Multiproduct Intermediaries*

Andrew Rhodes
Toulouse School of Economics

Makoto Watanabe
Department of Economics, VU University of Amsterdam

Jidong Zhou
School of Management, Yale University

May 2017
Very Preliminary

Abstract

This paper offers a framework for studying the optimal product range choice of a multiproduct intermediary, in an environment where consumers demand multiple products and face search frictions. We first demonstrate that the intermediary earns positive profit even if it is no more efficient than small firms at selling products. We then characterize its optimal stocking policy. The intermediary uses exclusively stocked high-value products as loss leaders to increase store traffic, and at the same time earns profit from non-exclusively stocked products which are relatively cheap to buy from manufacturers. We also show that relative to the social optimum, the intermediary tends to be too big and stock too many products exclusively.

Keywords: intermediaries, search, multiproduct shopping, loss leaders, exclusive contracts

JEL classification: D83, L42, L81

*Acknowledgment: to be added

1 Introduction

Many products are traded through intermediaries. Perhaps the most common intermediaries are retailers that buy from producers and sell to consumers.[1] One of the most important decisions for a retailer is to choose which products to stock. Typically consumers are interested in buying a large basket of products, but find it costly and time-consuming to visit several different retailers. Consequently when choosing whether or not to shop at a particular store, they will naturally take into account how much of their desired basket they can buy there, as well as the extent to which that store's products are exclusive (i.e. not available for purchase elsewhere). On the other hand retailers are often constrained in how many products they can stock, for example due to limited stocking space or the fact that stocking too many products can make the in-store shopping experience less pleasant. Indeed even big retailers like Walmart stock only a small fraction of the entire universe of products, and many consumers end up shopping at smaller specialist stores in order to buy some hard-to-find products.[2]

At the same time, to make themselves more attractive large retailers often negotiate exclusive rights to sell certain high-profile products. For instance, Home Depot have many exclusive brands such as American Woodmark in cabinets, and Martha Stewart in outdoor furniture and indoor organization. Many high-end fashion stores also sell unique colors or versions of certain labels. Of course manufacturers need to be compensated more if a retailer wants to stock their products exclusively. Another common way to acquire exclusivity is to through private labels. For instance, Walmart has many well-known private brands such as Sam's Choice and Great Value. Stocking exclusive or private brands has indeed become an important retailing strategy. For example in 2009 such products were responsible for some 40% of Macy's sales.[3]

Although product range and product exclusivity are important choices for retailers, few economic papers have studied them formally. This paper seeks to fill this gap, by building a general but tractable model of product selection. Our paper makes several contributions. Firstly, we show that when consumers have multiproduct demand, a multiproduct retailer can enter a market and make positive profit even if it is no

[1] There are also intermediaries such as brokers and platforms who do not own products but only help match producers and consumers and charge commission fees.

[2] For example in a survey over 60% of people had made a purchase from a small website, since they were unable to find the product they wanted at larger competitors. See http://goo.gl/MV6FRi

[3] This is similar for other major department stores in the US. See http://goo.gl/lfS9QP

more efficient in selling products than manufacturers. This offers a new rationale for the existence of intermediaries (as we discuss more in the literature review). Secondly, we characterize a retailer's optimal product stocking decision, and show how a product's demand curvature and elasticity influences whether it is stocked exclusively or non-exclusively. We also demonstrate that exclusively stocked products are often sold at a loss and used to drive store traffic. Thirdly, we show that a profit-maximizing retailer tends to be too big and stock too many exclusive products relative to the social optimum.

In more detail, we consider a model in which a continuum of manufacturers each produces a different product. Consumers view these products as independent and have elastic demand for all of them. Manufacturers may sell their product directly to consumers via their own retail outlet, or via an intermediary, or through both channels. This intermediary can sell multiple products, and offers two-part tariff contracts to manufacturers for the right to stock their product. The intermediary may demand the exclusive right to sell the product, or allow the manufacturer to keep selling to consumers as well. Consumers observe who sells what, but can only learn a firm's price(s) and buy its product(s) by incurring a search cost. This search cost is heterogeneous across consumers. We assume that the cost to any consumer of searching the intermediary is increasing in the number of products it stocks, as consistent for example with the idea that larger retailers are located further from consumers and may be harder to navigate. In light of this, and allowing for products to have different demand curves, we wish to understand what products the intermediary stocks and whether it chooses to stock them exclusively.

In our model, irrespective of the market structure each supplier of a given product always charges the usual monopoly price. Intuitively, with two-part tariffs the intermediary can get a wholesale price at the marginal cost and avoid double marginalization, and with search frictions the logic of Diamond (1971) implies no price competition even if a product is sold by both its manufacturer and the intermediary. This result greatly simplifies the pricing problem and enables us to focus on product selection. Given monopoly pricing, a sufficient statistic for a product is its monopoly profit π and monopoly consumer surplus v. We can then simply represent products as points in a two-dimensional (π, v) space. The intermediary's problem is then to choose a set of points within (π, v) space that it will stock exclusively, and another set of points which it will stock non-exclusively.

To illustrate the main economic forces in our model in a transparent way, we start by studying a simple benchmark case in which only exclusive contracts are possible, and the cost of searching the intermediary is the same as searching all of the manufacturers whose products it sells (i.e. the intermediary does not reduce consumer search costs). In this case, a consumer will visit the intermediary only if the average v of the stocked products compensates her unit search cost. Compared to direct sales by the manufacturer, this expands demand for the intermediary's low-v products but shrinks demand for its high-v products. We show that there always exists a set of products for which the demand expansion effect dominates, and hence the intermediary generates strictly positive profit by stocking them. Since high-v products are used as loss leaders to attract consumers[4], in the optimal solution the intermediary will only stock those with relatively low π. Conversely, since low-v products are profit generators, the intermediary will stock those with relative high π. Therefore, the intermediary's optimal product selection exhibits "negative correlation" in (π, v) space. We then link areas of (π, v) space back to underlying properties of demand curves such as elasticity and curvature, and argue that those products with a larger and more elastic or more convex demand are more likely to used as loss leaders.

We then proceed to solve for the general case, in which the intermediary can also use non-exclusive contracts, and may offer one-stop shopping convenience such that it is cheaper for consumers to search there compared to shopping around many smaller retailers. The trade-off between exclusively and non-exclusively stocking is that stocking more products exclusively makes it more attractive for consumers to visit because fewer products are available elsewhere, but at the same time is costlier since manufacturers need to be compensated more. The optimal product selection turns out to be qualitatively similar to that in the benchmark case, except that products with high-v and high-π are now stocked as well, though non-exclusively. These additional non-exclusive products make the intermediary more attractive to searchers due to economies of search, but enable non-searchers to buy them directly from manufacturers and so reduce the compensation needed by manufacturers.

Finally, we also compare the intermediary's stocking choice with what a social planner would choose if it seeks to maximize total welfare. An intermediary distorts con-

[4]In our paper the terminology "loss leaders" is used in the sense that the intermediary's revenue from a product is less than what it needs to pay to the manufacturer for the right to stock and sell that product. It does not have the usual connotation that a product is sold at a price below its unit cost.

sumers' search, because it forces them to buy a bundle of products including some low-v products which they ordinarily would not search for. On the other hand, consumers search too little from a welfare perspective, because they only account for their own surplus and ignore the profit earned by firms. Consequently under very weak conditions the social planner finds it optimal to have an intermediary. Nevertheless, the unfettered intermediary tends to stock more products than the social planner would like and often too many of them are stocked exclusively.

The paper proceeds as follows. After a literature review, Section 2 outlines the model. Section 3 examines the simple benchmark case with exclusive contracts and no economies of search, whilst Section 4 studies the general case. Section 5 discusses the interpretation of the (π, v) product space and some extensions of the model which include limited stocking space. We conclude and discuss avenues for further work in Section 6.

1.1 Related literature

There is already a substantial body of literature on intermediaries (see e.g. Spulber (1999)). An intermediary may exist because it improves the search efficiency between buyers and sellers (e.g. Rubinstein and Wolinsky (1987), Gehrig (1993), and Spulber (1996)), or because it acts as an expert or certifier that mitigates the asymmetric information problem between buyers and sellers (e.g. Biglaiser (1993), and Lizzeri (1999)).[5] We also study intermediaries in an environment with search frictions, but in our model an intermediary can profitably exist in the market even if it does not improve search efficiency. This new rationale for intermediaries relies on consumers demanding multiple different products, and this multiproduct feature distinguishes our model from existing work on intermediaries.

The mechanism by which an intermediary makes profit by stocking negatively correlated products in the (π, v) space is reminiscent of bundling (e.g. Stigler (1968), Adams and Yellen (1976), and McAfee, McMillan, and Whinston (1989)). By stocking some products consumers value highly, the intermediary forces consumers to visit and buy other low-value (but fairly profitable) products as well which consumers would other-

[5] In the context of retailers, other possible reasons for retailers to exist include that they may know more about consumer demand compared to manufacturers, they can internalize pricing externalities when products are complements or substitutes, or they may be more efficient in marketing activities due to economies of scale.

wise not buy.[6] However, in bundling models the firm often needs to adjust its prices after bundling to extract more consumer surplus and make bundling profitable. In our model a product's price remains the same no matter who sells it. More importantly our paper focuses on product selection, which is like which products should be "bundled", a question rarely discussed in the bundling literature. In a totally different context about information design, Rayo and Segal (2010) use this same bundling argument to show that an information provider often prefers partial information disclosure in the sense of pooling two negatively correlated prospects into one signal. They consider a discrete framework, and more importantly their information provider can send multiple signals (which is like our intermediary could organize and sell non-overlapped products in multiple stores). This makes the optimization problem in our paper very different from theirs. We also want to emphasize that the investigation of exclusivity arrangements in our paper has no counterpart in either the bundling literature or the above information design paper.

Our paper is also related to the growing literature on multiproduct search (e.g. McAfee (1995), Zhou (2014), Rhodes (2015), and Kaplan et al. (2016)). Existing papers usually investigate how multiproduct consumer search affects multiproduct retailers' pricing decisions when their product range is exogenously given. Our paper departs from this literature by focusing on product selection, another important decision for multiproduct retailers. Moreover our paper introduces manufacturers and so explicitly models the vertical structure of the retail market. In this sense it is also related to recent research on consumer search in vertical markets such as Janssen and Shelegia (2015), and Asker and Bar-Isaac (2016), though those works consider single-product search and address very different economic questions.

Finally, this paper also contributes to the literature on loss leaders (e.g. Lal and Matutes (1994), Chen and Rey (2012), and Johnson (2016)). Loss leaders are usually defined as products sold at a price below the unit cost, but we suggest a broader view of loss leaders: any product which generates a loss for the firm can be regarded as a loss leader if it enables the firm to make more profit from other products. (In our model the loss from a product is because its demand is decreased compared to direct sales such that its revenue is not enough to compensate the manufacturer.) In this broader sense of loss leading, our paper offers a framework that can help study which products

[6] Bundling models need consumers with heterogeneous valuations for each product. In our model consumers have the same valuation for a product but they differ in their search costs, so their net valuation after taking into account the search cost is actually heterogeneous.

should be used as loss leaders and what exclusivity arrangement should be made for them. These questions have not been systematically studied in the existing literature.

2 The Model

There is a continuum of manufacturers with measure one, and each produces a different product. Manufacturer i has a constant marginal cost $c_i \geq 0$. There is also a unit mass of consumers, who are interested in buying every product. The products are independent, such that each consumer wishes to buy $Q_i(p_i)$ units of product i when its price is p_i. We assume that $Q_i(p_i)$ is downward-sloping and has a unique monopoly price $p_i^m = \arg\max (p_i - c_i) Q_i(p_i)$. Per-consumer monopoly profit and consumer surplus from product i are respectively denoted by

$$\pi_i \equiv (p_i^m - c_i) Q_i(p_i^m) \quad \text{and} \quad v_i \equiv \int_{p_i^m}^{\infty} Q_i(p)dp \ . \tag{1}$$

Manufacturers can sell their products directly to consumers, for example via their own retail outlets. However there is also a single intermediary, which can buy products from manufacturers and resell them to consumers. The intermediary has no resale cost, and can stock as many products as it wishes.[7] The intermediary simultaneously makes take-it-or-leave-it offers to each manufacturer whose product it wishes to stock. These offers can be either 'exclusive' (meaning that only the intermediary sells the product to consumers) or 'non-exclusive' (meaning that both the relevant manufacturer and the intermediary sell the product to consumers). In both cases we restrict the intermediary to offer two-part tariffs. Hence if the intermediary wishes to stock product i, it proposes to manufacturer i a wholesale unit price τ_i and a lump-sum fee T_i. The intermediary also informs manufacturers about which products it intends to stock exclusively and non-exclusively.[8] Manufacturers then decide simultaneously whether or not to accept their offer.

Consumers know who sells what, but do not know the contractual arrangement (τ_i, T_i) between the intermediary and manufacturer i. In addition, consumers cannot observe a firm's price(s) or buy its product(s) without incurring a search cost.[9] If the

[7] Even with this assumption of unlimited stocking space, we will see that the intermediary chooses to stock a subset of the products. Limited stocking space will be discussed in Section 5.2.

[8] This assumption aims to capture the idea that in practice negotiations evolve over time, such that manufacturers can (roughly) observe what else the intermediary is going to stock.

