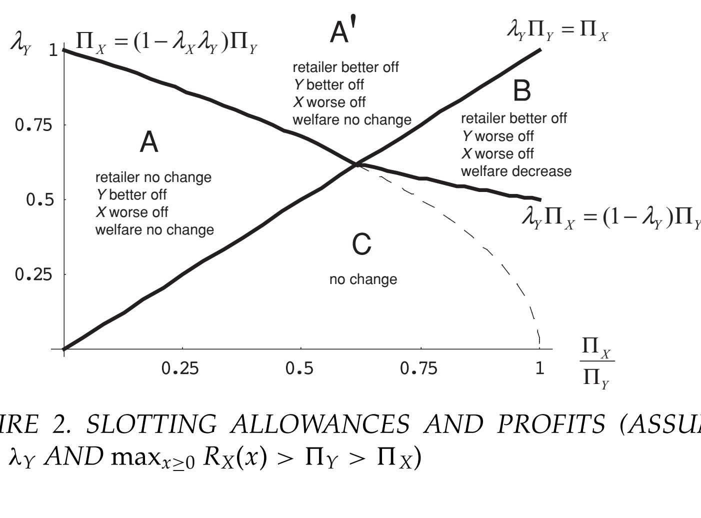

FIGURE 2. SLOTTING ALLOWANCES AND PROFITS (ASSUMING $\lambda_X = \lambda_Y$ AND $\max_{x \geq 0} R_X(x) > \Pi_Y > \Pi_X$)

profit in region C because shelf space is plentiful in this case with or without slotting allowances.

Manufacturer Y gains from slotting allowances in region A because manufacturer X loses, the retailer's profit is unchanged, and overall joint payoff is unchanged. In region A′, the retailer picks up some of the profit lost by manufacturer X, but manufacturer Y still gains because its payoff in this region, $\Pi_Y - \Pi_X$, exceeds its payoff in the absence of slotting allowances when shelf space is scarce, $\max\{0, \lambda_Y(\Pi_Y - \max_{x\geq 0} R_X(x))\}$. In region B, manufacturer Y earns $\Pi_Y - \Pi_X$ if slotting allowances are feasible and $\lambda_Y \Pi_Y$ otherwise. This implies that manufacturer Y's payoff is higher in region B with slotting allowances if and only if $(1 - \lambda_Y)\Pi_Y > \Pi_X$. But for all points in the interior of region B, $\lambda_x \Pi_X > (1 - \lambda_Y)\Pi_Y$. Hence, manufacturer Y loses in region B. However, there is no change in manufacturer Y's payoff in region C because shelf space is always plentiful in this case.

5. Conclusion

Slotting allowances have long been of concern in antitrust. They have been the subject of numerous congressional hearings and the focal point of two recent Federal Trade Commission reports on marketing practices in the retail grocery industry. Conventional wisdom suggests that the growth in slotting allowances can be attributed to the imbalance between the number of products (new and established) that are available for the retailer to choose from at any given time and the number of products that the retailer can profitably carry given its limited shelf space. In contrast, we show that slotting allowances may in fact also be contributing to the scarcity of retail shelf space.

In our model, slotting allowances allow the retailer to capture more efficiently the value of its shelf space when shelf space is scarce. By itself, this effect is welfare neutral. Slotting allowances in this case serve only to transfer rents from the weaker manufacturer to the stronger manufacturer and, for some parameters, from the weaker manufacturer to the retailer. However, the problem is that this same mechanism may also induce the retailer to limit its shelf space when shelf space would otherwise be plentiful. This effect tends to reduce welfare and can harm both manufacturers. Welfare decreases when shelf space is scarce because consumers suffer from reduced product choice but do not benefit from lower prices. The manufacturers are worse off when shelf space would otherwise be plentiful because overall joint profit is lower, but the retailer's profit is higher.

Our results point to a new source of welfare loss. Policy makers have previously been concerned with whether slotting allowances would lead to higher or lower retail prices, and whether dominant manufacturers might abuse slotting allowances by buying up scarce shelf space in order to foreclose smaller rivals with better products. These concerns have caused policy makers in specific cases to consider whether the excluded firms have equal access to capital markets, whether slotting allowances are coupled with exclusivity provisions (e.g., whether the dominant firm requires that the downstream firm not sell its competitor's product), and whether there are scale economies in production that may prevent an excluded firm from effectively competing elsewhere in the market. However, in our model, there is no effect on retail pricing, and the better product is always chosen conditional on the availability of shelf space. Instead, we find that welfare may be lower because slotting allowances may induce the retailer to limit its shelf space. This suggests that slotting allowances may be harmful even if they are not accompanied by exclusivity provisions, the market does not exhibit scale economies in production, and access to capital markets is not a concern.

