



Munich Personal RePEc Archive

Multiproduct Intermediaries and Optimal Product Range

Rhodes, Andrew and Watanabe, Makoto and Zhou, Jidong

Toulouse School of Economics, VU University of Amsterdam, Yale
University

October 2017

Online at <https://mpra.ub.uni-muenchen.de/82136/>

MPRA Paper No. 82136, posted 24 Oct 2017 09:05 UTC

Multiproduct Intermediaries and Optimal Product Range*

Andrew Rhodes

Makoto Watanabe

Toulouse School of Economics

VU University of Amsterdam

Jidong Zhou

Yale School of Management

October 2017

Abstract

This paper develops a framework for studying the optimal product range choice of a multiproduct intermediary when consumers demand multiple products. In the optimal product selection, the intermediary uses exclusively stocked high-value products to increase store traffic, and at the same time earns profit mainly from non-exclusively stocked products which are relatively cheap to buy from upstream suppliers. By doing this the intermediary can earn strictly positive profit, including in situations where it does not improve efficiency in selling products. A linkage between product selection and product demand features such as size and shape is established. It is also shown that relative to the social optimum, the intermediary tends to be too big and stock too many products exclusively.

Keywords: intermediaries, product range, multiproduct demand, search, exclusive contracts

JEL classification: D83, L42, L81

*We are grateful for helpful comments to Mark Armstrong, Heski Bar-Isaac, Joyee Deb, Paul Ellickson, Doh-Shin Jeon, Bruno Jullien, Fei Li, Barry Nalebuff, Martin Obradovits, Patrick Rey, Greg Shaffer, Greg Taylor, Raphael Thomadsen, Glen Weyl, Julian Wright and seminar participants in Bonn, MIT Sloan, NUS, Oxford, Tokyo, TSE, Yale, Zurich as well as the 8th Consumer Search and Switching Workshop (Vienna), Bristol IO Day, EARIE (Maastricht), EEA (Lisbon), ICT conference (Mannheim), SAET (Faro), SICS (Berkeley), and TNIT (Microsoft).

1 Introduction

Many products are traded through intermediaries. A leading example is that of retailers, who buy up products from manufacturers and resell them to consumers. Choosing which products to stock is an important decision for retailers. Consumers are usually interested in buying a basket of products, but find it costly to shop around and so tend to buy from a limited number of retailers whose product ranges closely match their needs. However, at the same time retailers are often constrained in how many products they can sell, for example due to limited stocking space or the fact that stocking too many products can make the in-store shopping experience less pleasant.¹ Moreover products differ in their desirability for consumers and profitability for sellers, and their demands can be interdependent in a multiproduct environment. This further complicates the product selection problem.

Very much related to the stocking problem is the issue of exclusivity. In particular, in order to make themselves more attractive to consumers, retailers are increasingly offering exclusive products that are not available for purchase elsewhere. They do this either by making large investments in their own private brands, or by paying manufacturers for exclusive rights to sell their products. For example in 2009, US departments stores such as Macy's and J.C. Penney generated over 40% of their sales from exclusive products.²

Surprisingly, there are very few papers which study a retailer's optimal choice of product range and product exclusivity. (This contrasts with the voluminous literature on other aspects of a retailer's problem, such as pricing and location choice.) Our paper seeks to fill this gap by developing a multiproduct intermediary framework which can help study these issues in a tractable way. Our paper makes several contributions. Firstly, we provide a new rationale for the existence of intermediaries. In particular we show that when consumers have multiproduct demand, a multiproduct retailer can use exclusivity to enter a market and make strictly positive profit, even if it is no more efficient in selling products than the smaller sellers which it displaces. Secondly and most importantly, we characterize the retailer's optimal product selection. Specifically, we show how all

¹Even large retailers like Walmart face such constraints. Many consumers have to go to smaller stores to buy some hard-to-find products. (See <http://goo.gl/MV6FRi> for some evidence on this.)

²See <http://goo.gl/lfS9QP> for further details. Exclusivity is also common in other parts of the retail market. For instance Home Depot has many exclusive brands such as American Woodmark in cabinets, and Martha Stewart in outdoor furniture and indoor organization. Target is well-known for offering exclusive brands in apparel and home goods. Many high-end fashion stores also sell unique colors or versions of certain labels.

information contained in a product’s demand curve can be represented by a simple two-dimensional sufficient statistic, which in turn determines whether the retailer chooses to stock that product, and whether it does so exclusively. We also show how these choices can be understood in terms of simple properties of the product’s demand curve, such as its size and shape. 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, Section 2 introduces our main model in which a continuum of manufacturers each produces a different product. Consumers view these products as independent and are interested in buying all of them, although different products are allowed to have different demands. A manufacturer’s product can be sold either through a single-product (specialist) store, a multiproduct (generalist) retailer, or both. The single-product retailer can be interpreted as either the manufacturer’s own retail outlet or a completely independent store, and both interpretations give rise to the same results in our model. We choose to frame the paper in terms of the former interpretation, given that with development of e-commerce manufacturers are increasingly selling their products direct to consumers.³ The multiproduct retailer offers to compensate manufacturers in exchange for the right to sell their products, and as part of this can demand exclusive sales rights. We also allow for the possibility that the retailer has a stocking constraint. Consumers are aware of who sells what, but have to pay a cost to learn a firm’s price(s) and buy its product(s). The cost of searching the intermediary is (weakly) increasing in the number of products it stocks, consistent for example with the idea that larger retailers are located further from consumers, or offer a worse instore shopping experience. Consumers also differ in their search costs, such that in equilibrium some end up buying more products than others.

Since the focus of our paper is product range choice, we intentionally simplify sellers’ pricing problems. In particular we assume that the intermediary can offer two-part tariff contracts to manufacturers. We then prove that irrespective of the market structure, each supplier of a given product always charges the usual monopoly price.⁴ This enables us to study product range choice in a tractable way, because it allows us to

³A 2016 Forbes article reports: “The number of manufacturers selling directly to consumers is expected to grow 71% this year to more than 40% of all manufacturers. And over a third of consumers report they bought directly from a brand manufacturer’s website last year”. (See <https://goo.gl/29uWSE>) Along the same lines, a 2017 report by the European Commission states that “many retailers... [now find] themselves competing against their own suppliers.” (See p. 288 of <https://goo.gl/Xg71n2>)

⁴Intuitively, 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.

summarize all information on a product's cost and demand characteristics via a simple two-dimensional statistic (π, v) , where π represents a product's monopoly profit and v represents its monopoly consumer surplus. 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.

In Section 3 we first solve a special case of the model in order to highlight some of the main economic forces at work. In particular we consider the situation in which the intermediary can stock as many products as it likes, but is restricted to using exclusive contracts, and offers no economies of search (i.e. the cost of searching the intermediary is the same as searching all of the manufacturers whose products it sells). We first prove that the intermediary earns strictly positive profit, and so will be active despite not improving search efficiency. We also prove that the intermediary stocks a strict subset of the product space i.e. it voluntarily limits its product range.

We then derive the intermediary's optimal stocking policy in this special case. One might expect the intermediary to sell products with relatively high values of π and v , but this turns out to be incorrect. Instead the intermediary's optimal product range exhibits a form of "negative correlation" in (π, v) space, consisting of two regions in the top-left and the bottom-right. Intuitively a consumer searches the retailer (respectively, an individual manufacturer) if its average (respectively, individual) v exceeds her unit search cost. Consequently demand for a low- v product increases when the intermediary stocks it, and since the manufacturer need only be compensated for its lost sales, these products are profit generators. Nevertheless the intermediary cannot stock too many low- v products otherwise it becomes less attractive to consumers, and therefore only stocks a limited number of the most profitable ones i.e. those with high π . Conversely demand for high- v products falls when the intermediary stocks them, and hence it makes a loss on them. These products are useful in attracting consumers, so the intermediary stocks some of them, but it manages its losses by choosing these products to have relatively low π .

In Section 4 we solve for the intermediary's optimal product range in the general case, where the intermediary can also use non-exclusive contracts and can provide economies of search. The intermediary faces the following tradeoff when deciding whether to stock a product exclusively or non-exclusively. On the one hand consumers are more likely to search it when it has many exclusive products which are not available for purchase elsewhere. On the other hand the intermediary also needs to compensate manufacturers more if it stocks their product exclusively, since manufacturers lose the ability to sell to

consumers who are not interested in shopping at the intermediary. We show that when the stocking space constraint does not bind, the optimal product selection is similar to the special case, except that the intermediary also stocks products in the top-right part of (π, v) space non-exclusively. Intuitively by stocking the latter products non-exclusively, the intermediary attracts more consumers due to economies of search, but still allows consumers who do not visit it to buy those products from their respective manufacturers, thus reducing how much those manufacturers need to be compensated. We also show that as the intermediary's stocking space becomes smaller, the intermediary's optimal product range contains fewer and fewer of these non-exclusive products and eventually again exhibits negative correlation in (π, v) space.

We also solve for a social planner's optimal product range in Section 5 and compare it with what the intermediary chooses. The intermediary distorts consumers' purchases, 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. We show that under weak conditions the social planner finds it optimal to have an intermediary. However the intermediary tends to stock more products than the social planner would like, and often too many of them are stocked exclusively.

Finally in Section 6 we discuss two issues. One issue is how to generate our (π, v) space and how to interpret different points within it. For instance, we argue that products with large and elastic/convex demands tend to have relatively high v and low π and so are stocked exclusively to attract consumers, whereas products with large and inelastic/concave demands tend to have relatively low v and high π and so are used by the intermediary as profit generators. Another issue we consider is upstream competition, which we do by assuming that each product has two manufacturers. Upstream competition does not qualitatively change the optimal product selection, but it reduces the intermediary's cost of buying products from manufacturers and so greatly improves its profit.

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 infor-

mation 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 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), McAfee, McMillan, and Whinston (1989), and Chen and Riordan (2013)). By stocking some products that consumers value highly, the intermediary forces consumers to visit and buy other low-value (but fairly profitable) products as well which consumers would otherwise 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, and so is more related to the question of which products a firm should bundle (however this question is 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 would be like the case where our intermediary could organize and sell non-overlapping products in multiple stores).⁷ This makes the optimization problem in our paper very different from theirs. In addition many other features of our model such as the investigation of non-exclusivity vs exclusivity arrangements, the economies of search, and the stocking space constraint, have no counterparts 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

⁵In 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.

⁶Bundling 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.

⁷Alternatively, if any subset of manufacturers could merge and use the same technology as the intermediary to sell their products together, the problem would then be more similar to Rayo and Segal (2010)'s.

(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 totally different economic questions.

Finally, this paper is related to the research on product assortment planning in operation research and marketing (see, e.g., the survey by Kök et al. (2015)). But that literature focuses on the optimal variety selection for a certain product when consumers have single-product demand. Our paper instead focuses on a retailer's optimal product range choice when consumers have multiproduct demand. We study this issue with explicit upstream manufacturers and consumer shopping frictions, neither of which is considered in the above mentioned literature.⁸

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 . When a consumer buys multiple products, her surplus is additive over these products. We assume that $Q_i(p_i)$ is downward-sloping and well-behaved such that $(p_i - c_i) Q_i(p_i)$ is single-peaked at the monopoly price p_i^m . Per-consumer monopoly profit and consumer surplus from product i are respectively denoted by

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

⁸In this aspect Bronnenberg (2017) is closer to our paper. He studies a free-entry model in a circular city with both manufacturers and retailers. Consumers have preferences for variety but shopping for variety is costly, so retailers can save consumers shopping costs by carrying multiple varieties. Bronnenberg's model is otherwise very different from ours and also focuses on different economic questions. In particular all varieties in his model are symmetric, so there is no meaningful way to study the composition of product selection which, however, is the focus of our paper.

Manufacturers can sell their products directly to consumers, for example via their own retail outlets (see below for an alternative interpretation). In addition there is a single intermediary, which can buy products from manufacturers and resell them to consumers. The intermediary has no resale cost, but can stock at most a measure $\bar{m} \leq 1$ of the products (which we call a “hard” constraint). An individual product can therefore be sold to consumers in one of three different ways: i) only by the manufacturer, ii) only by the intermediary, or iii) by both the intermediary and its manufacturer. We assume that the intermediary has all the bargaining power, and simultaneously makes take-it-or-leave-it offers to each manufacturer whose product it wishes to stock.⁹ These offers can be either ‘exclusive’ (meaning that only the intermediary can sell the product to consumers) or ‘non-exclusive’ (meaning that both the intermediary and the relevant manufacturer can sell the product to consumers). In both cases we suppose that the intermediary offers two-part tariffs, consisting of 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, and whether it intends to stock them exclusively or non-exclusively.¹⁰ Manufacturers then simultaneously decide whether or not to accept their offer.

Consumers know where each product is available, but do not observe (τ_i, T_i) in any contract between a manufacturer and the intermediary. In addition, consumers cannot observe a firm’s price(s) or buy its product(s) without incurring a search cost.¹¹ Consumers differ in terms of their ‘type’ or unit search cost s , which is distributed in the population according to a cumulative distribution function $F(s)$ with support $(0, \bar{s}]$. Suppose that the corresponding density function $f(s)$ is everywhere differentiable, strictly positive, and uniformly bounded with $\max_s f(s) < \infty$. One interpretation is that s is the opportunity cost of spending a unit time in shopping. If a consumer of type s visits a measure n of manufacturers, she incurs an aggregate search cost $n \times s$.¹² If the same consumer also visits the intermediary, and the intermediary stocks a measure m of products, she incurs

⁹Our results do not change qualitatively if instead the intermediary and manufacturer share any profits that are earned from sales of the latter’s product.