[9] As we will see this assumption greatly simplifies the pricing problem in our model and enables us

intermediary stocks a measure m of products, a consumer's total search cost is

$$s \times [n + \mathbf{1}_{\boldsymbol{int}} h(m)] \ , \tag{2}$$

where s is consumer-specific unit search cost, n is the measure of manufacturers searched by the consumer, and $\mathbf{1}_{\boldsymbol{int}}$ is an indicator function which is 1 if and only if the consumer searches the intermediary. We assume that the function $h(m)$ is weakly increasing, consistent for example with the idea that larger stores may take longer to navigate, and may also be located further out of town. When $h(m) < m$ we say that the intermediary generates economies of search, and when $h(m) > m$ we say that it generates diseconomies of search. Consumers are heterogeneous with respect to their 'unit' search cost s, which is distributed in the population according to a cumulative distribution function $F(s)$ with support $(0, \overline{s}]$. The corresponding density function $f(s)$ is everywhere differentiable, strictly positive, and uniformly bounded with $\max_s f(s) < \infty$. Once a consumer has searched a firm, she can recall its offer costlessly.

The timing of the game is as follows. At the first stage, the intermediary simultaneously makes offers to manufacturers whose product it would like to stock. An offer specifies (τ_i, T_i) and whether the intermediary will sell the product exclusively or not. The manufacturers then simultaneously accept or reject. At the second stage, all firms choose a price for each of their products. Both manufacturers and the intermediary are assumed to use linear pricing. At the third stage, consumers observe who sells what and form (rational) expectations about all prices. They then search sequentially among firms using passive beliefs, and then make their purchases.

As we will see, it will be convenient to index products by their per-consumer monopoly profit and consumer surplus as defined in (1) (rather than by their demand curve $Q_i(p_i)$). Therefore let $\Omega \subset \mathbb{R}^2_+$ be a two-dimensional product space (π, v), and suppose it is compact and convex.[10] Let $\underline{v} \geq 0$ and $\overline{v} < \infty$ be the lower and the upper bound of v. Then for each $v \in [\underline{v}, \overline{v}]$, there exist $\underline{\pi}(v) \leq \overline{\pi}(v) < \infty$ such that $\pi \in [\underline{\pi}(v), \overline{\pi}(v)]$. Let $(\Omega, \mathcal{F}, G)$ be a probability measure space where $\mathcal{F}$ is a σ-field which is the set of all measurable subsets of Ω according to measure G. (In particular, $G(\Omega) = 1$.) If a consumer buys a set $A \in \mathcal{F}$ of products at their monopoly prices, she obtains surplus $\int_A v dG$ before taking into account the search cost. When there is no

to study product selection in a tractable way. See a further discussion of this assumption in Section 5.3.

[10] In section 5 we will offer classes of demand functions which can generate this type of product space.

confusion, we also use G to denote the joint distribution function of (π, v), and let g be the corresponding joint density function. We assume that g is differentiable and strictly positive everywhere. To avoid trivial corner solutions and make the study interesting, we also assume that $\overline{v} \leq \overline{s}$.

Our aim is to characterize which products a profit-maximizing intermediary chooses to stock, and investigate whether or not it sells them exclusively. However before doing so, we briefly discuss what would happen if there were no intermediary. In this case, there is an equilibrium in which each manufacturer charges its monopoly price. (Consumers only observe a manufacturer's price after incurring the search cost, so it is (weakly) optimal for each manufacturer to charge the monopoly price.)[11] Therefore recalling equation (1), manufacturer i is searched only by consumers with $s \leq v_i$, such that it earns a total profit $\pi_i F(v_i)$.

3 Exclusive Contracts and No Search Economies

We now consider the case when the intermediary is active. We start with the special case where i) the intermediary can only offer exclusive contracts, and ii) $h(m) = m$ such that the cost of visiting the intermediary is the same as it would have cost to visit the manufacturers whose products it sells (i.e. no economies of scale in search). This relatively simple case will help illustrate the key economic forces that determine which products the intermediary should stock.

3.1 The intermediary's optimal product range

We have the following preliminary result concerning equilibrium contracts and their effect on equilibrium pricing. (All omitted proofs are available in the appendix.)

Lemma 1 *At the first stage, if the intermediary wishes to stock product i it offers $(\tau_i = c_i, T_i = \pi_i F(v_i))$ and the manufacturer accepts. At the second stage, all products are priced at the monopoly level.*

According to Lemma 1 each product is priced at its monopoly level, irrespective of whether it is sold by its original manufacturer or by the intermediary. Consequently we

[11] As is usual in search models, there also exist other equilibria in which consumers do not search (some) manufacturers because they are expected to charge very high prices. We do not consider these (uninteresting) equilibria.

can represent products using the (π, v) space Ω which was introduced earlier. Intuitively, the ability to offer two-part tariffs avoids double marginalization by the intermediary. In addition, because consumers only learn prices after incurring search costs, both manufacturers and the intermediary find it optimal to charge monopoly prices. If the intermediary stocks product i it compensates its manufacturer with a lump-sum payment $\pi_i F(v_i)$, which following earlier arguments is what the manufacturer would receive if it refused the offer and sold directly to consumers.

We now solve for a consumer's decision of whether or not to search the intermediary. Suppose the intermediary sells a positive measure of products $A \in \mathcal{F}$. Firstly, a consumer can cherry-pick from the products not stocked by the intermediary, and therefore will search any product $i \notin A$ if and only if $s \leq v_i$. Secondly though, a consumer *cannot* cherry-pick from amongst the intermediary's products – she must either search all or none of them. Therefore if a consumer visits the intermediary she incurs an additional search cost $s \int_A dG$, but also expects to receive additional utility $\int_A v dG$. Consequently a consumer visits the intermediary if and only if $s \leq k$, where

$$k = \frac{\int_A v dG}{\int_A dG} \tag{3}$$

is the average consumer surplus amongst the products sold at the intermediary.

The intermediary's problem is then

$$\max_{A \in \mathcal{F}} \int_A \pi \left[F(k) - F(v) \right] dG \ , \tag{4}$$

with k defined in (3).[12] In particular the intermediary earns $\pi \left[F(k) - F(v) \right]$ on each product it stocks. This is explained as follows. The intermediary attracts a mass of consumers $F(k)$, and so earns variable profit $\pi F(k)$ on each product it stocks. However from Lemma 1 the intermediary must also compensate a manufacturer with a lump-sum transfer $\pi F(v)$. The following simple observation will play an important role in subsequent analysis: among the products stocked by the intermediary, those with $v < k$ generate a profit while those with $v > k$ generate a loss. Intuitively a product with $v < k$ generates relatively few sales when sold by its manufacturer, since consumers anticipate receiving only a low surplus. When the same product is sold by the intermediary its sales increase, because more consumers search the intermediary (given its higher expected surplus k). The opposite is true for a product with $v > k$.

[12] Note that when $\int_A dG = 0$ the intermediary's profit is zero and it does not matter how we specify k. Some of our later analysis will consider limit cases where the measure of A goes to zero, and in those cases k will be well-defined via L'hopital's rule.

The following lemma is a useful first step in solving the intermediary's problem.

Lemma 2 *The intermediary makes a strictly positive profit. It optimally chooses* $\int_A dG \in (0,1)$ *i.e. it sells a strictly positive measure of product, but not all products.*

The intermediary earns strictly positive profit even though its search technology is no more efficient than that of the manufacturers whose products it resells. To understand why, recall that the intermediary always makes a gain on some products and a loss on others, and that these gains and losses are proportional to a product's per-customer profitability π. Now imagine that the intermediary selects its loss-making products from amongst those with low π, and selects its profit-making products from those with high π. This strategy seeks to minimize losses on the former, and maximize gains on the latter, and so might be expected to generate a net positive profit. In the appendix we show by construction that there is always some set A where this logic is correct.

We now solve explicitly for the set of products stocked by the intermediary. Instead of working directly with areas in Ω, it is more convenient to introduce a stocking policy function $q(\pi, v) \in \{0,1\}$. Then stocking products in a set $A \in \mathcal{F}$ is equivalent to adopting a measurable stocking policy function $q(\pi, v) = 1$ if and only if $(\pi, v) \in A$. The intermediary's problem then becomes

$$\max_{q(\pi,v)\in\{0,1\}} \int_\Omega q(\pi, v)\pi[F(k) - F(v)]dG \ ,$$

while the average consumer surplus k offered by the intermediary solves

$$\int_\Omega q(\pi, v)\,(v-k)\,dG = 0 \ . \tag{5}$$

We then proceed by treating (5) as a constraint and using the Lagrange method.[13] The Lagrange function is

$$\mathcal{L} = \int_\Omega q(\pi, v)[\pi(F(k) - F(v)) + \lambda(v-k)]dG \ , \tag{6}$$

where λ denotes the multiplier associated with the constraint (5). Since the integrand in (6) is linear in q, the optimal stocking policy is as follows:

$$q(\pi, v) = \begin{cases} 1 & \text{if } \pi(F(k) - F(v)) + \lambda(v-k) \geq 0 \\ 0 & \text{otherwise} \end{cases} \ .$$

[13] It can be shown that the optimization problem has a solution, and the legitimacy of the Lagrange method can be proved by the standard technique of calculus of variation. (The details will be added to the appendix soon.)

For given k and λ, we let $I(k, \lambda)$ denote the set of (π, v) for which $q(\pi, v) = 1$. It then remains to determine k and λ. Firstly, at the optimum we must have $F(k) \in (0, 1)$. To see why, note that Lemma 2 implies that $I(k, \lambda)$ must have a strictly positive measure, and therefore by the definition of k it must be true that $k \in (\underline{v}, \overline{v})$. Moreover by assumption $[\underline{v}, \overline{v}] \subseteq [0, \bar{s}]$ and so it follows that $F(k) \in (0, 1)$. Since k is interior, we can take the first-order condition of (6) with respect to k, and obtain

$$\int_{I(k,\lambda)} (f(k)\pi - \lambda) dG = 0 \ , \tag{7}$$

whereupon we observe that $\lambda > 0$. (Intuitively λ captures the impact on profit of a small decrease in k, and k can be decreased either by removing some loss-making products with high v, or adding some profitable products with low v.) Secondly, we have the original constraint (5), which we can rewrite as

$$\int_{I(k,\lambda)} (v - k) dG = 0 \ . \tag{8}$$

We therefore have a system of two equations (7) and (8) in two unknowns, and we prove in the appendix that this system has a solution.[14]

Proposition 1 *The intermediary optimally stocks products with*

$$v \leq k \quad \textit{and} \quad \pi \geq \lambda \frac{k - v}{F(k) - F(v)} \ ,$$

as well as products with

$$v \geq k \quad \textit{and} \quad \pi \leq \lambda \frac{k - v}{F(k) - F(v)} \ ,$$

where $k \in (\underline{v}, \overline{v})$ and $\lambda > 0$ jointly solve equations (7) and (8).

According to Proposition 1 the intermediary's optimal product selection consists of two "negatively correlated" regions in (π, v) space. We can divide Ω space into four quadrants, using a vertical locus $v = k$ and a horizontal locus $\pi = \lambda (k - v) / [F(k) - F(v)]$.[15] Firstly, the intermediary stocks products in the bottom-right quadrant: since products

[14] In numerical examples we find that the system has a unique solution with $k \in (\underline{v}, \overline{v})$, though we have been unable to formally prove uniqueness. If the system does ever has multiple solutions, the solution that generates the highest profit is the optimal one.

[15] The latter locus is everywhere continuous in v, including around the point $v = k$ where it equals $\lambda / f(k)$.

with $v > k$ make a loss, the intermediary chooses those with the lowest possible π. These products act as 'loss-leaders' – their high v attracts consumers to search the intermediary, and these consumers go on to buy other (profitable) products as well. Secondly, the intermediary also stocks products in the top-left quadrant: since products with $v < k$ make a profit, the intermediary chooses those with the highest possible π. On the other hand, products in the bottom-left quadrant are not stocked: their low π means that they would generate little direct profit, whilst their low v would dissuade some consumers from searching and thus reduce profit on other items. A similar intuition explains why products in the top-right quadrant are not stocked either.

Consider a uniform example where $F(s) = s$ and $G(\pi, v) = \pi v$ with support $\Omega = [0,1]^2$. One can check that in the optimal solution $k = \lambda = \frac{1}{2}$ and the intermediary stocks products with $(\pi - \frac{1}{2})(\frac{1}{2} - v) \geq 0$. The solid curves in Figure 1(a) below depict the optimal product range in this example. In this example the intermediary makes profit $\frac{1}{32}$ and improves industry profit by 12.5%.

Finally, it is also interesting to consider how the shape of the search cost distribution $F(s)$ influences the optimal product range. Observe that by Jensen's inequality, a concave $F(s)$ means that more consumers search the intermediary, than on average would search the manufacturers whose products it resells. Related, a concave $F(s)$ also implies that a larger intermediary tends to attract relatively more consumers, and so suggests that its optimal product range should be larger. This is partly borne out by Proposition 1, where we see that the horizontal locus $\pi = \lambda(k - v) / [F(k) - F(v)]$ increases in v when $F(s)$ is concave, such that ceteris paribus the top-left and bottom-right quadrants are large. (The opposite conclusions hold when $F(s)$ is convex.)

