Available online at www.sciencedirect.com
ScienceDirect
International Journal of Industrial Organization
26 (2008) 1090–1105
International Journal of Industrial Organization
www.elsevier.com/locate/econbase

Group purchasing, nonlinear tariffs, and oligopoly ☆

Howard P. Marvel *, Huanxing Yang

Ohio State University, 1945 North High Street, Columbus, OH 43210-1172, United States

Received 20 March 2007; received in revised form 22 June 2007; accepted 10 October 2007
Available online 17 November 2007

Abstract

Loyalty discounts are nonlinear tariffs that condition rebates or marginal prices on meeting aggregate purchase or market share targets. These discounts are widespread, and are often the impetus for consumers to form buying groups, or group purchase organizations (GPOs). This paper models the competitive effects of the introduction of a GPO into a market within which the preferences of the GPO's members are horizontally differentiated. While nonlinear tariffs are an effective way for a monopolist to extract consumer surplus, when two suppliers compete using such schedules, the results are far more competitive in comparison to simple Bertrand–Nash competition with linear tariffs. This result holds when the product of each of the suppliers is attractive to a substantial portion of consumers. In our model, the nonlinear schedule puts all customers "in play" to a degree that contrasts sharply with the competition at the margin characteristic of constant per unit prices. Moreover, competing in nonlinear tariffs removes allocative inefficiency that can result from single price competition.

JEL classification: L13; L42; D43

Keywords: Nonlinear tariffs; Loyalty discounts; Competition; Group purchasing organization

1. Introduction

This paper presents a model of the effect of nonlinear tariffs on competition among oligopolists. Tariff schedules that condition rebates or discounts on the volume or share of a consumer's purchases are commonplace when buying groups of retailers such as grocers, hardware stores, and auto parts dealers deal with suppliers. These conditional tariff schedules are often referred to as loyalty discounts. In recent years, they have attracted considerable attention as emergent group purchasing organizations (GPOs) have come to dominate the purchasing of pharmaceuticals and other medical supplies. The popularity of such arrangements appears to stem from their claim to reduce the acquisition costs of the products in question. Indeed, purchasing groups often form to extract such discounts by pitting rival suppliers against one another. This paper models the effect of aggregating consumers into a buying group on the competitiveness of suppliers of differentiated products. The role of group purchasing in our model is simply to enable suppliers to compete with rivals by means of nonlinear tariffs. Our model illustrates the impact of these tariffs on both the prices paid by group members and the efficiency

☆ We thank an anonymous referee, James D. Dana, Roman 1nderst, Daniel P. O'Brien, James Peck, and participants in an Ohio State University seminar, the Fall 2006 Midwest Economic theory Meetings, and the 2006 Duke/Northwestern/Texas IO theory Conference for comments and suggestions. All conclusions and/or errors are solely the responsibility of the authors.

* Corresponding author. Tel.: +1 614 292 1020.
E-mail address: marvel.2@osu.edu (H.P. Marvel).

0167-7187/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.ijindorg.2007.10.002

of allocation of the products in question among group members.

Nonlinear tariff schedules of the sort we model are both widespread and controversial. U.S. courts have typically refused to condemn such schedules under the antitrust laws, at least so long as the discounts are not "bundled" across products. European policy has been to outlaw such discount schedules whenever a supplier offering the schedule is deemed to be dominant. But neither the law nor the economics analysis of such discounts is settled. Indeed, the American antitrust authorities recently recommended that the U.S. Supreme Court defer taking up the question of the appropriate law governing bundled rebates, a complex form of nonlinear tariff, "… to allow… economic analysis to develop further…"[1]

In monopoly settings, nonlinear tariffs are an effective way to extract consumer surplus. But the existence of rebates conditioned on share targets, as is common in health care, means that the relevant setting is oligopoly. We show that as long as consumer preferences are not extraordinarily asymmetric, when firms compete using nonlinear tariff schedules, the results are far more competitive than those that result from simple Bertrand–Nash competition with linear tariffs. In our model, a nonlinear schedule puts all customers "in play" to a degree that contrasts sharply with the competition at the margin characteristic of constant per unit prices. Consumers who face oligopoly suppliers thus have a strong incentive to form GPOs when doing so can induce competition through nonlinear tariffs. In addition, nonlinear tariffs eliminate allocative inefficiency that can result from single price competition.

We assume that firms know only the distribution of consumer preferences, but cannot identify the preferences of a particular consumer. This incomplete information makes nonlinear tariffs infeasible in the absence of buyer groups. But when a buyer group forms, it can induce suppliers to compete with such tariffs. In so doing, a buyer group becomes both a mechanism for increasing the dimensions on which firms compete and a device for revealing (or aggregating) information.

In our model, buyer groups extract attractive prices from suppliers not through enhanced bargaining power, but rather through their ability to obtain tariffs that pit suppliers against one another more effectively. This basic result leads to an extension where we show that nonlinear tariffs raise the return to promotional or innovative activities designed to sharpen the advantage of a particular supplier's product over those of rivals, again in comparison to the linear tariff benchmark.

Not surprisingly, given the widespread use of such schedules and the controversy that surrounds them, we are not the first to take up this issue. Section 2 surveys the relevant economics literature and outlines briefly the governing case law, and provides a description of the stylized facts of medical GPOs that motivate our model. Section 3 presents a simple model with a uniform distribution and symmetric firms. Section 4 extends the results to more general distributions and allows firms to have asymmetric market positions. Finally, Section 5 summarizes and concludes the analysis.

2. Nonlinear tariffs with oligopoly in economics and law

2.1. Group purchasing and nonlinear tariffs in economics

Several papers have suggested that GPOs are a way for buyers to amass bargaining power (Chipty and Snyder, 1999; Inderst and Wey, 2003, 2007). When the supplier has increasing unit costs, buyer groups can capture a larger share of surplus than individual consumers do. In a multiple seller/multiple buyer setting, Snyder (1996, 1998) shows that, as the size of a buyer increases, sellers' ability to tacitly coordinate on price decreases, which leads to lower average prices. There is some reason to doubt that sheer size, as opposed to the ability to move business among suppliers, is the source of a GPO's low prices. In the market for healthcare supplies, GPO contracts often include purchase targets which require that a GPO be able to shift consumers between suppliers. Large pharmacy chains, unable to shift purchases among suppliers, do not obtain similar discounts (Ellison and Snyder, 2001). It thus appears that nonlinear tariffs, and not their size-driven bargaining power, determines the GPOs' ability to deliver lower prices to members.

A second strand of literature considers the effect of competition among sellers when a buyer or GPO can commit to purchase exclusively from one of the sellers. O'Brien and Shaffer (1997) show that the buyer can obtain a lower price through an exclusive commitment. The exclusive commitment comes with inefficiency in that the buyer does not receive its desired allocation of the suppliers' goods. For consumers, this inefficiency is justified by intensified competition between the rival suppliers.

[1] Antitrust Division, U.S. Department of Justice and Federal Trade Commission, "Brief for the United States as Amicus Curiae," 3M v. LePage's, Inc.,http://www.usdoj.gov/atr/cases/f203900/203900.pdf, page 19. For more on the bundled discounts in the LePage's case, see Marvel and Peck (2006).

Dana (2003) extends O'Brien and Shaffer by endogenizing the decisions of buyers to form groups. He shows that when buyers form groups that commit to buy exclusively from one supplier, they obtain a lower price equal to their supplier's marginal cost. Buyers thus have a powerful incentive to form groups and to enter into exclusive purchase commitments, though as in O'Brien and Shaffer, the resulting allocation is inefficient. But Dana assumes that sellers are able to identify the preferences of each individual customer, and thus supplier price discrimination is feasible even without GPOs. In contrast, suppliers in our model are unable to observe preferences of individual consumers, rendering price discrimination infeasible. The formation of a GPO makes nonlinear tariffs (price discrimination) feasible. It is this move to nonlinear tariffs that intensifies competition and yields lower prices, a very different mechanism from the one posited by Dana.

Economists have long understood that "perfect" (first-degree) price discrimination made possible through the use of nonlinear tariffs will yield a more efficient outcome than that possible with simple linear pricing.[2] Consumers are unlikely to be pleased with this increased efficiency, however, as all of their surplus is extracted by means of the tariff. But our approach allows suppliers to discriminate by means of take-it-or-leave-it offers of nonlinear tariff schedules, thereby combining first-degree price discrimination with oligopoly but without introducing any scope for bargaining. Previous efforts to integrate price discrimination into oligopoly have addressed how the advent of third-degree price discrimination affects prices and firm profitability. Holmes (1989) shows that under oligopoly, third-degree price discrimination may decrease firms' profits. This point is illustrated more explicitly by Corts (1998), who demonstrates that if two sellers rank consumer groups differently (or two sellers have different "strong markets"), then allowing third-degree price discrimination can yield lower prices in all markets. Our model demonstrates that similar results hold under general nonlinear tariffs.

There is a growing literature on competitive nonlinear pricing (second-degree price discrimination) that includes Stole (1995), Armstrong and Vickers (2001, 2006), Rochet and Stole (2002), Yang and Ye (2007), and Thanassoulis (2007). In these papers, consumers are heterogeneous both in horizontal and vertical dimensions, and nonlinear pricing is used to sort consumers in the vertical dimension only. With or without nonlinear pricing, in the horizontal dimension, competition only occurs for consumers at the margin. On the contrary, in our model consumers differ only in horizontal dimension, and nonlinear tariffs are used to influence the market share (in the horizontal dimension). With nonlinear tariffs competition becomes global in contrast to local competition at marginal consumers under linear tariffs. A more detailed comparison is provided below, at the conclusion of Section 3.

