KONKURRANSE-TILSYNET
Norwegian Competition Authority

Rapport 1/2018

When should Retailers Accept Slotting Allowances

Tommy Staahl Gabrielsen and Bjørn Olav Johansen, University of Bergen
Greg Shaffer, University of Rochester

Prosjektet har mottatt midler fra det alminnelige prisreguleringsfondet.

When Should Retailers Accept Slotting Allowances?*

Tommy Staahl Gabrielsen† Bjørn Olav Johansen‡

University of Bergen *University of Bergen*

Greg Shaffer§

University of Rochester

August 2017

Abstract

This paper looks at the use of slotting allowances as a signaling device when manufacturers have private information about their demands, and asks, when should retailers accept them, and when should they not? We find that accepting slotting allowances often comes with a tradeoff. On the one hand, by putting its shelf space up for bid and allowing manufacturers to pay for it with slotting allowances, a retailer can learn which is the best product. But it also allows the winning manufacturer to capture its product's value added (i.e., the difference in value between its product and the next best alternative). On the other hand, by not allowing manufacturers to use slotting allowances, the retailer remains uninformed about which is the best product — but, in this case the winning manufacturer is only able to capture the

*We thank the "alminnelige prisreguleringsfondet" through the Norwegian Competition Authority for financial support for this research. We also thank participants at the 4th CREST-ECODEC Conference, Paris, 2017, for valuable comments.

†University of Bergen, Department of Economics, tommy.gabrielsen@uib.no.

‡University of Bergen, Department of Economics, bjorn.johansen@uib.no.

§Simon School RC 270100, University of Rochester, Rochester, NY 14627; shaffer@simon.rochester.edu.

retailer's perceived value added of its product. We show that this tradeoff is such that the retailer is more likely to gain from accepting slotting allowances the higher is its opportunity cost of shelf space and the larger is the variance in the demand potential of the manufacturers' products. These findings have implications for how and why the incidence of slotting allowances varies across retail product categories.

"*The inspiration to create an energy bar occurred during a daylong, 175-mile ride with my buddy Jay. We'd been knawing on some "other" energy bars. Suddenly, despite my hunger, I couldn't take another bite. That's the moment I now call the "epiphany." Two years later, after countless hours in Mom's kitchen, Clif Bar became a reality. And the mission to create a better-tasting energy bar was accomplished.*"[1] Gary, Founder and Owner of Clif Bar

" *What I looked for was a whole new line. In the marketing line, the consumer has a demand, and if it is not being met, that's what I look for. If you walk down the aisle, all you'd see in the cookie business was fat, sugar, preservatives, etc. And I said to myself, 'Listen, if people wanted jam without sugar, they certainly want a cookie without sugar." And simply, that simple realization, and the appearance of that realization, started R.W. Frookies. It was pretty simple. There's fruit and there's cookies. Now there's Frookies. You sweeten it with fruit.*"[2] Richard, Founder and Owner of R.W. Frookies

1 Introduction

Imagine there are two products, but you can only choose one. One claims to be a "better tasting energy bar." The other claims to be a "healthier cookie." You are the retailer. Which do you choose? What if a seller is willing to offer you a lump-sum, upfront payment (i.e., a slotting allowance) to help you decide? Should you accept it? More generally, does your answer depend on whether the products are already established and merely seeking to upgrade their shelf positioning, or new and seeking to be carried for the first time?

In this paper, we consider the use of slotting allowances as a signaling device when manufacturers have private information about their demands and retailers have limited shelf space. We ask the following question: when should retailers accept slotting allowances, and when should they not? At first sight, this may seem like a simple question to answer because one might think that retailers should of course always want to accept slotting allowances, particularly when the payments can help them make better-informed decisions. But the evidence in practice suggests it is not that simple, and indeed we will find in this paper that retailers often face a tradeoff when they accept slotting allowances.

Having a clear policy one way or the other is important because grocery retailers are constantly making crucial decisions about how to allocate their scarce shelf space. Among

[1] Found on the package of a chocolate-chip Clif Bar (clifbar.com), purchased at retail.

[2] Richard Worth FNN (CNN) interview, July 11, 2012, www.youtube.com/watch?v=1QOTI9cXnCU.

other things, they must decide which of their slow-moving products to drop, and which new products to add. They must also decide which products to place at eye level, which products to put on display, and which products to place near the checkout lanes, etc. The constant is that there is never a shortage of choices, and the process never ends. The next day, there will be more products to consider, and more shelf-space decisions to make.

Mistakes are inevitable – as one would expect when retailers must make their decisions under uncertainty. Decisions that may have looked good ex-ante may turn out to be bad ex-post. The products that were added may turn out to be worse than the products that were dropped. The premium space that was given to one product may in retrospect have been more profitably given to another, and so on. Nowhere are the stakes higher than on supermarket grocery items, where profit margins are thin, shelf space is notoriously scarce, and many more products are available to be stocked than can possibly be carried.[3]

Mistakes also hurt the sellers, who may be frustrated by their inability to convince retailers of their products' worthiness. Gary, founder and owner of *Clif Bar*, succeeded. Many others have not. But, who can blame the retailers, who are constantly being told that the product or products being pitched to them "taste better" than the alternatives (see the above quote from Gary), or will be the "next big thing" (see the above quote from Richard). The problem is that all sellers have incentives to make these claims regardless of how confident they may actually be, or what their test-market data may be telling them.

If one could wave a magic wand and know for certain which products would sell the best, then of course one would want to wave the wand, right? Although they are not wands, it is widely believed that slotting allowances, which have been defined as "lump-sum, upfront payments from a manufacturer or producer (collectively, manufacturer) to a retailer to have a new product carried by the retailer and placed on its shelves" can go a long way toward mitigating the retailers' uncertainty and helping them make better-informed decisions.[4] A manufacturer who is more confident about its product's likely success can take advantage of the payments to offer more than a manufacturer who is less confident, thereby distinguishing itself and providing a valuable signal to the retailer about its product's likely selling potential. Indeed, no one questions that such payments

[3]Trade press articles are constantly hammering home this theme. Consider for example, the image conjured up by statements such as "The typical supermarket has room for fewer than 25,000 products. Yet there are some 100,000 available, and between 10,000 and 25,000 items are introduced each year," Tim Hammonds, past president of the U.S. Food Marketing Institute, as reported in Shaffer (2005).

[4]This is the definition used by the Federal Trade Commission in FTC (2001) and FTC (2003). Others have used a more expansive definition, encompassing all upfront, lump-sum payments from manufacturers to retailers. Our results are not sensitive to the precise definition. The key is that the payments are fixed and independent of quantity. In the UK, such payments are referred to as "reverse-fixed payments," because the lump-sum payments go in the reverse direction of what might expect (from sellers to buyers).

can help the retailers make better informed decisions and mitigate against costly errors.[5]

Here is how the FTC (2003) described it in their official report on slotting allowances:

> "When a manufacturer believes its product is highly likely to succeed, it will be willing to pay a significant slotting fee, knowing that it is likely to recover this expense through profits earned from future sales. By contrast, if a manufacturer has significant doubts about the likelihood of a product's success, it will not be willing to pay as high a slotting fee. Recognizing this, retailers rationally infer that a new product is more likely to succeed if the manufacturer is willing to pay a higher slotting fee. Based on this rational inference, retailers are more likely to stock products for which higher slotting fees are paid, not only because they receive a higher slotting fee, but also because they expect such products to be more highly valued by consumers ..."

Slotting allowances have received a lot of attention over the years. They have been the focus of Congressional Hearings and numerous government investigations (both in the U.S. and abroad), and there have been more than eighty academic articles, working papers, and surveys (not to mention countless trade press articles) written about them.[6] They are said to account for billions of dollars annually in the U.S. alone, and in the U.K., they have recently led to the creation of a dedicated ombudsman to adjudicate the practice, given the large amounts of money involved. Although there are many explanations for slotting allowances, signaling is perhaps the most prominent one (certainly it is the most prominent among the efficiency explanations).[7] And when the prerequisites for it hold (manufacturers have private information), signaling would seem to be a win-win for both manufacturers and retailers (and for consumers, who also want the best products carried).

Yet, despite all the scrutiny, there are still a number of unexplained anomalies. The FTC study cited above has found that slotting allowances are mostly used for new products (as opposed to being used for established products which may, for the first time, be seeking better shelf positioning). There is also widespread variation in their use. This variation is

[5] The signaling explanation for slotting allowances assumes that manufacturers have better information than retailers about the likely success of their products. While this is often the case, it is not always so.

[6] For an in-depth survey of the practice, see Shaffer (2013), which was commissioned by the UK Office of Fair Trading. The U.S. Federal Trade Commission conducted its own study of the practice (see FTC, 2001, 2003) as did the Canadian Bureau of Competition (2002) and the Norwegian Competition Authority (2005). Slotting fees are mentioned in the Competition Commission's investigation of grocery markets in April 2008, and they are the focus of the UK's *Groceries Supply Code of Practice*. See also Bloom et al. (2000) and Wilkie et al. (2002) for some insightful surveys regarding the views of practicing managers.

[7] The signaling explanation for slotting allowances is described on page 1 of the FTC's report, where it is the first item mentioned. It was also among the first explanations in the literature (see Kelly 1991).

reported to be along many dimensions. For a given retailer, for example, it is sometimes found that the incidence of slotting allowances is higher for small manufacturers than for large manufacturers. Moreover, the frequency of slotting allowances varies widely across product categories, with slotting allowances nearly always being observed in some categories (e.g., in the refrigerated and frozen food sectors) while seldom being seen in others. There are also differences in the incidence of slotting allowances across retailers, with some retailers (e.g., *Wal-Mart* and *Costco* at the time of the study) refusing to accept them, and other retailers seemingly going out of their way to embrace them. There are even differences in the incidence of slotting allowances at stores of the same retailer, with the incidence being higher for stores in urban areas, and lower for stores in rural areas.

These anomalies are hard to reconcile. Why would any retailer not want to accept slotting allowances? And when retailers do accept them, why do they accept them for some product categories but not for others? Are *Wal-Mart* and *Costco* simply better informed than the manufacturers about how well their new products will sell? Are some product categories (e.g. ice cream and TV dinners) systematically more prone to uncertainty about demand than others? Is it the case that for some product categories, signaling how well a new product will sell is typically done through other means (e.g., through wholesale prices only, where the manufacturer announces "Yes, my product has a high wholesale price, but you should carry it anyway, because if it were not as good as I say, I would not want to be offering you such a high price"), as some have argued (see Lariviere and Padmanabhan, 1997). And what about the incidence of slotting allowances differing across outlets of the same retailer? Is this retailer less well informed at some of its outlets than at others?

In this paper, we offer a potential explanation for some of the aforementioned anomalies by challenging the conventional wisdom that retailers necessarily benefit from having manufacturers reveal their private information. In a model with two manufacturers and one retailer, where the manufacturers are vying for a single slot of shelf space, we find that it is sometimes optimal for the retailer to adopt a policy of not accepting slotting allowances, thereby purposely remaining in the dark when making his shelf space decisions.

The reason for this is that accepting slotting allowances often comes with a tradeoff. On the one hand, it can lead to better informed decisions per the signaling reasons discussed above (in the model, we bias the results in favor of the retailer accepting slotting allowances by assuming that the manufacturers are perfectly informed about how well their products would sell — so that by allowing the manufacturers to signal their information, the retailer will always know with certainty which product is better). On the other hand, it may enable the better manufacturer to capture surplus that would otherwise have gone to the retailer.

The idea is that when manufacturers are allowed to signal their information with slotting allowances, the winning manufacturer only has to offer the retailer the maximum of what the retailer could earn by selling the losing manufacturer's product or sticking with his status quo (and earning his opportunity cost of shelf space), whereas if slotting allowances are not allowed, the manufacturers are subject to the retailer's preconception as to which product is better.[8] With slotting allowances, the winning manufacturer captures the value added of its product as perceived by the manufacturers (who have better information about how well their products will sell). Without slotting allowances, the winning manufacturer captures only the value added that is perceived by the retailer, prior to the resolution of his uncertainty. If the retailer perceives the products to be equally good, competition in wholesale prices ensures that neither manufacturer will earn positive profit.