¹⁰This assumption aims to capture the idea that in practice negotiations evolve over time, such that manufacturers can (roughly) observe what other products the intermediary stocks.

¹¹Our assumptions here try to capture the idea that a retailer’s product range is usually reasonably steady over time, whilst its prices fluctuate more frequently for example due to cost or demand shocks.

¹²Here we implicitly assume that visiting each manufacturer is equally costly. More generally, the cost of visiting different manufacturers may be different, and our framework can be extended to deal with that case. One possible way to do that is to use (π, v, θ) to characterize each product where θ captures the amount of time needed to visit a manufacturer.

an additional search cost of $h(m) \times s$.¹³ Once a consumer has searched a firm, she can recall its offer costlessly.

We assume that the function $h(m)$ is positive and weakly increasing, reflecting the idea that larger stores may take longer to navigate,¹⁴ and may also be located further out of town. (However notice that the case of $h(m)$ being a constant and so independent of the measure of stocked products is also allowed.) 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. When $h(m)$ is strictly increasing, the intermediary faces another “soft” constraint because as it stocks more products it becomes costlier for consumers to visit it. As we will see later on, when $h(m)$ increases fast enough this will cause the intermediary to voluntarily restrict its size even if its hard stocking space constraint is not binding.

Finally, 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 that sell to consumers choose a retail 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 retail prices. They then search sequentially among firms and make their purchases. We assume that if consumers observe an unexpected price at some firm, they hold passive beliefs about the retail prices they have not yet discovered.

2.1 Preliminary analysis

Our aim is to study which products a profit-maximizing intermediary should choose to stock, and whether or not it should sell them exclusively. However it is instructive to first briefly consider what would happen if there were no intermediary. In this case, the only equilibrium in which each product market is active has each manufacturer selling its product at the monopoly price. This follows from standard arguments concerning the hold-up problem in search models with only one firm (see, e.g., Stiglitz, 1979, and Anderson and Renault, 2006). In particular, since consumers only observe a manufacturer’s

¹³Considering a more general search cost function $h(m, s)$ would make our model less tractable but would not change the main insights.

¹⁴However we do not explicitly model in-store shopping process, since this would require us to analyze not only which products the intermediary stocks but also how it displays them to consumers.

price after incurring the search cost, their decision of whether to search a manufacturer depends only on the expected price there. Once a consumer arrives at the manufacturer, the search cost is already sunk and so the manufacturer optimally charges its monopoly price. Hence consumers should rationally expect monopoly pricing.¹⁵ Therefore recalling the notation introduced in (1), in equilibrium manufacturer i is searched only by consumers with $s \leq v_i$, and so it earns a profit $\pi_i F(v_i)$.

It turns out that we have a similar simple pricing outcome when the intermediary is active. (All omitted proofs can be found in the appendix.)

Lemma 1 *(i) In any equilibrium where each product market is active, each seller of a product charges consumers the relevant monopoly price.*

(ii) If product i is stocked exclusively by the intermediary, the intermediary offers the manufacturer $(\tau_i = c_i, T_i = \pi_i F(v_i))$. If product i is stocked non-exclusively by the intermediary, in terms of studying the optimal product range, it is without loss of generality to focus on the contracting outcome where the intermediary offers $(\tau_i = c_i, T_i)$ to manufacturer i , such that the manufacturer's total payoff is $\pi_i F(v_i)$.

To understand the intuition behind Lemma 1, recall from earlier that a product can be sold in three different ways. Firstly product i may be sold only by its manufacturer. The logic for why the manufacturer charges its monopoly price p_i^m is exactly the same as in the case of no intermediary. The intermediary then earns $\pi_i F(v_i)$, which forms its outside option if it receives an offer from the intermediary. Secondly product i may be sold exclusively by the intermediary. Since consumers do not observe the price before searching, the same hold-up argument implies that if the intermediary faces a wholesale price τ_i , it will charge the corresponding monopoly price $\arg \max (p - \tau_i) Q_i(p)$. Notice that joint profit earned on product i is maximized when the intermediary charges the monopoly price p_i^m , therefore in order to induce this outcome the intermediary proposes $\tau_i = c_i$ i.e. a bilaterally efficient two-part tariff. The intermediary then drives the manufacturer down to its outside option by offering it a lump-sum payment $T_i = \pi_i F(v_i)$. Thirdly product i may be sold by both its manufacturer and the intermediary. The analysis here is more complex. However the main idea is that the intermediary again avoids double-marginalization by proposing a contract with $\tau_i = c_i$, whilst search frictions eliminate price competition between the manufacturer and intermediary. In particular, following

¹⁵As 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, and given no consumers search these high prices can be trivially sustained. We do not consider these uninteresting equilibria in this paper.

Diamond's (1971) paradox if consumers expect both sellers to charge the same price for product i , they will search at most one of them and hence each finds it optimal to charge the monopoly price. The manufacturer is compensated for any sales that it loses in signing the contract by way of a lump-sum transfer.

Given Lemma 1, it is 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. Let $\underline{v} \geq 0$ and $\bar{v} < \infty$ be the lower and the upper bound of v . For each $v \in [\underline{v}, \bar{v}]$, there exist $\underline{\pi}(v) \leq \bar{\pi}(v) < \infty$ such that $\pi \in [\underline{\pi}(v), \bar{\pi}(v)]$. (In section 6.1 we provide examples of demand functions which can generate this type of product space.) Let (Ω, \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, $G(\Omega) = 1$.) When there is no 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. 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. To avoid trivial corner solutions, we also assume that $\bar{v} \leq \bar{s}$.

Discussion. Before we start solving for optimal product range, we discuss some of our modeling assumptions and their implications.

(i) *A continuum of products.* Considering a continuum of products is mainly for analytical convenience. A model with a discrete number of products $\{(\pi_i, v_i)\}_{i=1, \dots, n}$ would yield qualitatively similar insights but be messier to solve because the optimization problem would become a combinatorial one. (See footnote 22 later for the details. The case with only two products is easy to deal with, but is not rich enough to study the optimal product range choice in a meaningful way.)

(ii) *Homogeneous consumer demand.* Consumers are assumed to have demand for all products. In reality a consumer usually only buys a small fraction of the products available in a store, and some consumers want to buy more products than others (or similarly some products are needed more often than others). Our framework can be modified to capture this consumer demand heterogeneity,¹⁶ but it becomes less tractable because two consumers with the same s can have very different search patterns. On the other hand, we will show later that consumers with a lower search cost are more willing to shop around

¹⁶One possible way is to characterize each product by (π, v, α) where $\alpha \in [0, 1]$ is the probability that product (π, v) is needed by a consumer.

and buy more products. In this sense we have already allowed demand heterogeneity: consumers with a low/high search cost can be regarded as high/low-demand consumers.

(iii) *Direct sales from manufacturers to consumers.* Manufacturers are assumed to be able to sell direct to consumers. However nothing changes if instead the manufacturer faces a choice between selling via an independent single-product (specialist) or a multiproduct (generalist) retailer, or both. In particular consider the following modification of our set-up. Suppose that first the intermediary makes offers to manufacturers, who each accept or reject. Manufacturers are unable to sell direct to consumers. However second, if the manufacturer is not forbidden from doing so, it can make an offer to a relevant specialist retailer whose only option is to stock its product. As in Lemma 1, we can prove that equilibrium contracts are such that all sellers charge the relevant monopoly price, and the manufacturer fully extracts the single-product retailer. Consequently each manufacturer's profit is the same as it would earn if it could sell directly to consumers. Hence the intermediary's optimal product selection will be the same as in our main model.

(iv) *Lemma 1 and monopoly pricing.* The monopoly pricing outcome described in Lemma 1 enables us to represent products using (π, v) space, and hence study product range choice in a tractable way. However notice that monopoly pricing is not important per se - what really matters for our analysis is that the retail price of each product remains the same irrespective of where it is sold. Of course in practice prices usually differ across retail outlets, and a large literature already explores this. Our model abstracts from such price dispersion in order to make progress in understanding optimal product choice.

3 A Simple Case

We now turn to study the intermediary's optimal product range choice. We start with a special case where i) the intermediary can only offer exclusive contracts, 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), and iii) there is no stocking space limit (i.e. $\bar{m} = 1$). This relatively simple case is not meant to be realistic, but it helps to illustrate some of the economic forces influencing optimal product selection.

We first solve for a consumer's decision of whether or not to search the intermediary. Suppose the intermediary sells a positive measure of products $A \in \mathcal{F}$. A consumer can cherry-pick from the products not stocked by the intermediary (i.e. she will search any product $i \notin A$ if and only if $s \leq v_i$), but she *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 $\int_A v dG$ since she will buy all products available there. Consequently a consumer visits the intermediary if and only if $s \leq k$, where

$$k = \frac{\int_A v dG}{\int_A dG} \quad (2)$$

is the average consumer surplus amongst the products sold at the intermediary. (Note that the order in which the consumer searches through the intermediary and manufacturers does not matter.)

The intermediary's problem is then

$$\max_{A \in \mathcal{F}} \int_A \pi [F(k) - F(v)] dG, \quad (3)$$

with k defined in (2).¹⁷ In particular the intermediary earns a net profit $\pi [F(k) - F(v)]$ from product (π, v) if it stocks it. This is explained as follows. The intermediary attracts a mass of consumers $F(k)$, and so earns variable profit $\pi F(k)$. However from Lemma 1 the intermediary must also compensate the relevant 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 $v > k$, i.e. its demand is shrunk when sold through the intermediary.¹⁸

The following lemma is a useful first step in characterizing the intermediary's optimal product range.

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

¹⁷Note 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'hospital's rule.

¹⁸Notice that the same will be true for a general $h(m)$ if it increases in m fast enough. But if $h(m)$ is close to be constant and is sufficiently small, then k can be greater than any v in A . As we will see in Section 4, in the latter case the characterization of the optimal product selection will be significantly different and the problem will be more interesting with the hard stocking space constraint.

The intermediary earns strictly positive profit even though its search technology is no more efficient than that of the manufacturers whose products it resells.¹⁹ 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 proof we show by construction that there is always some set A where this logic is correct. On the other hand, even with no stocking space constraint, the intermediary does not stock all products. In the proof we show that starting from stocking all products, the intermediary can always do strictly better by excluding some loss-making products with high π together with some profit-generating products with low π .

We now solve explicitly for the optimal 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

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

where the average consumer surplus k offered by the intermediary solves

$$\int_{\Omega} q(\pi, v) (v - k) dG = 0 . \tag{4}$$

This is an optimization of functionals. It can be shown that this optimization problem has a solution, and the optimal solution can be derived by treating (4) as a constraint and using the following Lagrange method.

The Lagrangian function is

$$\mathcal{L} = \int_{\Omega} q(\pi, v) [\pi (F(k) - F(v)) + \lambda (v - k)] dG , \tag{5}$$

where λ is the Lagrange multiplier associated with the constraint (4). The first term $\pi (F(k) - F(v))$ is the direct effect on profit of stocking product (π, v) , and the second term $\lambda (v - k)$ reflects the indirect effect from the influence on consumer search behavior

¹⁹By continuity the same can be true even if the intermediary's search technology is less efficient than the manufacturers.

(where $\lambda > 0$ as we will see below). For the products with $v < k$, their direct effect is positive as we explained before, while their indirect effect is negative since stocking them reduces the average consumer surplus of the products in the intermediary. The opposite is true for the products with $v > k$. Since the integrand in (5) is linear in q , the optimal stocking policy is as follows:

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

For given k and λ , we let $I(k, \lambda)$ denote the set of (π, v) for which $q(\pi, v) = 1$. It consists of the following two regions:

$$v < k \quad \text{and} \quad \pi \geq \lambda \frac{k - v}{F(k) - F(v)} , \quad (6)$$

and

$$v > k \quad \text{and} \quad \pi \leq \lambda \frac{k - v}{F(k) - F(v)} . \quad (7)$$

(Notice that it is indifferent whether or not to stock products with $v = k$.)

Graphically we can divide Ω space into four quadrants, using a vertical locus $v = k$ and a horizontal locus $\pi = \frac{\lambda(k-v)}{F(k)-F(v)}$ (which is continuous in v , including at the point $v = k$). Then the intermediary's optimal product selection consists of two "negatively correlated" regions in (π, v) space. The intermediary stocks products in the bottom-right quadrant with high v and low π : since products with $v > k$ make a loss, the intermediary chooses those with the lowest possible π . These products are stocked to attract consumers to search the intermediary. The intermediary also stocks products in the top-left quadrant with low v and high π : since products with $v < k$ make a profit, the intermediary chooses those with the highest possible π . The products in the other regions are not stocked: those with low v and low π would generate little direct profit whilst dissuade consumers from searching, and those with high v and high π are too expensive to buy from their manufacturers.

