MANAGEMENT SCIENCE
Vol. 55, No. 12, December 2009, pp. 2028–2043
ISSN 0025-1909 | EISSN 1526-5501 | 09 | 5512 | 2028
DOI 10.1287/mnsc.1090.1073

Competition Under Generalized Attraction Models: Applications to Quality Competition Under Yield Uncertainty

Awi Federgruen
Graduate School of Business, Columbia University, New York, New York 10027, af7@columbia.edu
Nan Yang
Johnson School of Management, Cornell University, Ithaca, New York 14853, ny38@cornell.edu

We characterize the equilibrium behavior in a broad class of competition models in which the competing firms' market shares are given by an attraction model, and the aggregate sales in the industry depend on the aggregate attraction value according to a general function. Each firm's revenues and costs are proportional with its expected sales volume, with a cost rate that depends on the firm's chosen attraction value according to an arbitrary increasing function. Whereas most existing competition papers with attraction models can be viewed as special cases of this general model, we apply our general results to a new set of quality competition models. Here an industry has N suppliers of a given product, who compete for the business of one or more buyers. Each of the suppliers encounters an uncertain yield factor, with a given general yield distribution. The buyers face uncertain demands over the course of a given sales season. The suppliers compete by selecting key characteristics of their yield distributions, either their means, their standard deviations, or both. These choices have implications for their per-unit cost rates.

Key words: inventory production; uncertainty; games-group decisions; stochastic models
History: Received October 12, 2007; accepted April 25, 2009, by Paul H. Zipkin, operations and supply chain management. Published online in *Articles in Advance* October 16, 2009.

1. Introduction and Summary

Starting with a seminal paper by Friedman (1958), there is a well-established tradition in the economics, marketing, and operations literature to model competition in oligopolies by assuming that the firms' sales volumes are specified by a so-called attraction model; see, e.g., Lilien et al. (1992), Cooper (1993), Karnani (1985), Anderson et al. (1992), Besanko et al. (1998), So (2000), and Gallego et al. (2006). In an attraction model, each firm's strategic choices determine a single "*attraction value*" such that its market share is given by the ratio of this attraction value and the sum of the industry's values. Bell et al. (1975) have shown that this is, in fact, the only representation of market *shares* to satisfy four simple axioms. With market shares determined by the above ratio rule, the firm's expected sales volumes are completely determined by a specification of the *aggregate* sales in the industry. The latter is, most commonly, assumed to be a constant (i.e., independent of the strategic choices) or to increase toward this constant potential market size according to a very specific function $T(R)$ of the aggregate attraction value R (see (5) below).[1] However, starting with Kotler (1965), several papers have modeled aggregate sales as a general function of the aggregate attraction value in the industry; see also Bell et al. (1975), Karnani (1985), and Basuroy and Nguyen (1998).

Although in many settings aggregate sales should be represented as an *increasing* function of the aggregate attraction value, there are models in which aggregate sales decrease as this value improves. This includes the model that motivated our paper and to which we devote its second part. Here, alternative suppliers of a common good compete with each other in terms of certain characteristics of their uncertain yield processes, for example, the reliability of each manufacturing batch's yield factor. In this model, which considers either a single buyer or a finite set of buyers, it is in fact possible to *derive* the suppliers' sales volumes explicitly by identifying the optimal procurement policy of each of the buyers. (Most

[1] If X denotes the aggregate of the attraction values in the industry, the total sales in the industry is given by a function of the form $MX/(X+C)$, with $M, C > 0$ given constants.

competition models assume a specific functional form of the demand functions as *exogenously* given.) The resulting expected sales volumes imply market shares given by an attraction model with a reliability measure serving as the attraction value; however, as the suppliers improve their reliability, aggregate purchases by the buyers *decline*, the reduced supply risks reducing the need for safety stocks. Thus, when an individual supplier increases his yield reliability, this results in an increase of his market share, although not necessarily of his expected sales volume. Because the per-unit manufacturing cost increases with the selected yield reliability, the yield improvement also results in a reduced profit margin, thus giving rise to intricate sets of trade-offs.

In this paper, we analyze a general competition model, with market shares determined by an attraction model and aggregate sales specified as a general function of the aggregate attraction value. Firms compete with each other by selecting their attraction value, which impacts each firm's market share and aggregate sales, as well as the per-unit cost incurred. We show that the general competition model has a (pure) equilibrium provided the attraction intensity function, $A(R) = R/T(R)$, is log-concave, irrespective of whether the aggregate sales function $T(\cdot)$ is increasing or decreasing. Most importantly, the perspective in this paper is to provide a full characterization of the equilibrium behavior under arbitrary model parameters and cost functions. Often there are multiple equilibria, in which case we fully characterize their number and relative position vis-à-vis each other. We also show how, in the fully general model, the entry or exit of a supplier impacts the equilibria. We summarize our main results by describing their application to the above quality competition model. (See also §5 for a summary in the context of the general model.)

To motivate the quality competition model, note that in most industries, component suppliers or original equipment manufacturers increasingly compete in terms of product attributes other than direct cost. Many goods have become commoditized, and gross profit margins have shrunk, making it increasingly difficult to compete on the basis of price differentials (alone). The supplier's *quality* and his *yield effectiveness and reliability*, as measured by the percentage of effectively produced units, rank among the most critical of the various dimensions along which competing firms differentiate themselves. The same applies to suppliers of consumer goods to large department stores, retail organizations, or government agencies (in the latter case, e.g., vaccines or medical devices). The yield characteristics include the possibility of complete disruptions due to natural causes (such as fires or hurricanes) and man-made breakdowns (e.g., sabotage or terrorist attacks), as well as bankruptcies.[2] Many companies have adopted a multisourcing strategy, splitting orders among competing suppliers so as to mitigate various supplier risks.

Our quality competition model considers an industry with N potential suppliers competing for the business of B buyers in a single sales season. To facilitate the exposition, we initially consider a single purchasing firm or agency. However, almost all of our results carry over to the general oligopsony case with an arbitrary number of buyers (see §EC.4 of the online appendix, provided in the e-companion).[3] The purchasing firm faces an uncertain demand volume, whereas each of the suppliers experiences a given random yield factor. In the face of the combined demand and supply risks, the buyer determines a total order size and its allocation among the potential suppliers, minimizing purchasing costs while ensuring that a shortfall is avoided with a given minimum probability. As will become apparent in the analysis, the key characteristic of the yield distribution is its coefficient of variation (CV), or the supplier's *reliability*, defined as the reciprocal of the squared coefficient of variation. We show that this reliability measure is bounded from below by a positive constant. This endogenously determined minimum value may sometimes need to be increased due to standards imposed externally by the buyers or a regulatory agency (see §4 for examples).

A supplier can improve his reliability by (i) increasing the *yield predictability* via a reduction of the standard deviation of the yield factor, (ii) increasing the *yield target* (i.e., the mean of the yield distribution), or (iii) improving *both* the yield target and its standard deviation. Consequently, we distinguish between three types of competition, which we refer to as (I) yield predictability competition (YPC), (II) yield target competition (YTC), and (III) simultaneous yield target and predictability competition (YSC). To focus on the impact yield targets and reliabilities have on the suppliers' competitive positions, we initially assume that they charge a uniform price. As many industries become increasingly commoditized, this

[2] Even before the 2008 financial crisis, Babich et al. (2007, p. 123) described the severity of this type of risk: "The combined volume of defaults in 2001 and 2002 exceeded the total volume of defaults in the United States over the previous 20 years." In the automobile industry, for example, many suppliers routinely incur losses, with Delphi, the largest supplier of automotive parts in the United States, residing in Chapter 11 until recently. Choi and Hartley (1996) documented that, in this industry, purchasing managers consider the financial solvability of the suppliers a major selection criterion, along with criteria like consistency and reliability.

[3] An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

assumption often applies. However, see §4.4 for a discussion of the general case with supplier-dependent prices. (Even though the suppliers' goods are perfect substitutes, more reliable suppliers may be able to charge more than others.)

The competition models are Stackelberg games in which the suppliers compete by making yield choices, and the purchasing firm follows by determining how much she wants to order from each supplier. Starting with the yield predictability competition model, we show that this model always has an equilibrium, as long as the aggregate of the suppliers' minimum reliability standards is in excess of a threshold value, given by a simple function of the permitted shortfall probability. (If this condition is violated, the buyer may not be able to satisfy her service constraint under certain yield choices of the suppliers; see below for a characterization of the equilibrium behavior in this case.) Under multiple equilibria, there exists one that is componentwise smallest and one that is componentwise largest. Among all equilibria, the former is the most preferred, and the latter the least preferred, by all suppliers, whereas the opposite applies to the buyer. Also, under convex manufacturing cost rates, for example, all equilibria are completely ordered; i.e., if one supplier adopts a higher reliability measure under one equilibrium as opposed to an alternative equilibrium, the same applies to all of his competitors. In our numerical studies, we have observed that multiple equilibria arise frequently, and the largest and smallest equilibria are often far apart. Assuming suppliers dynamically adjust their yield choices (e.g., as best responses to the competitors' choices), the industry's equilibrium depends heavily on its initial choices: for example, if all suppliers start out with low (high) yield reliabilities, we show that the industry adopts the *smallest* (*largest*) equilibrium. This suggests that there is great permanent value to adopting *short-term* incentives (e.g., the imposition of the above *minimum* reliability standards) for suppliers to invest in yield improvements. Because of the competitive dynamics, such short-term incentives sustain themselves in the long run.

Any equilibrium is characterized by a set of suppliers that operate at their minimum standard level and a complementary set that choose to go beyond their minimum. We show that when each firm's per-unit production cost grows convexly with its chosen yield reliability, the set of "minimum performance" suppliers is consecutive in a specific supplier index. This index depends on the supplier's minimum reliability standard and his marginal cost rate and profit margin when operating at this reliability level. Finally, under an additional condition, broadly satisfied, we derive a bound for the number of distinct equilibria. We show that improving the minimum standards may eliminate a low-performance equilibrium and drive the suppliers to one in which they very significantly *outperform* these standards. Thereafter, the high-performance equilibrium is often self-sustaining, even when the minimum standards are no longer enforced.

We show that, under both the smallest and the largest equilibria, all suppliers react to a sales price decrease by investing in a lower yield reliability. (All of the comparative statistics results, described below, likewise refer to the smallest and largest equilibria.) In settings where the buyer has the bargaining power to reduce the sales price, exercising this power has the (perhaps unintended), consequence of incentivizing *all* suppliers to reduce their reliability investments. This phenomenon has been documented, for example, in the vaccine supply industry (see §4).