3.2 Comparison with the socially optimal solution

We now turn to the optimal product selection by a social planner who aims to maximize the sum of industry profit and consumer surplus. Notice that with no economies of scale in search (i.e., $h(m) = m$), consumers always prefer cherry-picking from manufacturers directly. In that case they buy a product if and only if it provides a positive net surplus $v - s > 0$. While in the case with the intermediary, they are forced to buy some low-v products with a negative net surplus in order to get other high-v products with a positive net surplus. This observation suggests that in the profit-maximizing solution, the intermediary might be "too big" relative to the socially optimal size.

Suppose the intermediary stocks a positive measure of products $A \in \mathcal{F}$. Then a consumer will visit the intermediary only if $s < k$, where k is the average v of the products in A as defined in (3) before. The welfare from these products is therefore $\int_A \int_0^k (\pi + v - s)dF(s)dG$. For those products sold directly by their manufacturers, a consumer will buy product (π, v) only if $s < v$. So the welfare from those products is $\int_{\Omega \setminus A} \int_0^v (\pi + v - s)dF(s)dG$. Maximizing the total welfare which is the sum of these two components is equivalent to

$$\max_A \int_A \int_v^k (\pi + v - s)dF(s)dG = \int_A \left(\pi[F(k) - F(v)] + \int_v^k (v - s)dF(s) \right) dG \ .$$

The objective function consists of two parts: The first part is the impact on industry profit of selling products in A through the intermediary. It is exactly the objective we have tried to maximize in the intermediary's problem. The second part is the impact on consumer surplus. This is always negative, reflecting the fact that selling through the intermediary harms consumers.

By a similar logic as in Lemma 2, we can show that the social planner will stock a positive measure of products, but not all products. This implies that although consumers prefer having no intermediary, the positive effect of selling through the intermediary on profit dominates at least for some product selection.

Following the Lagrange procedure as in the profit-maximizing problem, we can derive the optimal product selection $I(k, \lambda)$:

$$v \leq k \text{ and } \pi \geq \frac{\lambda(k - v) + \int_v^k (s - v)dF(s)}{F(k) - F(v)} \ ,$$

or

$$v \geq k \text{ and } \pi \leq \frac{\lambda(k - v) + \int_v^k (s - v)dF(s)}{F(k) - F(v)} \ ,$$

where k and λ solve the system of

$$\int_{I(k,\lambda)} (v - k)dG = 0 \quad \text{and} \quad \int_{I(k,\lambda)} (f(k)\pi - \lambda)dG = 0 \ .$$

Notice that the two equations for k and λ are exactly the same as those in the profit-maximizing problem. But $I(k, \lambda)$ here takes a different form, so the solution of (k, λ) can differ from that in the profit-maximizing problem.[16] Let (k_P, λ_P) and (k_W, λ_W) be the solution in the profit-maximizing and the welfare-maximizing problems respectively. Let I_P and I_W be the corresponding optimal product selections.

[16] By a similar argument as in Proposition 1, we can also show in the welfare-maximizing problem that the system of equations for k and λ must have a solution with $k \in (\underline{v}, \overline{v})$.

Since $\int_v^k (s-v)dF(s) > 0$ for any $v \neq k$, we can see that if $(k_P, \lambda_P) = (k_W, \lambda_W)$, then $I_W \subset I_P$. In other words, if the profit-maximizing and the welfare-maximizing problem generate the same k and λ, the intermediary is "too big" relative to the socially optimal size. In our running uniform example with $F(s) = s$ and $G(\pi, v) = \pi v$, this is indeed the case: $k = \lambda = \frac{1}{2}$ in both problems. Figure 1(a) below compares the product selections in this example (where the dashed curves are for the socially optimal solution).[17]

Nevertheless, (k, λ) can differ in these two problems such that some products are stocked by the intermediary but not by the social planner, and vice versa.[18] Consider the example with $F(s) = s$ and $G(\pi, v) = \pi^2 v$. One can numerically check that $k_P \approx 0.462$ and $\lambda_P \approx 0.652$, and $k_W \approx 0.424$ and $\lambda_W \approx 0.642$. Figure 1(b) compares the product selections in this example (where the dashed curves are for the socially optimal solution). For instance, when $v = 0.45$, the intermediary will only stock products with $\pi \gtrapprox 0.652$ while the social planner will only stock products with $\pi \lessapprox 0.630$.

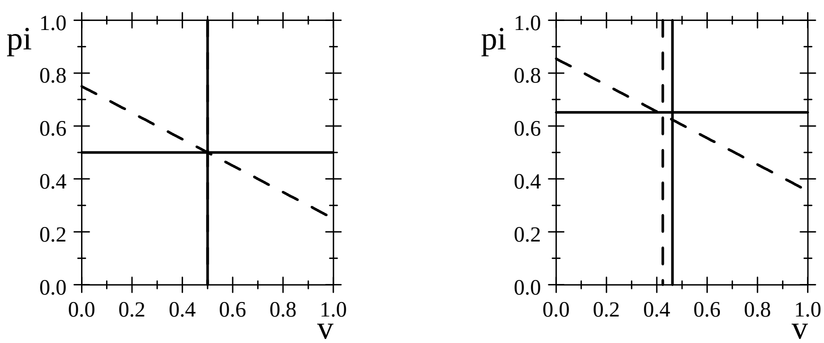

(a): $F(s) = s$ and $G(\pi, v) = \pi v$

(b): $F(s) = s$ and $G(\pi, v) = \pi^2 v$

Figure 1: Product range comparison with $h(m) = m$ and exclusive contracts

[17] Although the intermediary is too big relative to the socially optimal size in this example, it still improves total welfare relative to the case without the intermediary. Total welfare without the intermediary is about 0.177, and it goes up to about 0.188 in the profit-maximizing solution. In the socially optimal solution, it further goes up to about 0.191.

[18] Notice that $I_P \subseteq I_W$ is impossible. This is because given the constructions of $I(k, \lambda)$ in both problems, $I_P \subseteq I_W$ can happen only if $k_P = k_W = k$ and the boundary functions for π in the two cases have the same value at $v = k$. One can check that the latter further requires $\lambda_P = \lambda_W$. But we have known that when $(k_P, \lambda_P) = (k_W, \lambda_W)$, we must have $I_W \subset I_P$.

4 The General Case

We now consider a more general search cost function $h(m)s$, and also allow the intermediary to offer both exclusive and non-exclusive contracts. We make the following assumptions about $h(m)$: (i) $h(0) \geq 0$ and $h(1) \leq 1$, and (ii) $0 \leq h'(m) \leq 1$. Notice that we allow for the possibility $h(0) > 0$, such that a portion of the search cost is independent of the number of products stocked by the intermediary. More broadly we also allow for the possibility that $h(m) > m$ for some values of m, in which case the intermediary generates diseconomies of search.

4.1 The intermediary's optimal product range

The following lemma establishes that even in this more general environment, the (π, v) space remains a valid way to represent products.

Lemma 3 *There is an equilibrium in which i) the intermediary offers $\tau_i = c_i$ and a lump sum fee T_i to any manufacturer whose product it wishes to stock, such that the manufacturer's total payoff is $\pi_i F(v_i)$, and ii) all sellers of product i charge the monopoly price p_i^m.*

Allowing for non-exclusivity does not qualitatively change equilibrium contracts or downstream pricing.[19] In particular even when a product is stocked non-exclusively by the intermediary, both it and the manufacturer charge the monopoly price. Intuitively, the intermediary proposes a bilaterally-efficient two-part tariff with $\tau_i = c_i$ in order to avoid double marginalization. Since both the manufacturer and the intermediary have the same marginal input price, a Diamond Paradox argument then implies that both charge the monopoly price. In particular, the presence of strictly positive search frictions implies that consumers visit at most one of the two retailers, and do not compare prices, in such a way that neither has a incentive to try and undercut the other.

As before let $q(\pi, v) \in \{0, 1\}$ be the intermediary's stocking policy function. We now need an additional exclusivity policy function $\theta(\pi, v) \in \{0, 1\}$, which indicates whether product (π, v), conditional on being stocked, will be stocked exclusively by the intermediary. Henceforth whenever there is no confusion we will suppress the arguments in q and θ.

[19] Lemma 3 only proves existence of such an equilibrium. In fact we can prove a stronger result: all equilibria of the contracting game (if there is more than one) result in monopoly pricing.

Let us first investigate a consumer's optimal search rule. We already know from Lemma 3 that when a product is stocked by both its manufacturer and the intermediary, each charges the same (monopoly) price. Therefore no consumer visits both the intermediary and a manufacturer with a non-exclusive product. The payoff to a consumer of type s from searching the intermediary is then

$$u^1(s,q,\theta) = \int qvdG - h\left(\int qdG\right)s + \int_{v>s}(1-q)(v-s)\,dG\ , \tag{9}$$

where the final term is surplus obtained from products not available at the intermediary. The payoff to a consumer of type s from not searching the intermediary is

$$u^0(s,q,\theta) = \int_{v>s}(1-q\theta)(v-s)\,dG\ , \tag{10}$$

because she is able to purchase all products except those stocked exclusively by the intermediary. Observe that as the intermediary stocks more products exclusively i.e. as the function $\theta(\pi,v)$ takes value 1 for more products, visiting the intermediary becomes relatively more attractive. This suggests that even though the intermediary can now offer non-exclusive contracts, it may still use (more expensive) exclusive contracts because exclusively stocked products are more effective in attracting consumers.

To ease the exposition, we introduce the following tie-break rule: consumers visit the intermediary only if doing so strictly increases their payoff. Comparing equations (9) and (10) we then obtain the following result.

Lemma 4 *Consumers search the intermediary if and only if $s<k$, where*
(i) $k=0$ (nobody searches the intermediary) if $\int q\theta=0$ and $\int qdG \le h\left(\int qdG\right)$.
(ii) $k>\bar{s}$ (everybody searches the intermediary) if $\int qvdG > h\left(\int qdG\right)\bar{s}$.
(iii) $k\in(0,\bar{s}]$ otherwise and is the solution to

$$k = \frac{\int_{v<k} qvdG + \int_{v>k} q\theta vdG}{h(\int qdG) - \int_{v>k} q(1-\theta)dG}\ . \tag{11}$$

According to part (i) of the lemma, no consumer visits the intermediary when all its products are non-exclusive and it generates diseconomies of search. This is because consumers can acquire all of the intermediary's products elsewhere at lower cost. On the other hand, part (ii) of the lemma shows that all consumers visit the intermediary when it generates sufficiently strong economies of search. Finally, part (iii) of the lemma shows that in other cases consumers follow a cut-off strategy, and search the intermediary provided their search cost is sufficiently low. Intuitively the advantage of shopping at the intermediary is that it stocks some products exclusively and/or has a better search technology, while the disadvantage is that consumers may buy some products with low v which ordinarily would not interest them. However consumers with low s would like to buy most products anyway, and so the latter disadvantage is small.

Given the consumer search rule, the intermediary's profit, when it chooses a stocking policy (q, θ), is

$$\Pi(q, \theta) = \int_{v<k} q\pi[F(k) - F(v)]dG + \int_{v>k} q\theta\pi[F(k) - F(v)]dG \ , \tag{12}$$

where k is given in Lemma 4. For a product with $v < k$, the profit from it is independent of whether it is exclusively or non-exclusively stocked. This is because even under non-exclusivity the manufacturer makes zero sales, since consumers with $s < k$ buy from the intermediary, and consumers with $s \geq k$ find it too costly to search the manufacturer. Hence in both cases the intermediary earns revenue $\pi F(k)$, and must pay the manufacturer the full profit $\pi F(v)$ that it would earn if it rejected the offer. This explains the first term. The second term in (12) is profit earned on exclusive products where $v > k$. This takes the same form as in the previous section, and the explanation is the same. Finally, and most interestingly, products with $v > k$ which are stocked non-exclusively do not appear in equation (12), because they generate zero profit for the intermediary. The reason is that consumers with $s < k$ buy the product from the intermediary, whilst consumers with $s \in (k, v)$ buy it from the manufacturer. Hence the intermediary only needs to compensate the manufacturer by $\pi F(k)$, which is exactly the revenue that it earns from such a product. Although these products generate no direct revenue for the intermediary, we will show that the intermediary may stock them in order to influence consumers' search behavior.

The following lemma gives a first qualitative description of what the optimal product range looks like:

Lemma 5 *(i) The intermediary will always stock a positive measure of products and earn a strictly positive profit if $h(m) = m$ for all $m \in [0, 1]$ or if $h(m) < m$ for some $m \in (0, 1]$.*
(ii) When the intermediary optimally stocks a positive measure of products and consumers adopt a search rule with threshold k, (a) all products with $v > k$ (if any) must be stocked, and for each $v > k$ there exists $\pi^+(v)$ such that product (π, v) is stocked exclusively if and only if $\pi \leq \pi^+(v)$; (b) among the products with $v < k$ (if any), for each $v < k$ there exists $\pi^-(v)$ such that product (π, v) is stocked if and only if $\pi \geq \pi^-(v)$.

An important difference relative to the benchmark case in Section 3 is that now the intermediary will optimally stock *all* products with $v > k$. Intuitively, suppose to the contrary that some positive measure set of products B with $v > k$ are not stocked. Since consumers who search the intermediary have $s < k$, they intend to buy all products in B from their respective manufacturers. Now suppose that the intermediary deviates and stocks all products in B non-exclusively (as we saw earlier, these products will earn zero profit). Consumers who search B now buy these products direct from the intermediary, and since $h'(m) \leq 1$ this saves them search costs. Hence by stocking products in B the intermediary attracts more consumers, so can increase its profit.[20] Nevertheless similar to the benchmark model, products with $v > k$ that are stocked exclusively are loss-leaders, and so are chosen to have the lowest π possible in order to minimize that loss. Moreover, and again similar to the earlier benchmark model, products with $v < k$ make positive profit, and so are chosen to have the highest π possible in order to maximize these profits.