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 $k \in (\underline{v}, \bar{v})$. Moreover by assumption $[\underline{v}, \bar{v}] \subseteq [0, \bar{s}]$ and so it follows that $F(k) \in (0, 1)$. Since k is interior, we can take the first-order condition of (5) with respect to k , and obtain

$$\int_{I(k, \lambda)} (f(k)\pi - \lambda) dG = 0 , \quad (8)$$

whereupon we observe that $\lambda > 0$.²⁰ Secondly, we have the original constraint (4), which we can rewrite as

$$\int_{I(k,\lambda)} (v - k)dG = 0 . \quad (9)$$

We therefore have a system of two equations (8) and (9) in two unknowns.²¹

The following result summarizes the above analysis:²²

Proposition 1 *The intermediary optimally stocks products in the regions of (6) and (7), where $k \in (\underline{v}, \bar{v})$ and $\lambda > 0$ jointly solve equations (8) and (9).*

To illustrate, consider a uniform product space with $\Omega = [0, 1]^2$ and $G(\pi, v) = \pi v$. If $F(s) = s$ on $[0, 1]$, one can check that in the optimal solution the product space is divided by $v = k$ and $\pi = \lambda$ with $k = \lambda = \frac{1}{2}$. If $F(s) = \sqrt{s}$ on $[0, 1]$, one can check that in the optimal solution the product space is divided by $v = k$ and $\pi = \lambda(\sqrt{k} + \sqrt{v})$ with $k \approx 0.4876$ and $\lambda \approx 0.3515$. The shaded areas in Figure 1 below depict the optimal product range in these two examples. In the first example the intermediary makes profit $\frac{1}{32}$ and improves industry profit by 12.5% relative to the case of no intermediary, and in the second example the intermediary makes profit about 0.036 and improves industry profit by about 10.8%.

In this simple case it is clear that without improving search efficiency, the intermediary must harm consumers by restricting their opportunities to cherry-pick from all products. However, total welfare (which is the sum of industry profit and consumer surplus) could be improved. In fact, this is the case in both of the above examples: the intermediary improves total welfare by about 2.5% and 2.8%, respectively. This is because consumers search too little and buy too few products in the case of no intermediary: they search and

²⁰(8) implies that λ equals $f(k)$ times the average profit of the products stocked by the intermediary. 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 .

²¹If the system has multiple solutions, the solution that generates the highest profit is the optimal one.

²²If we consider a discrete number of products $\{(\pi_i, v_i)\}_{i=1, \dots, n}$, the intermediary's problem becomes

$$\max_{q_i \in \{0,1\}} \sum_i q_i \pi_i [F(k) - F(v_i)]$$

with $k = \sum_i q_i v_i / \sum_i q_i$. This is a combinatorial optimization problem. Given the number of possible stocking policies, 2^n , is very large even for dozens of products, this problem is usually not easy to solve. One approach is to make the problem smooth by allowing stochastic stocking policies with $q_i \in [0, 1]$. Then we can use the Lagrange method and will have bang-bang solutions. The additional complication is how to solve the two equations of k and λ usually depends in a messy way on the locations of the products in the discrete product space.

buy product (π, v) only if $s < v$, but from the perspective of total welfare they should search and buy if $s < \pi + v$. The intermediary forces consumers with $s < k$ to buy some low- v but high- π products which they would not buy otherwise. We will study the socially optimal product selection in Section 5.




Figure 1: Optimal product range: the simple case

Finally we briefly discuss how the shape of the search cost distribution $F(s)$ influences the optimal product range. Observe from Proposition 1 that the horizontal locus $\pi = \lambda(k - v) / [F(k) - F(v)]$ increases in v when $F(s)$ is concave (as we have seen in the above example with $F(s) = \sqrt{s}$) and decreases in v when $F(s)$ is convex. Hence the intermediary's optimal product range tends to contain more low- v and high- v items when $F(s)$ is concave, and the opposite when $F(s)$ is convex. To understand why, consider the case of a concave $F(s)$. Notice that the compensation paid by the intermediary to the manufacturer is $\pi F(v)$, which grows relatively sharply in v when v is low, but grows relatively slowly in v when v is large. Hence it makes sense for the intermediary to mainly stock products with very low v (where the extremely low compensation outweighs the negative effect on consumer search) and very high v (where the small additional compensation is outweighed by the beneficial effects of increased consumer search).

4 The General Case

We now return to the general case: the intermediary has a stocking space of size \bar{m} and can offer both *exclusive* and *non-exclusive* contracts, and the search cost of visiting the intermediary of size m is $h(m) \times s$, where $h(m)$ is weakly increasing. Let $\mathbf{q}(\pi, v) =$

$(q_E(\pi, v), q_{NE}(\pi, v))$ be the stocking policy function, where $q_E(\pi, v) \in \{0, 1\}$ indicates whether product (π, v) is stocked exclusively or not, and $q_{NE}(\pi, v) \in \{0, 1\}$ indicates whether product (π, v) is stocked non-exclusively or not. Note that for each product (π, v) , at most one of $q_E(\pi, v)$ and $q_{NE}(\pi, v)$ can be 1, but it is possible that both are 0 (which happens when the intermediary does not stock product (π, v)). Then

$$q(\pi, v) \equiv q_E(\pi, v) + q_{NE}(\pi, v)$$

indicates whether product (π, v) is stocked or not as before. Using the notation $q(\pi, v)$ is more convenient whenever the exclusivity arrangement does not matter. Henceforth whenever there is no confusion we will suppress the arguments (π, v) in the stocking policy function.

Let us first investigate a consumer's optimal search rule. Given all products are always sold at their monopoly prices, if a consumer decides to visit the intermediary, she will buy all the products available there regardless of whether they are exclusive or not, and will only buy those products not stocked there from the relevant independent manufacturers if $v > s$. (In other words, no consumer will search the same product twice.) Also notice that the order in which the consumer visits the various manufacturers and the intermediary does not matter. Therefore, if a consumer of type s chooses to visit the intermediary, her surplus is

$$u^1(s, \mathbf{q}) = \int qv dG - h \left(\int qdG \right) s + \int_{v>s} (1 - q)(v - s) dG, \quad (10)$$

where the first two terms are the surplus from visiting the intermediary and the final term is the surplus from products not available at the intermediary. Notice that exclusivity arrangement does not matter for consumer surplus in this case.

If a consumer of type s does not visit the intermediary, she will buy all products with $v > s$ available in manufacturers (i.e. not stocked exclusively by the intermediary). Thus her surplus is

$$u^0(s, \mathbf{q}) = \int_{v>s} (1 - q_E)(v - s) dG. \quad (11)$$

Observe that as the intermediary stocks more products exclusively i.e. as q_E 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 in order to attract more consumers.

To ease the exposition, we introduce the following tie-break rule: consumers visit the intermediary only if doing so strictly increases their payoff. As we show in the appendix,

the difference between (10) and (11) is non-negative at $s = 0$ and weakly concave in s . Then we obtain the following cut-off search rule.

Lemma 3 *Consumers search the intermediary if and only if $s < k$, where*

(i) $k = 0$ (nobody searches the intermediary) if $\int q_E dG = 0$ and $\int q dG \leq h(\int q dG)$.

(ii) $k > \bar{s}$ (everybody searches the intermediary) if $\int q v dG > h(\int q dG) \bar{s}$.

(iii) $k \in (0, \bar{s}]$ otherwise and is the solution to

$$k = \frac{\int_{v < k} q v dG + \int_{v > k} q_E v dG}{h(\int q dG) - \int_{v > k} q_{NE} dG} . \quad (12)$$

In this case $k < \bar{v}$ if and only if $\int q v dG < h(\int q dG) \bar{v}$.

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 simply because consumers can then acquire all of the intermediary's products elsewhere at a lower cost. On the other hand, part (ii) shows that all consumers visit the intermediary when it generates sufficiently strong economies of search. Finally, part (iii) shows that in other cases consumers follow a cut-off strategy, and search the intermediary provided their search cost is sufficiently low. Intuitively, in our model a consumer with a lower search cost is a high-demand consumer who is willing to buy more products, so has a higher incentive to visit the intermediary.^{23,24} Notice that in (iii) the non-exclusive products with $v > k$ affect consumer search behavior only by their mass but not by their values. This is because the only impact on consumers of buying them in the intermediary is the change of the search cost associated with them relative to directly buying from their manufacturers. We highlight the condition for $k < \bar{v}$ because if the search economies are sufficiently strong so that the opposite is true, the demand for any product sold by the intermediary will be greater than when it is sold directly by its manufacturer, so there will be no loss-making products.

Given the consumer search rule in Lemma 3 and the result of monopoly pricing from Lemma 1, the intermediary's profit, when it chooses a stocking policy \mathbf{q} , is

$$\Pi(\mathbf{q}) = \int_{v < k} q \pi[F(k) - F(v)] dG + \int_{v > k} q_E \pi[F(k) - F(v)] dG . \quad (13)$$

²³More precisely, 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.

²⁴This is consistent with the recent trend that more small local grocery stores are opened up to cater for consumers who only need a small basket of products and have no time to travel to big stores.

For a product with $v < k$, the profit from it is independent of its exclusivity (i.e., only $q = q_E + q_{NE}$ matters). 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 a manufacturer with $v < k$. Hence the intermediary always 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 (13) is profit earned on exclusive products with $v > k$. This takes the same form as in the previous section, and these products are stocked at a loss to drive store traffic. Note that this second term exists only if $k < \bar{v}$.

Finally, and most interestingly, products with $v > k$ which are stocked non-exclusively do not appear in equation (13), 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 directly from the manufacturer. Hence, to make up the manufacturer's lost revenue 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 for a given k , they can influence consumers' search behavior via k and so indirectly affect the intermediary's profit. As a result, the intermediary may still have an incentive to stock them.

The following lemma gives some sufficient conditions for the intermediary to make a profit.

Lemma 4 *The intermediary will always stock a strictly positive measure of products and earn a strictly positive profit if $h(m) = m$ for all $m \in [0, \bar{m}]$ or if $h(m) < m$ for some $m \in (0, \bar{m}]$.*

When the intermediary does not improve search efficiency, it can make a profit by stocking some products exclusively as in the simple case. When it improves search efficiency for some m , it can make a profit at least by stocking a measure m of products non-exclusively, though as we will see below using non-exclusive contracts only is usually not the optimal stocking policy unless \bar{m} is sufficiently large and economies of search are sufficiently strong.

In the following, we characterize the optimal product selection. The analysis turns out to be more transparent if we start with the case with no stocking space limit (i.e. $\bar{m} = 1$). We will investigate the case of $\bar{m} < 1$ afterwards. Henceforth, we assume $h'(m) \in [0, 1]$, i.e., there are (weakly) economies of scale in searching the intermediary when it expands marginally.

4.1 Unlimited stocking space

When the intermediary has no limit on how many products it can stock, the following lemma gives a first qualitative description of what the optimal product range looks like:

Lemma 5 *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 simple case in Section 3 is that now the intermediary will optimally stock *all* products with $v > k$. Suppose to the contrary that some positive measure set of products B with $v > k$ are not stocked. Then we show in the proof that stocking all products in B non-exclusively is a profitable deviation. As we saw earlier the intermediary earns zero profit from these products, but they induce more consumers (i.e., those with s slightly above k) to visit the intermediary since $h'(m) \leq 1$ implies that searching the products in B in the intermediary saves them search costs.²⁵ Once they visit the intermediary, they also buy other products available there which are on average profitable.

Nevertheless similar to the simple case, products with $v > k$ that are stocked exclusively make a loss, and so are chosen to have the lowest π possible in order to minimize that loss. Moreover, and again similar to the simple case, 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 (13), where k is given in Lemma 3. It is more convenient to introduce another parameter $m = \int qdG$, i.e., the measure of all 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 makes a strictly positive profit in the optimal solution (so $m > 0$ and $k > 0$), and not all consumers visit it (so $k < \bar{s}$). Lemma 4 has provided simple sufficient conditions for the former, and according to Lemma 3 a simple sufficient condition for the latter is $\int qvdG/h(\int qdG) < \bar{s}$

²⁵Note that in the knife-edge case where $h'(m) = 1$ the intermediary is indifferent in stocking products in B , since doing so does not change the search cost of marginal consumers, and so has no effect on the store traffic.

for any q , which is equivalent to $\max_x \int_x^{\bar{v}} v dG / h(\int_x^{\bar{v}} dG) < \bar{s}$.²⁶

Now the intermediary's problem is to maximize (13) subject to (12). It is more convenient to treat $m = \int qdG$ as another constraint. (This may become a real constraint when we introduce a limited stocking space in next subsection.) Notice that (12) can be rewritten as

$$\int_{v < k} qvdG + \int_{v > k} (q_E v + q_{NE} k) dG - h(m)k = 0. \quad (14)$$

Then the Lagrangian function of the problem is

$$\begin{aligned} \mathcal{L} = & \int_{v < k} q\pi [F(k) - F(v)] dG + \int_{v > k} q_E \pi [F(k) - F(v)] dG \\ & + \lambda \left(\int_{v < k} qvdG + \int_{v > k} (q_E v + q_{NE} k) dG - h(m)k \right) + \mu \left(m - \int_{v < k} qdG - \int_{v > k} qdG \right), \end{aligned}$$

where λ is the Lagrange multiplier associated with the constraint (14), and μ is the multiplier associated with the constraint $m = \int qdG$.²⁷ The intermediary maximizes \mathcal{L} by choosing q , k and m .

