OPERATIONS RESEARCH
Vol. 52, No. 6, November–December 2004, pp. 868–886
ISSN 0030-364X | EISSN 1526-5463 | 04 | 5206 | 0868
informs®
DOI 10.1287/opre.1040.0149

A General Equilibrium Model for Industries with Price and Service Competition

Fernando Bernstein
The Fuqua School of Business, Duke University, Durham, North Carolina 27708, fernando@mail.duke.edu
Awi Federgruen
Graduate School of Business, Columbia University, New York, New York 10027, af7@columbia.edu

This paper develops a stochastic general equilibrium inventory model for an oligopoly, in which all inventory constraint parameters are endogenously determined. We propose several systems of demand processes whose distributions are functions of all retailers' prices and all retailers' service levels. We proceed with the investigation of the equilibrium behavior of infinite-horizon models for industries facing this type of generalized competition, under demand uncertainty.

We systematically consider the following three competition scenarios. (1) Price competition only: Here, we assume that the firms' service levels are exogenously chosen, but characterize how the price and inventory strategy equilibrium vary with the chosen service levels. (2) Simultaneous price and service-level competition: Here, each of the firms simultaneously chooses a service level and a combined price and inventory strategy. (3) Two-stage competition: The firms make their competitive choices sequentially. In a first stage, all firms simultaneously choose a service level; in a second stage, the firms simultaneously choose a combined pricing and inventory strategy with full knowledge of the service levels selected by all competitors. We show that in all of the above settings a Nash equilibrium of infinite-horizon stationary strategies exists and that it is of a simple structure, provided a Nash equilibrium exists in a so-called reduced game.

We pay particular attention to the question of whether a firm can choose its service level on the basis of its own (input) characteristics (i.e., its cost parameters and demand function) only. We also investigate under which of the demand models a firm, under simultaneous competition, responds to a change in the exogenously specified characteristics of the various competitors by either: (i) adjusting its service level and price in the same direction, thereby compensating for price increases (decreases) by offering improved (inferior) service, or (ii) adjusting them in opposite directions, thereby simultaneously offering better or worse prices and service.

Subject classifications: inventory/production policies: marketing/pricing; games/group decisions: noncooperative.
Area of review: Manufacturing, Service, and Supply Chain Operations.
History: Received July 2001; revisions received June 2002, May 2003; accepted November 2003.

1. Introduction

We develop a stochastic general equilibrium inventory model for an oligopoly in which firms compete in terms of two strategic instruments, i.e., (i) their prices (or expected sales targets), and (ii) their service-level targets, in particular, their fill rates. The fill rate is the fraction of demand that can be met from existing inventory.[1] Each firm complements its choice of a service-level target and pricing strategy with an appropriate inventory strategy.

To build the general equilibrium model, we propose several demand models, where the stochastic demands are functions of all retailers' prices and all retailers' service levels. In a periodic review, infinite-horizon setting, the retailers face a stream of demands that are independent across time, but not necessarily across the retailers. End-of-the-period inventories are carried over to the next period. Stockouts are backlogged.[2] Customer sales are final, i.e., they cannot be cancelled when the customer needs to wait for delivery. Each retailer may, at the beginning of each period, place an order with his supplier, who fills the orders instantaneously or after a given lead time. The retailers incur retailer-specific inventory carrying costs. We consider both the case where direct (retailer-specific), out-of-pocket, backlogging costs prevail and the common case without such direct stockout costs. Our model applies both to retailers selling to the final consumer, as well as to vendors selling to retailers. (In describing and analyzing the model, we confine ourselves to the former setting.)

We observe an increasing number of industries in which some of the competing firms aggressively attempt to obtain larger market shares by providing higher levels of service. For example, in the fierce competition between amazon.com and barnesandnoble.com, the latter initiated a massive advertising campaign promising same business-day delivery in various parts of the country. The same firms are also examples of companies routinely posting expected waiting times by item or group of items. In the video rental industry, Blockbuster launched an advertising campaign in 1997, emphasizing high fill rates with its "Go Home Happy" slogan, and backing their promise up with a free rental guarantee under the slogan "I'll Be There" (see, e.g., Dana 2001).[3] Similarly, Domino's has offered a 30-minute delivery guarantee for its pizza sales, backed up with a free-of-charge delivery if the time limit is exceeded. In B2B settings, it is well known that vendors routinely specify allowable windows for order delivery times and they, too, often back these up with chargebacks in case their service targets are violated.[4]

Even if most retailers do not yet advertise their service levels, "customer service level as measured by % of time product is in stock" is now recognized by retail executives as one of the most important performance metrics, ahead of traditional key measures such as "sales per selling square foot."[5] An industry of "secret shoppers" has emerged to provide retail management with independent estimates of various customer service levels, while various Internet firms continuously gather and report information on online retailers.[6]

Ever since the seminal papers by Arrow et al. (1951) and Dvoretzky et al. (1952), a massive literature has addressed inventory systems under uncertainty. These papers provide a systematic trade-off for two competing risks, i.e., the risk of overstockage and that of understockage. It is relatively easy to assess the consequences of the former because the costs of carrying inventories can be measured or estimated in a fairly straightforward manner. At the same time, the consequences of shortages are much harder to quantify. Most standard inventory models assume that when shortages can be backlogged, the cost associated with a given backlog size can be assessed as an exogenous input to the model. Much has been written about the difficulty of specifying backlogging cost rates. Often, no out-of-pocket expense is associated with a stockout, and even if such out-of-pocket penalties prevail, it is generally agreed upon that backlogging cost rates or stockout penalties should, in addition, reflect the long-term or equilibrium impact on goodwill and market shares. The latter is hard to quantify, at least in the absence of a complete model describing the impact of the relative service levels offered by competing firms on each of their demand processes. As a consequence, many practitioners are more comfortable with a model in which (production, distribution, and carrying) costs are minimized subject to given service-level constraints, e.g., fill rates satisfying given minimum fill-rate bounds. Alternatively, an upper bound on the expected value or a given fractile of the customer waiting time may be imposed. However, here too, it is unclear whether a company should strive for a 90%, 95%, or 99% fill rate or promise a 24-hour, two-business-day, or one-week response time, say, and what the long-term revenue implications of this choice may be. The need to endogenize fill rate and stockout cost parameters via a model with competing firms was first articulated by Li (1992). An additional complication is that virtually all inventory models assume that a company's demands are not affected by its service levels, let alone by those of its competitors.

We systematically characterize the equilibrium behavior of the industry under three possible scenarios. (1) *price competition only*: Here, we assume that the firms' service levels are *exogenously* chosen, but characterize how the price equilibrium and inventory strategy vary with the chosen service levels. (2) *Simultaneous price and service-level competition*: Here, each of the firms simultaneously chooses a service level and a combined price and inventory strategy. (3) *Two-stage competition*: The firms make their competitive choices sequentially. In a first stage, all firms simultaneously choose a service level; in a second stage, the firms simultaneously choose a combined pricing and inventory strategy with full knowledge of the service levels selected by all competitors. Fudenberg and Tirole (1991) refer to the equilibria arising under (2) as *open-loop equilibria* and those under (3) as *closed-loop equilibria*. Our first principal result is to show that in all of the above settings a Nash equilibrium of infinite-horizon stationary strategies exists and that it is of a simple structure, provided a Nash equilibrium exists in a so-called reduced game. The reduced game is a single-stage, single-period game under scenarios (1) and (2) and a two-stage, single-period game under (3). More specifically, under scenario (1) each firm adopts a stationary price and a stationary base-stock policy to control its inventory, i.e., the firm increases its inventory to a given base-stock level whenever it is below this level and avoids placing replenishment orders otherwise. Under Scenarios (2) and (3), the equilibrium infinite-horizon strategies are of the same type, except that under (2) each firm selects a long-run service level along with the other choices and, under (3), in advance of selecting a price and base-stock policy combination. The critical control parameters for the equilibrium strategies in the infinite-horizon model are easily derived from the equilibrium of the reduced game. The equivalency results between the infinite horizon and the reduced games are obtained for fully general stochastic demand functions of the price and service vectors.[7]

The remainder of this paper is, therefore, devoted to the characterization of the equilibrium behaviors in the reduced games; these depend critically on the type of stochastic demand functions. We consider three of the most frequently used classes of demand functions in the marketing and industrial organization literature, appropriately adjusted for their dependence on the vector of service levels. The attraction models represent the first such class: with a fixed total potential market size, each firm acquires a market share that is proportional to an attraction value given by a (possibly firm-dependent) general function of its price and service level. Bell et al. (1975) have shown that this class of models, with the (generalized) MultiNomial Logit (MNL) models as a special case, is the unique class to satisfy a set of rather intuitive axioms. Linear demand functions constitute our second class of demand models. In a third and last class of so-called log-separable demand functions, each firm's demand is given by a general function of the price vector, multiplied with a scaling factor that is given by a general function of the industry's vector of service levels.

(The scaling factor equals one when the firms do not differentiate themselves on the basis of their service levels, i.e., when these levels are identical.)

We particularly focus on the question of whether a firm can choose its service level on the basis of its own (input) characteristics (i.e., its cost parameters and demand function) only, as is assumed in virtually all standard inventory models, including those that address combined pricing and inventory decisions. This implies, in particular, that a firm's service level can be chosen independently of its own equilibrium price or its competitors' equilibrium prices and service levels. This extreme form of robustness holds for the (generalized) MNL model: Here, each firm adopts the same equilibrium service level, whether engaged in simultaneous or two-stage competition (Scenarios (2) and (3), respectively) and this level only depends on the firm's own input characteristics. Under linear demand functions, each firm's equilibrium service level under simultaneous competition again depends only on its own input characteristics; however, in general, a different service level should be chosen under sequential competition, and this service level may well depend on the choices and characteristics of the competitors. A more limited robustness result applies to the general attraction model. (Recall that the MNL model is a special case of this class of models.) Here, each firm's equilibrium service level under simultaneous competition depends only on its own price and input characteristics; under two-stage competition, even this more limited robustness result breaks down. None of the above robustness results apply to the log-separable demand functions.

We also investigate under which of the demand models a firm, under simultaneous competition, responds to a change in the exogenously specified characteristics of the various competitors by either: (i) adjusting its service level and price in the same direction, thereby compensating for price increases (decreases) by offering improved (inferior) service, or (ii) adjusting them in opposite directions, thereby simultaneously offering better or worse prices and service. As stated, no adjustment in the service level is needed under the linear and MNL models. In the general attraction model with attraction functions that are separable in the firm's price and service level (or where the increase in the firm's attraction value due to an improvement of its service level is smaller when it charges a higher price), price and service levels always move in opposite directions. Within the general attraction model, adjustments in the same direction may arise only when the attraction functions are supermodular, i.e., the increase in a firm's attraction value due to an improvement of its service level is higher when the firm charges a higher price. Likewise, service and price levels may move in the same direction under the log-separable demand model.

The fact that many of the qualitative properties of the market equilibria (as well as their dynamics) depend on the type of demand model has important implications for empirical studies that attempt to estimate the system of stochastic demand functions. (See Leeflang et al. 2000, Basuroy and Nguyen 1998, and Besanko et al. 1998 for examples of such demand estimation studies under competitive pricing.) It suggests that a variety of demand models should be considered, instead of an upfront restriction to a particular type of model. Moreover, if in a given industry it is known, for example, that firms adjust their service levels in response to changes made by some of their competitors, this is a fundamental reason to consider demand model specifications that permit such responsive behavior.

The remainder of this paper is organized as follows: In §2, we review the relevant literature. In §3, we introduce our general model and notation. In §4, we show in full generality that under Scenarios (1) and (2) a Nash equilibrium exists of infinite-horizon stationary strategies under which each retailer adopts a stationary price, fill rate, and base-stock policy, provided a Nash equilibrium exists in a reduced single-stage game. Under (3), only the second stage is an infinite-horizon game. This scenario can, similarly, be reduced to a two-stage game with the second stage a "single-period" game. Sections 5–7 thus focus on the reduced single- or two-stage games and systematically characterize the equilibrium behavior under the three Scenarios (1)–(3), each section for one of the above three classes of demand functions. Section 8 summarizes our main results. All proofs are deferred to the appendix.

2. Literature Review

It appears that Schwartz (1966) was the first to reject the usual assumption in inventory models of a fixed penalty for stockouts. Schwartz (1970) explicitly models, in a single-location setting, the impact of stockouts on future sales: demand in a given period equals the potential demand multiplied with a factor that depends on the observed fill rate and the rate at which customers forgive the firms for the stockouts. See Hill (1976) for a generalization of this model. Caine and Plaut (1976) consider a periodic-review model with demands whose means are a function of the fill rates. Robinson (1991) provides a further generalization where the mean and variance of each period's demand varies linearly with the number of satisfied customers in the previous period. Ernst and Cohen (1992) and Ernst and Powell (1995) consider a single-supplier single-retailer system in which the demands faced by the retailers have a mean and standard deviation that depend on the steady-state service level. Ernst and Powell (1998) model this system as a Stackelberg game, with the supplier as the leader.

Several economists, starting with Dorfman and Steiner (1954), have realized that representing demands as a function only of the sales price(s) may oversimplify customer preferences. In the context of deterministic monopoly models, these and other authors (e.g., Spence 1975, 1976; Dixit 1979; Tirole 1988, §2.2) have considered a demand function with an additional attribute variable, referred to as "advertising," "service level," or "quality." Carlton (1989) and Carlton and Perloff (1999) argue that demand functions should be specified as a function of prices and customer service levels, which they quantify by the customer's expected waiting time. See Table 15.3 in Carlton (1989) for estimates of price and "delivery-lag" elasticities in a variety of industries. Karmarkar and Pitbladdo (1997) consider a multidimensional quality measure in the demand functions. They address perfect competition and oligopolistic markets, where the oligopoly equilibrium is specified by an entry and exit condition under which the profits of all firms are driven to zero.

In contrast, Banker et al. (1998) and Tsay and Agrawal (2000) characterize the equilibrium behavior of oligopolies with a fixed number of firms competing simultaneously with their price and a "quality" or service instrument. Both papers consider a single-period model for duopolies, with deterministic demand functions that are linear in all price and quality variables. The former (latter) characterize the equilibrium under sequential (simultaneous) competition. Both assume that each retailer's cost increases quadratically with the service or quality level provided. This exogenously specified quadratic cost function is substantiated by the fact that "under the assumptions of standard inventory models, moving from, say, 97% to 99% fill rate typically requires a greater incremental investment than does moving from 95% to 97%" (see Footnote 3 in Tsay and Agrawal 2000, p. 375). In our infinite-horizon model with an arbitrary number of competing firms, stochastic demand functions of a general (nonlinear) type and inventory carrying expenses, the cost/service relationship is endogenously determined; costs indeed increase convexly with the chosen service level, but the dependency fails to be quadratic. Anderson et al. (1992) consider a two-stage game with deterministic demands for retailers competing with a price and quality variable, similar to Scenario (3)'s reduced two-stage game. The authors restrict themselves to the case of identical retailers and a special case of the demand functions treated in §5.1.