Several other papers consider settings that loosely resemble ours. Che and Gale (1997) investigate methods by which sponsors of buying alliances manage competition among health insurance firms, actively choosing the format under which competition occurs. In contrast, our model treats GPOs as passive responders to take-it-or-leave-it offers whose only role is to facilitate nonlinear tariffs. Mathewson and Winter (1996) examine the effect of buyer groups in a monopolistic competition market, focusing on the behavior of parties excluded from groups. Kolay et al. (2004) study an all-units discount in a setting with a single buyer facing a single seller. Greenlee and Reitman (2005) study firms competing by offering loyalty discounts. Neither of these last two papers considers the possibility that buyers may form groups, which is the central topic of the current paper.

2.2. Antitrust policy for nonlinear discounts

Loyalty and bundled discounts have also featured prominently in case law. Concord Boat Corp. v. Brunswick Corp., 207 F.3d 1039 (8th Cir. 2000), and Virgin Atlantic Airways, Ltd. v. British Airways PLC., 60 F. Supp. 571 (S.D.N.Y. 1999), aff'd, 257 F.3d 256 (2d Cir. 2001), each involved nonlinear tariff discount structures that awarded consumers according to market share. In each case, the courts ultimately declined to rule that the tariff schedules were anticompetitive.[3]

The situation is very different in Europe. A judgment referred to as Michelin II has placed a nearly insurmountable burden of proof of efficiencies onto supposedly dominant firms wishing to adopt loyalty

[2] See Varian (1989) for a survey.

[3] Brunswick, whose MerCruiser stern drives were preferred by many customers to those of its principal rival, Volvo, offered discounts to boatbuilders according to the share of its stern drives they installed in their boats. Volvo countered with discounts of its own. The Concord Boat court concluded that this represented competition on the merits, and overturned a lower court award of in excess of $140 million to boatbuilder customers of Brunswick. The Court of Appeals determined that the discounts were themselves the products of competition and that since they were not shown to represent below cost pricing, they were not predatory.

In Virgin Atlantic, an attempt was made to show that the discounts that British Airways offered travel agents for making customer-specific sales targets resulted in below-cost pricing on transatlantic routes. The court rejected the proffered evidence on below-cost pricing, and thereby upheld the loyalty program.

rebates.[4] The position of European antitrust authorities appears to be that there are no positive effects of such discounts, and therefore that they are "abusive" unless adopted in response to competition from others. Accordingly, loyalty discounts are nearly, though not completely, per se illegal.[5]

2.3. Healthcare GPOs

In the United States, nonlinear tariffs are perhaps most controversial in markets for medical supplies. In 2003, Hillenbrand Industries agreed to pay $250 million to Kinetic Concepts, a rival manufacturer of specialty hospital beds, to settle a dispute over a discount schedule that awarded hospitals discounts on their bundled purchases of hospital beds of all types.[6] More recently, in March 2005, a jury awarded Masimo Corp. $140 million (before trebling) in damages in a dispute over a loyalty discount offered by the leading manufacturer of pulse oximeters, Nellcor,[7] though that case remains in litigation. Johnson & Johnson has sued Amgen for bundled discounts that lower the marginal price of its popular drug Aranesp when combined with purchases of related drugs.[8]

Our analysis is applicable to any oligopoly whose rival suppliers face purchasing aggregation sufficient to make nonlinear tariff competition feasible. But for purposes of exposition, we concentrate on medical supply purchasing because of the rising importance of GPOs in such markets and the controversy they have generated. There are hundreds of GPOs, and most hospitals belong to at least one. About 30 of the GPOs negotiate "sizeable contracts on behalf of their members."[9] Overall, it is estimated that "about 72% of purchases that hospitals make are done using GPO contracts,"[10] and that almost all (96 to 98% of) hospitals use GPOs for at least some of their purchasing.[11] Finally, GPO contracts fit our characterization of such contracts as nonlinear tariffs.[12]

The competitive concern most closely associated with the rise of GPOs appears to be one of agency failure (U.S. Department of Justice and Federal Trade Commission, 2004). GPOs could collude with suppliers to deny rival suppliers access to the healthcare market, pocketing fees from incumbent suppliers in return. We assume that the GPOs serve their members, and accordingly we do not introduce supplier/GPO collusion in our framework.

3. The model

Consider a standard Hotelling model of horizontal differentiation. Two suppliers offer competing, but not identical products. A consumer's location in product space is given by x, with supplier 1's location given by $x=0$, and the location of supplier 2 by $x=1$. Consumers are assumed to be uniformly distributed with density 1 along the interval [0, 1]. Each consumer is assumed to be willing to pay $v>0$ for a single unit of the good if that good's characteristic exactly matches the consumer's position on the interval. Willingness to pay is reduced by a factor $k>0$ per unit distance for a product whose characteristic does not match the consumer's preferred variety. Thus if the prices of the two firm's are denoted by p_1 and p_2, a consumer at x receives utility of $v-kx-p_1$ if the product is purchased from supplier 1, and utility of $v-(1-x)k-p_2$ when purchasing from supplier 2.

We assume that the distribution of customers is common knowledge, but that suppliers cannot identify the preferences of any individual customer. This limit on information prevents suppliers from engaging in price discrimination when dealing directly with consumers.[13]

[4] Manufacture française des pneumatiques Michelin v. Commission [2003] ECR II-4071.

[5] See European Commission, DG Competition, "DG Competition discussion paper on the application of Article 82 of the Treaty to exclusionary abuses," Brussels, December 2005.http://europa.eu.int/comm/competition/antitrust/others/discpaper2005.pdf.

[6] See Kinetic Concepts, Inc. v. Hillenbrand Industries, Inc., 262 F. Supp. 2d 722 (W.D. Texas, 2003).

[7] http://www.nellcor.com/legal/antitrust.aspx.

[8] The complaint against Amgen is similar to the loyalty discounts the Supreme Court struck down in SmithKline Corp. v. Eli Lilly & Co., 575 F.3d 1056 (3d Cir. 1978).

[9] The description in this paragraph is derived from a Unites States General Accounting Office report, "Group Purchasing Organization: Pilot Study Suggests Large Buying Groups Do Not Always Offer Hospitals Lower Prices," GAO-03-998T, July 16, 2003. The quotation is from page 6 of this report.

For a survey of the role of GPOs and the potential competition problems they pose, see U.S. Department of Justice and Federal Trade Commission (2004).

[10] Background data on health care GPOs are available from their trade association, HIGPA, http://www.higpa.org/about/about_faqs.asp.

[11] *Id.* GPO entry appears easy, with new GPOs such as Broadlane—formed in 1999 and a top seven GPO by 2002—able to grow rapidly. GPO members typically retain the outside option of dealing directly with suppliers. Many belong to multiple GPOs.

[12] The General Accounting Office (*Id.*, p. 9.) notes that GPO contracts have the characteristic that product prices fall as volume and share commitments are met. See also Mary Williams Walsh, "Senate Panel Weighs Tighter Rules for Hospital Suppliers, *New York Times*, September 15, 2004, http://www.nytimes.com/2004/09/15/business/15buyer.html?ex=1161403200&en=3f70190a62f85c9b&ei=5070.

[13] The restriction of suppliers to constant per unit prices also depends on the utility function we have assumed, namely that each consumer has unit demand. If consumers have elastic demands and a single crossing property is satisfied with respect to demand and type (x), then firms will be able to use nonlinear tariffs to screen consumers (that is, to practice second-degree price discrimination). See Spulber (1981) for the details.

Suppliers are assumed to have a common constant marginal cost, c. Assuming that the interval is covered (v is large relative to k), given p_1 and p_2 the demand functions are

$$D_i(p_i, p_j) = \frac{p_j - p_i + k}{2k},$$

The best-response functions are easily obtained and solved to yield the unique Bertrand–Nash equilibrium. Each firm offers a constant per unit price, $p_1 = p_2 = k + c$.

Thus far we have nothing more than a standard Hotelling model. Price competition does not lead to marginal cost pricing because each firm recognizes that when it cuts price, the lower price flows not only to consumers at the margin, but to every consumer. It would appear, then, that there are opportunities for lower prices if consumers can force the suppliers to compete for the business of all consumers, not just those at the margin.

3.1. Duopoly with the possibility of a GPO

Suppose that initially, all consumers in our market have the opportunity to form a single or universal GPO (cases with multiple GPOs are considered below). Should it form, the GPO is assumed to possess the same information as its suppliers. That is, the GPO knows the distribution of its members' preferences but cannot identify each consumer's individual preference. Given the prices offered by the suppliers, the GPO chooses to buy a proportion x of its needs from supplier 1, obtaining the remaining $1-x$ from supplier 2, $x \in [0, 1]$.[14] The decision variable of the GPO, x, becomes a continuous variable. With the GPO in place, suppliers can offer nonlinear tariffs offering volume discounts to influence the GPO's choice of x. As demonstrated below, this change in feasible tariff schedules changes the nature of competition.

The timing of the game is as follows. First, consumers decide whether to form a universal GPO. Second, if consumers have chosen to form a GPO, each supplier offers a nonlinear tariff schedule mapping from x to total payment. We use $r_i(\cdot)$, $i = 1, 2$ to denote these schedules, and we assume that the choice of schedules is noncooperative and that they are offered simultaneously. If no GPO is present, each supplier offers a single price p_i per unit with no quantity stipulation. The GPO, if present, then allocates its purchases by choosing x. Note that no bargaining occurs, as the schedules offered by suppliers to the GPO are on a take-it-or-leave-it basis. Alternatively, without a GPO, individual consumers choose their supplier. Finally, if present, the GPO allocates its purchases among its members. Note that in the first stage of the game, the decision whether to form the GPO is endogenous. To proceed, we first characterize the equilibria of the subgame where the GPO is formed in the first stage.