We can fix this idea more concretely by considering the following stylized example. Suppose the retailer believes ex-ante that it could earn an expected profit of 5 from selling each manufacturer's product (a markup of 1 plus unit sales of 5). However, suppose the truth is that one manufacturer's product would in fact generate a profit of 7 for the retailer, while the other manufacturer's product would only generate a profit of 3. Then, if the retailer "puts his shelf space up for bid" by accepting slotting allowances, the winning manufacturer only has to ensure that the retailer would earn a profit of 3 (plus epsilon) from stocking its product. But if the retailer does not put his shelf space up for bid, such that the better manufacturer has no credible means of signaling, each manufacturer would be forced to offer its product at cost, and the retailer would earn an expected profit of 5.[9] Not accepting slotting allowances in this instance is unambiguously better for the retailer.

We could also have tipped the decision in the other direction by choosing different numbers and incorporating a positive opportunity cost of shelf space (i.e., by incorporating a lower bound on what the retailer would be willing to accept if he put his shelf space up for bid). Indeed, our findings point to the retailer's opportunity cost of shelf space and the variance of the demand potential of the manufacturers' products as the key elements that determine whether the retailer would want to adopt a policy of accepting slotting allowances. We find that accepting slotting allowances is more likely to be agreeable to

[8] As discussed in the text, there is some literature that looks at the possibility of a manufacturer signaling its type through the use of wholesale price offers only, but this requires the retailer to know what the potential set of demand curves looks like, and in particular, to recognize when the high-demand firm's wholesale price would be above the vertical intercept of the low-demand firm's demand curve. We assume the retailer does not have this information in practice, nor would it be practical to implement.

[9] By not accepting slotting allowances, the retailer is essentially taking his chances on stocking the better product. He might get lucky, choose the defacto better product, and earn a profit of 7, or he might get unlucky, choose the weaker of the two products, and only earn a profit of 3. But notice that even if he turns out to be unlucky, he is still no worse off than what he would have been with slotting allowances.

the retailer the higher is the retailer's opportunity cost of shelf space and the greater is the variance in the manufacturers' demand potential. This then gives testable implications regarding which product categories, and for which types of products, we might expect to observe a higher incidence of slotting allowances. For a company such as *Wal-Mart*, which traditionally has had a store-wide policy of not accepting slotting allowances for any reason, we note that the effect of the policy is to not allow the winning manufacturers to capture the value added of their products. This is consistent with *Wal-Mart's* often stated corporate policy of always trying to obtain the lowest possible wholesale prices.

In extensions of the model, we consider how the retailer's decision of whether to adopt a policy of accepting slotting allowances depends also on whether two large manufacturers are competing for one slot, two small manufacturers are competing for one slot, or one large manufacturer and one small manufacturer are competing for one slot. Among the settings we consider is the case where the demand of the large manufacturer's product has both a higher mean and a lower variance than the demand of the small manufacturer's product, and the case where it has a lower variance but the same mean. These extensions allow us to consider, among other things, how the incidence of slotting allowances might be expected to vary depending on whether the new products in a given product category are dominated by a lot of me-too products, or alternatively by a lot of innovative products.

The rest of the paper proceeds as follows. In the next section, we provide an overview of the existing literature and discuss our relationship to it. In Section 3, we describe the model and define notation. In Section 4, we solve the model assuming both manufacturers' products are good. In Section 5, we consider the possibility that one or both products may be bad. In Section 6, we model asymmetry between the manufacturers and consider how this affects the incidence of slotting allowances. Section 7 concludes the paper. The proofs of the Propositions and Lemmas are contained in the Appendix. Throughout the paper, our focus is on understanding when retailers would be expected to adopt a policy of accepting slotting allowances, and when not, and then deriving testable implications.

2 Related literature

The literature that looks at slotting allowances does so from a number of different angles, both pro-competitive and anti-competitive. The anti-competitive stories primarily have to do with using slotting allowances to dampen competition among competing retailers by keeping wholesale prices elevated (see Shaffer, 1991; Foros and Kind, 2008; and Innes and Hamilton, 2006), foreclose upstream competitors by buying up the available shelf space (see Shaffer, 2005), foreclose downstream competitors by making it too costly for manufacturers to serve them (see Marx and Shaffer, 2007), and facilitate tacit collusion by making cheating easier for retailers to detect (see Piccolo and Miklos-Thal, 2012).

In contrast, the pro-competitive explanations focus mostly on signaling (see Kelly, 1991; Lariviere and Padmanabhan, 1997; and Desai, 2000) and its cousin, screening by the retailer (see Chu, 1992; Chu and Messenger, 1995; and DeVuyst 2005).[10] There is also work on using slotting allowances as a way of subsidizing demand-enhancing investments by the retailers (see Raju and Zhang, 2004; and Kolay and Shaffer, 2016), and/or increasing the incentives for demand-enhancing investments by the manufacturers (see Farrell, 2001).[11]

Of this literature, the anti-competitive stories are typically highly context dependent and require that the firms' contractual terms be observable across different manufacturers' products and across retailers. Moreover, they typically hinge on the competitiveness of the downstream retailing sector (which we abstract from by considering only a single retailer). And, unlike in our model, they do not attempt to address the anomalies identified above.

Of the pro-competitive stories, Kelly (1991), Lariviere and Padmanabhan (1997) and Desai (2000) are the closest to us in the sense that they were the first to argue that slotting allowances could be used to signal demand and would be expected to be observed in certain settings. Kelly did so eloquently, but without a formal model. Lariviere and Padmanabhan and Desai showed formally that slotting allowances could be used to signal demand, but only when signaling through wholesale prices would be insufficient to induce separation. In particular, Lariviere and Padmanabhan compare and contrast signaling via wholesale prices and slotting allowances, but do not consider the possibility of signaling through demand-enhancing advertising. Desai compares and contrasts all three methods of signaling. In both papers, it is assumed that there is only one manufacturer and one retailer, and that the manufacturer's product is one of two possible types, either good or bad. Their main finding is that there is a critical value of the retailer's opportunity cost of shelf space such that if the retailer's opportunity cost of shelf space is less than this critical value, the manufacturer will prefer to signal its type through its choice of wholesale price only, whereas if the retailer's opportunity cost of shelf space is greater than this critical

[10] Both deal with settings in which the buyers and sellers have different information. The difference between signaling and screening has to do with which party takes the action. Signaling requires that the party with the better information take the action (e.g., signal the product's quality). Screening requires that the lesser informed party takes the action (e.g., offer a menu that is designed to induce self selection).

[11] There are also competitively-neutral stories that have slotting allowances emerging in equilibrium as part of a bilateral bargaining outcome between a manufacturer and its retailer. These are essentially full information stories, where the slotting allowances arise in equilibrium either for the purpose of obtaining additional distribution (see Shaffer, 1990; and Kuksov and Pazgal, 2007), or for the purpose of extracting surplus away from a third-party (see Marx and Shaffer, 2010; and Caprice and von Schlippenbach, 2012).

value, the manufacturer will prefer to signal its type using all of its available instruments.

We differ from Lariviere and Padmanabhan (1997) and Desai (2000) in several important ways. First, we consider the use of slotting allowances from the retailer's perspective. Whereas in their models, the manufacturer always prefers to signal demand through its wholesale price only, the retailer always prefers that the signaling take place through slotting allowances. This means that if the retailer could set up the rules of the game as he sees fit, as he does in our model, or if the retailer did not have information about the vertical intercepts of the demand of the good and bad types (see footnote 8 above for more on this), slotting allowances would always be observed independently of the cost of shelf space in their models. Second, their models allow for two types, whereas we allow for a continuum of types. Third, and most importantly, they assume there is only manufacturer, whose only task is to signal the demand potential of its product. In contrast, we allow competing manufacturers to engage in signaling. As such, we are the first to identify the importance of "the competitive effect" as a reason why retailers may not want to accept slotting allowances, even if by accepting slotting allowances, demand would be revealed perfectly. This effect, previously missed, is one of the key drivers of our results.

3 The model

Our basic setup is one where two manufacturers face a single retailer. The retailer has limited shelf space[12] and can only accept at most one of the manufacturers' products in his store. Assuming the manufacturers have private information about the demands for their products (i.e., how well their products would sell if carried by the retailer),[13] the question we address in this paper is whether, when, and under what conditions will it be profitable for the retailer to adopt a policy of accepting (or not accepting) slotting allowances?

We assume for simplicity that the goods are produced at zero marginal cost, and that all consumers who are interested in product i, $i = 1, 2$, have a willingness to pay of at most $v > 0$ for a unit of the product. We let $q_i \geq 0$ denote the number of such consumers. We further assume that only manufacturer i knows the actual value of q_i at the start of the game. We thus assume that this information is private, and that the rival manufacturer

[12] Our assumption is that shelf space is scarce, which is well documented in grocery retailing.

[13] The manufacturers may have better information about their products' attributes and quality. They may also have conducted focus groups and/or collected data from test markets that might have revealed something about their product's likelihood of success. This is not to say that manufacturers are always better informed than their retail counterparts. Retailers may, for example, conduct their own market testing, and their predictive ability about which products will sell and which will not sell may be better, in some cases, than that of a relatively less sophisticated manufacturer. See the discussion in FTC (2001).

and retailer only know (prior to the retailer's decision of which product to carry) that q_i is uniformly distributed with density $1/[b-a]$ on the support $[a,b]$, where $b > a \geq 0$.[14]

Below we will distinguish between cases where the manufacturers offer simple linear wholesale contracts, and cases where slotting allowances are also offered. When simple linear wholesale contracts are used, the retailer's total payment for product i is given by $T_i(x_i) = w_i x_i$, where w_i is the per unit wholesale price and x_i is the quantity ordered (in equilibrium, x_i will either be zero or q_i). When slotting allowances are also offered, the retailer's total payment for product i is given by $T_i(x_i) = w_i x_i - S_i$, where w_i and x_i are as previously defined and $S_i \geq 0$ is the slotting fee given by manufacturer i to the retailer.

Let the highest and lowest value of q_1, q_2 be denoted by $q_{(1)} = \max\{q_1, q_2\}$ and $q_{(2)} = \min\{q_1, q_2\}$, respectively, and let $w_{(1)}$ and $w_{(2)}$ (and $S_{(1)}$ and $S_{(2)}$) denote the corresponding wholesale prices (and slotting fees) of the manufacturers that have these demands (i.e, $w_{(1)}$ and $S_{(1)}$ correspond to the wholesale price and slotting fee of the manufacturer of the product that has the higher demand, and similarly for $w_{(2)}$ and $S_{(2)}$).

We let $\pi_r \geq 0$ denote the retailer's profit, and $\pi_i \geq 0$ denote manufacturer i's profit. Furthermore, we let $s \geq 0$ denote the retailer's opportunity cost of shelf space. This cost is defined to be what the retailer could earn from the next best use of his available shelf space if he did not carry either manufacturer's product. It establishes a lower bound on the retailer's expected profit, and it means, among other things, that if the demand potential of a given manufacturer's product is sufficiently low, the retailer will not want to carry the product. As examples, one can think of the retailer's opportunity cost of shelf space as corresponding to the profit that the retailer can earn from an existing product that he is thinking about replacing, or to the profit that he can expect to earn if he were to use his shelf space to carry a store brand instead of buying from one of the manufacturers.

We formalize the game between the retailer and the manufacturers as follows:

1. The retailer decides whether to adopt a policy of accepting slotting allowances. Accepting slotting allowances means that the retailer will consider contracts that specify a wholesale price and a slotting allowance. Not accepting slotting allowances means that the retailer will consider contracts that specify a wholesale price only.[15]

2. Each manufacturer learns the value of its product's draw. The manufacturers then

[14] Our setup can alternatively be recast as one in manufacturer i has private information about the potential profitability of its product. Given zero marginal costs, this profit is given by vq_i for product i.