Li (1992) appears to be the first to model horizontal competition between firms facing uncertain demands (and) or supplies. The firms offer an identical product and charge equal prices, but compete in terms of their production/inventory strategies. Customers arrive according to a Poisson process and purchase with equal probability from any firm with positive inventory. If all firms are out of stock, the customer places the order with each firm, but buys from the firm that completes the order first, a practice common in the semiconductor industry. Van Mieghem and Dada (1999) develop a single-period model for firms producing a homogeneous good. Similar to our Scenario (3), the firms select a capacity level in a first stage and actual production quantities in a second stage; the common price is a linear function of the total quantity offered in the market, with a random intercept. Carr et al. (1999) develop a single-period model for firms competing in terms of their prices and facing linear stochastic demand functions. Bernstein and Federgruen (2001), similarly, develop a single-period model for firms facing price competition under a general class of stochastic demand functions. Kirman and Sobel (1974) and Bernstein and Federgruen (2004) appear to be the first infinite-horizon models for an industry with firms competing in terms of their prices only. We refer to Bernstein and Federgruen (2001) for a review of several (single-period) models in which retailers compete via their starting inventories. To our knowledge, this is the first paper to address an oligopoly under uncertain demands, with price and service-level competition.

3. Model and Notation

Consider an industry with N independent retailers facing random demands. We analyze a periodic-review infinite-horizon model in which each Retailer i positions himself on the market by selecting a steady-state fill rate f_i, where the fill rate is defined as the fraction of customer demands satisfied from on-hand inventory, as well as a retail price p_i. Without loss of practical generality, we restrict ourselves to service levels in the interval $[0.5, 1)$.[8] While the fill-rate target levels f are stationary choices, the retail price may, in principle, be varied in each period. Each retailer may, at the beginning of each period, place an order with his supplier, assumed to have ample capacity to fill any size order completely and in time for the retailer to meet this period's demand. Stockouts are backlogged. Thus, each Retailer i makes the following choices: (i) a one-time choice of f_i, at the beginning of the planning horizon; (ii) at the beginning of each period t, a retail price as well as the quantity to be ordered from the supplier.

One of the novel features of our model is that the demand faced by each Retailer i, in any period t, has a distribution that may depend on the entire vector of retail prices p in that period as well as the entire vector of fill-rate target levels f. Thus, let

$D_{it}(p, f) =$ the random demand faced by Retailer i in period t, under the retail price vector p and the service-level vector f,

where the cdf of D_{it}, denoted by $\tilde{G}_i(x|p, f)$, depends on the entire vector p as well as the complete vector of service levels f. Thus, demands in any period depend on the target fill rates, not on the actual inventory levels. We assume that the demand variables are of the multiplicative form, i.e.,

$$D_{it} = d_i(p, f)\epsilon_{it}, \tag{1}$$

with ϵ_{it} a general continuous random variable whose distribution is stationary and independent of the retail price vector p and the service-level vector f, i.e., for all $i = 1, \ldots, N$, the sequence $\{\epsilon_{it}\}$ has a common general cdf $G_i(\cdot)$ and density function $g_i(\cdot)$ such that $\tilde{G}_i(x|p, f) = G_i(x/d_i(p, f))$.

The multiplicative model implies that the coefficients of variation of the one-period demands are exogenous constants, independent of the price vector or the service levels, and hence of the expected sales volumes as well. Without loss of generality, we normalize $E(\epsilon_{it}) = 1, i = 1, \ldots, N$ and $t = 1, 2, \ldots$, so that $ED_{it}(p, f) = d_{it}(p, f)$. In other words, the functions $\{d_i(p, f)\}$ may be viewed as representing the expected one-period sales volumes. As in virtually all inventory models, we assume that the sequence of random variables $\{\epsilon_{it}: t = 1, 2, \ldots\}$, and hence the sequence $\{D_{it}: t = 1, 2, \ldots\}$, are independent for all $i = 1, \ldots, N$. At the same time, we allow for arbitrary correlations between the demands faced by the different retailers in any given period.

Information about the firms' service levels is not always as readily available as the unit price. As mentioned, in the B2B world, service-level guarantees are routinely provided by the vendors, often backed up with chargeback agreements for violations of these guarantees. Software systems allow retailers to monitor their vendors' compliance and provide them with comparative data regarding groups of vendors' service levels. We also mentioned several examples in the B2C world, where firms advertise service-level measures, as well as independent Internet services that rate online retailers in terms of their fill-rate performance (along with other service measures). Even when such information is not publicly available, consumers develop estimates on the basis of their own (repeat-purchase) experience as well as on the basis of word of mouth and other reputational information (see Tirole 1988, §2.3, for a general discussion of the consumers' ability to assess quality attributes of competing products and its implications for industrial organization models).

At this stage, we make minimal assumptions regarding the shape of the mean sales functions $\{d_i(p, f)\}$, other than the following basic monotonicity properties:

$$\begin{aligned} &\frac{\partial d_i(p,f)}{\partial p_i} \leqslant 0, \quad \frac{\partial d_i(p,f)}{\partial f_i} \geqslant 0; \\ &\frac{\partial d_i(p,f)}{\partial p_j} \geqslant 0, \quad \frac{\partial d_i(p,f)}{\partial f_j} \leqslant 0, \quad j \neq i. \end{aligned} \tag{2}$$

In other words, if a retailer increases his retail price (fill rate), this results in a decrease (increase) of his own expected sales while increasing (decreasing) those of his competitors.

No firm's sales are expected to increase under a uniform price increase:[9]

$$\text{(D)} \quad \sum_{j=1}^{N} \frac{\partial d_i}{\partial p_j} < 0 \quad \text{for all } i = 1, \ldots, N.$$

Decisions are made in the following sequence: At the beginning of each period, all retailers simultaneously determine their price and order quantity for that period; next, these orders are filled.

Each retailer pays the supplier a constant per-unit wholesale price, inclusive of delivery costs, or he incurs production costs at a constant rate. (We describe the remainder of the paper in terms of retailers purchasing their goods from outside suppliers.) Holding costs are proportional with end-of-period inventories. A retailer may incur direct, out-of-pocket, backlogging costs; if so, these are proportional with the backlog size. Thus, for each retailer $i = 1, \ldots, N$:

$w_i =$ the per-unit wholesale price paid by Retailer i,

$h_i^+ =$ the per-period holding cost for each unit carried in inventory, and

$h_i^- =$ the per-period direct backlogging cost for each unit backlogged at Retailer i.

Contrary to most standard inventory models, but more representative of actual cost/service trade-offs experienced in practice, our model does not require that direct backlogging costs exist, i.e., that $h_i^- > 0$. Even if $h_i^- = 0$, every firm is incentivized to carry appropriate safety stocks because a large backorder frequency or, equivalently, a low fill rate, reduces the retailer's average sales while increasing those of his competitors.

4. The Infinite-Horizon Model: Reduction to Single- or Two-Stage Games

In this section, we show that under each of the competitive scenarios, (1)–(3), the equilibrium behavior in the infinite-horizon model may be characterized by analyzing that of an equivalent single-stage or, in Scenario (3), two-stage game.

We start with the simultaneous price- and service-competition Scenario (2) and show that in the infinite-horizon retailer game, an N-tuple of stationary strategies arises as a Nash equilibrium, where each Retailer i adopts a fill rate f_i^*, a constant price p_i^*, and a base-stock policy with stationary base-stock level y_i^*. Moreover, the triplet (p^*, f^*, y^*) represents a Nash equilibrium in a single-stage game with the following profit function for Retailer i:

$$\begin{aligned} \Pi_i(p, f, y_i) = {} & (p_i - w_i) d_i(p, f) - h_i^+ E[y_i - D_i(p, f)]^+ \\ & - h_i^- E[D_i(p, f) - y_i]^+, \quad i = 1, \ldots, N, \end{aligned}$$

where Retailer i's action set is given by the following subset of $\Re^3$:

$$\begin{aligned} &\big\{(p_i, f_i, y_i)\colon p_i^{\min} \leqslant p_i \leqslant p_i^{\max}; 0 \leqslant f_i < 1, \\ &\Pr[D_i(p,f) \leqslant y_i] \geqslant f_i\big\} = \big\{(p_i, f_i, y_i)\colon p_i^{\min} \leqslant p_i \leqslant p_i^{\max}; \\ &\qquad 0 \leqslant f_i < 1; y_i \geqslant d_i(p,f) G_i^{-1}(f_i)\big\}, \end{aligned}$$

because

$$\Pr[D_i(p,f) \leqslant y_i] = \Pr\left[\epsilon_i \leqslant \frac{y_i}{d_i(p,f)}\right] = G_i\left(\frac{y_i}{d_i(p,f)}\right) \geqslant f_i$$

iff $y_i \geqslant d_i(p,f) G_i^{-1}(f_i)$.

While a retailer's price p and fill rate f have a potential impact on all retailers' profits, his choice of a base-stock level y affects his own profit function only. Thus, each Retailer i chooses y_i to maximize his profit, while ensuring a service level at least equal to the announced service level f_i. That is, for given p and f, Retailer i chooses an order-up-to level $y_i \geqslant d_i(p,f)G_i^{-1}(f_i)$, which maximizes $\Pi_i(\cdot,\cdot,y_i)$:

$$y_i(p,f) = d_i(p,f)G_i^{-1}\left(\max\left\{f_i, \frac{h_i^-}{h_i^- + h_i^+}\right\}\right). \tag{3}$$

This gives rise to an equivalent reduced single-stage retailer game in which each Retailer i competes with two instruments $(p_i, f_i) \in [p_i^{\min}, p_i^{\max}] \times [0,1)$ and reduced profit functions

$$\pi_i(p,f) = (p_i - w_i - k_i(f_i))d_i(p,f), \quad \text{where} \tag{4}$$

$$k_i(f_i) = h_i^+ E\left[G_i^{-1}\left(\max\left\{f_i, \frac{h_i^-}{h_i^- + h_i^+}\right\}\right) - \epsilon_i\right]^+ + h_i^- E\left[\epsilon_i - G_i^{-1}\left(\max\left\{f_i, \frac{h_i^-}{h_i^- + h_i^+}\right\}\right)\right]^+.$$

Note that $k_i(f)$ is the expected (end-of-period) inventory cost per unit of sales, required to guarantee a given service level of at least f. The reduced single-stage game is equivalent to an oligopoly model with deterministic demands $d(p,f)$ and cost functions $C_i(p,f) = k_i(f_i)d_i(p,f)$. It is the cost structure that sets the model apart from other oligopoly models with competition with two or more strategic instruments (see, e.g., Basuroy and Nguyen 1998, and the references therein). (The only exception is the price-quality two-stage model in Anderson et al. 1992, which considers identical retailers and a special case of the demand functions in §5.1.)

We now show that if a pair of vectors (p^*, f^*) is a Nash equilibrium of the (reduced) single-stage game, an N-tuple of infinite-horizon strategies can be constructed that is a Nash equilibrium in the original infinite-horizon game.

Theorem 1. *Consider Scenario* (2) (*simultaneous price and service competition*). *Assume that* (p^*, f^*) *is a Nash equilibrium in the reduced single-stage game. The N-tuple of infinite-horizon stationary strategies under which Retailer i adopts a stationary price p_i^*, a fill rate f_i^*, and a base-stock policy with* (*stationary*) *base-stock level $y_i(p^*, f^*)$ is a Nash equilibrium in the infinite-horizon game.*

We obtain similar reductions of the infinite-horizon model under Scenarios (1) and (3).

Corollary 1. (a) *Consider Scenario* (1) (*price competition only*), *i.e., assume that the vector of service levels $f = \hat{f}$ is exogenously given. Consider also the single-stage game in which each Retailer i competes with a single instrument $p_i \in [p_i^{\min}, p_i^{\max}]$ and reduced profit functions $\hat{\pi}_i(p) \stackrel{\text{def}}{=} \pi_i(p, \hat{f})$. Assume that $p^*(\hat{f})$ is a Nash equilibrium in this reduced single-stage game. The N-tuple of infinite-horizon stationary strategies under which Retailer i adopts a stationary price $p_i^*(\hat{f})$, a fill rate $\hat{f}_i$, and a base-stock policy with* (*stationary*) *base-stock level $y_i(p^*(\hat{f}), \hat{f})$ is a Nash equilibrium in the infinite-horizon game.*

(b) *Consider Scenario* (3) (*two-stage competition*) *and the following two-stage game*: *In the first stage, the retailers compete with a single instrument $f_i \in [0,1)$ and reduced profit functions $\hat{\pi}_i(f) \stackrel{\text{def}}{=} \pi_i(p^*(f), f)$, where $p^*(f)$ denotes the Nash equilibrium in the game considered in* (a). *Let $\hat{f}$ denote a Nash equilibrium of this first-stage game. In the second stage, the retailers face the game in* (a). *The N-tuple of infinite-horizon stationary strategies under which Retailer i adopts a stationary price $p_i^*(\hat{f})$, a fill rate $\hat{f}_i$, and a base-stock policy with base-stock level $y_i(p^*(\hat{f}), \hat{f})$ is a Nash equilibrium in the infinite-horizon game.*

In §§5–7, we exhibit several important classes of demand functions for which the reduced single-stage or two-stage games in Theorem 1 and Corollary 1 have an equilibrium so that Theorem 1 and Corollary 1 apply. Lemma 1 describes the shape of the $k_i(\cdot)$-functions.

Lemma 1. (a) *$k_i(\cdot)$ is increasing and differentiable.*

(b) *$k_i(\cdot)$ is convex, twice differentiable for all $f_i \neq h_i^-/(h_i^- + h_i^+)$, and $\lim_{f_i \to 1} k_i'(f_i) = \infty$ for all distributions G_i such that*:

$$(\mathrm{PF}_2) \quad \begin{array}{l} G_i \text{ is log-concave or, equivalently, is a} \\ \text{Polya Frequency function of order 2 for all} \\ x \geqslant G_i^{-1}(0.5), \text{ and } g_i \text{ has infinite support,} \end{array} \tag{5}$$

where $G_i^{-1}(\cdot)$ denotes the inverse of the G_i-distribution.