It is useful to rewrite the Lagrange function as

$$\begin{aligned} \mathcal{L} = & \int_{v < k} q[\pi(F(k) - F(v)) + \lambda v - \mu] dG \\ & + \int_{v > k} (q_E[\pi(F(k) - F(v)) + \lambda v - \mu] + q_{NE}(\lambda k - \mu)) dG - \lambda k h(m) + \mu m. \end{aligned} \quad (15)$$

This can be explained similarly as in the simple case by using the direct and indirect effect of stocking a product. In particular, $\lambda v - \mu$ reflects the indirect effect on consumer search incentive of stocking a product with $v < k$ or exclusively stocking a product with $v > k$, and $\lambda k - \mu$ reflects a similar effect of stocking a product with $v > k$ non-exclusively. As we show in the proof of the following proposition, $\mu = \lambda k h'(m) \leq \lambda k$ given $h'(m) \leq 1$. Therefore, unsurprisingly stocking a product with $v > k$ (regardless of its exclusivity) always increases consumers' incentive to visit the intermediary. If $h'(m) < 1$, even stocking a product with v slightly below k increases consumer search incentive as well.

²⁶More precisely, $\int_x^{\bar{v}} v dG = \int_x^{\bar{v}} \int_{\underline{\pi}(v)}^{\bar{\pi}(v)} v g(\pi, v) d\pi dv$. The equivalence result is because for any stocking policy q , $\exists x \in [v, \bar{v}]$ such that $\int qdG = \int_x^{\bar{v}} dG$, and in the same time $\int qvdG \leq \int_x^{\bar{v}} v dG$ since the average v improves when the product mass is allocated to the products with the highest possible v 's.

²⁷If $m = 1$, then we must have $q = 1$ everywhere and then the second constraint become redundant and the μ term disappears.

Proposition 2 *In the general case without stocking space limit, suppose the intermediary makes a strictly positive profit and $k \in (0, \bar{s})$ in the optimal solution (which is true if the conditions in Lemma 4 hold and $\max_x \int_x^{\bar{v}} v dG / h(\int_x^{\bar{v}} dG) < \bar{s}$). Then the optimal product selection features either*

(i) $m < 1$, and among the products with $v < k$, only those with

$$\pi \geq \lambda \frac{h'(m)k - v}{F(k) - F(v)} \quad (16)$$

are stocked and it does not matter whether they are stocked exclusively or non-exclusively, and among the products with $v > k$ (if $k < \bar{v}$), those with

$$\pi \leq \lambda \frac{k - v}{F(k) - F(v)} \quad (17)$$

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:

$$k = \frac{\int_{v < k} q v dG + \int_{v > k} q_E v dG}{h(m) - \int_{v > k} q_{NE} dG}, \quad (18)$$

$$\lambda = f(k) \frac{\int_{v < k} q \pi dG + \int_{v > k} q_E \pi dG}{h(m) - \int_{v > k} q_{NE} dG}, \quad (19)$$

$$m = \int q dG; \quad (20)$$

or

(ii) $m = 1$ (i.e., all products are stocked), and among the products with $v > k$ (if $k < \bar{v}$), 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 (18) and (19) with $q = 1$ and $m = 1$.

This characterization is consistent with the qualitative description of the optimal product range in Lemma 5. The main qualitative difference, compared to the simple case in Section 3, is that the intermediary will stock the products in the top-right corner non-exclusively (which were excluded when only exclusive contracts are available). Another difference is, if economies of scale in search is strong enough, the intermediary will stock all

products.²⁸ A subtler difference is that when $h'(m) < 1$, $\frac{h'(m)k-v}{F(k)-F(v)} \rightarrow -\infty$ when $v \rightarrow k^-$. This implies that for those products with v close to but smaller than k , they will always be stocked regardless of their π .

Notice that for the stocked products with $v < k$, the exclusivity arrangement does not matter. This is because even if such a product is also available for purchase in its manufacturer, the consumers who do not visit the intermediary (i.e., those with $s > k$) will not bother to visit the manufacturer either given $v < s$. This makes these products as if they were sold exclusively by the intermediary even if the contract is not exclusive. One way to tie-break this indifference is to introduce 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.) For this reason, in the following we claim that the products with $v < k$ are stocked non-exclusively.

To illustrate the optimal product selection, consider the uniform example with $G(\pi, v) = \pi v$ and $F(s) = s$. Suppose $h(m) = \alpha + \beta m$. Figure 2(a) and 2(b) below depict the optimal product selection when $h(m) = m$ and $h(m) = 0.7m$, respectively. (In the first example $k = \lambda = \frac{1}{2}$ and $m = 0.75$, and in the second $k = \lambda \approx 0.826$ and $m \approx 0.769$.) Now the products in the top-right corner are stocked non-exclusively,²⁹ and as economies of search improve the intermediary stocks more products overall but fewer exclusive products. With stronger economies of search the intermediary will rely less on exclusive products to attract consumers to visit.

Figure 2(c) and 2(d) below depict the optimal product selection when $h(m) = 0.4 + 0.5m$ and $h(m) = 0.4 + 0.2m$, respectively. (According to Lemma 4, the intermediary can make a positive profit in both examples. In the first example $k = \lambda \approx 0.487$ and $m \approx 0.964$, and in the second $k = \lambda \approx 0.832$ and $m \approx 0.985$.) Given there is a relatively large fixed component in the search cost, the intermediary needs to stock enough products to make consumers willing to visit. But similar as in the previous two examples, as economies of search become stronger it stocks more products overall but fewer exclusive

²⁸A simple sufficient condition for $m = 1$ is $\int vdG/h(1) > \bar{s}$. Under this condition, Lemma 3 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. A sufficient condition for $m < 1$ is: $\underline{\pi} = \underline{v} = 0$, $[0, \epsilon]^2 \subset \Omega$ for a sufficiently small $\epsilon > 0$, $h(1) < 1$, $h'(1) > 0$ and $\int vdG/h(1) < \bar{s}$. (The proof is available upon request.) In general, however, it appears hard to find a necessary and sufficient primitive condition for $m < 1$.

²⁹In the first example with no economies of search the intermediary only has a weak incentive to non-exclusively stock the products in the top-right corner $[0.5, 1]^2$.

products. Eventually if β is sufficiently close to zero, the intermediary will stock all products non-exclusively. In such a case, it will be more interesting to investigate the optimal product selection with a stocking space constraint.³⁰

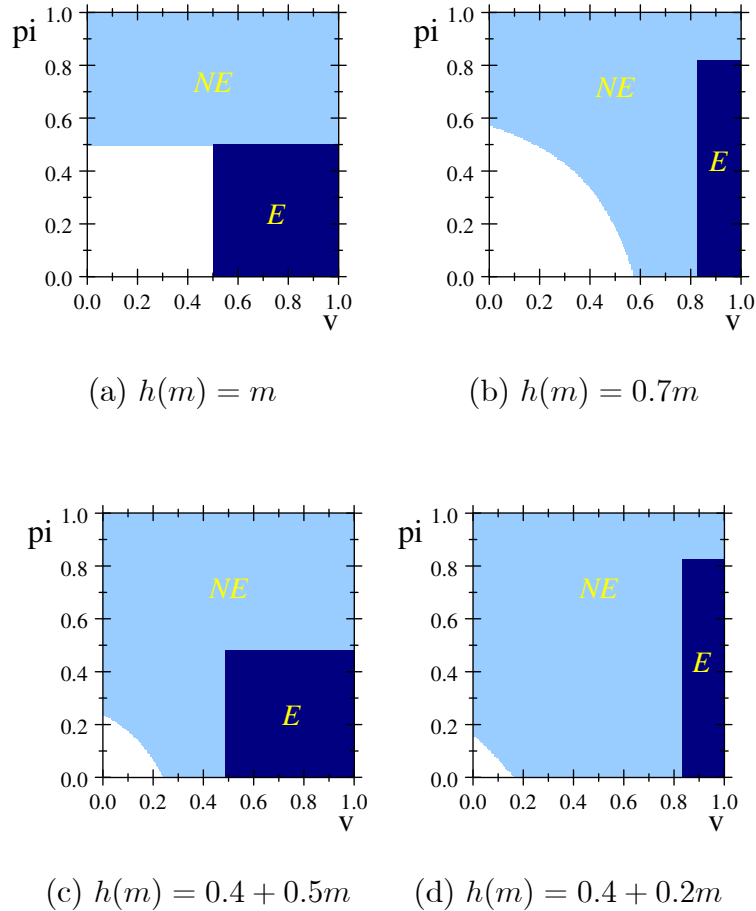


Figure 2: Optimal product range: the general case with $\bar{m} = 1$

4.2 Limited stocking space

We now introduce the stocking space limit $\bar{m} < 1$. If the constraint does not bind in the optimal solution, the characterization of the optimal product range is the same as in part (i) of Proposition 2. In the following, we focus on the case when the constraint binds in

³⁰Notice that stronger economies of search in visiting the intermediary can also be interpreted as more costly direct-to-consumer sales. Therefore, our discussion also suggests that when direct-to-consumer sales becomes easier (e.g., due to the online market), the retailer will become smaller and rely more on offering exclusive products.

the optimal solution. Then we have a real constraint $\bar{m} = \int qdG$, but the Lagrangian function is the same as (15) except that m is replaced by \bar{m} :

$$\begin{aligned} \mathcal{L} = & \int_{v < k} q[\pi(F(k) - F(v)) + \lambda v - \mu]dG \\ & + \int_{v > k} \{q_E[\pi(F(k) - F(v)) + \lambda v - \mu] + q_{NE}(\lambda k - \mu)\} dG - \lambda k h(\bar{m}) + \mu \bar{m}. \end{aligned} \quad (21)$$

Note that μ is now the Lagrangian multiplier associated with the hard stocking space constraint.

The following proposition reports the optimal product range in this case:

Proposition 3 *In the general case with a limited stocking space $\bar{m} < 1$, suppose the intermediary makes a strictly positive profit and $k \in (0, \bar{s})$ in the optimal solution (which is true if the conditions in Lemma 4 hold and $\int_x^{\bar{v}} v dG / h(\int_x^{\bar{v}} dG) < \bar{s}$ for any x such that $\int_x^{\bar{v}} dG \leq \bar{m}$). If the stocking space constraint binds in the optimal solution, then among the products with $v < k$, only those with*

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

are stocked and it does not matter whether they are stocked exclusively or non-exclusively, and among the products with $v > k$ (if $k < \bar{v}$ in the optimal solution), the optimal selection features either

(i) $\lambda k - \mu > 0$ and those with

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

are stocked exclusively and the others are stocked non-exclusively, or

(ii) $\lambda k - \mu = 0$ and those with

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

are stocked exclusively and some of the other products are stocked non-exclusively, or

(iii) $\lambda k - \mu < 0$ and only those with

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

are stocked exclusively. The parameters k , λ and μ solve (18)-(20) with m replaced by \bar{m} .

From (21), we can see that $\lambda k - \mu$ captures the effect on the intermediary's profit of stocking a product with $v > k$ non-exclusively. Its sign determines whether the intermediary should stock any such products. When $\lambda k - \mu = 0$ in the optimal solution, the intermediary is indifferent in which such products to select as long as the measure of them is such that $\lambda k - \mu = 0$. As a result, the product selection in this region is not uniquely pinned down.

It appears hard to find primitive conditions for the sign of $\lambda k - \mu$ in the optimal solution. But intuitively when the space constraint just starts binding, we have $\mu = \lambda k h'(\bar{m})$ from the previous analysis, and so $\lambda k - \mu > 0$ if $h'(\bar{m}) < 1$. If the constraint is tightened slightly from this point, what products should be removed? They should be the products with $v < k$ around the boundary $\pi(F(k) - F(v)) + \lambda v - \mu = 0$ because they contribute zero to the intermediary's profit while all other stocked products have a strictly positive contribution. This process continues until $\lambda k - \mu = 0$. Now if the stocking space limit further shrinks, some non-exclusive products with $v > k$ should be dropped because they have zero contribution. But they should not be dropped all at once because otherwise the constraint would be suddenly slack and we would have $\lambda k - \mu > 0$. Therefore, there should exist a range of \bar{m} in which $\lambda k - \mu = 0$ and the non-exclusive products with $v > k$ are removed gradually. Eventually we will reach the stage of $\lambda k - \mu < 0$ and there are no non-exclusive products with $v > k$ any more. In this stage if \bar{m} further shrinks, the least profitable products around the boundary $\pi(F(k) - F(v)) + \lambda v - \mu = 0$ (which applies for both $v < k$ and $v > k$) should be removed. Notice that when $\lambda k - \mu < 0$, we have $\lim_{v \rightarrow k^-} \frac{\mu - \lambda v}{F(k) - F(v)} = \infty$ and $\lim_{v \rightarrow k^+} \frac{\mu - \lambda v}{F(k) - F(v)} = -\infty$, so the products with v sufficiently close to k should be excluded regardless of their π .³¹ This intuitive discussion is confirmed in the numerical examples below.

Consider the running example with uniform product space $G(\pi, v) = \pi v$. To make it possible that $k > \bar{v}$ (which case we have not explored before) but in the same time $k < \bar{s}$, suppose $F(s) = s/2$, i.e. s is uniformly distributed on $[0, 2]$. The stocking space constraint is more likely to bind when economies of search are stronger. So let us consider the polar case where $h'(m) = 0$, i.e., $h(m)$ is a constant. Suppose $h(m) = \alpha$ and $\alpha > \frac{1}{4}[1 - (1 - \bar{m})^2]$ so that $k < \bar{s}$.³² Figure 3 below describes, when $\alpha = 0.4$, how the optimal product

³¹Intuitively, for the products with v slightly below k , their demand is only expanded a little via being sold through the intermediary, and for the products with v slightly above k , they contribute little in attracting more consumers to visit.