The GPO (a nonprofit organization) is assumed to have as its objective the maximization of the surplus of its consumer members. Though the GPO, like the suppliers, cannot identify the preferences of individual members, it can allocate its purchases efficiently for any given x. To see this, suppose that the GPO sets a price $p + k(1-2x)$ for good 1 and p for good 2. By means of this mechanism, all consumers located to the left of x select supplier 1's product and the remaining consumers pick supplier 2's product. These choices result in the efficient allocation of the products purchased by the GPO among its members. Moreover, the GPO has the flexibility to set p to balance its budget.[15]

Given this efficient allocation mechanism, the GPO's objective of maximizing the surplus of its members is equivalent to minimizing the sum of the payments to suppliers and the lost surplus that results from purchasing products that do not match consumer preferences. Thus the GPO's problem is

$$\max_x \left[-r_1(x) - r_2(1-x) - \int_0^x k\xi d\xi - \int_0^{1-x} k\xi d\xi \right],$$

which is equivalent to the following problem:

$$\min_x \left[r_1(x) + r_2(1-x) + k\left(\frac{x^2}{2} + \frac{(1-x)}{2}\right)^2 \right]. \qquad (1)$$

The GPO thus includes the welfare of all of its members in its computation, which changes the situation from the single price case, where the firms compete over the marginal customers. If the GPO is indifferent between $x = \frac{1}{2}$ and $x \neq \frac{1}{2}$, we assume that the GPO always picks $x = \frac{1}{2}$.[16]

Note that each revenue schedule $r_i(\cdot)$ is defined on the whole domain of [0, 1], and is, therefore, a complete schedule. Each schedule can also incorporate a fixed fee,

[14] The total purchase of the GPO is guaranteed to be 1, since v is assumed to be large relative to k.

[15] This setting resembles that in Anton and Yao (1989), who study split awards in a procurement auction.

[16] This tie-breaking rule is introduced to simplify the analysis.

in which case $r_i(0)>0$. We restrict attention to weakly increasing and differentiable revenue schedules. Now denote the marginal price schedule for each supplier by $a_i(\cdot)$,

$$a_1(x) = r_1'(x), a_2(1-x) = r_2'(1-x).$$

Then the first-order condition for the GPO's problem is

$$a_1(x) - a_2(1-x) + k(2x-1) = 0. \tag{2}$$

Notice that Eq. (2) is necessary but might not be sufficient, as a local maximizer might not be a global one.

Firm 1's problem can now be expressed as choosing its revenue schedule $r_1(\cdot)$ to maximize its profits, given $r_2(\cdot)$:

$$\max_{r_1(\cdot)} r_1(x) - cx,$$

subject to the requirement that the GPO's choice of x solves Eq. (1).

Lemma 1. *In any candidate equilibrium, the equilibrium market share must be symmetric. That is, in any candidate equilibrium each supplier must sell 1/2.*

Proof. Suppose that in a candidate equilibrium, the market share of supplier 1, denoted x^*, is such that $x^* \neq 1/2$. Without loss of generality, suppose $x^*>1/2$. First we show that there is a neighborhood of x^*, $[x^*-\epsilon, x^*]$ such that for $t\in[x^*-\epsilon, x^*]$, $a_1(t)\geq c$. Suppose the opposite, $a_1(t)<c$ for $t\in[x^*-\epsilon, x^*)$.
Then supplier 1 can deviate by setting very high marginal prices in $[x^*-\epsilon, x^*]$ (the marginal prices outside this region remain unchanged). This deviation can increase supplier 1's profit, since it no longer sells in the region $[x^*-\epsilon, x^*]$ in which it incurs strictly negative profits under the original tariff.

Now consider a deviation for supplier 2. Suppose that it charges $a_2(1-t)=c+\varepsilon$ for $t\in[x^*-\epsilon, x^*]$, with $\varepsilon>0$ but small, and retains the original tariff elsewhere. Given $a_1(t)\geq c$ for $t\in[x^*-\epsilon, x^*]$ and $x^*>1/2$, by Eq. (2) the GPO will buy $1-x^*+\epsilon$ from supplier 2. This deviation is profitable since supplier 2 obtains a profit margin ε from additional sales, and thus results in a contradiction. Therefore, in any candidate equilibrium $x^*=1/2$. □

From Lemma 1, in seeking equilibria we focus without loss of generality on tariff profiles for which market shares are symmetric. We refer to $x\in\left[0,\frac{1}{2}\right]$ as supplier 1's territory and $x\in\left[\frac{1}{2},1\right]$ as supplier 2's territory.

Lemma 2. *In any equilibrium, neither supplier can earn a profit strictly greater than* $\frac{k}{4}$.

Proof. Consider a tariff profile that results in a symmetric market share, and suppose supplier 1 earns a profit $\pi>\frac{k}{4}$. Now consider a deviation by supplier 2 in which it offers $a_2(1-x)=c+\varepsilon$ for $x\in[0, 1/2]$, with ε very small. Then the GPO has incentive to set $x^*=0$ (The GPO purchases nothing from supplier 1). To see this, compared to $x^*=1/2$, the GPO saves $\pi-\frac{k}{4}-\frac{1}{2}\varepsilon>0$ ($\frac{k}{4}$ is the increase in total transportation cost) by setting $x^*=0$. Since this is a profitable deviation, we can conclude that in any equilibrium each supplier cannot earn a profit strictly greater than $\frac{k}{4}$. □

In what follows, we first find an equilibrium, which will serve as a benchmark for other equilibria.

Lemma 3. *The following tariff profile is an equilibrium, with each firm selling 1/2 and earning profit* $\frac{k}{4}$:

$$a_1(x) = \begin{cases} c + k(1-2x), & x\leq\frac{1}{2} \\ c, & x>\frac{1}{2} \end{cases} \tag{3}$$

$$a_2(1-x) = \begin{cases} c + k(2x-1), & x\geq\frac{1}{2} \\ c, & x<\frac{1}{2} \end{cases} \tag{4}$$

Proof. Under this tariff profile, the GPO is indifferent among all $x\in[0, 1]$, since Eq. (2) is satisfied everywhere. It will set $x^*=1/2$, and each supplier earns profit $k/4$. Now we show that given the other supplier's tariff, no supplier has an incentive to deviate. Fix supplier 2's tariff schedule, and consider supplier 1's incentive to deviate. In this situation, supplier 1 has no incentive to change its marginal prices for $x\in\left[\frac{1}{2},1\right]$. If it increases its marginal prices in this range, the GPO will still buy 1/2 from it and its profit is unchanged. If it decreases its marginal prices (below the marginal cost c) in this range, the GPO could choose to continue to purchase 1/2 from it so that its profit remains unchanged. Alternatively, the GPO could increase its purchases from this supplier, buying more than 1/2 from it. In this case, the supplier incurs a profit loss as a consequence of selling additional units below marginal cost.

Next we show that supplier 1 has no incentive to change its tariff schedule for $x\in[0, 1/2]$. Under the original tariff schedule, $r_1(x)=cx+k(x-x^2)$ for

$x \in [0, 1/2]$, and $r_1\left(\frac{1}{2}\right) = \frac{c}{2} + \frac{k}{4}$. To show supplier 1 has no incentive to increase $r_1(x)$ for $x \in \left[0, \frac{1}{2}\right]$, note that the GPO's opportunity cost of getting x additional units from supplier 2 is $C(x) = cx + k(x - x^2)$, since supplier 2 is charging a marginal price equal to marginal cost, c for $x \in \left[0, \frac{1}{2}\right]$. Under the original tariff, $r_1(x) = C(x)$. If supplier 1 changes its tariff such that $r_1\left(\frac{1}{2}\right)$ strictly exceeds $\frac{c}{2} + \frac{k}{4}$,[17] the GPO will either buy nothing from supplier 1 or will only buy $x' < \frac{1}{2}$ if there is a x' such that $r_1(x') \le C(x')$. Note that the profit of supplier 1 under the original tariff if the GPO buys x from it is $k(x - x^2)$, which is increasing in x. Thus in either case, supplier 1's profit decreases. Therefore, no deviation can increase supplier 1's profit. Thus the original tariff profile is an equilibrium. □

The equilibrium specified in Lemma 3 has the feature that each firm in its own territory charges marginal prices equal to the GPO's opportunity cost of obtaining a marginal unit from the other firm at the marginal cost, thus the marginal profit margin equals to the GPO's preference advantage of one firm's product over the other firm's. However, the first order condition (2) is necessary only at the equilibrium market share x^*. There are many equilibrium tariff schedules that yield the same equilibrium outcome. For example, the following two-part tariff schedules also consist an equilibrium. Each supplier requires an up-front fee of $k/4$ and offers to sell additional units at a price of c per unit.[18] Other tariff schedules can achieve the same equilibrium outcome. For instance, each firm could set marginal prices equal to marginal cost in the other firms' territory, and set $r_i\left(\frac{1}{2}\right) = \frac{c}{2} + \frac{k}{4}$. Moreover,

$$r_1(x) \ge \int_0^x [c + k(1 - 2t)]\mathrm{d}t = cx + kx - kx^2 \text{ for } x \in \left[0, \frac{1}{2}\right)$$
$$r_2(1 - x) \ge \int_{1-x}^1 [c + k(2t - 1)]\mathrm{d}t = cx + k(1 - x) - k(1 - x)^2 \text{ for } x \in \left(\frac{1}{2}, 1\right].$$

The inequalities above ensures that the GPO has no incentive to set $x^* \neq \frac{1}{2}$. Note that in this type of equilibria, equilibrium tariff schedules can be asymmetric, and in each firm's own territory the marginal prices might be below the marginal cost in some range.

Note that thus far we have not imposed $a_i(\cdot) > c$. Henceforth, we adopt the assumption that each supplier's marginal price is greater than or equals to c in the other firm's territory. That is, $a_1(t) \ge c$ for $t \in \left[\frac{1}{2}, 1\right]$, and $a_2(1 - t) \ge c$ for $t \in \left[0, \frac{1}{2}\right]$.[19] With this assumption, we show that all equilibria yield to the same equilibrium outcome, as summarized in the following proposition.