Condition (PF_2) in (5) is trivially satisfied for all distributions whose density function decreases beyond the median, e.g., the normal and exponential distributions and many specifications of the gamma and Weibull distributions. The condition is closely related to the more common condition that the density functions be log-concave or PF_2 (see, e.g., Barlow and Proschan 1965). By Lemma 5.8 and Theorem 5.6 in Barlow and Proschan, the latter condition implies that the complements of the cdfs G_i are PF_2. Lemma 1 is intuitive: The operational costs increase convexly with the service level.[10]

5. Attraction Models

Many marketing models characterize the market shares obtained by competing retailers via a vector of attraction values $a = (a_1, \ldots, a_N)$. The market share achieved by a given firm i is given by its attraction value divided by the industry's total value, i.e., by

$$m_i = \frac{a_i}{\sum_{j=0}^N a_j}, \tag{6}$$

with a_0 the value of the no-purchase option. Bell et al. (1975) show that if market shares are exclusively determined by the attraction vector a, they must be given by the simple ratios (6), provided four general assumptions are satisfied. Attraction models are among the most commonly used market share models, both in empirical studies and theoretical models (see, e.g., Leeflang et al. 2000). In standard marketing models, the attraction values a may depend on the price vector and/or other marketing instruments such as advertising efforts. Here, we will assume that a_0 is a constant and the attraction value of any firm i is given by a general, twice-differentiable function of its price and service level, i.e., $a_i = a_i(p_i, f_i)$, with

$$\frac{\partial a_i}{\partial p_i} \leqslant 0, \quad \frac{\partial a_i}{\partial f_i} > 0. \tag{7}$$

With fixed total market size M, this results in the following mean demand functions:

$$d_i(p, f) = M\frac{a_i(p_i, f_i)}{\sum_{j=0}^N a_j(p_j, f_j)}. \tag{8}$$

Most attraction models assume a specific structure for the functions $a_i(\cdot, \cdot)$, most commonly the "Multiplicative Competitive Interaction" structure which, when applied to the pair of instruments (p_i, f_i), is of the form $a_i(p_i, f_i) = c_i p_i^{-\alpha_i} f_i^{\beta_i}$ or the MNL model, where

$$a_i(p_i, f_i) = \exp\{\beta_i f_i - \alpha_i p_i\} \tag{9}$$

(for constants $\alpha_i, \beta_i, c_i > 0$). See Anderson et al. (1992) or Mahajan and van Ryzin (1999) for a discussion of how the MNL model arises from either a random utility model or a set of choice axioms, similar to, though somewhat more restrictive than, those of Bell et al. (1975).[11]

Below we pay special attention to the following generalization of (9):

$$a_i(p_i, f_i) = \exp\{b_i(f_i) - \alpha_i p_i\}, \tag{10}$$

where the functions $b_i(\cdot)$ are twice differentiable, increasing, and concave, permitting us to represent settings where the marginal increase in a firm's attraction value due to an increase in its fill rate is nonnegative, but decreasing in f_i. Alternatively, if a firm wants to maintain a given attraction value, it needs to compensate for a price increase with ever larger increases in its service level.

5.1. Analysis of the General Attraction Model

For the single-stage price game under a given vector of service levels f, we need the following lemma, which follows by simple calculus. Let $\tilde{a}_i = \log a_i$ and $\tilde{d}_i = \log d_i$.

LEMMA 2. *Assume that the mean sales volumes $\{d_i(p_i, f_i)\}$ are given by the attraction model* (8).

$$\frac{\partial d_i}{\partial p_i} = \frac{\partial \tilde{a}_i}{\partial p_i} d_i\left(1 - \frac{d_i}{M}\right), \quad \frac{\partial d_i}{\partial p_j} = -\frac{\partial \tilde{a}_j}{\partial p_j}\frac{d_i d_j}{M}, \quad j \neq i;$$

$$\frac{\partial d_i}{\partial f_i} = \frac{\partial \tilde{a}_i}{\partial f_i} d_i\left(1 - \frac{d_i}{M}\right), \quad \frac{\partial d_i}{\partial f_j} = -\frac{\partial \tilde{a}_j}{\partial f_j}\frac{d_i d_j}{M}, \quad j \neq i.$$

Note that condition (D) requires that

$$\frac{a_i}{\sum_{j=0}^N a_j} < \frac{\partial a_i/\partial p_i}{\sum_{j\neq i} \partial a_j/\partial p_j} \quad \text{for all } i = 1, \dots, N. \tag{11}$$

THEOREM 2. *Consider under the general attraction model* (8) *the single-stage pricing game that arises when the vector of service levels f is fixed.*

(a) *The price game has a Nash equilibrium. The set of Nash equilibria is a lattice and, therefore, has a smallest $\underline{p}(f)$ and a largest $\bar{p}(f)$ element.*

(b) *The Nash equilibrium is unique, i.e., $\underline{p}(f) = \bar{p}(f) = p^*(f)$ under* (D) *if the function $\tilde{a}_i$ is concave in p_i for all $i = 1, \dots, N$.*

The proof of Theorem 2 shows that the single-stage price game is supermodular. One of the implications of this characterization is that a simple, so-called tatônnement or round robin scheme, converges to $\underline{p}(f)[\bar{p}(f)]$ provided it starts with an arbitrary price vector $p^0 \leqslant [\geqslant] \underline{p}(f)[\bar{p}(f)]$. In the kth iteration, each firm i determines which price level p_i^k maximizes its expected profits, assuming all competitors maintain their prices from the vector p^{k-1}. When all $a_i(\cdot)$ functions are log-concave in p and (D) holds, there is a unique equilibrium $p^*(f)$, by Theorem 2(b). In this case, starting with an arbitrary price vector p^0, the sequence p^k converges to $p^*(f)$. In view of (25) and the fact that the profit function $\tilde{\pi}_i$ is strictly concave in p_i (see (26)), this gives rise to the following simple iterative scheme. In iteration k, determine for each firm i the unique root of the equation

$$p = w_i + k_i(f_i) - \left(\frac{\partial \tilde{a}_i(p, f_i)}{\partial p_i}\right)^{-1} \frac{M}{M - d_i(p, p_{-i}^{k-1}, f_i)}. \tag{12}$$

The left-hand side of (12) is increasing and the right-hand side decreasing in p, because $M/(M - d_i(p, p_{-i}^{k-1}, f_i))$ is decreasing in p while the same is true for $-(\partial \tilde{a}_i(p, f_i)/\partial p_i)^{-1}$, as $\tilde{a}_i$ is concave in p_i. Thus, instead of solving a system of N nonlinear equations in N unknowns, p^* can be obtained by iteratively solving N independent single equations in a single unknown.

For general attraction functions $\{a_i\}$, it is difficult to predict how a change in the service levels f impacts the equilibrium prices. It is, for example, reasonable to expect that if a firm increases its service level, its equilibrium price will increase as well, but this can only be guaranteed for certain classes of attraction models, e.g., the generalized MNL model (10), analyzed below. We now turn our attention to the (simultaneous) reduced single-stage game, in which the retailers choose their prices and service levels simultaneously.

THEOREM 3. *Under the general attraction model, assume that all $\tilde{a}_i$-functions are jointly concave.*

(a) *The simultaneous single-stage* (*reduced*) *game has a Nash equilibrium* (p^*, f^*) *and any Nash equilibrium is a* *solution to the system of equations*

$$\frac{\partial \tilde{\pi}_i}{\partial p_i} = \frac{1}{p_i - w_i - k_i(f_i)} + \frac{\partial \tilde{a}_i}{\partial p_i}\left(1 - \frac{d_i}{M}\right) = 0, \quad i = 1, \ldots, N, \tag{13}$$

$$k_i'(f_i)\frac{\partial \tilde{a}_i}{\partial p_i} + \frac{\partial \tilde{a}_i}{\partial f_i} = 0, \quad i = 1, \ldots, N. \tag{14}$$

In particular, $f_i^* > h_i^-/(h_i^- + h_i^+)$ *for all* $i = 1, \ldots, N$.

(b) *Assume that all* $\tilde{a}_i$*-functions are submodular. Then, a change in some of Retailer* i*'s characteristics* $\{w_i, k_i(\cdot), a_i(\cdot, \cdot)\}$ *causes the equilibrium price and service level of any of the competitors to move in opposite directions, i.e., each firm either simultaneously increases its price and decreases its service level, or it decreases its price while increasing its service level.*

Remark. The proof of Theorem 3(b) and the fact that (14) is independent of w_i also establish that a change in Retailer i's wholesale price causes his own equilibrium price and service level to move in opposite directions, just like those of his competitors.

The simultaneous game fails to be supermodular, even for special cases such as the MNL model. Therefore, little can be said in general about the structure or cardinality of the set of Nash equilibria. The case where the functions $\{\tilde{a}_i\}$ are of the form $\tilde{a}_i(p_i, f_i) = b_i(f_i) + \alpha_i(p_i)$, with $\alpha_i(\cdot)$ and $b_i(\cdot)$ concave, represents one of the important special cases of the conditions in Theorem 3. Separable functions $\{\tilde{a}_i\}$ are appropriate when the percentage increase in the attraction value of a firm due to a marginal change in its price is independent of the prevailing service level and vice versa. Under separable functions $\{\tilde{a}_i\}$, we have that any retailer's equilibrium price and service level always move in opposite directions if an infinitesimal change occurs in any of the competitors' characteristics, i.e., the retailer responds either by being more or less competitive along both the price and service dimensions. In case of interdependence, i.e., when $\tilde{a}_i$ fails to be separable, this result continues to hold as long as $\partial^2\tilde{a}_i/\partial p_i\,\partial f_i \leqslant 0$, i.e., as long as $\tilde{a}_i$ is submodular. On the other hand, in the case of interdependence, one often expects that $\tilde{a}_i$ be supermodular ($\partial^2\tilde{a}_i/\partial p_i\,\partial f_i > 0$), reflecting decreased sensitivity of the attraction values to price changes when operating in a higher service-level regime. If $\tilde{a}_i$ is supermodular, $\partial f_i^*/\partial p_i^* > 0$ may occur; see (29). That is, the equilibrium price and service level of a retailer may move in the same direction in response to a change in the characteristics of some of the competitors: the retailer may respond by being more competitive along one strategic dimension but less along the other.

Theorem 3(b) answers the important question of whether each Retailer i can choose its service level f_i^* simply on the basis of its local characteristics only, i.e., its cost function $k_i(\cdot)$ and the attraction function a_i, which captures the revenue implications of improved service levels. This would imply, in particular, that f_i^* is invariant to changes in the characteristics of any of the competitors, an invariance result often assumed in standard inventory models. Theorem 3(b) shows that it applies if and only if (in equilibrium) $\partial^2\tilde{a}_i/\partial p_i^2 = 0$ and $\partial^2\tilde{a}_i/\partial p_i\,\partial f_i = 0$. The only attraction functions $\tilde{a}_i$ that satisfy these two conditions globally are those that are separable in p_i and f_i and linear in p_i. This gives rise, precisely, to the generalized MNL structure (10). Finally, for general attraction models, little can be said about the equilibrium behavior in the two-stage game where prices and service levels are chosen sequentially.

5.2. The (Generalized) MNL Model

We now give special attention to the generalized MNL model (10). Under this structure, each function $\tilde{a}_i$ is concave in p_i so that a unique Nash equilibrium $p^*(f)$ exists in the single-stage price game, under any fixed vector of service levels f, assuming (D) holds (as is the case, for example, when all $\alpha_i = \alpha$, $i = 1, \ldots, N$; see (11)).

Similarly, the simultaneous single-stage (reduced) game has a Nash equilibrium (p^*, f^*) because the conditions of Theorem 3 are satisfied. Moreover, by (14), for all $i = 1, \ldots, N$, the equilibrium service level f_i^* equals f_i^0, with f_i^0 the unique root of the function $\delta_i(f)$, where

$$\delta_i(f) = k_i'(f) - \frac{b_i'(f)}{\alpha_i}. \tag{15}$$

By Lemma 1 and the concavity of the b_i-functions, $\delta_i(f) = -b_i'(f)/\alpha_i < 0$ for any $f \leqslant h_i^-/(h_i^+ + h_i^-)$, and δ_i increases to $+\infty$, as $f_i \nearrow 1$. When an out-of-pocket rate h_i^- prevails, generally $h_i^- \geqslant h_i^+$, so that $h_i^-/(h_i^+ + h_i^-) \geqslant 0.5$. If $h_i^- = 0$ or $h_i^- < h_i^+$, we assume that $k_i'(0.5) < b_i'(0.5)/\alpha_i$ to preclude unrealistic settings where firm i is "best off" with a fill rate of less than 50%. Either way, we have that $f_i^0 > \max\{0.5, h_i^-/(h_i^+ + h_i^-)\}$.

The first term on the right of (15) represents the incremental operational costs associated with a marginal increase in the service level, while the second term denotes the incremental retail price value, i.e., the price increase that this marginal increase permits without altering the attraction value of firm i. f_i^0 thus represents the unique service level for which the incremental operational costs equal the incremental retail price value. We conclude:

Theorem 4. *Consider the MNL model and assume that* (D) *holds. For every fixed service-level vector* f, *there exists a unique Nash equilibrium* $p^*(f)$ *in the price game. The pair* $(p^*(f^0), f^0)$ *is the unique Nash equilibrium in the simultaneous single-stage* (*reduced*) *retailer game.*

Note that each of the unique equilibrium service levels $f_i^* = f_i^0$ depends only on Retailer i's characteristics $\{k_i(\cdot), b_i(\cdot)$ and $\alpha_i\}$; in particular, the equilibrium service level f_i^* is invariant to any changes in the characteristics of the competitors and resulting changes in their equilibrium service levels and prices. This is in contrast to the general attraction model; see Theorem 3(b). Similarly, the equilibrium service levels f^0 are entirely independent of the vector of wholesale prices w. Furthermore, an increase in any of the wholesale prices w_i causes all equilibrium retail prices p^* to go up, while leaving the equilibrium service levels unaltered. (This follows because $p^* = p^*(f^0)$ is the unique equilibrium in the price game under $f = f^0$ as the fixed vector of service levels. This price game is supermodular and $\partial^2 \tilde{\pi}_i / \partial p_i \, \partial w_i > 0$ for all i. The monotonicity result follows from Topkis 1998.)

Observe, however, that in the simultaneous game the service levels f^0 fail to be dominant. In particular, if some of firm i's competitors j choose retail prices $p_j \neq p_j^*(f)$, firm i's best corresponding service level may depend on all prices and service levels chosen by the competitors. This is demonstrated in Example 1.