We now characterize the details of the optimal product range. The intermediary's problem is to maximize (12), where k is given in Lemma 4. It is more convenient to introduce another parameter $m = \int qdG$, i.e., the measure of products stocked by the intermediary. In this general case, corner solutions with $m \in \{0, 1\}$ or $k \in \{0, \bar{s}\}$ can arise. In the following, we will focus on the case where the intermediary must make a strictly positive profit in the optimal solution (so $m > 0$ and $k > 0$), and not all consumers visit it (so $k < \bar{s}$). Result (i) in Lemma 5 has provided simple sufficient conditions for the former, and according to Lemma 4 a simple sufficient condition for the latter is $\int qvdG/h(\int qdG) < \bar{s}$ for any q, which is equivalent to $\max_x \int_x^{\bar{v}} vdG/h(\int_x^{\bar{v}} dG) < \bar{s}$.[21]

Now the intermediary's problem is to maximize (12) subject to (11). It is more convenient to treat $m = \int qdG$ as another constraint. Then we set up the Lagrange function:

$$\mathcal{L} = \int_{v<k} q\pi \left[F(k) - F(v)\right] dG + \int_{v>k} q\theta\pi \left[F(k) - F(v)\right] dG$$

$$+\lambda \left(\int_{v<k} qvdG + \int_{v>k} q[\theta v + (1-\theta)k]dG - h(m)k \right) + \mu \left(\int_{v<k} qdG + \int_{v>k} qdG - m \right),$$

[20] Note that in the knife-edge case where $h'(m) = 1$ the intermediary is indifferent between stocking products in B, since doing so does not change the search cost of marginal consumers, and so has no effect on its profit.

[21] More precisely, $\int_x^{\bar{v}} vdG = \int_x^{\bar{v}} \int_{\underline{\pi}(v)}^{\overline{\pi}(v)} vg(\pi, v)d\pi dv$. The equivalence result is because for any stocking policy q, $\exists\, x \in [\underline{v}, \bar{v}]$ such that $\int qdG = \int_x^{\bar{v}} dG$, and in the same time $\int qvdG \leq \int_x^{\bar{v}} vdG$ since the average v improves when the product mass is allocated to the products with the highest possible v's.

where λ is the Lagrange multiplier associated with the constraint (11), and μ is the multiplier associated with the constraint $m = \int qdG$. If $m = 1$, then we must have $q = 1$ everywhere and then the second constraint become redundant and the μ term disappears. The following proposition reports the optimal solution:

Proposition 2 *Suppose $m > 0$ and $k \in (0, \overline{s})$ in the optimal solution (which is true if the conditions in Lemma 5 hold and $\max_x \int_x^{\overline{v}} vdG/h(\int_x^{\overline{v}} dG) < \overline{s}$). Then the optimal product selection features either*
(i) $m < 1$, and among the products with $v < k$, only those with

$$\pi \geq \lambda \frac{h'(m)k - v}{F(k) - F(v)} \tag{13}$$

are stocked (and it does not matter whether they are stocked exclusively or non-exclusively), and among the products with $v \geq k$, those with

$$\pi \leq \lambda \frac{k - v}{F(k) - F(v)} \tag{14}$$

are stocked exclusively and the others are stocked non-exclusively. In this case, the parameters k, λ, and m solve the following system of equations:

$$k = \frac{\int_{v<k} qvdG + \int_{v>k} q\theta vdG}{h(m) - \int_{v>k} q(1-\theta)dG} \ , \tag{15}$$

$$\lambda = f(k) \frac{\int_{v<k} q\pi dG + \int_{v>k} q\theta\pi dG}{h(m) - \int_{v>k} q(1-\theta)dG}, \tag{16}$$

$$m = \int qdG \ ; \tag{17}$$

or
(ii) $m = 1$ (i.e., all products are stocked), and among the products with $v \geq k$, those with

$$\pi \leq \lambda \frac{k - v}{F(k) - F(v)}$$

are stocked exclusively, and it does not matter whether to stock the products with $v < k$ exclusively or non-exclusively. In this case, λ and k solve (15) and (16) with $q = 1$ and $m = 1$.

This characterization is consistent with the qualitative description of the optimal product range in Lemma 5. We have discussed, after Lemma 5, the major difference between this general case and the simple case in Section 3. That is, when non-exclusive contracts are available, the intermediary will stock the products in the top-right corner non-exclusively (which were excluded when only exclusive contracts are available). A subtler difference is that when $h'(m) < 1$, $\frac{h'(m)k-v}{F(k)-F(v)} \to -\infty$ when $v \to k^-$. This implies that for those products with v close to but smaller than k, they will always be stocked regardless of π. For instance, if Ω has a flat lower boundary $\underline{\pi} = 0$ in the π dimension, then all products with $v \in (h'(m)k, k)$ will be stocked.

Notice that for the stocked products with $v < k$, the exclusivity arrangement does not matter. This is because even if such a product is available for purchase in its manufacturer, the consumers who do not visit the intermediary (i.e., those with $s > k$) will not purchase given $v < s$. This is the same as if the product is stocked exclusively by the intermediary. We can tie-break this indifference by introducing some small-demand consumers who never visit the intermediary. In that case, the intermediary will strictly prefer to stock the products with $v < k$ non-exclusively in order to reduce the compensation to the manufacturers. (A formal proof is available upon request.)

In general it appears hard to find necessary and sufficient conditions for an interior solution with $m < 1$. The following result reports a sufficient condition for that.[22]

Proposition 3 *The intermediary will not stock all products (i.e. $m < 1$) if $\underline{\pi} = \underline{v} = 0$, $[0, \epsilon]^2 \subset \Omega$ for a sufficiently small $\epsilon > 0$, $h(1) < 1$, $h'(1) > 0$ and $\int v dG/h(1) < \overline{s}$.*

Consider our running uniform example with $G(\pi, v) = \pi v$ and $F(s) = s$. Consider first $h(m) = m$. One can solve that in this example $k = \lambda = \frac{1}{2}$ and $m = 0.75$, and products with $v \geq \frac{1}{2}$ and $\pi \leq \frac{1}{2}$ are stocked exclusively. The solid curves in Figure 2(a) below depict the optimal product range. Compared to the case with exclusively contracts, the only difference is that now the products in the top-right corner $[0.5, 1]^2$ are stocked non-exclusively (though the intermediary only has a weak incentive to do so given there are no economies of search). Consider then $h(m) = \alpha + \beta m$ with $\alpha, \beta \geq 0$ and $\frac{1}{2}(1-\beta)^2 < \alpha < 1 - \beta$. The imposed restrictions for α and β ensure $m \in (0, 1)$ and $k \in (0, 1)$ in the optimal solution. For instance, when $\alpha = 0.05$ and $\beta = 0.9$, one can solve that $k = \lambda \approx 0.515$ and $m \approx 0.823$. The solid curves Figure 2(b) below depict the optimal product range in this example.

[22] A simple sufficient condition for $m = 1$ is $\int v dG/h(1) > \overline{s}$. Under this condition, Lemma 4 implies that all consumers will visit the intermediary and buy if it stocks all products. This generates the highest possible industry profit and so also the highest possible intermediary profit. But this case does not satisfy $k < \overline{s}$ required in Proposition 2.

4.2 Comparison with the socially optimal solution

We now suppose that a social planner can choose the intermediary's stocking policy (q,θ). The consumer search rule is the same as in Lemma 4, and we again use $m = \int qdG$ to denote the measure of products stocked by the intermediary. Total welfare can be written as

$$TW(q,\theta) = \int \pi F(v)\,dG + \Pi(q,\theta) + \int_0^k u^1(s,q,\theta)\,dF(s) + \int_k^{\bar{s}} u^0(s,q,\theta)\,dF(s)\ . \quad (18)$$

The first term is the profits of manufacturers, who always earn $\pi F(v)$ regardless of whether they sell their product exclusively or whether they allow the intermediary to stock it (in which case the intermediary compensates them for lost profit). The second one is the intermediary's profit, which we defined earlier in equation (12). The third one is the surplus of consumers with $s < k$ who search the intermediary, where $u^1(s,q,\theta)$ was defined earlier in equation (9). The forth one is the surplus of consumers with $s \geq k$ who choose not to visit the intermediary, where $u^0(s,q,\theta)$ again was defined earlier in equation (10).

Notice that the consumers with $s \geq k$ are always made (weakly) worse off by the the presence of intermediary, because it restricts access to products with high v (if stocked exclusively) which ordinarily they would like to buy from the manufacturer. On the other hand, whether the presence of the intermediary benefits the consumers with $s < k$ depends on the strength of search economies generated by visiting the intermediary. Unlike in the benchmark case in Section 3, it is now possible that consumers overall prefer having the intermediary.

The social planner wishes to choose a stocking policy (q,θ) in order to maximize $TW(q,\theta)$. We have the following preliminary characterization of the social optimum:

Lemma 6 *(i) The social optimum always has a strictly positive measure of products if* $h(m) = m$ *for all* $m \in [0,1]$ *or if* $h(m) < m$ *for some* $m \in (0,1]$.
(ii) When the optimum has $m > 0$ *and consumers adopt a search rule with threshold* k, *(a) all products with* $v > k$ *(if any) must be stocked, and for each* $v > k$ *there exists* $w^+(v)$ *such that product* (π, v) *is stocked exclusively if and only if* $\pi \leq w^+(v)$*; (b) among the products with* $v < k$ *(if any), for each* $v < k$ *there exists* $w^-(v)$ *such that product* (π, v) *is stocked if and only if* $\pi \geq w^-(v)$.

Qualitatively the socially optimal stocking policy is like the one adopted by the intermediary, and the intuition is closely related to that of Lemma 5. For example it is again optimal for all products with $v > k$ to be stocked. Intuitively if some products with $v > k$ are not currently stocked by the intermediary, they can be added non-exclusively. This has no effect on the payoff of consumers who do not search the intermediary, but is weakly beneficial for those who do, given that $h'(m) \leq 1$ such that they save on search costs by buying a larger basket of products from the intermediary. Since stocking these products attracts weakly more consumers to search the intermediary, it is also beneficial for the intermediary's profit.

We now solve explicitly for the social planner's optimum. It turns out to be more convenient to work with the following alternative expression for total welfare[23]

$$\begin{aligned} TW(q,\theta) &= \int (1-q) \int_0^v (\pi + v - s) dF(s) dG + \int q(\pi+v)F(k)dG \\ &\quad - h(m) \int_0^k s dF(s) + \int_{v>k} q(1-\theta) \int_k^v (\pi+v-s)dF(s)dG \ . \end{aligned} \tag{19}$$

Maximizing this is the same as maximizing $TW(q,\theta) - TW(0,0)$, where $TW(0,0) = \int \int_0^v (\pi + v - s)\, dF(s)\, dG$ is the total welfare when there is no intermediary. To simplify the exposition, we look directly for an interior solution with $k \in (0, \overline{s})$ and $m \in (0,1)$. The two constraints are then the consumer search constraint in equation (11), and the stocking constraint $m = \int q dG$. As before let λ and μ be the respective multipliers associated with these two constraints. After some algebraic manipulations, we can write the Lagrange as follows:

$$\begin{aligned} \mathcal{L} &= \int_{v<k} q \left[(\pi+v)\left[F(k) - F(v)\right] + \int_0^v s dF(s) + \lambda v + \mu \right] dG \\ &\quad + \int_{v>k} q \left\{ \theta \left[(\pi+v)\left[F(k)-F(v)\right] + \int_k^v s dF(s) + \lambda(v-k) \right] + \int_0^k s dF(s) + \lambda k + \mu \right\} dG \\ &\quad - h(m) \int_0^k s dF(s) - \lambda k h(m) - \mu m \ . \end{aligned}$$

A useful preliminary observation is that the first-order condition with respect to m yields $\mu = -h'(m) \left[\lambda k + \int_0^k s dF(s) \right]$, which we can use to substitute out μ from the Lagrangean. Then proceeding as in the intermediary's problem, we first find that a

[23] This can be obtained by substituting the expressions for $\Pi(q,\theta)$, $u^1(s,q,\theta)$ and $u^0(s,q,\theta)$ into equation (18) and then rearranging. The first term is the surplus generated by the products not stocked by the intermediary. The second and third terms are the surplus generated by the products stocked in the intermediary and purchased by consumers with $s < k$ who visit the intermediary. The final term is the surplus generated by the products non-exclusively stocked in the intermediary and purchased by consumers with $s > k$ directly from their manufacturers.

product with $v < k$ is stocked if and only if

$$\pi > \frac{\lambda(kh'(m) - v) + \int_v^k (s-v)dF(s) + (h'(m)-1)\int_0^k s dF(s)}{F(k) - F(v)} \ . \tag{20}$$

As before, for products with $v < k$ exclusivity θ is unimportant. For products with $v > k$, they are stocked exclusively if

$$\pi < \frac{\lambda(k-v) + \int_v^k (s-v)dF(s)}{F(k) - F(v)} \ , \tag{21}$$

and are otherwise stocked non-exclusively. This is consistent with Lemma 6, and reflects the fact that exclusive products with $v > k$ are needed to make it attractive for consumers to search the intermediary, but at the same time should be chosen with a low π so as to minimize the losses incurred on them. Finally, it is straightforward to verify that the first-order condition with respect to k takes the same form as in the intermediary's problem, so the parameters k, λ and m solve the same system of equations as (15) - (17).