[15] We implicitly assume that the retailer can commit to its policy and not be tempted if a manufacturer were to offer a slotting allowance anyway. This seems reasonable to us because in practice, where shelf-space decisions are made on a daily basis, our assumption is tantamount to assuming that long run reputation/credibility concerns are often likely to trump any short run gains to the retailer from deviating.

simultaneously and independently choose their contract terms, taking the retailer's decision in stage one as given. We envision a stage-two game of alternating offers, where first manufacturer 1 chooses its terms, then manufacturer 2 chooses its terms (after observing manufacturer 1's terms), then manufacturer 1 gets to revise its terms, and so on, until the manufacturer which is to move decides not to change anything. The contract phase then ends, and the game proceeds to stage three.[16]

3. The retailer accepts at most one of the manufacturers' offers and sets the retail price.

At the conclusion of stage three, demand is revealed, payments are made, and all payoffs are realized. We will use Perfect Bayesian equilibrium as our solution concept, which means that: (i) at each stage of the game, each player's strategy specifies optimal actions, given her beliefs and the strategies of the other players, and (ii) given the strategy profile for all players, the beliefs are consistent with Bayes' rule whenever possible. That is, the players' strategies must be sequentially rational, and their beliefs must be consistent.

4 The case against accepting slotting allowances

To illustrate the main tradeoffs facing the retailer, and to understand why accepting slotting allowances is not always optimal, we will begin by assuming for now that both products are "good" in the sense that carrying either one would be an improvement over the retailer's status quo for all realizations of q_1 and q_2 (i.e., each would be profitable for the retailer to carry even in the unlikely event that only a consumers liked the product).

Formally, we assume that $0 \leq s \leq va$, and note that when both products are good, the problem for the retailer in essence boils down to choosing which one he thinks will be better in the sense of yielding the higher profit all things considered.[17] The question then arises how should the retailer structure the rules of the game with this task in mind? Should he adopt a policy of not accepting slotting allowances and simply choose whichever manufacturer offers him the lower wholesale price, or should he allow for offers that include slotting allowances and see which one gives him the better overall combination of terms?

[16] This game is equivalent to running an English auction (where the manufacturers bid against each other for the retailer's patronage until one of them is no longer willing to continue bidding). It is meant to capture a setting in which a retailer goes and back and forth between the manufacturers, attempting to secure the best deal possible, subject to the constraint that he cannot lie about the offers he receives.

[17] The reader may wonder why the retailer cannot carry both in this case. The reason is that we have in mind a setting in which shelf space limitations represent a hard constraint in the sense that the opportunity cost of stocking the second product given that the first one is stocked is prohibitively high (i.e., $s > vb$). For example, the retailer may only be able to promote one product at a time in its display.

Suppose first that the retailer adopts a policy of not accepting slotting allowances in stage one. In practice, this might take the form of a store or chain-wide policy of not accepting slotting allowances, as *Wal-Mart* and *Costco* allegedly have done, or it might simply reflect the retailer's norms and expectations in a given product category.[18] Either way, the policy forces the manufacturers to compete for the retailer's available shelf space with wholesale price offers only. With this in mind, the first thing to notice is that in the continuation game beginning in stage two, the manufacturers' ability to signal their products' demand potential is effectively eliminated. The reason is that the low-demand manufacturer can always profitably mimic the strategy of the high-demand manufacturer. Any attempt by the latter to distinguish itself from the former is therefore bound to fail, which means that in any Perfect Bayesian equilibrium, the equilibrium is not separating.

We can formalize this insight as follows:

Lemma 1 *Only pooling equilibria arise in the game in which the retailer adopts a policy of not accepting slotting allowances. The retailer's updated beliefs about each manufacturer's type (after observing the wholesale prices) will therefore be the same as his prior beliefs.*

Because the retailer's prior beliefs about each manufacturer's product are symmetric, and because the manufacturers are unable to signal their true demands (there can be no Bayesian updating of the retailer's prior beliefs about the demand potential of the manufacturers' products given that all equilibria are pooling), it will be optimal for the retailer in stage three to pick whichever product comes with the lower wholesale price. Knowing this, competition among the manufacturers then ensures that each manufacturer in stage two will offer a wholesale price that is equal to its product's marginal cost, i.e. $w_1^* = w_2^* = 0$. As a result, the retailer in stage three can do no better than to pick one of the products at random. His expected profit in stage three (and thus his expected profit in stage one from adopting a policy of not accepting slotting allowances) from choosing product i will therefore be equal to v times the expected value of product i's demand,

$$\bar{\pi}_r^W = (v - w_i^*)\, E\left[q_i\right] = vE[q_i] = v\left(\frac{a+b}{2}\right),$$

where we have used the fact that the retailer's profit-maximizing price in stage three is v. Here, the superscript W denotes the case in which only wholesale prices are permitted.

Now suppose that the retailer adopts a policy of accepting slotting allowances. As before, this might take the form of a store or chain-wide policy, or it might simply reflect

[18] FTC (2003) offers numerous examples of product categories in which slotting allowances play little or no role, and other product categories in which the retailer almost always asks for slotting allowances.

the retailer's norms and expectations in a particular product category. Either way, it makes a big difference for the analysis when the retailer accepts slotting allowances because this makes it possible for the manufacturers to signal the demand potential of their products. It is easy to show, for example, that there will always be separating equilibria in this case.

We can formalize this insight as follows:

Lemma 2 *Only separating equilibria arise in the game in which the retailer adopts a policy of accepting slotting allowances. The retailer will therefore be able to infer which manufacturer's product is better (and which product is worse) by observing their offers.*

Proof: See the appendix.

One might think that being able to infer which manufacturer's product is better would always benefit the retailer, but in fact, the opposite holds when both products are known to be good. To see this, note that although overall profits would be expected to increase when the retailer accepts slotting allowances (because the better product will always be carried), so too will the manufacturers' profits. In the absence of slotting allowances, the manufacturers earn zero profit in equilibrium (because the equilibrium wholesale prices are zero). However, with slotting allowances, the manufacturers' profits will not necessarily be competed down to zero because of the possibility of signaling. Indeed, signaling allows the manufacturer of the better product to earn rents equal to the added value of its product. Whether the retailer will be better off then depends on whether the expected increase in the overall profits is greater than or less than the increase in the manufacturers' profits.

To determine which is greater, and thus to determine the expected impact on the retailer's profit, we need to characterize the equilibrium offers in stage two. Using Lemma 2, we show in Appendix A that if attention is restricted to contract offers that would yield non-negative profit for a manufacturer if its product were carried (hereafter we will assume that this is the case),[19] then not only are all equilibria separating, they are also payoff equivalent (i.e., each player earns the same payoff in every equilibrium). Moreover, we also show in Appendix A that the better product (i.e., the one that has the greater demand potential) will always be the one carried, and that the manufacturer of this product will leave the retailer with just enough surplus that the retailer does not pick the rival manufacturer's product (in equilibrium, the retailer will be indifferent). It remains only to pin down the losing manufacturer's offer. But this too is straightforward. Under our refinement that rules out contract offers that would yield non-negative profit, it is easy

[19] For example, this rules out contract offers in which the manufacturer's wholesale price is below cost.

to show that the manufacturer that has the lower demand potential will offer the retailer a contract that would yield zero profit for the manufacturer if its product is carried.[20]

It follows that the retailer's expected profit in stage one from adopting a policy of accepting slotting allowances when both products are good will therefore be given by

$$\begin{aligned}\bar{\pi}_r^S &= E\left[\left(v - w_{(1)}^*\right) q_{(1)} + S_{(1)}^*\right] = vE\left[q_{(2)}\right] \\ &= v \sum_{i=1}^{2} \int_a^b \left(1 - \frac{q_i - a}{b-a}\right) \frac{q_i}{b-a} \mathrm{d}q_i \\ &= v \left(\int_a^b \left(1 - \frac{q_1 - a}{b-a}\right) \frac{q_1}{b-a} \mathrm{d}q_1 + \int_a^b \left(1 - \frac{q_2 - a}{b-a}\right) \frac{q_2}{b-a} \mathrm{d}q_2 \right) = v\left(\frac{2a+b}{3}\right).\end{aligned}$$

To see this, note that because the high-demand manufacturer only has to compensate the retailer for what he could earn if he carried the low-demand manufacturer's product, the retailer's expected profit when both products are good (and he accepts slotting allowances) is equal to his expected profit from selling the less profitable of the two products. Breaking this expected profit into its component parts, we have that $v \int_a^b \left(1 - \frac{q_2 - a}{b-a}\right) \frac{q_2}{b-a} dq_2$ is the expected profit the retailer can expect to receive from manufacturer 1 given all possible values for q_2, where for a given q_2, $\left(1 - \frac{q_2 - a}{b-a}\right)$ is the probability that $q_1 > q_2$ (to ensure that manufacturer 1 will be the winning manufacturer), vq_2 is the payoff manufacturer 1 would have to offer the retailer in order to win his patronage,[21] and $\frac{1}{b-a}$ is the density. And similarly for the expected profit the retailer can expect to receive from manufacturer 2. Summing up these expected profits yields the expression for $\bar{\pi}_r^S$ above. Here, the superscript S denotes the case in which the retailer accepts slotting allowances.

Comparing the retailer's expected profit when he does not accept slotting allowances with his expected profit when he does accept slotting allowances, we can easily see that $\bar{\pi}_r^W - \bar{\pi}_r^S = v\left(\frac{b-a}{6}\right) > 0$ for all $b > a$, which implies that the retailer will always be better off adopting a policy of not accepting slotting allowances when both products are good.

Proposition 1 *When the manufacturers have private information about the demand potential of their products, and the choice is between two good products (i.e., $0 \leq s \leq va$), the retailer will always be better off adopting a policy of not accepting slotting allowances.*

Proposition 1 seemingly contradicts much of the extant literature on the use of slot-

[20] If the losing manufacturer were to hold out for some positive profit in the out-of-equilibrium event its product were carried, the winning manufacturer would not have to give as good of terms to the retailer to make her indifferent, which implies that the losing manufacturer would then be able to profitably deviate.

[21] For example, if $w_{(1)}^* = v$, then the retailer can expect to receive a slotting allowance of $vE\left[q_{(2)}\right]$.

ting allowances as a signaling device, which finds not only that slotting allowances can profitably be used to signal private information about demand (Kelly, 1991; Chu, 1992; Desai and Srinivasan, 1995; Lariviere and Padmanabhan, 1997 and Desai, 2000), but also that it typically results in a win-win situation. In this literature, the manufacturers of the best products are better off when slotting allowances are used because they can then truthfully signal that their products will have higher demand; the retailers are better off because by carrying the products of the manufacturers that offer the best combination of wholesale prices and slotting allowances, they can avoid making costly stocking mistakes;[22] and consumers are better off because slotting allowances ensure that the best products are indeed carried. To the contrary, we have found that although it is true that slotting allowances can be used to signal a product's demand potential, and although it is true that the retailer would be informed about the best products (the use of slotting allowances ensures that all equilibria are separating), the retailer may nevertheless prefer to remain in the dark and force the manufacturers to compete using linear wholesale prices only.

The intuition for this is as follows. For the retailer, there is a trade-off between inducing competition between the manufacturers and having them truthfully reveal their products' sales potential. On the one hand, allowing the use of slotting allowances ensures that the better product will be stocked, but it is not efficient in extracting rent for the retailer. On the other hand, forcing the manufacturers to compete only on wholesale prices is efficient for rent extraction, because it increases the competition between the manufacturers, but it cannot guarantee that the product with the higher demand will be chosen. When both products are good, Proposition 1 implies that the competitive effect always dominates.

The reason why there is no counterpart to our finding in the previous literature is that the previous literature mainly studies a situation where a retailer faces a single manufacturer that has private information about whether its product is "good" (in the sense that it would always be profitable for the retailer to introduce) or "bad" (in the sense that it would never be profitable for the retailer to introduce). In this situation, it is not surprising that the retailer will always want the manufacturer to signal its product's demand potential — the retailer would incur a loss if it chose the bad product, and there is no countervailing competitive effect because the retailer cannot generate competition when there is only one manufacturer. By not allowing for more than one manufacturer with uncertain demand, the previous literature has missed this countervailing effect.

Our finding in Proposition 1 is a "negative result" in the sense that it specifies a set of circumstances under which we would *not* expect to observe slotting allowances.

[22] The implicit assumption here is that the manufacturers of the weaker-selling products would not be able to match the higher slotting allowances offered by the manufacturers of the better-selling products.