Proposition 1. *Assuming each supplier's marginal price is greater than or equal to c in the other firm's territory, there are multiple equilibria but all equilibria have the same equilibrium outcome. Each supplier sells 1/2 and earns a profit of $k/4$.*

Proof. First we show that $a_1(x) = c$ for $x \in \left[\frac{1}{2}, 1\right]$, and $a_2(1 - x) = c$ for $x \in \left[0, \frac{1}{2}\right]$ almost everywhere. Suppose there is an interval $(x_1, x_2) \in \left[\frac{1}{2}, 1\right]$ such that $a_1(x) > c$ for any $x \in (x_1, x_2)$. Now suppose that supplier 2 offers a tariff schedule consisting of an up-front fee of $\frac{k}{4} + \int_{x_1}^{x_2} [a_1(t) - c]\mathrm{d}t - \varepsilon$ with ε very small, and charges c for each additional unit. Then the GPO will buy at least 1/2 from supplier 2, since its total cost of acquiring the additional 1/2 from supplier 1 is $\frac{k}{4} + \int_{\frac{1}{2}}^1 a_1(t)\mathrm{d}t \ge \frac{k}{4} + \frac{1}{2}c + \int_{x_1}^{x_2} [a_1(t) - c]\mathrm{d}t$ (the inequality follows from the fact that $a_1(x) \ge c$ for $x \ge \frac{1}{2}$). But now supplier 2's profit becomes $\frac{k}{4} + \int_{x_1}^{x_2} [a_1(t) - c]\mathrm{d}t - \varepsilon > \frac{k}{4}$. This is a contradiction to Lemma 2. Thus in equilibrium there is no interval $(x_1, x_2) \in \left[\frac{1}{2}, 1\right]$ such that $a_1(x) > c$ for any $x \in (x_1, x_2)$. Combining this with the fact that $a_1(x) \ge c$ for $x \ge \frac{1}{2}$, in any equilibrium, $a_1(x) = c$ for $x \in \left[\frac{1}{2}, 1\right]$ almost everywhere. Similarly, we can show that $a_2(1 - x) = c$ for $x \in \left[0, \frac{1}{2}\right]$ almost everywhere.

Given that each firm always charges marginal price in the other firm's territory, Lemmas 2 and 3 imply that in any equilibrium, each supplier must earn a profit of $k/4$,

[17] This deviation can result from the following tariff change: adding a positive fixed fee and the marginal prices remain unchanged.

[18] Fix the other supplier's tariff schedule, if a supplier increases its fixed fee, then the GPO will buy nothing from the deviating firm and buy solely from the other firm.

[19] Absent this assumption, we have a continuum of unreasonable equilibria. For example, the following tariff profile is an equilibrium with each supplier earns zero profit:

$$r_1(x) = cx \text{ for } x \in \left[0, \frac{1}{2}\right], a_1(x) = c + k(1 - 2x) \text{ for } x \in \left[\frac{1}{2}, 1\right]$$
$$r_2(1 - x) = c(1 - x) \text{ for } x \in \left[\frac{1}{2}, 1\right], a_2(1 - x) = c + k(2x - 1) \text{ for } x \in \left[0, \frac{1}{2}\right]$$

This equilibrium has the feature that each firm's ability to extract rent in its own territory is restricted by the other firm's very aggressive (below marginal cost) pricing. A supplier does not benefit from aggressive pricing in its rival's territory. Such aggression merely dissipates its rival's profit. With such beggar-thy-neighbor aggression, we can have a continuum of equilibrium outcomes with a supplier's profit ranging from $\left[0, \frac{k}{4}\right]$, depending on how aggressive the other firm's marginal prices are set beyond its own territory.

We believe that all equilibria with a supplier's profit strictly less than $\frac{k}{4}$ are unreasonable. It is more plausible to suppose that suppliers will coordinate on the equilibrium with the highest available symmetric payoff, $k/4$. Moreover, each of the equilibria with payoffs strictly less than $\frac{k}{4}$ entails submarginal cost pricing in the rival's territory. If the GPO chooses asymmetric market shares from among its global maximizers with a positive probability, then below marginal cost pricing in the other firm's territory will reduce a firm's profit, which destroys this type of equilibria.

for otherwise a supplier can always earn at least this amount by deviating unilaterally to the tariff specified in Lemma 3. Therefore, though equilibrium tariff schedules are not unique, the equilibrium outcome is unique and symmetric. Each supplier earns $\frac{k}{4}$ and sells 1/2. □

Absent the GPO, each firm obtains an equilibrium profit of $k/2$, and each consumer pays $c+k$ for that consumer's preferred version of the good. But with the GPO, each firm's equilibrium profit becomes $k/4$, and each consumer pays $c+k/2$ for the preferred good. Thus we have the following proposition.

Proposition 2. *With the introduction of a GPO, each supplier's equilibrium profit decreases by* $\frac{k}{4}$ *and consumers obtain their desired good at a reduced price.*

Thus consumers have a strong incentive to band together by forming a GPO in the first stage. Doing so yields consumers their preferred good at a reduced price. Note that the role of the GPO is not to increase the bargaining power of its members, since it must accept the tariffs proffered independently by each supplier. The only change introduced by the addition of a GPO is that suppliers are able to employ nonlinear tariffs. The introduction of the GPO does, however, change the nature of supplier competition. Without the GPO, two firms simply compete for the marginal consumer located at 1/2. When facing the GPO, suppliers can compete for *all* consumers, leading to fierce competition and lower average prices.

To identify the force that amplifies competition, suppose that suppliers can commit to charging only one price in the presence of the GPO. Given p_1 and p_2, the GPO will set $x^* = \frac{1}{2} + \frac{p_2-p_1}{2k}$. Thus the demand functions are identical to those in the benchmark case without the GPO, yielding the equilibrium price $p_1 = p_2 = c+k$. This implies that when facing the GPO, the rival suppliers are better of if they can commit to charging a single per unit price. But for the suppliers, this commitment is not credible. Given that one supplier offers a single price $c+k$, its rival has an incentive to offer a nonlinear tariff to increase its market share and thereby its profit.

The above discussion implies that it is the availability of nonlinear tariffs, made possible by the presence of the GPO, that strengthens competition. The benefit of the GPO to its members comes directly from the ability of the GPO to pit suppliers against one another to a greater degree than occurs with simple per unit prices. Since firms cannot identify each individual consumer's preference, the GPO is a mechanism both for increasing the dimensions on which firms compete and for revealing (or aggregating) information.

More formally, suppose that supplier 1 is a monopolist. In the absence of a GPO, this supplier will set $p = \max\left\{\frac{v+c}{2}, v-k\right\}$, so that it covers the market as long as c is small in relation to the valuation of the consumers who least value firm 1's product. Firm 1's share of the potential market is $x = \min\left\{\frac{v-c}{2k}, 1\right\}$. The addition of the GPO ensures that the market is covered as long as $v-k>c$, but it also permits the supplier to set the marginal price $a(x)=v-kx$ and extract all consumer surplus.

Moving to duopoly, a nonlinear tariff provides the same opportunity to extract surplus as in the monopoly case. But with duopoly, this is only one part of the story. With a nonlinear tariff, a firm can compete for any consumer, without worrying about that such competition reduces the price charged for its inframarginal consumers. However, if both suppliers engage in these attempts to augment share through low prices at their margins, each supplier's ability to extract surplus is now limited by its rival's aggressive pricing. In our model, with a uniform distribution of consumer preferences and symmetric firms, suppliers are induced to defend *all* of their customers. This effect of nonlinear tariffs in intensifying competition always dominates the impact of the ability to extract surplus, yielding lower profits and higher consumer surplus.

Armstrong and Vickers (2006, Section 2) show that allowing volume discounts can lower consumer surplus and raise profits. Thanassoulis (2007) provides a similar result with firm-specific preferences. These results are directly opposite to ours. The contrast derives from differences in the nature of the admissible volume discounts. In their models, nonlinear pricing is used to sort different consumers in the vertical dimension, and it has only an indirect impact on competition in the horizontal dimension. With linear prices, the incentive for firms to compete for high demand consumers leads to lower prices for all consumers. Nonlinear prices permit suppliers to compete for different consumers in the vertical dimension freely, resulting in higher profits, higher average prices, and lower consumer welfare. Note that in their models with or without nonlinear prices, competition in the horizontal dimension occurs only for marginal consumers. By contrast, in our model nonlinear tariffs are employed to target each individual consumer according to its location in the horizontal dimension, where each consumer cares about the specific characteristics of the available products. With linear prices, competition only occurs for marginal consumers. With nonlinear tariffs competition occurs

globally and is more intense. This is why in our model, using nonlinear tariffs leads to lower profits and higher consumer surplus.

We acknowledge that a GPO as a whole can do better when it commits to buy exclusively from one supplier. Exclusive contracts can drive down the price all the way to the marginal cost, as shown by Dana (2003). But the added surplus capture is coupled by a loss in efficiency. Under exclusive contracts, consumers whose preference is very different from the chosen product can be worse off than without exclusive contracts. The customer who prefers most strongly the option that is given up when an exclusive contract is negotiated incurs a utility loss of $c+k$, in contrast to the utility loss of $c+k2$ with nonlinear tariffs and no exclusive contracts. Thus exclusive contracts will lead to internal conflict within a GPO, as consumers fight to buy exclusively from the product they prefer. This problem is absent without exclusive contracts, as each consumer gets the product he prefers.

