

Yale University

EliScholar – A Digital Platform for Scholarly Publishing at Yale

Cowles Foundation Discussion Papers

Cowles Foundation

5-11-2020

Multiproduct Intermediaries

Andrew Rhodes

Makoto Watanabe

Jidong Zhou

Follow this and additional works at: <https://elischolar.library.yale.edu/cowles-discussion-paper-series>



Part of the [Economics Commons](#)

Recommended Citation

Rhodes, Andrew; Watanabe, Makoto; and Zhou, Jidong, "Multiproduct Intermediaries" (2020). *Cowles Foundation Discussion Papers*. 2579.

<https://elischolar.library.yale.edu/cowles-discussion-paper-series/2579>

This Discussion Paper is brought to you for free and open access by the Cowles Foundation at EliScholar – A Digital Platform for Scholarly Publishing at Yale. It has been accepted for inclusion in Cowles Foundation Discussion Papers by an authorized administrator of EliScholar – A Digital Platform for Scholarly Publishing at Yale. For more information, please contact elischolar@yale.edu.

MULTIPRODUCT INTERMEDIARIES

By

Andrew Rhodes, Makoto Watanabe, and Jidong Zhou

May 2020

COWLES FOUNDATION DISCUSSION PAPER NO. 2263



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

Multiproduct Intermediaries*

Andrew Rhodes

Toulouse School of Economics
University of Toulouse Capitole

Makoto Watanabe

Department of Economics
VU Amsterdam

Jidong Zhou

School of Management

Yale University

May 2020

Abstract

This paper develops a new framework for studying multiproduct intermediaries when consumers demand multiple products and face search frictions. We show that a multiproduct intermediary is profitable even when it does not improve consumer search efficiency. In its optimal product selection, it stocks high-value products exclusively to attract consumers to visit, 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. As applications we use the framework to study the optimal design of a shopping mall, and the impact of direct-to-consumer sales by upstream suppliers on the retail market.

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

JEL classification: D83, L42, L81

*We are grateful for helpful comments to the Editor and an anonymous referee, as well as Mark Armstrong, Heski Bar-Isaac, Alessandro Bonatti, Joyee Deb, Paul Ellickson, Doh-Shin Jeon, Bruno Jullien, Fei Li, Barry Nalebuff, Martin Obradovits, Rob Porter, Jerome Renault, Patrick Rey, Mike Riordan, Greg Shaffer, Andy Skrzypacz, Guofu Tan, Greg Taylor, Raphael Thomadsen, Glen Weyl, Mike Whinston, Chris Wilson, Julian Wright and seminar participants in Bonn, Cambridge, Cornell, Durham, HKUST, MIT, MSU, NUS, NYU Stern, Oxford, Stanford, Tokyo, TSE, UC Davis, UCL, UCLA Anderson, UIUC, USC, Yale, Zurich as well as the 8th Consumer Search and Switching Workshop, Bristol IO Day, EARIE, EEA (Lisbon), ICT conference (Mannheim), SAET (Faro), SICS (Berkeley), TNIT (Microsoft), the 16th Annual Columbia/Duke/MIT/Northwestern IO Theory Conference, and the 19th CEPR IO Conference. Rhodes gratefully acknowledges funding from the European Research Council (grant No 670494) and the Agence Nationale de la Recherche (grant ANR-17-EURE-0010).

1 Introduction

Intermediaries are important players in the economy.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 new framework to study multiproduct intermediaries when consumers demand multiple products and have search frictions. We then use this framework to address several important questions. For example, how can a multiproduct intermediary create value and therefore profitably exist? Which products should the intermediary carry, and for which of them should it be the exclusive supplier in the market? Surprisingly there is very little existing research about this, even though in practice the decision of which products to carry is one of the most important choices faced by intermediaries.² We also ask, 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. What are the possible consequences of this new trend, and how should intermediaries respond?

Our framework can shed light on some important developments in retail (and other intermediary) markets. Traditionally, most manufacturers could only reach consumers via retailers. However in recent years this has changed. Manufacturers of many different types of product now sell direct-to-consumer (DTC).³ A survey by the European Commission (2017) found that the percentage of EU manufacturers selling direct is as high as 85% in clothing and shoes, and above 50% in five other product categories. Many established brands like Nike and Nestle sell direct via their own website or physical store, while smaller brands do so via online marketplaces like Amazon and Tmall.⁴ (Direct sales are

¹Krakovsky (2015) suggests that middlemen are responsible for around one third of US GDP.

²Writing in the Harvard Business Review, Fisher and Vaidyanathan (2012) argue that successful retailers must be “good on a number of dimensions... But assortment is number one.” A good product assortment is crucial for attracting shoppers, who are often time-constrained and therefore shop at a few retailers whose product range closely meets their needs. At the same time, even large retailers like Walmart cannot stock all products desired by consumers, either due to space constraints, or because stocking too many products makes the instore shopping experience less pleasant.

³As an example, in 2017 DTC wine sales in the US were worth \$2.7 billion, almost 10% of the total (US) market (see <https://bit.ly/30qNQVq>). DTC sales are also becoming popular in industries such as traditional packaged goods, consumer electronics, and travel services (see <https://bit.ly/2VKiB8e>).

⁴Nike’s DTC sales in 2017 were worth \$10.4 billion, and it set itself the goal of increasing them to \$16 billion by 2020 (see <https://on.wsj.com/2ITY5wN> and <https://bit.ly/2DgV706>). According to the European Commission (2017), 20% of respondent manufacturers sell via marketplaces. Meanwhile the consulting firm Oliver Wyman (2018) found that more than half of US online sales of general merchandise in 2016 were direct (via either the manufacturers’ own websites or marketplaces).

also gaining traction in other markets, such as those for TV platforms.) Commentators have argued that DTC will grow further in the coming years and risks undermining traditional retailers' one-stop shopping business model. Some have even linked DTC to the closure of certain retailers and the sluggish performance of others.⁵ Our paper offers a framework in which to assess these claims, and to study how retailers should respond as DTC sales become more prevalent.

Another important trend in retail markets is exclusivity. Exclusivity agreements between manufacturers and retailers are increasingly common, and already exist for a wide variety of products including toys, furniture, and even groceries.⁶ Macy's has signed exclusive deals with several brands, and set itself the goal of increasing the percentage of products that are unique to its stores from 29% in 2017 to 40% in 2020. Similarly, Target is well-known for its exclusive deals with designers of apparel and homeware.⁷ Our paper offers a framework to think about the incentives to enter into exclusivity arrangements, which products should be stocked exclusively, and how DTC and other changes in technology may affect the propensity to use exclusivity agreements.

Section 2 of our paper introduces the baseline model. It is mainly developed for retailers but, as we discuss later, the setup and some of the main insights also apply to a 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. Consumers have the same demand for a given product, but different products have different demands. There is also one 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 compensates a manufacturer whose product it carries by way of a two-part tariff, and can request exclusive sales rights. The intermediary may also be limited in how many products it can stock. Consumers must incur a search cost to learn a seller's price(s) and buy its product(s), and the search cost is proportional to the time needed to visit it. We normalize the time needed to search a measure one of manufacturers to one, and assume the time needed to search the intermediary is weakly increasing in how many products it carries. (In the retail example, the latter would be consistent with the idea that larger retailers are located further from consumers, or are harder to navigate.) An important feature of our model is that consumers differ in their cost of time and so differ in their search cost.

Since the focus of our paper is product range choice, we intentionally simplify sell-

⁵See for example <https://on.wsj.com/2V2i6lx> and <https://bit.ly/2DgV706>, and also Section 6.

⁶A survey by GfK (2007) for the UK Competition Commission found that 35% of grocery-product suppliers had been asked to enter an exclusivity agreement, and 19% had done so. Gielens et al (2014) give examples of exclusive products by well-known manufacturers like Procter and Gamble and Unilever.

⁷See <https://bit.ly/2IjQ82U> and <https://bit.ly/2DaS2ys> for further details.

ers' pricing problems. In our model, irrespective of the market structure each supplier of a given product always charges the usual monopoly price.⁸ We then argue that all information about 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 simple 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 does not improve consumer search efficiency. The underlying mechanism is related to the classic bundling argument, and it relies on consumers' multiproduct demand and heterogeneous search costs and the availability of exclusive contracts. 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.

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 .⁹ 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. Compared to the social optimum, we argue that the intermediary tends to stock too many products, and to stock too many of

⁸Intuitively, with two-part tariffs the intermediary buys marginal units from a manufacturer at cost. Given that consumers have homogeneous preferences and do not observe price before search, the logic of Diamond (1971) then implies that 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.

⁹The result that exclusive products help attract consumers to visit but may incur a loss for the retailer is similar to, but different from, the usual loss-leader argument. Here the loss is due to the high compensation required by the manufacturer, instead of below-cost pricing.

them exclusively. Nevertheless the intermediary can be welfare-enhancing due to the way it affects consumers' incentives to search.

Section 5 studies two applications of our framework. Section 5.1 discusses the optimal design of a shopping mall which acts as a platform and does not directly resell products. Our framework provides insights about which shops should join the mall and how much they should pay to do so, as well as the externalities they exert on other shops. Section 5.2 applies our framework to DTC sales. When DTC sales become easier the intermediary has to compensate manufacturers more for stocking their products, and so not surprisingly manufacturers are better off but the intermediary is worse off. More importantly, we also show that the intermediary should respond to easier DTC sales by stocking fewer products, but stocking a greater proportion of them exclusively. As we explain in more detail later on, we also find that the optimal product selection becomes more 'polarized' in the (π, v) space, and both a product's easiness to be sold direct and its location in the (π, v) space affect whether (and how much) its direct sales take off. Moreover, easier DTC sales in one sector can induce manufacturers in other (unaffected) sectors to also sell more direct, due to a spillover effect. Finally, we show by example that if the intermediary fails to adjust its product selection in this way, it can end up with negative profit and thus have to exit the market altogether.

The remainder of the paper is then structured as follows. Section 6 provides evidence for the key economic forces in our model and its main predictions. In Section 7 we first provide a foundation for the (π, v) product space and interpret different points within it. We then present two extensions and show that our main insights are qualitatively robust to demand heterogeneity and upstream competition. Finally Section 8 concludes and discusses the potential application of our framework to other types of intermediary. All proofs are available in the Online Appendix.

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. Our model features consumers demand-

¹⁰Other 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.

ing multiple products and having heterogeneous search costs and the possibility of using exclusive contracts.¹¹ These features 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 induces 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 (2019)). 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 (2020), although 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

¹¹In 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.

¹²Rhodes and Zhou (2019) also study retailers’ endogenous product range, but they consider a stylized merger setup with only two products.

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 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

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

Manufacturers can sell their product direct to consumers, for example via their own retail outlet.¹⁴ In addition there is one multiproduct 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

¹³Bronnenberg (2020) 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.

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

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

to stock them exclusively or 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 visit a measure one of manufacturers 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 $h(m) < m$ the intermediary generates economies of search, and when $h(m) > m$ 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 and 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.

2.1 Preliminary analysis

We start with the following useful result.

¹⁶This 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.

¹⁷Our assumptions here try to capture the idea that a retailer's product range is usually reasonably steady over time, whilst its prices are easier to adjust.

¹⁸In Section 5 we allow the time needed to buy from a manufacturer to vary across products.

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 a wholesale unit price $\tau_i = c_i$ and a lump-sum payment $T_i = \pi_i F(v_i)$. 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 the manufacturer $\tau_i = c_i$ and T_i at a level that ensures 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. Consumers then face a standard hold-up problem (see e.g. Stiglitz (1979) and Anderson and Renault (2006)). Since consumers only learn the manufacturer's price after they have sunk their search cost, the manufacturer optimally charges the monopoly price p_i^m .¹⁹ Consumers rationally anticipate this and therefore search if and only if $s \leq v_i$. Consequently the manufacturer earns a profit $\pi_i F(v_i)$. 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 subtler. Intuitively 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 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

¹⁹As 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.

suppose it is compact and convex. Let $\underline{v} \geq 0$ and $\bar{v} < \infty$ be the lower and the upper bound of v . For each $v \in [\underline{v}, \bar{v}]$, there exist $0 \leq \underline{\pi}(v) \leq \bar{\pi}(v) < \infty$ such that $\pi \in [\underline{\pi}(v), \bar{\pi}(v)]$. (In Section 7.1 we provide examples of demand functions which can generate this type of product space.)

Let (Ω, \mathcal{F}, G) be a probability measure space where \mathcal{F} is a σ -field which is the set of all measurable subsets of Ω according to measure G . (In particular, $G(\Omega) = 1$.) When there is no confusion, we also use G to denote the joint distribution function of (π, v) , and let g be the corresponding joint density function. We assume that g is differentiable and strictly positive everywhere. If a consumer buys a set $A \in \mathcal{F}$ of products at their monopoly prices, she obtains surplus $\int_A v dG$ before taking into account the search cost. To avoid trivial corner solutions, we also assume that $\bar{v} \leq \bar{s}$.