Nevertheless, before proceeding in the next section to focus on positive results (i.e., on conditions under which we *would* expect to observe slotting allowances), it should be noted that this negative result is of interest in its own right because it may shed light on why slotting allowances are typically *not* observed when retailers are deciding how best to allocate shelf space within their stores (FTC, 2003).[23] It is well known that not all shelf space is created equal. Some spaces within a store are scarcer than other spaces. Yet the evidence from the FTC's study on slotting allowances in the grocery industry suggests that slotting allowance payments typically do not guarantee any particular shelf placement, and that the premium spaces in the store (i.e., the area at eye level, near the cash register, or at the end of the aisles) tend to be auctioned off with the bidding focused instead on which of the manufacturers of the retailer's established products can offer the most attractive promotional deals (quantity discounts, discounts off the list price, etc).

This evidence is consistent with what we would predict from Proposition 1. It seems reasonable to assume that the retailer already knows which of his established products will sell the best (and thus that he already knows that the products that he is considering for his premium spaces would do better than, for example, some randomly chosen set of products), but he may be uncertain as to which of his top products would benefit the most from the increased visibility. In this case, by reinterpreting $q_i \geq 0$ as the additional quantity that manufacturer i can expect to sell as a result of obtaining the premium space (i.e., over and above what it would expect to sell under the status quo in its current location), our findings suggest that the retailer should indeed offer to allocate his premium spaces to the manufacturers that offer him the best promotional deals all else equal. That is, the retailer should not adopt a policy of accepting slotting allowances in these cases.

5 The case for accepting slotting allowances

In the previous section, the retailer's choice was between two good products. In this section, we extend the model to allow for the possibility that one or both products may be "bad." By a bad product, we mean a product that if carried would always earn less profit for the retailer than his opportunity cost of shelf space, even if he could procure the product at cost. This is an important extension to consider because the FTC study (mentioned above) also found that although slotting allowances are not commonly observed when

[23] An informal argument that is sometimes offered to explain this phenomenon is that the manufacturers of established products simply have more bargaining power and are thus better able to resist retailer demands for slotting allowances. However, this cannot explain why these same manufacturers do pay slotting allowances to the same retailers when they are trying to obtain shelf space for their *new* products.

retailers are deciding where in their stores to place their already established products, they are commonly observed when the retailers are deciding which new products to carry.

One of the main characteristics of new products is that they often fail.[24] We incorporate the possibility that one or both products may be bad by assuming now that the retailer's opportunity cost of shelf space lies between va and vb, $s \in [va, vb)$, which implies that his opportunity cost is less than what he could earn from carrying manufacturer i's product if q_i were to realize its maximum sales potential of b, but greater than what he could earn from carrying manufacturer i's product if q_i only obtained its lower bound of a. This complicates the model because now the winning manufacturer must not only beat its rival's best offer, it must also ensure that the retailer earns a profit of at least $s \geq va$.

As in the previous section, our focus is on whether the retailer will be better off adopting a policy of accepting slotting allowances or not accepting slotting allowances.

When the retailer adopts a policy of not accepting slotting allowances, the analysis is similar to that in the previous section. In the continuation game that follows, it is still the case that Lemma 1 holds (i.e., there can be no signaling in equilibrium when only wholesale price offers are allowed whether or not the products may be bad), and thus it is still the case that the second stage offers will be such that $w_1^* = w_2^* = 0$ (for the same reasons as before). The only difference is that now, in stage three, the retailer will pick one of the two products at random if and only if his expected profit from doing so exceeds s. Otherwise, he will reject both products and earn his opportunity cost. Given that q_i is uniformly distributed between a and b with density $1/[b-a]$, it follows that the retailer's expected profit in stage one from adopting a policy of not accepting slotting allowances is

$$\bar{\pi}_r^W(s) = \max\left\{s, v\left(\frac{a+b}{2}\right)\right\}. \tag{1}$$

In contrast, the manufacturers will be able to signal their demand potential when the retailer adopts a policy of accepting slotting allowances. In this case, in the continuation game that follows, all equilibria will be separating (i.e., Lemma 2 will continue to hold for the same reasons as before), and either both products will be rejected (which happens when the retailer perceives that neither product would allow it to earn a profit of at least s), or the losing manufacturer will offer its product at cost and the winning manufacturer i's offer will be such that the retailer will just be indifferent between carrying product i or going with his next best alternative.[25] Among other things, this means that any

[24] The FTC (2003) reports that the failure rate on new products can be as high as 70%.

[25] The retailer's next best alternative to manufacturer i's product will either be the outside option that gives him his opportunity cost of shelf space s (if manufacturer j's product is revealed to be bad), or

manufacturer that offers the retailer an expected profit of less than s will not be accepted. And, it means that when we allow for the possibility that one or both products may be bad, and the retailer adopts a policy of accepting slotting allowances, the retailer will earn s in stage three if at least one product is bad,[26] and $vq_{(2)}$ if both products are good.

To see what this implies about the retailer's expected profit in stage one, note that for all $s \in [va, vb]$, the ex-ante probability that at least one product will be bad is given by[27]

$$\Pr(vq_1 < s \vee vq_2 < s) \;=\; \left(\frac{s/v - a}{b - a}\right)^2 + 2\left(1 - \frac{s/v - a}{b - a}\right)\frac{s/v - a}{b - a},$$

while the ex-ante probability that both products will be good is given by

$$\Pr(vq_1 \geq s \cup vq_2 \geq s) \;=\; \left(1 - \frac{s/v - a}{b - a}\right)^2.$$

It follows that when one or both products may be bad (or good), the retailer's expected profit in the first stage from adopting a policy of accepting allowances is given by[28]

$$\begin{aligned} \bar{\pi}_r^S(s) \;=\; & \Pr(vq_i < s \vee vq_j < s) \times s \;+ \\ & \Pr(vq_1 \geq s \cup vq_2 \geq s) \times vE[q_{(2)}|q_1, q_2 \geq s], \end{aligned} \tag{2}$$

where

$$vE[q_{(2)}|q_1, q_2 \geq s] \;=\; v \sum_{j=1,2} \int_{s/v}^{b} \left\{ \left(1 - \frac{q_j - s/v}{b - s/v}\right) \frac{q_j}{b - s/v} \right\} dq_j$$

is the expected payoff of the retailer conditional on both products being good.

This expression captures the main upside of adopting a policy of accepting slotting allowances (or at least of entertaining the possibility of slotting allowances) as it is always weakly greater than s (strictly greater for all $s < vb$). Since the retailer is not disadvantaged in the other cases relative to the status quo (the retailer's payoff in these other cases is always s), it follows that an important benefit of accepting slotting allowances is that it allows the retailer to take a chance on both products being good, reaping the benefits

---

manufacturer j's product (if, like manufacturer i's product, manufacturer j's product turns out good).

[26] If both products are bad, the retailer will realize this from the offers he receives and go instead with his outside option. If only product j is bad, the retailer will carry manufacturer i's product in equilibrium, but manufacturer i will only need to offer the retailer an expected profit of s to induce him to do so.

[27] Here, $\left(\frac{s/v-a}{b-a}\right)^2$ denotes the probability that both will be bad, and $2\left(1 - \frac{s/v-a}{b-a}\right)\frac{s/v-a}{b-a}$ denotes the probability that only one will be bad, recognizing that this can be product 1 or 2 with equal probability.

[28] Note that this expression is valid only for values of $s \in [va, vb]$. If $s < va$, then $\bar{\pi}_r^S(s) = v(2a + b)/3$ (as we showed previously), and if $s > vb$, then both products are always bad and $\bar{\pi}_r^W(s) = \bar{\pi}_r^S(s) = s$.

if they are, while ensuring that he can earn s if one or both of them turn out to be bad.

We can push this insight further, and at the same time express the expected profit of the retailer in a more transparent way, by simplifying (2) and rewriting it as

$$\bar{\pi}_r^S(s) = s + v\left(\frac{(b - s/v)^3}{3(b-a)^2}\right), \tag{3}$$

where s comes from adding $\Pr(vq_1 \geq s \cup vq_2)s$ to the first line in (2), and $v\left(\frac{(b-s/v)^3}{3(b-a)^2}\right)$ comes from subtracting $\Pr(vq_1 \geq s \cup vq_2)s$ from the second line in (2) and simplifying.

Recall from (1) that the retailer can also always earn an expected profit of at least s and sometimes more by forcing the manufacturers to compete through wholesale prices, but the difference in (1) vis a vis the profit in (3) is that for all $s \in (va, vb)$, the retailer only earns more than s if his expected profit from picking one of the products at random exceeds s. In contrast, the retailer's expected profit from accepting slotting allowances is always greater than s in this range. This yields our first main result in this section: if $s \in [v\left(\frac{a+b}{2}\right), vb)$, the retailer should always adopt a policy of accepting slotting allowances.

In the case of $s \in (va, v\left(\frac{a+b}{2}\right))$, there is a tradeoff. On the one hand, the retailer earns an expected profit of $\bar{\pi}_r^W = v\left(\frac{a+b}{2}\right)$ with wholesale prices only. On the other hand, the retailer earns an expected profit of $\bar{\pi}_r^S = s + v\left(\frac{(b-s/v)^3}{3(b-a)^2}\right)$ with slotting allowances. It follows that the retailer will prefer slotting allowances only if $\bar{\pi}_r^S(s) \geq \bar{\pi}_r^W$, or whenever

$$s + v\left(\frac{(b - s/v)^3}{3(b-a)^2}\right) \;\geq\; v\left(\frac{a+b}{2}\right). \tag{4}$$

As we have seen previously, the tradeoff is always resolved in favor of not accepting slotting allowances when s is equal to its lower bound of va (lower bound in this section). But notice that for values of s that are above this level, the left-hand side of (4) is increasing in s, whereas the right-hand side of (4) is independent of s. This implies that the left-hand side will at some point be greater than the right-hand side as s continues to increase throughout its range, and that when this happens, the left-hand side will always continue to be greater. This gives us our second main result in this section: accepting slotting allowances is more likely to be preferred by the retailer the greater is the retailer's opportunity cost of shelf space. More formally, what the monotonicity in s tells us is that there necessarily exists a critical value of $s \in (va, \left(\frac{a+b}{2}\right)]$ such that for all values of s that exceed this critical value, the retailer will strictly prefer to accept slotting allowances.

Taken together, these results establish the case for accepting slotting allowances. Moreover, they suggest that the retailer's opportunity cost of shelf space will play a critical role in determining whether and when the retailer will want to accept slotting allowances.

The model also has implications for when the retailer might want to adopt a policy of accepting slotting allowances as a function of the mean and variance of the demand potential of the manufacturers' products. Thus, we can rewrite the inequality in (4) as:

$$s + v\left(\frac{\left(\overline{q} + \sqrt{3}\sigma - s/v\right)^3}{36\sigma^2}\right) \geq v\overline{q}, \tag{5}$$

where $\overline{q} = \frac{a+b}{2}$ denotes the mean and $\sigma = \frac{b-a}{\sqrt{12}}$ is the standard deviation. Conducting comparative statics with respect to $\overline{q}$, we can see that although both sides of (5) are increasing in $\overline{q}$, the right-hand side is increasing faster. This implies that, unlike with increases in s, accepting slotting allowances is less likely to benefit the retailer the higher is $\overline{q}$ all else equal. With respect to σ, we can see that while the right-hand side of (5) is independent of σ, the left-hand side of (5) is increasing in σ in the relevant range. All else equal, this implies that adopting a policy of accepting slotting allowances is more likely to benefit the retailer versus a policy of not accepting slotting allowances the higher is σ.

More formally, for a given $\overline{q}$, we can establish the relationship between s, σ and the retailer's decision whether to adopt a policy of accepting slotting allowances by letting $s^*(\sigma)$ denote the value of s that solves $\bar{\pi}_r^S(s) = \bar{\pi}_r^W$. Our comparative statics imply that $s^*(\sigma)$ is decreasing in σ. This can be seen by plotting $s^*(\sigma)$, as we do in Figure 1 below. Figure 1 illustrates the retailer's optimal decision whether to accept slotting allowances for a given a and b as a function of s and σ, where σ is along the horizontal axis, and s is along the vertical axis. As shown there, when s is sufficiently large, the retailer will always be better off adopting a policy of accepting slotting allowances, whereas when s is sufficiently small, the retailer will never want to adopt a policy of accepting slotting allowances. For intermediate values of s, the retailer will be better off adopting a policy of accepting slotting allowances if and only if σ is sufficiently large. Notice that the threshold of s above which the retailer will always be better off adopting a policy of accepting slotting allowances occurs at the point where $s = \overline{q}$ (which is obtained by evaluating $s^*(\sigma)$ in the limit as $\sigma \to 0$), while the threshold of s below which the retailer will never want to adopt a policy of accepting slotting allowances occurs approximately at the point $s \approx .88\overline{q}$ (which, holding $\overline{q}$ constant, is obtained by evaluating $s^*(\sigma)$ at the maximum value of σ).[29]

We can summarize our results in this section as follows:

[29] For a given $\overline{q}$, the maximum value of σ occurs where $b = 2\overline{q}$ and $a = 0$.