Our results also have testable implications. In a recent study on slotting allowances, the Federal Trade Commission (2003, p. vi) concluded that, "For the five product categories during the time periods of this study, the surveyed retailers' data on the frequency of slotting fees varied widely between and within product categories, across retailers, and across a particular retailer's regions." The Federal Trade Commission (2003, p. vii) further concluded that, "For the seven retailers during the time periods of this study, slotting allowances were more prevalent for ice cream and salad dressing products than for bread and hot dog products." In some instances, the Federal Trade Commission found that slotting allowances were paid on all new products in a particular category. In other categories, they found that few if any slotting allowances were paid on the new products.

Kuksov and Pazgal (2007) and Foros et al. (2009) have previously posited that these findings may be accounted for by differences in the intensity of retail competition, and by the importance of noncontractible demand-enhancing investments by manufacturers, respectively.[23] In contrast, we suggest that these findings may also (or instead) be due to differences in a retailer's bargaining power with respect to each manufacturer and with respect to each other. In particular, we have

23. An alternative explanation for why slotting allowances may be observed for branded goods but less so for unbranded goods is that suppliers of unbranded goods may face binding credit constraints. This accords with a concern of policy makers that slotting allowances may place small manufacturers at a disadvantage for reasons other than the merits of their products (see the discussion above in Section 1). We thank an anonymous referee for this insight.

shown that the incidence of slotting allowances may be inversely related to each retailer's bargaining power. If a retailer with market power has all the bargaining power with respect to each manufacturer, then the retailer extracts all the available surplus in the market and slotting allowances do not arise. But slotting allowances will arise if the manufacturers have all the bargaining power because then the retailer will want to limit its shelf space all else being equal in order to earn some profit. For intermediate levels of bargaining power, the retailer is more likely to limit its shelf space (and therefore slotting allowances are more likely to arise) the less bargaining power it has. This accords well with the observation that slotting allowances tend to arise for items such as ice cream and salad dressing, where strong brand names exist, but not for items where strong brand names do not exist, such as bread and hot dogs. This also accords well with the observation that Wal-Mart, a retailer that is thought to have a lot of bargaining power, allegedly does not accept slotting allowances.

We have also shown that, conditional on receiving slotting allowances, there is a nonmonotonic relationship between a retailer's bargaining power and the magnitude of any slotting allowances it receives. Retailers with low bargaining power vis-a-vis their suppliers will tend to negotiate roughly the same level of slotting allowances as retailers with high bargaining power, whereas retailers with moderate bargaining power will be able to negotiate the highest slotting allowances. In principle, these findings can be tested, but, to our knowledge, the data needed to do so is not yet available.

Appendix

Proof of Proposition 1. To complete the proof, it is sufficient to show that finding T_X that solves the program in (8)–(10) is equivalent to finding $(x_2, y_2, x_1, t_2, t_1)$ that solves the following program:

$$\max_{x_2 \geq 0, y_2 \geq 0, x_1 \geq 0, t_2, t_1} \Pi(x_2, y_2) - \tilde{\pi}_Y \tag{A1}$$

subject to

$$t_2 \geq c_X(x_2), \tag{A2}$$

$$\tilde{\pi}_Y = \lambda_Y(R(x_2, y_2) - t_2 - c_Y(y_2) - (R(x_1, 0) - t_1)), \tag{A3}$$

$$t_2 - c_X(x_2) = \lambda_X(\Pi(x_2, y_2) - \tilde{\pi}_Y - (1 - \lambda_Y)\Pi_Y), \tag{A4}$$

$$y_2 \in \arg\max_{y \geq 0} R(x_2, y) - c_Y(y), \tag{A5}$$

$$R(x_1, 0) - t_1 \geq R(x_2, 0) - t_2, \tag{A6}$$

$$R(x_2, y_2) - t_2 - c_Y(y_2) \geq \max_{y \geq 0} R(x_1, y) - t_1 - c_Y(y). \tag{A7}$$

Condition (A2) implies that manufacturer Y earns nonnegative profit. Conditions (A3) and (A4) imply that manufacturer X and manufacturer Y, respectively, earn their bargaining share of the retailer's gains from trade with them. Condition (A5) implies that y^{**} maximizes the retailer's joint payoff with manufacturer Y. Condition (A5) implies that the retailer chooses $(x^*, 0)$ over $(x^{**}, 0)$ when it only has a contract with manufacturer X. Condition (A7) implies that the retailer chooses (x^{**}, y^{**}) over (x^*, y) for any $y \geq 0$ when it has contracts in place with both manufacturers.