³²When the stocking space is \bar{m} , consumers have the highest incentive to visit the intermediary if it stocks all the products with $v \geq 1 - \bar{m}$ exclusively. Therefore, if $\int_{1-\bar{m}}^1 v dv < \alpha \bar{s}$ or equivalently $\alpha > \frac{1}{4}[1 - (1 - \bar{m})^2]$ given $\bar{s} = 2$, not all consumers will visit the intermediary (i.e. $k < \bar{s}$).

selection varies as \bar{m} shrinks.

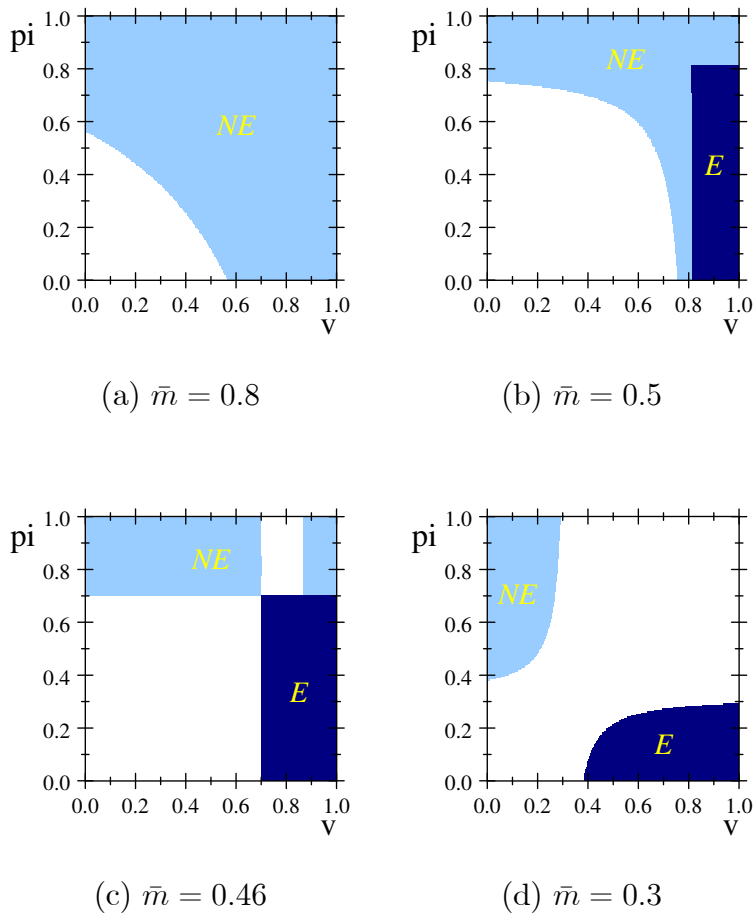


Figure 3: Optimal product range: the general case with $\bar{m} < 1$ and $h(m) = 0.4$

When \bar{m} is greater than about 0.65, $k > 1$ and so there is no region of $v > k$. In this case the demand for any stocked product is expanded compared to direct sales, and so there are no loss-making products. When \bar{m} is between about 0.65 and about 0.463, $k < 1$ and so the region of $v > k$ appears. In the same time, $\lambda k - \mu > 0$ and so result (i) in Proposition 3 applies: all the products in the region of $v > k$ are stocked, but only those with relatively low π are stocked exclusively. This is qualitatively similar to Figure 2 when there is no stocking space limit but economies of search are relatively weak. When \bar{m} is between about 0.463 and about 0.454, $k < 1$ and $\lambda k - \mu = 0$, so result (ii) in Proposition 3 applies: some non-exclusive products in the top-right corner start to be excluded, but there is flexibility in how to select products in this region. (In Figure 3(c) we remove those with relatively low v .) When \bar{m} is below about 0.454, $k < 1$ and $\lambda k - \mu < 0$. Then

result (iii) in Proposition 3 applies: now there are no non-exclusive products with $v > k$ any more. In this case as we already pointed out the products with v close to k will all be excluded regardless of their π . It is also worth mentioning that when \bar{m} becomes smaller, the intermediary tends to stock more exclusive products proportionally. (The fraction of exclusive products among all stocked products is 0, 0.3, 0.46, and 0.5, respectively, in the above four cases in Figure 3.) This is because when the store becomes smaller, the intermediary may need to use more exclusively available products to induce consumers to visit.

5 Comparison With the Social Optimum

We now turn to the optimal product selection by a social planner who aims to maximize total welfare which is defined as the sum of industry profit and consumer surplus. We assume that the social planner can control the stocking policy \mathbf{q} but not firm pricing and consumer search behavior.

As we pointed out before, if visiting the intermediary does not improve search efficiency (i.e., if $h(m) = m$), 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 the intermediary might be “too big” or stock too many products exclusively, relative to the socially optimal size. But this negative effect on consumers will be mitigated by the improved search efficiency when $h(m) < m$. On the other hand, as we also mentioned before, consumers search too little relative to the social optimum because they ignore the effect of their search decision on profit. When a product has v slightly below s , a consumer of type s will not search it in the case of no intermediary. But from the social planner’s view she should have searched it as long as it is socially efficient (i.e. if $\pi + v > s$). Therefore, the intermediary can improve market efficiency by forcing consumers to search some low- v but socially efficient products. The following analysis will illustrate these three effects. In general it is hard to compare them analytically, though numerical examples suggest that the first effect dominates.

Given a stocking policy \mathbf{q} , the consumer search rule is the same as in Lemma 3. Total welfare can then be written as

$$W(\mathbf{q}) \equiv \int \pi F(v) dG + \Pi(\mathbf{q}) + \int_0^k u^1(s, \mathbf{q}) dF(s) + \int_k^{\bar{s}} u^0(s, \mathbf{q}) dF(s) . \quad (22)$$

The first term is the profits of manufacturers, who always earn $\pi F(v)$ regardless of whether they sell their product by themselves or via the intermediary. The second one is the intermediary's profit, which we defined earlier in equation (13). The third one is the surplus of consumers with $s < k$ who search the intermediary, where $u^1(s, \mathbf{q})$ was defined earlier in equation (10). The fourth one is the surplus of consumers with $s \geq k$ who choose not to visit the intermediary, where $u^0(s, \mathbf{q})$ again was defined earlier in equation (11). Notice that the consumers with $s \geq k$ are always made (weakly) worse off by 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.

The social planner wishes to choose a stocking policy \mathbf{q} in order to maximize $W(\mathbf{q})$. In the following, we focus on the case with no stocking space limit (i.e., $\bar{m} = 1$). The analysis is then parallel to Section 4.1. (The case with a binding space constraint can be dealt with similarly as in Section 4.2.) We again use $m = \int qdG$ to denote the measure of products stocked by the intermediary. 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 in section 4.1. The intuition and the proof are both closely related to that of Lemma 5. We then solve explicitly for the social planner's optimum. As before, we treat the consumer search rule in equation (14) and $m = \int qdG$ as two constraints, and let λ and μ be the respective multipliers associated with these two constraints.

Proposition 4 *In the general case without stocking space limit, suppose the social optimum has $m > 0$ and $k \in (0, \bar{s})$ (which is true if the conditions in Lemma 6 hold and $\max_x \int_x^{\bar{v}} vdG/h(\int_x^{\bar{v}} dG) < \bar{s}$). Then the socially optimal product selection features either*

(i) $m < 1$, and among the products with $v < k$, only those with

$$\pi \geq \frac{\lambda(kh'(m) - v) + \int_v^k (s - v)dF(s) + (h'(m) - 1) \int_0^k s dF(s)}{F(k) - F(v)} \quad (23)$$

are stocked and the exclusivity arrangement does not matter, and among the products with $v > k$ (if $k < \bar{v}$), those with

$$\pi \leq \frac{\lambda(k - v) + \int_v^k (s - v)dF(s)}{F(k) - F(v)} \quad (24)$$

are stocked exclusively and the others are stocked non-exclusively. In this case, the parameters k , λ , and m solve the same system of equations as (18) - (20).

or

(ii) $m = 1$ (i.e., all products are stocked), and among the products with $v > k$ (if $k < \bar{v}$), those with

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

are stocked exclusively, and the exclusivity arrangement for the products with $v < k$ does not matter. In this case, λ and k solve (18) and (19) with $q = 1$ and $m = 1$.

This characterization is qualitatively similar to the optimal product range in Proposition 2. In particular, the parameters k , λ , and m solve the same system as in the intermediary's problem. But this does not imply that they will have the same solution in the two problems, because for given k , λ , and m the product selection takes different forms in the two problems. For this reason, a general comparison between the socially optimal selection and the intermediary's optimal selection is hard. Nevertheless for a fixed (k, λ) , by comparing (17) and (24) 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. Similarly, for a fixed (k, λ) , by comparing (16) and (23) we can see that if $h'(m) = 1$ (i.e., if there are no marginal economies of search), the intermediary stocks too many low- v products. But this effect can be reversed if $h'(m)$ is sufficiently small.

To illustrate, we return to our running example with $G(\pi, v) = \pi v$ and $F(s) = s$. We compare the socially optimal solution with the profit-maximizing solution when $h(m) = m$ and $h(m) = 0.4 + 0.5m$, respectively. In the first example, one can check that $k = \lambda = \frac{1}{2}$ (which is the same as in the intermediary's solution) and $m \approx 0.6875$ in the

socially optimal solution. Figure 4(a) below describes the socially optimal product range. Compared to the intermediary's solution, the social planner stocks fewer products overall and fewer products exclusively, and the social planner's product set is a strict subset of the intermediary's. In the second example, one can solve $k \approx 0.440$, $\lambda \approx 0.487$ and $m \approx 0.963$. Note that k and λ are now different from those solved in the intermediary's problem and m is slightly smaller. Figure 4(b) below describes the socially optimal product range in this example. Again, the social planner stocks fewer products overall and fewer products exclusively than the intermediary, though in this example the social planner's product set is not exactly a subset of the intermediary's.

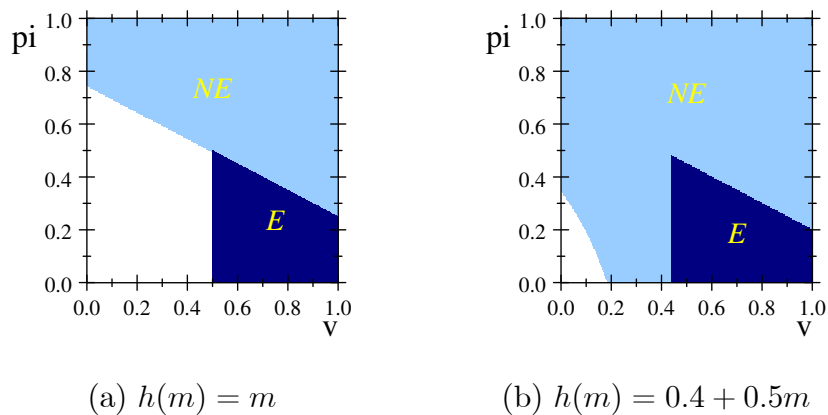


Figure 4: Socially optimal product range

Finally, it is worth mentioning that although the intermediary tends to stock too many products exclusively relative to the social optimum, banning exclusive products all together can harm efficiency unless economies of search are sufficiently strong. This can be easily seen from the extreme case of $h(m) = m$ where the intermediary will not exist if no exclusive contracts are allowed, and this can reduce total welfare as we have seen in the two examples in Figure 1.

6 Discussion

In this section, we first discuss the foundation of the (π, v) product space, and then study an extension with upstream competition (i.e., each product having more than one manufacturer).

6.1 Foundation of (π, v) Product Space

We provide two classes of demand functions which can generate the (π, v) product space. We also discuss how a product's demand curvature or demand elasticity affects where it is located in the product space.

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

$$Q_i(p_i) = a_i \left(1 - \frac{1 - \sigma_i}{2 - \sigma_i} (p_i - \mu_i) \right)^{\frac{1}{1 - \sigma_i}}, \quad (25)$$

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.³³ 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 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 2.³⁴

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 respectively

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

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-making products are those 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).³⁵

³³The curvature of demand function $Q(p)$ is defined as $Q''(p)Q(p)/[Q'(p)]^2$. It measures the elasticity of the slope of the inverse demand function.

³⁴It also includes constant elasticity demand when $\sigma_i = \frac{2 + \mu_i}{1 + \mu_i} \in (1, 2)$.

³⁵Anderson 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.

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

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

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 $p_i^m = \left(\frac{1}{1+\sigma_i}\right)^{\frac{1}{\sigma_i}}$, and monopoly profit and consumer surplus are respectively

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

Both π_i and v_i increase in a_i , and π_i/v_i increases in σ_i and so decreases in elasticity.³⁶ 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 to drive store traffic, 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, such as (a, σ) in the above examples (assuming in the first example that we fix $\mu = c = 0$). Then generically there is a one-to-one correspondence between (a, σ) and (π, v) , such that each point in the (π, v) space represents a single product.³⁷ 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 a continuum of different products. 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. This, however, does not affect our analysis because all the objective functions in this paper are linear in the stocking policy variables, and so we always have bang-bang solutions.

³⁶In this example $\pi_i/v_i > 1$ for any $\sigma_i > 0$, so it can only generate half of the quadrant \mathbb{R}_+^2 .

³⁷Notice, however, that even if products are uniformly distributed in the (a, σ) space, they can be non-uniformly distributed in the (π, v) space. That is why we consider a general distribution G .

