Multiproduct Intermediaries*

Andrew Rhodes
Toulouse School of Economics

Makoto Watanabe
VU University of Amsterdam

Jidong Zhou
Yale School of Management

May 2018

Abstract

This paper develops a new framework for studying multiproduct intermediaries. We show that a multiproduct intermediary is profitable even when it does not improve efficiency in selling products. In its optimal product selection, it stocks high-value products exclusively to attract consumers, then profits by selling non-exclusive products which are relatively cheap to buy from upstream suppliers. However, relative to the social optimum, the intermediary tends to be too big and stock too many products exclusively. We establish a link between product selection and product demand features such as size, shape and elasticity. As an application of the framework, we also study the impact of direct-to-consumer sales by upstream suppliers on the intermediary's product range and profitability.

Keywords: intermediaries, multiproduct demand, search, product range, exclusive contracts, direct-to-consumer sales

JEL classification: D83, L42, L81

*We are grateful for helpful comments to Mark Armstrong, Heski Bar-Isaac, Alessandro Bonatti, Joyee Deb, Paul Ellickson, Doh-Shin Jeon, Bruno Jullien, Fei Li, Barry Nalebuff, Martin Obradovits, Jerome Renault, Patrick Rey, Mike Riordan, Greg Shaffer, Andy Skrzypacz, Greg Taylor, Raphael Thomadsen, Glen Weyl, Mike Whinston, Chris Wilson, Julian Wright and seminar participants in Bonn, MIT, MSU, NUS, Oxford, Stanford, Tokyo, TSE, UCLA Anderson, 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), TNIT (Microsoft), and the 16th Annual Columbia/Duke/MIT/Northwestern IO Theory Conference.

1 Introduction

Intermediaries are important players in the economy, and according to some estimates are responsible for 34% of US GDP.[1] Many intermediaries carry multiple different products, and serve buyers with multiproduct demand. Examples include retailers such as supermarkets and department stores, shopping malls, TV platforms, travel agencies, and trade intermediaries. However much of the existing literature focuses on single-product intermediaries. Our paper builds a framework to study multiproduct intermediaries when consumers demand multiple products, and uses it to address some new and important questions. For example, in what ways can a multiproduct intermediary create value and therefore profitably exist? A multiproduct intermediary's profit depends crucially on the products it stocks. Which products should it carry, and for which of them should it be the exclusive supplier in the market? Is the intermediary too big or too small relative to the social optimum, and does it carry qualitatively the 'right' products? Further, it is increasingly easy for sellers to bypass traditional intermediaries and sell direct to buyers. How should intermediaries change the products they carry in order to deal with this threat?

There is surprisingly little research about which products an intermediary should carry, even though in practice it is one of the most important decisions they have to take. For example, in the case of retailers, product range is particularly important because consumers usually want to buy several different products but find it costly to shop around, and so they prefer retailers whose product ranges closely match their needs. At the same time retailers are often constrained in how many products they can stock, for example due to limited stocking space or the fact that stocking too many products can make the in-store shopping experience less pleasant.[2] Consequently retailers (and other intermediaries) must carefully choose their product range.[3]

At the same time, changes in technology are forcing intermediaries to rethink which products they carry. Returning again to the retail example, traditionally most manufacturers could only reach consumers via retailers. However in recent years this has changed. Many established brands like Nike, Nestle, and Proctor & Gamble now sell direct to consumers (DTC) through their own physical store or website.[4] Marketplaces such as

[1]See Krakovsky's book entitled "The Middleman Economy".

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

[3]For instance, a Harvard Business Review article states that: "Succeeding in the retail business means being good on a number of dimensions [including pricing and location] ... But assortment is number one." (See goo.gl/3STUFa.)

[4]By 2020 Nike aims to generate one third of its sales (worth about $16billion a year) via the di-

Amazon and Tmall enable smaller brands to do the same.[5] As a result, Forbes estimated that in 2016 more than 40% of US manufacturers would sell direct, an increase of 70% over the previous year. (A similar trend is emerging in other markets, such as those for TV platforms and travel agencies.) Many commentators have argued that direct sales weaken retailers' bargaining power and threaten their whole one-stop shopping business model.[6] One way retailers have responded to this threat is by streamlining their product range and stocking more exclusive products that are not available for purchase elsewhere. For example Macy's has signed exclusive deals with several brands, and by 2020 aims to have 40% of its products being unique to its stores.[7]

Our paper develops a model which seeks to capture many of these issues. The model is presented in Section 2. It is mainly developed for retailers but, as we discuss later, the setup and some of the main insights also apply to a much broader set of multiproduct intermediaries including shopping malls, TV platforms, and trade intermediaries. There is a unit mass of manufacturers, each of which produces a different product. Consumers view these products as independent and are interested in buying all of them, although different products have different demands. There is also a single multiproduct intermediary. A product can be sold to consumers either directly by the manufacturer, or by the intermediary, or through both these channels. The intermediary offers to compensate manufacturers in exchange for the right to stock their products, and as part of this can request exclusive sales rights. We allow for the possibility that the intermediary 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; in the retail example, this is consistent with the idea that larger retailers are located further from consumers, or offer a worse instore shopping experience. Consumers differ in their search cost, which generates heterogeneity in the products that they end up buying.

Since the focus of our paper is product range choice, we intentionally simplify sellers'

---

rect channel (see goo.gl/Y6e3w9). Nestle's direct to consumer brand Nespresso generates an estimated \$4.5billion in sales each year (see goo.gl/GrkEB6). Proctor & Gamble sell Tide Pods direct to consumers.

[5] Marketplaces provide logistical support for payments, packaging and delivery (see goo.gl/arSbuE). In 2016 half of the items that Amazon shipped were for third-party sellers (see goo.gl/z3tHjx).

[6] Some have even linked the demise of certain retail chains (and high street shopping more broadly) to DTC sales. See e.g. goo.gl/FnAqR2 and goo.gl/MkWk6j concerning the failure of Sports Authority.

[7] See goo.gl/46xgYA and goo.gl/mBTHNN 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.

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.[8] We then argue that all information on a product's cost and demand curve can be summarized using a simple two-dimensional statistic (π, v), where π represents a product's monopoly profit and v represents its monopoly consumer surplus. This enables us to study product range choice in a tractable way, since it reduces a potentially complicated product space into a simpler two-dimensional one. Specifically, the intermediary's problem is to choose a set of points within the (π, v) space that it will stock exclusively, and another set of points which it will stock non-exclusively.

In Sections 3 and 4 we solve for the intermediary's optimal product range, first in a special case to provide some initial intuition, and then in the general case. Unlike the standard single-product case, we show that a multiproduct intermediary can earn strictly positive profit even when it is no more efficient at selling products than the direct channel. We also show that the value from stocking a product can be split into a 'direct' and 'indirect' component. The direct component represents the profit earned through selling that product. It includes compensation paid to the manufacturer and can therefore be negative. The indirect component reflects cross-product externalities i.e. the way in which stocking a new product may either increase or decrease how many consumers search the intermediary, and thereby change the profitability of the intermediary's other products.[9]

The optimal stocking policy itself consists of three distinct 'regions' in the (π, v) space. Firstly, the intermediary stocks some products with high-v but low-π exclusively. These products make a direct loss but they help attract consumers due to their high v. Secondly, the intermediary recoups these losses by stocking some other very profitable products with low-v but high-π. However since these products have low v they may dissuade consumers from searching, and so in general the intermediary avoids stocking too many of them. Thirdly, depending on the intermediary's stocking constraint and search technology, it may also stock some products with high-v and high-π non-exclusively. These products break even but can have a positive indirect effect.

Section 5 of the paper then uses these results to understand the effects of DTC sales. To do this, we first solve for the intermediary's optimal stocking policy when DTC sales are

[8] Intuitively, with two-part tariffs the intermediary buys marginal units from a manufacturer at cost. The logic of Diamond (1971) then implies that with search frictions there is no price competition even if a product is sold by both its upstream supplier and the intermediary. As we discuss more in Section 2, the monopoly pricing itself is not crucial for our analysis - what matters is that the price of each product is the same across sellers.

[9] This indirect effect seems important in practice - see goo.gl/toxMuc for examples where dropping a product with low direct profit led to a sharp decline in sales of other (directly profitable) goods.

not feasible. Consistent with our earlier discussion, we find that manufacturers' ability to sell direct weakens the intermediary's bargaining power and decreases its profit. Although the intermediary can recoup some of this lost profit by reoptimizing its product range, it may have to significantly alter the products that it sells. Moreover, we show by example that failure to adjust its product range can lead the intermediary to earn negative profit and so exit the market altogether. The final part of Section 5 shows how our model applies to a different type of intermediary, namely a shopping mall which acts as a platform and does not directly resell products. The model provides insights about which retailers should join the mall and how much they should pay to do so, as well as the externalities they exert on other retailers, all of which are consistent with empirical evidence.

The remainder of the paper is then structured as follows. In Section 6 we show how to construct the (π, v) space and interpret different points within it. For example we show how the size, shape, and elasticity of a product's demand affect whether the retailer stocks it and whether the contract is an exclusive one. Section 7 then solves for the socially optimal stocking policy. Here we show that the intermediary tends to stock too many products, and to stock too many of them exclusively, but that nevertheless it can be welfare-enhancing due to the way it affects consumers' incentives to search. We also show that our main insights are qualitatively robust when there is upstream competition. Finally Section 8 concludes and discusses the potential application of our framework to other types of intermediary.

1.1 Related literature

There is already a substantial body of literature on intermediaries (see e.g. the book by Spulber (1999)). An intermediary may exist because it improves the search and matching efficiency between buyers and sellers (e.g. Rubinstein and Wolinsky (1987), Gehrig (1993), Spulber (1996), and Shevchenko (2004)), or because it acts as an expert or certifier that mitigates the asymmetric information problem between buyers and sellers (e.g. Biglaiser (1993), Lizzeri (1999), and Biglaiser and Li (2018)).[10] We study intermediaries in an environment with search frictions, but in our model an intermediary can profitably exist even if it does not improve search efficiency. This relies on consumers demanding multiple different products and the possibility of using exclusive contracts.[11] These two features

[10] Other reasons why retailers in particular may exist are i) they know more about consumer demand than manufacturers do, ii) they can internalize pricing externalities when products are substitutes or complements, and iii) they may be more efficient in marketing activities due to economies of scale.

[11] In Spiegler (2000) two agents create surplus when they interact. A third party which does not improve efficiency can extract this surplus through "exclusive-interaction" contracts, which force the agents into a Prisoner's Dilemma. Our paper studies a very different type of exclusivity arrangement.

distinguish our model from existing work on intermediaries. Since our main focus is the retail market, we also study optimal product range and the impact of DTC sales, neither of which are typically addressed by the intermediary literature.

The mechanism by which an intermediary makes profit from stocking multiple products is reminiscent of bundling (e.g. Stigler (1968), Adams and Yellen (1976), McAfee et al (1989), and Chen and Riordan (2013)). By stocking some products that consumers value highly but are not available elsewhere, the intermediary forces consumers to visit and buy other low-value (but fairly profitable) products as well which consumers would otherwise not buy. However since our paper focuses on product selection, it is more related to the question of which products a firm should bundle, something which is rarely discussed in the bundling literature. Rayo and Segal (2010) use the same bundling argument in a different setting with information design. They show that an information platform often prefers partial information disclosure, in the sense of pooling two negatively correlated prospects into one signal. (For example a search engine may pool a high-relevance but low-profit ad with a low-relevance but high-profit ad.) Unlike us, they consider a discrete framework and (more importantly) they allow the information platform to send an arbitrary number of signals (which in our framework, would be like allowing the intermediary to organize and sell non-overlapping products in multiple stores). Consequently their optimization problem is very different from ours. Moreover many other important features of our model, such as the role of exclusivity and the importance of search economies, have no counterpart in their paper (or in the wider bundling literature).