We would thus expect exclusive contracts to emerge for products for which consumers were nearly indifferent among the products of rival suppliers, while strong horizontal differentiation would yield multi-source contracting. In the health care marketplace, one might expect that preferences would be weaker among alternative medical devices than for pharmaceuticals. Nevertheless, even for medical device purchasing, sole source contracting is less common than dual- or multi-source purchasing. The pattern of contracting is broadly consistent with the presence of strong consumer preferences:

> The contracts negotiated by GPOs under which their members purchase medical devices are almost always much less than exclusive dealing… Typically, the agreements at issue offer buyers a discount of a certain percentage in exchange for the buyers' commitment to purchase a minimum percentage of its needs from a designated purchaser or perhaps a small number of designated purchasers.[20]

In practice, sole-source contracting is employed by GPOs, but the degree of reliance on such contracts varies considerably across GPOs. Sole-source contracts account for 19% of medical–surgical dollar purchasing volume for the largest GPO and 42% for the second. For the five other large GPOs studied by the U.S. General Accounting Office, the percentage of dollar volume varied from 2 to 46% of purchasing volume. One of the largest two GPOs reported using sole-source contracts extensively for "clinical preference" products—those for which hospitals and doctors expressed strong preferences for the products of particular suppliers. But that firm was an outlier. For the remainder of the GPOs, "commodity" products constituted the bulk of their sole-source purchasing.[21] Thus most GPO purchasing occurs under conditions consistent with those assumed in our model.

4. Generalizing the model

The simple model of the previous section assumed consumers are uniformly distributed and that preferences for the two firms' products are symmetric. In this section we study more general distributions. Consider two suppliers, 1 and 2, selling differentiated products to consumers who have different tastes. We continue to assume that consumers have unit demand. Let a consumer's valuation of supplier 1's and supplier 2's good be v_1 and v_2 respectively. In particular, $v_1=v-y/2$ and $v_2=v+y/2$. Thus $y=v_2-v_1$ measures a consumer's preference for good 2 over good 1. The variable y is distributed according to $F(y)$ on $[\underline{y}, \overline{y}]$ with $\underline{y}<0<\overline{y}$. We assume that the corresponding density function $f(y)$ is continuous and strictly positive everywhere on $[\underline{y}, \overline{y}]$. The total mass of consumers is normalized to 1. To guarantee the existence of a unique equilibrium in Bertrand competition, we impose the following monotone hazard rate assumptions:

$$\forall\, y\in[\underline{y}, \overline{y}], \frac{\mathrm{d}}{\mathrm{d}y}\left[\frac{F(y)}{f(y)}\right]>0>\frac{\mathrm{d}}{\mathrm{d}y}\left[\frac{1-F(y)}{f(y)}\right]. \tag{5}$$

Note that uniform distribution satisfies the above assumptions. As in the base model, we assume that a consumer's preference y is his private information, though the distribution of y is public information. Firms have the same constant marginal cost c. We assume that v, the basic value of the goods, is large enough to ensure that the entire market is covered.

4.1. Case 1: no GPO

We first consider the case without a GPO. Each supplier is confined to compete by setting a single price.

[20] Herbert Hovenkamp, *Group Purchasing Organization (GPO) Purchasing Agreements and Antitrust Law* (January 2004) (prepared on behalf of Health Industry Group Purchasing Association), pages 3–4.

[21] See *Group Purchasing Organizations: Use of Contracting Processes and Strategies to Award Contracts for Medical–Surgical Products: Before the Subcomm. on Antitrust, Competition Policy and Consumer Rights of the S. Comm. on the Judiciary, 108th Cong.* (2003) (testimony of U.S. General Accounting Office).

Let p_1 and p_2 be supplier 1's and 2's prices respectively. Consumers with $y<p_2-p_1$ will buy from firm 1 and those with $y>p_2-p_1$ will buy from firm 2. Firm i's profits can be written as:

$$\pi_1(p_1,p_2) = (p_1 - c)F(p_2 - p_1)$$
$$\pi_2(p_1,p_2) = (p_2 - c)[1 - F(p_2 - p_1)].$$

Firm i's objective is to maximize π_i, given p_j , by choosing p_i. The equilibrium prices satisfy the following first order conditions:

$$p_1^* = c + \frac{F(y^*)}{f(y^*)} \text{ and } p_2^* = c + \frac{1 - F(y^*)}{f(y^*)}$$

where $y^* \equiv p_2^* - p_1^*$ is the preference of the marginal consumer who is indifferent between buying two goods.

To show the existence and uniqueness of the equilibrium, note that y^* satisfies:

$$\frac{F(y^*)}{f(y^*)} - \frac{1 - F(y^*)}{f(y^*)} = -y^* \tag{6}$$

Define

$$H(y) \equiv \frac{F(y)}{f(y)} - \frac{1 - F(y)}{f(y)} + y. \tag{7}$$

Then y^* is defined by $H(y^*)=0$. Given assumption (5), $H(y)$ strictly increasing in y. Therefore, if an equilibrium exists, it must be unique. Also note that $H(\underline{y})<0$ and $H(\bar{y})>0$. Therefore, there is a unique equilibrium, since there is a unique y^* satisfying $H(y^*)=0$.

Suppose $F(0)=1/2$. That is, half of the consumers prefer good 1, so that the two firms' market positions are symmetric. It can be easily shown that $y^*=0$. Thus the equilibrium outcome is symmetric: $p_1^*=p_2^*$ and each firm's market share is 1/2.

Suppose $F(0)>1/2$. That is, more consumers prefer good 1, so that supplier 1 has a preference advantage. Now $H(0)>0$. Given the monotonicity of $H(y)$, we must have $y^*<0$. The equilibrium market share of supplier 1 is $F(y^*)>1/2$. To see this, note that if $F(y^*)\le 1/2$, then $H(y^*)<0$ since $y^*<0$, a contradiction. Given $F(y^*)>1/2$, the equilibrium prices exhibit $p_1^*>p_2^*$.

4.2. Case 2: with a GPO

Suppose that almost all consumers have the opportunity to form a universal GPO. The remaining consumers, with a total measure ε very small, do not join the GPO for some exogenous reason. We label these consumers as nonactive. We assume that nonactive consumers have the same distribution $F(y)$ as consumers as a whole. This implies that active consumers are distributed according to $F(y)$ as well. According to the previous analysis, firms charge single prices p_1^* and p_2^* to nonactive consumers. The prices p_1^* and p_2^* also serve as an active consumer's outside option if he does not join the GPO. From this point forward, we focus on active consumers with measure $1-\varepsilon$, and drop the ε in our derivations. We first study the equilibrium outcome with a universal GPO. We then derive a condition under which a universal GPO will form.

We first consider the equilibrium outcome if the GPO is formed. Suppliers now can offer nonlinear tariffs to influence the GPO's purchase shares. Denote $a_i(y)$ as firm i's marginal price schedule. The corresponding revenue schedules for the suppliers are:

$$r_1(y) = \int_{\underline{y}}^{y} a_1(\tilde{y})f(\tilde{y})\mathrm{d}\tilde{y};\ r_2(y) = \int_{y}^{\bar{y}} a_2(\tilde{y})f(\tilde{y})\mathrm{d}\tilde{y}$$

Again, we restrict consideration to continuously differentiable revenue schedules. The GPO cannot identify each consumer's preference. However, given that the GPO buys $F(y)$ from firm 1 and $1-F(y)$ from firm 2, it can efficiently allocate the products among its members. Specifically, the GPO can set a price p for good 1 and $p+y$ for good 2. Then consumers with their preferences to the left of y will buy good 1 and the remaining consumers will buy good 2. Moreover, p can be set properly to balance the GPO's budget.

The GPO's objective is to maximize its members' total welfare, given $r_1(\cdot)$ and $r_2(\cdot)$ ($a_1(\cdot)$ and $a_2(\cdot)$), by choosing y (equivalent to choosing $F(y)$). In particular, the GPO's problem becomes:

$$\max_{y} \int_{\underline{y}}^{y} \left(v - \frac{\tilde{y}}{2}\right) f(\tilde{y})\mathrm{d}\tilde{y} + \int_{y}^{\bar{y}} \left(v + \frac{\tilde{y}}{2}\right) f(\tilde{y})\mathrm{d}\tilde{y} - r_1(y) - r_2(y)$$

Supplier 1's problem is to choose $r_1(\cdot)$, given $r_2(\cdot)$, to maximize its profit $r_i(y)-cF(y)$, subject to y solving the GPO's problem. Supplier 2's problem is similar. Corresponding to Proposition 1 in the base model, we have the following proposition.

Proposition 3. *Assume $a_1(y)\ge c$ for $y\ge 0$ and $a_2(y)\ge c$ for $y\le 0$. There are multiple equilibria. However, all equilibria have the same equilibrium outcome:*

i) The market share of supplier 1 is F(0) and the market share of supplier 2 is $1-F(0)$.
ii) $a_1(y)=c$ for $y\in[0, \bar{y}]$ almost everywhere and $a_2=c$ for $y\in[\underline{y}, 0]$ almost everywhere.
iii) $\pi_1^*=\int_{\underline{y}}^{0} -y\, f(y)\, dy$ and $\pi_2^*=\int_{0}^{\bar{y}} y\, f(y)\, dy$.

The proof of this proposition is similar to that for the uniform distribution case. If the marginal (or indifferent) consumer y^* is not zero, say $y^*>0$, then firm 2 can deviate profitably by charging a marginal price $c+\varepsilon$ in some neighborhood $[y^*-\epsilon, y^*]$. Thus the equilibrium marginal consumer must be $y^*=0$. Now in the other firm's territory, each firm charges a marginal price equals to marginal cost, since otherwise at least one firm has an incentive to deviate. Finally, given that its rival is charging marginal cost in its territory, a firm can charge a marginal price equal to marginal cost plus its preference advantage in its own territory. This leads to the equilibrium profits in property (iii). The property of multiple equilibria is also familiar. The following two part linear tariffs are also an equilibrium: each supplier charges an up front fee π_i^* and then charging each additional unit at marginal cost c.