Example 1. Let $N = 3, M = 100$, and $a_0 = 0$, $w_i = 10$, $h_i^+ = 5$, while $h_i^- = 0$. Let $b_i(f_i) = 6f_i, i = 1, 2, b_3(f_3) = 3f_3$, and assume all ϵ-variables are Normally distributed with mean 1 and standard deviation 0.3. The equilibrium has $f^* = (0.866, 0.866, 0.698)$ and $p^* = (13.7, 13.7, 12.1)$. Assume all retail prices are frozen at their equilibrium as is the service level f_3^*. If Retailer 2 deviates from its equilibrium service level f_2^* to $\tilde{f}_2$, $p^* \neq p^*(f_1^*, \tilde{f}_2, f_3^*)$ and Retailer 1 is best off adjusting his service level f_1. For example, if Retailer 2 varies his service level from 75% to 95%, Retailer 1 should modify his from 83.8% to 88.2%.

Returning to the single-stage price game, Theorem 5 below shows that in the MNL model, it is possible to fully characterize whether a change in one of the firms' service levels will result in an increase or decrease in each of the equilibrium prices.

Theorem 5. *Under the MNL model, consider the single-stage price game for a fixed vector of service levels f. Assume that* (D) *holds.*

(a) *The price game has a unique equilibrium $p^*(f)$, where p^* is a differentiable function of the* (*fixed*) *vector of service levels f, with $(\partial p_i^* / \partial f_j)_{i,j=1}^N = A^{-1}B$. Here, A and B are $N \times N$ matrices with*

$$A_{ii} = \alpha_i^2 \frac{M - d_i}{M} \quad \text{and} \quad A_{ij} = -\alpha_i \alpha_j \frac{d_i d_j}{M^2}, \quad j \neq i,$$

$$B_{ii} = \frac{\alpha_i^2 (M - d_i)^2}{M^2} \delta_i(f_i) + \frac{b_i'(f_i)}{\alpha_i} A_{ii} \quad \text{and}$$

$$B_{ij} = \frac{b_j'(f_j)}{\alpha_j} A_{ij}, \quad j \neq i,$$

$$A^{-1} \geqslant 0 \quad \text{and} \quad A_{ii}^{-1} < \frac{M^2}{\alpha_i^2 (M - d_i)^2}. \tag{16}$$

(b) *p_i^* is strictly increasing in f_i; p_j^* is strictly decreasing in f_i for $f_i < f_i^0$; and strictly increasing in f_i for $f_i > f_i^0$.*

When firm i decides to increase its service level, it will always cause an increase in its own equilibrium retail price; i.e., the firm is able to charge a higher retail price in exchange for offering a higher service level. The impact of the service-level increase on the competitors' prices depends, however, on whether the service level is below or above the critical value f_i^0. As long as firm i's service level is below this critical value, the competitors react by decreasing their retail prices so as to "regain" their competitive edge. (In other words, the competitive impact of firm i's service-level increase exceeds that of the firm's price increase.) At the same time, when firm i's service level is already above the critical value f_i^0, the competitive impact of the simultaneous increase in firm i's price and service level allows the competitors to increase their prices along with (but not necessarily in the same amount as) firm i. Recall that when $f_i > f_i^0$, a further increase in firm i's service level results in a larger increase in per-unit operating costs than is recovered by the incremental "retail price value."

Theorem 6 below establishes that if firm i decides to increase its service level, this will result in an increased expected sales volume as well as increased expected profits as long as the service level f_i stays below the critical level f_i^0. Any increase of the service level f_i beyond this critical level always results in a decrease of the firm's expected sales volume and profits. Thus, assuming that the firms in the market adopt the unique equilibrium price vector $p^*(f)$ in response to any given combination of service levels f, it is optimal for each firm i to implement the service level f_i^0 regardless of the service levels offered by any of its competitors. In other words, assuming retail prices are in (the unique Nash) equilibrium, the optimal service level provided by any given firm can, under the MNL model, be determined as in standard inventory models, ignoring the competitive impact of service-level choices. The above implies that the vector of critical service levels f^0 is a dominant solution in the sequential two-stage game in which firms first (simultaneously) select a service level and in a second stage (simultaneously) decide on their retail prices. For each $i = 1, \ldots, N$, let $d_i^*(f) \stackrel{\text{def}}{=} d_i(p^*(f), f)$ and $\pi_i^*(f) \stackrel{\text{def}}{=} \pi_i(p^*(f), f)$ denote the expected sales volume and profit value under the vector of service levels f, assuming all firms adopt equilibrium retail prices.

Theorem 6. *Consider the MNL model and assume* (D).

(a) $\partial d_i^*(f) / \partial f_i > 0$ *for* $f_i < f_i^0$, $\partial d_i^*(f) / \partial f_i < 0$ *for* $f_i > f_i^0$, $i = 1, \ldots, N$.

(b) $\partial \pi_i^*(f) / \partial f_i > 0$ *for* $f_i < f_i^0$, $\partial \pi_i^*(f) / \partial f_i < 0$ *for* $f_i > f_i^0$, $i = 1, \ldots, N$.

(c) *The vector f^0 is a dominant solution in the two-stage game, with $p = p^*(f^0)$ the unique Nash equilibrium for its second stage.*

The above dominance result is in sharp contrast to settings where retail prices fail to be in equilibrium, i.e., where the vector of prevailing retail prices $p \neq p^*(f)$.

Here, each firm's optimal service level may critically depend on the service levels and/or retail prices offered by the competitors. As discussed in Fudenberg and Tirole (1991, Chapter 4), it is infrequent that a two-stage game can be shown to possess a Nash equilibrium. It is even more remarkable that the first-stage game has a unique Nash equilibrium of the strongest possible kind: f^0 is a dominant solution.

As mentioned in §2, Anderson et al. (1992) consider a special case of the model, with fully identical retailers and linear functions $b_i(f_i) = \theta f_i$ for some $\theta > 0$. These authors focus on the sequential game and establish the existence of a symmetric Nash equilibrium, which, in the case of two firms, is shown to be the unique Nash equilibrium. (It is easily verified that when the retailers have identical demand and cost characteristics, the vectors f^0 and $p^*(f^0)$ have identical components as well.) The authors also note (p. 241, Footnote 4) that the same symmetric solution is a Nash equilibrium in the simultaneous game as well.

6. The Linear Model

We now consider the case where the average demand functions $d(p, f)$ are linear in all prices and service levels. As mentioned in the introduction, this linear structure was considered in Banker et al. (1998) and Tsay and Agrawal (2000). Thus, let

$$d_i(p, f) = a_i - b_i p_i + \sum_{j\neq i} c_{ij} p_j + \beta_i f_i - \sum_{j\neq i} \gamma_{ij} f_j, \quad i = 1, \ldots, N, \tag{17}$$

with b_i, c_{ij}, β_i, and γ_{ij} positive constants, to ensure that the monotonicity properties in (2) are satisfied. Throughout this section, we assume that (D) holds, which translates to $b_i > \sum_{j\neq i} c_{ij}$, $i = 1, \ldots, N$. As in the attraction models, we start with the single-stage price game.

Theorem 7. *Consider, for the linear model, the single-stage price game that arises when the vector of service levels f is fixed.*

(a) *The price game has a unique equilibrium $p^*(f)$, with*

$$p_i^*(f) = \sum_{j=1}^{N} \alpha_{ij}[a_j + b_j w_j + b_j k_j(f_j) + \theta_{ij} f_j], \quad i = 1, \ldots, N, \tag{18}$$

with $[\alpha_{ij}]$ a matrix of nonnegative numbers and $[\theta_{ij}]$ a matrix of general constants. In particular, each of the equilibrium prices is a separable convex function of the vector f.

(b) *The equilibrium prices p^* are increasing in each of the wholesale prices $\{w_j\}$, demand function intercepts $\{a_j\}$, and each of the holding cost values $\{h_j^+\}$.*

(c) *For all $i, j = 1, \ldots, N$, there exist critical service levels $f_{ij}^* \geqslant 0.5$ such that p_i^* is decreasing in f_j for $0.5 \leqslant f_j \leqslant f_{ij}^*$ and increasing for $f_j \geqslant f_{ij}^*$.*

As in the MNL model, if firm j's service level f_j goes up, the retail price offered by every competing firm $i \neq j$ decreases, as long as the service level stays below a critical level f_{ij}^*, while resulting in a price increase for $f_j > f_{ij}^*$. Contrary to the MNL model, a different critical level may, however, prevail for each of firm j's competitors and $f_{ij}^* = 0.5$ may arise, in which case a service-level increase by firm j is always met by a price increase by competitor i. In contrast to the MNL model, when a firm increases its service level, this does not necessarily allow the firm to increase its price. This prima facie surprising phenomenon arises when the service-level increase causes the firm's competitors to reduce their prices by a significantly larger amount to offset the impact of the service-level increase. It appears that this is most likely to occur when a firm whose clientele is relatively price sensitive, but insensitive to service, needs to compete with other firms whose customers have the opposite attributes. As in the MNL model, $p^*(f)$ is easily computed via the tatônnement scheme.

We now proceed with the simultaneous single-stage game in which prices and service levels are chosen simultaneously. As in (15) for the MNL model, we assume that $k_i'(0.5) < \beta_i / b_i$ to preclude unrealistic settings where firm i is "best off" offering a fill rate of less than 50%.

Theorem 8. *The simultaneous (reduced) single-stage game for the linear model has a unique Nash equilibrium $(p^*(f^0), f^0)$, with f_i^0 the unique solution to the equation*

$$k_i'(f_i) = \frac{\beta_i}{b_i}, \quad i = 1, \ldots, N. \tag{19}$$

As in the MNL model, the simultaneous single-stage game has a unique Nash equilibrium, and the equilibrium service level f_i^0 for firm i is the level at which the marginal operational cost increase $k_i'(f_i)$ equals the marginal increase in retail price value. (The impact of an increase of firm i's service level by a basis point on the average sales volume is identical to a decrease of the price by β_i / b_i units.) In particular, the (unique) equilibrium service level of a firm only depends on that firm's demand and cost functions. Thus, as in the MNL model, but in contrast to the general attraction models, a firm's equilibrium service level is not affected by the attributes or equilibrium prices and service levels of any of its competitors.

The following represents a major contrast with the MNL model. Recall that in the latter, f^0 arises as the unique Nash equilibrium, and in fact the dominant solution in the sequential two-stage game in which the firms first choose their service levels and in the second stage select their retail prices. In the linear model, this robustness result breaks down. In fact, even though the closed-form solution for $p^*(f)$ in (18) permits us to specify the profit functions $\hat{\pi}(f) \stackrel{\text{def}}{=} \pi(p^*(f), f)$ in the first-stage game in closed form as well, these functions do not appear to have any structure

Table 1. Comparison of sequential and simultaneous games.

	d_1^*	d_2^*	d_3^*	π_1	π_2	π_3
Sequential	19.77	13.63	13.63	55.86	46.45	46.41
Simultaneous	18.87	12.58	12.58	50.88	39.55	39.55

that guarantees the existence of an equilibrium. The following example shows that f^0 fails to be a Nash equilibrium in the first-stage game. We conclude, in particular, that the equilibrium behavior of the industry may vary fundamentally, depending upon whether service levels and price levels are chosen simultaneously or sequentially.

EXAMPLE 2. Let $N=3$, with $d_1(p,f)=20-7p_1+2p_2+4p_3+4f_1-f_2-f_3$, $d_2(p,f)=20+p_1-4p_2+2p_3-8f_1+10f_2-f_3$, and $d_3(p,f)=20+p_1+2p_2-4p_3-8f_1-f_2+10f_3$. Assume that all ϵ-variables are Normally distributed with mean 1 and standard deviation 0.3. Let $w_i=5$, $i=1,2,3$, while $h_1^+=1$, $h_2^+=h_3^+=2$, and $h_i^-=0$ for $i=1,2,3$.

The simultaneous game has $(p^*(f^0),f^0)$ as its unique Nash equilibrium, with $p^*=(7.90,8.86,8.86)$ and $f^0=(0.68,0.87,0.87)$. At the same time, f^0 fails to be an equilibrium in the sequential game. The first-stage profit functions $\hat{\pi}(f)$ can, as mentioned, be obtained in closed form and are, in this example, concave. The sequential game has a unique service vector equilibrium $f^*=(0.50,0.90,0.90)$, with corresponding price vector $p^*(f^*)=(7.94,9.19,9.20)$. Table 1 compares the average sales volumes and expected profits for the retailers in the sequential and simultaneous games.

Retailer 1's clientele is relatively price sensitive, but insensitive to service. Retailers 2 and 3, at the same time, cater to a segment of the market that is willing to pay higher prices in exchange for better service. Indeed, both in the sequential and the simultaneous games, Retailer 1 adopts a lower service level and a lower price than its more "upscale" competitors. Observe, however, that simultaneous determination of prices and service levels results in significantly less differentiation between the competing retailers than in the sequential setting. Interestingly, Retailer 1 adopts a significantly lower service level and a somewhat higher price in the sequential, as compared to the simultaneous game setting, correspondingly realizing higher profits. Retailers 2 and 3, similarly, adopt higher prices in the sequential setting, but they compensate for their higher prices by offering better service. In the sequential setting, all retailers exploit their knowledge about all competitors' service levels to charge significantly larger prices, to achieve somewhat larger expected sales volumes and significantly larger profits. Finally, Figure 1 exhibits the dependency of $p_1^*(f)$ on f_1, fixing $f_2=f_2^*$ and $f_3=f_3^*$. It illustrates Theorem 7(c), showing that for sufficiently low service levels, i.e., $f_1 \leqslant f_{11}^*=0.68$, the equilibrium price of Retailer 1 decreases in response to a service level increase, but increases for service levels above the critical level f_{11}^*.

Figure 1. Dependence of p_1^* on f_1.