Our paper is also related to the growing literature on multiproduct search (e.g. McAfee (1995), Shelegia (2012), 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 range choice, another important decision for multiproduct retailers.[12] 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 (2017), though those works consider single-product search and address very different economic questions.

Finally, our paper is also related to research on product assortment in operations research and marketing (see e.g. the survey by Kök et al (2015)). Typically this literature focuses on a situation where consumers demand a single product, and studies the optimal number of (symmetric) varieties of that product to stock. Our paper focuses instead on a

[12] Rhodes and Zhou (2017) also study retailers' endogenous product range, but they consider a stylized merger setup with only two products.

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 usually considered in the above mentioned literature.[13]

2 The Model

There is a continuum of manufacturers with measure one, and each produces a different product. Manufacturer i has 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, and 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

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

Manufacturers can sell their product direct to consumers, for example via their own retail outlet.[14] (We consider the case where DTC sales are infeasible in Section 5.1.) 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. 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.[15] 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

[13] Bronnenberg (2017) considers a model where consumers like variety but dislike shopping, so retailers stock multiple varieties to reduce consumers' shopping costs. His model is otherwise very different from ours, as are the questions it studies. For example in his model varieties are symmetric, and so he does not look at the optimal composition of a retailer's product line.

[14] Alternatively, we can interpret direct sales as a manufacturer selling through an independent single-product retailer. If the manufacturer can make a take-it-or-leave-it offer to the independent retailer, all our analysis and results remain unchanged.

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

non-exclusively.[16] Manufacturers who received an offer then simultaneously accept or reject.

Consumers know where each product is available, but do not observe the terms of any upstream contracts. Moreover consumers cannot observe a firm's price(s) or buy its product(s) without incurring a search cost.[17] Consumers differ in terms of their 'unit' search cost s, which is distributed in the population according to a cumulative distribution function $F(s)$ with support $(0, \bar{s}]$. The corresponding density function $f(s)$ is everywhere differentiable, strictly positive, and uniformly bounded with $\max_s f(s) < \infty$. If a consumer searches a measure n of manufacturers, she incurs a total search cost $n \times s$. If a consumer also searches the intermediary, and the intermediary stocks a measure m of products, she incurs an additional search cost $h(m) \times s$. We can thus interpret s as the opportunity cost of time, with the time needed to buy from a manufacturer normalized to 1, and the time needed to visit the intermediary equal to $h(m)$.[18] We assume that the function $h(m)$ is positive and weakly increasing, reflecting the idea that larger stores may take longer to navigate, and may also be located further out of town. (Notice that we allow for the case where $h(m)$ is constant and so independent of the intermediary's size. Below we provide a microfoundation for why $h(m)$ might be strictly increasing.) We also introduce the following notation: when $h(m)/m < 1$ the intermediary generates economies of search, and when $h(m)/m > 1$ it generates diseconomies of search. Finally as is standard, we assume that after searching consumers may costlessly recall past offers.

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. 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. The intermediary uses 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.

[16] This 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.

[17] Our 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.

[18] More generally, the time needed to buy from a manufacturer can vary across products. One possible way to deal with that case is to use a triplet (π, v, θ) to characterize a product where θ is the amount of time needed for a consumer to buy from its manufacturer. Intuitively, all else equal a product with higher θ is more likely to be carried by the intermediary.

2.1 Preliminary analysis

We start with the following useful result. (All omitted proofs are 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 that a product can be sold in three different ways. Firstly product i may be sold only by its manufacturer. Since consumers only learn the manufacturer's price after they have sunk their search cost, it is optimal for the manufacturer to charge the monopoly price p_i^m.[19] This is just the standard hold-up problem which arises in search models (see e.g. Stiglitz (1979) and Anderson and Renault (2006)). Consumers therefore rationally anticipate monopoly pricing and so, using the notation introduced in (1), search if and only if $s \leq v_i$. Consequently the manufacturer's equilibrium profit is $\pi_i F(v_i)$. Notice that this is also the manufacturer's outside option if the intermediary makes it an offer.

Continuing with the intuition for Lemma 1, 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 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

[19] As is usual in search models, there also exist other equilibria in which consumers do not search (some) sellers 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.

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 $Q_i(p_i)$ and cost c_i). This helps convert the potentially complicated product space into a two-dimensional one. Henceforth 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 $\overline{v} < \infty$ be the lower and the upper bound of v. For each $v \in [\underline{v}, \overline{v}]$, there exist $0 \leq \underline{\pi}(v) \leq \overline{\pi}(v) < \infty$ such that $\pi \in [\underline{\pi}(v), \overline{\pi}(v)]$. (In section 6 we provide examples of demand functions which can generate this type of product space.) Let $(\Omega, \mathcal{F}, G)$ be a probability measure space where $\mathcal{F}$ is a σ-field which is the set of all measurable subsets of Ω according to measure G. (In particular, $G(\Omega) = 1$.) 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 $\overline{v} \leq \overline{s}$.

Discussion. Before solving for the intermediary's optimal product range, we first 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.[20]

(ii) Homogeneous consumer demand. Consumers are assumed to demand all products. However in reality some consumers want to buy more products than others (and similarly some products are needed more often than others). We could explicitly add this heterogeneity into the model, but the analysis would be less tractable. Moreover our model already generates demand heterogeneity in the sense that consumers with a lower search cost will end up buying more products.

(iii) An instore-search microfoundation for $h(m)$. Suppose there is a fixed cost $\kappa_0 s$ of traveling to the intermediary, and also a cost $\kappa_1 s$ to search each product in the store. Also suppose the intermediary can influence the consumer search order e.g. via where it places different products within the store. We can show that by forcing consumers to search exclusive products with high v last, the intermediary can induce every consumer

[20] See footnote 27 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.

who visits to search all its products.[21] Consumers anticipate this, and so the total cost of searching the intermediary is $h(m)s$ with $h(m) = \kappa_0 + \kappa_1 m$, which is strictly increasing provided $\kappa_1 > 0$.

(iv) Lemma 1 and monopoly pricing. The monopoly pricing outcome described in Lemma 1 enables us to represent products using the (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. (For instance this could be the case when there is a resale price agreement between the manufacturer and the retailer.) Of course in practice prices can 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 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 intermediary generates no economies of search, and iii) $\bar{m} = 1$ such that there is no stocking space limit. This relatively simple case 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 $A \in \mathcal{F}$. If a consumer visits the intermediary, she will buy all products available there and so obtain an additional utility $\int_A v dG$, but at the same time she also incurs an additional search cost $s \int_A dG$. Consequently a consumer visits the intermediary if and only if $s \le k$, where

$$k = \frac{\int_A v dG}{\int_A dG} \tag{2}$$

is the average consumer surplus amongst the products sold at the intermediary. (Note that the consumer will search any product $i \notin A$ if and only if $s \le v_i$, and that the order in which she searches through the intermediary and manufacturers does not matter.)

The intermediary's problem is then to

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

[21] Further details are available on request. This is similar to the idea of search diversion in Hagiu and Jullien (2011).

with k defined in (2).[22] 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.[23]

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. $\int_A dG \in (0,1)$).*

The intermediary earns strictly positive profit even though its search technology is no more efficient than that of the manufacturers whose products it resells.[24] To understand why, recall that the intermediary always makes a gain on the low-v products and a loss on the high-v products, and that these gains and losses are proportional to a product's per-customer profitability π. Now imagine that the intermediary selects its profit-making products amongst those with high π, and selects its loss-making products amongst those with low π. This strategy seeks to maximize gains on the former, and minimize losses 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 never stocks all products.

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 where $q(\pi, v) = 1$ if and only if $(\pi, v) \in A$.

[22] When $\int_A dG = 0$ intermediary profit is zero regardless of how we specify k. In some later analysis we consider limit cases where the measure of A goes to zero, and k will be well-defined via L'hopital's rule.

[23] The same is true for a general $h(m)$ if it increases fast enough in m. However if $h(m)$ is (close to) constant and sufficiently small, k can be greater than any v in A. We consider this case in Section 4.

[24] By continuity the same is true even when the intermediary's search technology is slightly less efficient.

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

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

where λ is the Lagrange multiplier associated with the constraint (4). The first term $\pi\left[F(k)-F(v)\right]$ is the direct effect on profit of stocking product (π, v). The second term $\lambda(v-k)$ reflects the indirect effect from the influence on consumer search behavior (where $\lambda > 0$ as we will see below), and it captures the cross-product externalities in our model. For products with $v < k$ their direct effect is positive as we explained before, while their indirect effect is negative because stocking them reduces the average surplus offered by the intermediary and so causes less consumers to search it. 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:

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

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

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

and

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

(Notice that the intermediary is indifferent about whether or not to stock products with $v = k$.) Therefore the intermediary's optimal product selection consists of two "negatively correlated" regions in the product space.[25] First, the intermediary stocks products in

[25] The horizontal line $\lambda\,(k-v)\,/\,[F(k)-F(v)]$ which divides up the product space is continuous in v, and upward (downward) sloping when $F\,(s)$ is concave (convex).

the bottom-right corner with high v and low π. These products are used to induce consumers to search, but since they have $v > k$ they make a loss and so the intermediary minimizes this by choosing those with the lowest possible π. Second, the intermediary stocks products in the top-left corner with low v and high π. Since these products have $v < k$ they make a profit, and so the intermediary maximizes this by choosing those with the highest possible π. Products in other parts of the product space are not stocked: those with low v and low π would generate little direct profit yet 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 $F(k) \in (0,1)$. This is because Lemma 2 shows that $I(k,\lambda)$ has strictly positive measure, which from the definition of k implies $k \in (\underline{v}, \overline{v})$, and moreover by assumption $s \in (0, \bar{s}]$ where $\bar{s} \geq \bar{v}$. Since k is interior, we can take the first-order condition of (5) with respect to k, and obtain

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

whereupon we observe that $\lambda > 0$.[26] Secondly, we have the original constraint (4), which we can rewrite as

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

We therefore have a system of two equations (8) and (9) in two unknowns. If the system has multiple solutions, the solution that generates the highest profit is the optimal one.

The following result summarizes the above analysis:[27]

Proposition 1 *The intermediary optimally stocks products in the regions of (6) and (7), where $k \in (\underline{v}, \overline{v})$ 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 $[0,1]$, in the optimal solution the product space is divided by $v = k$ and $\pi = \lambda$

[26] Using (8) λ equals $f(k)$ multiplied by the average profit of the intermediary's products. Intuitively it captures the impact on profit of a small decrease in k.

[27] If we consider a discrete number of products $\{(\pi_i, v_i)\}_{i=1,\ldots,n}$, 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. In general it is not easy to solve because there are 2^n possible stocking policies (which is very large even for a few dozen products). One approach is to make the problem smooth by allowing stochastic stocking policies with $q_i \in [0,1]$, such that we can use the Lagrange method and obtain bang-bang solutions. However the solutions for k and λ will then be a complicated function of the locations of individual products in the product space.

with $k = \lambda = \frac{1}{2}$. If $F(s) = \sqrt{s}$ on $[0, 1]$, 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 0.03125 and improves industry profit by 12.5% relative to the case of no intermediary, and in the second example the intermediary makes profit of about 0.036 and improves industry profit by about 10.8%.