It seems difficult to make comparisons between the socially optimal selection and the intermediary's selection. Nevertheless for a fixed (k, λ) we can make a couple of remarks. Firstly, by comparing (14) and (21) and using the fact $\int_v^k (s-v)dF(s) > 0$ for $v > k$, we can deduce that the intermediary stocks too many products exclusively relative to the socially optimal size. Intuitively when the intermediary considers stocking some products exclusively, it neglects the negative impact it has on consumers with high search costs, who choose not to search it and therefore lose the ability to buy those products. Secondly, due to economies of search it is more ambiguous whether the intermediary under- or over-stocks products with low v. Nevertheless by comparing the numerators of (13) and (20), we find that if $\underline{v}$ is sufficiently close to 0, the intermediary stocks too many products with the lowest values of v. It does this because it does not fully internalize the negative effect this has on consumers who search it, who are forced to incur extra search costs to buy products which ordinarily they would avoid buying.

We conclude this section by returning to our running example with $G(\pi, v) = \pi v$ and $F(s) = s$. Consider first the example of $h(m) = m$. One can check that $k_W = \lambda_W = \frac{1}{2}$ and $m_W \approx 0.6875$ in the socially optimal solution. The dashed curves in Figure 2(a) below describe the socially optimal product range in this example. Compared to the intermediary's solution, k and λ are the same, but the social planner stocks fewer products overall and fewer products exclusively. (The comparison is the same as in the benchmark case with exclusive contracts only except that now the top-right corner is included.) Consider then $h(m) = \alpha + \beta m$. When $\alpha = 0.05$ and $\beta = 0.9$, one can solve $k_W \approx 0.499$, $\lambda_W \approx 0.510$ and $m_W \approx 0.795$. The dashed curves in Figure 2(b) below describe the socially optimal product range in this example. Again, the social planner stocks fewer products overall and fewer products exclusively than the intermediary (though the social planner's product set is not a subset of the intermediary's).

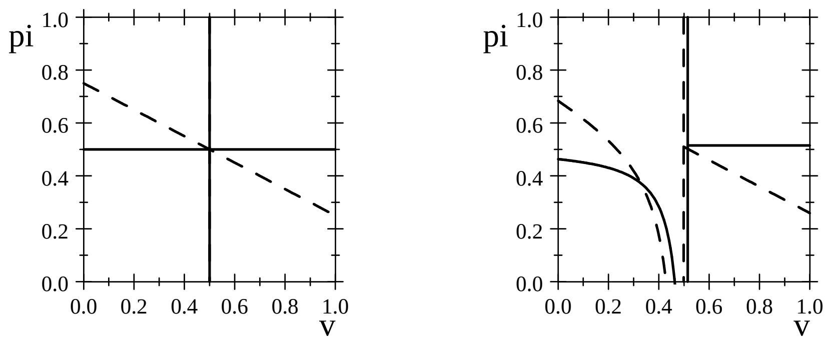

(a): $h(m) = m$ (b): $h(m) = 0.05 + 0.9m$

Figure 2: Product range comparison in the general case

5 Discussion

5.1 Construction and interpretation of (π, v) space

We have thus far represented the intermediary's optimal stocking decision using (π, v) space. In this section we provide two classes of demand functions which can generate this (π, v) space, and then discuss how a product's demand curvature or demand elasticity affects where it is located in (π, v) space.

Demand curvature: Suppose that product i has a constant-curvature demand function:

$$Q_i(p_i) = a_i \left(1 - \frac{1-\sigma_i}{2-\sigma_i}(p_i - \mu_i)\right)^{\frac{1}{1-\sigma_i}}, \tag{22}$$

where $a_i > 0$ denotes the scale of demand, $\mu_i \geq 0$ is the minimum allowed price, and $\sigma_i \in (-\infty, 2)$ is the curvature of the demand curve.[24] When $\sigma_i < 1$, the support of price is $[\mu_i, \mu_i + \frac{2-\sigma_i}{1-\sigma_i}]$; when $1 \leq \sigma_i < 2$, the support of price is $[\mu_i, \infty)$. This is a rich class which includes very concave 'rectangular-shaped' demand when σ_i is sufficiently

[24]The curvature of demand function $Q(p)$ is defined as $Q''(p)\, Q(p) / [Q'(p)]^2$.

negative, linear demand when $\sigma_i = 0$, exponential demand when $\sigma_i = 1$, and very convex demand close to the original point when σ_i is close to 2.[25]

When unit cost is $c_i \geq \mu_i$, monopoly price is $p_i^m = 1 + \mu_i \frac{1-\sigma_i}{2-\sigma_i} + \frac{c_i}{2-\sigma_i}$. Then monopoly profit and consumer surplus are

$$\pi_i = a_i \left(\frac{1}{2-\sigma_i} \right)^{\frac{1}{1-\sigma_i}} \left(1 + (\mu_i - c_i) \frac{1-\sigma_i}{2-\sigma_i} \right)^{\frac{2-\sigma_i}{1-\sigma_i}} ,$$

and

$$v_i = a_i \left(\frac{1}{2-\sigma_i} \right)^{\frac{2-\sigma_i}{1-\sigma_i}} \left(1 + (\mu_i - c_i) \frac{1-\sigma_i}{2-\sigma_i} \right)^{\frac{2-\sigma_i}{1-\sigma_i}} .$$

Notice that both π_i and v_i are increasing in the demand scale parameter a_i, and $\pi_i/v_i = 2 - \sigma_i$. For each fixed σ_i, we can generate a ray from the original point by varying a_i. By varying σ_i, we can change the slope of the ray to cover the whole quadrant $\mathbb{R}_+^2$. (Intuitively, when σ_i is lower demand is more concave and 'rectangular-shaped', such that the firm can appropriate more of the available surplus and so $\frac{\pi_i}{v_i}$ becomes higher.) Consequently, in this example, the high-v and low-π 'loss leaders' which are stocked exclusively are the products with a relatively large and convex demand (i.e. those with relatively high a_i and σ_i). While the profitable low-v and high-π products are those with a relatively large and concave demand (i.e. those with relatively high a_i and low σ_i).[26]

Demand elasticity: Suppose that product i's demand function is

$$Q_i(p_i) = a_i \left(1 - p_i^{\sigma_i}\right)$$

for $p_i \in [0, 1]$, where $a_i > 0$ is the scale parameter as before, and $\sigma_i > 0$ is now an elasticity parameter. For any $p_i \in (0, 1)$, the demand elasticity is

$$\frac{\sigma_i p_i^{\sigma_i}}{1 - p_i^{\sigma_i}} ,$$

and it decreases in σ_i. When σ_i is close to 0, the demand is very convex and price sensitive; when σ_i is large, the demand is very concave and price insensitive.

To get analytical solutions, let us assume $c_i = 0$. The monopoly price is then $p_i^m = (\frac{1}{1+\sigma_i})^{\frac{1}{\sigma_i}}$, and monopoly profit and consumer surplus are

$$\pi_i = \frac{a_i \sigma_i}{1+\sigma_i} \left(\frac{1}{1+\sigma_i} \right)^{\frac{1}{\sigma_i}} ,$$

[25] It also includes constant elasticity demand when $\sigma_i = \frac{2+\mu_i}{1+\mu_i} \in (1, 2)$.

[26] Anderson and Renault (2003) and Weyl and Fabinger (2013) show that this insight extends beyond the class of demands discussed here. In particular they show that in general demands that are 'more concave' are associated with a higher π_i/v_i ratio.

and

$$v_i = \frac{a_i \sigma_i}{1+\sigma_i}\left(1 - \frac{2+\sigma_i}{1+\sigma_i}\left(\frac{1}{1+\sigma_i}\right)^{\frac{1}{\sigma_i}}\right) .$$

Both π_i and v_i increase in a_i, and π_i/v_i increases in σ_i and so decreases in elasticity.[27] Intuitively when demand is more elastic the monopoly price is lower, such that profit is lower and consumer surplus is higher. Hence viewed in light of this class of demands, the intermediary tends to use the products with a relatively large and elastic demand as loss-leaders, and earns profit from the products with a relatively large and inelastic demand.

Discussion: Suppose π and v are determined by two product-specific parameters, say, (a, σ) like in the second example or the first example with $\mu = c = 0$. Then generically there is a one-to-one correspondence between (a, σ) and (π, v), and so each point in the (π, v) space represents a single product. (Notice, however, that even if products are uniformly distributed in the (a, σ) space, they can be non-uniformly distributed in the (π, v) space.) Nevertheless, if π and v are determined by more than two parameters like in the first example with product specific μ and c, then generically each point in the (π, v) space represents multiple different products (but with a measure of zero). In this case, the stocking policy function $q(\pi, v)$ can take a continuous value in $[0, 1]$ with the interpretation that $q(\pi, v)$ fraction of the products at point (π, v) are stocked. Similarly, $\theta(\pi, v)$ can also take a continuous value in $[0, 1]$ with the interpretation that $\theta(\pi, v)$ fraction of the stocked products at point (π, v) are exclusive products. This does not affect our analysis because the objective functions in all our optimization problems are linear in q and θ, and so we always have bang-bang solutions.

5.2 Limited stocking space

We have assumed so far that the intermediary can stock an unlimited number of products. In many cases they have to respect a stocking space constraint. Our analysis can carry over to that case with small modifications. To illustrate the idea in a simple case, let us consider the model in Section 3 with $h(m) = m$ and exclusive contracts. Suppose now the intermediary cannot stock more than a measure $\bar{m}$ of products, and suppose $\bar{m}$ is less than the measure of stocked products in the unconstrained optimization problem.

[27] One disadvantage of this example is that $\pi_i/v_i > 1$ for any $\sigma_i > 0$, so it can only genearate half of the quadrant $\mathbb{R}^2_+$.

The intermediary will then exclude the products which contribute the least to profit. Intuitively, the least profitable products are those with v close to k: for the products with v slightly below k, their demand is only expanded a little by being sold through the intermediary; for the products with v slightly above k, they contribute little in attracting more consumers to visit.

Formally, the intermediary's problem is now

$$\max_{q(\pi,v)\in\{0,1\}} \int q(\pi,v)\,\pi(F(k)-F(v))dG$$

subject to $\int q(\pi,v)(v-k)dG = 0$ and $\bar{m} - \int q(\pi,v)\,dG = 0$. This problem can be solved by the usual Lagrange procedure. Let λ and μ the respective Lagrange multipliers associated with these two constraints. The Lagrange function is then

$$\mathcal{L} = \int q(\pi,v)\left[\pi(F(k)-F(v)) + \lambda(v-k) - \mu\right] dG + \mu\bar{m}\ .$$

Therefore, for $v < k$, $q(\pi,v) = 1$ if and only if $\pi \geq \frac{\lambda(k-v)+\mu}{F(k)-F(v)}$; for $v > k$, $q(\pi,v) = 1$ if and only if $\pi \leq \frac{\lambda(k-v)+\mu}{F(k)-F(v)}$. Notice that μ must be positive since it is the shadow value of the stocking space. Then we can deduce that $\lim_{v\to k^-} \frac{\lambda(k-v)+\mu}{F(k)-F(v)} = \infty$ and $\lim_{v\to k^+} \frac{\lambda(k-v)+\mu}{F(k)-F(v)} = -\infty$. That is, all products with v sufficiently close to k should be excluded regardless of their π. This confirms the intuition above. The first-order condition with respect to k is the same as in the unconstrained problem. From this and the two constraints, we can solve k, λ and μ. The following graph depicts the optimal product range in the running uniform example with $\bar{m} = 0.4$, where the dashed curves are for the unconstrained problem. (Recall that the measure of stocked products in the unconstrained problem is 0.5.)

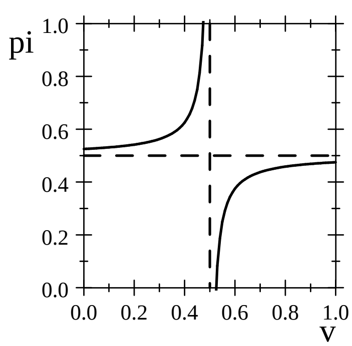

Figure 3: Optimal product range with limited stocking space

5.3 Observable prices

An important assumption we have made is that consumers do not observe the retail prices before they visit the intermediary or manufacturers. This has greatly simplified our analysis: due to this assumption, each product will be charged at its monopoly price regardless of where it is sold, and then we are able to study the intermediary's product selection problem in a tractable way based on the (π, v) product space. This assumption makes sense when prices fluctuate frequently over time (e.g. due to frequent change in cost conditions).