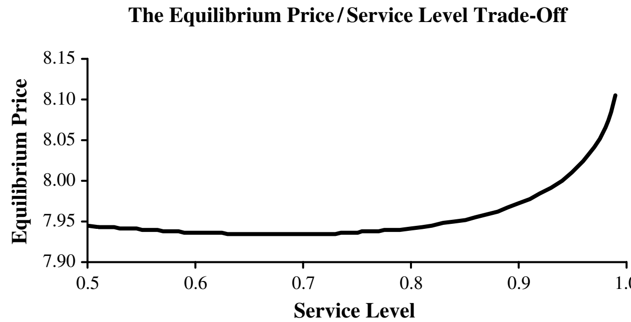

7. The Log-Separable Model

Our third and final model for the average sales functions $d(p,f)$ assumes that a regular system of standard demand functions $\{q_i(p)\colon i=1,\ldots,N\}$ is scaled up or down as a function of the service levels f offered by the different firms. This gives rise to the specification

$$d_i(p,f)=\psi_i(f)q_i(p), \tag{20}$$

with $\partial\psi_i(f)/\partial f_i>0$, $\partial\psi_i(f)/\partial f_j<0$, $j\neq i$, and with the normalization $\psi_i(f,f,\ldots,f)=1$. We further assume that the functions $\{\psi_i(f)\colon i=1,\ldots,N\}$ are log-supermodular in (f_i,f_j), while the functions $\{q_i(p)\}$ are log-supermodular in (p_i,p_j), i.e.,

$$\frac{\partial^2\log\psi_i(f)}{\partial f_i\,\partial f_j}\geqslant 0 \quad\text{and}\quad \frac{\partial^2\log q_i(p)}{\partial p_i\,\partial p_j}\geqslant 0 \quad\text{for all}\quad i\neq j. \tag{21}$$

Finally, to ensure a unique equilibrium in the reduced price game, we need

$$\frac{\partial^2\log q_i(p)}{\partial p_i^2}>\sum_{j\neq i}\frac{\partial^2\log q_i(p)}{\partial p_i\,\partial p_j}\quad\text{for all } i. \tag{22}$$

As shown in Milgrom and Roberts (1990), many systems of standard demand functions, i.e., demand functions of price variables only, have the above log-supermodularity property, for example:

(Linear) $q_i(p)=a_i-\beta_ip_i+\sum_{j\neq i}\beta_{ij}p_j$

with $a_i>0,\beta_i,\beta_{ij}\geqslant 0$ for all i and $j\neq i$;

(Logit) $q_i(p)=(k_ie^{-\lambda p_i})\Big/\left(C_i+\sum_{j=1}^{N}k_je^{-\lambda p_j}\right)$

with $\lambda>0$ and $C_i,k_i>0$ for all i;

(Cobb-Douglas) $q_i(p)=a_ip_i^{-\beta_i}\prod_{j\neq i}p_j^{\beta_{ij}}$

with $a_i>0,\beta_i>1,\beta_{ij}\geqslant 0$ for all i and $j\neq i$;

(CES) $q_i(p)=\gamma p_i^{r-1}\Big/\left(\sum_{j=1}^{N}p_j^{r}\right)$ with $r<0$ and $\gamma>0$.

Moreover, (22) is satisfied in all four cases, except that in the linear case condition (D) needs to be satisfied, i.e., $\beta_i > \sum_{j\neq i}\beta_{ij}$ for all i.

One possible specification of the functions $\psi(f)$, analogous to (Logit) is

$$\psi_i(f) = \frac{(\sum_{j=1}^N \xi_j)e^{\mu f_i}}{\sum_{j=1}^N \xi_j e^{\mu f_j}}, \quad \text{with } \mu, \xi > 0. \tag{23}$$

The price game that arises under a fixed vector of service levels f is well known. Because, by assumption the functions $\{q_i(p)\}$ are log-supermodular, the price game has a unique Nash equilibrium $p^*(f)$ (see Milgrom and Roberts 1990). It is also easily verified that the equilibrium price vector is increasing in all of the service levels as it is in all of the wholesale prices w and all of the holding cost rates h^+. For example, the fact that $\partial p_i^*/\partial f_j \geqslant 0$ for all i, j, follows immediately from Topkis (1998) because

$$\frac{\partial^2 \tilde{\pi}_i}{\partial p_i\,\partial f_j} = \frac{\partial^2 \log[p_i - w_i - k_i(f_i)]}{\partial p_i\,\partial f_j} = 0 \quad \text{if } j \neq i \quad \text{and}$$

$$\frac{\partial^2 \tilde{\pi}_i}{\partial p_i\,\partial f_i} = \frac{k_i'(f_i)}{p_i - w_i - k_i(f_i)} \geqslant 0.$$

We thus proceed immediately to the analysis of the simultaneous reduced single-stage game in which prices and service levels are selected simultaneously.

Theorem 9. *Consider the log-separable model and assume* (21).

(a) *The reduced price game under a given vector f, has a unique Nash equilibrium $p^*(f)$ that is increasing in all of the service levels.*

(b) *The simultaneous single-stage game has a Nash equilibrium (p^*, f^*).*

Guaranteeing uniqueness of the equilibrium is more complex in this case. While a sufficient condition similar to (24) and involving various second-order derivatives of π_i can be derived, this condition does not appear to translate into simple conditions for the $q_i(\cdot)$ and $\psi_i(\cdot)$ functions. However, as the single-stage game is log-supermodular, even if multiple equilibria arise, they form a lattice, and there is a largest equilibrium that is preferred by all retailers.

In contrast to the MNL and the linear model, but in accordance with the general attraction models, the equilibrium vector of service levels f^* can no longer be determined by identifying for each firm i at what service level the marginal operational cost $k_i'(f_i)$ equals the marginal increase in "retail price value." Indeed, the equilibrium service level f_i^* for firm i may now depend on the demand and cost characteristics of all of its competitors. However, because the simultaneous game is log-supermodular, an equilibrium pair (p^*, f^*) may be computed by a tatônnement scheme similar to the one in §5.

Figure 2. Changes in prices and service levels.

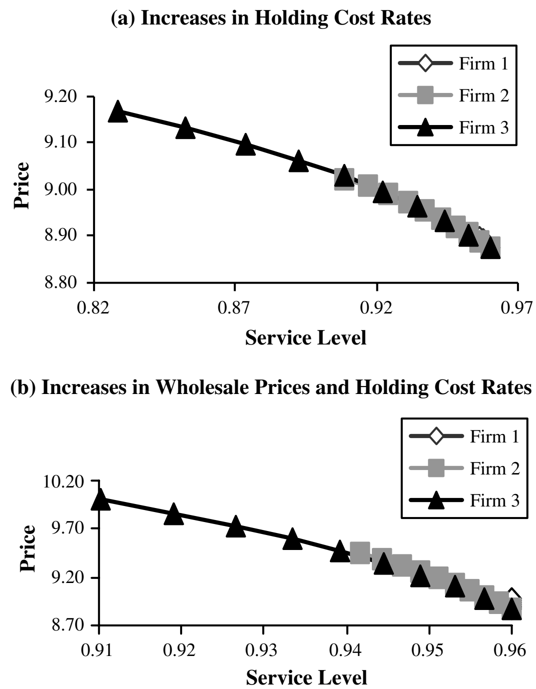

Example 3. Let $N = 3$. Assume that the mean demand functions are of the form $d_i(p, f) = (a_i - 8p_i + \sum_{j\neq i} p_j)\cdot(1 + f_i - 0.5\sum_{j\neq i} f_j)$ and that all ϵ-variables are Normally distributed with mean 1 and standard deviation 0.3. We consider a symmetric base scenario in which $w_i = 5$, $h_i^- = 0$, $h_i^+ = 1$, and $a_i = 80$ for $i = 1, 2, 3$. This system gives rise to a (unique) symmetric equilibrium in which all retailers adopt a price $p^* = 8.85$ and a service level $f^* = 0.942$.

In Figure 2(a) we show for 10 scenarios how the equilibrium evolves when the holding cost rate for Retailer 1 (2, 3) increases by 0 (7.5, 15)%, from one scenario to the next. All three retailers react to the increased holding cost rates by increasing their price and reducing their service level. As expected, Retailer 3's adjustments are larger than those of Retailer 2, which in turn exceed those of Retailer 1. In Figure 2(b) we repeat the same experiment, increasing the wholesale prices from one scenario to the next by 0%, 5%, and 10% for Retailers 1, 2, and 3, respectively, keeping the holding cost rates at 20% of the wholesale price values. In Figure 3, we show how the equilibrium evolves from the base scenarios when the intercepts a_i increase, from one scenario to the next, by 2.5%, 5%, and 10% for Retailers 1, 2, and 3, respectively. The retailers respond with a simultaneous increase of their price and service level. This applies even to Retailer 3, whose own demand and cost characteristics remain unaltered; his moderate price and service-level increases are clearly in response to the more extensive changes by the competitors. We conclude that under the log-separable model, a retailer's equilibrium price and service level may either change in the same or in opposite

Figure 3. Increases in the base demands a_i.

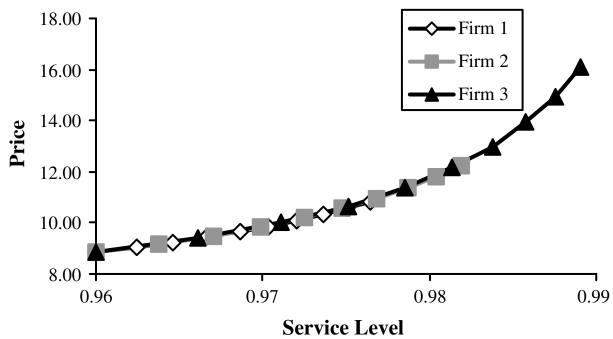

directions in response to a change in the model's inputs. This is in contrast to the general attraction models where, under the conditions of Theorem 3(b), a firm's equilibrium price and service level always change in opposite directions. The result is also in contrast to the (generalized) MNL and linear models in which each retailer's equilibrium service level is invariant to any changes in any of its competitors' characteristics.

As with the general attraction and linear models, little can be said about the equilibrium in the two-stage game. In general, f^* fails to be an equilibrium in the first-stage game.

8. Conclusions

We have systematically characterized the infinite-horizon equilibrium behavior under the three competition scenarios: (1) price competition only, (2) simultaneous price- and service-level competition, and (3) two-stage competition. We have shown that in each of these scenarios a Nash equilibrium of infinite-horizon stationary strategies exists under which each retailer adopts a stationary price, fill rate, and base-stock policy, provided a Nash equilibrium exists in a (reduced) single-stage (Scenarios (1) and (2)) or two-stage game (Scenario (3)).

As far as the reduced-price game characterizing Scenario (1) is concerned, its equilibrium behavior is similar across all of the classes of demand functions considered. The price game always has a Nash equilibrium, and the equilibrium is unique under widely satisfied conditions. The dependence of the equilibrium prices on the vector of service levels f varies, however, from one class of demand function to the next: little can be said for the general attraction model, but for the (generalized) MNL model, Retailer i's price p_i^* is always increasing in its own service level and decreasing in the service level f_j offered by any of his competitors $j \neq i$, as long as $f_j < f_j^0$ (the unique service level for which the incremental operational costs equal the incremental retail price value), and increasing thereafter. In the linear demand model, each Retailer i's equilibrium price p_i^* may no longer be increasing in his service level; instead $-p_i^*$ is unimodal in its own and every competitor's service levels. Finally, in the log-separable model, each equilibrium price is increasing in all service levels.

Turning to the simultaneous single-stage game that characterizes the simultaneous competition Scenario (2), a Nash equilibrium (p^*, f^*) exists for all of the considered classes of demand functions under mild conditions: in the general attraction model, it is sufficient that the attraction functions be log-concave (which they are in the specification of the MNL model), no restrictions are required in the linear model, while in the log-separable model, it is sufficient that the functions ψ_i and q_i be log-supermodular in (f_i, f_j) and (p_i, p_j), respectively. In the generalized MNL and linear models, $f^* = f^0$, the unique break-even service-level vector, so that each retailer's equilibrium service level is completely invariant to changes in the characteristics of any of the competitors. In the general attraction and the log-separable models, at the same time, each retailer needs to adjust his equilibrium service level in response to changes in the competitors' characteristics; in the former case, the adjustment is always in the opposite direction of the adjustment of the equilibrium price, implying that the retailer either becomes more or less competitive along both the price and service dimensions. Numerical examples show that under the log-separable model, the equilibrium service level may move in the opposite or in the same direction.

Little can be said about the two-stage competition Scenario (3) except in the (generalized) MNL model, where the vector of break-even service levels f^0 is again a Nash equilibrium and, in fact, a dominant solution in the first-stage game. Numerical examples for some of the other demand classes show that the equilibrium service level in the two-stage game may differ significantly from that which arises as part of the simultaneous game equilibrium.

We conclude that when estimating a system of (stochastic) demand equations, the proper class of demand models must be chosen with great care, because the choice has profound implications for the industry's equilibrium behavior in some or all of the considered competition scenarios. We hope that future empirical work, based on equilibrium models as in this paper, will characterize how firms in different industries position themselves in terms of their prices and (service) quality levels and how the equilibrium in a given industry evolves in response to an external change. Graham et al. (1983) and Bailey et al. (1985) provide examples of descriptive studies of this type in the airline industry. Here, quality of service on a given route is measured by the frequency of an airline's flights on this route and the probability of finding an available seat on the flight closest to the passenger's preferred departure time.[12] Graham et al. (1983) and Bailey et al. (1985) focus on the impact of deregulation of the industry in 1978. They report that most airlines responded to the industry's deregulation by increasing their load factors (thus reducing this measure of service quality), while simultaneously decreasing their prices.[13]

Our model can be extended to the case where some of the customers switch to a substitute retailer when they encounter a stockout. One possible model with substitution demand incorporates an $N \times N$ matrix P. Here, P_{ij} is the probability that a customer of Retailer i switches to Retailer j when faced with a stockout. The switching probabilities in P would themselves depend on part or all of the price and service-level vectors, i.e., $P = P(p, f)$. It remains an open question whether this or a similar model with substitution demands is tractable.

Appendix. Proofs

Proof of Theorem 1. Fix $i = 1, \ldots, N$. Assume that all retailers $j \neq i$ adopt a stationary retailer price p_j and a fill rate f_j throughout the infinite horizon. Retailer i then faces an infinite-horizon combined pricing and inventory control problem with, in each period, stochastic demand functions $D_i(\cdot, p_{-i}, f_{-i}) = d_i(\cdot, p_{-i}, f_{-i})\epsilon_i$ and with stationary and linear procurement and holding costs. An immediate adaptation of Federgruen and Heching (1999) shows that it is optimal for Retailer i to adopt a stationary price p_i, a fill rate f_i, and stationary base-stock policy with base-stock level y_i such that (p_i, f_i, y_i) maximizes the profit function $\Pi_i(\cdot \mid p_{-i} f_{-i})$, i.e., (p_i, f_i) maximizes $\pi_i(\cdot \mid p_{-i}, f_{-i})$. The theorem thus follows from the fact that (p^*, f^*) is a Nash equilibrium of the reduced single-stage game. □

Proof of Lemma 1. (a) For $f_i \leqslant h_i^-/(h_i^- + h_i^+)$, $k_i(f_i)$ is constant in f_i; for $f_i > h_i^-/(h_i^- + h_i^+)$,

$$k_i'(f_i) = \frac{(h_i^+ + h_i^-)f_i - h_i^-}{g_i(G_i^{-1}(f_i))} \geqslant 0,$$

so k_i is an increasing and differentiable function of f_i.