In this simple case the intermediary harms consumers, because it does not generate any search efficiencies and at the same time restricts consumer choice. Specifically, consumers with $s < k$ who search the intermediary are forced to buy some low-v products which they otherwise would not buy. Similarly consumers with $s > k$ who do not search the intermediary are unable to buy some attractive high-v products which are now only available at the intermediary. Nevertheless the intermediary may improve total welfare, defined as the sum of industry profit and consumer surplus. Indeed this is the case in both of the above numerical examples, where total welfare increases by about 2.5% and 2.8% respectively. Intuitively, in the absence of an intermediary consumers search too little: they search and buy a product (π, v) only if $s < v$, but from a total welfare perspective they should do so whenever $s < \pi + v$. Moreover amongst products with the same $\pi + v$, this problem of 'under search' is more severe for those with low-v and high-π. Since the intermediary's optimal product selection leads to an expansion in demand for these products, its presence can increase total welfare. Nevertheless the intermediary does not account for the harm it imposes on consumers with $s > k$ and so its product selection is not socially optimal. We will discuss the socially optimal product selection in Section 7.1.

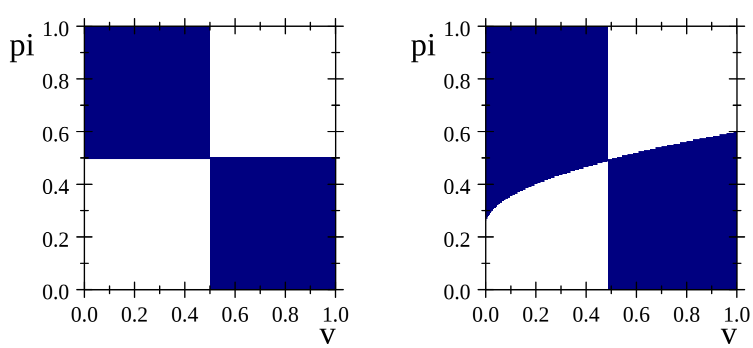

(a) $F(s) = s$ (b) $F(s) = \sqrt{s}$

Figure 1: Optimal product range: the simple case

4 The General Case

We now characterize the optimal product selection in the general case in which i) the intermediary can offer both *exclusive* and *non-exclusive* contracts, ii) the cost of searching the intermediary is $h(m) \times s$ with $h(m)$ weakly increasing, and iii) there is a limit $\bar{m}$ on the measure of products the intermediary can stock. Let $\mathbf{q}(\pi, v) = (q_E(\pi, v), q_{NE}(\pi, v))$ denote the stocking policy function, where $q_E(\pi, v) \in \{0, 1\}$ and $q_{NE}(\pi, v) \in \{0, 1\}$. In particular $q_E(\pi, v) = 1$ if and only if product (π, v) is stocked exclusively, and similarly $q_{NE}(\pi, v) = 1$ if and only if product (π, v) is stocked non-exclusively. (Note that for each product (π, v) at most one of $q_E(\pi, v)$ and $q_{NE}(\pi, v)$ can be 1, but both can be 0 in which case the product is not stocked at all.) It is again convenient to let

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

denote whether or not product (π, v) is stocked. Henceforth whenever there is no confusion we will suppress the arguments (π, v) in the stocking policy function.

4.1 Consumer search behavior

We first solve for consumers' optimal search rule given a stocking policy $\mathbf{q}$. Recall from Lemma 1 that all sellers of a product charge the same price. Hence a consumer will never search both the intermediary and a manufacturer whose product is stocked there. Moreover if a consumer does consider searching a manufacturer, she will do so only if $v > s$. It is also straightforward to see that the order in which a consumer visits the various manufacturers and the intermediary does not matter. Therefore a consumer who searches the intermediary gets an expected surplus

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

where the first two terms are surplus obtained directly from the intermediary, and the final term is surplus obtained by searching products not available at the intermediary. Notice that only $q = q_E + q_{NE}$ matters, and not whether products are stocked exclusively or non-exclusively.

At the same, a consumer who does not search the intermediary gets expected surplus

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

because she can only buy products which are available from their manufacturers (and so are not stocked exclusively by the intermediary). Notice that when the intermediary stocks more products exclusively (precisely, for a fixed q more products have $q_E = 1$), $u^1(s, \mathbf{q}) - u^0(s, \mathbf{q})$ increases and so consumers find it more attractive to search the intermediary. In order to ease the exposition, we suppose that a consumer searches the intermediary if and only if doing so strictly improves her payoff. We then obtain the following result:

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 $\int q dG \leq h\left(\int q dG\right)$.
(ii) $k > \bar{s}$ (everybody searches the intermediary) if $\int qv dG > h\left(\int q dG\right)\bar{s}$.
(iii) $k \in (0, \bar{s}]$ otherwise and is the unique solution to

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

According to part (i) of the lemma, no consumer visits the intermediary when all its products are non-exclusive and it generates diseconomies of search. This is because in that case all the intermediary's products can be acquired elsewhere at a lower search 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, and so has a higher incentive to visit the intermediary.[28] Notice that consumer search behavior is affected only by the measure (and not the identity) of non-exclusive products with $v > k$. This is because consumers with $s < k$ would buy these products anyway, so making them available at the intermediary only changes the search cost associated with buying them.

4.2 Optimal product range

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

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

We can break up profit into three components. Firstly, products with $v < k$ generate a profit of $\pi[F(k) - F(v)] > 0$ and this is independent of whether they are stocked exclusively (i.e. only $q = q_E + q_{NE}$ matters). This is because even under non-exclusivity

[28] More precisely, the advantage of shopping at the intermediary is that it stocks some exclusive products and/or has a better search technology, while the disadvantage is having to buy some low-v products. Since consumers with low s would like to buy most products anyway, the latter disadvantage is small.

the manufacturer makes zero direct sales, since consumers with $s < k$ buy its product from the intermediary and consumers with $s \geq k$ are not willing to search it. Hence the intermediary earns revenue $\pi F(k)$, and must pay the manufacturer the full profit $\pi F(v)$ that it would earn by rejecting the offer. Secondly, products with $v > k$ that are stocked exclusively generate a negative profit of $\pi [F(k) - F(v)] < 0$, and the explanation is the same as in the simple case. Lastly, products with $v > k$ that are stocked non-exclusively generate *zero* profit and so do not appear in equation (13). The reason is that when a manufacturer signs a non-exclusive contract, consumers with $s < k$ switch and buy its product from the intermediary, but consumers with $s \in (k, v)$ continue to buy direct from it. Hence the intermediary only needs to compensate the manufacturer by $\pi F(k)$, which is exactly its own revenue from selling that product. Note that although the intermediary breaks even on these products, it may have an incentive to stock them in order to influence consumer search behavior. (We show later in Section 7.2 that with upstream competition even these non-exclusive products can generate a profit.)

The following lemma shows that the intermediary is guaranteed to earn strictly positive profit provided there is some feasible store size $\tilde{m}$ where it does not generate strict diseconomies of search.[29]

Lemma 4 *The intermediary stocks a strictly positive measure of products and earns a strictly positive profit if there exists an* $\tilde{m} \in (0, \bar{m})$ *such that* $h(\tilde{m}) \leq \tilde{m}$*.*

We now proceed to characterize the optimal product selection when the intermediary can profitably exist with $k > 0$ (e.g. when the condition in Lemma 4 is satisfied). The intermediary wishes to maximize its profit from equation (13) given the space limit $\bar{m} \leq 1$, where k was defined earlier in Lemma 3. When $k \in (0, \bar{s}]$ we know that k satisfies equation (12), which we can rewrite as

$$\int_{v<k} qv dG + \int_{v>k} (q_E v + q_{NE} k) dG - h(m)k = 0 \ , \tag{14}$$

where m denotes the measure of products stocked by the intermediary and satisfies

$$m = \int q dG \ . \tag{15}$$

The stocking space constraint can be written as

$$m \leq \bar{m} \ . \tag{16}$$

[29]Note that this is only a simple sufficient condition for the existence of the intermediary. In the numerical examples below we provide weaker sufficient conditions.

It is again convenient to solve the intermediary's problem using the Lagrangian method, where we treat (14)-(16) as constraints. Let λ, μ and η be their respective Lagrange multipliers. After some simple manipulations we can write the (Kuhn-Tucker) Lagrange function as

$$\begin{aligned}\mathcal{L} &= \int_{v<k} q\left\{\pi\left[F(k)-F(v)\right]+\lambda v-\mu\right\} dG \\ &\quad +\int_{v>k}\left\{q_E\left[\pi\left[F(k)-F(v)\right]+\lambda v-\mu\right]+q_{NE}\left(\lambda k-\mu\right)\right\} dG \\ &\quad -\lambda k h(m)+\mu m+\eta(\bar{m}-m)\ . \end{aligned} \tag{17}$$

(Note that if $k>\bar{s}$, the constraint (14) does not apply and so we set $\lambda=0$.) It is again useful to understand this Lagrange function in terms of the direct and indirect effects of stocking a product. The direct effect is the profit generated by this product: recall from earlier that it is zero if the product is non-exclusive and $v>k$, and otherwise equals $\pi\left[F(k)-F(v)\right]$. The indirect effect captures the effect on consumer search behavior: it equals $\lambda k-\mu$ if the product is non-exclusive and $v>k$, and otherwise is $\lambda v-\mu$. Since the integrands are again linear in q, or (q_E, q_{NE}), we have the following characterization of the optimal product selection:

Proposition 2 *Suppose the intermediary earns a strictly positive profit (e.g. the condition in Lemma 4 is satisfied). The optimal product selection is as follows (with $k>0$, $\lambda\geq 0$ and $\mu\geq 0$ defined in the appendix):*
(i) Products with $v<k$ are stocked if and only if

$$\pi \geq \frac{\mu-\lambda v}{F(k)-F(v)}\ , \tag{18}$$

and it does not matter whether these products are exclusive or not.
(ii) Products with $v>k$ are stocked exclusively if and only if

$$\pi \leq \frac{\max\{\lambda k,\mu\}-\lambda v}{F(k)-F(v)}\ . \tag{19}$$

At the same time, if $\lambda k>\mu$ then all other products with $v>k$ are stocked non-exclusively, if $\lambda k=\mu$ some of the other products with $v>k$ are stocked non-exclusively, and if $\lambda k<\mu$ none of the other products with $v>k$ are stocked.

As in the previous simple case, the intermediary stocks some high-v and low-π products exclusively to attract consumers, and some low-v and high-π products to generate profits. A major difference is that now the intermediary may also stock high-v and high-π products non-exclusively. Whether or not that happens depends on the sign of $\lambda k - \mu$, the indirect effect of stocking products with $v > k$ non-exclusively.[30,31]

In general it appears hard to find primitive conditions for the sign of $\lambda k - \mu$, but we have the following useful observation:

Corollary 1 *If the space constraint is not binding ($m < \bar{m}$) in the optimal product selection, then $\mu = \lambda k h'(m)$. In the region of $v > k$, all products are stocked (i.e. $\lambda k > \mu$) if $h'(m) < 1$, but only low-π products are stocked exclusively (i.e. $\lambda k < \mu$) if $h'(m) > 1$.*

This result implies that with (marginal) diseconomies of search in visiting the intermediary the optimal product selection features two negatively correlated regions as in the previous simple case: one with low-v and high-π products in the top-right corner and the other with high-v and low-π products in the bottom-right corner.[32]

Notice that in the region of $v < k$ the intermediary is indifferent about whether to stock a product exclusively or non-exclusively, since exclusivity has no effect on either the direct profit generated (c.f. equation (13)) nor on consumer search behavior (c.f. Lemma 3). One way to break this indifference is to introduce some small-demand consumers who never visit the intermediary. In that case the intermediary strictly prefers to stock products with $v < k$ non-exclusively, because it reduces the compensation paid to manufacturers. (A formal proof is available upon request.) Hence in the following we say that products with $v < k$ are stocked non-exclusively.