When a single supplier is able to increase his yield target, i.e., the mean of his yield distribution, *all* suppliers increase their yield reliability. We also show that every new entrant to the market causes all incumbents to improve their yield reliability; conversely, every departure from the industry induces all remaining firms to reduce it. When the *mean* demand volume goes up, the suppliers react by *reducing* their yield reliability, but when the *standard deviation* of the demand volume goes up, they respond by *increasing* their yield reliability. More comprehensively, we show that suppliers find it in their competitive interest to respond to increased *volatility* of the buyer's demand volume (i.e., an increased CV of the demand distribution) by increasing their yield reliability; at the same time, they exploit increased *risk averseness* of the buyer (i.e., a lower tolerance for the shortfall probability) by *reducing* their equilibrium yield reliability so as to force the buyer to increase purchase orders.

In symmetric models, there exists a critical number of suppliers $N^0(\underline{x})$—which depends on the minimal standard $\underline{x}$—such that the equilibrium is unique (and larger than the minimum standard) if the number of suppliers N is in excess of $N^0(\underline{x})$. If N is smaller than this critical number of firms, the minimum standard $\underline{x}$ represents one equilibrium, possibly in conjunction with one or two symmetric equilibria in which all suppliers adopt a common higher reliability value.

We obtain similar characterizations of the equilibrium behavior for the other two competition models. The remainder of this paper is organized as follows: In §2, we provide an overview of the related literature. In §3, we characterize the equilibrium behavior in the above defined general class of competition models. Section 4 applies the results of §3 to the aforementioned quality competition models. Section 5 concludes the paper with a summary of important conclusions. All proofs are relegated to §EC.1 of the online appendix.

2. Literature Review

In §1, we surveyed the literature on competition under attraction models. In this section, we give a brief review of the literature directly relevant to the quality competition model.

The economics literature treats quality as a differentiating attribute of the product rather than its procurement process; see §EC.2 of the online appendix for a brief review. Very few papers analyze the impacts of uncertain yields in decentralized supply chains. The recent paper by Zhu et al. (2007, p. 421) comments in its opening paragraph, "Although researchers in operations management have long realized the importance of operations beyond the walls of a firm and explored various management issues for better coordination along supply chains, research on quality improvement has been largely limited to operations inside the walls of a firm." Likewise, in their survey paper, Tsay et al. (1999, p. 327) note that only a few models consider the *choice* of the quality level, and, if so, primarily "from the vantage point of a single organization contemplating how to design its internal practices in light of its own costs of quality." Zhu et al. (2007) analyze a model with a single supplier and a single buyer facing a deterministic demand process in which the buyer and the supplier sequentially decide to invest in an improvement of the yield characteristics of the supplier's production process. Babich et al. (2007) consider an industry with two suppliers and one buyer. Particularly motivated by the risk of suppliers' defaulting and therefore not being able to deliver on their orders, the authors assume that each supplier's random yield factor is a Bernoulli random variable, which is equal to zero, with a probability given by the firm's likelihood of default. The two firms compete by selecting a unit price.

Deo and Corbett (2009) assume that an arbitrary number of suppliers, offering a homogenous good, engage in Cournot competition, where the (common) per-unit price is a linear function of the total actual supply offered to the market. (A firm's actual supply is its intended production volume multiplied with a random yield factor, which is independently generated from a common distribution.) The firms compete by selecting their intended production volumes. Deo and Corbett (2009) use their model to explain the number of flu vaccine suppliers in the United States. Chick et al. (2008) consider a supply chain with a single buyer and a single supplier, whose random yield factor follows a general distribution. The buyer derives a benefit from its order, the magnitude of which grows as a concave function of the order size.[4] The authors characterize how the buyer and the supplier sequentially determine their order size and intended production volume. To our knowledge, ours is the first model to analyze a supply chain in which suppliers compete by targeting key characteristics of their uncertain yield processes.

3. Competition Under a General Attraction Model

Consider an industry with N firms, each selling a specific good or service at a given unit price. As in standard attraction models, each firm i's market share is proportional to its attraction value x_i, which can be selected within a given range $[\underline{x}_i, \bar{x}_i]$ $(i = 1, \ldots, N)$. The attraction values sometimes denote a single strategic choice, for example, the firm's advertising budget in Friedman (1958), or the firm's manufacturing yield reliability in the quality competition model of §4. In other settings, it is a function of several strategic choices: To give but a few examples, in the combined price and quality multinomial logit competition model in Anderson et al. (1992), the attraction value is an exponential function of a linear combination of the firm's price and quality level. In Bernstein and Federgruen (2004), the attraction value is a general function of the firm's price and service level, characterized by its fill rate. Kotler (1965) specifies the attraction value as a function of the firm's price and advertising and distribution budget, whereas Carpenter et al. (1988) model it as a function of a variety of marketing instruments.

As to the aggregate expected sales in the industry, we assume that it is determined by the aggregate attraction value. Thus,

$$s_i \stackrel{\text{def}}{=} \text{the expected sales volume of firm } i, \quad i = 1, \ldots, N, \quad \text{satisfies}$$

$$s_i = T\left(\sum_{j=1}^{N} x_j\right) \cdot \frac{x_i}{\sum_{j=1}^{N} x_j}, \quad i = 1, \ldots, N, \qquad (1)$$

where $T(\cdot)$ is assumed to be twice differentiable. See the discussion after Theorem 1 for some examples.

The cost incurred by a firm is proportional to its sales volume, with a cost rate that is nondecreasing in the firm's attraction value. This assumption is made in many competition papers with attraction models, for example, the quality competition model in Anderson et al. (1992, §7.5.2), or the price–service level competition model of Bernstein and Federgruen (2004),

[4] The authors are again particularly motivated by the flu vaccine supply problem, where this benefit function relates to the national cost savings due to a larger fraction of the population being vaccinated.

in which all costs are shown to be proportional with the expected sales volume at a rate that is increasing and convex in the chosen fill rate or attraction value. Gallego et al. (2006) (Besanko et al. 1998, So 2000) assume that the operational cost of each supplier is given by a convex [linear] function of his demand volume, the shape of which is *independent* of the strategic choices.[5] Thus, let

$w_i =$ the per-unit sales price of firm i, $i=1,\dots,N$;
$c_i(x_i) =$ the per-unit cost rate of firm i, a nondecreasing, twice-differentiable function of the attraction value x_i, $i=1,\dots,N$.

The lower bounds $\{\underline{x}_i\}$ are sometimes endogenous to the model. In other settings, these bounds are exogenously imposed by company policies or government regulations (see, e.g., §4). As to the upper bounds $\{\bar{x}_i\}$, if $\lim_{x_i\to\infty} c_i(x_i) > w_i$, only $x_i \le \inf\{x_i\colon c_i(x_i)\ge w_i\}$ represent relevant choices for firm i, to ensure a nonnegative profit value. In the analysis below, we therefore assume $\bar{x}_i \stackrel{\text{def}}{=} \inf\{x_i\colon c_i(x_i)\ge w_i\}$; all of our results are easily extended when these upper bounds need to be specified at lower levels.

Firm i's expected profit function is given by

$$\pi_i(\mathbf{x}) = (w_i - c_i(x_i))s_i = (w_i - c_i(x_i))\left(\frac{x_i}{\sum_{j=1}^N x_j}\right) T\left(\sum_{j=1}^N x_j\right).$$

With $x_{-i} = \sum_{j\ne i} x_j$, it is easier to employ

$$\tilde{\pi}_i(\mathbf{x}) \stackrel{\text{def}}{=} \log \pi_i(\mathbf{x}) = \log(w_i - c_i(x_i)) + \log x_i + \log\left[\frac{T(x_i + x_{-i})}{x_i + x_{-i}}\right], \tag{2}$$

with

$$\frac{\partial \tilde{\pi}_i}{\partial x_i} = G_i(x_i) - H(x_i + x_{-i}), \tag{3}$$

where

$$G_i(x_i) = \frac{-c_i'(x_i)}{w_i - c_i(x_i)} + \frac{1}{x_i} \quad \text{and} \quad H(R) = \left\{\log\left[\frac{R}{T(R)}\right]\right\}'. \tag{4}$$

Thus, the marginal profit increase of a firm due to a marginal increase in its attraction value depends on the competitors' strategic choices only via their sum, x_{-i}. The dependence is captured by the function $H(\sum_{j=1}^N x_j) = \partial \log(\sum_{j=1}^N x_j / T(\sum_{j=1}^N x_j))/\partial x_i$, the marginal *increase* in the logarithm of the aggregate required attraction value per unit sold, because of an increase in firm i's attraction value (or that of any firm, for that matter). When the total sales function $T(\cdot)$ is monotone—as is the case in all above examples—$H(R)$ represents the elasticity of the industry's aggregate attraction value with respect to its sales. The function $A(R) \stackrel{\text{def}}{=} R/T(R)$ denotes the required attraction intensity, i.e., the aggregate required attraction value per unit sold. The following property of this function plays a fundamental role in the equilibrium behavior of the competition model:

(A): The attraction intensity $A(R) = R/T(R)$ is log-concave in R.

Theorem 1 (Existence of Equilibria). *Assume* (A).

(a) *The competition game is* (*log-*)*supermodular and has at least one equilibrium. The set of equilibria is a lattice; in particular, there exists a componentwise smallest equilibrium* $\underline{\mathbf{x}}^*$ *and a componentwise largest equilibrium* $\bar{\mathbf{x}}^*$.

(b) *If the attraction intensity* $A(R)$ *is increasing, we have that among all equilibria,* $\underline{\mathbf{x}}^*$ [$\bar{\mathbf{x}}^*$] *maximizes* [*minimizes*] *the expected profit for all firms.*

Remark 1. Many models specify $T(\cdot)$ as a constant M. Other papers use the specification

$$T(R) = \frac{MR}{(R+C)}. \tag{5}$$

This total sales function arises when assuming a given population of M individuals, each of whom either purchase one unit from one of the firms or nothing at all, C represents the attraction value of the no-purchase option, and the likelihood of an individual choosing a specific option is proportional to its attraction value. Several marketing papers, including Kotler (1965), Karnani (1985), and Basuroy and Nguyen (1998), assume $T(R) = MR^\beta$ for some $0 \le \beta < 1$ to represent a market that increases with the aggregate attraction value, however, at a decreasing marginal rate. Although all of the above specifications use an increasing $T(\cdot)$ function, the endogenously derived total sales function in the quality competition model of §4 is, in fact, decreasing for reasons explained in the Introduction; see (12) for its closed-form expression.