Discussion. Before 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.* We assume that consumers need all products and have the same demand for a given product. (Different consumers will buy different products in our model, but only because they differ in their search cost.) This ensures that all sellers of a given product charge the same (monopoly) price, thereby allowing us to represent products using the (π, v) approach. In practice, however, consumers differ in which products they want to buy (and how much). Section 7.2 shows how to incorporate demand heterogeneity while still using the (π, v) approach, and demonstrates that the key insights are unchanged.

(iii) *An alternative interpretation of the demand function.* We assume that consumers have elastic demand for each product. Suppose instead that they have unit demand but heterogeneous valuations which are drawn independently across products. Suppose also that consumers make an upfront decision of whether or not to search the intermediary, and that they can only learn their valuation for a product by searching (one of) its seller(s).²¹ Then the (π, v) approach is still valid, as is the equilibrium we derive below. (That is, if the valuation distribution of product i is D_i , the same analysis works with $Q_i(p_i) = 1 - D_i(p_i)$.) Under this interpretation, consumers with the same search cost will buy different products.

²⁰See footnote 25 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.

²¹If consumers observe their valuations before they search, the market can collapse due to the well-known hold-up problem unless the seller sells enough products (see e.g. Rhodes (2015)).

(iv) *An instore-search microfoundation for $h(m)$.* Suppose there is a fixed cost κ_0 of traveling to the intermediary, and a cost κ_1 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 high- v products last, the intermediary can induce every consumer who visits to search all its products.²² 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 when $\kappa_1 > 0$.

(v) *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 matters is that each product's retail price is the same irrespective of where it is sold. (For instance there could be 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 range choice.

3 A Simple Case

We now study the intermediary's optimal product range choice. We start with a simple 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 important 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 $A \in \mathcal{F}$ exclusively. If a consumer of type s visits the intermediary, she will buy all products available there and so obtain an additional utility $\int_A v dG$. At the same time since the size of the intermediary is $m = \int_A dG$, the consumer also incurs an additional search cost $s \int_A dG$ given $h(m) = m$. Consequently a consumer of type s visits the intermediary if and only if $s \leq \hat{v}$, where

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

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

We now write down the intermediary's optimization problem. The intermediary's net profit from stocking product (π, v) is $\pi [F(\hat{v}) - F(v)]$: it attracts a mass $F(\hat{v})$ of

²²Further details are available on request. This is related to the idea of search diversion in Hagiu and Jullien (2011).

consumers and so earns variable profit $\pi F(\hat{v})$, but from Lemma 1 makes a lump-sum transfer $\pi F(v)$ to the product's manufacturer. Consequently the intermediary wishes to²³

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

The following simple observation will play an important role in subsequent analysis: among the products stocked by the intermediary, those with $v < \hat{v}$ generate a profit while those with $v > \hat{v}$ generate a loss. Intuitively a product with $v < \hat{v}$ 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 average surplus \hat{v}). The opposite is true for a product with $v > \hat{v}$, i.e. its demand is shrunk when sold through the intermediary.

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 $h(m) = m$, i.e. its search technology is no more efficient than that of the manufacturers whose products it resells. To understand why, recall that the intermediary makes a gain on low- v products but a loss on 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 product space 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 $(\pi, v) \in A$. The intermediary's problem then becomes

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

²³When $\int_A dG = 0$ intermediary profit is zero regardless of how we specify \hat{v} . In some later analysis we consider limit cases where the measure of A goes to zero, and \hat{v} will be well-defined via L'hospital's rule.

where the average consumer surplus \hat{v} offered by the intermediary solves

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

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. In particular, letting λ be the Lagrange multiplier associated with the constraint (4), the Lagrangian function is

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

The ‘value’ of stocking product (π, v) can be decomposed into a direct and an indirect effect. The direct effect is the profit $\pi [F(\hat{v}) - F(v)]$ earned on the product. The indirect effect $\lambda(v - \hat{v})$ captures cross-product externalities. Specifically, it measures how stocking product (π, v) changes the number of consumers who search the intermediary and thereby affects the profits earned from other products. We show below that $\lambda > 0$ and hence the direct and indirect effects have opposite signs. For example products with $v > \hat{v}$ have a negative direct effect (as explained above), but a positive indirect effect because stocking them increases the average surplus offered by the intermediary and so attracts more consumers to search it and buy other products.

The integrand in (5) is linear in q so the optimal stocking policy is as follows:

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

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

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

and

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

(The intermediary is indifferent about whether or not to stock products with $v = \hat{v}$.) The intermediary’s optimal product selection consists of two “negatively correlated” regions in the product space. First, the intermediary stocks products with high v and low π . Products with $v > \hat{v}$ induce consumers to search but also generate a loss. The intermediary minimizes this loss by stocking those with the lowest possible π , as this entails lower payments to manufacturers. Second, the intermediary stocks products with low v and high π . Products with $v < \hat{v}$ make a profit, which the intermediary maximizes by stocking those with the highest possible π . Products in other parts of the (π, v) 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 \hat{v} and λ . Firstly, at the optimum \hat{v} is interior and $F(\hat{v}) \in (0, 1)$.²⁴ Then we can take the first-order condition of (5) with respect to \hat{v} , and obtain

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

whereupon we observe that $\lambda > 0$. (Note that λ captures the impact on profit of a small increase in \hat{v} , and it equals $f(\hat{v})$ multiplied by the average profit of the intermediary's products.) Secondly, we have the original constraint (4), which we can rewrite as

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

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.²⁵

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

Figure 1 depicts the intermediary's optimal product range when the product space is $\Omega = [0, 1]^2$, the distribution over it is $G(\pi, v) = \pi v$, and the search cost distribution is $F(s) = s$. The intermediary stocks the products in the shaded areas, and total industry profit is 12.5% higher than it would be with no intermediary.²⁶

²⁴This is because Lemma 2 shows that $I(\hat{v}, \lambda)$ has strictly positive measure, which from the definition of \hat{v} implies $\hat{v} \in (\underline{v}, \bar{v})$, and moreover by assumption $s \in (0, \bar{s}]$ where $\bar{s} \geq \bar{v}$.

²⁵If we consider a discrete number of products $\{(\pi_i, v_i)\}_{i=1, \dots, N}$, the intermediary's problem becomes $\max_{q_i \in \{0, 1\}} \sum_i q_i \pi_i [F(\hat{v}) - F(v_i)]$ with $\hat{v} = \sum_i q_i v_i / \sum_i q_i$. 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 \hat{v} and λ are then complicated functions of the locations of individual products in the product space.

²⁶The curve $\lambda \frac{\hat{v} - v}{F(\hat{v}) - F(v)}$ which divides up the product space is constant when search costs are uniformly distributed. Otherwise it slopes up when $F(s)$ is concave, and slopes down when $F(s)$ is convex.




Figure 1: Optimal product range in the simple case

Finally, we discuss the welfare impact of the intermediary in this simple case. The intermediary harms consumers because they do not enjoy any search efficiencies (due to $h(m) = m$) and because they have less choice (due to exclusivity). Specifically, consumers with $s < \hat{v}$ who search the intermediary buy some low- v products which they otherwise would not buy. Similarly consumers with $s > \hat{v}$ 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 the above example, where total welfare increases by about 2.5%. Intuitively, consumers search too few manufacturers: they search and buy only if $s < v$, but from a 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 and so its product selection is not socially optimal. We discuss the socially optimal product selection in the next section.

4 The General Case

We now consider 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 $q_E(\pi, v)$ be an indicator function which is 1 if and only if product (π, v) is stocked exclusively, and let $q_{NE}(\pi, v)$ be an indicator function which is 1 if and only if product (π, v) is stocked non-exclusively. Let $\mathbf{q}(\pi, v) = (q_E(\pi, v), q_{NE}(\pi, v))$ denote the stocking policy function, and note that $\mathbf{q}(\pi, v) \in \{(0, 0), (0, 1), (1, 0)\}$. 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 only do so 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 and has a unit search cost s gets an expected surplus

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

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 time, a consumer of type s who does not search the intermediary gets expected surplus

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

because she can only buy products which are available from their manufacturers (i.e. those not stocked exclusively by the intermediary). This surplus is lower when the intermediary stocks more products exclusively (precisely, for a fixed q more products have $q_E = 1$). In order to ease the exposition, we suppose that a consumer searches the intermediary if and only if doing so strictly improves her payoff (i.e. if $u^1(s, \mathbf{q}) > u^0(s, \mathbf{q})$). The optimal search rule is as follows:

Lemma 3 *Consumers search the intermediary if and only if $s < \hat{s}$, where*

- (i) $\hat{s} = 0$ (nobody searches the intermediary) if $\int q_E dG = 0$ and $\int qdG \leq h(\int qdG)$.
- (ii) $\hat{s} > \bar{s}$ (everybody searches the intermediary) if $\int qvdG > h(\int qdG)\bar{s}$.
- (iii) $\hat{s} \in (0, \bar{s}]$ otherwise and is the unique strictly positive solution to

$$\hat{s} = \frac{\int_{v<\hat{s}} qvdG + \int_{v>\hat{s}} q_E vdG}{h(\int qdG) - \int_{v>\hat{s}} q_{NE} dG}. \quad (12)$$

Intuitively in this general case a consumer's decision of whether to search the intermediary is influenced by both the products it stocks (and their exclusivity) and the economy of search it generates. 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.

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, 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. Consumers with low s would like to buy most products anyway and so the latter disadvantage is small.²⁷

Finally, notice that \hat{s} in (12) degenerates to the average surplus \hat{v} defined in (2) if all stocked products are exclusive and $h(m) = m$. We use a different notation in this general case to emphasize that the threshold is now no longer the simple average surplus of the products sold at the intermediary.

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

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

Firstly, the intermediary earns $\pi[F(\hat{s}) - F(v)] > 0$ on products with $v < \hat{s}$ irrespective of whether they are stocked exclusively or non-exclusively (i.e. only $q = q_E + q_{NE}$ matters). Intuitively, even under non-exclusivity the manufacturer makes zero direct sales, because consumers with $s < \hat{s}$ buy its product from the intermediary and consumers with $s \geq \hat{s}$ are not willing to search it. Hence the intermediary earns variable profit $\pi F(\hat{s})$ and pays the manufacturer its outside option $\pi F(v)$. Secondly, the intermediary earns $\pi[F(\hat{s}) - F(v)] < 0$ from stocking products with $v > \hat{s}$ exclusively, and the explanation is the same as in the simple case. Lastly, the intermediary earns zero profit from stocking products with $v > \hat{s}$ non-exclusively. The reason is that when a manufacturer signs a non-exclusive contract, consumers with $s < \hat{s}$ switch and buy its product from the intermediary, but consumers with $s \in (\hat{s}, v)$ continue to buy direct. Hence the intermediary only needs to compensate the manufacturer by $\pi F(\hat{s})$ which equals its own revenue from

²⁷Notice that consumer search behavior is affected only by the measure (and not the identity) of non-exclusive products with $v > \hat{s}$. This is because consumers with $s < \hat{s}$ would buy these products anyway, so making them available at the intermediary only changes the search cost associated with buying them.

selling that product. Note that although the intermediary breaks even on these products, it may stock them in order to influence consumer search behavior.

The following lemma provides a simple sufficient condition under which the intermediary is guaranteed to earn strictly positive profit.

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

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

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

$$m = \int q dG. \quad (15)$$

The stocking space constraint can be written as

$$m \leq \bar{m}. \quad (16)$$

It is again convenient to solve the intermediary's problem using the Lagrangian method, treating (14)-(16) as constraints. Let λ , μ and η be their respective Lagrange multipliers (and note that $\lambda = 0$ if $\hat{s} > \bar{s}$ because in that case the constraint (14) does not apply). After some manipulations we can write the (Kuhn-Tucker) Lagrange function as

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

It is again useful to decompose the Lagrange function into direct and indirect effects of stocking product (π, v) . The direct effect is the profit generated by the product: recall from earlier that it is zero if the product is non-exclusive and $v > \hat{s}$, and otherwise equals $\pi [F(\hat{s}) - F(v)]$. The indirect effect is the change in profits from other products due to changes in consumer search behavior: it equals $\lambda \hat{s} - \mu$ if the product is non-exclusive and $v > \hat{s}$, and otherwise is $\lambda v - \mu$. Since the integrands are again linear in stocking variables, 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 $\hat{s} > 0$, $\lambda \geq 0$ and $\mu \geq 0$ defined in the proof):*

(i) *Products with $v < \hat{s}$ are stocked if and only if*

$$\pi \geq \frac{\mu - \lambda v}{F(\hat{s}) - F(v)}, \quad (18)$$

and it does not matter whether these products are exclusive or not.