(b) Because g_i has infinite support, $\lim_{f_i \nearrow 1} k_i'(f_i) = \infty$. For $f_i > h_i^-/(h_i^- + h_i^+)$,

$$k_i''(f_i) = \big[(h_i^+ + h_i^-)g_i^2(G_i^{-1}(f_i)) - ((h_i^+ + h_i^-)f_i - h_i^-) \cdot g_i'(G_i^{-1}(f_i))\big]/g_i^3(G_i^{-1}(f_i)) \geqslant 0.$$

By condition (PF$_2$) in (5), $g_i'(G_i^{-1}(f_i)) \leqslant g_i^2(G_i^{-1}(f_i))/f_i$: Changing variables $x_i = G_i^{-1}(f_i) \in [G_i^{-1}(0.5), \infty)$, this inequality can be written as $g_i^2(x_i) \geqslant G_i(x_i)g_i'(x_i) \Leftrightarrow (\log G_i)''(x_i) = [G_i(x_i)g_i'(x_i) - g_i^2(x_i)]/G_i^2(x_i) \leqslant 0$. As $-((h_i^- + h_i^+)f_i - h_i^-) < 0$, $k_i''(f_i) \geqslant [(h_i^+ + h_i^-)f_i - ((h_i^+ + h_i^-)f_i - h_i^-)]/(f_i g_i(G_i^{-1}(f_i))) = h_i^-/(f_i g_i(G_i^{-1}(f_i))) > 0$, so that $k_i(\cdot)$ is convex on the entire interval $(0, 1)$. We conclude that k_i is twice differentiable for all $f_i \neq h_i^-/(h_i^- + h_i^+)$. □

Proof of Theorem 2. (a) We show that the profit functions $\pi_i(p, f)$ are log-supermodular in p, which means that $\tilde{\pi}_i(p, f) \stackrel{\text{def}}{=} \log \pi_i(p, f) = \log[p_i - w_i - k_i(f_i)] + \tilde{d}_i(p, f)$ is supermodular in p. Because these functions are twice differentiable, supermodularity is equivalent to showing that $\partial^2 \tilde{\pi}_i(p, f)/\partial p_i \partial p_j = \partial^2 \tilde{d}_i(p, f)/\partial p_i \partial p_j \geqslant 0$ for all $j \neq i$. It follows from Lemma 2 that $\partial^2 \tilde{d}_i/\partial p_i \partial p_j = (\partial \tilde{a}_i/\partial p_i)(\partial \tilde{a}_j/\partial p_j)(d_i/M)(d_j/M) \geqslant 0$. Because each firm's set of feasible price values is a compact lattice, it follows that the price game is log-supermodular, establishing (a); see Topkis (1998) or Milgrom and Roberts (1990).

(b) A unique Nash equilibrium $p^*(f)$ is guaranteed, by Milgrom and Roberts (1990), if

$$-\frac{\partial^2 \tilde{\pi}_i}{\partial p_i^2} > \sum_{j \neq i} \frac{\partial^2 \tilde{\pi}_i}{\partial p_i \, \partial p_j}. \tag{24}$$

Note from (4) and Lemma 2 that

$$\frac{\partial \tilde{\pi}_i}{\partial p_i} = \frac{1}{p_i - w_i - k_i(f_i)} + \frac{\partial \log d_i}{\partial p_i} \tag{25}$$

$$\frac{\partial^2 \log d_i}{\partial p_i^2} = -\frac{1}{(p_i - w_i - k_i(f_i))^2} + \frac{\partial^2 \tilde{a}_i}{\partial p_i^2}\left(1 - \frac{d_i}{M}\right) - \left(\frac{\partial \tilde{a}_i}{\partial p_i}\right)^2 \frac{d_i}{M}\left(1 - \frac{d_i}{M}\right). \tag{26}$$

Thus, (24) holds if

$$-\frac{\partial^2 \log d_i}{\partial p_i^2} > \sum_{j \neq i} \frac{\partial^2 \log d_i}{\partial p_i \partial p_j}.$$

Note that

$$-\frac{\partial^2 \log d_i}{\partial p_i^2} > \left(\frac{\partial \tilde{a}_i}{\partial p_i}\right)^2 \frac{d_i}{M}\left(1 - \frac{d_i}{M}\right) > \frac{\partial \tilde{a}_i}{\partial p_i} \sum_{j \neq i} \frac{\partial \tilde{a}_j}{\partial p_j} \frac{d_i}{M} \frac{d_j}{M} = \sum_{j \neq i} \frac{\partial^2 \log d_i}{\partial p_i \, \partial p_j}, \tag{27}$$

where the second inequality is equivalent to

$$\frac{\partial \tilde{a}_i}{\partial p_i}\left(1 - \frac{d_i}{M}\right) < \sum_{j \neq i} \frac{\partial \tilde{a}_j}{\partial p_j} \frac{d_j}{M} \quad \text{or}$$

$$\frac{\partial \tilde{a}_i}{\partial p_i} = \frac{\partial a_i/\partial p_i}{a_i} < \sum_{j=1}^{N} \frac{\partial \tilde{a}_j}{\partial p_j} \frac{d_j}{M} = \left(\sum_{j \neq i} \frac{\partial a_j}{\partial p_j}\right) \Big/ \left(\sum_{l=0}^{N} a_l\right).$$

Thus, (27) follows from (D) and (11). □

Proof of Theorem 3. (a) By Fudenberg and Tirole (1991), it suffices to show that all functions $\tilde{\pi}_i(p, f) = \log(p_i - w_i - k_i(f_i)) + \tilde{d}_i(p, f)$ are jointly strictly concave in (p_i, f_i). By Lemma 1, this property is immediate for the first term. To verify concavity for $\tilde{d}_i$, note from Lemma 2 that $\partial \tilde{d}_i/\partial p_i = \partial \tilde{a}_i/\partial p_i(1 - d_i/M)$ and $\partial \tilde{d}_i/\partial f_i = \partial \tilde{a}_i/\partial f_i(1 - d_i/M)$. Thus,

$$\frac{\partial^2 \tilde{d}_i}{\partial p_i^2} = \left(1 - \frac{d_i}{M}\right)\left[\frac{\partial^2 \tilde{a}_i}{\partial p_i^2} - \left(\frac{\partial \tilde{a}_i}{\partial p_i}\right)^2 \frac{d_i}{M}\right] \leqslant 0 \quad \text{and}$$

$$\frac{\partial^2 \tilde{d}_i}{\partial f_i^2} = \left(1 - \frac{d_i}{M}\right)\left[\frac{\partial^2 \tilde{a}_i}{\partial f_i^2} - \left(\frac{\partial \tilde{a}_i}{\partial f_i}\right)^2 \frac{d_i}{M}\right] \leqslant 0$$

by the concavity of $\tilde{a}_i$ in p_i and f_i, while

$$\frac{\partial^2 \tilde{d}_i}{\partial p_i \partial f_i} = \left(1-\frac{d_i}{M}\right)\left[\frac{\partial^2 \tilde{a}_i}{\partial p_i \partial f_i} - \frac{\partial \tilde{a}_i}{\partial p_i}\frac{\partial \tilde{a}_i}{\partial f_i}\frac{d_i}{M}\right].$$

The determinant of the Hessian of $\tilde{d}_i$ equals

$$\begin{aligned}
&\left(1-\frac{d_i}{M}\right)^2\left\{\left[\frac{\partial^2 \tilde{a}_i}{\partial p_i^2}\frac{\partial^2 \tilde{a}_i}{\partial f_i^2} - \left(\frac{\partial^2 \tilde{a}_i}{\partial p_i \partial f_i}\right)^2\right] - \left(\frac{\partial \tilde{a}_i}{\partial p_i}\right)^2 \frac{\partial^2 \tilde{a}_i}{\partial f_i^2}\frac{d_i}{M}\right.\\
&\qquad\left. - \left(\frac{\partial \tilde{a}_i}{\partial f_i}\right)^2 \frac{\partial^2 \tilde{a}_i}{\partial p_i^2}\frac{d_i}{M} + 2\frac{\partial^2 \tilde{a}_i}{\partial p_i \partial f_i}\frac{\partial \tilde{a}_i}{\partial f_i}\frac{\partial \tilde{a}_i}{\partial p_i}\frac{d_i}{M}\right\}\\
&= \left(1-\frac{d_i}{M}\right)^2\left\{\left[\frac{\partial^2 \tilde{a}_i}{\partial p_i^2}\frac{\partial^2 \tilde{a}_i}{\partial f_i^2} - \left(\frac{\partial^2 \tilde{a}_i}{\partial p_i \partial f_i}\right)^2\right]\right.\\
&\qquad + \frac{d_i}{M}\left(\frac{\partial \tilde{a}_i}{\partial p_i}\sqrt{-\frac{\partial^2 \tilde{a}_i}{\partial f_i^2}} + \frac{\partial \tilde{a}_i}{\partial f_i}\sqrt{-\frac{\partial^2 \tilde{a}_i}{\partial p_i^2}}\right)^2\\
&\qquad\left. -2\frac{d_i}{M}\frac{\partial \tilde{a}_i}{\partial p_i}\frac{\partial \tilde{a}_i}{\partial f_i}\left[\sqrt{\frac{\partial^2 \tilde{a}_i}{\partial f_i^2}\frac{\partial^2 \tilde{a}_i}{\partial p_i^2}} - \frac{\partial^2 \tilde{a}_i}{\partial p_i \partial f_i}\right]\right\}\\
&\geqslant 0.
\end{aligned}$$

(The terms in square brackets are nonnegative by the joint concavity of $\tilde{a}_i$.) Because of the joint concavity of the functions $\{\tilde{\pi}_i\}$, a Nash equilibrium satisfies Equation (13) and

$$\frac{\partial \tilde{\pi}_i}{\partial f_i} = -\frac{k_i'(f_i)}{p_i - w_i - k_i(f_i)} + \frac{\partial \tilde{a}_i}{\partial f_i}\left(1-\frac{d_i}{M}\right) = 0, \quad i = 1, \ldots, N. \tag{28}$$

Multiplying (13) by $k_i'(f_i)$ and adding (28), we obtain (14). Because $k_i'(f_i) = 0$ for $f_i \leqslant h_i^-/(h_i^- + h_i^+)$ and $\partial \tilde{a}_i/\partial f_i > 0$, it follows from (14) that $f_i^* > h_i^-/(h_i^- + h_i^+)$.

(b) Note that each of the equations in (14) involves only a single pair (p_i, f_i). Thus, (14) implicitly determines f_i^* as a function of p_i^*, and by the Implicit Function Theorem,

$$\frac{\partial f_i^*}{\partial p_i^*} = -\frac{k_i'(f_i^*)\dfrac{\partial^2 \tilde{a}_i}{\partial p_i^2} + \dfrac{\partial^2 \tilde{a}_i}{\partial p_i \partial f_i}}{k_i''(f_i^*)\dfrac{\partial \tilde{a}_i}{\partial p_i} + k_i'(f_i^*)\dfrac{\partial^2 \tilde{a}_i}{\partial p_i \partial f_i} + \dfrac{\partial^2 \tilde{a}_i}{\partial f_i^2}} \leqslant 0, \tag{29}$$

because all second-order derivatives are negative, $k_i'(f_i^*) \geqslant 0$, and $0 \leqslant k_i''(f_i^*)$, which exists because $f_i^* > h_i^-/(h_i^- + h_i^+)$. □

Proof of Theorem 5. (a) The price game has a unique Nash equilibrium in view of Theorem 2, because $\tilde{a}_i(p_i, f_i) = (b_i(f_i) - p_i)/\mu$ is linear in p_i and (D) holds. Because p^* is a Nash equilibrium, it follows that p_i^* maximizes the function $\tilde{\pi}_i(p_i | p_{-i}^*, f)$, which is strictly concave in p_i. Moreover,

$$\frac{\partial \tilde{\pi}_i(p, f)}{\partial p_i} = \frac{1}{p_i - w_i - k_i(f_i)} - \alpha_i + \alpha_i\frac{d_i}{M},$$

so that

$$\lim_{p_i \searrow [w_i + k_i(f_i)]} \frac{\partial \tilde{\pi}_i(\cdot \mid p_{-i}^*, f)}{\partial p_i} = +\infty, \quad \text{while}$$

$$\lim_{p_i \nearrow +\infty} \frac{\partial \tilde{\pi}_i(\cdot \mid p_{-i}^*, f)}{\partial p_i} = -\alpha_i < 0,$$

because $\lim_{p_i \nearrow +\infty} d_i = 0$. Thus, p_i^* is the unique root of the equation

$$\frac{\partial \tilde{\pi}_i(p, f)}{\partial p_i} = \frac{1}{p_i - w_i - k_i(f_i)} - \alpha_i + \alpha_i\frac{d_i}{M} = 0. \tag{30}$$

Because the functions $\{\partial \tilde{\pi}_i(p, f)/\partial p_i\colon i = 1, \ldots, N\}$ are continuously differentiable in p and f, it follows from the Implicit Function Theorem that p^* is a differentiable function of f. Thus, with $\delta_i(f)$ defined as in (15), we have from (30) that

$$\begin{aligned}
A_{ii} &= -\frac{\partial^2 \tilde{\pi}_i}{\partial p_i^2} = \alpha_i^2\frac{(M-d_i)^2}{M^2} + \alpha_i^2\frac{d_i}{M}\frac{M-d_i}{M}\\
&= \alpha_i^2\frac{M-d_i}{M}, \quad i = 1, \ldots, N,
\end{aligned} \tag{31}$$

$$A_{ij} = -\frac{\partial^2 \tilde{\pi}_i}{\partial p_i \partial p_j} = -\alpha_i\alpha_j\frac{d_i d_j}{M^2}, \quad j \neq i, \tag{32}$$

$$\begin{aligned}
B_{ii} &= \frac{\partial^2 \tilde{\pi}_i}{\partial p_i \partial f_i} = \frac{k_i'(f_i)}{(p_i - w_i - k_i(f_i))^2} + b_i'(f_i)\frac{\alpha_i d_i}{M}\left(1-\frac{d_i}{M}\right)\\
&= k_i'(f_i)\alpha_i^2\left(1-\frac{d_i}{M}\right)^2 + b_i'(f_i)\frac{\alpha_i d_i}{M}\left(1-\frac{d_i}{M}\right)\\
&= \alpha_i^2\frac{M-d_i}{M^2}\left(\delta_i(f_i)(M-d_i) + M\frac{b_i'(f_i)}{\alpha_i}\right),
\end{aligned} \tag{33}$$

$$B_{ij} = \frac{\partial^2 \tilde{\pi}_i}{\partial p_i \partial f_j} = -\alpha_i\frac{b_j'(f_j)d_i d_j}{M^2}, \quad j \neq i. \tag{34}$$