Relaxing this assumption makes our model intractable even in the simple setting with exclusive contracts.[28] Suppose prices are observable before consumers visit a firm. If manufacturer i sells its product exclusively to consumers at price p_i, a consumer will visit it and buy $Q_i(p_i)$ units if $s < S_i(p_i) \equiv \int_{p_i}^{\infty} Q_i(p)dp$. It will thus charge the optimal price $p_i^* = \arg\max_{p_i} F(S_i(p_i))(p_i - c_i)Q_i(p_i)$. An important difference here is that now the distribution of search costs affects pricing. Denote $\pi_i^* \equiv (p_i^* - c_i)Q_i(p_i^*)$ and $v_i^* = S_i(p_i^*)$. For a given search cost distribution $F(s)$, we can construct a product space based on (π^*, v^*). When the intermediary stocks a positive measure of products exclusively, its pricing problem is even more complicated because of a complementarity effect: when the intermediary reduces the prices of a subset of its products, more consumers will visit and so this increases the demand for the other products as well. This makes the intermediary's pricing problem no longer separable across products. Such a pricing problem is hard to deal with in general. However, if the intermediary were to charge the same prices as the manufacturers, then our analysis can apply to the (π^*, v^*) product space. This offers a lower bound of the real optimum. The intermediary can actually do better by adjusting its retail prices, and this gives the intermediary an additional incentive to be active in the market.

6 Conclusion

Product range is an important choice for retailers who intermediate between manufacturers and consumers. This paper has developed a framework for studying the optimal product range choice of a multiproduct intermediary when consumers need a basket

[28] In the general case with non-exclusive contracts, another complication is that there will be potential price competition between the intermediary and a manufacturer if its product is non-exclusively stocked by the intermediary.

of products and face shopping frictions (both of which are natural features of retail markets). We have shown that (i) whenever the intermediary can use exclusive contracts, it exists profitably even if it does not improve search efficiency for consumers; (ii) the intermediary uses exclusively stocked products that consumers value highly as loss leaders and makes profit from non-exclusively stocked products that are relatively cheap to buy from manufacturers; (iii) the intermediary tends to be too big and stock too many products exclusively compared to the socially optimal size.

This paper clearly has a few limitations which we hope to address in future work. First, we have intentionally simplified the pricing decisions of manufacturers and the intermediary by assuming two-part-tariff contracts and unobservability of prices before consumers search. This has enabled us to study the optimal product range and exclusivity in a tractable way. Second, we have focused on a monopoly intermediary. Thus we have not studied how competition among intermediaries might shape their product range choice, which is certainly an important dimension in reality. Third, we have assumed that each product has only one manufacturer, so we only studied the breadth of an intermediary's product assortment. It will be fruitful to consider multiple manufacturers for each product which supply differentiated versions. We will then be able to study both the breadth and depth of an intermediary's product range choice.

Appendix

Proof of Lemma 1. Clearly given the search friction, any manufacturer i selling its own product optimally charges p_i^m. Now define $p_i^*(\tau) = \arg\max (p - \tau) Q_i(p)$, and consider the intermediary. If the intermediary wishes to stock product i and offers a wholesale price τ_i, it will optimally charge $p_i^*(\tau_i)$ and hence must pay the manufacturer $T_i = \pi_i F(v_i) - (\tau_i - c_i) Q_i(p_i^*(\tau_i))$. Hence the intermediary's profit would be

$$[p_i^*(\tau_i) - c_i] Q_i(p_i^*(\tau_i)) - \pi_i F(v_i) \ ,$$

but this is maximized at $\tau_i = c_i$ such that the intermediary charges $p_i^*(c_i) = p_i^m$. ■

Proof of Lemma 2. (i) To prove that the optimum has $\int_A dG > 0$, it suffices to construct such an A which generates a strictly positive profit. Consider two interior points in Ω: $(\pi_1, \tilde{v})$ and $(\pi_2, \tilde{v})$ with $\pi_1 > \pi_2$. Let $A_1 = [\pi_1 - \delta, \pi_1] \times [\tilde{v} - \epsilon, \tilde{v}]$ and

$A_2 = [\pi_2, \pi_2 + \Delta(v)] \times [\tilde{v}, \tilde{v} + \epsilon]$, where $\Delta(v)$ is uniquely defined for each $v \in [\tilde{v}, \tilde{v} + \epsilon]$ by

$$\int_{\pi_1-\delta}^{\pi_1} g(\pi, 2\tilde{v} - v)\, d\pi = \int_{\pi_2}^{\pi_2+\Delta(v)} g(\pi, v)\, d\pi \ . \tag{23}$$

Convexity of Ω implies that we have $A_1, A_2 \subset \Omega$ for sufficiently small $\epsilon \geq 0$ and $\delta > 0$. Notice that $\Delta(v)$ is constructed in such a way that for each v in A_2, the mass of products stocked is the same as that of the 'mirror' valuation $2\tilde{v} - v$ in A_1. This implies that the average v of the products in $A_1 \cup A_2$ is always $\tilde{v}$, and so a consumer will visit the intermediary, when it stocks $A = A_1 \cup A_2$, if and only if $s < \tilde{v}$.

Fix a sufficiently small δ such that $\pi_1 - \delta > \pi_2 + \Delta(v)$ for all $v \in [\tilde{v}, \tilde{v} + \epsilon]$. The intermediary's profit from stocking $A = A_1 \cup A_2$ is

$$\Pi(\epsilon) = \int_{\tilde{v}-\epsilon}^{\tilde{v}} \int_{\pi_1-\delta}^{\pi_1} \pi [F(\tilde{v}) - F(v)]\, dG + \int_{\tilde{v}}^{\tilde{v}+\epsilon} \int_{\pi_2}^{\pi_2+\Delta(v)} \pi [F(\tilde{v}) - F(v)]\, dG \ .$$

Straightforward calculations reveal that $\Pi(0) = \Pi'(0) = 0$. However,

$$\begin{aligned} \Pi''(0) &= f(\tilde{v}) \left[\int_{\pi_1-\delta}^{\pi_1} \pi g(\pi, \tilde{v})\, d\pi - \int_{\pi_2}^{\pi_2+\Delta(\tilde{v})} \pi g(\pi, \tilde{v})\, d\pi \right] \\ &> f(\tilde{v}) [(\pi_1 - \delta) - (\pi_2 + \Delta(\tilde{v}))] \int_{\pi_1-\delta}^{\pi_1} g(\pi, \tilde{v})\, d\pi > 0 \ , \end{aligned}$$

where the first inequality used (23) evaluated at $v = \tilde{v}$. Therefore, $\Pi(\epsilon) > 0$ for ϵ in a neighborhood of 0.

(ii) Let $\hat{v} = \int v dG$. Consider $B_1 = [\pi_1 - \delta, \pi_1] \times [\hat{v}, \hat{v} + \epsilon]$ and $B_2 = [\pi_2, \pi_2 + \Delta(v)] \times [\hat{v} - \epsilon, \hat{v}]$, where $\pi_1 > \pi_2$, and where $\Delta(v)$ is uniquely defined for each $v \in [\hat{v} - \epsilon, \hat{v}]$ by

$$\int_{\pi_1-\delta}^{\pi_1} g(\pi, 2\hat{v} - v)\, d\pi = \int_{\pi_2}^{\pi_2+\Delta(v)} g(\pi, v)\, d\pi \ . \tag{24}$$

Convexity of Ω implies that $B_1, B_2 \subset \Omega$ for sufficiently small $\epsilon \geq 0$ and $\delta > 0$. Similarly as above, the average v of the products in $B_1 \cup B_2$ is always $\tilde{v}$, and so the average v in $A = \Omega \backslash (B_1 \cup B_2)$ is $\hat{v}$ as well. Then a consumer will visit the intermediary, when it stocks $A = \Omega \backslash (B_1 \cup B_2)$, if and only if $s < \hat{v}$.

Fix a sufficiently small δ such that $\pi_1 - \delta > \pi_2 + \Delta(v)$ for all $v \in [\hat{v} - \epsilon, \hat{v}]$. The intermediary's profit from stocking $A = \Omega \backslash (B_1 \cup B_2)$ is

$$\hat{\Pi}(\epsilon) = \hat{\Pi} - \int_{\hat{v}}^{\hat{v}+\epsilon} \int_{\pi_1-\delta}^{\pi_1} \pi [F(\hat{v}) - F(v)]\, dG - \int_{\hat{v}-\epsilon}^{\hat{v}} \int_{\pi_2}^{\pi_2+\Delta(v)} \pi [F(\hat{v}) - F(v)]\, dG \ ,$$

where $\hat{\Pi} = \hat{\Pi}(0)$ is the profit from stocking Ω. Simple calculations reveal that $\hat{\Pi}'(0) = 0$. However, similar as in (i),

$$\hat{\Pi}''(0) = f(\hat{v}) \left[\int_{\pi_1 - \delta}^{\pi_1} \pi g(\pi, \hat{v})\, d\pi - \int_{\pi_2}^{\pi_2 + \Delta(\hat{v})} \pi g(\pi, \hat{v})\, d\pi \right] > 0$$

by using (24) evaluated at $v = \hat{v}$. Therefore, $\hat{\Pi}(\epsilon) > \hat{\Pi}$ for ϵ in a neighborhood of 0. ■

Proof of Proposition 1. It remains to prove that (7) and (8) have a solution with $k \in (\underline{v}, \overline{v})$. Let $\phi(v;k) \equiv \frac{k-v}{F(k)-F(v)}$.

We first claim that for any $k \in (\underline{v}, \overline{v})$, (8) has a unique solution

$$\lambda(k) \in \left(\frac{\underline{\pi}}{\max_v \phi(v;k)}, \frac{\overline{\pi}}{\min_v \phi(v;k)} \right)$$

and $\lambda'(k) \in (0, \infty)$. The proof is as follows. The left-hand side of (8) is strictly negative when $\lambda \max_v \phi(v;k) \leq \underline{\pi}$, because then $v \leq k$ for all products in $I(k,\lambda)$ and $v < k$ for a strictly positive measure of them. The left-hand side of (8) is strictly positive when $\lambda \min_v \phi(v;k) \geq \overline{\pi}$ and the reasoning is the same. The left-hand side of (8) is also strictly increasing in λ in the above range, since as λ increases the top-left region in $I(k,\lambda)$ with $v - k < 0$ shrinks while the bottom-right region with $v - k > 0$ expands. Uniqueness of $\lambda(k)$ then follows. Define $\lambda(\underline{v}) = \lim_{k \to \underline{v}} \lambda(k)$ and $\lambda(\overline{v}) = \lim_{k \to \overline{v}} \lambda(k)$. We must have $\lambda(\underline{v})\phi(v;\underline{v}) \leq \underline{\pi}$ and $\lambda(\overline{v})\phi(v;\overline{v}) \geq \overline{\pi}$ for any v (or except for a zero-measure set). Notice also that the left-hand side of (8) is $\mathcal{C}^1$ in (λ, k), so the implicit function theorem implies that $\lambda(k)$ is differentiable. $\lambda'(k) \in (0, \infty)$ can be verified by direct computation.

Now consider (7) with λ replaced by $\lambda(k)$:

$$\int_{I(k,\lambda(k))} (f(k)\pi - \lambda(k))dG = 0 \ . \tag{25}$$

We show that it has a solution $k \in (\underline{v}, \overline{v})$. Consider the following differentiable function of k:

$$\Phi(k) = \int_{I(k,\lambda(k))} \left[\pi(F(k) - F(v)) + \lambda(k)(v-k)\right] dG \ . \tag{26}$$

When $k = \underline{v}$ or $\overline{v}$, $I(k, \lambda(k))$ is an empty set and so $\Phi(\underline{v}) = \Phi(\overline{v}) = 0$. According to the construction of $I(k,\lambda)$ and the definition of $\lambda(k)$, $\Phi(k) > 0$ for $k \in (\underline{v}, \overline{v})$. Therefore by the mean-value theorem $\Phi'(k) = 0$ must have a solution in $(\underline{v}, \overline{v})$. On the other hand, one can verify that $\Phi'(k)$ equals the left-hand side of (25) by using the definition of $\lambda(k)$ and the construction of $I(k,\lambda)$. Then (25) must have a solution $k \in (\underline{v}, \overline{v})$. ■

Proof of Lemma 3. First, consider products sold exclusively by the manufacturer: following arguments from earlier proofs, it is clear the intermediary offers $\tau_i = c_i$ and charges p_i^m. Similarly, it is clear that on products where manufacturer i has exclusive access, it optimally charges p_i^m.

Second, consider contracts and pricing on a non-exclusive product i. Denote by $p_i^*(\tau)$ the monopoly price on product i when marginal cost is τ, and for simplicity assume it is unique. Note that $p_i^*(\tau)$ increases in τ for any well-behaved demand function. In the putative equilibrium where consumers expect the intermediary and manufacturer i to both charge p_i^m, note that no consumer searches both, and let x^I and x^{M_i} denote the measures of consumers who search the intermediary and manufacturer i respectively. (Note it is possible that $x^I = 0$ for example.)

The proof proceeds as follows. We solve for equilibrium pricing given any τ_i, and given that consumers expect both firms to charge p_i^m. We then examine the intermediary's continuation profit following any offer τ_i, and show it to be weakly higher when $\tau_i = c_i$.