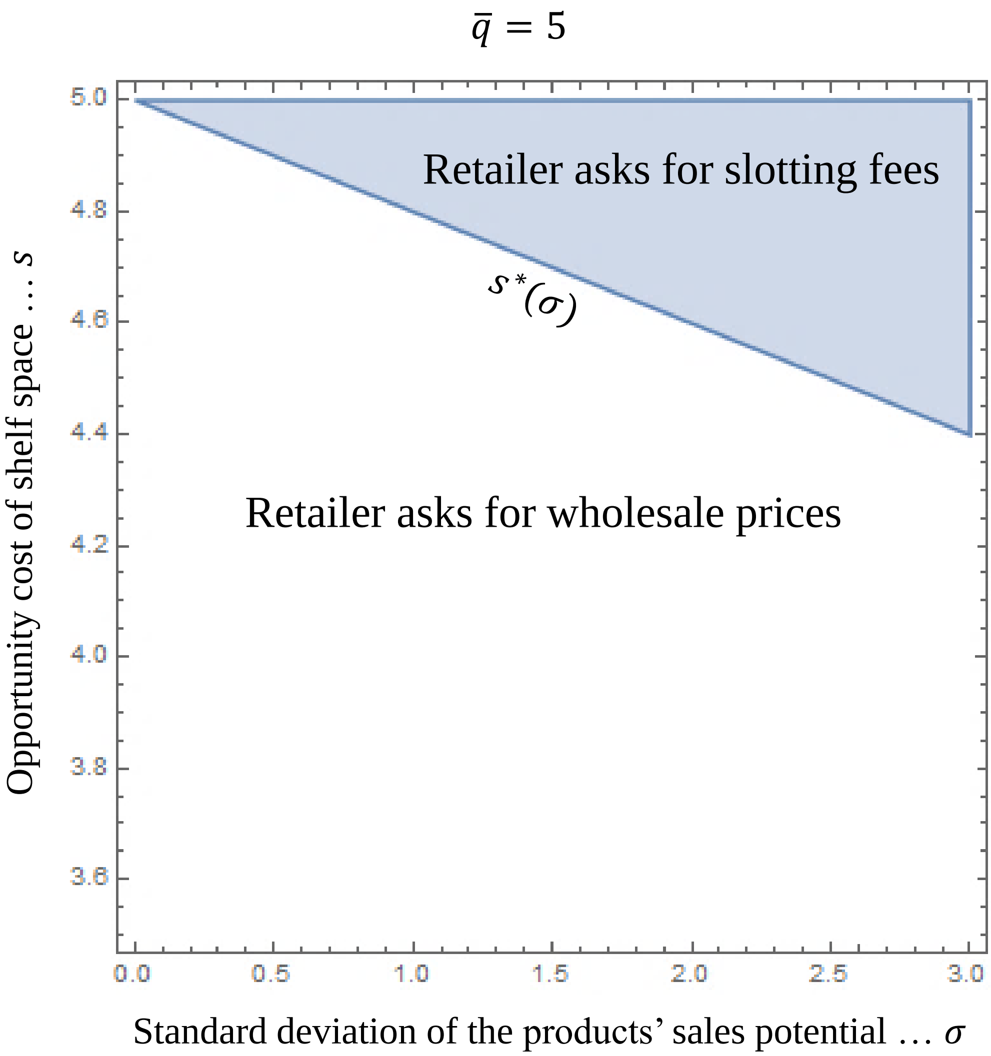

Figure 1: The critical $s^*(\sigma)$

Proposition 2 *When the manufacturers have private information about the demand potential of their products, and the retailer does not know ex-ante whether the individual products are good or bad relative to his opportunity cost of shelf space (i.e., $va < s < vb$),*

- *there exists a critical threshold $\overline{s} = v\overline{q}$ such that for all $s \in [\overline{s}, vb)$, the retailer will be better off adopting a policy of accepting slotting allowances;*
- *there exists a critical threshold $\underline{s} > va$ such that for all $s \in (va, \underline{s}]$, the retailer will be better off adopting a policy of not accepting slotting allowances; and*
- *for all $s \in (\underline{s}, \overline{s})$, the retailer will be better off adopting a policy of accepting slotting allowances if and only if $s > s^*(\sigma)$, where $s^*(\sigma) \in (\underline{s}, \overline{s})$ is decreasing in σ.*

Proposition 2 shows that when the retailer does not know ex-ante whether the individual products he is considering are good or bad, there are two key factors that determine whether he will want to accept slotting allowances: his opportunity cost of shelf space and the variance of the demand potential of the manufacturers' products. Of these, the retailer's opportunity cost is more important in the sense that there are settings in which the retailer's decision is independent of the variance of the manufacturers' demand potential, whereas the retailer will always want to adopt a policy of accepting slotting allowances if his opportunity cost is high enough. How high it needs to be depends on the variance of the manufacturers' demand potential. The greater the variance, the lower is the threshold.

Specifically, we have shown that when the opportunity cost of shelf space is greater than $v\overline{q}$, the retailer will always want to accept slotting allowances. This enables the manufacturers to signal their demand potential, and the product with the higher demand potential, assuming it is greater than s/v, will be stocked. On the other hand, when the opportunity cost of shelf space is less than $v\overline{q}$, the retailer faces a tradeoff. For a low enough opportunity cost, a policy of not accepting slotting allowances will be profitable, whereas if it is high enough (but still below $v\overline{q}$) slotting allowances should be accepted.

One can best understand the intuition for these results by first understanding why the retailer may sometimes prefer to remain in the dark — the reason is because accepting slotting allowances comes with both an upside and a downside. The upside is that both products may turn out to be better than expected, in which case the retailer benefits because even the losing manufacturer would then have to offer the retailer better than expected profits. The downside is that one or both products may turn out to be bad, in which case the offer by the manufacturer of the better product would be weakly worse (strictly worse if $s < v\overline{q}$) for the retailer than what he could have obtained in the absence of slotting allowances. Balancing these two sides determines the retailer's optimal strategy.

Notice, however, that as the opportunity cost of shelf space increases, less weight is placed on the downside, making accepting slotting allowances relatively more desirable. Less weight is placed on the downside as s increases because increases in s raise the retailer's floor: if both products turn out to be bad, the retailer can simply say "no" to both products, earn s, and limit the downside damage, if any, relative to $v\overline{q}$. It is as if increases in s provide insurance for the retailer against a bad outcome. At the same time, increases in s have no effect on the retailer's expected profit from not accepting slotting allowances whenever this profit exceeds s. Thus, accepting slotting allowances becomes more and more desirable relative to not accepting them as s increases for all $s < v\overline{q}$. For $s \geq v\overline{q}$, the retailer's expected profit from not accepting slotting allowances is fixed at s, whereas his expected profit is always greater than s when he accepts slotting allowances.

These findings, along with our findings in Proposition 1, may explain why one tends to observe slotting allowances on new products but not on established products, and why, when one does observe slotting allowances on new products, usage is not uniform. In particular, our results have implications for some of the anomalies identified in the 2003 FTC staff report cited above. For example, they suggest the following testable implication:

Testable implication: *A higher opportunity cost of shelf space will, all else equal, be associated with a higher incidence of slotting allowances on new product introductions.*

It is well known that the incidence of slotting allowances varies across product categories, and that some of the highest incidences of slotting allowances occur on frozen and refrigerated products, where shelf space is typically more costly to expand than usual. This observation can be explained by our result simply because the opportunity cost of shelf space may be higher for these products than they are for non-refrigerated food items. Our results may also explain why one might observe differences in the incidence of slotting allowances between different outlets of the same retailer.[30] It is also well known that stores located in urban areas tend to have higher opportunity costs of shelf space than stores located in rural areas, and similarly, small convenience stores tend to have higher opportunity costs of shelf space than large supermarket stores. An implication of our result, therefore, is that we would expect to see a higher incidence of slotting allowances in convenience stores than we would in other stores of the same retailer (e.g., a higher incidence of slotting allowances in *Tesco Express* than in *Tesco Extra*, and in stores that are located in the city than in stores of the same retailer that are located in the suburbs.

Our findings also give rise to some counterintuitive implications. Consider a product category that is doing well in the sense that there are no slow-moving products, and compare it to a product category that is doing poorly in the sense that there are many slow-moving products. Then, our results suggest that we are more likely to see the retailer adopting a policy of accepting slotting allowances for the category that is doing well (where the opportunity cost of dropping a product is relatively high) than for the category that is doing poorly (where the opportunity cost of dropping a product is relatively low). But this is counterintuitive because one might have guessed that it is precisely when the category is performing poorly that the retailer should place more emphasis on choosing the best products and not making mistakes. One might even be tempted to ascribe the causality incorrectly, believing that the poorly performing category is performing poorly *because* the retailer has chosen not to allow the manufacturers to signal their private information.

[30] Store by store variation of the same retailer is difficult to explain from a bargaining perspective.

In fact, the opposite may be true; the retailer may have chosen not to allow the manufacturers to signal their private information *because* the category was performing poorly.

Another implication of Proposition 2 has to do with the relationship between the incidence of slotting allowances and the variance in the demand potential of the products:

Testable implication: *A higher variance in the demand potential of the manufacturers' products will, all else equal, be associated with a higher incidence of slotting allowances.*

The prediction from Proposition 2 is that categories that consistently have a higher demand variance will tend to have a higher incidence of slotting allowances. This may explain why the incidence of slotting allowances typically varies across product categories that differ in their mix of large and small manufacturers. On the one hand, in a category that is dominated by a few large, well known manufacturers, the retailer may expect a lower demand variance, which tends to increase the likelihood that he will opt for using simple wholesale prices. On the other hand, in a category that is more fragmented and tends to consist of new product offerings from small, lesser known manufacturers, the retailer may expect a higher demand variance, and more often ask for slotting allowances.[31] These predictions are indeed consistent with claims by small manufacturers that they are disproportionately asked to pay slotting allowances relative to their larger counterparts.

There have also been other attempts to explain the observed variations in the incidence of slotting allowances. One informal argument that is sometimes made to explain the variation between small and large manufacturers is that the observed variation has to do with differences in the manufacturers' bargaining power. According to this argument, smaller, lesser known producers tend to have weaker bargaining power relative to larger, more established producers when it comes to introducing new products. Consequently, they may find it difficult to resist the retailers' demands when they are asked to pay slotting allowances, even when the larger producers are able to resist. Notice that the presumption is that the retailers always want slotting allowances, something that we disagree with. Here, we ascribe the observed variation in the incidence of slotting allowances between small and large manufacturers not to differences in the manufacturers' bargaining power, but to the characteristics of the new products that these manufacturers tend to offer.

Another informal argument that we have heard, this one relating the variation in the incidence of slotting allowances to the retailer's opportunity cost of shelf space, is that the higher is the retailer's opportunity cost of shelf space, the more costly it is for the retailer

[31] At this point, we are simply pointing out the implications that would follow if the retailers happen to have these beliefs about large and small manufacturers. In the next section, we will take a closer look at what the underlying primitives behind such beliefs might look like and suggest when they may hold.

to make a mistake. According to this argument, the higher is the retailer's opportunity cost, the more important it will be for the retailer to know the demand potential of the manufacturers' products all else equal, and consequently the more likely it will be that the retailer will ask to be compensated with slotting allowances. In other words, retailers ask for slotting allowances because they wish to avoid making costly mistakes when deciding which new products to carry. The problem with this argument is that it always holds in the sense that there would then never be any reason for the retailer not to ask for slotting allowances: having more information is always better for the retailer than not having the information according to this story. This explanation would need to be supplemented with a story such as ours to understand why slotting allowances are sometimes not observed.