To verify that the matrix A is indeed invertible for all pairs of vectors $(p^*(f), f)$ and that A^{-1} is a nonnegative matrix, note that A may be written as $A = \Lambda(I - T)$, with Λ a diagonal matrix with

$$\Lambda_{ii} = \alpha_i\frac{M-d_i}{M}, \quad i = 1, \ldots, N,$$

and T a substochastic matrix, with $T_{ii} = 0$ and

$$T_{ij} = \frac{\alpha_j}{\alpha_i}\frac{d_i d_j}{M(M-d_i)} \quad \text{for } j \neq i.$$

(T is substochastic, i.e., it is nonnegative and its row sums

$$\sum_{j=1}^{N} T_{ij} = \frac{d_i}{\alpha_i M}\frac{\sum_{j\neq i}\alpha_j d_j}{M-d_i} < \frac{d_i}{M} < 1 \quad \text{for all } i = 1, \ldots, N$$

because $\sum_{j\neq i}\alpha_j d_j < \alpha_i(M - d_i)$ by (D).) Thus, A is invertible and

$$A^{-1} = (I-T)^{-1}\Lambda^{-1} = \sum_{n=0}^{\infty} T^n\Lambda^{-1} \geqslant 0. \tag{35}$$

To verify the bound on the diagonal elements of A^{-1}, let c_{ij} denote the cofactor of the (i, j)th element of the matrix $(I-T)$ and $\det(I-T)$ its determinant. It follows from (35) and Cramer's rule that

$$(A^{-1})_{ii} = \frac{M}{\alpha_i^2(M-d_i)}\frac{c_{ii}}{\det(I-T)}. \tag{36}$$

Because the matrix T is substochastic, it follows from Theorem 2.3 in Seneta (1973) that $c_{ij} \leqslant c_{ii}$ for all i, j. Thus, developing $\det(I-T)$ along the ith row of the matrix, we obtain

$$\det(I-T) = c_{ii} - \sum_{j\neq i}\frac{\alpha_j}{\alpha_i}\frac{d_i d_j}{M(M-d_i)}c_{ij}$$
$$\geqslant c_{ii}\left(1-\sum_{j\neq i}\frac{\alpha_j}{\alpha_i}\frac{d_i d_j}{M(M-d_i)}\right).$$

Substituting this lower bound into (36), we obtain

$$(A^{-1})_{ii} \leqslant \frac{M}{\alpha_i^2(M-d_i)}\frac{1}{1-\sum_{j\neq i}(\alpha_j/\alpha_i)(d_i d_j/M(M-d_i))}$$
$$= \frac{M^2}{\alpha_i[\alpha_i M(M-d_i) - d_i\sum_{j\neq i}\alpha_j d_j]}$$
$$< \frac{M^2}{\alpha_i[\alpha_i M(M-d_i)-d_i\alpha_i(M-d_i)]} = \frac{M^2}{\alpha_i^2(M-d_i)^2},$$

where the second inequality follows from $\sum_{j\neq i}\alpha_j d_j < \alpha_i(M-d_i)$.

(b) Note that

$$B = A\begin{pmatrix} b_1'(f_1)/\alpha_1 & \cdots & 0 \\ & \ddots & \\ 0 & \cdots & b_N'(f_N)/\alpha_N \end{pmatrix}$$
$$+\begin{pmatrix} \alpha_1^2\frac{(M-d_1)^2}{M^2}\delta_1 & \cdots & 0 \\ & \ddots & \\ 0 & \cdots & \alpha_N^2\frac{(M-d_N)^2}{M^2}\delta_N \end{pmatrix}.$$

Thus,

$$\frac{\partial p_j^*}{\partial f_i} = (A^{-1}B)_{ji} = (A^{-1})_{ji}\frac{\alpha_i^2(M-d_i)^2\delta_i(f_i)}{M^2}, \quad i\neq j, \tag{37}$$

$$\frac{\partial p_i^*}{\partial f_i} = \frac{b_i'(f_i)}{\alpha_i} + (A^{-1})_{ii}\frac{\alpha_i^2(M-d_i)^2\delta_i(f_i)}{M^2}. \tag{38}$$

Thus, for all $j\neq i$, $\partial p_j^*/\partial f_i$ has the same sign as $\delta_i(f_i)$; i.e., $\partial p_j^*/\partial f_i < 0$ for $f_i < f_i^0$ and $\partial p_j^*/\partial f_i > 0$ for $f_i > f_i^0$, which shows that p_j^* is decreasing (increasing) in f_i for all $f_i < (>) f_i^0$. As for $\partial p_i^*/\partial f_i$, it is clearly positive for $f_i > f_i^0$ because $b_i' \geqslant 0$, A^{-1} is a nonnegative matrix (see part (a)), and $\delta_i(f_i) > 0$ for $f_i > f_i^0$. However, for $f_i < f_i^0$, $\delta_i(f_i) < 0$. Thus, invoking the upper bound for $(A^{-1})_{ii}$ in (16), we obtain $\partial p_i^*/\partial f_i \geqslant b_i'(f_i)/\alpha_i + \delta_i(f_i) = k_i'(f_i) > 0$ for $f_i < f_i^0$, as well. Thus, p_i^* is strictly increasing in f_i over the complete interval $[0, 1)$. □

Proof of Theorem 6. (a) Consider the matrix A defined in Theorem 5. We first prove the following identity for all $i = 1, \ldots, N$:

$$d_i\sum_{j\neq i}\alpha_j\frac{d_j}{M}(A^{-1})_{ji} = \alpha_i(M-d_i)(A^{-1})_{ii} - \frac{M}{\alpha_i}. \tag{39}$$

This identity is immediate from $\sum_{j=1}^N (A_{ij})(A^{-1})_{ji} = A_{ii}(A^{-1})_{ii} + \sum_{j\neq i}^N (A_{ij})(A^{-1})_{ji} = 1$ by substituting (31) and (32). It follows from Lemma 2 that

$$\frac{\partial d_i^*}{\partial f_i} = \frac{\partial d_i}{\partial f_i} + \sum_{j=1}^N\frac{\partial d_i^*}{\partial p_j}\frac{\partial p_j^*}{\partial f_i}$$
$$= \frac{b_i'}{M}d_i^*(M-d_i^*) - \alpha_i d_i\left(1-\frac{d_i^*}{M}\right)\frac{\partial p_i^*}{\partial f_i}$$
$$+\frac{d_i^*}{M}\sum_{j\neq i}\alpha_j d_j^*\frac{\partial p_j^*}{\partial f_i}.$$

Let $\Delta_{ii} = \alpha_i^2(M-d_i)^2\delta_i/M^2$. Substituting (37) and (38), we obtain

$$\frac{\partial d_i^*}{\partial f_i} = b_i'\frac{d_i^*}{M}(M-d_i^*) - \alpha_i\frac{d_i^*}{M}(M-d_i^*)$$
$$\cdot\left[\frac{b_i'}{\alpha_i} + (A^{-1})_{ii}\Delta_{ii}\right] + \frac{d_i^*}{\mu M}\sum_{j\neq i}\alpha_j d_j^*(A^{-1})_{ji}\Delta_{ii}.$$

Thus, substituting (39), we obtain

$$\frac{\partial d_i^*}{\partial f_i} = \Delta_{ii}(A^{-1})_{ii}\left[-\frac{\alpha_i d_i^*(M-d_i^*)}{M} + \alpha_i(M-d_i^*)\right] - \frac{M}{\alpha_i}\Delta_{ii}$$
$$= \Delta_{ii}\left[(A^{-1})_{ii}\alpha_i\frac{(M-d_i^*)^2}{M} - \frac{M}{\alpha_i}\right]. \tag{40}$$

In view of (16), the second factor in (40) is negative. Thus, $\partial d_i^*(f)/\partial f_i$ has the opposite sign of Δ_{ii} and hence the opposite sign of $\delta_i(f_i)$. Part (a) now follows immediately from the fact that $\delta_i(\cdot)$ is increasing with $f_i = f_i^*$ as its unique root.

(b) Note that

$$\frac{\partial \pi_i^*(f)}{\partial f_i} = \left[\frac{\partial p_i^*}{\partial f_i} - k_i'(f_i)\right]d_i^* + [p_i^* - w_i - k_i(f_i)]\frac{\partial d_i^*}{\partial f_i}.$$

Substituting (38) and (40), we obtain

$$\frac{\partial \pi_i^*}{\partial f_i} = \left[\frac{b_i'}{\alpha_i} - k_i' + (A^{-1})_{ii}\Delta_{ii}\right]d_i^* + \frac{M}{\alpha_i(M-d_i^*)}\frac{\partial d_i^*}{\partial f_i}$$
$$= -\frac{\delta_i d_i^*}{\alpha_i} + (A^{-1})_{ii}\Delta_{ii}d_i^*$$
$$+\frac{\Delta_{ii}}{M-d_i^*}\left[(A^{-1})_{ii}(M-d_i^*)^2 - \frac{M^2}{\alpha_i^2}\right]$$
$$= (A^{-1})_{ii}\Delta_{ii}M - \frac{\delta_i d_i^*}{\alpha_i} - \frac{\Delta_{ii}}{\alpha_i^2}\frac{M^2}{M-d_i^*}$$
$$= (A^{-1})_{ii}\alpha_i\frac{(M-d_i^*)^2}{M^2}\delta_i - \frac{\delta_i d_i^*}{\alpha_i} - (M-d_i^*)\frac{\delta_i}{\alpha_i}$$
$$= \delta_i\left[(A^{-1})_{ii}\alpha_i\frac{(M-d_i^*)^2}{M^2} - \frac{M}{\alpha_i}\right], \tag{41}$$

where the second equality follows from (40) and the fourth equality by the definition of Δ_{ii}. By (16), the second factor in (41) is always negative. Thus, $\partial \pi_i^*(f)/\partial f_i$ has, once again, the opposite sign of δ_i.

(c) Immediate from part (b). $\square$

Proof of Theorem 7. (a) We first show that the price game is log-supermodular. Because the feasible strategy space $\Omega(f) = \{p\colon\ w_i + k_i(f_i) \leqslant p_i \leqslant p_i^{\max}, i = 1, \ldots, N\}$ is a compact sublattice of $\Re^N$, it suffices to show that each of the functions $\tilde{\pi}_i \stackrel{\text{def}}{=} \log \pi_i = \log[p_i - w_i - k_i(f_i)] + \log d_i(p, f)$ has the property

$$\frac{\partial^2 \tilde{\pi}_i}{\partial p_i\, \partial p_j} = \frac{\partial^2 \log d_i(p, f)}{\partial p_i\, \partial p_j} \geqslant 0 \quad \text{for all } i \neq j.$$

Note that,

$$\frac{\partial \log d_i(p, f)}{\partial p_i} = -\frac{b_i}{d_i(p, f)},$$

so that, indeed,

$$\frac{\partial^2 \log d_i(p, f)}{\partial p_i\, \partial p_j} = \frac{b_i c_{ij}}{d_i^2(p, f)}. \tag{42}$$

To prove that a unique equilibrium exists, it suffices to show (Milgrom and Roberts 1990) that

$$\begin{aligned} -\frac{\partial^2 \tilde{\pi}_i(p, f)}{\partial p_i^2} &= \frac{1}{(p_i - w_i - k_i(f_i))^2} + \frac{b_i^2}{d_i^2(p, f)} \\ &> \frac{b_i^2}{d_i^2(p, f)} \geqslant \frac{b_i \sum_{j\neq i} c_{ij}}{d_i^2(p, f)} \\ &= \sum_{j\neq i} \frac{\partial^2 \log d_i(p, f)}{\partial p_i\, \partial p_j} = \sum_{j\neq i} \frac{\partial^2 \tilde{\pi}_i}{\partial p_i\, \partial p_j}. \end{aligned} \tag{43}$$

(The second inequality follows from $b_i > \sum_{j\neq i} c_{ij}$ and the second equality from (42).)