We note that the attraction intensity function $A(R) = R/T(R)$ is both log-concave and increasing in all of the above examples. This is easily verified for the case where the aggregate sales function $T(R)$ is constant, or of the form (5) or a power function $T(R) = MR^\beta$ with $0 \le \beta < 1$. We refer to Lemma 1 below for a verification of these properties in the quality competition model of §4. Because the competition model is a supermodular game, both the smallest

[5] A natural generalization of our model would allow for *fixed* costs that are dependent on the chosen attraction values as well. However, this generalization significantly complicates the equilibrium analysis. (Among competition papers with attraction demand models, Karnani (1985) and Anderson et al. (1992, §7.5.3) consider fixed costs that depend on the strategic choices, however, assuming that the variable cost rates are completely independent of these.)

and largest equilibria can be computed by a simple tatônnement scheme (with $\underline{\mathbf{x}}$ and $\bar{\mathbf{x}}$ as the starting point, respectively), in which, in each iteration, each firm determines its best response to the competitors' choices.[6]

As will be shown in §4, multiple equilibria often arise. Whereas the shape of the aggregate sales function $T(R)$, via that of the associated attraction intensity function $A(R)$, determines whether the game is supermodular or not, and in particular whether a pure Nash equilibrium exists, additional information about the number of equilibria and their structure (beyond the lattice structure mentioned in Theorem 1(a)) depends on the properties of the G_i-functions, $i = 1, \ldots, N$.[7] The shape of the G_i-functions only depends on the shape of the cost-rate functions as well as the magnitude of the variable profit margins $\{w_i - c_i(x_i)\}$. Among these structural properties, the following are of particular importance: (i) Which of the firms adopt their minimally feasible attraction value, and which choose to invest in a larger value? (ii) When is the set of equilibria an ordered set, i.e., for any pair of equilibria $\mathbf{x}^*$ and $\mathbf{x}^{**}$, either $\mathbf{x}^* \leq \mathbf{x}^{**}$ or $\mathbf{x}^* \geq \mathbf{x}^{**}$? Assuming the attraction intensity function $A(R)$ is increasing, one implication of the set of the equilibria being ordered is that *all* equilibria can be *uniformly* ordered in terms of *each* of the firms' preferences: All firms are worse off as we move from one equilibrium to another with a larger attraction value for some, and hence for all firms. (Theorem 2 below shows that this situation arises, for example, whenever the cost-rate functions $c_i(\cdot)$ are convex.) Whereas the firms uniformly prefer equilibria with lower attraction values, often the consumer has the opposite preference ranking, generating industrial policy challenges (see §4 for a discussion of the latter).

We focus on the case where the G_i-functions are strictly decreasing with inverse functions $G_i^{-1}(\cdot)$, $i = 1, \ldots, N$. This property applies, for example, when the cost-rate functions $c_i(\cdot)$ are convex, because

$$G_i'(x_i) = \frac{-c_i''(x_i)[w_i - c_i(x_i)] - [c_i'(x_i)]^2}{[w_i - c_i'(x_i)]^2} - \frac{1}{x_i^2}. \quad (6)$$

See, however, Remark 2 below for a discussion of the case where all G_i-functions are increasing. Consider a starting point where all firms operate at their minimum attraction levels $\{\underline{x}_i\}$. Let $\underline{R}(k) \stackrel{\text{def}}{=} \sum_{i=1}^k \underline{x}_i$ denote the aggregate minimum attraction value of the first k firms. Let $\underline{S}$ denote the set of firms who would be worse off by making marginal improvements to their minimum level; i.e., by (3), $\underline{S} \stackrel{\text{def}}{=} \{i\colon (\partial \log \pi_i / \partial \underline{x}_i)(\underline{\mathbf{x}}) = G_i(\underline{x}_i) - H(\underline{R}(N)) \leq 0\}$. Thus, each firm i is characterized by an index $I_i \stackrel{\text{def}}{=} G_i(\underline{x}_i)$; note that this index value depends only on the firm's own cost-rate function, his minimum attraction value, and his sales price. Without loss of generality, number the suppliers in increasing order of their index values, i.e., $I_1 \leq I_2 \leq \cdots \leq I_N$. With this numbering, $\underline{S} = \{1, \ldots, |\underline{S}|\}$ and $|\underline{S}|$ is the highest indexed supplier whose index value I_i is below $H(\underline{R}(N))$; i.e., $|\underline{S}| = \max\{i\colon I_i \leq H(\underline{R}(N))\}$. For any equilibrium $\mathbf{x}^*$, let $\underline{S}(\mathbf{x}^*) \stackrel{\text{def}}{=} \{i\colon x_i^* = \underline{x}_i\}$ and $S^+(\mathbf{x}^*) \stackrel{\text{def}}{=} \{i\colon x_i^* > \underline{x}_i\}$ denote the set of suppliers that operate at and above their minimum reliability standards, respectively.

Theorem 2 (Characterization of the Set of Equilibria). *Assume* (A) *and all G_i-functions are decreasing.*

(a)(i) *For every equilibrium $\mathbf{x}^*$, there exists some $k^*(\mathbf{x}^*)$ $(0 \leq k^* \leq |\underline{S}|)$ such that $\underline{S}(\mathbf{x}^*) = \{1, \ldots, k^*\}$ and $S^+(\mathbf{x}^*) = \{k^*+1, \ldots, N\}$.*

(ii) *$\mathbf{x}^*$ is of the following form:*

$$x_i^* = \begin{cases} \underline{x}_i & i = 1, \ldots, k^*, \quad (7a) \\ G_i^{-1}(H(\underline{R}(k^*) + \rho)) & i = k^*+1, \ldots, N, \quad (7b) \end{cases}$$

where ρ is a root of the characteristic equation

$$\sum_{i=k^*+1}^{N} G_i^{-1}(H(\underline{R}(k^*) + \rho)) - \rho = 0. \quad (8)$$

(iii) *Any pair of equilibria $\mathbf{x}^*$ and $\mathbf{x}^{**}$ is completely ordered, i.e., either $\mathbf{x}^* \leq \mathbf{x}^{**}$ or $\mathbf{x}^{**} \leq \mathbf{x}^*$. Moreover, if $\mathbf{x}^* \leq \mathbf{x}^{**}$, $k^*(\mathbf{x}^*) \geq k^*(\mathbf{x}^{**})$, and, assuming the attraction intensity $A(R)$ is increasing, all firms are better off under $\mathbf{x}^*$ compared to $\mathbf{x}^{**}$.*

(b) *Assume the following condition applies.*

Condition (G): *For all $i = 1, \ldots, N$, $G_i(\cdot)$ is decreasing and $G_i^{-1} \circ H(\cdot)$ is strictly concave, so that* (8) *has at most two roots, $\underline{\rho}$ and $\bar{\rho}$, where $\underline{\rho} \leq \bar{\rho}$.*

(i) *For any $1 \leq k \leq |\underline{S}|$, there exists at most one equilibrium $\mathbf{x}^*$ such that $k^*(\mathbf{x}^*) = k$. Such an equilibrium $\mathbf{x}^*(k)$ satisfies* (7a) *and* (7b) *with $\rho = \bar{\rho}$.*

(ii) *There exist at most two interior equilibria $\mathbf{x}^*(0)$ and $\bar{\mathbf{x}}^*(0)$, with $x_i^*(0) = G_i^{-1}(H(\underline{\rho}))$ and $\bar{x}_i^*(0) = G_i^{-1}(H(\bar{\rho}))$, $i = 1, \ldots, N$.*

(iii) *There exist at most $|\underline{S}| + 2$ equilibria.*

The concavity condition (G) is satisfied, for example, when all G_i-functions are concave themselves and the $H(\cdot)$-function is convex. (Because $G_i(\cdot)$ is concave

[6] Topkis (1998) considers two variants of the tatônnement scheme: (i) In *simultaneous optimization*, in each iteration, all firms assume their competitors stay with their choices in the previous iteration; (ii) In round robin, in each iteration, one cycles through the N suppliers and each determines a best response to the most recently adopted choices of the competitors.

[7] Note that each firm's profit function, in general, fails to be (log-)concave, or even quasi-concave, so that supermodularity arises as an essential tool in establishing the existence of pure Nash equilibria.

and strictly decreasing, its inverse $G_i^{-1}(\cdot)$ is concave and strictly decreasing as well. $G_i^{-1} \circ H(\cdot)$ is therefore strictly concave as the composition of a concave and strictly decreasing function with a decreasing and strictly convex function.)

The following Theorem considers the special case where the model is symmetric, i.e., all firms have identical characteristics. In this case, let $c(\cdot) \stackrel{\text{def}}{=} c_1(\cdot) = \cdots = c_N(\cdot)$ and $G(\cdot) \stackrel{\text{def}}{=} G_1(\cdot) = \cdots = G_N(\cdot)$. Define $N^1(\underline{x}) \stackrel{\text{def}}{=} \min\{N \geq 2\colon\ H(N\underline{x}) < G(\underline{x})\} \leq \infty$. Also, define $\underline{x}^0$ as the unique root of $G(\cdot)$, which exists because $G(\cdot)$ is strictly decreasing, whereas $\lim_{x\downarrow 0} G(x) = \infty$ and $\lim_{x\uparrow \bar{x}} G(x) = -\infty$. Let $\mathbf{e} \stackrel{\text{def}}{=} (1, 1, \ldots, 1) \in \mathbb{R}^N$.

Theorem 3 (Symmetric Case). *Assume identical firms, condition* (A), *and* $G(\cdot)$ *decreasing.*

(a) *There exists at least one equilibrium. All equilibria are ordered and symmetric.*

(b) *Assume in addition that the attraction intensity function* $A(R) = R/T(R)$ *is increasing, and the minimum reliability standard* $\underline{x} \geq \underline{x}^0$. *The vector* $\underline{\mathbf{x}}$ *is the unique equilibrium, irrespective of the number of firms in the industry.*

(c) *Assume condition* (G). *There exists a number of firms* $N^0(\underline{x}) \leq N^1(\underline{x}) \leq \infty$ *such that the following applies*:

(i) *If* $N \geq N^0(\underline{x})$, *there exists a unique equilibrium* $\mathbf{x}^*$ *that is symmetric and interior and whose common component* x^* *is the larger* (*or unique*) *root of the characteristic equation*

$$\theta^{(N)}(x) \stackrel{\text{def}}{=} G^{-1} \circ H(Nx) - x = 0. \tag{9}$$

This unique equilibrium increases with every new entering firm.

(ii) *If* $N < N^0(\underline{x})$, *the set of equilibria consists of* $\underline{\mathbf{x}}$, *possibly in conjunction with one or two symmetric and interior equilibria,* $\mathbf{x}^* = x^*\mathbf{e}$ *and* $\mathbf{x}^{**} = x^{**}\mathbf{e}$, *with* x^* (x^{**}) *one of the* (*at most two*) *roots of the characteristic Equation* (9).