Lastly, an argument that appears in the literature (see Lariviere and Padmanabhan, 1997; and Desai, 2000) to explain why the incidence of slotting allowances would be expected to be increasing in the retailer's opportunity cost of shelf space is that manufacturers of good products will find it more difficult to signal the demand potential of their products when the retailer's opportunity cost of shelf space is high than when it is low. According to this argument, manufacturers prefer to signal the demand potential of their products whenever possible by distorting their wholesale prices upward. And when this is enough to create a separating equilibrium, slotting allowances will not be observed. Only when it is not enough, which occurs when the retailer's opportunity of cost shelf space is sufficiently high, will they be observed. In other words, slotting allowances will only be used as a signaling device when attempts to signal through wholesale prices are exhausted.

In these papers, the presumption is that the manufacturers are the ones dictating how and when the retailers will receive slotting allowances, as opposed to the retailers being the ones who can choose their policy. If the retailer had the power to determine his policy, he would always prefer to receive slotting allowances (because this literature assumes away upstream competition). In contrast, we have found that, even when wholesale prices cannot be used as a signaling device, the retailer may still not want to accept slotting allowances — given that more than one manufacturer is competing for the same space.

6 Asymmetric manufacturers

We have thus far assumed that the manufacturers are symmetric in the sense that the demands for both products are drawn from the same distribution. We now extend the model to allow the manufacturers to be asymmetric. In particular, we have in mind a setting in which a "small" manufacturer competes against a "large" manufacturer for the retailer's available slot. In such settings, small manufacturers often argue that they are at a disadvantage relative to large manufacturers, because the retailers are biased in favor of accepting the large manufacturers' products, all else equal (see FTC, 2001; FTC, 2003).

One way to make sense of this argument is to observe that large manufacturers often have more financial resources to devote to developing new products than small manufacturers do, and that this can lead to offerings that have both a higher mean and a lower variance. In other words, retailers may be biased in favor of accepting the large manufacturers' products simply because the large manufacturers tend to spend more.

We model this by assuming that one firm, manufacturer 1, can have multiple draws from the support $[a, b]$, while continuing to assume that the other firm, manufacturer 2, can have only one draw. A natural interpretation of this is that the large manufacturer is simultaneously working on several ideas for new products, some of which may work out, and some of which may not, and that at the end of the day, it gets to pick the best of its draws before presenting its case to the retailer. The small manufacturer, on the other hand, does not have the luxury of working on several ideas simultaneously, and therefore only has "one shot." Its product may still turn out to be better in the sense of having a higher demand potential than the best of the large manufacturer's products, but this will be less likely the more draws from the distribution the large manufacturer is able to take.

The question we now ask is how does this interaction between small and large firms affect the incidence of slotting allowances? Would we expect to see a higher or lower incidence when the large firm has more resources to devote to new-product development?

As we did in Sections 3 and 4, we begin with the case in which the retailer has adopted a policy of not accepting slotting allowances. Both manufacturers draw from the same uniform distribution with support $[a, b]$, and the only difference between the manufacturers is that now manufacturer 1 can draw more than once. We let $k \geq 1$ denote the number of draws, and we assume for convenience that the draws are independent. This has the effect of flattening out the distribution of the demand potential of the product that the large manufacturer ultimately presents to the retailer (the best of its draws),[32] as depicted in Figure 2 below, and it means that whereas manufacturer 2's density function will still be $f_2(q_2) = \frac{1}{b-a}$, the density function for manufacturer 1 will now be $f_1(q_1) = \frac{k}{b-a}\left(\frac{q_1-a}{b-a}\right)^{k-1}$.[33] The retailer will thus face a different expected demand and variance on firm 1's best product. Formally, the expected demand and variance of manufacturer 1's best product

[32] It is easy to see that manufacturer 1 will never want to present more than one product to the retailer.

[33] To see this, note that manufacturer 1 has k values for q_1 that are equally likely to be highest. Take one of these values, q_1. Then $\left(\frac{q_1-a}{b-a}\right)^{k-1}$ denotes the probability that all other draws are below q_1.

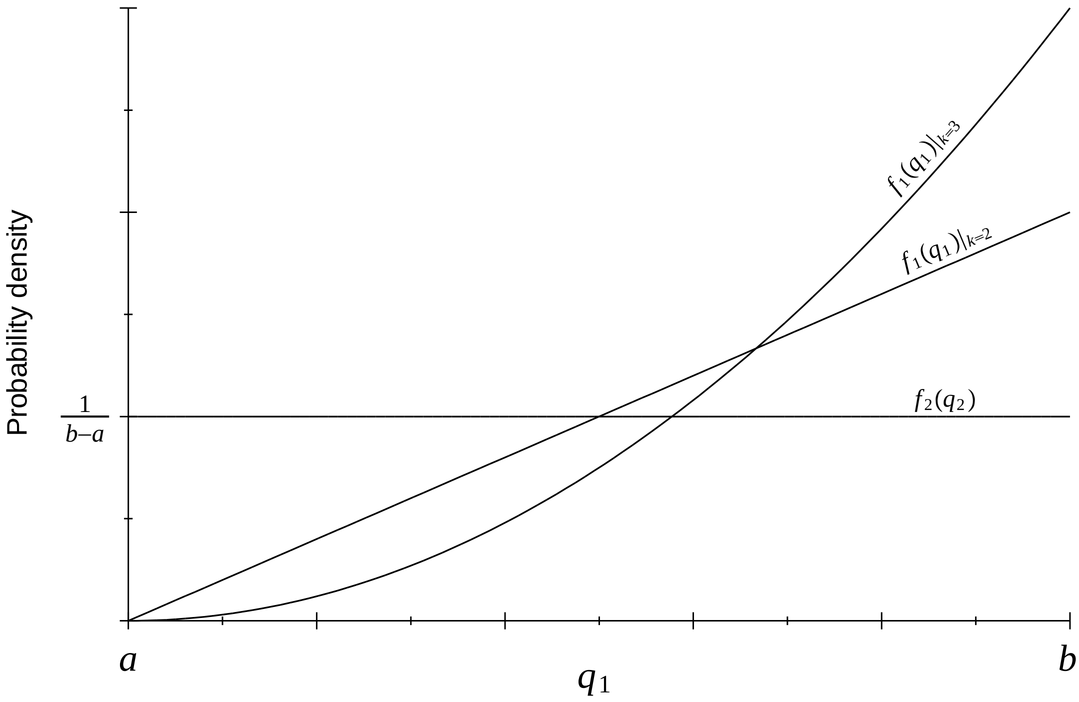

Figure 2: Density function for $k = 1, 2, 3$

when manufacturer 1 can make k draws from a uniform distribution with support $[a, b]$ is

$$E[q_1|k] = k \int_a^b \left\{ \left(\frac{q_1 - a}{b - a} \right)^{k-1} \frac{q_1}{b - a} \right\} \mathrm{d}q_1 = \frac{a + kb}{k + 1},$$

$$\sigma[q_1|k] = k \int_a^b \left\{ \left(\frac{q_1 - a}{b - a} \right)^{k-1} \frac{(q_1 - E[q_1|k])^2}{b - a} \right\} \mathrm{d}q_1 = \frac{k(a - b)^2}{(2 + k)(1 + k)^2},$$

respectively. It follows from this that $E[q_1|k] > E[q_1|1]$ and $\sigma[q_1|k] < \sigma[q_1|1]$ for all $k > 1$. That is, it follows that relative to the demand potential of manufacturer 2's product, the expected demand of manufacturer 1's offering will be higher and its variance will be lower.

This provides some support for the common belief among retailers that the new offerings of larger, more established manufacturers tend to have a lower demand variance than the new offerings of smaller, less established manufacturers (we discussed the comparative-statics implications of this in the previous section). More importantly, for this section, it also implies that in the absence of signaling, the large manufacturer will always be able to offer the retailer a better deal, in expectation, than the small manufacturer can offer.

Solving for the equilibria in the absence of signaling, we can establish the following:

Lemma 3 *When the retailer does not accept slotting allowances, the large manufacturer's product is chosen in all equilibria, the small manufacturer earns zero, and the large manufacturer earns a payoff equal to the expected value added of its product,* $E[q_1|k] - E[q_1|1]$.

Lemma 3 formalizes the oft expressed claims by small manufacturers that retailers are biased in favor of accepting large manufacturers' products, all else equal. What it says is that in the absence of slotting allowances, the small manufacturer has no chance of competing and winning. It also implies that all gains accrue to the large manufacturer in the sense that it gets to keep the value gained from having k draws from the distribution.

We now consider what the retailer earns. Depending on the value of the retailer's opportunity cost of shelf space, there are three cases to consider. Suppose first that $\overline{q} > \frac{s}{v}$. Then, with wholesale price competition, manufacturer 1 always wins, and in equilibrium it will charge a wholesale price w_1^* that is just high enough that the retailer's expected profit will be equal to his expected profit from buying manufacturer 2's product at cost:

$$(v - w_1^*)\,E[q_1|k] \;=\; v\left(\frac{a+b}{2}\right).$$

Suppose next that $\frac{s}{v}$ lies between the expected demand potential of the two manufacturers' products, $\overline{q} \leq \frac{s}{v} \leq E[q_1|k]$. Then, manufacturer 1 will still win the retailer's favor, but now its wholesale price will be such that the retailer's expected profit is equal to s:

$$(v - w_1^*)\,E[q_1|k] \;=\; s.$$

Suppose last that s exceeds the expected demand potential of the two manufacturers' products. Then, both products will be rejected and the retailer will also earn s in profit.

It follows that the retailer's expected profit in the absence of slotting allowances is

$$\bar{\pi}_r^W(s) = \max\left\{s, v\left(\frac{a+b}{2}\right)\right\}, \tag{6}$$

which is the same as it was in the previous section when the manufacturers were symmetric. Although a better product will be carried in expectation, the retailer does not gain from this because the large manufacturer's wholesale price will have increased commensurately.

Whether or not the retailer's expected profit will also be the same when slotting allowances are accepted is a different matter. Things are more complicated then because signaling becomes possible. To consider this case, note first that when $\frac{s}{v} \geq \frac{a+b}{2}$, the retailer necessarily earns s with wholesale price offers only. Hence, when $\frac{s}{v} \geq \frac{a+b}{2}$, the retailer will always be weakly better off when it adopts a policy of accepting slotting allowances. This follows because he cannot be any worse off by doing that, but he might be better off if both manufacturers' products turn out to be good (for the reasons discussed previously).

Suppose now that the retailer's opportunity cost of shelf space is such that $a \leq \frac{s}{v} < \frac{a+b}{2}$. Then, the probability that at least one of the two products will be bad is given by[34]

$$\begin{aligned}
\Pr\left(vq_1 < s \vee vq_2 < s\right) &= \frac{s/v-a}{b-a}\left(\frac{s/v-a}{b-a}\right)^k + \left(1-\frac{s/v-a}{b-a}\right)\left(\frac{s/v-a}{b-a}\right)^k \\
&\quad + \left(\frac{s/v-a}{b-a}\right)\left(1-\left(\frac{s/v-a}{b-a}\right)^k\right),
\end{aligned}$$

while the probability that both manufacturers' products will be good is given by

$$\Pr(vq_1 \geq s \cup vq_2 \geq s) = \left(1-\frac{s/v-a}{b-a}\right)\left(1-\left(\frac{s/v-a}{b-a}\right)^k\right).$$

Since the retailer can expect to earn s in the former case, and v times the expected quantity of the losing manufacturer's product conditional on both products being good,

$$v\left(E\left[q_{(2)}|k, q_1 \geq q_2 \geq s/v\right] + E\left[q_{(2)}|k, q_2 \geq q_1 \geq s/v\right]\right),$$

in the latter case, it follows that the retailer's expected profit in the first stage from adopting a policy of accepting allowances when manufacturer 1 has k draws is given by

$$\begin{aligned}
\bar{\pi}_r^S(s,k) &= \Pr(vq_i < s \vee vq_j < s) \times s \; + \\
&\quad \Pr(vq_1 \geq s \cup vq_2 \geq s) \times vE\left[q_{(2)}|k, q_1 \geq q_2 \geq s/v\right] + \\
&\quad \Pr(vq_1 \geq s \cup vq_2 \geq s) \times vE\left[q_{(2)}|k, q_2 \geq q_1 \geq s/v\right],
\end{aligned} \tag{7}$$

where

$$E\left[q_{(2)}|k, q_1 \geq q_2 \geq s/v\right] = \int_{s/v}^{b}\left\{\left(1-\left(\frac{q_2-s/v}{b-s/v}\right)^k\right)\left(\frac{q_2}{b-s/v}\right)\right\}dq_2$$

[34]The first term on the right-hand side of the expression denotes the probability that both manufacturers' products will be bad, the second term denotes the probability that only manufacturer 1's product will be bad, and the last term denotes the probability that only manufacturer 2's product will be bad.

and

$$E\left[q_{(2)}|k, q_2 \geq q_1 \geq s/v\right] \;=\; k\int_{s/v}^{b}\left\{\left(1-\frac{q_1-s/v}{b-s/v}\right)\left(\frac{q_1-s/v}{b-s/v}\right)^{k-1}\left(\frac{q_1}{b-s/v}\right)\right\}dq_1.$$

Similar to what we did in the previous section, we can express the retailer's expected profit in a more transparent way by letting $\theta \equiv \frac{s/v-a}{b-a}$ and rewriting the expression in (7) as

$$\bar{\pi}_r^S(s,k) \;=\; s+\frac{vk(3+k)(1-\theta^k)(b-s/v)^2}{2(b-a)(2+3k+k^2)}. \tag{8}$$

This makes it easier to compare the retailer's expected profit when slotting allowances are accepted (given in (8)) with his expected profit when they are not accepted (given in (6)).