(ii) *Products with $v > \hat{s}$ are stocked exclusively if and only if*

$$\pi \leq \frac{\max\{\lambda\hat{s}, \mu\} - \lambda v}{F(\hat{s}) - F(v)}. \quad (19)$$

Of the other products with $v > \hat{s}$: if $\lambda\hat{s} > \mu$ all of them are stocked non-exclusively, if $\lambda\hat{s} = \mu$ some of them are stocked non-exclusively, and if $\lambda\hat{s} < \mu$ none of them are stocked.

As in the 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. An important difference is that now the intermediary may also stock high- v and high- π products non-exclusively.²⁸ Whether or not that happens depends on the sign of the indirect effect $\lambda\hat{s} - \mu$.²⁹ In general it appears hard to find primitive conditions for the sign of $\lambda\hat{s} - \mu$, 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 $\mu = \lambda\hat{s}h'(m)$. Hence in the region of $v > \hat{s}$, all products are stocked if $h'(m) < 1$, but only low- π products are stocked (exclusively) if $h'(m) > 1$.*

Corollary 1 shows that in an interesting special case of the model, when the intermediary generates (marginal) diseconomies of search its optimal product selection has two negatively correlated regions in the (π, v) space (as in the simple case).³⁰

According to Proposition 2 the intermediary is indifferent about stocking products with $v < \hat{s}$ exclusively or non-exclusively. (This is because exclusivity has no effect on either the direct profit in equation (13) or on consumer search behavior in Lemma 3). One way to break this indifference is to introduce some small-demand consumers who never

²⁸Another subtler difference is that unlike in the simple case it is possible that $\hat{s} < \underline{v}$ (when $\underline{v} > 0$ and diseconomies of search are strong) or $\hat{s} > \bar{v}$ (when economies of search are strong). In these cases only one part of Proposition 2 is relevant.

²⁹If $\lambda\hat{s} - \mu = 0$ the intermediary is indifferent about which of these non-exclusive products to stock (and the equation $\lambda\hat{s} - \mu = 0$ pins down their measure). In this case the optimal selection is not unique.

³⁰These two regions are also disconnected, because given $h'(m) > 1$, as $v \rightarrow \hat{s}$ the thresholds in (18) and (19) tend respectively to ∞ and $-\infty$. Intuitively products with $v \approx \hat{s}$ generate only a small direct profit or loss, so their negative effect on consumers' incentives to search the intermediary dominates.

visit the intermediary. In that case the intermediary strictly prefers to stock products with $v < \hat{s}$ non-exclusively, because it reduces the compensation paid to manufacturers. Therefore in subsequent analysis we interpret products with $v < \hat{s}$ as being non-exclusive.

We now illustrate the optimal product selection using some examples. First, we illustrate how the intermediary's search technology affects its product selection. Figure 2 considers an example with $F(s) = s$ and $G(\pi, v) = \pi v$ and no space constraint (i.e. $\bar{m} = 1$).³¹ The search technology is $h(m) = 1.2m$ in part (a), $h(m) = m$ in part (b), and $h(m) = 0.8m$ in part (c). Consistent with Corollary 1, when $h(m) = 1.2m$ only the low- π products with $v > \hat{s}$ are stocked (exclusively), whereas when $h(m) = 0.8m$ all products with $v > \hat{s}$ are stocked. (When $h(m) = m$ the intermediary is indifferent about stocking each product in $[0.5, 1]^2$ non-exclusively. The figure depicts the case where the intermediary stocks all these products.) Comparing across the three cases, as the intermediary's search technology becomes more efficient it stocks more products (m increases from 0.29 to 0.75 and then to 0.76) but a smaller proportion of them are exclusive (the percentage decreases from 50% to 33% and then to 27%). Intuitively, the intermediary relies less on exclusive products to attract consumers when it already helps to reduce their search costs.




Figure 2: Optimal product range and search technology

Second, we illustrate how the space limit \bar{m} affects product selection. Figure 3 considers an example with $F(s) = s$ and $G(\pi, v) = \pi v$ (as above), and also $h(m) = 0.4$ such that the cost of visiting the intermediary is independent of its size. In this example marginal economies of search are so strong that the intermediary always uses all its stocking space. In part (a) $\bar{m} = 0.3$ and the optimal solution has $\lambda \hat{s} - \mu < 0$, hence the intermediary stocks products in two negatively correlated and disconnected regions.

³¹This example with $h(m) = a + bm$ can be fully solved. A sufficient condition for the intermediary to exist is $a + b \leq 3/2$, which is weaker than the condition given in Lemma 4.

In part (b) $\bar{m} = 0.46$ and the optimal solution has $\lambda\hat{s} - \mu = 0$, hence the intermediary stocks some but not all of the products in the top-right corner non-exclusively. (As discussed in footnote 29, when $\lambda\hat{s} - \mu = 0$ the optimal solution in the top-right corner is not unique. In particular, only the measure of products that are stocked non-exclusively can be determined. Conditional on stocking the correct measure of products, the intermediary is indifferent over exactly which products in the top-right corner to stock. The figure depicts the case where the intermediary stocks those with the highest v .) In part (c) $\bar{m} = 0.5$ and the optimal solution has $\lambda\hat{s} - \mu > 0$, hence all products with $v > \hat{s}$ are stocked.³² Comparing across the three cases, as the intermediary becomes able to stock more products, it attracts more consumers (\hat{s} increases from 0.32 to 0.7 and then to 0.81) but stocks proportionately fewer exclusive products (the percentage decreases from 50% to 45.6% and then to 30.5%), while sales outside the intermediary fall (from 0.37 to 0.22, and then to 0.175). Notice that qualitatively the effect of a better search technology (Figure 2) and larger space limit (Figure 3) are the same.




Figure 3: Optimal product range and stocking space constraint

4.3 Socially optimal product range

We now briefly study the socially optimal product range. Suppose a social planner wishes to maximize total welfare, and it can choose the stocking policy \mathbf{q} but has no direct control over firm pricing or how consumers search. This problem can be solved in a similar way to the intermediary's. We therefore report the main result here and relegate the details to the online appendix.

³²The product selection is similar to Figure 3(a) for \bar{m} below around 0.454, similar to Figure 3(b) for \bar{m} between around 0.454 and 0.463, and similar to Figure 3(c) for \bar{m} between around 0.463 and 0.65. When \bar{m} exceeds around 0.65 the intermediary's size and strong economies of search enable it to attract all consumers, none of its products are exclusive, and only low- v and low- π manufacturers make direct sales.

Proposition 3 *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 socially optimal policy is characterized as follows (with $\hat{s} > 0$, $\lambda \geq 0$ and $\mu \geq 0$ defined in the proof):*

(i) *Products with $v < \hat{s}$ are stocked if and only if*

$$\pi + v \geq \frac{\mu - \lambda v - \int_0^v s dF(s)}{F(\hat{s}) - F(v)}, \quad (20)$$

and it does not matter whether they are exclusive or not.

(ii) *Products with $v > \hat{s}$ are stocked exclusively if and only if*

$$\pi + v \leq \frac{\max\{\mu, \lambda \hat{s} + \int_0^{\hat{s}} s dF(s)\} - \lambda v - \int_0^v s dF(s)}{F(\hat{s}) - F(v)}. \quad (21)$$

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

The welfare-optimal stocking policy is qualitatively the same as the one used by the intermediary. Firstly, exclusive products with $v > \hat{s}$ are again chosen to have low π , and products with $v < \hat{s}$ 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 $v > \hat{s}$ is further reduced when it is sold exclusively by the intermediary, but conversely demand for a product with $v < \hat{s}$ 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, we can show that when the stocking constraint is slack at the optimum all products with $v > \hat{s}$ are stocked if $h'(m) < 1$, but only those with low π are stocked (exclusively) if $h'(m) > 1$. Intuitively, when the intermediary adds some non-exclusive products with $v > \hat{s}$ all the consumers with $s < \hat{s}$ will switch their purchases of those products away from manufacturers towards the intermediary, and their total search cost falls if $h'(m) < 1$ but increases if $h'(m) > 1$.

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 (\hat{s}, λ, μ) 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 4(a) 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 (and is a subset of) that in Figure 2(a). Figure 4(b) plots the socially optimal product range for the case $h(m) = m$. It turns out that in this case (\hat{s}, λ, μ) are the same as in Figure 2(b). The social planner's product range is also a strict subset of the intermediary's, and it stocks fewer products exclusively. Figure 4(c) plots the socially optimal product range for the case $h(m) = 0.8m$. Here (\hat{s}, λ, μ) differ from Figure 2(c), and the social planner's selection is not a subset of the intermediary's. Nevertheless the intermediary again stocks fewer products overall and also fewer products exclusively.




Figure 4: 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. For example in the simple case with $h(m) = m$, a ban on exclusivity prevents the intermediary from existing and so strictly reduces total welfare.

5 Two Applications

We now apply our framework to examine the optimal design of a shopping mall (which acts as a platform and does not directly set prices) and the effect of DTC sales on retail markets. Here in each application we lay out various theoretical predictions and then document relevant empirical evidence in the following section of the paper.

5.1 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 mall can host up to a measure \bar{m}

of sellers, and charges each of them a fixed fee. Consumers pay $n \times s$ to search a measure n of independent shops and $h(m) \times s$ to search a mall which contains a measure m of shops. The timing is exactly as in Section 4 except that now the mall is a platform and therefore individual sellers set prices.

Our earlier analysis applies straightforwardly. Firstly, it is easy to see that seller i should 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 search the mall if and only if $s < \hat{s}$ where \hat{s} was defined earlier in Lemma 3. Thirdly, sales of a product (π, v) at the mall generate a gross profit of $\pi F(\hat{s})$. Therefore a seller with $v < \hat{s}$ will pay $\pi [F(\hat{s}) - F(v)]$ to join the mall, whilst a seller with $v > \hat{s}$ is willing to join the mall for free if it also maintains an independent store elsewhere, but must be paid $\pi [F(v) - F(\hat{s})]$ to exclusively join the mall. Consequently the mall owner's profit is the same as in equation (13).

The optimal mix of stores within the mall is exactly as in Proposition 2. Using some terminology from existing literature, we can interpret sellers with $v > \hat{s}$ 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 \hat{s}), which allows the mall to charge a higher fee to other 'non-anchor' stores.

5.2 Direct-to-consumer sales

We argued in the introduction that it has become increasingly easy for manufacturers to sell direct to consumers. Our framework can be extended to look at the consequences of easier DTC sales. To this end consider the general case from Section 4, but now index products by (π, v, θ) where $\theta > 0$ represents how easy it is for consumers to buy direct from the product's manufacturer. Specifically, a consumer with unit search cost s , if she visits a measure n of manufacturers with DTC index θ , incurs a total search cost $n \times \theta s$. Importantly we assume that θ can differ across products (so that direct sales are easier for products with a lower θ), and we say that all products with a common θ belong to the same 'sector'. We let $G(\pi, v, \theta)$ be the joint distribution over the product space, and $[\underline{\theta}, \bar{\theta}]$ be the support of θ . We also assume that conditional on θ the (π, v) space is compact and convex, and $\bar{v} \leq \underline{\theta} \bar{s}$ to avoid corner solutions.

We start by characterizing the intermediary's optimal product selection. We then consider what happens as DTC sales become easier in certain sectors. All the omitted details in this section can be found in the online appendix.

5.2.1 Optimal stocking policy

An important difference with the baseline model is that if the manufacturer of a product (π, v, θ) only sells direct to consumers, it sells to all consumers for whom $v > \theta s$. Hence

the manufacturer's outside option is $\pi F(v/\theta)$. As one would expect this outside option is higher when θ is lower i.e. when direct sales are easier. This already suggests that the intermediary will have less incentive to carry products which are easier to sell direct.

Analysis of consumer search and intermediary profit closely follows that from Section 4. Consumers visit the intermediary if and only if their search cost is below a threshold \hat{s} which is defined in a similar way to in Lemma 3 from earlier. The intermediary earns a positive profit on products with $v < \theta\hat{s}$ and exclusivity does not matter; for products with $v > \theta\hat{s}$ it breaks even when stocking non-exclusively, and makes negative profit when stocking exclusively.³³ We then obtain the following characterization of the optimal stocking policy:

Proposition 4 *The optimal stocking policy is characterized as follows (with $\hat{s} > 0$, $\lambda \geq 0$ and $\mu \geq 0$ defined in the proof):*

(i) *Products with $v < \theta\hat{s}$ are stocked if and only if*

$$\pi \geq \frac{\mu - \lambda v}{F(\hat{s}) - F(\frac{v}{\theta})}, \quad (22)$$

and it does not matter whether these products are exclusive or not.

(ii) *Products with $v > \theta\hat{s}$ are stocked exclusively if and only if*

$$\pi \leq \frac{\max\{\lambda\theta\hat{s}, \mu\} - \lambda v}{F(\hat{s}) - F(\frac{v}{\theta})}. \quad (23)$$

Of the other products with $v > \theta\hat{s}$: if $\lambda\theta\hat{s} - \mu > 0$ they are stocked non-exclusively, if $\lambda\theta\hat{s} = \mu$ some of them are stocked non-exclusively, and if $\lambda\theta\hat{s} - \mu < 0$ none of them are stocked.