Thus, when $\underline{\mathbf{x}} > \underline{\mathbf{x}}^0$, the minimum attraction level is set at a high enough level that $\underline{\mathbf{x}}$ arises as the unique equilibrium, irrespective of the number of firms in the industry, as long as the function $A(R) = R/T(R)$ is increasing, a condition trivially satisfied in all of the above reviewed models. When $\underline{\mathbf{x}} \leq \underline{\mathbf{x}}^0$, the shape of the $G(\cdot)$ function impacts on the equilibrium behavior. However, under condition (G) and assuming $N^1(\underline{x}) < \infty$, a unique equilibrium is again guaranteed, as long as the number of competitors is sufficiently *large*, and under this unique equilibrium, all firms exceed the minimum standard and increase their attraction value as the competition becomes fiercer, i.e., as the number of firms grows.

Remark 2. We complete this section with a discussion of the case where, for each firm $i = 1, \ldots, N$, $G_i(x_i)$ is *increasing* on the feasible range $[\underline{x}_i, \bar{x}_i]$. This case cannot occur when, as hitherto assumed, the upper bound value $\bar{x}_i = \inf\{x_i\colon\ c_i(x_i) \geq w_i\}$. (It follows from (6) that $G_i'(x_i) < 0$ for x_i sufficiently close to $\bar{x}_i$.) However, uniformly increasing G_i-functions may arise when the upper bound values $\{\bar{x}_i\}$ are chosen at lower levels and the cost-rate functions $\{c_i(\cdot)\}$ are concave. Because, by condition (A), $H(\cdot)$ is decreasing, we have in this case that, for all $i = 1, \ldots, N$ and all $\mathbf{x}_{-i}$, $\partial\tilde{\pi}_i(x_i, x_{-i})/\partial x_j$ is an increasing function of x_i on the complete interval $[\underline{x}_i, \bar{x}_i]$. This implies that either $\underline{x}_i$ or $\bar{x}_i$ arises as each firm i's best response to any combination of attraction values chosen by its competitors. In particular, in any equilibrium $\mathbf{x}^*$, each firm positions itself either at the lower or at the upper bound of its feasible range.

4. Quality Competition Model

Consider an industry with N suppliers of a given product competing for the business of a single buyer in a specific sales season. (However see the online appendix, §EC.4, for a generalization of our results to allow for any number of buyers.) Each of the suppliers encounters an uncertain yield factor with a given, general, and supplier-dependent yield distribution. The buyer faces uncertain demand over the course of the season, with a Normal distribution. Her challenge is to select a set of suppliers as well as a total order quantity and its allocation among the selected suppliers so as to ensure that her demand is met with a given minimum probability, while minimizing procurement costs. (An alternative representation of the buyer's procurement problem, discussed in §4.5, involves explicit shortage and overage costs.)[8]

The suppliers compete by selecting key characteristics of their yield distributions, either their means, their standard deviations, or both.[9] Depending on

[8] The above representation assumes a single round of sales, without any recourse options. This assumption reflects many practical situations with long lead times. Nevertheless, its relaxation to allow for multiple procurement rounds would be valuable, although it would result in much more complex dynamic Stackelberg games. It is therefore beyond the scope of this paper.

[9] In the electronic design automation industry, manufacturers focus their competitive strategies on "design for yield." Similarly, Pisano (1996) documents that in the pharmaceutical and biotechnology industries, firms control the characteristics of their yield distributions by deciding how much time and effort to allocate to the product and process phases. Özer et al. (2007) describe how suppliers in the semiconductor industry strategize on how much time and effort to put into the design phase to improve their yield characteristics; they report a graph by Hitachi GST exhibiting the dependence of the yield characteristics with respect to the length of the design phase. Firms are also able to (partially) control their perceived reliability and estimated financial default probabilities by adopting an appropriate financial structure.

the source(s) behind his random yields, a supplier may improve the CV of his yield distribution by investing more time and effort into the design phase; by adopting appropriate technologies, materials, manufacturing, and logistical processes or a more secure financial structure; or improving his facilities' security. The supplier's choice has implications for his per-unit cost rate. We initially assume that the suppliers' prices are identical. This assumption often applies because many industries have become commoditized.[10] (In §4.4, we discuss the general case, where the suppliers differentiate themselves on the basis of their prices as well.) The suppliers operate in a make-to-order or purchase-to-order environment, which permits them to defer commitment to all material and other variable costs until the buyer's orders are received.

We also assume that the buyer only pays for delivered units that are useable, i.e., they satisfy the quality standards. This assumption is adopted by the majority of the literature on random yield supply systems.[11] The practitioner-oriented outsourcing literature (e.g., Brown and Wilson 2005) refers to this as *fixed* pricing schemes. (In §4.4, we analyze other settings, where the buyer is required to pay for all units ordered and entered into the production process, or where the buyer incurs *two* cost rates, one that applies to all units ordered and the other which is charged only for the useable ones.) Thus, let

$X_i =$ the random yield factor at supplier i's facility, with mean p_i, variance ς_i^2, CV $\gamma_i = \varsigma_i/p_i$, and reliability measure $x_i \stackrel{\text{def}}{=} \gamma_i^{-2} = p_i^2/\varsigma_i^2$, $i = 1, \ldots, N$;

$D =$ the uncertain demand during the season, which is Normally distributed with mean μ, variance σ^2, and CV $\gamma_D = \sigma/\mu$;

$c_i(p_i, \varsigma_i) =$ the expected cost for supplier i to procure an *effective* unit, a twice-differentiable function, with $\lim_{\varsigma_i \downarrow 0} c_i(p_i, \varsigma_i) = \lim_{p_i \uparrow 1} c_i(p_i, \varsigma_i) = +\infty$;

$w =$ the price charged to the buyer for every effectively delivered unit;

$\alpha =$ maximum permitted probability of a shortfall;

$U =$ a standard Normal random variable with cdf $\Phi(\cdot)$ and complementary cdf $\overline{\Phi}(\cdot)$;

$z_\alpha = \Phi^{-1}(1-\alpha)$;

$y_i(y_i^*) =$ the (optimal) order placed with supplier i, $i = 1, \ldots, N$;

$Y_E(Y_E^*) = \sum_{i=1}^N p_i y_i (\sum_{i=1}^N p_i y_i^*) =$ the (optimal) expected aggregate sales of the suppliers.

The general yield distributions allow for a positive probability mass at zero to reflect the possibility of a total supply disruption or breakdown, financial defaults, batch failures or acceptance sampling, or supplier delays resulting in untimely deliveries. The Normal distribution provides an adequate and frequently used specification of the demand distribution. We assume

$$\mu \geq z_\alpha \sigma, \tag{10}$$

ensuring that the likelihood of the demand volume D assuming negative values is no larger than the permitted shortfall probability α.[12]

The cost-rate functions $c_i(p_i, \varsigma_i)$ may be derived from an underlying more primitive description of the cost structure: For example, assume first that all of supplier i's labor and material costs are incurred for every *attempted* unit, whether ultimately resulting in an *effective* unit or not, and the cost per unit is given by $\bar{c}_i(p_i, \varsigma_i)$. The supplier's cost associated with an order of size y_i is then given by $\bar{c}_i(p_i, \varsigma_i) y_i = c_i(p_i, \varsigma_i)(p_i y_i)$, where $c_i(p_i, \varsigma_i) = \bar{c}_i(p_i, \varsigma_i)/p_i$ may be interpreted as the expected cost incurred for each *effective* unit that is procured. However, some cost components (e.g., packaging, warehousing, and shipping costs) may be incurred only for effective units that satisfy the required quality specifications. Assume these cost components amount to $\bar{c}_i^{(2)}(p_i, \varsigma_i)$ per (effective) unit. In this case, the total variable cost incurred by supplier i is given by $\bar{c}_i(p_i, \varsigma_i) y_i + \bar{c}_i^{(2)}(p_i, \varsigma_i)(p_i y_i) = c_i(p_i, \varsigma_i)(p_i y_i)$, where $c_i(p_i, \varsigma_i) = \bar{c}_i(p_i, \varsigma_i)/p_i + \bar{c}_i^{(2)}(p_i, \varsigma_i)$. Note that if $\lim_{\varsigma_i \downarrow 0} \bar{c}_i(p_i, \varsigma_i) = \lim_{p_i \uparrow 1} \bar{c}_i(p_i, \varsigma_i) = +\infty$, the same limiting behavior applies to the cost rates $c_i(p_i, \varsigma_i)$ as well.

An important assumption in our model is that the buyer determines all gross (production or purchasing) orders from the various suppliers. This assumption is shared with all of the literature on inventory systems with random yields discussed above. One might envision a setting where the buyer specifies a (maximum) purchase quantity of useable units from a supplier, who proceeds to determine a gross order quantity that optimally balances the *supplier*'s risk of overage and underage, vis-à-vis the desired purchase quantity. However, the supplier is often not in a position to target a specific supply of useable units, particularly when there is a significant likelihood of a complete disruption, i.e., when the yield distribution has a

[10] See also surveys like Choi and Hartley (1996, p. 333) for the automobile industry, concluding that "price is one of the least important selection items, regardless of the position on the supply chain."

[11] It may be implemented by initially charging for all produced units, and providing a rebate for units found to be defective at the buyer's site or because of external failures reported by the end consumer.

[12] A general distribution may be used to describe the demand volume. However, in the analysis below, its impact on the equilibrium choices of the suppliers is identical to those arising under a Normal distribution, with matching first and second moments.

positive mass at zero. (Recall that such yield distributions are required to model supply chain disruptions, financial defaults, batch failures, acceptance sampling; and untimely deliveries, among others. In such cases, the supplier fails to meet any desired purchase quantity with this probability, regardless of what gross order size he initiates.) Finally, the supplier can only commit himself to a given purchase quantity of good units if (full) inspection of all produced units takes place at the supplier's site. This is often impossible or impractical (see, for example, Baiman et al. 2000, Balachandran and Radhakrishnan 2005), particularly when failures occur during the distribution process to the buyer or failures can only be observed externally by the consumer during the consumption process. (The buyer is often committed to exchanging any failing unit for a good one, and her purchase quantities from the suppliers must adequately include an appropriate stock of replacement units.)