Our results in this section then follow with the help of three observations.

First, note that $\bar{\pi}_r^S(s,k)$ exceeds s for all $k \geq 1$ and values of s in the feasible range, i.e., for all $a \leq s/v < b$.[35] Here, as in the previous section, we can see that adopting a policy of accepting slotting allowances has option value relative to the status quo. If one or both products turn out to be bad, the retailer's fallback is simply s, same as it is in the status quo. But, there is a chance that both products may turn out to be good, in which case, the retailer can expect to earn more than s. Since the profit the retailer can expect to earn in the absence of slotting allowances obtains its maximum at s for all $\frac{s}{v} \geq \frac{a+b}{2}$, it follows that there is an upper bound on s at $s = v\left(\frac{a+b}{2}\right) = v\bar{q}$ such that the retailer will strictly prefer to accept slotting allowances for all s greater than $v\bar{q}$ in the feasible range.

Second, note that in the limit, as k goes to infinity, $\bar{\pi}_r^S(s,k)$ approaches

$$s \;+\; \frac{v(b-s/v)^2}{2(b-a)},$$

which weakly exceeds $\bar{\pi}_r^W(s)$ for all s in the relevant range, and strictly exceeds it for all $s/v > a$. It follows that there is an upper bound of $\bar{k}$, such that for all $s/v > a$ and k greater than this upper bound, the retailer will always prefer to accept slotting allowances.

Third, note that $\bar{\pi}_r^S(s,k)$ is increasing in k. It is easy to verify that $\bar{\pi}_r^S(s,1)$ collapses to the $\bar{\pi}_r^S(s)$ from the previous section, and that for all $k>1$, $\bar{\pi}_r^S(s,k)$ is strictly higher. Among other things, this means that the payoff from accepting slotting allowances will be higher for the retailer when one of the manufacturers can take more draws, and increasingly so, the more draws the manufacturer can take. Since the payoff from not accepting slotting allowances is independent of k, it follows that for all $k < \bar{k}$, the threshold level of s that

[35] Note that the profit in (8) is valid only for values of $s \in [va, vb]$. If $s < va$, then $\bar{\pi}_r^S(s,k) = v(2a+b)/3$ (as we showed previously), and if $s > vb$, then both products are always bad and $\bar{\pi}_r^W(s) = \bar{\pi}_r^S(s) = s$.

determines when accepting slotting allowances will begin to be profitable will be strictly lower than the corresponding threshold level of s when the number of draws is lower.

To see this, let $\tilde{s}(k)$ denote the value of s that solves $\bar{\pi}_r^S(s,k) = \bar{\pi}_r^W(s)$ for all $s < v\overline{q}$ and $k < \overline{k}$. Then, our observations above imply that $\tilde{s}(k)$ will be strictly decreasing in k. This is illustrated in Figure 3 below for a representative number of draws with $k \geq 1$. Figure 3 illustrates the retailer's optimal decision whether to accept slotting allowances

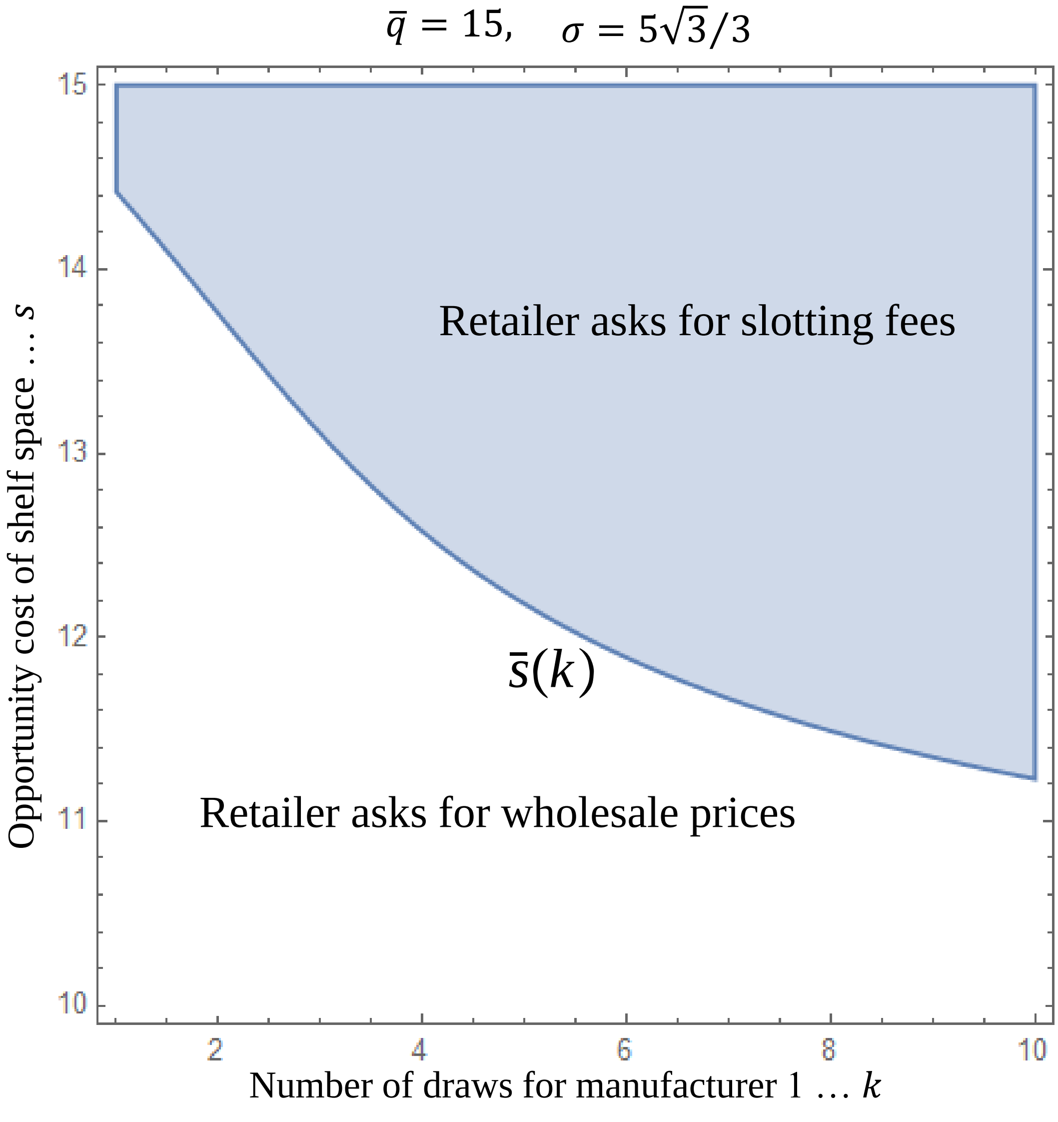

Figure 3: The critical $\bar{s}(k)$

for a given a and b as a function of s and k, where k is given along the horizontal axis, and s is given along the vertical axis. As shown in the figure, when s is sufficiently large, the retailer will always be better off adopting a policy of accepting slotting allowances, whereas when s is sufficiently small, the retailer will never want to adopt a policy of accepting slotting allowances. For intermediate values of s, the retailer will be better off adopting a policy of accepting slotting allowances if and only if k is sufficiently large.

We can summarize our results in this section as follows:

Proposition 3 *When the manufacturers have private information about the demand potential of their products, and the retailer does not know ex-ante whether the individual products are good or bad relative to his opportunity cost of shelf space (i.e., $va < s < vb$),*

- *there exist critical thresholds $\overline{s} = v\overline{q}$ and $\overline{k} > 1$ such that for all $s \in [\overline{s}, vb)$ or $k \geq \overline{k}$, the retailer will be better off adopting a policy of accepting slotting allowances; and*
- *for all $s < \overline{s}$ and $k < \overline{k}$, the retailer will be better off having a policy of accepting slotting allowances if and only if $s > \tilde{s}(k)$, where $\tilde{s}(k) \in (va, \overline{s})$ is decreasing in k.*

One way to understand the intuition behind Proposition 3 is to notice that the retailer's tradeoff turns on whether he is able to capture any of the surplus created when he decides to accepts slotting allowances. We have shown that when both products are known to be good ex-ante, the winning manufacturers always gains more than the actual surplus created (causing the retailer to be worse off with slotting allowances), but that when the individual products may be either good or bad ex-ante, the tradeoff can go either way. Until now, however, the starting point for the (symmetric) manufacturers was zero profit. With asymmetric manufacturers, a difference is that the large manufacturer expects to earn strictly positive profit in the absence of slotting allowances. If the retailer were instead to open its shelf space up for bidding, this manufacturer may or may not see an increase in its profit — because the actual realization of its product's value added may turn out to be higher or lower than the ex-ante expected value added. It may even be the case that the small manufacturer ends up winning the bid. Either way, it is less likely (relative to the symmetric case) that the winning manufacturer will be able to earn more than what the large manufacturer would be able to earn in the absence of slotting allowances, and even when it is able to earn more, it is less likely that its profit will increase by more than the increase in the overall surplus (a necessary condition for the retailer to be worse off with slotting allowances). For these reasons, accepting allowances becomes more likely.

Another way to understand the intuition is to note that the retailer's incentive for accepting slotting allowances turns on its ability capture supra-competitive returns in the event that both products turn out to be good. When one manufacturer is able to take additional draws from the common distribution, the chance that both manufacturers' products turn out to be good increases, and moreover, conditional on both products turning out to be good, the value of the retailer's expected profit will be higher as well. Both a higher expected value and a higher probability that both products will be good make it more likely that accepting slotting allowances will be desirable for the retailer.

The bottom line is that we would expect to see a higher incidence of slotting allowances when large manufacturers can take multiple draws – in effect, take multiple shots at winning the retailer's favor – all else equal. This suggests the following testable implication.

Testable implication: *The incidence of slotting allowances will be higher when small manufacturers are competing against large manufacturers for the retailer's shelf space and the large manufacturers are observed to be spending more on new product development.*

This implication is consistent with the views expressed by small manufacturers that they are more likely to be asked to pay slotting allowances when they are competing against large manufacturers than when they are competing against other small manufacturers (and more likely to pay slotting allowances the greater is the size disparity between the small and large manufacturer, an implication which is also predicted by the model). Typically, however, the small manufacturers express this in terms of the retailers being biased in favor of the large manufacturers, and they express dismay that this puts them at a disadvantage because they cannot afford to pay as much as the large manufacturers can pay. Our results suggest, however, that things would be even worse for them in the absence of slotting allowances. With slotting allowances, a small manufacturer has a chance to earn positive profit in the event that its new product beats the odds and turns out to be exceptionally good, whereas in the absence of slotting allowances, its product will always be passed over.