6.2 Upstream competition

We now extend our model by introducing upstream competition between manufacturers and show that our main insights are still valid. In particular we assume now that each product is supplied by two homogeneous manufacturers. To simplify the exposition we focus on the case of no search economies i.e. $h(m) = m$, where m denotes the measure of distinct products stocked by the intermediary. (Therefore if the intermediary contracts with two manufacturers supplying the same product, the cost of searching the intermediary only increases by one unit.) We also assume that the intermediary has no stocking constraint i.e. $\bar{m} = 1$, and is able to offer both exclusive and non-exclusive contracts. The timing closely follows that of the main model. At the first stage the intermediary announces to all manufacturers its stocking intentions, and then makes public (possibly discriminatory) offers which specify both a two-part tariff and (non-)exclusivity. Manufacturers simultaneously accept or reject their offers, and (when appropriate) believe that the other manufacturer of their product will accept. At the following stages firms set prices and consumers search sequentially with passive beliefs and randomize whenever indifferent.

Closely following Lemma 1 from earlier, we can prove that in equilibrium all sellers of a product charge the monopoly price. The intuition is the same as before: the intermediary uses bilaterally-efficient two-part tariffs to avoid double marginalization, and the search friction nullifies direct pricing competition between sellers just like in Diamond (1971).³⁸ Consequently we can still represent products as points in a two-dimensional (π, v) space. In the spirit of our earlier analysis, $q_E(\pi, v) \in \{0, 1\}$ indicates whether product (π, v) is stocked exclusively by the intermediary i.e. consumers cannot buy it elsewhere. That is, $q_E(\pi, v) = 1$ if and only if the intermediary contracts with both manufacturers exclusively. Similarly, $q_{NE}(\pi, v) \in \{0, 1\}$ indicates whether product (π, v) is stocked non-exclusively by the intermediary i.e. consumers also have the opportunity to buy it from a manufacturer. That is, $q_{NE}(\pi, v) = 1$ if and only if the intermediary contracts with at least one manufacturer non-exclusively. Hence

$$q(\pi, v) = q_E(\pi, v) + q_{NE}(\pi, v) \in \{0, 1\}$$

indicates whether or not the intermediary stocks product (π, v) . It is straightforward to see

³⁸One subtle difference is that here we need $f(0) = 0$ to sustain monopoly pricing when the manufacturers both sell direct to consumers. This condition is satisfied as long as s is bounded away from 0. In the uniform example below, for convenience we still assume $s \sim U[0, 1]$, but this can be regarded as the limit case of $s \sim U[\epsilon, \epsilon + 1]$ with $\epsilon \rightarrow 0$.

that given monopoly pricing, the payoffs from respectively searching and not searching the intermediary are the same as those in equations (10) and (11) with $h(m) = m$. Therefore Lemma 3 implies that provided $\int q_E dG > 0$ and $\bar{s} \geq \bar{v}$, there is a unique cutoff $k \in (\underline{v}, \bar{v})$ satisfying (12) with $h(m) = m$, such that consumers search the intermediary if and only if $s < k$. For convenience, we rewrite it here as

$$\int_{v < k} q(v - k) dG + \int_{v > k} q_E(v - k) dG = 0 . \quad (26)$$

Notice that given that the two manufacturers for each product are homogeneous, here only the stocking policy at the product level (instead of at the manufacturer level) matters for consumer search decision.

Now consider how much the intermediary must compensate manufacturers (on top of the production cost) in order to stock their product. The following is a useful preliminary result:

Lemma 7 (i) *If the intermediary contracts with both manufacturers of a product non-exclusively, it does not need to compensate them.*

(ii) *If the intermediary contracts with both manufacturers of a product exclusively, it needs to compensate each by an amount $\max\{0, \pi[F(v) - F(k)]\}$.*

(iii) *It is (weakly) dominated for the intermediary to contract with only one manufacturer of a product, or to contract with both but only exclusively with one of them.*

We explain parts (i) and (ii), and leave the details of part (iii) to the appendix. First consider products with $v < k$. Provided one manufacturer supplies the intermediary, the other manufacturer expects to make no sales and is therefore willing to also supply the intermediary at marginal cost. Second consider products with $v > k$. Provided one manufacturer supplies the intermediary, the other manufacturer is unable to sell its product to consumers with $s < k$. If the intermediary contracts with both manufacturers non-exclusively, each manufacturer earns $\frac{1}{2}\pi[F(v) - F(k)]$ irrespective of whether it accepts or rejects the intermediary's contract. Hence in this case each manufacturer is willing to provide its product to the intermediary at marginal cost. If instead the intermediary wishes to have exclusive sales rights, a manufacturer that rejects its contract becomes a monopolist over consumers with $s \in (k, v)$ and therefore earns $\pi[F(v) - F(k)]$. Hence in this case each manufacturer must be compensated by that amount.

Lemma 7 implies that the intermediary only needs to consider three options: either not stock a product, or stock it according to options (i) or (ii). Whenever it stocks a product, the intermediary contracts with both the manufacturers so as to induce them to

accept lower compensation. It is now without loss of generality to use $q_{NE}(\pi, v) \in \{0, 1\}$ to indicate whether the intermediary contracts with both manufacturers non-exclusively.

The intermediary's profit function is then

$$\int_{v < k} q\pi F(k)dG + \int_{v > k} q_{NE}\pi F(k)dG + \int_{v > k} q_E\pi[3F(k) - 2F(v)]dG . \quad (27)$$

According to Lemma 7, no compensation beyond the production cost is needed for stocking a product with $v < k$ (regardless of whether it is stocked exclusively or non-exclusively) or stocking a product with $v > k$ non-exclusively. For a given k , these products are now cheaper to stock than in the basic model due to the upstream competition. For an exclusively stocked product with $v > k$, the intermediary earns gross profit $\pi F(k)$ but needs to pay $2\pi [F(v) - F(k)]$ to its manufacturers. The compensation to each manufacturer is also lower than in the basic model, but the intermediary now needs to compensate two manufacturers instead of one. Whether it is now cheaper or more expensive to stock these exclusive products depends on v . Those with $F(v) < 2F(k)$ become cheaper to stock (and can even become profit generators), while those with $F(v) > 2F(k)$ (if any) become more expensive to stock.

The intermediary maximizes (27) subject to the search constraint (26). The Lagrangian function of this optimization problem is

$$\begin{aligned} \mathcal{L} = & \int_{v < k} q [\pi F(k) + \lambda(v - k)] dG + \int_{v > k} q_{NE}\pi F(k) dG \\ & + \int_{v > k} q_E \{ \pi [3F(k) - 2F(v)] + \lambda(v - k) \} dG . \end{aligned} \quad (28)$$

It is easy to argue that some products must be stocked exclusively in the optimal solution, so we must have an interior solution $k \in (\underline{v}, \bar{v})$. Then the first-order condition of (28) with respect to k yields

$$\lambda = f(k) \frac{\int_{v < k} q\pi dG + \int_{v > k} q_{NE}\pi dG + 3 \int_{v > k} q_E\pi dG}{\int_{v < k} q dG + \int_{v > k} q_E dG} , \quad (29)$$

from which we deduce $\lambda > 0$. Then it is straightforward to derive the following result.

Proposition 5 *The optimal product selection with upstream competition is characterized as follows:*

(i) *The intermediary buys products with $v < k$ from both manufacturers if*

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

and otherwise does not stock them.

(ii) The intermediary buys all products with $v > k$ from both manufacturers. It contracts with the manufacturers exclusively if

$$\pi \leq \frac{\lambda}{2} \frac{v - k}{F(v) - F(k)},$$

and otherwise contracts with them non-exclusively. The parameters k and λ solve equations (26) and (29).

Our predictions about the intermediary's optimal product range are thus qualitatively robust to the introduction of upstream competition. Nevertheless notice that for a fixed (k, λ) , both the stocking region of $v < k$ and the non-exclusive region of $v > k$ expands compared to the basic model. This is because these products are now cheaper to acquire as we have explained. While the exclusive region of $v > k$ shrinks since the boundary in part (ii) is lower than (17). This is because stocking the products with $v > k$ exclusively is now relatively less profitable compared to stocking them non-exclusively. Of course considering upstream competition will change (k, λ) as well. Once that is taken into account, it appears no general conclusions can be drawn. Figure 5 plots the optimal product range in the example with uniform product space when $F(s) = s$ and $F(s) = \sqrt{s}$, respectively. In former case, the intermediary stocks fewer exclusive products and also fewer products overall compared to the basic model, while the opposite is true in the latter case.

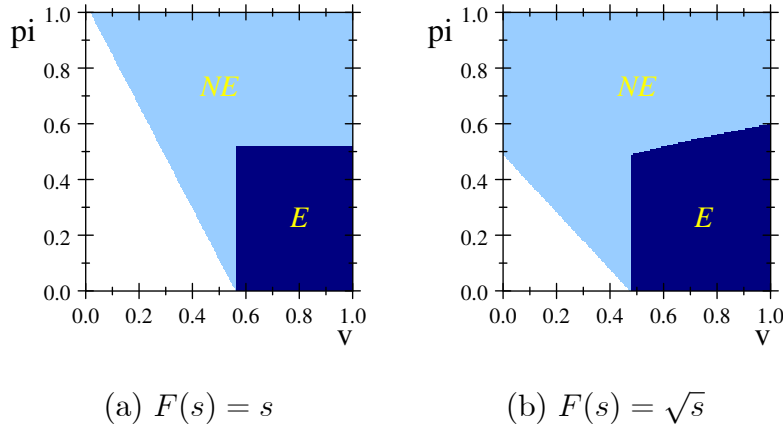


Figure 5: Optimal product range with upstream competition

However, as expected in both cases the intermediary's profit is higher than in the basic model, since the upstream competition brings down the overall compensation needed

for the manufacturers. The profit increases from $\frac{1}{32}$ to about 0.2 in the first case, and from about 0.036 to about 0.305 in the second. The main source of the significant profit improvement is the reduced compensation to the manufacturers, instead of the reoptimization of the product selection. For example, in the first case if the intermediary simply adopted the same stocking policy as in the basic model, its profit would already increase to 0.1875. Reoptimization further improves the profit, but only by a relatively small amount. Finally, notice that this also implies that the intermediary has a strong incentive to produce a private label of a product, since it has a similar effect as introducing upstream competition in reducing the compensation to the manufacturer.

7 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 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 in order to increase search, 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 or two homogenous manufacturers. It will be interesting 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.

Finally, we want to point out that the framework developed in this paper could be modified to study other types of intermediary. For example, a key decision for a shopping

mall is what stores it should include. If we regard each store as a product category, the problem becomes similar to product range choice and it is important to take into account the externalities each store has on other stores. Of course a shopping mall is more like a platform which does not possess the products and allows decentralized pricing. But the situation is actually similar to our model if each store in the mall sets the same price as in their own outlets. Another possible application is intermediaries in international trade.³⁹ In that case retailers in a destination country act as consumers in our model and have demand for multiple products. Direct trade can be too costly for some manufacturers and retailers, and so they choose to use trade intermediaries. The range of products an intermediary handles can be an important factor retailers care about.

Appendix

Proof of Lemma 1. (i) Consider an equilibrium in which a set A_M of products are sold only by their manufacturers, a set A_E of products are stocked exclusively by the intermediary, and a set A_{NE} of products are stocked non-exclusively by the intermediary. Let p_l be the equilibrium price of product $l \in A_M$, p_j be the equilibrium price of product $j \in A_E$, and $p_{i,M}$ and $p_{i,I}$ be the equilibrium price of product $i \in A_{NE}$ at its manufacturer and the intermediary, respectively. Note that if $p_{i,I} > p_{i,M}$ it is possible that a consumer visits the intermediary which stocks product i but buys product i from its manufacturer. However if $p_{i,I} \leq p_{i,M}$ it is impossible that in equilibrium a consumer visits both the intermediary and the manufacturer.

(i-1) As in the case of no intermediary, it is easy to see $p_l = p_l^m$ for $l \in A_M$ given our informational assumption.

(i-2) We then show $p_j = p_j^m$ for $j \in A_E$. Suppose the wholesale price of product j is τ_j . The hold-up logic implies that the intermediary must charge $p_j^*(\tau_j) = \arg \max_p (p - \tau_j) Q_j(p)$. (Note that $p_j^*(c_j) = p_j^m$.) Since the intermediary makes a take-it-or-leave-it offer, it will optimally offer a lump-sum fee $T_j = \pi_j F(v_j) - (\tau_j - c_j) Q_j(p_j^*(\tau_j)) \times \eta$ to manufacturer j , where η is the measure of consumers who visit the intermediary and which only depends on the *expected* surplus from visiting the intermediary. (In particular, given consumers do not observe the contract details, η is independent of the actual wholesale price τ_j .) Hence the intermediary's profit from stocking product j exclusively

³⁹See, e.g., Bernard et al. (2010) and Ahn et al. (2011) for empirical evidence on trade intermediaries in the US and China, respectively.

is

$$\eta \times [p_j^*(\tau_j) - \tau_j] Q_j(p_j^*(\tau_j)) - T_j = \eta \times [p_j^*(\tau_j) - c_j] Q_j(p_j^*(\tau_j)) - \pi_j F(v_j) . \quad (30)$$

This is maximized at $p_j^*(\tau_j) = p_j^m$ such that the intermediary should offer a wholesale price $\tau_j = c_j$.