The buyer's service level constraint can be formulated as

$$\Pr\left(\sum_{i=1}^{N} X_i y_i \geq D\right) \geq 1-\alpha. \tag{11}$$

We assume that the end of the period inventory level ($\sum_{i=1}^{N} X_i y_i - D$) is Normal. This applies, of course, when all yield distributions are Normal themselves. In addition, it applies as a close approximation, when the yield distributions are nearly Normal or when there are at least four suppliers, even when these distributions have a fundamentally different shape. Diversification among four or more suppliers occurs, for example, in the vaccine industry, discussed below, as well as for distributors of food items (see Sheffi 2005, pp. 216–217, for a discussion of the banana industry). We refer to Federgruen and Yang (2008, 2009) for an extensive discussion, theoretical foundations, as well as numerical investigations of the adequacy of the Normal approximation.[13]

With unreliable suppliers, even the existence of a feasible procurement strategy is questionable, as simple examples exhibit. (For example, when all yield factors have a probability mass at zero, with $p = \Pr(X_i = 0)$, a feasible solution fails to exist if $p^N > \alpha$.) When the end-of-the-season inventory is (assumed to be) Normal, a set of suppliers is feasible if and only if its aggregate reliability measure ($\sum_{i=1}^{N} x_i$) is in excess of z_α^2, a simple function of the permitted shortfall probability α, only.

Lemma 1 (Characterization of the Optimal Expected Sales Volumes).

(a) *A feasible solution exists if and only if* $\sum_{i=1}^{N} x_i > z_\alpha^2$.

(b) *If* $\sum_{i=1}^{N} x_i > z_\alpha^2$, *the expected sales volumes are given by*

$$p_i y_i^* = \left(\frac{x_i}{\sum_{j=1}^{N} x_j}\right) Y_E^*,$$

where

$$Y_E^* = T\left(\sum_{i=1}^{N} x_i\right) \stackrel{\text{def}}{=} \mu\left(1 - \frac{z_\alpha^2}{\sum_{i=1}^{N} x_i}\right)^{-1} \cdot \left[1 + z_\alpha \sqrt{\frac{1}{\sum_{i=1}^{N} x_i} + \gamma_D^2\left(1 - \frac{z_\alpha^2}{\sum_{i=1}^{N} x_i}\right)}\right], \tag{12}$$

with $T(\cdot)$ *a decreasing function.*

(c) $H(R)$ *is a positive, strictly decreasing, and strictly convex function.*

We conclude that the expected sales volumes have the structure of the generalized attraction model in §3, with a *decreasing* total sales function $T(\cdot)$. Moreover, because $H(\cdot)$ is decreasing, the attraction intensity function $A(\cdot)$ is log-concave; i.e., assumption (A) is satisfied, by itself guaranteeing that any of the competition models below are (log-)supermodular, and hence have either a unique pure Nash equilibrium or a componentwise smallest and a componentwise largest Nash equilibrium (see Theorem 1). Moreover, because $T(\cdot)$ is decreasing, $A(R) = R/T(R)$ is increasing, guaranteeing that, in the case of multiple equilibria, all firms are uniformly better (worse) off under the componentwise smallest (largest) equilibrium, among all possible equilibria.

In the remainder, we analyze various types of competition between the suppliers. These are represented as Stackelberg games in the sense that suppliers first engage in a noncooperative game to select specific yield characteristics, followed by the buyer's decisions about how much to order from each. Our model assumes *symmetric* information among all parties concerned. In particular, the buyer knows the mean and standard deviation of each of the suppliers' yield factors on the basis of qualification processes, declared standards, or prior experience in earlier sales seasons.[14] As to the suppliers, we will show that each needs to know only the total reliability measure in the industry to determine his best response function. Thus, the information requirements for the various firms are limited, and in many applications, it

[13] Federgruen and Yang (2009) report, for example, on a numerical study of 80 instances, all with $N = 4$ suppliers and uniform yield distributions, in which the average error in the order sizes, based on a Normal approximation of the shortfall distribution compared with the optimal values, is on average 0.49%, with a maximum error of 4.31%.

[14] Several companies, including Chrysler, Eastman Kodak, Motorola, Texas Instruments, and Xerox, have jointly established the Consortium for Supplier Training programs, both to encourage suppliers to improve their yield characteristics and to monitor their progress (see Zhu et al. 2007).

is reasonable to assume that they are met. Nevertheless, future work should address settings where some of the distributional parameters may only be known imperfectly, thus calling for game theoretical models with *asymmetric* information.

4.1. Yield Predictability Competition Model

In this subsection, we model competition between the suppliers assuming they select a predictability level for their yield distribution. A predictability level can be targeted by adopting appropriate design and technology choices or quality control processes. Because competition is restricted to the choices of the standard deviations of the yield distributions, we assume here that the yield targets $\{p_i\}$ are exogenously given at levels $p_i = p_i^0 \geq 0$.

As far as the per-unit cost-rate functions $c_i(\cdot, \cdot)$ are concerned, in this model, we merely assume

$$\frac{\partial c_i(p_i, \varsigma_i)}{\partial \varsigma_i} < 0 \tag{13}$$

to reflect the fact that a less volatile yield distribution can only be achieved by adopting better materials, technologies, and quality processes, as well as higher investments in the design phase.

For any supplier i, selecting the yield standard deviation ς_i is equivalent to selecting the CV. value $\gamma_i = \varsigma_i / p_i^0$ or the supplier's reliability level, i.e., $x_i = \gamma_i^{-2} = (p_i^0)^2/\varsigma_i^2$. (We include the possibility of $p_i^0 = 0$, and hence $x_i = 0$, to enable the modeling of firms entering the industry.) To highlight the dependence of any supplier i's cost of manufacturing an effective unit on x_i, define

$$c_i^P(x_i \mid p_i^0) \stackrel{\text{def}}{=} \begin{cases} c_i(p_i^0, p_i^0/\sqrt{x_i}) & \text{if } p_i^0 > 0 \text{ and hence } x_i > 0, \\ c_i(0, 0) & \text{if } p_i^0 = x_i = 0, \end{cases}$$

which is clearly strictly increasing in x_i because $\partial c_i(p_i, \varsigma_i)/\partial \varsigma_i < 0$. We assume, in addition, that $c_i^R(x_i \mid p_i^0)$ is decreasing in p_i^0; i.e., it is less costly to procure an *effective* unit with a given reliability measure x_i when the supplier's expected yield is larger:

$$\frac{\partial c_i^R(x_i \mid p_i^0)}{\partial p_i^0} \leq 0. \tag{14}$$

In choosing a reliability level x_i, firm i faces a natural upper limit:

$$x_i \leq \bar{x}_i(p_i^0) \stackrel{\text{def}}{=} \begin{cases} \max\{x_i\colon c_i^P(x_i \mid p_i^0) \leq w\} < \infty & \text{if } p_i^0 > 0, \\ 0 & \text{if } p_i^0 = 0. \end{cases} \tag{15}$$

(The gross profit margin per effectively delivered unit for supplier i is given by $w - c_i^P(x_i \mid p_i^0)$, because $c_i^P(x_i \mid p_i^0)$ is continuously increasing and $\lim_{\varsigma_i \downarrow 0} c_i(p_i^0, \varsigma_i) = \lim_{x_i \uparrow \infty} c_i^P(x_i \mid p_i^0) = \infty$, $\bar{x}_i < \infty$, is well defined.) Note from (14) and (15) that $\bar{x}_i(p_i^0)$ is increasing in p_i^0 and w. Similarly, it is easily verified that $\varsigma_i \leq \sqrt{p_i^0(1-p_i^0)}$, i.e., the standard deviation of the yield distribution is maximally large when the support of this distribution is confined to the extreme values $X_i = 1$ and $X_i = 0$ (see Müller and Stoyan 2002, p. 57, Example 1.10.5). This upper bound implies

$$x_i \geq p_i^0/(1-p_i^0). \tag{16}$$

In addition, a lower bound $\underline{x}_i^e$, independent of the yield target p_i^0, may be imposed, either by the buyer or by external stipulations, such as government regulations.[15] Thus, let

$$\underline{x}_i(p_i^0) = \text{the minimum reliability level for supplier } i \stackrel{\text{def}}{=} \max(\underline{x}_i^e, p_i^0/(1-p_i^0)), \quad i = 1, \ldots, N. \tag{17}$$

Like the upper bound $\bar{x}_i(p_i^0)$, the lower bound $\underline{x}_i(p_i^0)$ is increasing in p_i^0 as well. Finally, to exclude situations where no feasible solution exists, under some of the suppliers' choices, we assume

$$\sum_{i=1}^{N} \underline{x}_i > z_\alpha^2. \tag{18}$$

(We revisit this assumption at the end of this subsection.) To simplify the notation, we generally suppress the dependence of the parameters with respect to p_i^0. As in (4), we define $G_i^P(x) \stackrel{\text{def}}{=} -c_i^{P\prime}(x)/(w - c_i^P(x)) + 1/x$, $i = 1, \ldots, N$. We conclude the following:

Theorem 4 (Yield Predicatability Competition Model). *Assume* (18).

(a) *Condition* (A) *holds and* $A(R)$ *is increasing in the yield predictability competition model*.

(b) *The results of Theorem* 1 *apply*.

(c) *Assuming the* $G_i^P(\cdot)$*-functions are decreasing, the results of Theorems* 2(a), 3(a), *and* 3(b) *apply*.

[15] For example, the Center for Disease Control and Prevention (CDC) purchases more than 50% of all routinely administered vaccines in the United States through the Vaccine Assistance Act (Section 317 of the Public Health Service Act, 1963) and the Vaccines For Children program, which was established in 1994. To enforce minimum reliability standards, the CDC together with the U.S. Food and Drug Administration established current Good Manufacturing Practices, which required many of the vaccine manufacturers to renovate their facilities (see Klein and Myers 2006). Many manufacturers institute qualification processes for which any potential supplier must compete to become part of the supplier base; see Gerling et al. (2002) for an example of such a qualification process prepared by semiconductor companies such as Motorola, Infineon Technologies, Phillips, and Texas Instruments. Terwiesch et al. (2001) describe the qualification processes in the data storage industry, and Özer et al. (2007) describe those employed by Hitachi. Also, many firms require suppliers to comply with qualification processes such as ISO 9000.

(d) *Assuming condition* (G) *holds, the results of Theorems* 2(b) *and* 3(c) *apply.*

Only the second part of condition (G) may be challenging to verify. As mentioned in §3, it holds when the G_i^P-functions are concave. In the YPC model, condition (G) also holds when the cost rate functions are linear, i.e., $c_i^P(x_i) = c_i x_i$ for all i, despite the fact that the G_i^P-functions fail to be concave. In this case,

$$(G_i^P)^{-1} \circ H(R) = \frac{1}{H(R)} + \frac{w}{2c_i} - \sqrt{\frac{w^2}{4c_i^2} + \frac{1}{H^2(R)}}. \quad (19)$$

See §EC.3 of the online appendix for the proof that $(G_i^P)^{-1} \circ H(\cdot)$ is concave for all i.