Thus in equilibrium the GPO buys $F(0)$ from firm 1 and $1-F(0)$ from firm 2. To ensure efficient allocation in the GPO, the GPO sets equal internal prices for two goods, that is, $p_1^{\mathrm{I}}=p_2^{\mathrm{I}}=p^{\mathrm{I}}$. Therefore, the presence of the GPO always leads to the efficient allocation of goods. Denote firm i's profit under GPO as π_i^{GPO}.

Specifically,

$$\pi_1^{\mathrm{GPO}} = \int_{\underline{y}}^{0} -y f(y)\mathrm{d}(y) = -yF(y)|_{\underline{y}}^{0} + \int_{\underline{y}}^{0} F(y)\mathrm{d}y$$
$$= \int_{\underline{y}}^{0} F(y)\mathrm{d}y = \int_{\underline{y}}^{0} \frac{F(y)}{f(y)} f(y)\mathrm{d}y$$
$$\pi_2^{\mathrm{GPO}} = \int_{0}^{\bar{y}} y f(y)\mathrm{d}y = yF(y)|_{0}^{\bar{y}} - \int_{0}^{\bar{y}} F(y)\mathrm{d}y$$
$$= \int_{0}^{\bar{y}} [1-F(y)]\mathrm{d}y = \int_{0}^{\bar{y}} \frac{1-F(y)}{f(y)} f(y)\mathrm{d}y$$

Note that we use integration by parts in the derivations. The internal price p^{I} can be expressed as:

$$p^{\mathrm{I}} = c + \pi_1^{\mathrm{GPO}} + \pi_2^{\mathrm{GPO}}$$
$$= c + \int_{\underline{y}}^{0} \frac{F(y)}{f(y)} f(y)\mathrm{d}y + \int_{0}^{\bar{y}} \frac{1-F(y)}{f(y)} f(y)\mathrm{d}y.$$

4.3. Comparison

Our results thus far permit us to determine whether active consumers have incentives to form a GPO in the first stage of our game. If the GPO forms, we can then determine its impact on market outcomes. We consider symmetric and asymmetric preferences in turn.

4.3.1. Symmetry

We first consider the case where $F(0)=1/2$, so that the two firms' market positions are symmetric. Note that the distribution of consumers can be asymmetric with respect to $y=0$. Based on our previous analysis, the equilibrium market shares are 1/2 for each supplier, with or without a GPO. We show that each firm's profit decreases under the GPO.

$$\pi_1^{\mathrm{GPO}} - \pi_1^* = \int_{\underline{y}}^{0} \frac{F(y)}{f(y)} f(y)\mathrm{d}y - \frac{F(0)}{f(0)} F(0)$$
$$= \int_{\underline{y}}^{0} \left[\frac{F(y)}{f(y)} - \frac{F(0)}{f(0)}\right] f(y)\mathrm{d}y < 0$$

$$\pi_2^{\mathrm{GPO}} - \pi_2^* = \int_{0}^{\bar{y}} \frac{1-F(y)}{f(y)} f(y)\mathrm{d}y - \frac{1-F(0)}{f(0)} [1-F(0)]$$
$$= \int_{0}^{\bar{y}} \left[\frac{1-F(y)}{f(y)} - \frac{1-F(0)}{f(0)}\right] f(y)\mathrm{d}y < 0$$

The inequalities are due to the monotone hazard rate property (assumption (5)). Since the market shares are the same with or without GPO, we must have $p^{\mathrm{I}}<p_i^*$. Thus active consumers have incentives to join the GPO. The following proposition summarizes the results.

Proposition 4. *If two firms' market positions are symmetric ($F(0)=1/2$) and assumption (5) is satisfied, then consumers have incentives to form a GPO. In the presence of a GPO, each firm's profit decreases and each consumer is strictly better off.* □

4.3.2. Asymmetry

Now we turn to the asymmetric case. Without loss of generality, we assume that $F(0)>1/2$, that is, supplier 1 has the preference advantage. We first show that the disadvantaged firm, firm 2, has a lower profit under the GPO. From previous analysis, we know that without the GPO the marginal consumer $y^*<0$. Thus supplier 2's market share is smaller under the GPO. Now we compare the average prices or average profit margins.

$$\frac{\pi_2^{\mathrm{GPO}}}{1-F(0)} - \frac{1-F(y^*)}{f(y^*)}$$
$$= \frac{1}{1-F(0)} \int_{0}^{\bar{y}} \left[\frac{1-F(y)}{f(y)} - \frac{1-F(y^*)}{f(y^*)}\right] f(y)\mathrm{d}y < 0.$$

The inequality results from assumption (5). Hence supplier 2's average price is lower under the GPO. Combining the lower price with its decreased market share, supplier 2 has a lower profit under the GPO.

For supplier 1, the advantaged firm, the impact of GPO on profit is ambiguous. This is because under GPO it gets a bigger market share. For consumers to have incentive to form the universal GPO, each individual consumer should benefit from forming the GPO. By joining a universal GPO, a consumer is able to purchase that consumer's preferred product at the internal price, p^{I}. Without the GPO, the consumer pays either p_1^* or p_2^*

and may not purchase the preferred product. Thus a sufficient condition for a consumer's willingness to join a universal GPO is $p^{\mathrm{I}} \le p_1^*$ and $p^{\mathrm{I}} \le p_2^*$. Since $p_1^* > p_2^*$, the binding constraint is $p^{\mathrm{I}} \le p_2^*$. Therefore, each consumer has an incentive to join the GPO if and only if

$$p^{\mathrm{I}} \le p_2^* \Leftrightarrow \int_{\underline{y}}^{0} \frac{F(y)}{f(y)} f(y)\mathrm{d}y + \int_{0}^{\bar{y}} \frac{1-F(y)}{f(y)} f(y)\mathrm{d}y \le \frac{1-F(y^*)}{f(y^*)}. \tag{8}$$

One sufficient condition for Eq. (8) to be satisfied is that

$$\int_{\underline{y}}^{0} \frac{F(y)}{f(y)} f(y)\mathrm{d}y + \int_{0}^{\bar{y}} \frac{1-F(y)}{f(y)} f(d)\mathrm{d}y \le \frac{1-F(0)}{f(0)},$$

since $\frac{1-F(0)}{f(0)} < \frac{1-F(y^*)}{f(y^*)}$. Suppose Eq. (8) is satisfied, then the formation of GPO not only makes consumers better of, but also leads to socially efficient allocation of goods. This is because nonlinear tariffs remove the price distortion caused by single price competition.

Note that if $F(0)=1/2$, then Eq. (8) is always satisfied.[22] This suggests that symmetry of two firms' market positions might affect Eq. (8). For this purpose, we provide the following comparative statics. Let the initial distribution of consumers be $F(y)$ on $[\underline{y}, \bar{y}]$ with $\underline{y}<0<\bar{y}$ and $F(0)>1/2$. Now suppose each consumer's preference over firm 1's product is increased by a small amount $\Delta v>0$ with $\bar{y}-\Delta v>0$. Define $y-\Delta v=y'$. It follows that y' distributes on $[\underline{y}-\Delta v, \bar{y}-\Delta v]$ with a cumulative distribution $G(y')$ and density function $g(y')$. Note that $F(y)=G(y')$. It is easy to see that two firms' market positions are more asymmetric under $G(y')$ than under $F(y)$.

Lemma 4. *Eq. (8) is more difficult to satisfy under $G(y')$ than under $F(y)$.*

Proof. It is sufficient to show that

$$\int_{\underline{y}}^{0} \frac{F(y)}{f(y)} f(y)\mathrm{d}y + \int_{0}^{\bar{y}} \frac{1-F(y)}{f(y)} f(y)\mathrm{d}y < \int_{\underline{y}-\Delta v}^{0} \frac{G(y')}{g(y')} g(y')\mathrm{d}y' + \int_{0}^{\bar{y}-\Delta v} \frac{1-G(y')}{g(y')} g(y')\mathrm{d}y', \tag{9}$$

and

$$\frac{1-F(y^*)}{f(y^*)} > \frac{1-G(y'^*)}{g(y'^*)}. \tag{10}$$

We first prove Eq. (9).

$$\begin{aligned}
&\int_{\underline{y}}^{0} \frac{F(y)}{f(y)} f(y)\mathrm{d}y + \int_{0}^{\bar{y}} \frac{1-F(y)}{f(y)} f(y)\mathrm{d}y \\
&\quad - \int_{\underline{y}-\Delta v}^{o} \frac{G(y')}{g(y')} g(y')\mathrm{d}y' - \int_{0}^{\bar{y}-\Delta v} \frac{1-G(y')}{g(y')} g(y')\mathrm{d}y' \\
&\quad = \int_{\underline{y}}^{0} \frac{F(y)}{f(y)} f(y)\mathrm{d}y + \int_{0}^{\bar{y}} \frac{1-F(y)}{f(y)} f(y)\mathrm{d}y \\
&- \int_{\underline{y}}^{\Delta v} \frac{F(y)}{f(y)} f(y)\mathrm{d}y - \int_{\Delta v}^{\bar{y}} \frac{1-F(y)}{f(y)} f(y)\mathrm{d}y \\
&\quad = - \int_{0}^{\Delta v} \left[\frac{F(y)}{f(y)} - \frac{1-F(y)}{f(y)}\right] f(y)\mathrm{d}y < 0.
\end{aligned}$$

The last inequality comes from the fact that $F(0)>1/2$ combined with our monotone hazard rate condition. We then show Eq. (10), which is equivalent to $y^* < y'^* + \Delta v$. Note that y'^* is defined by $H_G(y'^*)=0$, where $H_G(\cdot)$ is the function defined similarly to Eq. (7) under distribution $G(y')$. Like $H(y)$, $H_G(y')$ is increasing in y'.

$$\begin{aligned}
H_G(y^* - \Delta v) &= \frac{G(y^* - \Delta v)}{g(y^* - \Delta v)} - \frac{1 - G(y^* - \Delta v)}{g(y^* - \Delta v)} \\
&\quad + y^* - \Delta v = \frac{F(y^*)}{f(y^*)} - \frac{1-F(y^*)}{f(y^*)} \\
&\quad + y^* - \Delta v = -y^* + y^* - \Delta v = -\Delta v < 0.
\end{aligned}$$

Given the monotonicity of $H_G(\cdot)$, the above inequality implies that $y'^* > y^* - \Delta v$. This further implies that

$$\frac{1-G(y'^*)}{g(y'^*)} < \frac{1-G(y^* - \Delta v)}{g(y^* - \Delta v)} = \frac{1-F(y^*)}{f(y^*)}. \qquad \square$$

Lemma 4 is based on two results. Inequality (9) says that firms' joint profit under the GPO increases as the two firms' market positions become more asymmetric. This means that the GPO's internal price increases in response to more asymmetric market positions. Inequality (10) says that the disadvantaged firm's price under single price competition is lower when two firms' market positions become more asymmetric. Combining these two effects, we can see that consumers who prefer the disadvantaged firm's product have less incentive to join the GPO as firms' market positions become more asymmetric.