(i-3) We finally show $p_{i,I} = p_{i,M} = p_i^m$ for $i \in A_{NE}$. The proof consists of a few steps.

Step 1: $p_{i,M} \leq p_i^m$.

If in contrast $p_{i,M} > p_i^m$ in equilibrium, then reducing $p_{i,M}$ slightly will be a profitable deviation. First, the number of consumers who buy product i from respectively the intermediary and manufacturer i does not change. For those consumers who visit the intermediary and buy product i there, they do not observe manufacturer i 's price reduction and so still buy from the intermediary. For those consumers who visit the intermediary first and then come to manufacturer i , they will be surprised by the price reduction but will still buy from manufacturer i as originally planned. The number of such consumers does not increase since their search decision is based on expected equilibrium prices. For those who visit manufacturer i first, their initial plan must be to buy product i at the manufacturer (otherwise they would have no reason to visit it). Again a private price reduction will not increase the number of such consumers, and once they arrive they buy as planned (given passive beliefs). Second then, manufacturer i earns strictly more profit from its direct sales to consumers, and earns the same profit from sales made through the intermediary.

Step 2: $p_{i,M} = p_i^m$.

If in contrast $p_{i,M} < p_i^m$ in equilibrium, then increasing $p_{i,M}$ slightly will be a profitable deviation. Consider the following two cases separately:

(a) $p_{i,I} > p_{i,M}$. Consider a slight increase to $p_{i,M} + \varepsilon < \min\{p_{i,I}, p_i^m\}$. For those who visit the intermediary first and then come to manufacturer i (based on the expected price), they will be surprised by manufacturer i 's price increase but will still buy from it since its price remains strictly below $p_{i,I}$. For those who visit manufacturer i first (again, based on the expected price), they will buy as planned given the new price is still lower than $p_{i,I}$. Therefore, the number of consumers who buy at manufacturer i remains unchanged, but the profit from each of them is now higher.

(b) $p_{i,I} \leq p_{i,M}$. For those who visit the intermediary first, they will not come to manufacturer i according to their beliefs, so they are irrelevant for a private price deviation. For those who plan to visit manufacturer i , they must not visit the intermediary on equilibrium path. If $p_{i,M}$ is slightly increased, will some of them switch to visiting

the intermediary? The answer is no, because in our continuum framework this single price deviation has a zero-measure impact on the consumer surplus from not visiting the intermediary and so will not change consumer search behavior.⁴⁰ Therefore again a small price increase will improve manufacturer i 's profit.

Step 3: $p_{i,I} \leq p_i^m$.

Suppose in contrast $p_{i,I} > p_i^m (= p_{i,M})$ in equilibrium. In this case, there are two possible types of consumer who buy product i . Let $\eta_{i,I}$ be the measure of consumers who buy i at the intermediary, and let $\eta_{i,M}$ be the measure of consumers who buy i at manufacturer i . (Some of the latter consumers may visit the intermediary but buy from the manufacturer.) Consider two cases separately:

(a) $\tau_i \leq c_i$. Then a small reduction of $p_{i,I}$ will be a profitable deviation. Slightly decreasing $p_{i,I}$ will weakly increase $\eta_{i,I}$. At the same time the intermediary makes a higher profit from each such consumer given that $p_i^*(\tau_i) \leq p_i^m < p_{i,I}$.

(b) $\tau_i > c_i$. In this case we argue that a deviation to $p'_{i,I} = p_i^m$ (together with an adjustment of the two-part tariff) will be profitable. In the hypothetical equilibrium, we must have

$$T_i + \eta_{i,I} \times (\tau_i - c_i)Q_i(p_{i,I}) + \eta_{i,M} \times \pi_i = \pi_i F(v_i) .$$

Then the intermediary's profit from product i is

$$\eta_{i,I} \times (p_{i,I} - \tau_i)Q_i(p_{i,I}) - T_i = \eta_{i,I} \times (p_{i,I} - c_i)Q_i(p_{i,I}) + \eta_{i,M} \times \pi_i - \pi_i F(v_i) .$$

If $p_{i,I}$ is reduced to p_i^m , $(p_{i,I} - c_i)Q_i(p_{i,I})$ will increase to π_i , the per-consumer monopoly profit, and $\eta_{i,I} + \eta_{i,M}$ will increase at least weakly.⁴¹ Then the profit must be improved.

Step 4: $p_{i,I} = p_i^m$.

Suppose in contrast $p_{i,I} < p_i^m (= p_{i,M})$ in equilibrium. Then if a consumer visits the intermediary, she will not visit manufacturer i . In this case it is then impossible that

⁴⁰With a discrete number of products, the same result holds by a slightly different argument. Consider a consumer who is *ex ante* indifferent between whether or not to visit the intermediary. If she visits manufacturer i and finds $p_{i,M}$ slightly higher than expected, will she now want to visit the intermediary? Since the cost of visiting the manufacturer is already sunk, she actually would have a strict preference for not visiting the intermediary if $p_{i,M}$ remained the same as expected. Therefore the same is true if $p_{i,M}$ is only slightly higher than she expected.

⁴¹In fact, it can be shown that $\eta_{i,I} + \eta_{i,M}$ remains unchanged. The consumers who buy product i can be divided into three groups: some don't visit the intermediary and buy i at manufacturer i ; some visit the intermediary but buy i at manufacturer i ; the rest visit the intermediary and buy i there. The deviation does not affect the first group. The deviation may affect the distribution of consumers between the second and the third group, but does not affect the total number of consumers who visit the intermediary which only depends on the expected prices.

$\tau_i \geq c_i$. Otherwise the intermediary could improve its profit from product i by raising $p_{i,I}$ slightly. (Note that this deviation does not affect the number of consumers who visit the intermediary, and once they arrive they will still buy product i at the intermediary as long as $p_{i,I}$ is still below $p_{i,M}$.)

Now consider the possibility of $\tau_i < c_i$. Then we must have $p_{i,I} = p_i^*(\tau_i)$ in an equilibrium. Then a deviation to $\tau'_i = c_i$ and $p'_{i,I} = p_i^m$ will be profitable. (Given the contract details are unobservable to consumers, such a deviation will not affect the number of consumers who visit the intermediary and buy i .)

This completes the proof for $p_{i,I} = p_{i,M} = p_i^m$ for $i \in A_{NE}$.

(ii) The equilibrium two-part tariff for product $j \in A_E$ has been proved in (i-2) above. Now consider the equilibrium two-part tariff for product $i \in A_{NE}$. It is easy to see that $\tau_i < c_i$ is impossible. Otherwise the intermediary would have an incentive to reduce its price for product i to $p_i^*(\tau_i)$. However, we cannot rule out the possibility of $\tau_i > c_i$ (together with T_i such that manufacturer i 's profit is $\pi_i F(v_i)$). The reason is that if the intermediary raises its price for product i above p_i^m , some consumers who visit the intermediary and initially planned to buy i there may then switch to buying from manufacturer i . If the number of such consumers is large enough (which requires $f(s)$ to be large enough for small s), the intermediary does not dare to raise its price.

Fortunately, this indeterminacy of the contract details does not matter for our subsequent analysis of optimal product selection. Suppose in an equilibrium $\tau_i \neq c_i$ for some $i \in A_{NE}$. The lump-sum fee T_i satisfies

$$T_i + \eta_{i,I} \times (\tau_i - c_i) Q_i(p_{i,I}) + \eta_{i,M} \times \pi_i = \pi_i F(v_i) .$$

Note that given the monopoly pricing result, $\eta_{i,I}$ is also the number of consumers who visit the intermediary which is denoted by η_I . Then the intermediary's profit from stocking product i is

$$\begin{aligned} \eta_I \times (p_i^m - \tau_i) Q_i(p_i^m) - T_i &= \eta_I \times (p_i^m - c_i) Q_i(p_i^m) + \eta_{i,M} \times \pi_i - \pi_i F(v_i) \\ &= \pi_i [\eta_I - (F(v_i) - \eta_{i,M})] . \end{aligned} \tag{31}$$

Since consumer search and purchase behavior only depends on the retail prices, this profit is the same as if $\tau_i = c_i$. Therefore, without loss of generality, we can focus on a contracting outcome with $\tau_i = c_i$. ■

Proof of Lemma 2. (i) We first show that the intermediary can make a positive profit by stocking a positive measure of products. Consider two interior points in Ω : (π_1, \tilde{v})

and (π_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

$$\int_{\pi_1 - \delta}^{\pi_1} g(\pi, 2\tilde{v} - v) d\pi = \int_{\pi_2}^{\pi_2 + \Delta(v)} g(\pi, v) d\pi . \quad (32)$$

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,

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

where the second equality used (32) evaluated at $v = \tilde{v}$. Therefore, $\Pi(\epsilon) > 0$ for ϵ in a neighborhood of 0.

(ii) We then show that stocking all the products is not the most profitable strategy. Let $\hat{v} = \int_{\Omega} v dG$. 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

$$\int_{\pi_1 - \delta}^{\pi_1} g(\pi, 2\hat{v} - v) d\pi = \int_{\pi_2}^{\pi_2 + \Delta(v)} g(\pi, v) d\pi . \quad (33)$$

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 \hat{v} , and so the average v in $A = \Omega \setminus (B_1 \cup B_2)$ is \hat{v} as well. Then a consumer will visit the intermediary, when it stocks $A = \Omega \setminus (B_1 \cup B_2)$, 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 $A = \Omega \setminus (B_1 \cup B_2)$ 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 (33) 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 (8) and (9) have a solution with $k \in (\underline{v}, \bar{v})$.⁴² Let $\phi(v; k) \equiv \frac{k-v}{F(k)-F(v)}$.

We first claim that for any $k \in (\underline{v}, \bar{v})$, (9) has a unique solution

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

and $\lambda'(k) \in (0, \infty)$. The proof is as follows. The left-hand side of (9) 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 (9) is strictly positive when $\lambda \min_v \phi(v; k) \geq \bar{\pi}$ and the reasoning is the same. The left-hand side of (9) 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 $\lambda(\underline{v}) = \lim_{k \rightarrow \underline{v}} \lambda(k)$ and $\lambda(\bar{v}) = \lim_{k \rightarrow \bar{v}} \lambda(k)$. We must have $\lambda(\underline{v})\phi(v; \underline{v}) \leq \underline{\pi}$ and $\lambda(\bar{v})\phi(v; \bar{v}) \geq \bar{\pi}$ for any v (or except for a zero-measure set). Notice also that the left-hand side of (9) 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 (8) with λ replaced by $\lambda(k)$:

$$\int_{I(k, \lambda(k))} (f(k)\pi - \lambda(k)) dG = 0 . \quad (34)$$

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

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

⁴²In numerical examples we find that the system has a unique solution with $k \in (\underline{v}, \bar{v})$, though we have been unable to formally prove uniqueness.

When $k = \underline{v}$ or \bar{v} , $I(k, \lambda(k))$ is an empty set and so $\Phi(\underline{v}) = \Phi(\bar{v}) = 0$. According to the construction of $I(k, \lambda)$ and the definition of $\lambda(k)$, $\Phi(k) > 0$ for $k \in (\underline{v}, \bar{v})$. Therefore by the mean-value theorem $\Phi'(k) = 0$ must have a solution in (\underline{v}, \bar{v}) . On the other hand, one can verify that $\Phi'(k)$ equals the left-hand side of (34) by using the definition of $\lambda(k)$ and the construction of $I(k, \lambda)$. Then (34) must have a solution $k \in (\underline{v}, \bar{v})$. ■

Proof of Lemma 3. The difference in payoff between (10) and (11) is

$$\Delta(s) = \int qvdG - h\left(\int qdG\right)s - \int_{v>s} q_{NE}(v-s)dG. \quad (35)$$

(We have used $q - q_E = q_{NE}$.) Notice that $\Delta(0) \geq 0$, and $\Delta(s)$ is weakly concave because

$$\Delta'(s) = -h\left(\int qdG\right) + \int_{v>s} q_{NE}dG$$

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 $\Delta(k) = 0$. (12) is just a rewriting of $\Delta(k) = 0$. In this case, $k < \bar{v}$ if and only if $\Delta(\bar{v}) < 0$ which equals $\int qvdG - h(\int qdG)\bar{v} < 0$. ■

Proof of Lemma 4. When $h(m) = m$ for all $m \in [0, \bar{m}]$, by a similar argument as in the simple case we can show that the intermediary can make a strictly positive profit by stocking some products exclusively. (Note that the set of exclusive products constructed in Lemma 2 can be arbitrarily small.)

Now consider the case of $h(m) < m$ for some $m \in (0, \bar{m}]$.⁴³ 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 $\min_{v \in A} v = \underline{v}$).

⁴³In this case, it is possible that $h(0) > 0$. Then the approach in Lemma 2 does not apply because $k \rightarrow 0$ when the measure of stocked products goes to 0. That is why we adopt a different approach.