We now illustrate the optimal product selection using some examples. First, we illustrate the result in Corollary 1 about how $h'(m)$ affects the optimal stocking policy. Consider the case with $F(s) = s$ and $G(\pi, v) = \pi v$, and suppose there is no space constraint *(i.e. $\bar{m} = 1$)* and $h(m) = a + bm$ with $a, b \geq 0$.[33] In Figure 2(a) $h(m) = 0.8m$ and so the intermediary generates economies of search. As shown in Corollary 1 we have $\lambda k > \mu$ and so all products with $v > k$ are stocked. In Figure 2(b) $h(m) = m$ as in the simple case we

[30] If $\lambda k - \mu = 0$, the intermediary is indifferent about which of these non-exclusive products to select, and the equation $\lambda k - \mu = 0$ is used to determine their measure. (Notice that in that case only the measure of non-exclusively stocked products, instead of the exact composition, matters for both (12) and (17). Consequently the optimal selection is not unique.)

[31] Another subtle difference is that in this general case it is possible that $k < \underline{v}$ when $\underline{v} > 0$ and diseconomies of search are sufficiently strong, or $k > \overline{v}$ when economies of search are sufficiently strong. When this happens, only one part of the characterization in Proposition 2 is relevant.

[32] When $\lambda k < \mu$ these two regions are also disconnected, because as $v \to k$ the two thresholds in (18) and (19) tend respectively to ∞ and $-\infty$. Intuitively products with v close to k generate only a small direct profit or loss, but when for example $h'(m) > 1$ their negative effect on consumers' incentives to search the intermediary dominates.

[33] This example can be fully solved. (Details are available upon request.) A sufficient condition for the intermediary to exist in this example is $a + b \leq \frac{3}{2}$, which is weaker than Lemma 4.

studied earlier. Here we have $\lambda k = \mu$ and so the intermediary is indifferent about whether to stock each of the products in the top-right corner $[0.5, 1]^2$ non-exclusively. We have drawn the figure for the case where all those products are stocked. Finally in Figure 2(c) $h(m) = 1.2m$ and so the intermediary generates diseconomies of search. We now have $\lambda k < \mu$ and so there are only exclusive products in the region of $v > k$. Comparing across these three cases, as the intermediary's search technology becomes less efficient it stocks fewer products (m decreases from 0.76 to 0.75 and then to 0.29) but a larger proportion of them are exclusive (the percentage increases from 27% to 33% and then to 50%).

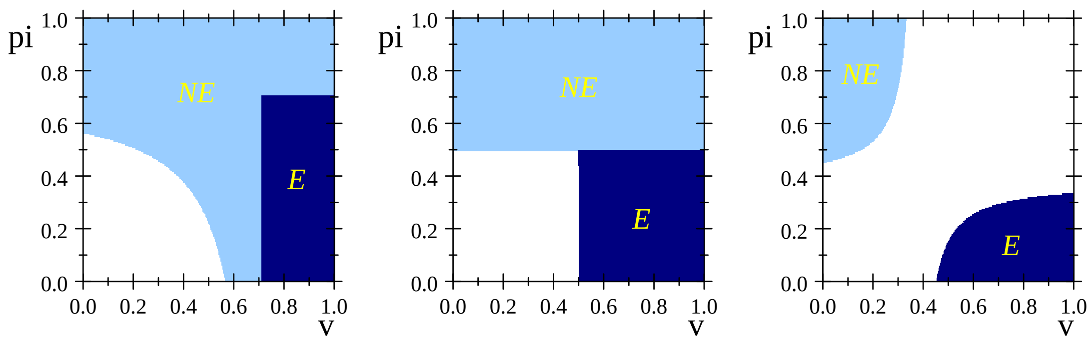

(a) $h(m) = 0.8m$ (b) $h(m) = m$ (c) $h(m) = 1.2m$

Figure 2: Optimal product range with no space constraint

Second, we illustrate how the space limit $\bar{m}$ affects product selection. Consider again the case with $G(\pi, v) = \pi v$ and $F(s) = s$, but now suppose $h(m) = 0.4$. The marginal economies of search are so strong in this case that the intermediary will stock all products if $\bar{m} = 1$, and so the space constraint is always binding. Starting from $\bar{m} = 1$, as we reduce $\bar{m}$ the intermediary first drops the products with low-v and low-π. When $\bar{m}$ is above approximately 0.65 we have that $k > 1$ and so all products are stocked non-exclusively. When $\bar{m}$ is between around 0.65 and 0.463 we find that $k < 1$, and so as depicted in Figure 3(a) the intermediary starts to stock some products exclusively. At the same time $\lambda k - \mu > 0$ and so all products with $v > k$ are stocked. On the other hand, when $\bar{m}$ is between around 0.463 and 0.454 we find that $\lambda k - \mu = 0$ in the optimal solution, and so the intermediary is indifferent about stocking any product in the top-right corner non-exclusively. The condition $\lambda k - \mu = 0$ then determines exactly how many are stocked. Figure 3(b) illustrates this situation. (Recall that there is flexibility over which products in the top-right corner should be stocked. In the figure we choose those with the highest v.) Finally when $\bar{m}$ is below around 0.454 we have $\lambda k - \mu < 0$ and so the optimal product selection consists of two negatively correlated regions. This is illustrated in Figure 3(c).

Comparing Figures 3(a)-(c), we find that as the space constraint becomes tighter a higher proportion of the intermediary's products are stocked exclusively. This is similar to what we found in Figure 2 when the intermediary's search technology became less efficient.

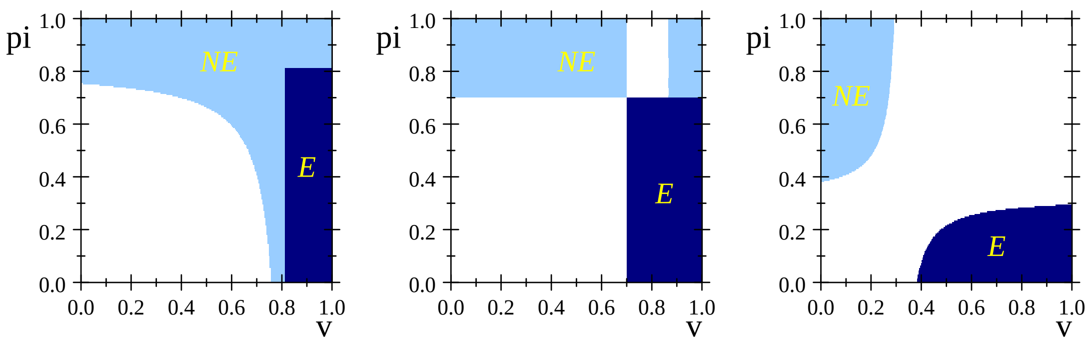

(a) $\bar{m} = 0.5$ (b) $\bar{m} = 0.46$ (c) $\bar{m} = 0.3$

Figure 3: Optimal product range with binding space constraint and $h(m) = 0.4$

5 Applications

We now apply our framework to examine the effect of DTC sales on retail markets, and to discuss the optimal design of a shopping mall which acts a platform and does not possess products directly.

5.1 The impact of direct-to-consumer sales

We argued in the introduction that it has become increasingly easy for manufacturers to sell direct to consumers. In this section we use our framework to assess the impact of DTC sales on the retail market.

As a first step, we solve for the intermediary's optimal product selection when DTC sales are infeasible. In particular, suppose that manufacturers can only sell their product to consumers via the intermediary. A consumer's payoff from searching the intermediary is $\int qvdG - sh(m)$, where as usual $m = \int qdG$. Hence the intermediary attracts all consumers with $s < k$, where

$$k = \frac{\int qvdG}{h\left(\int qdG\right)} \ . \tag{20}$$

Clearly if the intermediary wishes to stock product i it will offer the manufacturer ($\tau_i = c, T_i = 0$) and de facto becomes the exclusive supplier. The intermediary's profit is therefore $\int q\pi F(k)dG$, which is strictly positive for any stocking policy except $\mathbf{q} = \mathbf{0}$. The optimal product selection is reported in the following result:

Proposition 3 *With no DTC sales, the intermediary stocks all the products with*

$$\pi \geq \frac{\mu - \lambda v}{F(k)} \ , \tag{21}$$

where $k > 0$, $\lambda \geq 0$ and $\mu \geq 0$ are defined in the appendix.

We can show that $\lambda > 0$ whenever the intermediary stocks only some products (i.e. $m < 1$) and does not attract all consumers (i.e. $k < \bar{s}$). In this case the right-hand side of (21) is decreasing in v, and so the intermediary only stocks products with both a high v and a high π. This contrasts with our earlier results, where we found that when DTC sales are feasible the optimal product range includes two negatively correlated regions, namely a high-v-and-low-π region and a low-v-and-high-π region. Indeed we also found that these two regions are the only ones that are stocked when, for example, the intermediary creates marginal diseconomies of scale or has a sufficiently binding stocking constraint. Hence the feasibility of DTC sales can fundamentally change what the intermediary stocks.

To illustrate this point further, we return to our running example with $G(\pi, v) = \pi v$ and $F(s) = s$. Figure 4(a) plots the optimal stocking policy when $h(m) = 1.2m$ and $\bar{m} = 1$, in which case the stocking constraint is not binding in the optimal solution. Figure 4(b) does the same when $h(m) = 0.4$ and $\bar{m} = 0.3$, in which case the stocking constraint is binding in the optimal solution. In both cases we observe that the intermediary stocks very different products than it did in Figures 2(c) and 3(c) respectively, which are the comparable diagrams when DTC sales are feasible.

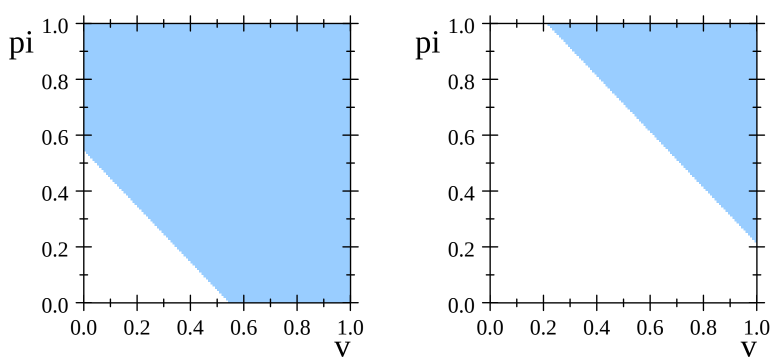

(a) $h(m) = 1.2m$ and $\bar{m} = 1$ (b) $h(m) = 0.4$ and $\bar{m} = 0.3$

Figure 4: Optimal product range with no DTC sales

Now consider the impact of DTC sales on industry profits:

Corollary 2 *After DTC sales become feasible, manufacturer profit and total industry profit both increase, whilst the intermediary's profit decreases.*

Manufacturers clearly benefit when DTC sales become feasible since they earn a positive profit $\pi F(v)$. However the intermediary's profit is lower with DTC sales. To understand why, imagine that initially manufacturers can sell direct but then suddenly they can't. Even if the intermediary does not respond by changing its product selection, its profit increases because i) it no longer has to compensate manufacturers, and ii) weakly more consumers search it, since they no longer have the option to buy direct. It is also easy to show that the negative effect of DTC sales on the intermediary is outweighed by its positive effect on manufacturers, and so DTC sales raise industry profit.