Combining Lemma 4 and that fact that Eq. (8) is satisfied when firms' market positions are symmetric, we conclude that the GPO will form as long as two firms'

[22] By the monotone hazard rate property (5),

$$\begin{aligned}
&\int_{\underline{y}}^{0} \frac{F(y)}{f(y)} f(y)\mathrm{d}y + \int_{0}^{\bar{y}} \frac{1-F(y)}{f(y)} f(y)\mathrm{d}y \le \int_{\underline{y}}^{0} \frac{F(0)}{f(0)} f(y)\mathrm{d}y \\
&\quad + \int_{0}^{\bar{y}} \frac{1-F(0)}{f(0)} f(y)\mathrm{d}y = \int_{\underline{y}}^{0} \frac{1-F(0)}{f(0)} f(y)\mathrm{d}y \\
&\quad + \int_{0}^{\bar{y}} \frac{1-F(0)}{f(0)} f(y)\mathrm{d}y = \frac{1-F(0)}{f(0)},
\end{aligned}$$

where the first equality holds since $F(0)=1/2$.

market positions are not too asymmetric. The above results are summarized in the following proposition.

Proposition 5. *When two firms' market positions are asymmetric ($F(0)>1/2$), then consumers have incentives to form a GPO as long as their suppliers' market positions are not too asymmetric. If the GPO is formed, then each consumer is better off, the disadvantaged firm is worse off, and the allocative inefficiency present with single price competition is eliminated.*

Proposition 5 shows that nonlinear tariffs intensify competition under general distributions as long as two firms' market positions are not too asymmetric. This result can be understood in the following way. As firms' market positions become more asymmetric, nonlinear tariffs enable the advantaged firm to extract more consumer surplus. This surplus extraction effect can mitigate the competition intensifying effect, and may outweigh it when firms' market positions become very asymmetric. Proposition 5 also implies that a GPO is more likely to arise when two suppliers market positions become more symmetric.

4.4. Implications and discussion

Our analysis generates clear implications for antitrust policy. If nonlinear tariff offers to GPOs were to be banned, so that each supplier was restricted to a single price, firms would be better off and consumers worse off, with total social surplus decreased. Thus, nonlinear tariffs should be allowed.[23] This is the opposite of the conclusion for nonlinear tariffs under monopoly, where such tariffs increase efficiency at the expense of reducing consumers' surplus. Our conclusion is also the reverse of that obtained by Armstrong and Vickers (2006) for second-degree price discrimination with vertical differentiation.

In our model, the feasibility of nonlinear tariffs is controlled by consumers. When nonlinear tariffs intensify competition, consumers make them available by forming a GPO. If nonlinear tariffs hurt consumers, they avoid forming a GPO to make such tariffs infeasible. GPOs, by obtaining lower prices from suppliers and limiting their ability to extract surplus, have the side effect of rendering entry by new rival suppliers more difficult. To see this, consider the simple model of uniform distribution. Suppose initially only firm 1 is present in the market. Firm 2, then, is a potential entrant with a irreversible entry cost that we denote by f. When only firm 1 is present, there is no benefit for customers to form a GPO. Should firm 2 enter, however, consumers will form a GPO, with the result that each firm gets a lower profit $k/4$ compared to $k/2$ in the absence of the GPO. Therefore, if $f\in(k/4, k/2]$, entry occurs without GPO, but firm 2 is deterred given that consumers will form a GPO on the event of entry.

4.5. GPOs, promotion, and innovation

One common complaint about GPOs is that they stifle innovation by disadvantaging new products.[24] Prospective medical device manufacturers claim to be at a disadvantage because of the stiff rebates that established rivals' offer. Some theoretical models also have the feature that innovation is suppressed by buyer power. In a model featuring bargaining between oligopolistic manufacturers and buyers, Inderst and Wey (2003) demonstrate that sellers' incentives to innovate can increase when buyers in the downstream market combine. They generalize their result in a subsequent paper (Inderst and Wey, 2005). Our model generates a similar result, but our result does not stem from changes in bargaining power. In our model, though suppliers face heightened competition due to nonlinear tariffs, these tariffs also increase incentives to promote and innovate.

We demonstrate this effect by endogenizing consumers' willingness to pay for the products of firm 1 and firm 2, denoted as v_1 and v_2 respectively. Suppose the basic value of each good to a consumer is v, but each firm i can increase v_i by $\Delta v_i=v_i-v$ if it promotes its product at a cost $C(\Delta v_i)$, where $C(\cdot)$ is increasing and strictly convex (enough). Suppose firms make their promotional decisions (choose Δv_i) before they set prices. For simplicity, we assume that y is distributed symmetrically around 0, thus $\underline{y}=-\overline{y}$ and $F(0)=1/2$. We compare two settings: without and with a GPO.

In the absence of a GPO, we first derive the equilibrium prices given Δv_1 and Δv_2. Given p_1 and p_2, the marginal consumer is characterized by

$$y^* = p_2 - p_1 + (\Delta v_1 - \Delta v_2).$$

[23] Our model assumes that a GPO maximizes the welfare of its members. Some of the controversy over GPOs arises from concerns that payments by suppliers to GPOs are not transmitted to GPO members. This possibility is outside of the confines of our model.

[24] See, e.g., the Medical Device Manufacturers Association discussion of GPOs, http://www.medicaldevices.org/public/issues/gpo.asp.

And firms' profits are

$$\pi_1 = (p_1 - c)F(y^*); \pi_2 = (p_2 - c)[1 - F(y^*)].$$

The first order conditions give us the following equilibrium characterization

$$p_1^* = c + \frac{F(y^*)}{f(y^*)}; p_2^* = c + \frac{1 - F(y^*)}{f(y^*)}; \tag{11}$$
$$y^* = \frac{1 - F(y^*)}{f(y^*)} - \frac{F(y^*)}{f(y^*)} + (\Delta v_1 - \Delta v_2).$$

Diferentiating Eq. (11) with respect to Δv_1 yields

$$\frac{\partial y^*}{\partial \Delta v_1} = \frac{1}{3 + \frac{[1-2F(y^*)]f'(y^*)}{f^2(y^*)}}.$$

Now consider firm 1's incentive to increase Δv_1 (given Δv_2) in the first stage. We assume that the objective function $\pi_1(\Delta v_1, \Delta v_2) - C(\Delta v_1)$ is concave in Δv_1 ($C(\Delta v_1)$ is sufficiently convex). Then the following first order condition is sufficient:

$$C'(\Delta v_1) = \frac{\partial \pi_1}{\partial \Delta v_1} = \frac{\mathrm{d}}{\mathrm{d}y^*}\left[\frac{F^2(y^*)}{f(y^*)}\right]\frac{\partial y^*}{\partial \Delta v_1}$$
$$= \frac{2f(y^*)F(y^*) - F^2(y^*)f'(y^*)}{3f^2(y^*) + [1 - 2F(y^*)]f'(y^*)} \tag{12}$$

Similarly, we can derive the first order condition for Δv_2. Given that two firms are symmetric, in equilibrium we have $\Delta v_1^* = \Delta v_2^* \equiv \Delta v^*$, and $y^* = 0$. Thus by Eq. (12) the equilibrium Δv^* is characterized by

$$C'(\Delta v^*) = \frac{2f(0)F(0) - F^2(0)f'(0)}{3f^2(0)}. \tag{13}$$

Next consider the case with a GPO and firms compete in nonlinear tariffs. Given Δv_1 and Δv_2, the marginal consumer in equilibrium is $y^* = \Delta v_1 - \Delta v_2$. And $\pi_1(\Delta v_1, \Delta v_2)$ becomes

$$\pi_1(\Delta v_1, \Delta v_2) = \int_{\underline{y}}^{\Delta v_1 - \Delta v_2} [-y + \Delta v_1 - \Delta v_2] f(y)\mathrm{d}y$$
$$= (\Delta v_1 - \Delta v_2)F(\Delta v_1 - \Delta v_2) + \int_{\underline{y}}^{\Delta v_1 - \Delta v_2}$$
$$-yf(y)\mathrm{d}y = \int_{\underline{y}}^{\Delta v_1 - \Delta v_2} F(y)\mathrm{d}y. \tag{14}$$

In deriving the last equality we used integral by parts. Now by Eq. (14) the first order condition characterizing Δv_1 is

$$C'(\Delta v_1) = \frac{\partial \pi_1}{\partial \Delta v_1} = F(\Delta v_1 - \Delta v_2) \tag{15}$$