Suppose that $\hat{T}_X$ solves (8)–(10). Then we claim that

$$(x_2, y_2, x_1, t_2, t_1) \equiv (x^{**}(\hat{T}_X), y^{**}(\hat{T}_X), x^*(\hat{T}_X), \hat{T}_X(x^{**}), \hat{T}_X(x^*)),$$

solves the program in (A1)–(A7). To see this, note that (6) implies (A3), (9) implies (A5), and (10) implies (A4). The definitions of x^{**} and x^* imply that (A2), (A5), and (A7) are satisfied. Thus, $(x_2, y_2, x_1, t_2, t_1)$ is a feasible solution to (A1)–(A7). Suppose $(x_2, y_2, x_1, t_2, t_1)$ does not solve this program. Then there exists $(x_2', y_2', x_1', t_2', t_1')$ satisfying (A2)–(A7) such that the maximand in (A1) is greater at $(x_2', y_2', x_1', t_2', t_1')$ than at $(x_2, y_2, x_1, t_2, t_1)$. Consider contract T_X' defined by:

$$T_X'(x) \equiv \begin{cases} t_2', & \text{if } x = x_2' \\ t_1', & \text{if } x = x_1' \\ \infty, & \text{otherwise.} \end{cases}$$

Because $(x_2', y_2', x_1', t_2', t_1')$ satisfies (A2)–(A7) it follows that $(x^{**}(T_X'), y^{**}(T_X')) = (x_2', y_2')$ and $x^*(T_X') = x_1'$, and so T_X' satisfies (9)–(10). Thus, T_X' is a feasible solution to (9)–(10) and gives the retailer and manufacturer X higher joint payoff than $\hat{T}_X$, a contradiction. Thus, $(x_2, y_2, x_1, t_2, t_1) \equiv (x^{**}(\hat{T}_X), y^{**}(\hat{T}_X), x^*(\hat{T}_X), \hat{T}_X(x^{**}), \hat{T}_X(x^*))$ solves (A1)–(A7).

Now suppose $\hat{T}_X$ is not an equilibrium contract. Then $\hat{T}_X$ does not solve (8)–(10). Because $\hat{T}_X$ is not an equilibrium contract, the contract T_X'', where

$$T_X''(x) \equiv \begin{cases} \hat{T}_X(x^{**}(\hat{T}_X)), & \text{if } x = x^{**}(\hat{T}_X) \\ \hat{T}_X(x^*(\hat{T}_X)), & \text{if } x = x^*(\hat{T}_X) \\ \infty, & \text{otherwise,} \end{cases}$$

is also not an equilibrium contract. If T_X'' is not a feasible solution to (8)–(10), then at least one of (9)–(10) is violated. Consider $(x_2, y_2, x_1, t_2, t_1) \equiv (x^{**}(\hat{T}_X), y^{**}(\hat{T}_X), x^*(\hat{T}_X), \hat{T}_X(x^{**}), \hat{T}_X(x^*))$. If T_X'' violates (9), then $(x_2, y_2, x_1, t_2, t_1)$ violates at least one of (A2)–(A7). If T_X'' violates (10), then $(x_2, y_2, x_1, t_2, t_1)$ violates (A4). If T_X'' is a feasible solution to (8)–(10), then there exists T_X''' also feasible but giving a higher value of the maximand in (8). Then $(x_2, y_2, x_1, t_2, t_1)$ and

$$\left(x_2''', y_2''', x_1''', t_2''', t_1'''\right) \equiv \left(x^{**}\left(T_X'''\right), y^{**}\left(T_X'''\right), x^*\left(T_X'''\right), T_X'''(x^{**}), T_X'''(x^*)\right),$$

both satisfy (A2)–(A7), but $(x_2''', y_2''', x_1''', t_2''', t_1''')$ results in a higher value of the maximand in (A1) than $(x_2, y_2, x_1, t_2, t_1)$. Thus, $(x^{**}(\hat{T}_X), y^{**}(\hat{T}_X), x^*(\hat{T}_X), \hat{T}_X(x^{**}), \hat{T}_X(x^*))$ does not solve (A1)–(A7). $\square$

References

Aghion, P. and P. Bolton, 1987, "Contracts as a Barrier to Entry," *American Economic Review*, 77, 388–401.