Interestingly, our model suggests that if an intermediary is to survive the threat posed by DTC sales, it may need to qualitatively change the mix of products that it stocks. To illustrate this, consider again the examples in Figure 4. It turns out that if DTC sales become feasible but the intermediary wishes to stock (exclusively) the products in the shaded regions, it will end up making a loss (profits become respectively -0.027 and -0.036). However in both cases the intermediary can make positive profit by dropping many of the high-v and high-π products that it previously stocked.

Finally, the impact of DTC sales on consumer surplus and total welfare is difficult to investigate analytically since, as we have seen, the feasibility of DTC sales can significantly change the intermediary's product selection. However in the examples we have studied above DTC sales improve both consumer surplus and total welfare. (In the example in Figure 4(a), DTC sales improve consumer surplus from 0.110 to 0.145 and total welfare from 0.329 to 0.413; in the example in Figure 4(b), DTC sales improve consumer surplus from 0.062 to 0.144 and total welfare from 0.186 to 0.406.) In addition, if $\bar{m}$ is sufficiently small then DTC sales must improve both consumer surplus and total welfare since without DTC sales consumers are only able to buy very few products.

5.2 Shopping malls

Our framework can also shed light on the design of shopping malls. Suppose there is a unit mass of sellers each of which can either join a shopping mall and/or set up its own independent shop in the same area. The shopping mall can host up to $\bar{m}$ sellers and charges each of them a fixed fee. Consumers pay s to search an independent shop and $h(m) \times s$ to search a mall which contains m shops. The timing of offers to join the mall, pricing by sellers, and search decisions by consumers, are the same as in Section

2. (Notice that since the shopping mall acts as a platform, it does not set retail prices directly.)

This is isomorphic to our main model. Firstly, it is straightforward to see that seller i will charge p_i^m wherever it sells its product. (We do not even need Lemma 1.) Hence our (π, v) representation remains valid. Secondly, consumers will search the mall if and only if $s < k$ where k was defined earlier in Lemma 3. Thirdly, sales of a product (π, v) at the mall generate a gross profit of $\pi F(k)$. Therefore a seller with $v < k$ will pay $\pi [F(k) - F(v)]$ to join the mall, whilst a seller with $v > k$ is willing to join the mall for free if it also maintains an independent store elsewhere, but must be paid $\pi [F(v) - F(k)]$ to exclusively join the mall. Consequently the mall owner's profit is the same as in equation (13), and Proposition 2 characterizes which sellers join the mall.

We can interpret sellers with $v > k$ who are exclusive to the mall as 'anchor stores'. According to our model, the mall owner should subsidize anchor stores to encourage them to join. This is worthwhile because their presence attracts more consumers (i.e. increases k), which allows the mall to charge a higher fee to other 'non-anchor' stores. Consistent with this, Gould et al (2005) show empirically that 73% of anchor stores pay zero rent (and the mall developer often pays for things like development and maintenance costs). Even among stores which pay rents, anchor stores pay substantially lower rent than non-anchor stores ($4.13 vs $29.37 per square foot). They also find that adding more anchor stores to a mall leads to an increase in both the sales of, and the rents charged to, non-anchor stores.

6 Foundation of the (π, v) Product Space

Up to now we have characterized the intermediary's optimal stocking policy in terms of the (π, v) space. In this section we explain how to construct the (π, v) space, and also how to interpret different points within in.

We first discuss how to generate the (π, v) space. Suppose that demand for product i can be written as $Q_i(p_i) \equiv Q(p_i, \boldsymbol{\delta}_i)$ where $\boldsymbol{\delta}_i$ is a vector of demand parameters. We can then let $(\boldsymbol{\delta}, c)$ denote the vector of product-specific parameters. Using the definitions introduced in equation (1), one can calculate $\pi_i = \pi(\boldsymbol{\delta}_i, c_i)$ and $v_i = v(\boldsymbol{\delta}_i, c_i)$, and then derive the (π, v) space from the parameter space $(\boldsymbol{\delta}, c)$.[34]

[34] If products differ in exactly two parameters (like in the second example below), it is possible to have a one-to-one correspondence between the parameter space and the (π, v) space. If products differ in more than two parameters (like in the first example below), generically each point in the (π, v) space represents a continuum of different products. We can then let $q_E(\pi, v)$ and $q_{NE}(\pi, v)$ denote the fraction of products at point (π, v) which are respectively stocked exclusively and non-exclusively. Since all the

We now discuss two particular classes of demand which illustrate how different points in the (π, v) space can be interpreted in terms of the scale, curvature and elasticity of demand.

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

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

where $a_i > 0$ is the scale and $\sigma_i \in (-\infty, 2)$ is the curvature of demand. (Curvature is defined as the elasticity of the slope of inverse demand.) This is a rich class which includes concave demand when $\sigma_i < 0$, linear demand when $\sigma_i = 0$, convex demand when $\sigma_i > 0$, and exponential demand when $\sigma_i = 1$. One can check that if $c_i \frac{1-\sigma_i}{2-\sigma_i} < 1$, the monopoly price for product i is $p_i^m = 1 + \frac{c_i}{2-\sigma_i}$, and

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

Notice that π_i and v_i are both increasing in the scale parameter a_i and decreasing in the cost parameter c_i, whilst the ratio $\pi_i/v_i = 2 - \sigma_i$ is decreasing in the amount of curvature. (Intuitively a lower σ_i means that demand is more concave and 'rectangular-shaped', and hence the firm can extract a greater share of the available surplus.) Consequently each product in (π, v) space lies on a ray from the origin, where σ_i determines the slope of the ray and (a_i, c_i) determine how far the product is along the ray.[35] Therefore in this example, products with relatively large and convex demand should be stocked exclusively, whilst products with relatively large but concave demand should be stocked non-exclusively and used to generate profit.

Demand elasticity: As another example, suppose that product i's demand function is

$$Q_i(p_i) = a_i \left(1 - p_i^{\sigma_i}\right) , \tag{23}$$

where $a_i > 0$ is again a scale parameter, but $\sigma_i > 0$ is now an elasticity parameter. Specifically, it is easy to show that for any $p_i \in (0, 1)$ the demand elasticity is strictly decreasing in σ_i. When products all have the same marginal cost $c = 0$, we can compute $p_i^m = (\frac{1}{1+\sigma_i})^{\frac{1}{\sigma_i}}$, and

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

objective functions in our paper are linear in stocking policy variables, we obtain bang-bang solutions where $q_E(\pi, v), q_{NE}(\pi, v) \in \{0, 1\}$ and so all our analysis is unchanged.

[35] This insight holds beyond the particular demand class studied here. Anderson and Renault (2003) and Weyl and Fabinger (2013) show that in general 'more concave' demands have a higher π_i/v_i ratio.

Again both π_i and v_i are increasing in a_i, whilst π_i/v_i is increasing in σ_i and therefore decreasing in demand elasticity.[36] (Intuitively when demand is more elastic the monopoly price is lower and consumers enjoy a greater share of surplus.) Consequently in this example there is a one-to-one correspondence between (a, σ) space and (π, v) space. Moreover products with relatively large and elastic demands should be stocked exclusively, whilst products with relatively large and inelastic demand should be stocked non-exclusively to earn positive profit.

7 Extensions

In this section we return to the baseline model and discuss two extensions: one with a social planner who aims to maximize total welfare, and the other with upstream competition where each product is supplied by multiple manufacturers.

7.1 The socially optimal stocking policy

Consider the baseline model with DTC sales. We study the optimal product range of a social planner that wishes to maximize total welfare (defined as the sum of industry profit and consumer surplus). The social planner chooses the stocking policy $\mathbf{q}$, but has no direct control over firm pricing or how consumers search.

Since the consumer search rule is the same as in Lemma 3, total welfare can be written as

$$W(\mathbf{q}) \equiv \int \pi F(v)\, dG + \Pi(\mathbf{q}) + \int_0^k u^1(s, \mathbf{q})\, dF(s) + \int_k^{\bar{s}} u^0(s, \mathbf{q})\, dF(s)\ . \qquad (24)$$

The first term is profit earned by manufacturers. Note that a manufacturer earns $\pi F(v)$ regardless of whether or not it sells via the intermediary. The second term is the intermediary's profit, which we defined earlier in equation (13). The remaining two terms are consumer surplus. Consumers with $s < k$ search the intermediary and earn $u^1(s, \mathbf{q})$, which we defined in equation (10). Meanwhile consumers with $s \geq k$ do not search the intermediary and so earn $u^0(s, \mathbf{q})$, which we defined in equation (11). When some products with $v > k$ are stocked exclusively, consumers with $s \geq k$ are made worse off because they cannot buy all the products they would like to. At the same time, consumers with $s < k$ are forced to buy some low-v products that ordinarily they would not purchase, and so depending on the strength of economies of search may or may not be better off.

[36] Notice that $\pi_i/v_i > 1$ for any $\sigma_i > 0$, and so the (π, v) space can only lie in half of the quadrant $\mathbb{R}_+^2$.

The social planner wishes to maximize (24). We can solve its optimization problem using the same techniques as we did earlier for profit maximization. We again use $m = \int q dG$ to denote the measure of products stocked by the intermediary.

Proposition 4 *A sufficient condition for the socially optimal stocking policy to have $m > 0$ is that $h(\tilde{m}) \leq \tilde{m}$ for some $\tilde{m} \in (0, \bar{m})$. The optimal policy is characterized as follows (with $k > 0$, $\lambda \geq 0$ and $\mu \geq 0$ defined in the appendix):*
(i) Products with $v < k$ are stocked if and only if

$$\pi + v \geq \frac{\mu - \lambda v - \int_0^v s dF(s)}{F(k) - F(v)} , \tag{25}$$

and it does not matter whether they are exclusive.
(ii) Products with $v > k$ are stocked exclusively if and only if

$$\pi + v \leq \frac{\max\{\mu, \lambda k + \int_0^k s dF(s)\} - \lambda v - \int_0^v s dF(s)}{F(k) - F(v)} . \tag{26}$$

Of the remaining products with $v > k$, all of them are stocked non-exclusively if $\lambda k + \int_0^k s dF(s) > \mu$, some of them are stocked non-exclusively if $\lambda k + \int_0^k s dF(s) = \mu$, and none of them are stocked if $\lambda k + \int_0^k s dF(s) < \mu$.