Given a particular θ the conditions for whether and how a product is stocked are similar to those from the general case. Using Proposition 4 we can derive two implications for how the stocking policy differs across sectors. The first is:

Corollary 2 *Consider products with the same (π, v) . If the intermediary stocks (exclusively or non-exclusively) one in sector θ' , then it also stocks all those in sectors $\theta > \theta'$.*

Corollary 2 further implies that in the special case where the distribution of θ is independent of (π, v) then more products are stocked by the intermediary in higher- θ sectors. One reason (as alluded to above) is that manufacturers with higher- θ products have a worse outside option and so are cheaper for the intermediary to contract with. Another reason is that the indirect benefit to the intermediary of stocking higher- θ products is also higher, because by stocking these products the intermediary saves consumers who visit it more search costs. A second implication is the following:

³³Closely following Lemma 4 we can also show that a sufficient condition for the intermediary to profitably exist is that there is a positive measure of θ for which $h(\tilde{m}) \leq \theta\tilde{m}$ for some $\tilde{m} \in (0, \bar{m})$.

Corollary 3 *Suppose $m < \bar{m}$ at the optimum and consider products with $\theta \hat{s} \in (v, \bar{v})$.*

(i) Among those with $\theta < h'(m)$ none are stocked for v in a neighborhood of $\theta \hat{s}$, and none with $v > \theta \hat{s}$ are stocked non-exclusively. That is, in sectors where DTC sales are relatively easy, the intermediary only stocks products in two disconnected and negatively-correlated regions.

(ii) Among those with $\theta > h'(m)$ all are stocked for v in a neighborhood of $\theta \hat{s}$, and all with $v > \theta \hat{s}$ are stocked. That is, in sectors where DTC sales are relatively hard, the intermediary's product range is a connected set which includes all high- v products.

Corollary 3 suggests that in an interesting special case of the model, product selection is more ‘polarized’ in sectors where DTC is easier. Intuitively, the intermediary finds it relatively expensive to stock products with $\theta < h'(m)$ and also generates search economies on them. Hence it stocks exclusively a small measure of high- v and low- π products which are very successful in attracting consumers, and stocks a small measure of low- v and high- π products which are particularly profitable. On the other hand, the intermediary finds it relatively cheap to stock products with $\theta > h'(m)$ and also generates search economies on them. Hence the optimal product selection is very different.³⁴

5.2.2 Changes in the ease of DTC sales

We now consider the effect of improvements in technology which make DTC sales easier. Formally, DTC sales become easier when for each sector θ the DTC sales cost index for product (π, v, θ) decreases to $\tau(\theta; \pi, v) \leq \theta$, with a strict inequality for a positive measure of sectors.³⁵ We assume throughout that the cost of visiting the intermediary is unchanged, which is plausible when the intermediary is an offline retailer such as a supermarket or department store.

We begin with the following straightforward result:

Corollary 4 *When DTC sales become easier, manufacturer profit increases while the intermediary's profit decreases.*

Manufacturers clearly benefit since they always earn their outside option $\pi F(v/\theta)$, and this strictly increases for any manufacturer whose product becomes strictly easier

³⁴In the limit case where $\theta \rightarrow \infty$ for all products (and so DTC sales are impossible), the intermediary stocks products with $\pi \geq (\mu - \lambda v)/F(\hat{s})$, i.e. those with high π and high v . Intuitively the intermediary does not need to compensate manufacturers beyond their marginal cost, and so carries products with high v to attract consumers and high π to maximize variable profits.

³⁵Here we implicitly assume that DTC sales become easier for exogenous reasons e.g. due to the development of online marketplaces. In practice manufacturers can also invest in direct channels. Fixing the investment cost, and under some regularity conditions, manufacturers with high- v and high- π products are more likely to invest i.e. $\tau(\theta; \pi, v)$ falls more for those products.

to sell direct. On the other hand the intermediary is made worse off. It is easy to see that if the intermediary does not change its stocking policy it becomes worse off, both because it has to compensate manufacturers more, and because easier DTC sales leads to a reduction in the measure of consumers that use the intermediary; Corollary 4 then shows that the intermediary is still worse off even after it adjusts its stocking policy.

We now discuss in more detail how the intermediary should adjust its stocking policy as DTC sales become easier. Unfortunately in general it is impossible to provide analytical results. Therefore we use an example to illustrate the main insight. Suppose that $\bar{m} = 1$, $h(m) = m$, (π, v) is on $[0, 1]^2$ uniformly, $F(s) = s/2$ with support $[0, 2]$,³⁶ and θ has a binary distribution, independent of (π, v) , with $\Pr(\theta = 1/2) = z$ and $\Pr(\theta = 3/2) = 1 - z$, where $z \in [0, 1]$. We vary the proportion z of products for which DTC sales are easy. As z increases more products fall in the low- θ sector, and therefore the intermediary carries fewer products as described in Figure 5(a), but a higher proportion of them are exclusive as described in Figure 5(b). Intuitively, when it becomes easier to sell certain products direct, their manufacturers must be compensated more and so it is less profitable for the intermediary to still carry them. At the same time, however, easier direct sales reduce consumers' need to search the intermediary, and so the intermediary relies more on product exclusivity to remain attractive. Figure 5(c) then plots the total measure of consumers buying direct (the increasing dashed curve), the total measure of consumers buying from the intermediary multiplied by the measure of products it stocks (the decreasing dashed curve), as well as the sum of these two things (the increasing solid curve). As expected, the plot suggests that easier DTC sales lead to an expansion of total sales but also diversion away from the intermediary towards the manufacturers.³⁷

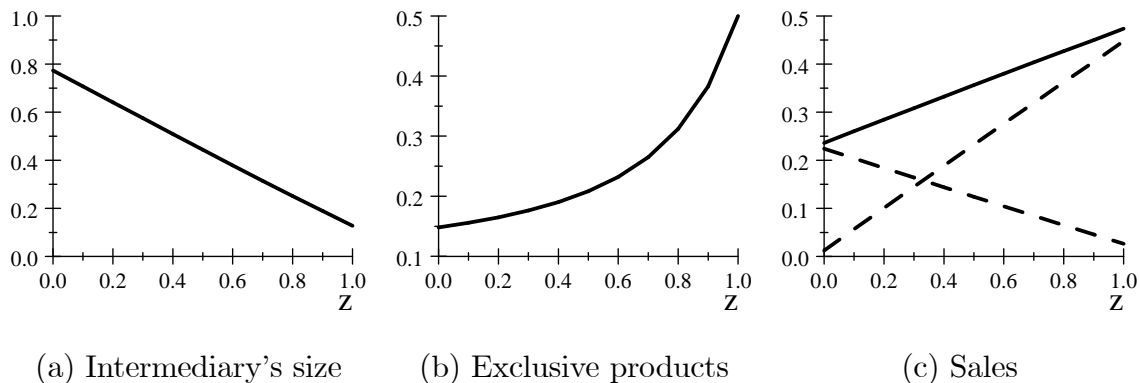


Figure 5: The impact of easier DTC sales

³⁶The support of s ensures that $\bar{v} \leq \theta \bar{s}$ in this example, and so avoids corner solutions.

³⁷The plot only imperfectly capture sales however, because different products in the (π, v, θ) space have different demands and so the amount bought differs.

Finally, we also point out some other interesting implications that come out of the above example. The first implication is that easier DTC sales in some sectors can have spillovers on other sectors. Specifically, we find that as z increases the intermediary attracts fewer consumers i.e. \hat{s} decreases. (Intuitively easier direct sales make the intermediary’s one-stop shopping service less attractive to consumers.) This implies that even in sectors where the DTC technology does not change, more consumers buy non-exclusive high- v and high- π products direct rather than via the intermediary. The second implication is that if the intermediary does not adjust its product selection, easier DTC sales may force it to close down entirely. For instance, if in the above example z increases from $1/5$ to $4/5$ but the intermediary does not change its stocking policy then it will make a loss of -0.049 , whereas if it reoptimizes it will make positive profit.³⁸

6 Empirical Evidence

In this section we document some evidence for the main economic forces in our model and the main predictions.

6.1 Cross-product effects

As we have already discussed, indirect or ‘cross-product’ effects play an important role in our model. Here we offer some evidence of their practical relevance.

Firstly, some studies have documented significant cross-category spillovers at the retail level. For example, Sen et al (2013) estimate that when a supermarket adds a gas station total trips to the store increase by 14%-15%, and gas-buyers spend 7.7%-9.3% more on groceries. They also argue that the resulting profit increase on groceries outweighs that from gas, since the former has significantly higher profit margins.³⁹ (Thus in the spirit of our model we might interpret gas as a low- π but high- v product, and groceries overall as a high- π product.) Meanwhile Gielens et al (2014) estimate the impact of a Dutch supermarket exclusively stocking five national brands. They find that in three of those cases the (absolute) cross-category sales effect dominates the own-category sales effect by a significant amount.

Secondly, and related, our model shows that some categories should be used to attract consumers and others to generate profit. Consistent with this, Briesch et al (2013) state

³⁸The impact of DTC sales on industry profit, consumer surplus and total welfare is difficult to investigate analytically. In the above example, when z increases from $1/5$ to $4/5$, all three increase.

³⁹Sen et al (2013) also point out in their footnote 3 that “conversations with managers at the retailer indicated that there was no effort to position gasoline as a loss-leader or undercut the prices of nearby gas stations”, and so the cross-category effect here is unlikely to be part of a usual loss-leader pricing strategy.

that “A fundamental tenet of category management is that individual categories play different roles”. They also estimate the utility that consumers get from different categories and find it to be very heterogeneous (see their Table 8), labeling those with highest utility as ‘destination categories’ (which in our model could be interpreted as having high v). In a similar spirit, Kök et al (2015) argue that a retailer like Best Buy might offer more variety in certain categories than would be optimal in isolation, so as to attract more consumers and increase sales from other categories. Similarly Fisher and Vaidyanathan (2012) give the example of a supermarket which dropped low-selling (and, one might infer, low-profit) products and then promptly lost a large number of customers and hence sales of other profitable goods.

Finally, evidence from shopping mall contracts demonstrates the importance of cross-seller effects. Recall Section 5 where we applied our framework to mall design. Consistent with our discussion there, Gould et al (2005) show that adding more ‘anchor stores’ to a mall increases both the sales and rents paid by ‘non-anchor stores’. (Thus, again in the spirit of the model, we might interpret anchor stores as high- v sellers and non-anchor stores as low- v sellers.) Their results also suggest that anchor stores are subsidized to join the mall in anticipation of these positive spillovers. Specifically, they document that 73% of anchor stores pay no rent while others pay substantially reduced rents (\$4.13 vs \$29.37 per square foot), while at the same time the mall often pays for things like development and maintenance costs.

6.2 Product exclusivity

As we demonstrated in the Introduction, exclusivity is increasingly important in retail markets. Here we offer some evidence in favor of the trade-offs generated by exclusivity in our model, and offer some more speculative evidence concerning our model’s predictions about which (and how many) products should be offered exclusively.

In our model, a manufacturer typically sells fewer units if it accepts an exclusive contract. However when used optimally, exclusivity brings more consumers to the intermediary and raises joint bilateral profit. Gielens et al (2014) study empirically the effects of five product exclusivity arrangements in a Dutch supermarket. They find that in each case the exclusivity arrangement benefits the retailer through higher store-wide sales, and that in all but one case harms the manufacturer through lower total sales. They also find that in four out of the five cases joint profit is increased by the arrangement.

Our model also makes some predictions about which products should be sold exclusively, and when the intermediary’s product range should contain more exclusive products. Unfortunately we are not aware of any empirical research that directly speaks to these predictions. What does seem to be true is that traditionally exclusivity was associated only with luxury goods, whereas it is increasingly being used in a wide variety of cate-

gories including in groceries. (See, e.g., Gielens et al (2014) and Upshaw et al (2017).) We hope that our (π, v) framework offers a new way to think about exclusivity. According to our model, what matters for whether a product should be sold exclusively is its ability to attract consumers and its profit margin when sold alone, instead of whether it is a necessity or a luxury good.

One type of exclusive product which has been studied in the literature is private labels. Private labels are increasingly high-quality and are often argued to increase store traffic,⁴⁰ and so play the same role as exclusive products in our model. Our model suggests that exclusive products should have high- v and low- π , and as we discuss in more detail in Section 7.1 below, this suggests they should have higher cost pass-through. Consistent with this, Hong and Li (2017) find that private labels have higher pass-through. (However this should be interpreted with caution since pass-through rates depend on many factors.) Our model also suggests that smaller retailers should stock proportionately more exclusive products. Consistent with this, some successful small retailers like Trader Joe’s have very large high-quality private label shares. (Trader Joe’s stocks around 4000 items compared to 50000 at a typical grocery store, and 80% of them are private label. See Sorescu et al (2011).) Along the same lines, Briesch et al (2013) present some descriptive statistics showing that Walmart (a mass merchandiser) has a lower level of private labels and higher level of national brands compared to four smaller supermarkets.