This implication is also consistent with conventional wisdom which suggests that the pace of new product introductions is one of the key drivers behind the incidence of slotting allowances (see, for example, Sullivan, 1997). According to this view, an increase in the pace of new product introductions puts increasing pressure on the retailer's limited shelf space, causing the retailers to demand slotting allowances to help them sort things out and make better decisions. Our analysis suggests two distinct mechanisms through which this may happen. To the extent that the pace of new product introductions does indeed put increasing pressure on the retailers' limited shelf space, we would expect to see the retailer's opportunity cost of shelf space increasing. All else equal, this would lead to a higher incidence of slotting allowances per the reasons discussed previously. But even without a change in s, we find that the incidence of slotting allowances would still be expected to increase to the extent that the increase in the pace of new product introductions can be attributed to a subset of the manufacturers taking additional draws from the distribution.

7 Conclusion

This paper has explored the implications of accepting slotting allowances when manufacturers have private information about their demands. Contrary to conventional wisdom, which focuses exclusively on how accepting slotting allowances can help retailers make better informed decisions when they are used as a signaling device, we have taken a broader view and noted that accepting slotting allowances often involves a tradeoff for the retailer. On the one hand, we have shown that accepting slotting allowances can indeed increase overall surplus by ensuring that the best products are likely to be sold (in the model, the best product is always sold because the manufacturers have perfect information about how well their products will sell). On the other hand, we have shown that accepting slotting allowances has consequences for how this surplus is distributed. Although the former unambiguously favors accepting slotting allowances, we have shown that the latter can go either way, and may even be enough to outweigh the former. It may be, for example, that slotting allowances enable the winning manufacturer to increase its profit by more than the increase in the overall surplus, in which case the retailer would always be worse off.

We have found that the tradeoff is such that the retailer is more likely to gain from accepting slotting allowances the higher is his opportunity cost of shelf space and the larger is the variance in the demand potential of the manufacturers' products. We have also shown that the retailer is more likely to gain from accepting slotting allowances when small manufacturers are competing against large manufacturers for the retailer's shelf space and the large manufacturers are observed to be spending more on new product development. Lastly, we have offered an explanation for why retailers typically do not accept slotting allowances on established products that are seeking access to their premium shelf space (as opposed to accepting them on new products that are being sold for the first time).

These findings have implications for how and why the incidence of slotting allowances would be expected to vary across small and large manufacturers, retail product categories, and different stores of the same retailer, and are consistent with the observations of these incidences in the well-known FTC (2001) and FTC (2003) reports. The findings also shed light on how the incidence of slotting allowances would be expected to vary depending on whether the new products in a given product category are dominated by a lot of me-too products (lower incidence all else equal), or by a lot of truly innovative products (higher incidence all else equal). They are consistent with the views expressed by small manufacturers that they are more likely to be asked to pay slotting allowances when they are competing against large manufacturers than when they are competing against other small manufacturers, and they are consistent with the view that an increase in the pace of new product introductions can lead to an increase in the incidence of slotting allowances. Our findings also suggest caution when attempting to ascribe blame for a poorly performing category to whether or not the retailer accepts slotting allowances.

We have assumed in the model that the retailer's policy of accepting or not accepting allowances was the same for both manufacturers. This was done to highlight the fact that, ultimately, the retailer must be able to commit to his policy if he is to have credibility, and that this commitment is likely to be easier when it is uniformly applied. That said, it would be interesting to extend the model to allow the retailer to adopt an asymmetric policy, whereby he can choose to treat asymmetric manufacturers differently. This would allow one to answer questions such as whether it is more or less likely that a small manufacturer would be asked to pay slotting allowances relative to a large manufacturer being asked to pay when the small and large manufacturer are competing for the same shelf space slot.

It would also be interesting to extend the model to allow slotting allowances to be used for exclusionary purposes, where the manufacturer who is willing to pay the most when slotting allowances are allowed is not necessarily the one that has the best product. This would then mix pro-competitive motives for slotting allowances with anti-competitive motives for slotting allowances, possibly leading to a more complete set of implications.

We leave these extensions to future research.

Appendix

Proof of Lemma 2.

Recall the structure of the game: The retailer approaches and receives offers from the suppliers in turn. Given that supplier j currently has made the best offer – i.e., supplier j's offer is the standing offer – then, if the rival supplier i does not make a strictly better offer, then the retailer declares supplier j the winner.

Step 1. We show that an equilibrium in which both suppliers are playing pooling strategies, does not exist.

A pooling strategy implies that supplier $i \in \{1, 2\}$, irrespective of her type $q_i \in [a, b]$, makes the same contract offer $\left(w_i^P, S_i^P\right)$. Given that both suppliers play pooling strategies, the retailer's expected profit from accepting either offer is

$$\begin{aligned} E\left[\pi_r^P\right] &= \left(v - w_i^P\right) E\left[q_i\right] + S_i^P \\ &= \left(v - w_i^P\right) \frac{a+b}{2} + S_i^P \geq 0 \end{aligned} \tag{9}$$

The condition that supplier i (irrespective of type) earns non-negative profits, implies that

$$a w_i^P - S_i^P \geq 0 \tag{10}$$

Suppose supplier i's offer $\left(w_i^P, S_i^P\right)$ is the standing offer. Notice that, given supplier i's pooling strategy, supplier $j \neq i \in \{1, 2\}$ could deviate from her own pooling strategy (w_j^P, S_j^P) and offer to the retailer a contract that has $w_j = v$ and $S_j = \left(v - w_i^P\right) \frac{a+b}{2} + S_i^P + \varepsilon$, where $\varepsilon > 0$ is a very small value. The retailer would accept this offer, and supplier j's deviation profit is therefore equal to (ignoring ε)

$$\begin{aligned} \pi_j^{dev} &= v q_j - \left(\left(v - w_i^P\right) \frac{a+b}{2} + S_i^P\right) \\ &= v \left(q_j - \frac{a+b}{2}\right) + \frac{a+b}{2} w_i^P - S_i^P \end{aligned}$$

Given (10), we have that

$$\pi_i^{dev} \geq v \left(q_j - \frac{a+b}{2}\right) + \frac{b-a}{2} w_i^P \tag{11}$$

Given that $q_j > (a+b)/2$ and $w_i^P > 0$, the right hand side of (11) is positive. Thus, all supplier types with demands that are strictly above average would like to deviate from their pooling strategies. And some supplier types with demands that are average and below would also like to deviate (given that $w_i^P > 0$). Hence, an equilibrium in which both suppliers play pooling strategies, does not exist.

The argument above extends to any sub-interval on $[a, b]$, or to any set of sub-intervals on $[a, b]$. We therefore know that two suppliers i and j of types q_i and q_j, will not play the same strategy in equilibrium, unless they are the same type, $q_i = q_j = q$.

Step 2. We show that an equilibrium in which only one of the suppliers is playing a pooling strategy, does not exist.

The logic is the same as above. Suppose supplier i plays a pooling strategy (w_i^P, S_i^P), and supplier j does not. Suppose that supplier j's offer (w_j, S_j) is the standing offer, which gives the retailer a (certain) profit equal to $\pi_r = E\left[\pi_r^P\right] + \varepsilon$ if accepted, where $\varepsilon > 0$ again is a small but positive value, and $E\left[\pi_r^P\right]$ is the profit given by (9). Given (10), π_r has a maximum equal to (ignoring ε)

$$\pi_r = v\frac{a+b}{2} - \frac{b-a}{2}w_i^P$$

Thus, given that i is a type $q_i = a$, and given that $w_i^P < v$, i has no incentive to deviate from her offer (w_i^P, S_i^P). However, given that i is a type $q_i > (a+b)/2$, supplier i could profitably deviate by offering $w_i = v$ and $S_i = v(a+b)/2$. The retailer would accept this offer, and supplier i would therefore earn a profit of

$$vq_i - \frac{a+b}{2} > 0,$$

given that supplier j does not respond with a better offer. Thus, because the two types $q_i = a$ and $q_i > (a+b)/2$ have different best-responses to supplier j's standing offer, supplier i's equilibrium strategy cannot be a pooling strategy.

Again, the argument above extends to any sub-interval on $[a, b]$, or to any set of sub-intervals on $[a, b]$. We therefore know that two types $\bar{q}_i$ and $\underline{q}_i$, will not play the same strategy in equilibrium, unless they are both the same type, $\bar{q}_i = \underline{q}_i$.

*

Given that (semi-) pooling will not occur in any equilibrium, there will be no payoff uncertainty for the retailer when choosing between the suppliers' offers. This implies that either the retailer is able to infer the suppliers' quantities by looking at their offers, or the suppliers offers are such that they remove all payoff uncertainty for the retailer (which implies $w_i = v$ for each retailer, $i \in \{1, 2\}$). *Q.E.D.*

References

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

Canadian Bureau of Competition (2002), *The Abuse of Dominance Provisions (Sections 78 and 79 of the Competition Act) as Applied to the Canadian Grocery Sector*, Ottawa: Canada.

Caprice, S. and V. von Schlippenbach (2012), One-Stop Shopping as a Cause of Slotting Fees: A Rent-Shifting Mechanism, *Journal of Economics & Management Strategy*, 22: 468-487.

Chu, W. (1992), Demand Signaling and Screening in Channels of Distribution, *Marketing Science*, 11: 327-347.

Chu, W. and P. Messinger (1995), Product Proliferation and the Determination of Slotting and Renewal Allowances, *Seoul Journal of Business*, 1: 93-115.

Desai, P. (2000), Multiple Messages to Retailers: Signaling New Product Demand, *Marketing Science*, 19: 381-389.

DeVuyst, C. (2005), Demand Screening with Slotting Allowances and Failure Fees, *Journal of Agricultural & Food Industrial Organization*, Vol 3: No. 2, Article 5, http://www.bepress.com/jafio/ vol3/iss2/art5.

Farrell, J. (2001), Some Thoughts on Slotting Allowances and Exclusive Dealing, remarks prepared for the American Bar Association Section of Antitrust Law, 49th Annual Spring Meeting, Washington, D.C., March 28, 2001.

Federal Trade Commission (2001), *Report on the Federal Trade Commission Workshop on Slotting Allowances and Other Marketing Practices in the Grocery Industry*, www.ftc.gov/opa/ 2001/02/slottingallowancesreportfinal.pdf.

Federal Trade Commission (2003), *Slotting Allowances in the Retail Grocery Industry: Selected Case Studies in Five Product Categories*, www.ftc.gov/opa/2003/11/slottingallowancerpt031114.pdf.

Foros, Ø. and H. Kind (2008), Do Slotting Allowances Harm Retail Competition?, *Scandinavian Journal of Economics*, 110: 367-384.

Innes, R. and S. Hamilton (2006), Naked Slotting Fees for Vertical Control of Multi-Product Retail Markets, *International Journal of Industrial Organization*, 24: 303-318.

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

Kolay, S. and G. Shaffer (2016), Reverse-Fixed Payments in Distribution Channels with a Dominant Retailer, mimeo.

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. and G. Shaffer (2007), Upfront Payments and Exclusion in Downstream Markets, *Rand Journal of Economics*, 38: 823-843.

Marx, L. and G. Shaffer (2010), Slotting Allowances and Scarce Shelf Space, *Journal of Economics & Management Strategy*, 19: 575-603.

Norwegian Competition Authority (2005), Betaling for Hylleplass (Payment for Shelf Space), Report 2/05, Bergen.

Piccolo, S. and J. Miklos-Thal (2012), Colluding Through Suppliers, *Rand Journal of Economics*, 43: 492-513.

Raju, J. and Z.J. Zhang (2005), Channel Coordination in the Presence of a Dominant Retailer, *Marketing Science*, 24: 254-262.

Shaffer, G. (1990), Paying for Shelf Space, working paper.

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

Shaffer, G. (2005), Slotting Allowances and Optimal Product Variety, *The B.E. Journal of Economic Analysis & Policy*, Vol 5: Iss. 1 (Advances), Article 3, www.bepress.com/bejeap/ advances/vol5/iss1/art3.

Shaffer, G. (2013), *Impact of Reverse-Fixed-Payments on Competition*, Office of Fair Trading, London: UK, http://webarchive.nationalarchives.gov.uk/20140402142426/ http://www.oft.gov.uk/OFTwork/research/economic-research/completed-research.

Sullivan, M. (1997), Slotting Allowances and the Market for New Products, *Journal of Law and Economics*, 40: 461-493.

Wilkie, W., Desrochers, D., and G. Gundlach (2002), Marketing Research and Public Policy: The Case of Slotting Fees, *Journal of Public Policy & Marketing*, 21: 275-288.