It is socially optimal for the intermediary to stock a positive measure of products, provided that it generates (weak) economies of search for at least some $\tilde{m} \in (0, \bar{m})$. Qualitatively the welfare-optimal stocking policy is similar to what an unfettered intermediary would choose. Firstly, exclusive products with $v > k$ are again chosen to have low π, and products with $v < k$ are chosen to have high π. Intuitively this is because, as we noted earlier, consumers do not take into account sellers' profit, and therefore search (and buy) too little from a welfare perspective. Demand for a product with $v > k$ is further reduced when it is sold exclusively by the intermediary, but conversely demand for a product with $v < k$ is increased when it is sold by the intermediary. Choosing the former products to have low π minimizes the additional welfare loss, and choosing the latter to have a high π maximizes the welfare gains. Secondly, and mirroring Corollary 1 from earlier, when the stocking constraint is slack at the optimum all products with $v > k$ are stocked if $h'(m) < 1$, but only those with low π are stocked (and done so exclusively) if $h'(m) > 1$.[37]

We would now like to compare the stocking policies chosen by the intermediary and social planner. Unfortunately this is hard to do analytically, because in general (k, λ, μ)

[37] The indirect value of non-exclusively stocking a product with $v > k$ is $\lambda k - \mu + \int_0^k s dF(s)$, which after solving for μ simplifies to $[1 - h'(m)] \left(\lambda k + \int_0^k s dF(s)\right)$. Intuitively non-exclusivity of a product leads consumers with $s < k$ to buy it from the intermediary rather than from its manufacturer, and at the margin this increases these consumers' search cost by $h'(m) - 1$ units.

differ across the two solutions and are determined by a complex system of equations. Intuitively though, one would expect the social planner to stock fewer high-v products exclusively since this harms consumers that do not search the intermediary. Similarly, if economies of search are not too strong, one would expect the social planner to stock fewer low-v products because consumers who search the intermediary end up buying them even though they provide little surplus. We confirm this intuition using our running example with $G(\pi, v) = \pi v$ and $F(s) = s$. Suppose there is no stocking space constraint. Figure 5(a) plots the socially optimal product range for the case $h(m) = 0.8m$. Here (k, λ, μ) differ from Figure 2(a), and the social planner's selection is not a subset of the intermediary's. Nevertheless the intermediary stocks fewer products overall and also fewer products exclusively. Figure 5(b) plots the socially optimal product range for the case $h(m) = m$. It turns out that in this case (k, λ, μ) are the same as in Figure 2(b). The social planner's product range is therefore a strict subset of the intermediary's, and it also stocks fewer products exclusively. Figure 5(c) plots the socially optimal product range for the case $h(m) = 1.2m$. Since the intermediary has diseconomies of search, this product range is now much smaller than in Figure 2(c).

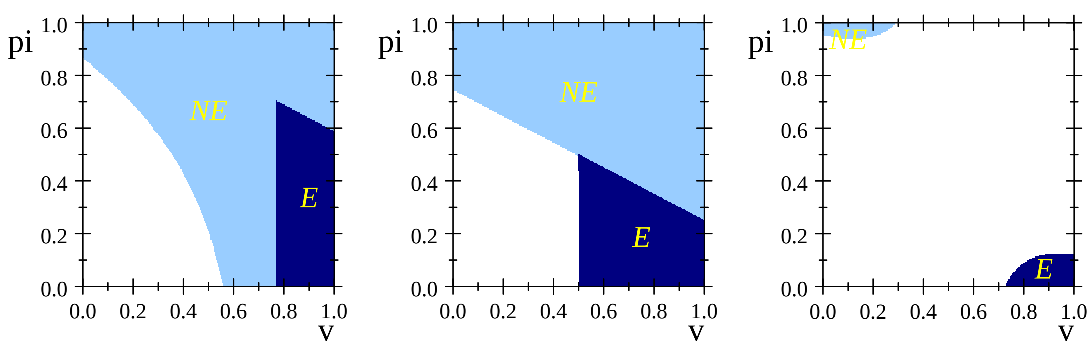

(a) $h(m) = 0.8m$ (b) $h(m) = m$ (c) $h(m) = 1.2m$

Figure 5: Socially optimal product range

Finally, even though our results suggest that the intermediary stocks too many exclusive products, a complete ban on exclusivity is not necessarily welfare enhancing. This is easily seen from the simple case where $h(m) = m$. If the intermediary cannot offer exclusive contracts it is unable to earn positive profit and therefore will not exist. However if it can offer exclusive contracts, we saw earlier that it exists and can raise welfare.

7.2 Upstream competition

We now show that our main insights are robust when there is upstream competition. To do this as simply as possible, we assume that each product is supplied by two homogeneous manufacturers, and that there are no economies of search in visiting the intermediary. Precisely, we assume that $h(m) = m$ where m denotes the measure of *distinct* products stocked by the intermediary. We further 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, whilst consumers search sequentially with passive beliefs and randomize when indifferent.

Closely following Lemma 1, we can prove that in equilibrium all sellers of a product charge the monopoly price.[38] Intuitively the intermediary again 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). Consequently we can still represent products using a two-dimensional (π, v) space. Adapting our earlier notation, we now let $q_E(\pi, v)$ be an indicator function which equals 1 if and only if product (π, v) is available exclusively from the intermediary. (This means that the intermediary signs an exclusive contract with both manufacturers of the product.) Similarly, and in the spirit of our earlier notation, we now let $q_{NE}(\pi, v)$ be an indicator function which equals 1 if and only if product (π, v) is available at both the intermediary and at least one manufacturer. (Note that this can be generated by several different contracting arrangements between the intermediary and the two manufacturers.) Hence $q(\pi, v) = q_E(\pi, v) + q_{NE}(\pi, v)$ is again an indicator function which equals 1 if and only if the intermediary stocks product (π, v).

The optimal consumer search rule can then be derived as follows. Since the two manufacturers of a product are homogeneous, what matters from a consumer's respective is whether the intermediary has exclusivity at the level of a product, rather than with respect to an individual manufacturer. Therefore given our new definitions of q_E and q_{NE}, the payoffs from respectively searching and not searching the intermediary are the

[38] One subtlety is that we need $f(0) = 0$ to sustain monopoly pricing when the manufacturers both sell direct to consumers. This holds provided that s is bounded away from 0. (For convenience we assume in our numerical examples that $s \sim [0, 1]$, but this can be regarded as the limit case where $s \sim [\epsilon, \epsilon + 1]$ and $\epsilon \to 0$.)

same as those in equations (10) and (11) with $h(m) = m$. Lemma 3 implies that provided $\int q_E dG > 0$, 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

$$\int_{v<k} q\,(v-k)\,dG + \int_{v>k} q_E\,(v-k)\,dG = 0\ . \tag{27}$$

The following result shows how much the intermediary must compensate manufacturers (on top of the production cost) in order to stock their product:

Lemma 5 *(i) If the intermediary wants $q_{NE}(\pi, v) = 1$ it should offer each manufacturer of product (π, v) a non-exclusive contract with zero compensation.*
(ii) If the intermediary wants $q_E(\pi, v) = 1$ it must offer each manufacturer of product (π, v) an exclusive contract with compensation $\max\{0, \pi[F(v) - F(k)]\}$.

According to Lemma 5 when the intermediary wants to induce $q_{NE}(\pi, v) = 1$ it can achieve this at zero cost by offering each manufacturer a non-exclusive contract. Although there are other contract arrangements which also lead to $q_{NE}(\pi, v) = 1$, they must entail (weakly) higher compensation and are therefore (weakly) dominated.[39] Henceforth when we say that the intermediary stocks a product non-exclusively, we mean that it signs a non-exclusive contract with each manufacturer.

When the intermediary contracts with one manufacturer it affects the 'prominence' and hence bargaining power of the other manufacturer. Firstly, notice that if one manufacturer supplies the intermediary, all consumers with $s < \min\{v, k\}$ buy its product through the intermediary rather than from the other manufacturer. This explains why in Lemma 5 each manufacturer whose product has $v < k$ is willing to supply the intermediary (exclusively or non-exclusively) with zero compensation. Secondly, if the intermediary offers a non-exclusive contract to both manufacturers of a product with $v > k$, the manufacturers are duopolists over consumers with $s \in (k, v)$, and each earns $\frac{1}{2}\pi[F(v) - F(k)]$ irrespective of whether it accepts the intermediary's offer. Hence each is willing to supply the intermediary with zero compensation. Thirdly, if the intermediary offers an exclusive contract to both manufacturers of a product with $v > k$, a manufacturer who rejects will become a monopolist over consumers with $s \in (k, v)$, and therefore earn $\pi[F(v) - F(k)]$. Hence each manufacturer must be compensated by this amount.

Using Lemma 5 we can write the intermediary's profit as

$$\int_{v<k} q\pi F(k)dG + \int_{v>k} q_{NE}\pi F(k)dG + \int_{v>k} q_E \pi[3F(k) - 2F(v)]dG\ . \tag{28}$$

[39] Compensation is strictly higher under any alternative contracting arrangement when $v > k$. Details are available on request.

According to Lemma 5, the intermediary only pays marginal cost when it stocks a product with $v < k$ (regardless of exclusivity) and when it stocks a product with $v > k$ non-exclusively. Hence for a given k these products are cheaper to stock than in the basic model due to upstream competition. On the other hand, when the intermediary stocks a product with $v > k$ exclusively it must pay $\pi \left[F\left(v\right) - F\left(k\right) \right]$ to each manufacturer. Compensation per manufacturer is therefore lower than in the basic model, but the intermediary now needs to compensate two manufacturers instead of one. Products with $F(v) < 2F(k)$ become cheaper to stock (and can even generate a profit), while those with $F(v) > 2F(k)$ (if any) become more expensive to stock.

Adapting the proof of Lemma 2 it is simple to show that there exists a stocking policy which generates strictly positive profit. We then have the following result:

Proposition 5 *The optimal product selection with upstream competition, $h\left(m\right) = m$ and $\bar{m} = 1$ is as follows (with $k \in \left(\underline{v}, \bar{v}\right)$ and $\lambda > 0$ defined in the appendix):*
(i) Products with $v < k$ are stocked (and exclusivity does not matter) if and only if

$$\pi \geq \lambda \frac{k - v}{F\left(k\right)} \ .$$

(ii) Products with $v > k$ are all stocked, and are stocked exclusively if

$$\pi \leq \frac{\lambda}{2} \frac{v - k}{F\left(v\right) - F\left(k\right)} \ .$$

Our predictions about the intermediary's optimal product range are thus qualitatively robust to the introduction of upstream competition. One difference compared to the main model is that even without economies of search it is strictly optimal to stock high-π and high-v products non-exclusively. Figure 6 plots the optimal product range in our running example with a uniform product space, when $F(s) = s$ and $F(s) = \sqrt{s}$ respectively. In the former case the intermediary stocks fewer exclusive products and also fewer products overall compared to the main model, while the opposite is true in the latter case. However in both cases the introduction of upstream competition raises the intermediary's profit: in the first example it increases from $1/32$ to about 0.2, and in the second example it increases from about 0.036 to about 0.305. This happens primarily because upstream competition reduces the compensation that must be paid to manufacturers.[40]

[40] To illustrate this, even if the intermediary did not change its product selection in the first example its profit would increase to 0.1875, and so the additional profit earned from reoptimizing its product selection is relatively small.

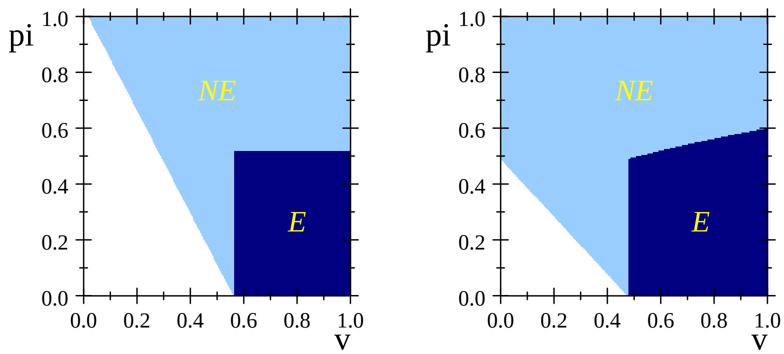

(a) $F(s) = s$ (b) $F(s) = \sqrt{s}$

Figure 6: Optimal product range with upstream competition

Finally, notice that there is a strong connection between upstream competition and private labels. Specifically, suppose each product can be produced by one independent manufacturer and also the intermediary. If the intermediary stocks a private label product with $v < k$ it captures all demand for that product. Similarly by stocking a private label with $v > k$, the intermediary sells it to all consumers with $s < k$ and so reduces the amount of compensation it must give the manufacturer for exclusivity on that product.