6.3 Retail technology and stocking space

As pointed out by Hortaçsu and Syverson (2015), starting in the late 1990s improvements in IT and logistics led to a rapid increase in the number of very large retailers such as warehouse clubs and supercenters. At the same time, sales at smaller traditional retailers decreased. The comparative static exercise with respect to \bar{m} (as in the example in Figure 3) offers a convenient way to think through these developments.⁴¹ Specifically, imagine that better technology enables the intermediary to better manage and so stock more products, and suppose we interpret direct sales more broadly as sales by small specialist retailers. The example in Figure 3 suggests that sales by specialist retailers do indeed fall as \bar{m} increases, because consumers switch their purchases towards the intermediary. A novel feature is that this occurs even though the intermediary in our model does not use aggressive pricing strategies against the smaller competitors - the result is driven purely by consumer search frictions and the improvement in the intermediary’s retail technology.

⁴⁰See, e.g., <https://bit.ly/2vEzVNq>.

⁴¹We thank the referee for pointing out this application.

6.4 DTC sales

In Section 5 we used our framework to study how the retail landscape changes as DTC sales become easier. Traditionally direct channels were rare, except in a few sectors like apparel and high-end technology (Oliver Wyman, 2018). However, in recent years, direct channels have appeared for a wide variety of products. According to a report by the OECD (2019), “growth in e-commerce has prompted a concomitant expansion into the online retail sphere by manufacturers, including many manufacturers without an equivalent brick-and-mortar retail presence.” One reason is that the Internet directly reduces consumer search costs, and makes it easier and cheaper for manufacturers to advertise their products. Another reason is that the Internet allows manufacturers to sell to consumers without making large investments (European Commission, 2017). Moreover if we interpret the search frictions more broadly to include delivery waiting times, improvements in distribution networks also make direct sales easier. Nevertheless some products are easier to sell direct than others. For example, selling direct is easier when products are standardized and easy to deliver, as is the case for traditional packaged goods, software, and financial and travel services. One contribution of our model is to show that whether a particular product should be sold direct depends not only on the ease of direct sales θ , but also on cost and demand factors as reflected in its (π, v) value.

We showed earlier that as DTC sales become easier, the market expands and sales are diverted away from the intermediary. Consistent with this, Duch-Brown et al (2017) estimate that the e-commerce channel reduced sales of portable media players at traditional retailers in Europe by 13.2%, but raised overall sales by 5.3%. (They provide similar evidence for digital cameras and portable computers). Unfortunately it is unclear what proportion of online sales in their study are direct or by specialist retailers. However other evidence of the expansion and diversion effects comes from Aguirregabiria et al’s (2016) study of the Ontario liquor market, where consumers can buy wine either from two specialist outlets (selling local wines) or from a large (generalist) liquor shop. They estimate that if the specialist wine stores were to close, the liquor shop would sell more wine yet total wine sales would fall.

Our model also suggests that as DTC sales become easier, the intermediary should become smaller but stock proportionally more exclusive products so as to remain attractive to consumers. We are not aware of any formal empirical evidence on the link between DTC sales and product selection. However, anecdotal evidence seems consistent with the model’s prediction. According to a World Economic Forum report (2017), newly opened supermarkets are 18% smaller than existing ones, and newly opened mass merchandise stores are typically half as large as existing ones. Many commentators have also emphasized that in order to succeed retailers “cannot afford to be just one-stop shops”, and

should have more exclusive brands.⁴² To this end, prominent retailers such as Macy’s and Target have been simultaneously shrinking their assortments and adding more exclusive products.⁴³

Finally, our model also suggests that easier DTC sales make it harder for the intermediary to attract customers and also reduces its profitability. Again we are not aware of formal empirical evidence about this impact of DTC. However some commentators have linked aggressive pursuit of DTC sales by manufacturers such as Nike to the closure of sports retailers like Boston’s City Sports, and declines in the share price of others such as Foot Locker and Dick’s. Consistent with our model, they have also argued that Dick’s has performed better than expected due to increased usage of private brands and other exclusively stocked products.⁴⁴

7 Discussion and Extensions

In this section we first discuss the foundation of the (π, v) product space, and then present two extensions: one with demand heterogeneity, and the other with upstream competition where each product has multiple varieties supplied by different manufacturers.

7.1 Foundation and interpretation of the (π, v) product space

We now explain how to construct the (π, v) space and interpret different points within it. Suppose that demand for product i can be written as $Q_i(p_i) \equiv Q(p_i, \delta_i)$ where δ_i is a vector of demand parameters. Using the definitions introduced in equation (1), one can calculate $\pi_i = \pi(\delta_i, c_i)$ and $v_i = v(\delta_i, c_i)$, and then derive the (π, v) space from the parameter space (δ, c) .⁴⁵

We now introduce a class of demand functions to illustrate how different points in the (π, v) space can be interpreted in terms of the size and shape of demand. Suppose that

⁴²See, e.g., <https://bit.ly/2VJyjk4> and <https://bit.ly/2DOPVkk>.

⁴³See, e.g., <https://bit.ly/2IjQ82U> and <https://on.wsj.com/2B2oCSq> for the product strategy adjustment of Macy’s and <https://on.wsj.com/2TojcfC> for Target.

⁴⁴See references in footnote 5 for Boston’s City Sports, as well as <https://bit.ly/2JGUZee> for discussion of Foot Locker and Dick’s.

⁴⁵If products differ in exactly two parameters, it is possible to have a one-to-one correspondence between the parameter space and the (π, v) space. However if products differ in more than two parameters, 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 respectively the fraction of exclusive and non-exclusive products at point (π, v) . Since our objective functions are linear in stocking policy variables, the solution is bang-bang with $q_E(\pi, v), q_{NE}(\pi, v) \in \{0, 1\}$ and so our analysis is unchanged.

product i has a constant-curvature demand function

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

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 $\sigma_i = 0$, and convex demand when $\sigma_i > 0$. It is straightforward to show that provided marginal cost is not too large, the monopoly price for product i is $p_i^m = 1 + \frac{c_i}{2 - \sigma_i}$, and π_i and v_i are both increasing in a_i and decreasing in c_i , and moreover $\pi_i/v_i = 2 - \sigma_i$. (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.⁴⁶) 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) determines how far the product is along the ray. Notice also that in this example the rate of cost pass-through equals $1/(2 - \sigma_i)$. Our model therefore suggests that for this particular demand class, products with large and convex demands (so high pass-through rates) should be stocked exclusively to attract consumers to visit, while products with large and concave demands (so low pass-through rates) should be stocked non-exclusively as profit generators.

7.2 Demand heterogeneity

In our baseline model consumers differ in their search costs and can end up purchasing different products. (As explained in discussion (iii) in Section 2.1, the baseline model also permits an interpretation where consumers have unit demand but heterogeneous valuations for each product.) In this section we introduce some additional demand heterogeneity and show that the main insights concerning product selection remain unchanged, as long as there are no extra price effects caused by the new demand heterogeneity. Suppose there are n different types of consumer, and let $\rho_j > 0$ denote the probability of being type $j = 1, \dots, n$. A consumer of type j demands an amount $\alpha_{i,j} Q_i(p_i)$ of good i when faced with a price p_i , where $\alpha_{i,j} \geq 0$ is independent of price.⁴⁷ (One example would be the case where $\alpha_{i,j} \in \{0, 1\}$, and different types of consumer want different subsets of the products. Another example would be the case where $\alpha_{i,j} \equiv \alpha_j$ but α_j varies across types, such that some consumers have larger demands for everything than others.) For simplicity here we focus on the simple case as in Section 3, and assume that a consumer's

⁴⁶This 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.

⁴⁷A more general approach would be to assume that type j has demand $Q_i(p_i, \alpha_j)$ for product i . However, because the intermediary and manufacturer might be searched by different consumer types, they would not necessarily have the same monopoly price, and so we would not be able to use the (π, v) approach. We thank the referee for encouraging us to think more about this.

unit search cost s is independent of her demand type. (The general case is reported in a separate supplementary document.)

Given our specification of demand heterogeneity, the monopoly price for good i is independent of which (and how many) consumer types are served. Hence we can again represent products as points in a two-dimensional (π, v) space, Ω . Moreover, slightly abusing notation, we can also define a measurable function $\alpha_j(\pi, v)$ such that when a consumer of type j buys the product (π, v) at its monopoly price, she gets surplus $\alpha_j(\pi, v) \times v$ and generates gross profit $\alpha_j(\pi, v) \times \pi$. To avoid corner solutions we assume that $\alpha_j(\pi, v) \times v < \bar{s}$ for all j and $(\pi, v) \in \Omega$.

Suppose the intermediary uses a stocking policy $q(\pi, v)$. Closely following our earlier analysis, type j consumers will search a manufacturer who sells direct if and only if $s \leq \alpha_j(\pi, v)v$, and will search the intermediary if and only if $s < \hat{v}_j$ where

$$\hat{v}_j = \frac{\int_{\Omega} q(\pi, v) \alpha_j(\pi, v) v dG}{\int_{\Omega} q(\pi, v) dG} \quad (24)$$

is the average surplus of the products sold at the intermediary for type j consumers. The intermediary must therefore pay a manufacturer $\sum_{j=1}^n \rho_j \alpha_j(\pi, v) \pi F(\alpha_j(\pi, v)v)$ in order to stock its product. The intermediary's total profit is then

$$\int_{\Omega} q(\pi, v) \left\{ \sum_{j=1}^n \rho_j \alpha_j(\pi, v) \pi [F(\hat{v}_j) - F(\alpha_j(\pi, v)v)] \right\} dG . \quad (25)$$

To keep the exposition clean we henceforth drop the dependence of q and α_j on (π, v) . The intermediary chooses a stocking policy function to maximize its profit (25) subject to the definition of \hat{v}_j in equation (24). The Lagrangian function is

$$\mathcal{L} = \int_{\Omega} q \left\{ \underbrace{\sum_{j=1}^n \rho_j \alpha_j \pi [F(\hat{v}_j) - F(\alpha_j v)]}_{\text{Direct effect}} + \underbrace{\sum_{j=1}^n \lambda_j (\alpha_j v - \hat{v}_j)}_{\text{Indirect effect}} \right\} dG , \quad (26)$$

where $\lambda_j \geq 0$ denotes the Lagrange multiplier for equation (24). As usual we can decompose the intermediary's value of stocking a product into the direct profit it earns, and the indirect effect it has on the profitability of other products via consumers' search decisions. Similar to our earlier analysis, for any given consumer type with $\alpha_j > 0$ and $\hat{v}_j > 0$ the direct and indirect effects of stocking a product have opposite signs. However as is clear from (26), it is possible that after summing over all types, the direct and indirect effects actually have the same sign. It is therefore possible that by stocking a product the intermediary earns a direct profit and also increases its total profit from other products (and vice versa). This is the main difference caused by demand heterogeneity.

We now characterize the intermediary’s optimal stocking policy.⁴⁸ It is clear from (26) that fixing $\hat{\mathbf{v}} = (\hat{v}_1, \dots, \hat{v}_n)$ and $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_n)$, a product (π, v) is stocked if either

$$\sum_{j=1}^n \rho_j \alpha_j [F(\hat{v}_j) - F(\alpha_j v)] > 0 \quad \text{and} \quad \pi \geq \frac{\sum_{j=1}^n \lambda_j (\hat{v}_j - \alpha_j v)}{\sum_{j=1}^n \rho_j \alpha_j [F(\hat{v}_j) - F(\alpha_j v)]}, \quad (27)$$

or

$$\sum_{j=1}^n \rho_j \alpha_j [F(\hat{v}_j) - F(\alpha_j v)] < 0 \quad \text{and} \quad \pi \leq \frac{\sum_{j=1}^n \lambda_j (\hat{v}_j - \alpha_j v)}{\sum_{j=1}^n \rho_j \alpha_j [F(\hat{v}_j) - F(\alpha_j v)]}. \quad (28)$$

The optimal stocking policy is more complicated than in our earlier one-type analysis. Specifically, because the α_j ’s depend on (π, v) , it is clear from equations (27) and (28) that the intermediary may no longer use a simple threshold rule when deciding what to stock. Nevertheless the main insights from our earlier analysis are robust. According to equation (27) the intermediary stocks profitable products if and only if their π is sufficiently high given the values of $(\alpha_1, \dots, \alpha_n)$. Similarly according to equation (28) the intermediary stocks loss-making products if and only if their π is sufficiently low given the values of $(\alpha_1, \dots, \alpha_n)$. Intuitively, as we have seen earlier in the paper, this maximizes gains from profitable products and minimizes losses from unprofitable products while simultaneously making it attractive for consumers to search.