Similarly, we can derive the first order condition for Δv_2. In symmetric equilibrium, $\Delta v_1^{\mathrm{GPO}} = \Delta v_2^{\mathrm{GPO}} \equiv \Delta v^{\mathrm{GPO}}$ and $y^* = 0$. Thus by Eq. (15) the equilibrium Δv^{GPO} is characterized by

$$C'(\Delta v^{\mathrm{GPO}}) = F(0) = \frac{1}{2}. \tag{16}$$

Compare the right hand sides of Eqs. (13) and (16),

$$F(0) > \frac{2f(0)F(0) - F^2(0)f'(0)}{3f^2(0)} \tag{17}$$
$$\Leftrightarrow [f^2(0) + (1 - F(0))f'(0)]F(0)$$
$$+2f(0)F(0)[f(0) - 1] > 0.$$

By the monotonic hazard ratio property (5), $f^2(0) + (1 - F(0))f'(0) > 0$. Thus one of the sufficient conditions for Eq. (17) to hold is that $f(0) \geq 1$. Note that uniform distribution satisfies Eq. (17). Given that Eq. (17) holds, then $\Delta v^{\mathrm{GPO}} > \Delta v^*$, since $C(\cdot)$ is increasing and convex. Therefore, as long as $f(0) \geq 1$ the introduction of a GPO leads to a higher level of demand-increasing activity, whether promotion or product improvement. The GPO increases the return on investment in demand-enhancing activities, since the exploitation of any preference advantage that results is not hindered by any distortion in prices.

4.6. Multiple GPOs

As Section 2 shows, the healthcare market is stocked with multiple GPOs whose membership can overlap. We have assumed to this point that any GPO that arises must represent all consumers. We now relax this assumption.

First consider *proportional* GPOs, GPOs consisting of a positive measure of consumers and with distribution function $F(y)$ identical to that of consumers as a whole. Proportional GPOs can be of different sizes. Note that proportional GPOs obtain the same outcome as the universal GPO. This is because for GPOs, only the composition of consumer preferences and willingness to substitute among suppliers matters. In particular, tariff schedules can be tied to market share targets.

Alternatively, consumers could choose to affiliate with others holding similar preferences, so that proportionality would be lost. Suppose customers with y close to 0 form a GPO of their own, which we term a marginal GPO. Assume that such consumers are able to identify one another and to convince suppliers that their preferences are indeed concentrated around zero. A

credible marginal GPO can obtain average prices very close to the suppliers' marginal cost, c. Thus, if feasible, consumers with y close to zero have an incentive to form a *marginal* GPO. However, informational difficulties are likely to limit severely the role of marginal GPOs. Any GPO claiming to represent marginal consumers is likely to attract consumers who misrepresent their preferences in order to join. In addition, suppliers will find it difficult to confirm claims about GPO composition, since every GPO has an incentive to claim that it is a marginal GPO in order to induce lower prices. In the presence of such informational issues, it is likely that suppliers will treat GPOs equally.[25] The resulting equal treatment in turn makes consumers indifferent among joining different GPOs and form GPOs randomly, which justifies suppliers' belief that each GPO is proportional.

Ideally, it would be desirable to model the first stage as a coalition-formation game and to characterize the coalition proof Nash equilibria. However, the specific context and incomplete information makes the problem intractable.[26] Two layers of information asymmetry complicate the analysis. First, given that each consumer has a unit demand and consumers are not able to identify each other's preferences, it is impossible for a GPO to induce its members to report their preferences truthfully. This means that transfers among GPO members are infeasible. It also implies that marginal consumers subsidize consumers with more extreme preferences in any GPO. Second, the inability of firms to identify individual preferences makes it difficult for them to identify the composition of a specific GPO, as each GPO wants to claim that it is a marginal GPO. Instead, firms should hold beliefs about each GPOs composition, which are hard to pin down under incomplete information. For these reasons, we restrict our attention to universal GPO or proportional GPOs, and leave the coalition proof GPOs under incomplete information for future research.

5. Summary and conclusions

GPOs have captured a large and increasing share of healthcare supply purchasing. They have attracted health-care providers as members by obtaining lower average prices from suppliers. The price advantage that they generate for their members is not merely the consequence of buying power arising from the size of aggregated purchasing groups. Ellison and Snyder (2001) provide evidence that "price discounts depend more on the ability to substitute among alternative suppliers than on sheer buyer size. In particular, hospitals and HMOs, which can use restrictive formularies to enhance their substitution opportunities beyond those available for drugstores, obtain substantially lower prices. Chain drugstores only receive a small size discount relative to independents, at most 2% on average, and then only for products for which drugstores have some substitution opportunities (i.e., not for on-patent branded drugs). Our model demonstrates why this is so. We show that the same nonlinear tariffs that allow monopolists to extract all available surplus have a very different impact under oligopoly. By enabling rival suppliers to compete in nonlinear tariffs, GPOs generate efficient product allocations that nonetheless leave substantial surplus in the pockets of consumers.[27]

When suppliers' market positions are not too asymmetric, nonlinear tariffs increase the degree of competition very substantially in comparison to the already competitive alternative of Bertrand–Nash pricing in markets populated by heterogeneous consumers. Nonlinear pricing places all—not just marginal—consumers in play. By inducing tariff schedules that squeeze surplus out of suppliers, GPOs reduce the profits available to potential entrants. They also increase the returns to promotion and innovation that strengthen consumer preferences for a supplier's product, an effect that encourages investment in promotion and innovation. But the effects on entry and market dominance come in the context of greatly enhanced competition that GPOs and their associated nonlinear tariffs facilitate. Accordingly, any attempt to limit use of such discount schedules is likely to impose substantial costs on consumers together with inefficiency in consumption.

References

Anton, J.J., Yao, D.A., 1989. Split awards, procurement, and innovation. RAND Journal of Economics 20, 538–552.

Armstrong, M., Vickers, J., 2001. Competitive price discrimination. Rand Journal of Economics 32, 579–605.

Armstrong, M., Vickers, J., 2006. Competitive Nonlinear Pricing and Bundling. University College London. November.

Che, Y.-K., Gale, I., 1997. Buyer alliances and managed competition. Journal of Economics & Management Strategy 6, 175–200.

[25] Note that many healthcare providers can and do join multiple GPOs, further homogenizing them in appearance to suppliers.

[26] Dana (2006) studies coalition proof buyer groups under complete information.

[27] Evidence on GPO behavior suggests that commitment programs that tie prices to share levels are used most intensively by the largest and most successful of GPOs. See United States General Accounting Office, "Group Purchasing Organization: Use of Contracting Processes and Strategies to Award Contracts for Medical–Surgical Products", Study GAO-03-ss8T, July 16, 2003.

Chipty, T., Snyder, C.M., 1999. The role of firm size in bilateral bargaining: a study of the cable television industry. Review of Economics and Statistics 81, 326–340.

Corts, K., 1998. Third-degree price discrimination in oligopoly: all-out competition and strategic commitment. RAND Journal of Economics 29, 306–323.

Dana, J., 2006. Buyer Groups as Strategic Commitments. Northwestern University. December.

Ellison, S.F., Snyder, C.M., 2001. Countervailing power in wholesale pharmaceuticals. MIT Dept. of Economics Working Paper No. 01-27.

Greenlee, P., Reitman, D., 2005. Competing with loyalty discounts. United States Department of Justice EAG Discussion Paper 04-2 (revised).

Holmes, T.J., 1989. The effects of third-degree price discrimination in oligopoly. American Economic Review 79, 244–250.

Inderst, R., Wey, C., 2003. Bargaining, mergers, and technology choice in bilaterally oligopolistic industries. RAND Journal of Economics 34, 1–19.

Inderst, R., Wey, C., 2005. How strong buyers spur upstream innovation. DIW Berlin (German Institute for Economic Research) Discussion Paper.

Inderst, R., Wey, C., 2007. Buyer power and supplier incentives. European Economic Review 51, 647–667.

Kolay, S., Shaffer, G., Ordover, J.A., 2004. All-units discounts in retail contracts. Journal of Economics & Management Strategy 13, 429–459.

Marvel, H.P., Peck, J., 2006. Inventory Turnover and Product Variety. Ohio State University.

Mathewson, F., Winter, R.A., 1996. Buyer groups. International Journal of Industrial Organization 15, 137–164.

O'Brien, D.P., Shaffer, G., 1997. Nonlinear supply contracts, exclusive dealing, and equilibrium market foreclosure. Journal of Economics & Management Strategy 6, 755–785.

Rochet, J.-C., Stole, L.A., 2002. Nonlinear pricing with random participation. Review of Economic Studies 69, 277–311.

Snyder, C.M., 1996. A dynamic theory of countervailing power. RAND Journal of Economics 27, 747–769.

Snyder, C.M., 1998. Why do larger buyers pay lower prices? Intense supplier competition. Economics Letters 58, 205–209.

Spulber, D.F., 1981. Spatial nonlinear pricing. American Economic Review 71, 923–933.

Stole, L., 1995. Nonlinear pricing and oligopoly. Journal of Economics & Management Strategy 4, 529–562.

Thanassoulis, J., 2007. Competitive mixed bundling and consumer surplus. Journal of Economics & Management Strategy 16, 437–467.

U.S. Department of Justice and Federal Trade Commission, 2004. Improving health care: a dose of competition. Technical Report. URL http://www.ftc.gov/reports/healthcare/040723healthcarerpt.pdf.

Varian, H., 1989. Price discrimination. In: Schmalensee, R., Willig, R.D. (Eds.), Handbook of Industrial Organization, vol. 1. North-Holland, Amsterdam, pp. 597–654. chapter 10.

Yang, H. Ye, L., 2007. Nonlinear Pricing, Contract Veriety, and Competition. The Ohio State University.