Because $p^*(f)$ is a Nash equilibrium, it follows that for all $i = 1, \ldots, N$, p_i^* is a global maximum of the single variable function $\tilde{\pi}_i(p \mid p_{-i}^*, f) = \log[p - w_i - k_i(f_i)] + \log[a_i - b_i p + C]$, with $C = \sum_{j\neq i} c_{ij} p_j^* + \beta_i f_i - \sum_{j\neq i} \gamma_{ij} f_j$. It follows that p_i^* is the unique solution to the equation

$$\frac{\partial \tilde{\pi}_i}{\partial p_i} = \frac{1}{p - w_i - k_i(f_i)} - \frac{b_i}{a_i - b_i p + C} = 0.$$

In other words, p^* is the unique solution to the system of equations

$$\frac{1}{p_i - w_i - k_i(f_i)} - \frac{b_i}{d_i(p, f)} = 0 \quad \text{or}$$
$$b_i p_i - b_i w_i - b_i k_i(f_i) = a_i - b_i p_i + \sum_{j\neq i} c_{ij} p_j + \beta_i f_i - \sum_{j\neq i} \gamma_{ij} f_j,$$
$$i = 1, \ldots, N. \tag{44}$$

This is a linear system of equations in p, which in matrix form can be written as

$$Ap = \left[a_i + b_i w_i + b_i k_i(f_i) + \beta_i f_i - \sum_{j\neq i} \gamma_{ij} f_j\right], \quad \text{where}$$

$$A = \begin{pmatrix} 2b_1 & \cdots & -c_{1N} \\ & \ddots & \\ -c_{N1} & \cdots & 2b_N \end{pmatrix} = \Lambda(I - T),$$

with $\Lambda = \operatorname{diag}(2b_1, \ldots, 2b_N)$, and

$$T = \begin{pmatrix} 0 & \cdots & \dfrac{c_{1N}}{2b_1} \\ & \ddots & \\ \dfrac{c_{N1}}{2b_N} & \cdots & 0 \end{pmatrix} \text{ is substochastic because}$$

$$\frac{\sum_{j\neq i} c_{ij}}{2b_i} \leqslant \frac{b_i}{2b_i} < 1.$$

Thus, A is invertible with $A^{-1} = (I - T)^{-1}\Lambda^{-1} = (\sum_{n=0}^{\infty} T^n)\Lambda^{-1} \geqslant 0$ and

$$\begin{aligned} p_i^*(f) &= \sum_{j=1}^{N} A_{ij}^{-1}[a_j + b_j w_j + b_j k_j(f_j)] \\ &\quad + \Big(A_{ii}^{-1}\beta_i - \sum_{j\neq i} A_{ij}^{-1}\gamma_{ji}\Big) f_i \\ &\quad + \sum_{j\neq i}\Big(A_{ij}^{-1}\beta_j - \sum_{l\neq j} A_{il}^{-1}\gamma_{lj}\Big) f_j, \end{aligned}$$

proving the claims in part (a). Parts (b) and (c) follow from (18) and $A^{-1} \geqslant 0$. $\square$

Proof of Theorem 8. We first verify that the profit functions

$$\tilde{\pi}_i(p, f) = \log[p_i - w_i - k_i(f_i)] + \log d_i(p, f) \tag{45}$$

are jointly concave in (p_i, f_i). Because $k_i(f)$ is convex, both terms to the right of (45) are jointly concave as the composition of an increasing concave function and a jointly concave function. This implies that any solution to the system of first-order conditions

$$0 = \frac{\partial \tilde{\pi}_i(p, f)}{\partial p_i} = \frac{1}{p_i - w_i - k_i(f_i)} - \frac{b_i}{d_i(p, f)}, \tag{46}$$

$$0 = \frac{\partial \tilde{\pi}_i(p, f)}{\partial f_i} = \frac{-k_i'(f_i)}{p_i - w_i - k_i(f_i)} + \frac{\beta_i}{d_i(p, f)}, \tag{47}$$

is a Nash equilibrium. Substituting (46) into (47), we obtain (19), which has a unique solution f^0 because $k_i'(0.5) < \beta_i/b_i$ and $\lim_{f_i \nearrow 1} k_i'(f_i) = +\infty$. Fixing $f = f^0$ in (46), we obtain (44) for which the proof of Theorem 5 shows that it has the unique solution $p^*(f^0)$. Thus, $(p^*(f^0), f^0)$

is the unique equilibrium on the interior of the action space $X_{i=1}^{N}\{(p_i, f_i): w_i + k_i(f_i) \leqslant p_i \leqslant p_i^{\max}; 0.5 \leqslant f_i < 1\}$. One easily verifies that no point on its boundary is an equilibrium. $\square$

Proof of Theorem 9. Only part (b) requires a proof. To prove the existence of a Nash equilibrium, it suffices to show that the game is log-supermodular in the price and service-level variables. Let $\tilde{\pi}_i = \log \pi_i$, $i = 1, \ldots, N$. First, note that the space of feasible strategies $X_{i=1}^{N}\{(p_i, f_i): p_i \geqslant w_i + k_i(f_i)\}$ is a sublattice of $\Re^{2N}$ because the set $\Omega_i = \{(p_i, f_i): p_i \geqslant w_i + k_i(f_i)\}$ is a sublattice of $\Re^2$: if $(p_i^1, f_i^1) \in \Omega_i$ and $(p_i^2, f_i^2) \in \Omega_i$, then $(\max(p_i^1, p_i^2), \max(f_i^1, f_i^2)) \in \Omega_i$ and $(\min(p_i^1, p_i^2), \min(f_i^1, f_i^2)) \in \Omega_i$. The game is therefore supermodular provided that for all $i = 1, \ldots, N$,

$$\frac{\partial^2 \tilde{\pi}_i}{\partial p_i \partial f_i} = \frac{\partial^2 \log[p_i - w_i - k_i(f_i)]}{\partial p_i \partial f_i} = \frac{k_i'(f_i)}{(p_i - w_i - k_i(f_i))^2} \geqslant 0,$$

$$\frac{\partial^2 \tilde{\pi}_i}{\partial p_i \partial f_j} = 0, \quad \frac{\partial^2 \tilde{\pi}_i}{\partial p_j \partial f_i} = 0,$$

$$\frac{\partial^2 \tilde{\pi}_i}{\partial p_i \partial p_j} = \frac{\partial^2 \log q_i(p)}{\partial p_i \partial p_j} \geqslant 0, \quad \frac{\partial^2 \tilde{\pi}_i}{\partial f_i \partial f_j} = \frac{\partial^2 \log \psi_i(f)}{\partial f_i \partial f_j} \geqslant 0. \quad \square$$

Endnotes

1. Alternatively, the expected value or a given fractile of the waiting time experienced by the customer may be used as the service-level target; see Bernstein and Federgruen (2002) for a treatment of this variant.
2. The model can be extended to the case where stockouts result in lost sales; the profit functions and resulting analysis become somewhat complex.
3. Their comparatively high fill rates were enabled by a novel revenue-sharing contract with the movie studios, reducing the wholesale price per tape by a factor of 10, approximately. Blockbuster's strategic move to increase its fill rates dramatically is generally credited for the company's acquiring a dominating market share of close to 35% and for its parent company Viacom seeing its stock doubled during the first year of this initiative.
4. Independent software firms, such as Compliance Networks, provide retailers with tracking systems to monitor the vendors' compliance with the prespecified service targets as well as with data comparing groups of vendors' fill rates and on-time shipment performance (see *Chain Store Age* 2002b).
5. See, for example, the recent survey of Bearing Point (formerly KPMG Consulting) in *Chain Store Age* (2002a, p. 5a), conducted in cooperation with Washington Inventory Services.
6. Two examples are bizrate.com and resellerrating.com. The former guides retailers on the basis of a number of "post-fulfillment satisfaction" measures, in particular, "availability of product you wanted" (defined as "product was in stock at time of expected delivery") and "on-time delivery" (defined as "product arrived when expected").
7. As with dynamic programs with a single decision maker, it is rare that infinite-horizon stochastic games reduce to a single-period stochastic game. When they do, the equivalency greatly simplifies the identification and structure of the infinite-horizon Nash equilibrium.
8. If a service level below 0.5 is used, customers are more likely to experience a backlog than not; in this case, characteristics of the customer waiting time, e.g., its expected value or a given fractile of the waiting time, should be used to characterize the customer service level (see Bernstein and Federgruen 2002).
9. Similarly, aggregate sales usually decrease if one of the firms increases its price: $\sum_{j=1}^{N} \partial d_j / \partial p_i < 0$ for all $i = 1, \ldots, N$.
10. As mentioned, this shape of the cost-service trade-off function has been assumed ex ante in the service/quality competition models of Anderson et al. (1992, p. 239), Besanko et al. (1998), and Tsay and Agrawal (2000).
11. Recent econometric studies based on the MNL model include Berry et al. (1995), Villas-Boas and Winer (1999), and Besanko et al. (1998); see McFadden (1980, 1986), Schmalensee and Thisse (1988), and Urban and Hauser (1980) for reviews of earlier applications.
12. Empirical studies usually aggregate both dimensions of quality via the average load factor, i.e., the total number of passengers divided by the seats available on a route.
13. See Jaskow and Rose (1989, §25.8) for a general survey of the impact of economic regulation on equilibria in a variety of industries.

Acknowledgments

The authors are indebted to the associate editor and two referees for their careful review and numerous suggestions, which have improved the exposition of this paper.

References

Anderson, S. P., A. de Palma, J. F. Thisse. 1992. *Discrete Choice Theory of Product Differentiation.* MIT Press, Cambridge, MA.

Arrow, K. J., T. Harris, J. Marschak. 1951. Optimal inventory policy. *Econometrica* **19**(3) 250–272.

Bailey, E., D. Graham, D. Kaplan. 1985. *Deregulating the Airlines*. MIT Press, Cambridge, MA.

Banker, R. D., I. Khosla, K. K. Sinha. 1998. Quality and competition. *Management Sci.* **44**(9) 1179–1192.

Barlow, R. E., F. Proschan. 1965. *Mathematical Theory of Reliability.* John Wiley and Sons, New York.

Basuroy, S., D. Nguyen. 1998. Multinomial logit market share models: Equilibrium characteristics and strategic implications. *Management Sci.* **44**(10) 1396–1408.

Bell, D., R. Keeney, J. D. Little. 1975. A market share theorem. *J. Marketing Res.* 136–141.

Bernstein, F., A. Federgruen. 2001. Decentralized supply chains with competing retailers under demand uncertainty. *Management Sci.* Forthcoming.

Bernstein, F., A. Federgruen. 2002. Coordination mechanisms for supply chains under price and service competition. Working paper, Duke University, Durham, NC and Columbia University, New York.

Bernstein, F., A. Federgruen. 2004. Dynamic inventory and pricing models for competing retailers. *Naval Res. Logist.* **51**(2) 258–274.

Berry, S., J. Levinsohn, A. Pakes. 1995. Automobile prices in market equilibrium. *Econometrica* **63** 841–890.

Besanko, D., S. Gupta, D. Jain. 1998. Logit demand estimation under competitive pricing behavior: An equilibrium framework. *Management Sci.* **44**(11) 1533–1547.

Caine, G. J., R. H. Plaut. 1976. Optimal inventory policy when stockouts alter demand. *Naval Res. Logist. Quart.* **23** 1–13.

Carlton, D. 1989. The theory and the facts of how markets clear. R. Schmalensee, R. Willig, eds. *Handbook of Industrial Organization*, Vol. 2. Elsevier Science, Amsterdam, The Netherlands.

Carlton, D., J. Perloff. 1999. *Modern Industrial Organization.* Longman Science & Technology, London, U.K.

Carr, S., I. Duenyas, W. Lovejoy. 1999. The harmful effects of capacity improvements. Working paper, University of Michigan, Business School, Ann Arbor, MI.

Chain Store Age. 2002a. Inventory management 2002: Data, detail & discipline. (December) Murray Forseter Publisher/Editor, New York.

Chain Store Age. 2002b. Just what the buyer ordered? (January) Murray Forseter Publisher/Editor, New York.

Dana, J. 2001. Competition in price and availability when availability is unobservable. *RAND J. Econom.* **32**(3) 497–513.

Dixit, A. 1979. Quality and quantity competition. *Rev. Econom. Stud.* **46**(4) 587–599.

Dorfman, R. B., P. D. Steiner. 1954. Optimal advertising and optimal quality. *Amer. Econom. Rev.* **44**(December) 826–836.

Dvoretzky, A., J. Kiefer, J. Wolfowitz. 1952. The inventory problem: I. Case of known distributions of demand. *Econometrica* **20** 187–222.

Ernst, R., M. Cohen. 1992. Coordination alternatives in a manufacturer/dealer inventory system under stochastic demand. *Production Oper. Management* **1** 254–268.

Ernst, R., S. Powell. 1995. Optimal inventory policies under service-sensitive demand. *Eur. J. Oper. Res.* **87** 316–327.

Ernst, R., S. Powell. 1998. Manufacturer incentives to improve retail service levels. *Eur. J. Oper. Res.* **104** 437–450.

Federgruen, A., A. Heching. 1999. Combined pricing and inventory control under uncertainty. *Oper. Res.* **47**(3) 454–475.

Fudenberg, D., J. Tirole. 1991. *Game Theory.* MIT Press, Cambridge, MA.

Graham, D., D. Kaplan, D. Sibley. 1983. Efficiency and competition in the airline industry. *Bell J. Econom.* **14** 118–198.

Hill, T. W., Jr. 1976. A simple perturbed demand inventory model with ordering cost. *Management Sci.* **23**(1) 38–42.

Jaskow, P., N. Rose. 1989. The effects of economic regulation. R. Schmalensee, R. Willig, eds. *Handbook of Industrial Organization*, Vol. 2. Elsevier Science, Amsterdam, The Netherlands.

Karmarkar, U., R. Pitbladdo. 1997. Quality, class, and competition. *Management Sci.* **43**(1) 27–39.

Kirman, A., M. Sobel. 1974. Dynamic oligopoly with inventories. *Econometrica* **42**(2) 279–287.

Leeflang, P., D. Wittink, M. Wedel, P. Naert. 2000. *Building Models for Marketing Decisions.* Kluwer Academic Publishers, Dordrecht/Boston/London.

Li, L. 1992. The role of inventory in delivery-time competition. *Management Sci.* **38**(2) 182–197.

Mahajan, S., G. van Ryzin. 1999. Retail inventories and consumer choice. S. Tayur, M. Magazine, R. Ganeshan, eds. *Quantitative Models for Supply Chain Management.* Kluwer, Dordrecht, The Netherlands, 491–551.

McFadden, D. 1980. Econometric models of probabilistic choice among products. *J. Bus.* **53** 513–552.

McFadden, D. 1986. The choice theory of market research. *Marketing Sci.* **5** 275–297.

Milgrom, P., J. Roberts. 1990. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* **58** 1255–1277.

Robinson, L. 1991. Appropriate inventory policies when service affects future demands. Working paper, Johnson Graduate School of Business, Cornell University, Ithaca, NY.

Schmalensee, R., J. F. Thisse. 1988. Perceptual maps and the optimal location of new products: An integrative essay. *Internat. J. Res. Marketing* **5** 225–249.

Schwartz, B. L. 1966. A new approach to stockout penalties. *Management Sci.* **12**(12) B538–B544.

Schwartz, B. L. 1970. Optimal inventory policies in perturbed demand models. *Management Sci.* **16**(8) B509–B518.

Seneta, E. 1973. *Non-Negative Matrices: An Introduction to Theory and Applications.* John Wiley and Sons, New York.

Spence, M. 1975. Monopoly, quality, and regulation. *Bell J. Econom.* **6**(2) 417–429.

Spence, M. 1976. Product selection, fixed costs, and monopolistic competition. *Rev. Econom. Stud.* **43**(2) 217–235.

Tirole, J. 1988. *The Theory of Industrial Organization.* MIT Press, Cambridge, MA.

Topkis, D. 1998. *Supermodularity and Complementarity.* Princeton University Press, Princeton, NJ.

Tsay, A., N. Agrawal. 2000. Channel dynamics under price and service competition. *Manufacturing Service Oper. Management* **2**(4) 372–391.

Urban, G. L., J. R. Hauser. 1980. *Design and Marketing of New Products.* Prentice-Hall, Englewood Cliffs, NJ.

Van Mieghem, J., M. Dada. 1999. Price versus production postponement: Capacity and competition. *Management Sci.* **45**(12) 1631–1649.

Villas-Boas, J. M., R. S. Winer. 1999. Endogeneity in brand choice models. *Management Sci.* **45**(10) 1324–1338.