7.3 Upstream competition

We now extend our baseline model by introducing upstream manufacturer competition. When a product has multiple varieties produced by different upstream manufacturers, stocking some of them will make the non-contracted ones less ‘prominent’ for consumers who visit the intermediary first. This makes manufacturers more willing to sell their products through the intermediary and has interesting effects on the contract terms. In particular, due to upstream competition the intermediary can even make direct profit from exclusively stocked products. Nevertheless the intermediary’s optimal product selection is qualitatively similar to earlier. This extension also illustrates how our framework can be used to study both product breadth and product depth.

Suppose that each product (category) i has a unit-continuum of symmetric differentiated varieties, each produced by a different manufacturer at the same marginal cost c_i . We capture differentiation by assuming that each variety is either a ‘match’ or else has no value for the consumer.⁴⁹ A consumer wishes to buy $Q_i(p_i)$ units of one ‘matched’

⁴⁸In the spirit of Lemma 2 we can derive simple conditions under which the intermediary’s profit is strictly positive. For example, when $\alpha_j(\pi, v) \in \{0, 1\}$ it is sufficient that there exists a convex set $\Omega' \subset \Omega$ such that $\alpha_j(\pi, v) = 1$ for all j and $(\pi, v) \in \Omega'$, i.e. each consumer type needs all the products in Ω' .

⁴⁹As we describe below, this ‘double-continuum’ framework is a tractable way to introduce upstream product differentiation into the (π, v) approach. Considering a discrete number of varieties would be messy except in the special case where they are all homogeneous.

variety of each product. We assume that matches occur with a Poisson rate $\gamma > 0$, and so the probability that a consumer matches with one or more of $q \leq 1$ varieties is

$$\phi(q) \equiv 1 - e^{-\gamma q} ,$$

which is increasing and concave in q . In the spirit of our earlier analysis, we assume that consumers cannot observe whether a variety is a match until they have searched its seller, and also assume that the cost of searching the intermediary depends on the total measure of varieties that it stocks (counted both within and across categories).⁵⁰ We also simplify the analysis by assuming that before they have searched anything, consumers make a one-off decision of whether or not to visit the intermediary, and do so if it strictly increases their payoff. We also assume that \bar{s} is sufficiently large that not all consumers search the intermediary in equilibrium. The model and timing are otherwise the same as earlier. For simplicity here we focus on the simple case as in Section 3.

Given the way we model product differentiation, it is still an equilibrium for each seller to charge its monopoly price. In particular, closely following Lemma 1, the intermediary uses bilaterally-efficient two-part tariffs to avoid double marginalization, and the search friction nullifies competition between sellers of two varieties of the same product.⁵¹ Consequently we can still represent products as points in a two-dimensional (π, v) space, although at each point there is now a continuum of varieties. Henceforth we let $q(\pi, v) \in [0, 1]$ denote the measure of varieties of product (π, v) stocked by the intermediary.

Consumer search behavior. We prove that a consumer's expected surplus, if she chooses to visit the intermediary, is

$$u^1(s, \mathbf{q}) = \int \phi(q)v dG - s \int q dG + \int_{v > s/\gamma} [1 - \phi(q)] \phi(1 - q) \left(v - \frac{s}{\gamma} \right) dG . \quad (29)$$

The first two terms are the surplus obtained at the intermediary: with probability $\phi(q)$ a consumer finds a matched variety of product (π, v) and gets surplus v , while the total cost of searching the intermediary is $s \int q dG$ in the simple case. The third term is the surplus obtained from manufacturers: with probability $1 - \phi(q)$ a consumer does not find a matched variety of product (π, v) at the intermediary, and (it can be shown) obtains $\phi(1 - q)(v - s/\gamma)$ in expectation from searching the remaining $1 - q$ manufacturers, such that it is worthwhile to search among them only if $v > s/\gamma$. Following a similar logic, a consumer's expected surplus from not visiting the intermediary can be shown to equal

⁵⁰Notice that compared to the main model, the costs of searching the intermediary and a manufacturer are scaled down by an order, but that this does not qualitatively affect consumer search behavior.

⁵¹Once a consumer finds a match she has no incentive to search other varieties of that product, because all firms charge the same price and her search cost is strictly positive. (One subtlety is that we now require $f(0) = 0$ to ensure monopoly pricing when a positive fraction of varieties is stocked by manufacturers.)

$$u^0(s, \mathbf{q}) = \int_{v > s/\gamma} \phi(1 - q) \left(v - \frac{s}{\gamma} \right) dG . \quad (30)$$

Consumers follow a cut-off search rule. Specifically, there exists a unique \hat{s} such that $u^1(\hat{s}, \mathbf{q}) = u^0(\hat{s}, \mathbf{q})$. Consumers with $s < \hat{s}$ first search the intermediary and buy a product (π, v) if they find a matched variety, and then search among manufacturers for unmatched varieties of products with $v > s/\gamma$. Consumers with $s \geq \hat{s}$ do not visit the intermediary and instead search among manufacturers for products with $v > s/\gamma$.

Intermediary profit. We prove that the intermediary's profit from stocking a measure q of varieties of product (π, v) when $v \leq \hat{s}/\gamma$ is

$$\Pi_l(q, \pi, v) = \pi \left[F(\hat{s}) \phi(q) - F(\gamma v) \frac{\phi(1) - \phi(q)}{1 - q} q \right] . \quad (31)$$

The first term is gross profit: the intermediary is searched by consumers with $s < \hat{s}$ and a match occurs with probability $\phi(q)$. The second term is manufacturer compensation: it can be shown that the intermediary makes a lump-sum transfer $F(\gamma v) [\phi(1) - \phi(q)] / (1 - q)$ to each of the q manufacturers whose variety it stocks. Interestingly, this per-manufacturer compensation is decreasing in q due to a novel “prominence effect”. Intuitively, a stand-alone manufacturer is only searched by consumers with $s \leq v\gamma$ who did not find a matched variety at the intermediary. As q increases so does the probability of a consumer finding a match at the intermediary, leading to a fall in demand (and profit) of manufacturers who sell direct. It is straightforward to show that the intermediary's profit (31) is strictly positive and increasing in q for all $q > 0$ i.e. as in our earlier analysis, low- v products are profit-generators.

We also prove that the intermediary's profit from stocking a measure q of varieties of product (π, v) when $v > \hat{s}/\gamma$ is

$$\Pi_h(q, \pi, v) = \pi \left[F(\hat{s}) \phi(q) - F(\hat{s}) \frac{\phi(1) - \phi(q)}{1 - q} q - [F(\gamma v) - F(\hat{s})] \frac{\phi(1 - q)}{1 - q} q \right] . \quad (32)$$

The first term is again the intermediary's gross profit, while the remaining terms are lump-sum transfers to manufacturers. Concerning the latter, it can be shown that if a manufacturer rejects the intermediary's offer it sells to consumers with $s \leq \hat{s}$ with probability $[\phi(1) - \phi(q)] / (1 - q)$, and sells to consumers with $s \in (\hat{s}, \gamma v)$ with probability $\phi(1 - q) / (1 - q)$. The former probability is the same as in equation (31) and so is decreasing in q due to the prominence effect. However, the latter probability is increasing in q , because consumers with $s \in (\hat{s}, \gamma v)$ only search manufacturers, and when q is higher each manufacturer faces less competition. An important difference with earlier analysis is that high- v products can now be profitable – (32) is strictly positive whenever v is close to \hat{s}/γ , because upstream competition reduces the lump-sum transfer paid to manufacturers.

Optimal product selection. As usual we can solve the intermediary’s problem using the Lagrangian method. The optimal stocking policy for products with $v \leq \hat{s}/\gamma$ is straightforward, and depends on two thresholds $\underline{\pi}_l(v) < \bar{\pi}_l(v)$. The fraction of varieties of a given product (π, v) stocked by the intermediary is weakly increasing in π , and equals 0 for $\pi \leq \underline{\pi}_l(v)$ and equals 1 for $\pi \geq \bar{\pi}_l(v)$. Intuitively the (positive) profit earned from stocking a product is concave in q due to concavity of the matching function $\phi(q)$, and so the intermediary maximizes gains from these products by stocking more varieties of those with higher π , much like in our earlier analysis.

The optimal stocking policy for products with $v > \hat{s}/\gamma$ is more complicated, and it is helpful to characterize it using two thresholds $v^* < v^{**}$. Among products with $v < v^*$, the fraction of varieties stocked by the intermediary is weakly increasing in $\pi > 0$, and equals 1 for π sufficiently high. Intuitively, upstream competition turns these products into profit-generators, so by the usual logic the intermediary maximizes these gains by stocking more of them as π increases. On the other hand, among products with $v > v^{**}$, the fraction of varieties stocked by the intermediary is weakly decreasing in $\pi > 0$, and equals 0 for π sufficiently high. Intuitively, as in our earlier analysis these products are loss-makers, and the losses become more severe as q increases due to the second “prominence effect” identified earlier. Hence the intermediary minimizes its losses by stocking fewer varieties of products with larger π . The characterization is more complex for products with $v \in (v^*, v^{**})$, but we can show that the optimal fraction of varieties stocked lies within two bounds, which tighten monotonically as π increases and converge. (The details can be found in a separate supplementary document.)

Finally, our analysis in this section may also provide some new insights into private labels. Empirical studies typically find that retailers are able to negotiate higher margins on national brands when they also stock private labels (see e.g. Ailawadi et al (2010)). Since private labels can be regarded as exclusive products, this would be consistent with our first ‘prominence’ effect. To our knowledge such an explanation for a retailer’s improved bargaining position is new to the literature.

8 Conclusion

This paper has developed a new and tractable framework for studying multiproduct intermediaries. 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; and (iii) the intermediary tends to be too big, and stock too many products exclusively, compared to the social optimum. Moreover (iv) our framework can also shed light on some important trends in intermediary markets, such as increased use of exclusivity arrangements, growth of large retailers and decline of small ones, and the rise of DTC sales. For instance, we show how easier DTC sales harm the intermediary by reducing its bargaining power, and induce the intermediary to redesign its product selection.

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.⁵² 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.⁵³ 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 applied is international trade, because intermediaries play an important role in helping domestic manufacturers to export.⁵⁴ 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.

Throughout the paper we have focused on a single intermediary. It would be interesting to study how competition between intermediaries (or a merger of two intermediaries) shapes their product ranges. Our (π, v) approach may not apply straightforwardly to this setting, in which case some modified techniques would need to be developed. We plan to think more about this in future work.

⁵²See <https://tcrn.ch/2AsRRlm> for details on Disney and ESPN, and <https://bit.ly/2JqS4XU> for details on HBO's 'HBO Now' service.

⁵³See <https://bit.ly/2JsDfEb> for arguments why these costs can be much lower when content is bundled rather than sold direct.

⁵⁴See e.g. Bernard et al (2010) and Ahn et al (2011) for empirical evidence on trade intermediaries in the US and China respectively.

References

- ADAMS, W., AND J. YELLEN (1976): “Commodity Bundling and the Burden of Monopoly,” *Quarterly Journal of Economics*, 90(3), 475–498.
- AGUIRREGABIRIA, V., D. ERSHOV, AND J. SUZUKI (2016): “Estimating the Effects of Deregulation in the Ontario Wine Retail Market,” mimeo.
- AHN, J., A. KHANDELWAL, AND S. WEI (2011): “The Role of Intermediaries in Facilitating Trade,” *Journal of International Economics*, 84(1), 73–85.
- AILAWADI, K., E. BRADLOW, M. DRAGANSKA, V. NIJS, R. ROODERKERK, K. SUDHIR, K. WILBUR, AND J. ZHANG (2010): “Empirical Models of Manufacturer-Retailer Interaction: A Review and Agenda for Future Research,” *Marketing Letters*, 21(3), 273–285.
- 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 (2020): “Vertical Information Restraints: Pro- and Anti-Competitive Impacts of Minimum Advertised Price Restrictions,” *Journal of Law and Economics*, 63(1), 111–148.
- 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.
- BRIESCH, R., W. DILLON, AND E. FOX (2013): “Category Positioning and Store Choice: The Role of Destination Categories,” *Marketing Science*, 32(3), 488–509.
- BRONNENBERG, B. (2020): “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.