8 Conclusion

This paper has developed a new and tractable framework in which to study a multiproduct intermediary's product range choice. We hope that this framework will be useful in future work, especially in situations where product heterogeneity and multiproduct demand are important. Using our framework, we have shown that (i) the intermediary earns strictly positive profit even if it does not improve search efficiency, provided that it can use exclusive contracts and consumers demand multiple products; (ii) the optimal product range is relatively simple to characterize, and reflects endogenous cross-product externalities which arise due to the search friction. Specifically, some exclusive products are used to attract consumers who then generate profit for the intermediary by buying its other non-exclusive products; (iii) DTC sales harm the intermediary by reducing its bargaining power, and may require it to significantly adjust its product range; and (iv) the intermediary tends to be too big, and stock too many products exclusively, compared to the social optimum.

Our paper has mainly focused on retailers, but its insights apply to a much broader set of multiproduct intermediaries. One example is shopping malls, which we discussed earlier in Section 5. Another pertinent example is TV platforms. Traditionally, content providers mainly reached viewers through TV subscription services. However in recent years Disney, HBO, ESPN and several others have launched (or committed to launch) DTC channels.[41] Our model can be applied to understand how this might affect the industry: the manufacturers in our model can be interpreted as content providers, whilst the search friction could be interpreted more broadly as encompassing the costs of finding, subscribing to, and paying for various services.[42] Our model then provides implications about, for example, what channels a cable TV company will carry and how it is affected by the introduction of DTC channels. Another example where our framework could be fruitfully applied is international trade, because intermediaries play an important role in helping domestic manufacturers to export.[43] In that case foreign importers are like the consumers in our model, whilst trade frictions are similar to our search cost. Depending on how easy it is for individual manufacturers to export on their own (which is like DTC in our model), they may choose to sell via an intermediary. Hence by modifying our model one could obtain insights about which products should be exported, and by whom.

Finally, we believe that our paper opens up some promising avenues for future research. Firstly, we have focused on a monopoly intermediary. It would be interesting to allow for multiple intermediaries, and study how competition between them (or a merger of two intermediaries) shapes their optimal product range. Secondly, we have assumed that each product has only one manufacturer or two homogenous manufacturers. It would be interesting to investigate the case where each product has several manufacturers, each of whom produces a differentiated variety. This would make it possible to study both the intermediary's optimal breadth (i.e. how many different types of product to stock) and depth (i.e. how many varieties of a particular product to stock).

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

[41] See goo.gl/7rZ4LW for details on Disney and ESPN, and goo.gl/8Mwf2i for details on HBO's 'HBO Now' service.

[42] See goo.gl/7eQGek for arguments why these costs can be much lower when content is bundled rather than sold direct.

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

$j \in A_E$, 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 is

$$\eta \times [p_j^*(\tau_j) - \tau_j] Q_j(p_j^*(\tau_j)) - T_j = \eta \times \left[p_j^*(\tau_j) - c_j\right] Q_j(p_j^*(\tau_j)) - \pi_j F(v_j) \ . \tag{29}$$

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.[44] 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

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

[44] With 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.

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.[45] 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 $\tau_i \geq c_i$. 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

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

[45]In 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.

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

$$\begin{aligned}\eta_I \times (p_i^m - \tau_i)Q_i(p_i^m) - T_i &= \eta_I \times (p_i^m - c_i)Q_i(p_i^m) + \eta_{i,M} \times \pi_i - \pi_i F(v_i) \\ &= \pi_i \left[\eta_I - (F(v_i) - \eta_{i,M})\right] . \qquad (30)\end{aligned}$$

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: $(\pi_1, \tilde{v})$ and $(\pi_2, \tilde{v})$ with $\pi_1 > \pi_2$. Let $A_1 = [\pi_1 - \delta, \pi_1] \times [\tilde{v} - \epsilon, \tilde{v}]$ and $A_2 = [\pi_2, \pi_2 + \Delta(v)] \times [\tilde{v}, \tilde{v} + \epsilon]$, where $\Delta(v)$ is uniquely defined for each $v \in [\tilde{v}, \tilde{v} + \epsilon]$ by

$$\int_{\pi_1 - \delta}^{\pi_1} g(\pi, 2\tilde{v} - v)\, d\pi = \int_{\pi_2}^{\pi_2 + \Delta(v)} g(\pi, v)\, d\pi . \qquad (31)$$

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

$$\Pi(\epsilon) = \int_{\tilde{v}-\epsilon}^{\tilde{v}} \int_{\pi_1 - \delta}^{\pi_1} \pi \left[F(\tilde{v}) - F(v)\right] dG + \int_{\tilde{v}}^{\tilde{v}+\epsilon} \int_{\pi_2}^{\pi_2 + \Delta(v)} \pi \left[F(\tilde{v}) - F(v)\right] dG . \qquad (32)$$

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

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

where the second equality used (31) 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

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

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 $A = \Omega \backslash (B_1 \cup B_2)$ is $\hat{v}$ as well. Then a consumer will visit the intermediary, when it stocks $A = \Omega \backslash (B_1 \cup B_2)$, if and only if $s < \hat{v}$.

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

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

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

$$\Delta(s) = \int qv dG - h\left(\int q dG\right)s - \int_{v>s} q_{NE}(v - s)\, dG \ . \tag{34}$$

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

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

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 $\Delta(k) = 0$. ■

Proof of Lemma 4. Recall the areas A_1 and A_2 that were defined in the proof of Lemma 2, and let $v^*(\epsilon)$ be the unique solution to

$$\tilde{m} = \int_{A_1 \cup A_2} dG + \int_{v > v^*(\epsilon)} dG .$$

Notice that since $\tilde{m}$ is bounded away from 1 and Ω is convex, we can find $\tilde{v}$ sufficiently close to $\underline{v}$, and $\delta > 0$ sufficiently small, such that for $\epsilon > 0$ sufficiently small we have both $A_1, A_2 \subset \Omega$ and $v^*(\epsilon) > \tilde{v} + \epsilon$.

Suppose the intermediary stocks products in A_1 and A_2 exclusively and stocks products with $v > v^*(\epsilon)$ non-exclusively. Firstly notice that for any $\epsilon > 0$ we have $k \geq \tilde{v}$. To see this, note that by the definition of $\Delta(s)$ in equation (34) we have

$$\begin{aligned} \Delta(\tilde{v}) &= \int qvdG - h(\int qdG)\tilde{v} - \int_{v > \tilde{v}} q_{NE}(v - \tilde{v})\, dG \\ &\geq \int_{A_1 \cup A_2} vdG - \tilde{v} \int_{A_1 \cup A_2} dG = 0 , \end{aligned}$$

where the inequality uses $h\left(\int qdG\right) = h(\tilde{m}) \leq \tilde{m}$, and the final equality follows since $A_1 \cup A_2$ is constructed to have average valuation $\tilde{v}$. The hypotheses of Lemma 3(iii) are satisfied, and the proof of that lemma shows that $\Delta(\tilde{v}) \geq 0$ implies $k \geq \tilde{v}$. Secondly, recall that non-exclusive products with $v < k$ earn positive profit, and non-exclusive products with $v > k$ earn zero profit. Therefore since $k \geq \tilde{v}$ the intermediary's profit is weakly larger than in equation (32) in the proof of Lemma 2, and in that proof we showed that this profit is strictly positive for some $\epsilon > 0$. ■

Proof of Proposition 2. As a preliminary step, the slackness conditions for the space constraint are $\eta(\bar{m} - m) = 0$ and $\eta \geq 0$. Hence, $\eta = 0$ if $m < \bar{m}$ in the optimal solution, and η is irrelevant if $m = \bar{m}$ since in that case the binding space constraint is the same as (15). Also, when the intermediary makes a strictly positive profit, we must have $k > 0$ so that some consumers visit it.

First, we characterize the optimal product selection for given parameters (k, λ, μ). From (17) it is clear that for $v < k$, $q = 1$ if and only if

$$\pi\left[F(k) - F(v)\right] + \lambda v - \mu \geq 0 \Leftrightarrow \pi \geq \frac{\mu - \lambda v}{F(k) - F(v)} , \tag{35}$$

and the exclusivity arrangement does not matter. From (17) it is also clear that for $v > k$, $q_E = 1$ if and only if

$$\pi\left[F(k) - F(v)\right] + \lambda v - \mu \geq \max\{0, \lambda k - \mu\} \Leftrightarrow \pi \leq \frac{\max\{\lambda k, \mu\} - \lambda v}{F(k) - F(v)} . \tag{36}$$

When the opposite is true, $q_{NE} = 1$ if $\lambda k - \mu > 0$, $q_{NE} = 0$ if $\lambda k - \mu < 0$, and the intermediary is indifferent between $q_{NE} = 0$ and $q_{NE} = 1$ if $\lambda k - \mu = 0$. (In the third case, there is flexibility over which products should be stocked, but the selection has to satisfy $\lambda k - \mu = 0$ as we explain more below.)

Second, we give the conditions for determining the parameters. It is useful to consider a few cases separately:

(a) Suppose the optimal solution has $k < \bar{s}$.[46] Then k satisfies (14). Since k is interior, the first-order condition of (17) with respect to k gives

$$\lambda = f(k)\frac{\int_{v<k} q\pi dG + \int_{v>k} q_E \pi dG}{h(m) - \int_{v>k} q_{NE} dG} > 0 \ . \tag{37}$$

(Notice that $\lambda > 0$ since λ and k defined in (12) share the same denominator.) If $m < \bar{m}$ in the optimal solution, then $\eta = 0$ as we have pointed out and so the first-order condition of (17) with respect to m gives

$$\mu = \lambda k h'(m) \geq 0 \ , \tag{38}$$

and (k, λ, μ, m) jointly solve equations (14), (15), (37) and (38). If $m = \bar{m}$ in the optimal solution, η is irrelevant as we have explained. Note that if $m = \bar{m} < 1$ then (k, λ, μ) jointly solve equations (14), (15) and (37) with m replaced by $\bar{m}$ everywhere, and we can conclude that $\mu \geq 0$.[47] Note that if instead $m = \bar{m} = 1$ then μ does not matter so without loss of generality we can set $\mu = 0$, and (k, λ) jointly solve equations (14) and (37) with m replaced by $\bar{m}$.[48] (Regardless of whether (16) binds or not, when $\lambda k - \mu = 0$ we have an additional equation, but this determines the measure of products with $v > k$ that are stocked non-exclusively.)

(b) Suppose the optimal solution has $k > \bar{s}$.[49] Then there is no region of $v > k$ given our focus on $\bar{v} \leq \bar{s}$. The value of k now does not affect the product selection which is characterized by (35) with $F(k) = 1$. As stated in the text $\lambda = 0$ in this case. We can also deduce that $m = \bar{m}$ in the optimal solution. Suppose to the contrary that $m < \bar{m}$: (38) implies $\mu = 0$, and so (17) is increasing in q for all products which contradicts $m < \bar{m}$. If $m = \bar{m} < 1$ in the optimal solution then $\mu \geq 0$ is determined by $\bar{m} = \int qdG$, where $q = 1$

[46] According to Lemma 3, a sufficient condition for $k < \bar{s}$ in the optimal solution is that $\int qvdG/h(\int qdG) < \bar{s}$ for any q, or equivalently $\max_x \int_x^{\bar{v}} vdG/h(\int_x^{\bar{v}} dG) < \bar{s}$. This condition is easy to hold when economies of search are weak and $\bar{s} \gg \bar{v}$.