The phenomenon of multiple equilibria is not just a theoretical possibility. We have encountered many instances with either two or three equilibria, even when all of the procurement cost functions are linear. Moreover, the equilibria are often far apart. Assume that firms dynamically adjust their choices before converging to an equilibrium, perhaps by iteratively selecting best responses to the choices made by their competitors. The adopted equilibrium is then critically dependent on the starting conditions of the industry. As mentioned, because the game is supermodular, we know that the (componentwise) *smallest* equilibrium is adopted when the firms start off at or close to the vector of *minimum* reliability standards $\underline{\mathbf{x}}$, whereas the (componentwise) *largest* equilibrium arises when the firms start off at high levels of reliability close to the $\bar{\mathbf{x}}$-values. By Theorem 2(a), assuming convex cost-rate functions $c_i^P(\cdot)$ (or more generally, decreasing G_i^P-functions), the different equilibria are progressively more beneficial to *all* suppliers, as we move from the largest equilibrium to smaller ones; conversely, the buyer is progressively worse off because her total cost is given by wY_E^*, which by Lemma 1 is decreasing in $\sum_{i=1}^N x_i^*$.

The above observations have the following public policy implication: to ensure that the industry adopts a long-term equilibrium with relatively high reliability measures, it may pay to provide short-term incentives, via tax credits, subsidies, or the like, for the firms to invest in reliability improvements, thus inducing a high-performance equilibrium. Even if the incentives are eliminated after a while, firms are likely to readjust to a high-performance equilibrium, given their starting conditions. In addition, an increase of the minimum reliability standards $\underline{\mathbf{x}}$ may be used to induce a much larger impact on the industry's equilibrium behavior. This is demonstrated by the following example:

Example 1. Let $N = 3$, $\alpha = 0.1\%$, $p_1^0 = p_2^0 = p_3^0 = 0.77$, $\underline{x}_i = p_i^0/(1-p_i^0) = \underline{x} = 3.35$, $w = 1{,}000$, $\mu = 100$, and $\sigma = 20$; $c_i^P(x_i) = ix_i$, $i = 1, 2, 3$. Thus, the three suppliers differ only in terms of their manufacturing cost functions, and the minimum reliability values are the lowest possible choices for these standards, which arise under maximally unreliable suppliers (see (17)). There are three equilibria: $\mathbf{x}^{*L} = \underline{\mathbf{x}}$, $\mathbf{x}^{*M} = (4.93, 4.90, 4.88)$, and $\mathbf{x}^{*H} = (334.77, 205.12, 146.36)$. Thus, under the intermediate equilibrium $\mathbf{x}^{*M}$, the suppliers operate with only slightly higher than minimum reliability measures. At the same time, the reliability choices under the largest equilibrium $\mathbf{x}^{*H}$ are two orders of magnitude larger, with the suppliers reducing the CV of their yield distribution from $\gamma = (0.55, 0.55, 0.55)$ to $\gamma = (0.05, 0.07, 0.08)$. Three equilibria continue to prevail when the minimum standard $\underline{x}$ is increased by up to 45%. When it is increased by 46% (to $\underline{x} = 4.89$), $\mathbf{x}^{*L} = (4.89, 4.89, 4.89)$ and $\mathbf{x}^{*H} = (334.77, 205.12, 146.36)$ are the *only two* equilibria. Finally, when the minimum standard is increased by 47% or more (but by less than 4,272%), $\mathbf{x}^{*H} = (334.77, 205.12, 146.36)$ is the only equilibrium. In other words, by enforcing minimum standards only 47% higher than the bare minimum, the industry is induced to adopt reliability measures approximately 43.72 times the initial minimum values.

The following theorem shows that the smallest and the largest equilibria, $\mathbf{x}^{*L}$ and $\mathbf{x}^{*H}$, are a monotone function of a number of the model parameters; under these equilibria, any new entrant to the industry causes all firms to improve their reliability and the buyer to enjoy a cost reduction. (In the context of the YPC model, this generalizes Theorem 3(c) to general asymmetric industries.) As shown in Theorem 1(b), this pair of equilibria are especially important because, among all equilibria, *all* suppliers are best (worst) off under $\mathbf{x}^{*L}$ ($\mathbf{x}^{*H}$), whereas the opposite applies to the buyer.

Theorem 5 (Comparative Statics with Respect to the Equilibria).

(a) *All equilibria depend on the parameters of the demand distribution only via its CV* γ_D.

(b) $\mathbf{x}^{*L}$ *and* $\mathbf{x}^{*H}$ *are componentwise increasing in* w, γ_D, *and the maximum permitted shortfall probability* α. *In particular, for fixed* $\sigma(\mu)$, *the two equilibria are componentwise decreasing (increasing) in* $\mu(\sigma)$. *Under linear* $c_j^P(\cdot)$-*functions with* $c_j^P(x_j) = c_j x_j$, *these equilibria are decreasing in any of the marginal costs* $\{c_j\}$ *as well.*

(c) *Assume* $c_i^{P\prime}(x_i \mid p_i^0)$ *is decreasing in* p_i^0.

(i) $\mathbf{x}^{*L}$ *and* $\mathbf{x}^{*H}$ *are componentwise increasing in any of the firms' yield target* p_j^0, $j = 1, \ldots, N$.

(ii) *In both the smallest and largest equilibria, a new entrant (firm* $N+1$*) causes all incumbent firms to increase their reliability measures, resulting in a decrease of the buyer's cost.*

One implication of the first monotonicity result is that the buyer "pays" for a lower unit price by having to cope with less reliable yield processes at *all* suppliers. For example, in the vaccine supplier industry, the Center for Disease Control and Prevention (CDC) is chartered to pay as little for *established* vaccines as it is able to negotiate. Indeed, Table 2 in Klein and Myers (2006) shows that the federally contracted prices are, on average, 40% lower than the catalog price that applies to the private sector sales. The National Vaccine Advisory Committee has identified this fact and the resulting reduced profit margins as one of the primary reasons why suppliers have left the industry. In the United States, the number of vaccine manufacturers has dropped from 26 in 1967 to a mere 6 in 2006. Indeed, it follows from Theorem 5(c)(ii) that the exit of many suppliers causes the equilibrium reliability choices to go down by itself. However, not recognized in the committee's report is the fact that the highly reduced prices may well have eliminated incentives to improve yield reliabilities among those suppliers that chose to stay in the market. Thus, vaccine supplies may have become increasingly unreliable, not just because the number of suppliers decreased, but also because the federal contracts incentivized the remaining suppliers to adopt low levels of yield reliability, a phenomenon explained by Theorem 5(b). In contrast, if *new* vaccines become covered by the Vaccines for Children program, the CDC is required to purchase them at a price close to the supplier's catalog price. This policy has the unintended effect of incentivizing the industry to concentrate on new vaccines rather than to exploit the learning curve and improve the manufacturing processes for more established products.

To illustrate Theorem 5, consider in Example 1 a reduction of the price w. As long as $w \geq 75$, three equilibria continue to prevail. For example, when $w = 75$, the three equilibria are $\mathbf{x}^{*L} = \underline{\mathbf{x}}$, $\mathbf{x}^{*M} = (6.02, 5.47, 4.96)$, and $\mathbf{x}^{*H} = (20.95, 13.95, 10.29)$; each of the smallest and largest equilibria is componentwise smaller than its counterpart when $w = 1{,}000$. When $w = 74$, there are only *two* equilibria, i.e., $\mathbf{x}^{*L} = \underline{\mathbf{x}}$ and $\mathbf{x}^{*H} = (20.58, 13.73, 10.13)$, once again demonstrating the componentwise monotonicity of the smallest and largest equilibria. Finally, the theorem is silent about whether the equilibria are monotone in the minimum reliability standards $\underline{\mathbf{x}}$. Indeed, the following example shows that, for example, the largest equilibrium may fail to be monotone.

Example 2. Let $N = 8$, $\alpha = 0.1\%$, p_i be randomly generated on the interval $[0.550, 0.999]$ ($\mathbf{p} = [0.58, 0.9937, 0.81, 0.74, 0.78, 0.70, 0.74, 0.65]$), $\underline{x}_i = p_i^0/(1 - p_i^0)$, $w = 1{,}000$, $\mu = 100$, and $\sigma = 20$; $c_i^P(x_i) = ix_i$, $i = 1, \ldots, 8$. There is a unique equilibrium, which is an interior point: $\mathbf{x} = [399.29, 224.04, 155.06, 118.46, 95.81, 80.42, 69.29, 60.86]$. When the minimum standards $\mathbf{x}$ are increased by 43%, there is again a unique equilibrium, $\mathbf{x} = [378.38, 225.55, 152.41, 116.95, 94.84, 79.75, 68.79, 60.48]$. In this case, supplier 2's minimum standard is increased from 157.73 to 225.55, forcing him to increase his reliability to the new minimum standards. However, all other suppliers decrease their reliabilities.

We conclude this subsection with a discussion of what happens when condition (18) is violated but

$$\sum_{j=1}^{N} \bar{x}_j > z_\alpha^2, \tag{20}$$

i.e., under some but not all reliability measure vectors $\mathbf{x}$, the buyer is incapable of meeting her service constraint. Under such vectors $\mathbf{x}$, no orders will be placed, resulting in zero profit for each supplier. It is easily verified that *no* (pure) equilibrium exists under which the buyer is serviced if

$$\sum_{j=1}^{N} \bar{x}_j - \max_{1 \leq j \leq N} (\bar{x}_j - \underline{x}_j) \leq z_\alpha^2. \tag{21}$$

(Let $\bar{x}_i - \underline{x}_i = \max_{1 \leq j \leq N} (\bar{x}_j - \underline{x}_j) > 0$. Under any equilibrium $\mathbf{x}^*$ under which the buyer is serviced, $\sum_{j=1}^{N} x_j^* > z_\alpha^2$. If firm i decreases his reliability measure from x_i^* to $z_\alpha^2 + \epsilon - \sum_{j \neq i} x_j^* \geq z_\alpha^2 + \epsilon - \sum_{j \neq i} \bar{x}_j \geq \underline{x}_i + \epsilon$ by (21), the new total reliability value is $z_\alpha^2 + \epsilon$. It follows from Lemma 1 that as ϵ continues to decrease, the total order placed by the buyer goes to infinity, as does the order received by firm i, because his market share approaches $(z_\alpha^2 - \sum_{j \neq i} x_j^*)/z_\alpha^2 > \underline{x}_i/z_\alpha^2$. Finally, firm i's profit margin approaches $w - c_i^P(z_\alpha^2 - \sum_{j \neq i} x_j^*) > 0$. In other words, as $\epsilon \downarrow 0$, firm i's profit grows infinitely large, contradicting the assumption that $\mathbf{x}^*$ is an equilibrium.) Under (21), at least one of the suppliers is an *essential market maker* in the sense that, irrespective of his competitors' choices, this firm is capable of creating an infeasible situation for the buyer.