- DUCH-BROWN, N., L. GRZYBOWSKI, A. ROMAHN, AND F. VERBOVEN (2017): “The Impact of Online Sales on Consumers and Firms: Evidence from Consumer Electronics,” *International Journal of Industrial Organization*, 52, 30–62.
- EUROPEAN-COMMISSION (2017): “Report from the Commission to the Council and the European Parliament: Final Report on the E-commerce Sector Inquiry,” Brussels.
- FISHER, M., AND R. VAIDYANATHAN (2012): “Which Products Should You Stock?,” *Harvard Business Review*, November.
- GEHRIG, T. (1993): “Intermediation in Search Markets,” *Journal of Economics and Management Strategy*, 2(1), 97–120.
- GFK-REPORT (2007): “Research on Suppliers to the UK Grocery Market,” prepared for the UK Competition Commission.
- GIELENS, K., E. GIJSBRECHTS, AND M. DEKIMPE (2014): “Gains and Losses of Exclusivity in Grocery Retailing,” *International Journal of Research in Marketing*, 31(3), 239–252.
- 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.
- HONG, G., AND N. LI (2017): “Market Structure and Cost Pass-Through in Retail,” *Review of Economics and Statistics*, 99(1), 151–166.
- HORTAÇSU, A., AND C. SYVERSON (2015): “The Ongoing Evolution of US Retail: A Format Tug-of-War,” *Journal of Economic Perspectives*, 29(4), 89–112.
- 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 (2019): “Relative Price Dispersion: Evidence and Theory,” *American Economic Journal: Microeconomics*.
- KÖK, 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.
- KRAKOVSKY, M. (2015): *The Middleman Economy*. Palgrave Macmillan US.

- 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.
- OECD (2019): “Implications of E-commerce for Competition Policy - Background Note,” Organization for Economic Co-operation and Development.
- OLIVER-WYMAN (2018): “Retail’s Revolution: How Retail and Consumer Goods Companies Can Adapt,” New York: Oliver Wyman.
- 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 (2019): “Consumer Search and Retail Market Structure,” *Management Science*, 65(6), 2607–2623.
- RUBINSTEIN, A., AND A. WOLINSKY (1987): “Middlemen,” *Quarterly Journal of Economics*, 102(3), 581–593.
- SEN, B., J. SHIN, K. SUDHIR, AND N. YANG (2013): “Demand Externalities from Co-Location,” mimeo.
- 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.
- SORESCU, A., R. FRAMBACH, J. SINGH, A. RANGASWAMY, AND C. BRIDGES (2011): “Innovation in Retail Business Models,” *Journal of Retailing*, 87, S3–S16.
- 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.
- UPSHAW, D., D. AMYX, AND M. HARDY (2017): “The Nature of Exclusivity,” *Journal of Marketing Development and Competitiveness*, 11(2), 46–59.
- 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.
- WORLD-ECONOMIC-FORUM (2017): “Shaping the Future of Retail for Consumer Industries,” Cologny-Geneva: World Economic Forum.
- ZHOU, J. (2014): “Multiproduct Search and the Joint Search Effect,” *American Economic Review*, 104(9), 2918–2939.

Online Appendix

In this online appendix, we report all the omitted details and proofs of the paper (except those for Section 7 which can be found in a separate supplementary document).

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

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

(i-2) We then show $p_j = p_j^m$ for $j \in A_E$. Suppose the wholesale price of product j is τ_j . The hold-up logic as explained in the main text 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

$$\eta \times [p_j^*(\tau_j) - \tau_j] Q_j(p_j^*(\tau_j)) - T_j = \eta \times [p_j^*(\tau_j) - c_j] Q_j(p_j^*(\tau_j)) - \pi_j F(v_j) . \quad (33)$$

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.⁵⁵ 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 $p_i^*(\tau_i) \leq p_i^m < p_{i,I}$.

⁵⁵With 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.

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

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

Then the intermediary's profit from product i is

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

If $p_{i,I}$ is reduced to p_i^m , $(p_{i,I} - c_i)Q_i(p_{i,I})$ will increase to π_i , the per-consumer monopoly profit, and $\eta_{i,I} + \eta_{i,M}$ will increase at least weakly.⁵⁶ 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 wish 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

⁵⁶In 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.

$i \in A_{NE}$. The lump-sum fee T_i satisfies

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

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

$$\begin{aligned} \eta_I \times (p_i^m - \tau_i)Q_i(p_i^m) - T_i &= \eta_I \times (p_i^m - c_i)Q_i(p_i^m) + \eta_{i,M} \times \pi_i - \pi_i F(v_i) \\ &= \pi_i [\eta_I - (F(v_i) - \eta_{i,M})] . \end{aligned} \quad (34)$$

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

Proof of Lemma 2. (i) We first show that the intermediary can make a positive profit by stocking a positive measure of products. Consider two interior points in Ω : (π_1, \tilde{v}) and (π_2, \tilde{v}) with $\pi_1 > \pi_2$. Let $A_1 = [\pi_1 - \delta, \pi_1] \times [\tilde{v} - \epsilon, \tilde{v}]$ and $A_2 = [\pi_2, \pi_2 + \Delta(v)] \times [\tilde{v}, \tilde{v} + \epsilon]$, where $\Delta(v)$ is uniquely defined for each $v \in [\tilde{v}, \tilde{v} + \epsilon]$ by

$$\int_{\pi_1 - \delta}^{\pi_1} g(\pi, 2\tilde{v} - v) d\pi = \int_{\pi_2}^{\pi_2 + \Delta(v)} g(\pi, v) d\pi . \quad (35)$$

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

$$\tilde{\Pi}(\epsilon) = \int_{\tilde{v} - \epsilon}^{\tilde{v}} \int_{\pi_1 - \delta}^{\pi_1} \pi [F(\tilde{v}) - F(v)] dG + \int_{\tilde{v}}^{\tilde{v} + \epsilon} \int_{\pi_2}^{\pi_2 + \Delta(v)} \pi [F(\tilde{v}) - F(v)] dG . \quad (36)$$

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

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

where the second equality used (35) evaluated at $v = \tilde{v}$. Therefore, $\tilde{\Pi}(\epsilon) > 0$ for ϵ in a neighborhood of 0.

(ii) We then show that stocking all the products is not the most profitable strategy. Let $v^* = \int_{\Omega} v dG$. Consider $B_1 = [\pi_1 - \delta, \pi_1] \times [v^*, v^* + \epsilon]$ and $B_2 = [\pi_2, \pi_2 + \Delta(v)] \times [v^* - \epsilon, v^*]$, where $\pi_1 > \pi_2$, and where $\Delta(v)$ is uniquely defined for each $v \in [v^* - \epsilon, v^*]$ by

$$\int_{\pi_1 - \delta}^{\pi_1} g(\pi, 2v^* - v) d\pi = \int_{\pi_2}^{\pi_2 + \Delta(v)} g(\pi, v) d\pi. \quad (37)$$

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 v^* , and so the average v in $A = \Omega \setminus (B_1 \cup B_2)$ is v^* as well. Then a consumer will visit the intermediary, when it stocks $A = \Omega \setminus (B_1 \cup B_2)$, if and only if $s < v^*$.

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

$$\Pi^*(\epsilon) = \Pi^*(0) - \int_{v^*}^{v^* + \epsilon} \int_{\pi_1 - \delta}^{\pi_1} \pi [F(v^*) - F(v)] dG - \int_{v^* - \epsilon}^{v^*} \int_{\pi_2}^{\pi_2 + \Delta(v)} \pi [F(v^*) - F(v)] dG,$$

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

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

by using (37) evaluated at $v = v^*$. Therefore, $\Pi^*(\epsilon) > \Pi^*(0)$ for ϵ in a neighborhood of 0. ■

Proof of Lemma 3. The difference in payoff between $u^1(s, \mathbf{q})$ defined in (10) and $u^0(s, \mathbf{q})$ defined in (11) is

$$\Delta(s) = \int qv dG - h \left(\int qdG \right) s - \int_{v > s} q_{NE} (v - s) dG. \quad (38)$$

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

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

is weakly decreasing in s . No consumer visits the intermediary (i.e. $\hat{s} = 0$) 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 (i). All consumers visit the intermediary (i.e. $\hat{s} > \bar{s}$) 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 (ii). 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 $s < \hat{s}$, where \hat{s} solves $\Delta(\hat{s}) = 0$. (12) is just a rewriting of $\Delta(\hat{s}) = 0$.⁵⁷ ■

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 $\hat{s} \geq \tilde{v}$. To see this, note that by the definition of $\Delta(s)$ in equation (38) we have

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

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} . Then the proof of Lemma 3 shows that $\Delta(\tilde{v}) \geq 0$ implies $\hat{s} \geq \tilde{v}$. Secondly, recall that non-exclusive products with $v < \hat{s}$ earn positive profit, and non-exclusive products with $v > \hat{s}$ earn zero profit. Therefore since $\hat{s} \geq \tilde{v}$ the intermediary's profit is weakly larger than in equation (36) 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 $\hat{s} > 0$ so that some consumers visit it.

First, we characterize the optimal product selection for given parameters (\hat{s}, λ, μ) . From (17) it is clear that for a product with $v < \hat{s}$, $q = 1$ if and only if

$$\pi [F(\hat{s}) - F(v)] + \lambda v - \mu \geq 0 \Leftrightarrow \pi \geq \frac{\mu - \lambda v}{F(\hat{s}) - F(v)} , \quad (39)$$

⁵⁷Notice that $\Delta(0) = 0$ if the intermediary only stocks products non-exclusively (i.e. if $q = q_{NE}$). In that case, $\Delta(\hat{s}) = 0$ has one zero solution and another strictly positive solution under our conditions. The zero solution should be discarded.

and the exclusivity arrangement does not matter. From (17) it is also clear that for a product with $v > \hat{s}$, $q_E = 1$ if and only if

$$\pi [F(\hat{s}) - F(v)] + \lambda v - \mu \geq \max\{0, \lambda \hat{s} - \mu\} \Leftrightarrow \pi \leq \frac{\max\{\lambda \hat{s}, \mu\} - \lambda v}{F(\hat{s}) - F(v)}. \quad (40)$$

When the opposite is true, $q_{NE} = 1$ if $\lambda \hat{s} - \mu > 0$, $q_{NE} = 0$ if $\lambda \hat{s} - \mu < 0$, and the intermediary is indifferent between $q_{NE} = 0$ and $q_{NE} = 1$ if $\lambda \hat{s} - \mu = 0$. (As we explain below, in the third case there is flexibility over which products are stocked.)

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

(a) Suppose the optimal solution has $\hat{s} < \bar{s}$.⁵⁸ Then \hat{s} satisfies (14). Since \hat{s} is interior, the first-order condition of (17) with respect to \hat{s} gives

$$\lambda = f(\hat{s}) \frac{\int_{v < \hat{s}} q \pi dG + \int_{v > \hat{s}} q_E \pi dG}{h(m) - \int_{v > \hat{s}} q_{NE} dG} > 0. \quad (41)$$

(Notice that $\lambda > 0$ since λ and \hat{s} 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

$$\mu = \lambda \hat{s} h'(m) \geq 0, \quad (42)$$

and $(\hat{s}, \lambda, \mu, m)$ jointly solve equations (14), (15), (41) and (42). If $m = \bar{m}$ in the optimal solution, η is irrelevant as we have explained. Note that if $m = \bar{m} < 1$ then (\hat{s}, λ, μ) jointly solve equations (14), (15) and (41) with m replaced by \bar{m} everywhere, and we can conclude that $\mu \geq 0$.⁵⁹ Note that if instead $m = \bar{m} = 1$ then μ does not matter so without loss of generality we can set $\mu = 0$, and (\hat{s}, λ) jointly solve equations (14) and (41) with m replaced by \bar{m} .⁶⁰ (Regardless of whether (16) binds or not, when $\lambda \hat{s} - \mu = 0$ we have an additional equation, and this determines the measure of products with $v > \hat{s}$ that are stocked non-exclusively.)

(b) Suppose the optimal solution has $\hat{s} > \bar{s}$.⁶¹ Then there is no region of $v > \hat{s}$ given our focus on $\bar{v} \leq \bar{s}$. The value of \hat{s} now does not affect the product selection which is characterized by (39) with $F(\hat{s}) = 1$. As stated in the text $\lambda = 0$ in this case. We can also

⁵⁸According to Lemma 3 a stocking policy \mathbf{q} induces $\hat{s} < \bar{s}$ when $\int q v dG < h(\int q dG) \bar{s}$. A sufficient condition for this to hold is $\max_x \int_x^{\bar{v}} v dG / h(\int_x^{\bar{v}} dG) < \bar{s}$, which is easily satisfied when economies of search are weak or when $\bar{s} \gg \bar{v}$.

⁵⁹Notice that if to the contrary $\mu < 0$, then firstly the righthand side of (39) is strictly negative, and secondly $\lambda \hat{s} - \mu > 0$. Hence all products should be stocked, which contradicts $m = \bar{m} < 1$. The same argument applies in parts (b) and (c) below when $m = \bar{m} < 1$.

⁶⁰Note that the righthand side of (39) is weakly negative and that $\lambda \hat{s} - \mu \geq 0$, so this is consistent with $m = 1$. A similar argument applies in part (b) below when $m = \bar{m} = 1$.