[47] Notice that if to the contrary $\mu < 0$, then by (35) all products should be stocked, which contradicts $m = \bar{m} < 1$. The same argument applies in parts (b) and (c) below when $m = \bar{m} < 1$.

[48] Note that the righthand side of (35) is weakly negative and that $\lambda k - \mu \geq 0$, so this is consistent with $m = 1$. A similar argument applies in part (b) below when $m = \bar{m} = 1$.

[49] According to Lemma 3, $k > \bar{s}$ in the optimal solution requires $\int qvdG > h(\bar{m})\,\bar{s}$.

for the products satisfying (35). If instead $m = \bar{m} = 1$ then μ does not matter and so we can without loss of generality set $\mu = 0$.

(c) Suppose the optimal solution has $k = \bar{s}$.[50] Again, there is no region of $v > k$ in this case. If $m < \bar{m}$ in the optimal solution, (λ, μ, m) jointly solve equations (14) with $k = \bar{s}$, (15) and (38). If $m = \bar{m} < 1$ in the optimal solution, λ and $\mu \geq 0$ jointly solve (14) with $k = \bar{s}$ and $\bar{m} = \int qdG$, where $q = 1$ for the products satisfying (35). If instead $m = \bar{m} = 1$ then μ does not matter so we can set $\mu = 0$, and any $\lambda \geq 0$ is consistent with the right hand side of (35) being weakly negative such that every product is indeed stocked. ■

Proof of Corollary 1. This is immediate from cases (a) and (c) in the proof of Proposition 2. (Note that we must have $m = \bar{m}$ in case (b).) ■

Proof of Proposition 3. We only give a sketch of the proof since it is very similar to that of Proposition 2. As before we treat $m = \int qdG$ as a constraint and set up the (Kuhn-Tucker) Lagrange function

$$\mathcal{L} = \int q(\pi F(k) + \lambda v - \mu)dG - \lambda k h(m) + \mu m + \eta(\bar{m} - m) \ , \tag{39}$$

where we again set $\lambda = 0$ if $k > \bar{s}$ in the optimal solution.

For given parameters (k, λ, μ) it is clear from (39) that the intermediary will stock product (π, v) if and only if

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

To determine the parameters it is again useful to consider a few cases separately. (a) Suppose the optimal solution has $k < \bar{s}$. If $m < \bar{m}$ in the optimal solution, we have $\eta = 0$, $\mu = \lambda k h'(m)$, and (k, λ, m) solve

$$k = \frac{\int qvdG}{h(m)}, \ \ \lambda = f(k)\frac{\int q\pi dG}{h(m)}, \ \ m = \int qdG \ . \tag{40}$$

If $m = \bar{m}$ in the optimal solution, η is irrelevant. If $\bar{m} < 1$ then (k, λ, μ) solve (40) with $m = \bar{m}$. If instead $\bar{m} = 1$ then without loss of generality we can set $\mu = 0$, and (k, λ) solve (40) with $m = 1$ and without the m equation. (b) Suppose the optimal solution has $k > \bar{s}$. Then as before we have $\lambda = 0$ and $m = \bar{m}$ in the optimal solution. If $\bar{m} < 1$ then μ is determined by $\bar{m} = \int qdG$. If instead $\bar{m} = 1$ we can without loss of generality set $\mu = 0$. (c) Suppose the optimal solution has $k = \bar{s}$. If $m < \bar{m}$ in the optimal solution,

[50] According to Lemma 3, $k = \bar{s}$ in the optimal solution requires $\int qvdG = h\,(m)\,\bar{s}$.

$\mu = \lambda k h'(m)$ and (λ, m) solve (40) with $k = \overline{s}$ and without the λ equation. If instead $m = \bar{m} < 1$ then (λ, μ) solve (40) with $k = \overline{s}$ and $m = \bar{m}$ and without the λ equation. If instead $m = \bar{m} = 1$ then λ solves the first equation in (40) with $k = \overline{s}$, $\mu = 0$ and $m = 1$. ■

Proof of Corollary 2. First, note that without DTC sales each manufacturer earns 0, whereas with DTC sales each manufacturer earns $\pi F(v)$, so total manufacturer profit strictly increases. Let $\mathbf{q^D}$ and $\mathbf{q^{ND}}$ be the optimal stocking policies respectively with and without DTC sales. Second, consider total industry profit and suppose DTC sales are feasible. Notice that if $\{\mathbf{q_E} = \mathbf{q^{ND}}, \mathbf{q_{NE}} = \mathbf{0}\}$ then compared with the case where DTC is unavailable: i) equations (34) and (20) imply that k is the same, and therefore industry profit on products with $q^{ND} = 1$ is also the same, but ii) industry profit is higher on goods with $q^{ND} = 0$. However in general the intermediary chooses $\mathbf{q^D} \neq \mathbf{q^{ND}}$, so its profit is higher than under $\{\mathbf{q_E} = \mathbf{q^{ND}}, \mathbf{q_{NE}} = \mathbf{0}\}$ whilst manufacturer profit is the same. Thus industry profit is (weakly) higher when DTC sales are feasible. Third, consider the intermediary's profit and suppose DTC sales are infeasible. Notice that if the intermediary stocks a product if and only if $q^D = 1$ then i) k is weakly higher because using equation (34) $\Delta\left(k^{ND}\right) \leq 0$ where k^{ND} is the solution to equation (20), and ii) compensation paid to manufacturers is weakly lower. Hence intermediary profit must be (weakly) higher when DTC sales are infeasible. ■

Proof of Proposition 4. The proof that the social optimum has $m > 0$ uses a similar technique to the one in the proof of Lemma 4 and so is omitted. To characterize the optimal stocking policy, substitute the expressions for $\Pi(\mathbf{q})$, $u^1(s, \mathbf{q})$ and $u^0(s, \mathbf{q})$ into equation (24) to get

$$\begin{aligned} W(\mathbf{q}) &= \int (1-q) \int_0^v (\pi + v - s) dF(s) dG + \int q(\pi + v) F(k) dG \\ &\quad - h(m) \int_0^k s dF(s) + \int_{v>k} q_{NE} \int_k^v (\pi + v - s) dF(s) dG \, . \end{aligned} \tag{41}$$

Maximizing $W(\mathbf{q})$ is the same as maximizing $W(\mathbf{q}) - W(\mathbf{0})$ i.e. the welfare improvement due to the intermediary. We therefore maximize $W(\mathbf{q}) - W(\mathbf{0})$ subject to the constraints (14), (15) and (16), and let λ, μ and η be the associated Lagrange multipliers.

After some algebraic manipulations, we can write the Lagrange function as follows:

$$\mathcal{L} = \int_{v<k} q \underbrace{\left[(\pi + v) \left[F(k) - F(v) \right] + \lambda v - \mu + \int_0^v s dF(s) \right]}_{[1]} dG$$

$$+ \int_{v>k} \left\{ q_E \underbrace{\left[(\pi + v) \left[F(k) - F(v) \right] + \lambda v - \mu + \int_0^v s dF(s) \right]}_{[2]} + q_{NE} \underbrace{\left[\lambda k - \mu + \int_0^k s dF(s) \right]}_{[3]} \right\} dG$$

$$- h(m) \left[\lambda k + \int_0^k s dF(s) \right] + \mu m + \eta(\bar{m} - m) \ . \quad (42)$$

This problem is similar to the profit-maximizing one. As before $\eta = 0$ if $m < \bar{m}$ in the optimal solution, and η is irrelevant if $m = \bar{m}$.

For given parameters (k, λ, μ), products with $v < k$ will be stocked if $[1] \geq 0$, which leads to (25). In this region, the exclusivity arrangement does not matter. Products with $v > k$ will be stocked exclusively if $[2] \geq \max\{0, [3]\}$, which leads to (26). When the opposite is true, they will be stocked non-exclusively if $[3] > 0$.

We now explain how to determine the parameters (k, λ, μ). Here we only consider the case where the optimal solution has $k < \bar{s}$. (The case of $k > \bar{s}$ or $k = \bar{s}$ can be easily dealt with like we did in the proof of Proposition 2.) Taking the first order condition of (42) with respect to k and then using equation (12) to cancel terms, we obtain

$$\lambda = f(k) \frac{\int_{v<k} q \pi dG + \int_{v>k} q_E \pi dG}{h(m) - \int_{v>k} q_{NE} dG} > 0 \ . \quad (43)$$

This is the same as equation (37) from the proof of Proposition 2. If $m < \bar{m}$ in the optimal solution, then $\eta = 0$ and so the first-order condition of (42) with respect to m gives

$$\mu = h'(m) \left[\lambda k + \int_0^k s dF(s) \right] \geq 0 \ , \quad (44)$$

Then (k, λ, μ, m) jointly solve equations (14), (15), (43) and (44). If $m = \bar{m}$ in the optimal solution, η is irrelevant as we have explained. Note that if $m = \bar{m} < 1$ then (k, λ, μ) jointly solve equations (14), (15) and (43) with m replaced by $\bar{m}$ everywhere, and we can conclude that $\mu \geq 0$ as in the proof of Proposition 2. If instead $m = \bar{m} = 1$ then μ does not matter so without loss of generality we can set $\mu = 0$, and (k, λ) jointly solve equations (14) and (43) with m replaced by $\bar{m}$. ■

Proof of Proposition 5. The intermediary maximizes (28) subject to the search constraint (27). The Lagrangian function of this optimization problem is

$$\begin{aligned}\mathcal{L} &= \int_{v<k} q\left[\pi F(k)+\lambda(v-k)\right]dG+\int_{v>k} q_{NE}\pi F(k)\,dG \\ &\quad +\int_{v>k} q_E\left\{\pi\left[3F(k)-2F(v)\right]+\lambda(v-k)\right\}dG\ . \qquad (45)\end{aligned}$$

The intermediary must stock some products exclusively otherwise from Lemma 2 no consumer will search it and it will earn zero profit. From equation (27) the optimum must have $k \in (\underline{v}, \bar{v})$, and so taking the first order condition of (45) with respect to k yields

$$\lambda = f(k)\,\frac{\int_{v<k} q\pi dG+\int_{v>k} q_{NE}\pi dG+3\int_{v>k} q_E\pi dG}{\int_{v<k} qdG+\int_{v>k} q_E dG}\ , \qquad (46)$$

from which we deduce $\lambda > 0$. Equations (27) and (46) jointly determine k and λ, and the remainder of the proposition follows immediately from equation (45). ■

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 (2017): "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.

Biglaiser, G., and F. Li (2018): "Middlemen: the good, the bad, and the ugly," *RAND Journal of Economics*, 49(1), 3–22.

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.

Gould, E., P. Pashigian, and C. Prendergast (2005): "Contracts, Externalities, and Incentives in Shopping Malls," *Review of Economics and Statistics*, 87(3), 411–422.

Hagiu, A., and B. Jullien (2011): "Why Do Intermediaries Divert Search?," *RAND Journal of Economics*, 42(2), 337–362.

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.

Rhodes, A., and J. Zhou (2017): "Consumer Search and Retail Market Structure," *Management Science*, forthcoming.

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

Shelegia, S. (2012): "Multiproduct Pricing in Oligopoly," *International Journal of Industrial Organization*, 30(2), 231–242.

Shevchenko, A. (2004): "Middlemen," *International Economic Review*, 45(1), 1–24.

Spiegler, R. (2000): "Extracting Interaction-Created Surplus," *Games and Economic Behavior*, 30(1), 142–162.

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.