The most complex situation arises in the intermediate case where (18) is violated, i.e., some reliability choices result in an infeasible solution, but no single firm is an essential market maker, i.e., $\sum_{j=1}^{N} \bar{x}_j - \max_{1 \leq j \leq N} (\bar{x}_j - \underline{x}_j) > z_\alpha^2$. Assuming condition (G) holds, however, the following is a necessary and sufficient condition for a vector $\mathbf{x}^*$ to be an equilibrium ($\partial \tilde{\pi}_i(\cdot, x_{-i}^*)/\partial x_i$ has, under (G), at most two roots, so that $\tilde{\pi}_i(\cdot, x_{-i}^*)$ has at most two local maxima on $[\underline{x}_i, \bar{x}_i]$; call x_i' the second local maximum of firm i, if any):

(a) $\mathbf{x}^*$ is a *local* Nash equilibrium, i.e., every firm's choice is a *local* maximum of his profit function;
(b) $\tilde{\pi}_i(x_i', x_{-i}^*) \leq \tilde{\pi}_i(\mathbf{x}^*)$ for all $i = 1, \ldots, N$; and
(c) $\sum_{j=1}^{N} x_j^* - \max_{1 \leq j \leq N} (x_j^* - \underline{x}_j) > z_\alpha^2$.

To verify the sufficiency, note that under (c), no individual firm can create an infeasible situation by deviating. Moreover, by (a), x_i^* is a *local* maximum, and by (b), the only other possible local maximum has an inferior profit value. The necessity of each of the parts (a), (b), and (c) is immediate.

4.2. Yield Target Competition Model

Assume, now, that each supplier i selects his yield *target* p_i, under a given standard deviation ς_i^0 of its yield distribution. Analogous to (13), we again need a single condition with respect to the shape of the unit cost-rate functions $c_i(\cdot,\cdot)$. In fact, instead of requiring monotone behavior, all we require is that the function $c_i(\cdot,\varsigma_i^0)$ is quasi-convex, i.e., it is either increasing or it decreases first until a point $p_i^1(\varsigma_i^0)$ and is increasing thereafter.[16]

The YTC model may be specified as one in which each supplier selects his reliability level $x_i = (p_i)^2/(\varsigma_i^0)^2$, now with a cost-rate function $c_i^T(x_i \mid \varsigma_i^0) \stackrel{\text{def}}{=} c_i(\varsigma_i^0\sqrt{x_i}, \varsigma_i^0)$, an increasing function of x_i for $x_i \geq \underline{x}_i^1 \stackrel{\text{def}}{=} [p_i^1(\varsigma_i^0)]^2/(\varsigma_i^0)^2$, because $c_i(p_i,\varsigma_i^0)$ is increasing in p_i for $p_i \geq p_i^1(\varsigma_i^0)$. Clearly, no choice with $p_i \leq p_i^1(\varsigma_i^0)$ (i.e., $x_i \leq \underline{x}_i^0$) is sensible for firm i; after all, by moving from $p_i < p_i^0(\varsigma_i^0)$ to $p_i^0(\varsigma_i^0)$, firm i simultaneously improves its profit margin and market share. Therefore, we specify $\underline{x}_i$ as the larger of any externally specified minimum reliability standard $\underline{x}_i^e$, or $\underline{x}_i^0$. Analogous to (15), we specify $\bar{x}_i(\varsigma_i^0) \stackrel{\text{def}}{=} \max\{x_i \leq 1/\varsigma_i^0\colon c_i^T(x_i \mid \varsigma_i^0) \leq w\} < \infty$, because $\lim_{x_i\uparrow 1/\varsigma_i^0} c_i^T(x_i \mid \varsigma_i^0) = \lim_{p_i\uparrow 1} c_i(p_i,\varsigma_i^0) = \infty$.

It is again easily verified that $G_i^T(x) = -c_i^{T\prime}(x)/(w - c_i^T(x)) + 1/x$. Because, for each firm i, $G_i^T(\cdot)$ has the same properties as the function $G_i^P(\cdot)$ in §4.1, almost all of the results obtained from the YPC model continue to apply.

Theorem 6 (Yield Target Competition Model).
(a) *All results in Theorems* 4, 5(a), *and* 5(b) *continue to apply, assuming the G_i^T-functions have the same properties as the G_i^P-functions there.*
(b) *Assume $\underline{x}_i(\varsigma_i^0)$ is decreasing in ς_i^0, whereas $c_i^{T\prime}(x_i \mid \varsigma_i^0)$ is increasing in ς_i^0. The smallest and the largest equilibria are componentwise decreasing in the standard deviation of any of the yield distributions.*

The only result that we have not been able to extend for the general asymmetric model is the fact that, under both the smallest and the largest equilibria, all incumbent firms increase their reliability measures whenever a new supplier enters the industry (see Theorem 3(c)(i)). In the YTC model, a firm's entry cannot be modeled as if the firm were "present" both before and after entry, the firm's entry being characterized by a change in one or more parameters, with respect to which the equilibria can be shown to be monotone.

[16] This assumption is motivated by the following considerations: When the mean yield approaches zero, the effective variable cost per unit of product delivered often goes to infinity. Therefore, increasing the mean yield often decreases the variable cost at the beginning. When the mean yield is sufficiently high, any improvement in the yield target will result in a higher cost rate.

4.3. Simultaneous Yield Target and Predictability Competition Model

Consider now a setting where each firm i is capable of selecting both the mean and the standard deviation of his yield distribution, or equivalently, both his yield target and reliability measure (p_i, x_i). Analogous to (2), the profit functions can be written as $\tilde{\pi}_i(\mathbf{p},\mathbf{x}) \stackrel{\text{def}}{=} \log\pi_i(\mathbf{p},\mathbf{x}) = \log(w - c_i(p_i, p_i/\sqrt{x_i})) + \log x_i - \log(x_i + x_{-i}) + \log T(x_i + x_{-i})$, $i = 1,\ldots,N$. Observe, for any given firm i, that whereas the choice of x_i affects not just its own profit function but those of all of his competitors, the choice of p_i—given a value for x_i—only impacts this firm's *own* profit. To allow for a simple analysis, we assume beyond the monotonicity of $c_i(\cdot,\cdot)$ in its second argument, and the quasi-convexity in its first, that

$$c_{i1}(\cdot,\cdot) + \frac{c_{i2}(\cdot,\cdot)}{\sqrt{x_i}} \leq 0, \tag{22}$$

where c_{i1} and c_{i2} refer to the partial derivatives of the c_i-function with respect to their first and second arguments. (This additional condition is trivially satisfied on the range where the cost-rate functions decrease with the target level; on the other range, it requires that the cost reduction due to an increase in the yield standard deviation dominates the cost increase due to the improved yield target.) For a given choice of x_i, in view of (22), it is optimal to select as high as possible a yield target p_i, i.e., $p_i = x_i/(1+x_i)$ (see (16)). It is therefore possible to reduce the two-dimensional competition model to an equivalent one-dimensional competition model, with profit functions

$$\begin{aligned}\tilde{\pi}_i^S(\mathbf{x}) &\stackrel{\text{def}}{=} \max_{0\leq p_i\leq x_i/(x_i+1)} \tilde{\pi}_i(\mathbf{p},\mathbf{x})\\ &= \log(w - c_i^S(x_i)) + \log x_i - \log(x_i + x_{-i})\\ &\quad + \log T(x_i + x_{-i}),\end{aligned} \tag{23}$$

where $c_i^S(x_i) \stackrel{\text{def}}{=} c_i(x_i/(x_i+1), \sqrt{x_i}/(x_i+1))$. Without loss of practical generality, $p_i \geq 0.5$, i.e., $x_i \geq 1$. It is easily verified that the function $c_i^S(\cdot)$ is again increasing. (When $x_i \geq 1$ increases, the first argument of the $c_i(\cdot,\cdot)$ function increases, whereas its second argument $1/(1/\sqrt{x_i} + \sqrt{x_i})$ is decreasing in x_i because its numerator is increasing in x_i; the monotonicity of $c_i^S(\cdot)$ now follows from $\partial c_i/\partial p_i \geq 0$ and $\partial c_i/\partial\varsigma_i \leq 0$.)

The YSC model is, therefore, another special case of the general model of §3, with $G_i^S(x_i) = -c_i^{S\prime}(x_i)/(w - c_i^S(x_i)) + 1/x$.

Theorem 7 (Simultaneous Yield Target and Predictability Competition Model). *All of the results in Theorems* 4, 5(a), *and* 5(b) *continue to apply, assuming the functions $G_i^S(\cdot)$ have the same properties as the functions $G_i^P(\cdot)$ there.*

4.4. Competition Under Price Differentiation

Hitherto we have assumed that suppliers differentiate themselves only in terms of their yield distributions. In this case, each supplier i achieves a positive market share, irrespective of its reliability, which (in terms of billable units) is strictly proportional to its reliability measure x_i (see Lemma 1(b)). More generally, firms may, in addition, charge different prices, with w_i the price of supplier i. (Renumber the suppliers such that $w_1 \leq \cdots \leq w_N$.) Alternatively, the buyer may be required to pay for all units ordered and entered into the production process, or she may incur *two* cost rates, w_i^o and w_i^e, one that applies to all units ordered and the other which is charged only for the useable ones.[17] It is easily verified that this setting is equivalent to one in which each supplier i charges a single cost rate $w_i = w_i^e + w_i^o/p_i$ for each useable unit delivered. Under price differentiation, only the suppliers with a sufficiently low price are patronized by the buyer (see Dada et al. 2007, Federgruen and Yang 2009).

Although the suppliers' sales functions can still be computed efficiently (by Algorithm (SCM) in Federgruen and Yang 2009), they are not available in closed form. Example EC.1 in §EC.5 of the online appendix shows, in fact, that the profit functions in the YPC model may fail to be supermodular. Moreover, the model may in some instances fail to have a (pure) Nash equilibrium, albeit that in our investigations, to date, this only occurs under unrealistically large supplier markups.