⁶¹According to Lemma 3, $\hat{s} > \bar{s}$ in the optimal solution requires $\int q v dG > h(\bar{m}) \bar{s}$.

deduce that $m = \bar{m}$ in the optimal solution. Suppose to the contrary that $m < \bar{m}$: (42) 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 $q = 1$ for the products satisfying (39). 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 $\hat{s} = \bar{s}$.⁶² Again, there is no region of $v > \hat{s}$ in this case. Notice that if $m < 1$ in the optimal solution then it must be that $\lambda \geq 0$.⁶³ If $m < \bar{m}$ in the optimal solution, (λ, μ, m) jointly solve equations (14) with $\hat{s} = \bar{s}$, (15) and (42). If $m = \bar{m} < 1$ in the optimal solution, λ and $\mu \geq 0$ jointly solve (14) with $\hat{s} = \bar{s}$ and $\bar{m} = \int qdG$, where $q = 1$ for the products satisfying (39). 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 (39) 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).) ■

Omitted details in Section 4.3 and Proof of Proposition 3. In the problem of socially optimal product range choice, the consumer search rule is the same as in Lemma 3. Total welfare can then be written as

$$W(\mathbf{q}) \equiv \int \pi F(v) dG + \Pi(\mathbf{q}) + \int_0^{\hat{s}} u^1(s, \mathbf{q}) dF(s) + \int_{\hat{s}}^{\bar{s}} u^0(s, \mathbf{q}) dF(s) . \quad (43)$$

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 $s < \hat{s}$ search the intermediary and earn $u^1(s, \mathbf{q})$, which we defined in equation (10). Meanwhile consumers with $s \geq \hat{s}$ do not search the intermediary and so earn $u^0(s, \mathbf{q})$, which we defined in equation (11). When some products with $v > \hat{s}$ are stocked exclusively, consumers with $s \geq \hat{s}$ are made worse off because they cannot buy all the products they would like to. At the same time, consumers with $s < \hat{s}$ 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.

⁶²According to Lemma 3, $\hat{s} = \bar{s}$ in the optimal solution requires $\int qvdG = h(m)\bar{s}$.

⁶³Suppose to the contrary that $\lambda < 0$. Then the profit $\pi[F(\hat{s}) - F(v)] = \pi[1 - F(v)]$ of products along the boundary of the set defined by (39) equals $\mu - \lambda v$ which is now strictly increasing in v . Suppose the intermediary replaces a strictly positive measure of low- v products close to the boundary which it stocks under \mathbf{q} , with a strictly positive measure of higher- v products close to the boundary which it does not stock under \mathbf{q} . This can be done so that m does not change given $m < 1$, and by Lemma 3 we would have $\hat{s} > \bar{s}$ so that all consumers still visit the intermediary. At the same time, the newly added products are more profitable than the removed ones because the boundary is increasing. Therefore, the above change must improve the intermediary's profit, which yields a contradiction.

The social planner wishes to maximize (43). 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. The optimization problem can be solved by using the same techniques as we did earlier for profit maximization. We again use $m = \int qdG$ to denote the measure of products stocked by the intermediary.

Substituting the expressions for $\Pi(\mathbf{q})$, $u^1(s, \mathbf{q})$ and $u^0(s, \mathbf{q})$ into equation (43), we get

$$W(\mathbf{q}) = \int (1-q) \int_0^v (\pi + v - s) dF(s) dG + \int q(\pi + v) F(\hat{s}) dG - h(m) \int_0^{\hat{s}} s dF(s) + \int_{v > \hat{s}} q_{NE} \int_{\hat{s}}^v (\pi + v - s) dF(s) dG. \quad (44)$$

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:

$$\begin{aligned} \mathcal{L} = & \int_{v < \hat{s}} q \underbrace{\left[(\pi + v) [F(\hat{s}) - F(v)] + \lambda v - \mu + \int_0^v s dF(s) \right]}_{[1]} dG \\ & + \int_{v > \hat{s}} \left\{ q_E \underbrace{\left[(\pi + v) [F(\hat{s}) - F(v)] + \lambda v - \mu + \int_0^v s dF(s) \right]}_{[2]} + q_{NE} \underbrace{\left[\lambda \hat{s} - \mu + \int_0^{\hat{s}} s dF(s) \right]}_{[3]} \right\} dG \\ & - h(m) \left[\lambda \hat{s} + \int_0^{\hat{s}} s dF(s) \right] + \mu m + \eta(\bar{m} - m). \quad (45) \end{aligned}$$

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 (\hat{s}, λ, μ) , products with $v < \hat{s}$ will be stocked if $[1] \geq 0$, which leads to (20). In this region, the exclusivity arrangement does not matter. Products with $v > \hat{s}$ will be stocked exclusively if $[2] \geq \max\{0, [3]\}$, which leads to (21). When the opposite is true, they will be stocked non-exclusively if $[3] > 0$.

We now explain how to determine the parameters (\hat{s}, λ, μ) . Here we only consider the case where the optimal solution has $\hat{s} < \bar{s}$. (The case of $\hat{s} > \bar{s}$ or $\hat{s} = \bar{s}$ can be easily dealt with like we did in the proof of Proposition 2.) Taking the first order condition of (45) with respect to \hat{s} and then using equation (12) to cancel terms, we obtain

$$\lambda = f(\hat{s}) \frac{\int_{v < \hat{s}} q \pi dG + \int_{v > \hat{s}} q_E \pi dG}{h(m) - \int_{v > \hat{s}} q_{NE} dG} > 0. \quad (46)$$

This is the same as equation (41) from the proof of Proposition 2. If $m < \bar{m}$ in the optimal solution, then $\eta = 0$ and so the first-order condition of (45) with respect to m

gives

$$\mu = h'(m) \left[\lambda \hat{s} + \int_0^{\hat{s}} s dF(s) \right] \geq 0 , \quad (47)$$

Then $(\hat{s}, \lambda, \mu, m)$ jointly solve equations (14), (15), (46) and (47). If $m = \bar{m}$ in the optimal solution, η is irrelevant as we have explained. Note that if $m = \bar{m} < 1$ then (\hat{s}, λ, μ) jointly solve equations (14), (15) and (46) 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 (\hat{s}, λ) jointly solve equations (14) and (46) with m replaced by \bar{m} .

Omitted details in Section 5.2 and Proofs for Proposition 4 and Corollaries 2-4. Since much of the analysis is similar to the baseline model in Section 4, here we sketch the details and then prove the main results.

As in the baseline model, given a stocking policy $\mathbf{q}(\pi, v, \theta)$ a consumer's expected utilities from respectively searching and not searching the intermediary are

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

and

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

As usual one can verify that $u^1(s, \mathbf{q}) - u^0(s, \mathbf{q})$ is weakly concave and non-negative at $s = 0$. Using the same steps as in the proof of Lemma 3, we can then show that consumers use a cut-off search rule and will visit the intermediary if and only if $s < \hat{s}$. For example, in the interior solution case with $\hat{s} \in (0, \bar{s}]$, \hat{s} is determined by

$$\hat{s} = \frac{\int_{v < \theta \hat{s}} qv dG + \int_{v > \theta \hat{s}} q_E v dG}{h \left(\int qdG \right) - \int_{v > \theta \hat{s}} q_{NE} \theta dG} . \quad (48)$$

Following similar steps as in the baseline model and noting that a manufacturer's outside option from not selling at all through the intermediary is $\pi F(v/\theta)$, we can write the intermediary's profit as:

$$\int_{v < \theta \hat{s}} q\pi \left[F(\hat{s}) - F\left(\frac{v}{\theta}\right) \right] dG + \int_{v > \theta \hat{s}} q_E \pi \left[F(\hat{s}) - F\left(\frac{v}{\theta}\right) \right] dG . \quad (49)$$

Proof of Proposition 4. Solving the intermediary optimization problem is very similar to the baseline model, so again we only report the main steps. The definition of \hat{s} can be rewritten as

$$\int_{v < \theta \hat{s}} qv dG + \int_{v > \theta \hat{s}} (q_E v + q_{NE} \theta \hat{s}) dG - h(m) \hat{s} = 0 ,$$

where $m = \int qdG$, and the stocking space constraint is $m \leq \bar{m}$. Then the Lagrangian of the problem is

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

It is then straightforward to derive the characterization of optimal product selection reported in Proposition 4. As in the baseline model, both λ and μ are non-negative in the solution. In particular, whenever \hat{s} is an interior solution (i.e. $\hat{s} < \bar{s}$), the first-order condition with respect to \hat{s} yields

$$\lambda = f(\hat{s}) \frac{\int_{v < \theta \hat{s}} q \pi dG + \int_{v > \theta \hat{s}} q_E \pi dG}{h(m) - \int_{v > \theta \hat{s}} q_{NE} \theta dG} .$$

■

We now prove the three corollaries on DTC sales.

Proof of Corollary 2. Consider two sectors with θ_1 and $\theta_2 > \theta_1$ respectively. It suffices to show that if a product (π, v) in sector θ_1 is stocked, then its counterpart in sector θ_2 must be stocked.

Suppose $v < \theta_1 \hat{s}$ and the sum of the direct and indirect effects in the first line of \mathcal{L} with θ_1 is positive. Then clearly $v < \theta_2 \hat{s}$, and also the sum of the direct and indirect effects with θ_2 is positive as well.

If instead $\theta_1 \hat{s} < v < \theta_2 \hat{s}$ and product (π, v) in sector θ_1 is stocked exclusively, then product (π, v) in sector θ_2 must be stocked. If on the other hand $\theta_1 \hat{s} < v < \theta_2 \hat{s}$ and product (π, v) in sector θ_1 is stocked non-exclusively (which implies $\lambda \theta_1 \hat{s} - \mu > 0$), then $\lambda v - \mu > \lambda \theta_1 \hat{s} - \mu > 0$. This, together with $F(\hat{s}) - F(v/\theta_2) > 0$, implies that the total effect in the first line of \mathcal{L} with θ_2 is positive and so product (π, v) in sector θ_2 must be stocked.

Finally, if $v > \theta_2 \hat{s}$ and product (π, v) in sector θ_1 is stocked (exclusively or non-exclusively), then at least one of the terms after q_E and q_{NE} in the second line of \mathcal{L} remain positive, so product (π, v) in sector θ_2 must also be stocked. ■

Proof of Corollary 3. Firstly, since by assumption the space constraint is not binding, the first order condition of \mathcal{L} with respect to m (after noting that η does not matter) yields $\mu = \lambda \hat{s} h'(m)$.

Secondly then, we have that $\lambda \theta \hat{s} - \mu \propto \theta - h'(m)$. It is then clear that if $\theta > h'(m)$ then all products with $v > \theta \hat{s}$ are stocked (either exclusively or non-exclusively), whereas if $\theta < h'(m)$ then no products with $v > \theta \hat{s}$ are stocked non-exclusively.

Thirdly, consider the limit as $v \rightarrow \theta \hat{s}$. As v approaches from the left the stocking threshold tends to

$$\lambda \hat{s} \frac{h'(m) - \theta}{F(\hat{s}) - F(\frac{v}{\theta})},$$

which equals $-\infty$ when $\theta > h'(m)$ and equals $+\infty$ when $\theta < h'(m)$. As v approaches from the right the stocking threshold for exclusivity tends to

$$\frac{\lambda \hat{s} \max\{\theta, h'(m)\} - \lambda \theta \hat{s}}{F(\hat{s}) - F(\frac{v}{\theta})},$$

which equals $-\infty$ when $\theta < h'(m)$. The remainder of the claimed results in the corollary then follow. ■

Proof of Corollary 4. First, consider manufacturer profits. When DTC sales become easier the manufacturer of product (π, v, θ) earns $\pi F(\frac{v}{\tau(\theta; \pi, v)})$, regardless of where its product is sold. Since $\tau(\theta; \pi, v) \leq \theta$ this is weakly higher than before.

Second, consider the intermediary's profit. Let $\mathbf{q}(\pi, v, \theta)$ and $\tilde{\mathbf{q}}(\pi, v, \theta)$ be the optimal stocking policies before and after DTC sales become easier, respectively. Let \hat{s} and \tilde{s} be the associated search thresholds. Consider the situation before DTC sales become easier. Suppose the intermediary adopts a stocking policy $\mathbf{q}'(\pi, v, \theta) = \tilde{\mathbf{q}}(\pi, v, \theta)$ (which may not be optimal), and let \hat{s}' be the associated search threshold. Then from $u^1(s, \mathbf{q}) - u^0(s, \mathbf{q})$, it is clear that consumers' incentive to visit the intermediary must be higher than after DTC sales become easier, which implies $\hat{s}' \geq \tilde{s}$. From (49), it is then clear that the intermediary's profit under $\mathbf{q}'(\pi, v, \theta)$ must be higher than that after DTC sales become easier. The intermediary can do even better by choosing $\mathbf{q}(\pi, v, \theta)$ before DTC sales become easier. Consequently the intermediary's profit must fall when DTC sales become easier. ■