Suppose now the intermediary stocks all products in A *non-exclusively* (i.e., $q = q_{NE} = 1$ only for $(\pi, v) \in A$). Then from (35) we can see $\Delta(0) = 0$, and

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

for all $s \in [0, \underline{v}]$. This implies $k > \underline{v}$. From (13), it is ready to see that the intermediary's profit is $\int_{A \cap \{v < k\}} \pi [F(k) - F(v)] dG > 0$. ■

Proof of Lemma 5. (a) Suppose $k < \bar{v}$ so that there are products with $v > k$. Suppose in contrast that in the optimal solution \mathbf{q} , $q = 0$ for a strictly positive measure of products with $v > k$. Denote this set of products by B . Consider a new stocking policy $\tilde{\mathbf{q}}$ such that

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

and $\tilde{q}_E = q_E$. (That is, the products in B are now stocked non-exclusively.) Let \tilde{k} be the new consumer search threshold associated with $\tilde{\mathbf{q}}$. We aim to show that this new stocking policy is more profitable than \mathbf{q} and so a contradiction arises. We can see from (13) that this is true if $\tilde{k} \geq k$, or equivalently if $\tilde{\Delta}(k) \geq \Delta(k)$, where $\tilde{\Delta}(\cdot)$ is (35) associated with the new stocking policy. Using the construction of $\tilde{\mathbf{q}}$ and the definition of $\Delta(\cdot)$ in (35), one can check that

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

Since $\int \tilde{q} dG - \int q dG = \int_B (1 - q) dG$ and $h'(m) \leq 1$ for all m , we have $\tilde{\Delta}(k) - \Delta(k) \geq 0$. Therefore, the proposed new stocking policy is a profitable deviation, and so in the optimal solution we must have $q = 1$ for all $v > k$.

We now prove the second part in result (a). Suppose in contrast that in the optimal solution \mathbf{q} , there is a strictly positive measure of $v > k$ such that for each of these v , there exist $\pi' > \pi''$ such that $q_E(\pi', v) = 1$ and $q_{NE}(\pi'', v) = 1$ (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 policy $\tilde{\mathbf{q}}$ with

$$\tilde{q}_E(\pi, v) = 1 \text{ if } \pi \leq \tilde{\pi}(v) \text{ and } \tilde{q}_{NE}(\pi, v) = 1 \text{ if } \pi > \tilde{\pi}(v) ,$$

where $\tilde{\pi}(v)$ is the unique solution to $\int_{\tilde{\pi}(v)}^{\tilde{\pi}(v)} g(\pi, v) d\pi = \int_{\tilde{\pi}(v)}^{\tilde{\pi}(v)} q_E(\pi, v) g(\pi, v) d\pi$. (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

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

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.

(b) Suppose that $k > \underline{v}$ so that there are products with $v < k$. Suppose in contrast that in the optimal solution \mathbf{q} , 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 $q(\pi'', v) = 0$ (i.e., some low- π products are stocked while some high- π products are not). 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 policy with

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

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

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

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. We first consider the case where $m < 1$ in the optimal solution. Then the first-order condition with respect to m is $\mu = \lambda k h'(m)$. We use this to replace μ in our analysis. The first-order condition with respect to k yields (19). The other two equations (18) and (20) are simply the two constraints. Both k and λ are positive. From (15) 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 exclusivity arrangement does not matter. For $v > k$, notice that $q_E = 1$ and $q_{NE} = 1$ are mutually exclusive, and $\lambda k - \mu = \lambda k(1 - h'(m)) \geq 0$. Then we deduce that

$q_E = 1$ if

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

and $q_{NE} = 1$ otherwise.

The case with $m = 1$ (so $q = 1$ everywhere) is simple. For $v < k$, again the exclusivity arrangement does not matter. For $v > k$, the optimal exclusivity is determined the same as above. ■

Proof of Proposition 3. From (21) it is ready to see that for $v < k$, $q = 1$ if and only if

$$\pi \geq \frac{\mu - \lambda v}{F(k) - F(v)},$$

and the exclusivity arrangement does not matter. Now consider $v > k$ if $k < \bar{v}$ in the optimal solution:

(i) If $\lambda k - \mu > 0$ in the optimal solution, q_E and q_{NE} are determined exactly the same as in the case with an unlimited stocking space. That is, $q_E = 1$ if

$$\pi \leq \frac{\lambda(k - v)}{F(k) - F(v)}, \quad (36)$$

and $q_{NE} = 1$ otherwise.

(ii) If $\lambda k - \mu = 0$ in the optimal solution, $q_E = 1$ if $\pi(F(k) - F(v)) + \lambda v - \mu \geq 0$, or equivalently if (36) holds as $\mu = \lambda k$. The intermediary is indifferent in how to select the non-exclusive products among the others, as long as the mass of them can satisfy $\lambda k - \mu = 0$.

(iii) If $\lambda k - \mu < 0$ in the optimal solution, it is clear that $q_{NE} = 0$, and $q_E = 1$ if $\pi(F(k) - F(v)) + \lambda v - \mu \geq 0$, or equivalently if

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

Finally, the parameters μ , λ and k solve the system of the k constraint, the first-order condition with respect to k , and the space constraint. That is just (18)-(20) with m replaced by \bar{m} . (In the case of $\lambda k - \mu = 0$, we have an additional equation, but in that case to pin down the region of non-exclusive products with $v > k$, we also have another parameter to determine if the region is parameterized by one parameter like we will do in the numerical example below.) ■

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 \mathbf{q} , $q = 0$ for a strictly positive measure of products with $v > k$. Denote this set of products by B . Consider a new stocking policy $\tilde{\mathbf{q}}$ such that

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

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

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

which is weakly positive since $\int \tilde{q} dG - \int q dG = \int_B (1 - q) dG$ and $h'(m) \leq 1$ for all m . Since all parties weakly benefit from $\tilde{\mathbf{q}}$ we have a contradiction.

We now prove the second part in result (a). Suppose in contrast that in the optimal solution \mathbf{q} , there is a strictly positive measure of $v > k$ such that for each of these v , there exist $\pi' > \pi''$ such that $q_E(\pi', v) = 1$ and $q_{NE}(\pi'', v) = 1$. 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 policy $\tilde{\mathbf{q}}$ with

$$\tilde{q}_E(\pi, v) = 1 \text{ if } \pi \leq \tilde{\pi}(v) \text{ and } \tilde{q}_{NE}(\pi, v) = 1 \text{ if } \pi > \tilde{\pi}(v),$$

where $\tilde{\pi}(v)$ is the unique solution to $\int_{\tilde{\pi}(v)}^{\tilde{\pi}(v)} g(\pi, v) d\pi = \int_{\underline{\pi}(v)}^{\tilde{\pi}(v)} q_E(\pi, v) g(\pi, v) d\pi$. By construction $u^0(s, \cdot)$ and $u^1(s, \cdot)$ are unchanged since consumers only care about v , hence consumer surplus is unchanged. However in the proof of Lemma 5 we showed that the intermediary's profit is higher under $\tilde{\mathbf{q}}$, 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 \mathbf{q} , 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 $\int_{\tilde{\pi}(v)}^{\tilde{\pi}(v)} g(\pi, v) d\pi = \int_{\tilde{\pi}(v)}^{\tilde{\pi}(v)} q(\pi, v) g(\pi, v) d\pi$. Similarly as before, $u^0(s, \cdot)$ and $u^1(s, \cdot)$ are unchanged hence consumer surplus is unchanged. However in the proof of Lemma 5 we showed that the intermediary's profit is higher under $\tilde{\mathbf{q}}$, hence we have a contradiction. ■

Proof of Proposition 4. Substituting the expressions for $\Pi(\mathbf{q})$, $u^1(s, \mathbf{q})$ and $u^0(s, \mathbf{q})$ into equation (22) yields

$$W(\mathbf{q}) = \int (1-q) \int_0^v (\pi+v-s) dF(s) dG + \int q(\pi+v) F(k) dG - h(m) \int_0^k s dF(s) + \int_{v>k} q_{NE} \int_k^v (\pi+v-s) dF(s) dG. \quad (37)$$

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.

Maximizing $W(\mathbf{q})$ is the same as maximizing $W(\mathbf{q}) - W(\mathbf{0})$, the welfare improvement by the intermediary, where $W(\mathbf{0}) = \int \int_0^v (\pi+v-s) dF(s) dG$ is the total welfare when there is no intermediary. After some algebraic manipulations, we can write the Lagrange function as follows:

$$\begin{aligned} \mathcal{L} = & \int_{v<k} q \underbrace{[(\pi+v)[F(k)-F(v)] + \lambda v - \mu + \int_0^v s dF(s)]}_{[1]} dG \\ & + \int_{v>k} \underbrace{\{q_E[(\pi+v)[F(k)-F(v)] + \lambda v - \mu + \int_0^v s dF(s)] + q_{NE}[\lambda k - \mu + \int_0^k s dF(s)]\}}_{[2]} dG \\ & - h(m) \underbrace{[\lambda k + \int_0^k s dF(s)]}_{[3]} + \mu m. \end{aligned}$$

Consider first the case of $m < 1$ in the optimal solution. The first-order condition with respect to m yields $\mu = h'(m)[\lambda k + \int_0^k s dF(s)]$. Proceeding as in the intermediary's problem, we can see that $q = 1$ for $v < k$ if and only if $[1] \geq 0$. Using $F(k) - F(v) = \int_v^k dF(s)$ and the above expression for μ , one can verify that this is equivalent to (23). The exclusivity arrangement does not matter. For $v > k$, notice that $q_E = 1$ and $q_{NE} = 1$ are mutually exclusive, and $[3] \geq 0$ given $h'(m) \leq 1$. Then $q_E = 1$ if $[2] \geq [3]$, which is equivalent to (24), and $q_{NE} = 1$ otherwise. 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 (18) - (20).

The case with $m = 1$ is simple. For $v > k$, the conditions for $q_E = 1$ and $q_{NE} = 1$ remain the same. ■

Proof of Lemma 7. Let (i, j) denote the intermediary’s stocking policy for manufacturer i and manufacturer j of a product. Let $i, j \in \{\phi, E, NE\}$, where ϕ means the intermediary does not stock the manufacturer’s product, E means it contracts with the manufacturer exclusively, and NE means it contracts with the manufacturer non-exclusively.

In the main text we already proved results (i) and (ii) which specify the compensation in the cases of (NE, NE) and (E, E) , respectively. We now derive compensation in the remaining cases and then prove result (iii). First, consider (ϕ, E) and (E, ϕ) . If the relevant manufacturer accepts the exclusive contract, it earns 0 directly from consumers, whereas if it rejects it earns $\frac{\pi}{2}F(v)$ since consumers with $s < v$ will randomly pick one manufacturer to visit. Hence, the manufacturer needs to be compensated by $\frac{\pi}{2}F(v)$. Second, consider (ϕ, NE) and (NE, ϕ) . If the relevant manufacturer accepts the non-exclusive contract, it earns $\max\{0, \frac{1}{2}\pi[F(v) - F(k)]\}$ directly from consumers, whereas if it rejects it earns $\frac{1}{2}\pi F(v)$. Therefore, it needs to be compensated by $\frac{1}{2}\pi F(\min\{v, k\})$. Third, consider (E, NE) . If manufacturer i accepts the exclusive contract, it earns 0 directly from consumers, whereas if it rejects it earns $\max\{0, \frac{1}{2}\pi[F(v) - F(k)]\}$. Manufacturer j earns $\max\{0, \pi[F(v) - F(k)]\}$ regardless of whether it accepts or rejects the non-exclusive contract. Hence, total compensation is $\max\{0, \frac{1}{2}\pi[F(v) - F(k)]\}$. (The same is true for the case of (NE, E) .) Notice that in all three cases, compensation is (weakly) higher than that in the case of (NE, NE) , but the effect on the search constraint (26) is the same because all the options lead to $q = 1$ and $q_E = 0$. As a result, these cases are dominated by (NE, NE) . (When $v < k$, (E, NE) and (NE, E) are weakly dominated.) ■

References

- ADAMS, W., AND J. YELLEN (1976): “Commodity Bundling and the Burden of Monopoly,” *Quarterly Journal of Economics*, 90(3), 475–498.
- AHN, J., A. KHANDELWAL, AND S. WEI (2011): “The Role of Intermediaries in Facilitating Trade,” *Journal of International Economics*, 84(1), 73–85.

- ANDERSON, S., AND R. RENAULT (2003): “Efficiency and Surplus Bounds in Cournot Competition,” *Journal of Economic Theory*, 113(2), 253–264.
- (2006): “Advertising Content,” *American Economic Review*, 96(1), 93–113.
- ASKER, J., AND H. BAR-ISAAC (2016): “Vertical Information Restraints: Pro- and Anti-Competitive Impacts of Minimum Advertised Price Restrictions,” mimeo.
- BERNARD, A., B. JENSEN, S. REDDING, AND P. SCHOTT (2010): “Wholesalers and Retailers in US Trade,” *American Economic Review*, 100(2), 408–413.
- BIGLAISER, G. (1993): “Middlemen as Experts,” *RAND Journal of Economics*, 24(2), 212–223.
- BRONNENBERG, B. (2017): “Innovation and Distribution: A General Equilibrium Model of Manufacturing and Retailing,” mimeo.
- CHEN, Y., AND M. RIORDAN (2013): “Profitability of Product Bundling,” *International Economic Review*, 54(1), 35–57.
- 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.
- KAPLAN, G., G. MENZIO, L. RUDANKO, AND N. TRACHTER (2016): “Relative Price Dispersion: Evidence and Theory,” mimeo.
- KOK, G., M. FISHER, AND R. VAIDYANATHAN (2015): “Assortment Planning: Review of Literature and Industry Practice,” in *Retail Supply Chain Management*, ed. by N. Agrawal, and S. Smith. Springer.
- 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.
- STIGLITZ, J. (1979): “Equilibrium in Product Markets with Imperfect Information,” *American Economic Review*, 69(2), 339–345.
- 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.