(a) Suppose they agree on $\tau_i \leq c_i$ and some T_i. We first prove there is an equilibrium in which manufacturer i charges p_i^m and the intermediary charges $p_i^*(\tau_i)$. Recall that since consumers expect both firms to charge p_i^m, they expect to search at most one of manufacturer i and the intermediary. Moreover because $p_i^*(\tau_i) \leq p_i^m$, in this putative pricing equilibrium they still visit at most one of them, such that manufacturer i earns $T_i + (\tau_i - c_i)\, x^I Q_i(p_i^*(\tau_i)) + x^{M_i}\pi_i$, and the intermediary earns $x^I [p_i^*(\tau_i) - \tau_i]\, Q_i(p_i^*(\tau_i)) - T_i$ from product i. (I) If manufacturer i charges $p < p_i^m$ it cannot induce more consumers to search it, and so it earns $T_i + (\tau_i - c_i)\, x^I Q_i(p_i^*(\tau_i)) + x^{M_i}(p - c_i)\, Q_i(p)$ which is less than what it earns by charging $p = p_i^m$. If manufacturer i charges $p > p_i^m$ it induces some mass $\epsilon \in \left[0, x^{M_i}\right]$ of consumers to search and buy from the intermediary, hence it earns $T_i + (\tau_i - c_i)\left(x^I + \epsilon\right) Q_i(p_i^*(\tau_i)) + \left(x^{M_i} - \epsilon\right)(p - c_i)\, Q_i(p)$. However since $\tau_i \leq c_i$, this profit is less than what it earns by charging $p = p_i^m$. (II) If the intermediary charges $p \neq p_i^*(\tau_i)$ it earns weakly less than $x^I (p - \tau_i)\, Q_i(p) - T_i$, since in case $p > p_i^m$ it may induce some consumers to search and buy from the manufacturer. However $x^I (p - \tau_i)\, Q_i(p) - T_i$ is less than what can be earned by charging $p = p_i^*(\tau_i)$. (III) Given the pricing equilibrium, the intermediary must offer T_i such that manufacturer i earns $\pi_i F(v_i)$ i.e. T_i solves

$$T_i + (\tau_i - c_i)\, x^I Q_i(p_i^*(\tau_i)) + x^{M_i}\pi_i = \pi_i F(v_i) \ . \tag{27}$$

Substituting back in, the intermediary then earns $x^I [p_i^*(\tau_i) - c_i]\, Q_i(p_i^*(\tau_i)) + \pi_i \left[x^{M_i} - F(v_i)\right]$, which is maximized when $p_i^*(\tau_i) = p_i^m$, which requires $\tau_i = c_i$. Hence the intermediary earns a profit $\pi_i \left[x^I + x^{M_i} - F(v_i)\right]$.

(b) Suppose now they agree on $\tau_i > c_i$ and some T_i. We first prove there is an equilibrium in which manufacturer i charges p_i^m and the intermediary charges a price $p_i^{I*} \geq p_i^m$. Notice again that since consumers expect the same price p_i^m at both firms, they expect to visit only one of them. Hence in the putative equilibrium, any consumer who first searches manufacturer i buys; on the other hand, some fraction $\epsilon(p)$ of the consumers who first search the intermediary will subsequently search manufacturer i, where p is the price charged by the intermediary, and $\epsilon(p) = 0$ for all $p \leq p_i^m$, but $\epsilon'(p) \geq 0$ for all $p > p_i^m$ with the added restriction that $\epsilon(p) \leq x^I$. Consequently in the putative equilibrium, the intermediary's price is

$$p_i^{I*} = \arg\max \left[x^I - \epsilon(p)\right](p - \tau_i) Q_i(p) - T_i \ . \tag{28}$$

Manufacturer i charges p_i^m and earns $\left[x^{M_i} + \epsilon\left(p_i^{I*}\right)\right]\pi_i + T_i + \left[x^I - \epsilon\left(p_i^{I*}\right)\right](\tau_i - c_i) Q_i\left(p_i^{I*}\right)$. (I) The manufacturer cannot strictly gain by deviating from p_i^{I*}, given the definition (28). (II) If manufacturer i charges $p < p_i^m$ it cannot induce more consumers to search it, so its profit is clearly lower than if it charges p_i^m. Similarly if manufacturer i charges $p_i \in \left(p_i^m, p_i^{I*}\right)$, all consumers who search it buy it, and clearly its profit is lower than if it charges p_i^m. If manufacturer i charges $p_i > p_i^m, p_i \geq p_i^{I*}$ then it induces $\delta(p)$ consumers to search and buy from the intermediary, where $\delta(p) \in \left[0, \pi^{M_i} + \epsilon\left(p_i^{I*}\right)\right]$ with $\delta'(p) \geq 0$. Hence the manufacturer's profit is

$$T_i + \left[\pi^I - \epsilon\left(p_i^{I*}\right) + \epsilon(p)\right](\tau_i - c_i) Q_i\left(p_i^{I*}\right) + \left[\pi^{M_i} + \epsilon\left(p_i^{I*}\right) - \epsilon(p)\right](p - c_i) Q_i(p) \ , \tag{29}$$

which is less than what is earned when $p = p_i^m$, because $(p - c_i) Q_i(p) < \pi_i$, and because $(\tau_i - c_i) Q_i\left(p_i^{I*}\right) \leq \left(p_i^{I*} - c_i\right) Q_i\left(p_i^{I*}\right) \leq \pi_i$. (III) Given the pricing equilibrium, the intermediary must offer T_i such that manufacturer i earns $\pi_i F(v_i)$ i.e. T_i solves

$$\left[x^{M_i} + \epsilon\left(p_i^{I*}\right)\right]\pi_i + T_i + \left[x^I - \epsilon\left(p_i^{I*}\right)\right](\tau_i - c_i) Q_i\left(p_i^{I*}\right) = \pi_i F(v_i) \ . \tag{30}$$

Substituting back in, the intermediary then earns $\left[x^I - \epsilon\left(p_i^{I*}\right)\right]\left(p_i^{I*} - c_i\right) Q_i\left(p_i^{I*}\right) + \left[x^{M_i} + \epsilon\left(p_i^{I*}\right) - F(v_i)\right]\pi_i$, which is weakly less than what we found in part (a) when the intermediary offers $\tau_i = c_i$.

(c) Summarizing then, the intermediary's continuation payoff is weakly larger when it offers $\tau_i = c_i$. In the resulting equilibrium the intermediary and manufacturer i both charge p_i^m. ■

Proof of Lemma 4. The difference in payoff between (9) and (10) is

$$\Delta(s) = \int qvdG - h(\int qdG)s - \int_{v>s} q(1-\theta)(v-s)\,dG\ . \tag{31}$$

Notice that $\Delta(0) \geq 0$, and $\Delta(s)$ is weakly concave because

$$\Delta'(s) = -h(\int qdG) + \int_{v>s} q(1-\theta)dG \tag{32}$$

is weakly decreasing in s. (i) No consumer visits the intermediary (i.e. $k = 0$) if and only if $\Delta(s) \leq 0$ for all $s > 0$. A necessary and sufficient condition for this is $\Delta(0) = 0$ and $\Delta'(0) \leq 0$, which is equivalent to the conditions stated in the lemma. (ii) All consumers visit the intermediary (i.e. $k > \bar{s}$) if and only if $\Delta(s) > 0$ for all $s > 0$. A necessary and sufficient condition for this is $\Delta(\bar{s}) > 0$, which simplifies to the condition in the lemma. (iii) Finally in all other cases, $\Delta(s) > 0$ for s in a neighborhood of 0, and $\Delta(\bar{s}) \leq 0$, so given that $\Delta(s)$ is weakly concave consumers use a cut-off strategy. Consumers strictly prefer visiting the intermediary if they have $s < k$, where k solves $\Delta(s) = 0$. (11) is just a rewriting of $\Delta(s) = 0$. ■

Proof of Lemma 5. (i) When $h(m) = m$ for all $m \in [0,1]$, we already know from the previous section that the intermediary can make a strictly positive profit by stocking some products exclusively. Now consider the case of $h(m) < m$ for some $m \in (0,1]$.[29] We show that the intermediary can now makes a strictly positive profit by stocking some products non-exclusively. Consider a product set $A \subset \Omega$ such that $\int_A dG = m$ and $\int_{A\cap\{v<a\}} dG > 0$ for any $a > \underline{v}$. Such a set A always exists (e.g. when A is convex and $\min_{v\in A} v = \underline{v}$). Suppose now the intermediary stocks all products in A non-exclusively. Then $\Delta(0) = 0$ and

$$\Delta'(s) = -h(m) + m > 0$$

for all $s \in [0, \underline{v}]$. This implies $k > \underline{v}$. From (12), it is ready to see that the intermediary's profit is $\int_{A\cap\{v<k\}} \pi[F(k) - F(v)]dG > 0$.

(ii-a) Suppose $k < \bar{v}$ so that there are products with $v > k$. Suppose in contrast that in the optimal solution (q, θ), $q = 0$ for a strictly positive measure of products with $v > k$. Denote this set of products by B. Consider a new stocking policy $\left(\tilde{q}, \tilde{\theta}\right)$ such that

$$\tilde{q}(\pi, v) = \begin{cases} 1 & \text{if } (\pi, v) \in B \\ q(\pi, v) & \text{otherwise} \end{cases}$$

[29] In this case, it is possible that $h(0) > 0$. Then the approach in Lemma 2 does not apply because $k \to 0$ when the measure of stocked products goes to 0. That is why we adopt a different approach.

and $\tilde{q}\tilde{\theta} = q\theta$. Let $\tilde{k}$ be the new consumer search threshold associated with $\left(\tilde{q}, \tilde{\theta}\right)$. We aim to show that this new stocking policy is more profitable than (q, θ) and so a contradiction arises. We can see from (12) that this is true if $\tilde{k} \geq k$, or equivalently if $\tilde{\Delta}(k) \geq \Delta(k)$, where $\tilde{\Delta}(\cdot)$ is (31) associated with the new stocking policy. Using the construction of $\left(\tilde{q}, \tilde{\theta}\right)$, one can check that

$$\begin{aligned}
\tilde{\Delta}(k) - \Delta(k) &= \int_B (1-q)vdG - [h(\int \tilde{q}dG) - h(\int qdG)]k - \int_B (1-q)(v-k)dG \\
&= \left(\int_B (1-q)dG - \ [h(\int \tilde{q}dG) - h(\int qdG)]\right) k \ .
\end{aligned}$$

Since $\int \tilde{q}dG - \int qdG = \int_B (1-q)dG$ and $h'(m) \leq 1$ for all m, we have $\tilde{\Delta}(k) - \Delta(k) \geq 0$. Therefore, in the optimal solution we must have $q(\pi, v) = 1$ for all $v > k$.

We now prove the second part in result (a). Suppose in contrast that in the optimal solution (q, θ), there is a strictly positive measure of $v > k$ such that for each of these v, there exist $\pi' > \pi''$ such that $\theta(\pi', v) = 1$ and $\theta(\pi'', v) = 0$ (i.e., some high-π products are stocked exclusively while some low-π products are stocked non-exclusively). Denote this set of v by V. Now fix the stocking policy for all products with $v < k$, but for those with $v > k$ define a new exclusivity policy

$$\tilde{\theta}(\pi, v) = \begin{cases} 1 & \text{if } \pi \leq \tilde{\pi}(v) \\ 0 & \text{if } \pi > \tilde{\pi}(v) \end{cases} ,$$

where $\tilde{\pi}(v)$ is the unique solution to $\int_{\underline{\pi}(v)}^{\tilde{\pi}(v)} g(\pi, v)\, d\pi = \int_{\underline{\pi}(v)}^{\bar{\pi}(v)} \theta(\pi, v)\, g(\pi, v)\, d\pi$. (That is, for each $v > k$, the mass of exclusively stocked products in the original stocking policy is shifted to the products with the lowest possible π.) By construction this does not affect consumer search behavior (so $\tilde{k} = k$) since they only care about v. Then for each $v > k$, we have

$$\int_{\underline{\pi}(v)}^{\bar{\pi}(v)} \pi\, [F(k) - F(v)]\, \theta(\pi, v)\, g(\pi, v)\, d\pi \leq \int_{\underline{\pi}(v)}^{\bar{\pi}(v)} \pi [F(\tilde{k}) - F(v)] \tilde{\theta}(\pi, v)\, g(\pi, v)\, d\pi \ ,$$

with strict inequality for $v \in V$. That is, the intermediary makes less loss from those products with $v > k$ under the new policy. This improves its profit, and so we have a contradiction.

(ii-b) Suppose that $k > \underline{v}$ so that there are products with $v < k$. Suppose in contrast that in the optimal solution (q, θ), there is a strictly positive measure of $v < k$ such that for each of these v, there exists some $\pi' < \pi''$ such that $q(\pi', v) = 1$ and $q(\pi'', v) = 0$.