Bernheim, B.D. and M. Whinston, 1985, "Common Marketing Agency as a Device for Facilitating Collusion," *RAND Journal of Economics*, 15, 269–281.

Bloom, P., G. Gundlach and J. Cannon, 2000, "Slotting Allowances and Fees: Schools of Thought and the Views of Practicing Managers," *Journal of Marketing*, 64, 92–108.

Canadian Competition Bureau, 2002, "The Abuse of Dominance Provisions as Applied to the Canadian Grocery Sector," Available at http://strategis.ic.gc.ca/pics/ct/ct02465e.pdf (accessed May 10, 2010).

Chu, W., 1992, "Demand Signaling and Screening in Channels of Distribution," *Marketing Science*, 11, 327–347.

Dana, J., 2004, "Buyer Groups as Strategic Commitments," Working Paper, Northwestern University.

Desai, P., 2000, "Multiple Messages to Retain Retailers: Signaling New Product Demand," *Marketing Science*, 19, 381–389.

Federal Trade Commission, 2001, "Report on the Federal Trade Commission Workshop on Slotting Allowances and Other Marketing Practices in the Grocery Industry," Washington, DC. Available at http://www.ftc.gov/opa/2001/02/slotting.htm (accessed May 10, 2010).

Federal Trade Commission, 2003, "Slotting Allowances in the Retail Grocery Industry: Selected Case Studies in Five Product Categories," Washington, DC. Available at http://www.ftc.gov/opa/2003/11/slottingallowance.htm (accessed May 10, 2010).

Foros, Ø. and H. Kind, 2008, "Do Slotting Allowances Harm Retail Competition?" *Scandinavian Journal of Economics*, 110, 367–384.

——, —— and J. Sand, 2008, "Slotting Allowances and Manufacturers' Retail Sales Efforts," *Southern Economic Journal*, 76, 266–282.

Gundlach, G. and P. Bloom, 1998, "Slotting Allowances and the Retail Sale of Alcohol Beverages," *Journal of Public Policy and Marketing*, 17, 173–184.

Inderst, I. and G. Shaffer, 2007, "Retail Mergers, Buyer Power, and Product Variety," *Economic Journal*, 117, 45–67.

Israilevich, G., 2004, "Assessing Supermarket Product-Line Decisions: The Impact of Slotting Fees," *Quantitative Marketing and Economics*, 2, 141–167.

Kelly, K., 1991, "The Antitrust Analysis of Grocery Slotting Allowances: The Procompetitive Case," *Journal of Public Policy and Marketing*, 10, 187–198.

Kuksov, D. and A. Pazgal, 2007, "The Effects of Costs and Competition on Slotting Allowances," *Marketing Science*, 26, 259–267.

Lariviere, M. and V. Padmanabhan, 1997, "Slotting Allowances and New Product Introductions," *Marketing Science*, 16, 112–128.

Marx, L.M. and G. Shaffer, 2004, "Rent Shifting, Exclusion, and Market-Share Discounts," Working Paper, Duke University.

—— and ——, 2007a, "Rent Shifting and the Order of Negotiations," *International Journal of Industrial Organization*, 25, 1109–1125.

—— and ——, 2007b, "Upfront Payments and Exclusion in Downstream Markets," *RAND Journal of Economics*, 38, 823–843.

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

Rao, A. and H. Mahi, 2003, "The Price of Launching a New Product: Empirical Evidence on Factors Affecting the Relative Magnitudes of Slotting Allowances," *Marketing Science*, 22, 246–268.

Shaffer, G., 1991, "Slotting Allowances and Resale Price Maintenance: A Comparison of Facilitating Practices," *Rand Journal of Economics*, 22, 120–135.

——, 2005, "Slotting Allowances and Optimal Product Variety," *The B.E. Journal of Economic Analysis & Policy*, Vol. 5: Iss. 1 (Advances), Article 3, 2005. Available at http://www.bepress.com/bejeap/advances/vol5/iss1/art3/ (accessed May 10, 2010).

Sudhir, K. and V. Rao, 2006, "Do Slotting Allowances Enhance Efficiency or Hinder Competition?," *Journal of Marketing Research*, 43, 137–155.

Sullivan, M., 1997, "Slotting Allowances and the Market for New Products," *Journal of Law and Economics*, 40, 461–493.