4.5. Procurement Decisions Based on Costs of Over- and Understockage

As a final variant of our model, assume the buyer determines her procurement decisions on the basis of a traditional trade-off analysis of the *cost* of over- and understockage, rather than the service-level constraint (11). Thus, let h denote the carrying cost of any unsold unit at the end of the season, and b the cost of any unmet demand. Federgruen and Yang (2009) have shown that under identical supplier prices, all suppliers attain a positive market share, with the share for supplier i again given by $x_i/(\sum_{j=1}^N x_j)$. Also, the buyer's optimal cost under a given targeted effective supply Y_E is

$$\Psi^T(Y_E) = wY_E + h(Y_E - \mu) + (b+h) \cdot \int_{Y_E}^{\infty} \bar{\Phi}\left(\frac{u-\mu}{\sqrt{\sigma^2 + Y_E^2/(\sum_{i-1}^N x_i)}}\right) du,$$

which depends on the suppliers' reliability levels only via the single measure, $R = \sum_{i=1}^N x_i$, the total reliability measure in the industry. Federgruen and Yang (2009, Theorems 4.1 and 4.3) also show that Ψ^T is a differentiable, strictly convex function with a unique minimum Y_E^*. Thus, as in our basic model, the optimal effective total order Y_E^* depends on the vector $\mathbf{x}$ only via $\sum_{i=1}^N x_i$, i.e., $Y_E^* = T^C(\sum_{i=1}^N x_i)$.[18] It follows that the logarithms of the profit functions $\{\tilde{\pi}_i\}$ continue to be of the form (2), so that, like the other models, this variant of the quality competition model is a special case of the general model of §3 as well. Thus, the full characterization of the equilibrium behavior in §3 continues to apply, provided that one can show that the attraction intensity function $A^C(R) \stackrel{\text{def}}{=} R/T^C(R)$ is log-concave and increasing. Unlike (12), the function $T^C(\cdot)$ is no longer available in closed form. The first and second derivatives of $R/T^C(R)$ *are* available in closed form via the implicit function theorem, but their expressions are very complex. Although we have not been able to prove log-concavity and monotonicity of the attraction intensity function $A^C(\cdot)$, these properties have consistently been verified in an extensive numerical exploration. In other words, the structural properties of the equilibrium behavior of the model are the same, regardless of whether the buyer determines her procurement decisions on the basis of a service-level constraint or to minimize purchase, inventory carrying, and shortage costs.

5. Conclusions

We have characterized the equilibrium behavior in a broad class of competition models in which the competing firms' market shares are given by an attraction model, and the aggregate sales in the industry depend on the aggregate attraction value according to a general function. Each firm's revenues and costs are proportional with its expected sales volume, with a cost rate that depends on the firm's chosen attraction value according to an arbitrary increasing function. We have shown that most existing competition papers with attraction models in the economics, marketing, and operations literature can be viewed as special cases of this general model. The general model also includes a new series of quality competition models among suppliers with uncertain yield characteristics, developed in the previous section. Unlike existing applications, in these competition models, the total sales function is decreasing with the aggregate attraction value.

We have shown that the general competition model can be guaranteed to be (log-)supermodular based on

[17] The outsourcing literature calls these *variable* pricing schemes, which may incorporate any desired level of risk sharing between the supplier and the buyer.

[18] $T^C(\cdot)$ is, in general, decreasing, like its counterpart in the service-level-based model. However, the function may be locally increasing. Figure 2(a) in Federgruen and Yang (2009) shows that aggregate sales may sometimes increase as $(\sum_{i=1}^N x_i) \uparrow \infty$, i.e., as one approaches full reliability.

a single property of the so-called attraction intensity function, defined as the ratio of the aggregate attraction value and aggregate sales. The required property is the log-concavity of this attraction intensity function, which is satisfied in all of the reviewed applications of the general model. Log-supermodularity of the competition model implies that it has a pure Nash equilibrium and, in the case of multiple equilibria, a componentwise smallest and a componentwise largest equilibrium. Assuming the attraction intensity function is increasing (as well as log-concave), *all* firms are better (worse) off under the former (latter) equilibrium, among all possible equilibria. Additional structural properties of the set of equilibria and the number of possible equilibria depend only on the shape of the G_i-functions, which in turn depends only on the the shape of the cost-rate functions and the expected revenue of each unit sold. For example, if each firm's cost rate depends convexly on its attraction value, and multiple equilibria prevail, these are completely ordered with respect to each other: when comparing a pair of equilibria, either *all* firms offer a higher attraction value under the second equilibrium or *all* of their measures are lower. (Assuming again an increasing attraction intensity function, *all* firms are progressively better off when one moves from one equilibrium to another equilibrium with lower attraction values.) In this case, we have developed a firm index such that, in any equilibrium, the firms with the k lowest index values, for some $1 \leq k \leq N$, choose to adopt the minimum attraction value, whereas all other firms compete with higher values. (This index only depends on the firm's minimum attraction value, and the corresponding unit cost rate and its derivative.) Finally, under a slightly stronger condition (G), we have derived an upper bound for the number of equilibria.

Applying our results to the quality competition models, we have shown that multiple equilibria arise frequently, and they are often far apart. Assuming a simple dynamic adjustment process, we have shown that which of the equilibria is adopted may be influenced by the starting conditions of the industry. This suggests that temporary incentives to invest in high-quality processes result in long-term adherence to higher-quality equilibria, even after the incentives are dropped. Selection of the minimum standards provides a second mechanism to induce significant reliability improvements: examples show how an increase of the minimum standards by less than 50% can induce all competitors to reduce the CV value of their yield factor by a factor of 10. Reducing barriers to industry entry provides a third such mechanism: any new entrant causes both the smallest and largest equilibria to go up componentwise. Finally, buyers may face a low reliability equilibrium and end up incurring higher total costs when forcing their suppliers to accept low sales prices. Also, we have systematically exhibited how the (smallest and largest) equilibria depend on the various model parameters.

6. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

References

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

Babich, V., A. N. Burnetas, P. H. Ritchen. 2007. Competition and diversification effects in supply chains with supplier default risk. *Manufacturing Service Oper. Management* **9**(2) 123–146.

Baiman, S., P. E. Fischer, M. V. Rajan. 2000. Information, contracting, and quality costs. *Management Sci.* **46**(6) 776–789.

Balachandran, K. R., S. Radhakrishnan. 2005. Quality implications of warranties in a supply chain. *Management Sci.* **51**(8) 1266–1277.

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

Bell, D. E., R. L. Keeney, J. D. C. Little. 1975. A market share theorem. *J. Marketing Res.* **12**(2) 136–141.

Bernstein, F., A. Federgruen. 2004. A general equilibrium model for industries with price and service competition. *Oper. Res.* **52**(6) 868–886.

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

Brown, D., S. Wilson. 2005. *The Black Book of Outsourcing: How to Manage the Changes, Challenges and Opportunities*. John Wiley & Sons, Hoboken, NJ.

Carpenter, G. S., L. G. Cooper, D. M. Hanssens, D. F. Midgley. 1988. Modeling asymmetric competition. *Marketing Sci.* **7**(4) 393–412.

Chick, S. E., H. Mamani, D. Simchi-Leviz. 2008. Supply chain coordination and influenza vaccination. *Oper. Res.* **56**(6) 1493–1506.

Choi, T. Y., J. L. Hartley. 1996. An exploration of supplier selection practices across the supply chain. *J. Oper. Management* **14** 333–343.

Cooper, L. G. 1993. Market-share models. J. Eliashberg, G. L. Lilien, eds. *Handbooks in Operations Research and Management Science*, Vol. 5, Chap. 6. Elsevier Science Publishers B.V., Amsterdam 259–314.

Dada, M., N. C. Petruzzi, L. B. Schwarz. 2007. A newsvendor's procurement problem when suppliers are unreliable. *Manufacturing Service Oper. Management* **9**(1) 9–32.

Deo, S., C. J. Corbett. 2009. Cournot competition under yield uncertainty: The case of the U.S. influenza vaccine market. *Manufacturing Service Oper. Management* **11**(4) 563–576.

Federgruen, A., N. Yang. 2008. Selecting a portfolio of suppliers under demand and supply risks. *Oper. Res.* **56**(4) 916–936.

Federgruen, A., N. Yang. 2009. Optimal supply diversification under general supply risks. *Oper. Res.*, ePub ahead of print June 25, http://or.journal.informs.org/cgi/content/abstract/opre.1080.0667v1.

Friedman, L. 1958. Game-theory models in the allocation of advertising expenditures. *Oper. Res.* **6**(5) 699–709.

Gallego, G., W. T. Huh, W. Kang, R. Phillips. 2006. Price competition with the attraction demand model: Existence of unique equilibrium and its stability. *Manufacturing Service Oper. Management* **8**(4) 359–375.

Gerling, W. H., A. Preussger, F. W. Wulfert. 2002. Reliability qualification of semiconductor devices based on physics-of-failure and risk and opportunity assessment. *Quality Reliability Engrg. Internat.* **18**(2) 81–98.

Karnani, A. 1985. Strategic implications of market share attraction models. *Management Sci.* **31**(5) 536–547.

Klein, J. O., M. G. Myers. 2006. Strengthening the supply of routinely administrered vaccines in the United States: Problems and proposed solutions. *Clinical Infectious Diseases* **42**(Supplement 3) S97–S103.

Kotler, P. 1965. Competitive strategies for new product marketing over the life cycle. *Management Sci.* **12**(4) 104–119.

Lilien, G. L., P. Kotler, K. S. Moorthy. 1992. *Marketing Models*. Prentice Hall, Englewood Cliffs, NJ.

Müller, A., D. Stoyan. 2002. *Comparison Methods for Stochastic Models and Risks*. John Wiley & Sons, Chichester, UK.

Özer, Ö, O. Uncu, D. Cheng. 2007. An integrated framework to optimize time-to-market and production decisions for Hitachi Global Storage Technologies, Working paper, University of Texas at Dallas, Richardson.

Pisano, G. P. 1996. Learning-before-doing in the development of new process technology. *Res. Policy* **25** 1097–1119.

Sheffi, Y. 2005. *The Resilient Enterprise: Overcoming Vulnerability for Competitive Advantage*. MIT Press, Cambridge, MA.

So, K. C. 2000. Price and time competition for service delivery. *Manufacturing Service Oper. Management* **2**(4) 392–409.

Terwiesch, C., R. E. Bohn, K. S. Chea. 2001. International product transfer and production ramp up: A case study from the data storage industry. *R&D Management* **31**(4) 435–451.

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

Tsay, A. A., S. Nahmias, N. Agrawal. 1999. Modeling supply chain contracts: A review. S. Tayur, R. Ganeshan, M. Magazine, eds. *Quantitative Models for Supply Chain Management*, Chap. 10. Kluwer Academic Publishers, Boston, 299–336.

Zhu, K., R. Q. Zhang, F. Tsung. 2007. Pushing quality improvement along supply chains. *Management Sci.* **53**(3) 421–436.