Denote this set of v by V. Now fix the stocking policy for products with $v > k$, but for products with $v < k$ define a new stocking policy

$$\tilde{q}(\pi, v) = \begin{cases} 1 & \text{if } \pi \geq \tilde{\pi}(v) \\ 0 & \text{if } \pi < \tilde{\pi}(v) \end{cases} ,$$

where $\tilde{\pi}(v)$ is the unique solution to $\int_{\tilde{\pi}(v)}^{\bar{\pi}(v)} g(\pi, v)\, d\pi = \int_{\underline{\pi}(v)}^{\bar{\pi}(v)} q(\pi, v)\, g(\pi, v)\, d\pi$. (That is, for each $v < k$, the mass of stocked products in the original stocking policy is shifted to the products with the highest possible π.) Similarly as before, by construction this does not affect consumer search behavior (so $\tilde{k} = k$). Then for each $v < k$, we have

$$\int_{\underline{\pi}(v)}^{\bar{\pi}(v)} \pi \left[F(k) - F(v)\right] q(\pi, v)\, g(\pi, v)\, d\pi \leq \int_{\underline{\pi}(v)}^{\bar{\pi}(v)} [\pi F(\tilde{k}) - F(v)] \tilde{q}(\pi, v)\, g(\pi, v)\, d\pi ,$$

with strict inequality for $v \in V$. That is, the intermediary makes higher profit from those products with $v < k$ under the new policy. This is a contradiction. ■

Proof of Proposition 2. Let us rewrite the Lagrange function as

$$\mathcal{L} = \int_{v<k} q[\pi(F(k) - F(v)) + \lambda v + \mu] dG$$

$$+ \int_{v>k} q \left\{\theta[\pi(F(k) - F(v)) + \lambda(v - k)] + \lambda k + \mu\right\} dG - \lambda h(m)k - \mu m . \tag{33}$$

We first consider the case where $m < 1$ in the optimal solution. Then the first-order condition with respect to m is $\mu = -\lambda h'(m)k$. We use this to replace μ in our analysis. The first-order condition with respect to k yields (16). The other two equations (15) and (17) are simply the two constraints. Both k and λ are positive. From (33) it is ready to see that for $v < k$, $q = 1$ if and only if

$$\pi(F(k) - F(v)) + \lambda(v + \mu) \geq 0 \Leftrightarrow \pi \geq \lambda \frac{h'(m)k - v}{F(k) - F(v)} ,$$

and the value of θ does not matter. For $v > k$, $\theta = 1$ if and only if

$$\pi(F(k) - F(v)) + \lambda(v - k) \geq 0 \Leftrightarrow \pi \leq \lambda \frac{k - v}{F(k) - F(v)} .$$

When $\theta = 1$, $q = 1$ if and only if

$$\pi(F(k) - F(v)) + \lambda(v - k) + \lambda k + \mu \geq 0 \Leftrightarrow \pi \leq \lambda \frac{h'(m)k - v}{F(k) - F(v)} .$$

But this is automatically satisfied given $h'(m) \leq 1$ and the condition for $\theta = 1$. When $\theta = 0$, $q = 1$ if and only if $\lambda k(1 - h'(m)) \geq 0$. This is also automatically satisfied given $h'(m) \leq 1$ and $\lambda k \geq 0$.

The case with $m = 1$ is simpler. We replace $q = 1$ and $\mu = 0$ in (33). For $v < k$, the value of θ does not matter. For $v > k$, the condition for $\theta = 1$ is the same as before. ■

Proof of Proposition 3. Suppose in contrast that the intermediary stocks all products. We show that for any fixed θ, under the stated conditions excluding products in $B = [\underline{\pi}, \underline{\pi} + \epsilon] \times [\underline{v}, \underline{v} + \epsilon]$ improves profit for a sufficiently small $\epsilon > 0$.

Let $k(\epsilon)$ be the threshold in the search rule when B is excluded. We first show that $k(0) \in (\underline{v}, \overline{s})$. Notice that when all products are stocked,

$$\begin{aligned} \Delta(\underline{v}) &= \int v dG - h(1)\underline{v} - \int (1-\theta)(v - \underline{v}) dG \\ &\geq \int v dG - h(1)\underline{v} - \int (v - \underline{v}) dG \\ &= (1 - h(1))\underline{v} \ , \end{aligned}$$

and

$$\Delta'(\underline{v}) = -h(1) + \int (1-\theta) dG \ .$$

If $\int \theta dG = 0$, then $\Delta(\underline{v}) \geq 0$ (where the equality holds if $\underline{v} = 0$) and $\Delta'(\underline{v}) > 0$ since $h(1) < 1$. If $\int \theta dG > 0$, then $\Delta(\underline{v}) > 0$. In either case, we have $k(0) > \underline{v}$. This also implies $k(0) > 0$, so we must have $k(\epsilon) > \epsilon$ for a sufficiently small ϵ. On the other hand, $\Delta(\overline{s}) = \int v dG - h(1)\overline{s} < 0$ given $\int v dG / h(1) < \overline{s}$. So $k(0) < \overline{s}$.

From (31), we know that if ϵ is sufficiently small such that $k(\epsilon) > \epsilon$, then $k(\epsilon)$ is determined by

$$\int v dG - \int_B v dG = h\left(1 - \int_B dG\right) k(\epsilon) + \int_{v > k(\epsilon)} (1-\theta)\left(v - k(\epsilon)\right) dG \ .$$

(In the last term, we have used $q = 1$ for $v > k(\epsilon) > \epsilon$.) One can verify that $k'(0) = 0$ and

$$k''(0) = 2g(\underline{\pi}, \underline{v}) \frac{h'(1)k(0) - \underline{v}}{h(1) - \int_{v > k(0)} (1-\theta) dG} \ .$$

(Notice that the denominator in $k''(0)$ must be positive because it is equal to $-\Delta'(k(0))$.)

The intermediary's profit, when ϵ is sufficiently small such that $k(\epsilon) > \epsilon$, is

$$\Pi(\epsilon) = \int_{v < k(\epsilon)} \pi [F(k(\epsilon)) - F(v)] dG - \int_{\underline{v}}^{\underline{v}+\varepsilon} \int_{\underline{\pi}}^{\underline{\pi}+\varepsilon} \pi [F(k(\epsilon)) - F(v)] dG$$

$$+ \int_{v > k(\epsilon)} \theta \pi [F(k(\epsilon)) - F(v)] dG \ .$$

One can check that $\Pi'(0) = 0$ and

$$\Pi''(0) = f(k(0))k''(0)\left(\int_{v<k(0)} \pi dG + \int_{v>k(0)} \theta\pi dG\right) - 2g(\underline{\pi}, \underline{v})\underline{\pi}[F(k(0)) - F(\underline{v})] \ .$$

Given $k(0) > \underline{v}$, the bracket term in the first part must be strictly positive. When $\underline{\pi} = 0$, the second term is zero. Then $\Pi''(0) > 0$ for any θ whenever $k''(0) > 0$. This is true if $\underline{v} = 0$ and $h'(1) > 0$. ■

Proof of Lemma 6. (i) The proof for the case $h(m) = m$ for all $m \in [0, 1]$ is very similar to the proof of Lemma 2 and hence is omitted. In the case where $h(m) < m$ for some $m \in (0, 1]$, note that if the intermediary stocks a mass m of products non-exclusively, its profit is strictly higher (by Lemma 5) and consumers are no worse off since they can still buy every product from the manufacturer.

(ii-a) Suppose $k < \bar{v}$ so that there are products with $v > k$. Suppose in contrast that in the optimal solution (q, θ), $q = 0$ for a strictly positive measure of products with $v > k$. Denote this set of products by B. Consider a new stocking policy $\left(\tilde{q}, \tilde{\theta}\right)$ such that

$$\tilde{q}(\pi, v) = \begin{cases} 1 & \text{if } (\pi, v) \in B \\ q(\pi, v) & \text{otherwise} \end{cases}$$

and $\tilde{q}\tilde{\theta} = q\theta$. In the proof of Lemma 5 we showed that the intermediary's profit is weakly higher under $\left(\tilde{q}, \tilde{\theta}\right)$. Observe also that $u^0(s, q, \theta)$ is unchanged, so consumers with $s > k$ are weakly better off under $\left(\tilde{q}, \tilde{\theta}\right)$. Hence it remains to show that $u^1(s, q, \theta)$ is weakly higher under $\left(\tilde{q}, \tilde{\theta}\right)$ for all $s < k$. To prove this, notice that following the logic of the proof of Lemma 5,

$$u^1\left(s, \tilde{q}, \tilde{\theta}\right) - u^1(s, q, \theta) = \left(\int_B (1-q)dG - \ [h(\int \tilde{q}dG) - h(\int qdG)]\right) s \ ,$$

which is weakly positive since $\int \tilde{q}dG - \int qdG = \int_B (1-q)dG$ and $h'(m) \leq 1$ for all m. Since all parties weakly benefit from $\left(\tilde{q}, \tilde{\theta}\right)$ we have a contradiction. We now prove the second part in result (a). Suppose in contrast that in the optimal solution (q, θ), there is a strictly positive measure of $v > k$ such that for each of these v, there exist $\pi' > \pi''$ such that $\theta(\pi', v) = 1$ and $\theta(\pi'', v) = 0$. Denote this set of v by V. Now fix the stocking policy for all products with $v < k$, but for those with $v > k$ define a new exclusivity policy

$$\tilde{\theta}(\pi, v) = \begin{cases} 1 & \text{if } \pi \leq \tilde{\pi}(v) \\ 0 & \text{if } \pi > \tilde{\pi}(v) \end{cases} \ ,$$

where $\tilde{\pi}(v)$ is the unique solution to $\int_{\underline{\pi}(v)}^{\tilde{\pi}(v)} g(\pi, v)\, d\pi = \int_{\underline{\pi}(v)}^{\bar{\pi}(v)} \theta(\pi, v)\, g(\pi, v)\, d\pi$. By construction $u^0(s, .)$ and $u^1(s, .)$ are unchanged, hence consumer surplus is unchanged. However in the proof of Lemma 5 we showed that the intermediary's profit is higher under $\left(\tilde{q}, \tilde{\theta}\right)$, hence we have a contradiction.

(ii-b) Suppose that $k > \underline{v}$ so that there are products with $v < k$. Suppose in contrast that in the optimal solution (q, θ), there is a strictly positive measure of $v < k$ such that for each of these v, there exists some $\pi' < \pi''$ such that $q(\pi', v) = 1$ and $q(\pi'', v) = 0$. Denote this set of v by V. Now fix the stocking policy for products with $v > k$, but for products with $v < k$ define a new stocking policy

$$\tilde{q}(\pi, v) = \begin{cases} 1 & \text{if } \pi \geq \tilde{\pi}(v) \\ 0 & \text{if } \pi < \tilde{\pi}(v) \end{cases},$$

where $\tilde{\pi}(v)$ is the unique solution to $\int_{\tilde{\pi}(v)}^{\bar{\pi}(v)} g(\pi, v)\, d\pi = \int_{\underline{\pi}(v)}^{\bar{\pi}(v)} q(\pi, v)\, g(\pi, v)\, d\pi$. Similarly as before, $u^0(s, .)$ and $u^1(s, .)$ are unchanged hence consumer surplus is unchanged. However in the proof of Lemma 5 we showed that the intermediary's profit is higher under $\left(\tilde{q}, \tilde{\theta}\right)$, hence we have a contradiction. ■

References

ADAMS, W., AND J. YELLEN (1976): "Commodity Bundling and the Burden of Monopoly," *Quarterly Journal of Economics*, 90(3), 475–498.

ANDERSON, S., AND R. RENAULT (2003): "Efficiency and Surplus Bounds in Cournot Competition," *Journal of Economic Theory*, 113(2), 253–264.

ASKER, J., AND H. BAR-ISAAC (2016): "Vertical Information Restraints: Pro- and Anti-Competitive Impacts of Minimum Advertised Price Restrictions," mimeo.

BIGLAISER, G. (1993): "Middlemen as Experts," *RAND Journal of Economics*, 24(2), 212–223.

CHEN, Z., AND P. REY (2012): "Loss Leading as an Exploitative Practice," *American Economic Review*, 102(7), 3462–3482.

DIAMOND, P. (1971): "A Model of Price Adjustment," *Journal of Economic Theory*, 3(2), 156–168.

Gehrig, T. (1993): "Intermediation in Search Markets," *Journal of Economics and Management Strategy*, 2(1), 97–120.

Janssen, M., and S. Shelegia (2015): "Consumer Search and Double Marginalization," *American Economic Review*, 105(6), 1683–1710.

Johnson, J. (2016): "Unplanned Purchases and Retail Competition," *American Economic Review*, forthcoming.

Kaplan, G., G. Menzio, L. Rudanko, and N. Trachter (2016): "Relative Price Dispersion: Evidence and Theory," mimeo.

Lal, R., and C. Matutes (1994): "Retail Pricing and Advertising Strategies," *Journal of Business*, 67(3), 345–370.

Lizzeri, A. (1999): "Information Revelation and Certification Intermediaries," *RAND Journal of Economics*, 30(2), 214–231.

McAfee, P. (1995): "Multiproduct Equilibrium Price Dispersion," *Journal of Economic Theory*, 67(1), 83–105.

McAfee, P., J. McMillan, and M. Whinston (1989): "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values," *Quarterly Journal of Economics*, 104(2), 371–383.

Rayo, L., and I. Segal (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118(5), 949–987.

Rhodes, A. (2015): "Multiproduct Retailing," *Review of Economic Studies*, 82(1), 360–390.

Rubinstein, A., and A. Wolinsky (1987): "Middlemen," *Quarterly Journal of Economics*, 102(3), 581–593.

Spulber, D. (1996): "Market Making by Price-Setting Firms," *Review of Economic Studies*, 63(4), 559–580.

——— (1999): *Market Microstructure: Intermediaries and the Theory of the Firm.* Cambridge University Press.

Stigler, G. (1968): "A Note on Block Booking," in *The Organization of Industry*, ed. by G. Stigler. University of Chicago Press.

Weyl, G., and M. Fabinger (2013): "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition," *Journal of Political Economy*, 121(3), 528–583.

Zhou, J. (2014): "Multiproduct Search and the Joint Search Effect," *American Economic Review*, 104(9), 2918–2939.