Operations Research
Vol. 59, No. 4, July–August 2011, pp. 1033–1039
ISSN 0030-364X | EISSN 1526-5463 | 11 | 5904 | 1033
http://dx.doi.org/10.1287/opre.1110.0935

TECHNICAL NOTE

Procurement Strategies with Unreliable Suppliers

Awi Federgruen
Graduate School of Business, Columbia University, New York, New York 10027, af7@columbia.edu

Nan Yang
Olin Business School, Washington University in St. Louis, St. Louis, Missouri 63130, yangn@wustl.edu

We propose and analyze a general periodic-review model in which the firm has access to a set of potential suppliers, each with specific yield and price characteristics. Assuming that unsatisfied demand is backlogged, the firm incurs three types of costs: (i) procurement costs, (ii) inventory-carrying costs for units carried over from one period to the next, and (iii) backlogging costs. A procurement strategy requires the specification, in each period, of (i) the *set* of suppliers to be retained, (ii) their respective *shares* in this period's replenishments, as well as (iii) the traditional *aggregate* order placed (among the various suppliers).

We show how the optimal procurement strategy can be obtained with an efficient algorithm. A base-stock policy is no longer optimal, but in each period there exists a maximum inventory level, such that orders are placed if and only if the starting inventory is below this threshold. In each period it is optimal to retain a given number of suppliers that are cheapest in terms of that period's effective cost rates, i.e., the expected cost per usable unit. The optimal *number* of suppliers to be retained in a given period depends on all current and future parameters and distributions, but this dependence can be aggregated into a single so-called benchmark cost measure. Under Normal yield and demand distributions, the suppliers' market shares are determined by a single aggregate score, itself the product of a simple reliability score and a cost score.

Subject classifications: supplier selections; unreliable yields; multiperiod procurement strategies.
Area of review: Manufacturing, Service, and Supply Chain Operations.
History: Received November 2008; revisions received April 2009, July 2009; accepted August 2009.

1. Introduction and Summary

Fifty years of multiperiod inventory models have focused on the challenge of managing demand risks, assuming that all orders are filled completely. In practice, firms often deal with less than perfectly reliable suppliers who deliver only a random fraction of any given order. This type of supply risk arises because of uncertain yields, disruptions such as fires, hurricanes, strikes, sabotages and terrorist attacks, or failure of a replenishment batch to satisfy quality standards.

One strategy for dealing with such supply risks is to spread orders over multiple suppliers.[1] Toyota, for example, started to seek multiple suppliers after a fire at its almost-exclusive provider (Aisin) caused its assembly plants to shut down in 1997; see Treece (1997). Suppliers may be willing to assume part of the supply risks, but only if orders are placed well in advance, allowing them to engage in multiple production rounds and draw down inventory pools. However, purchasing firms are often unable to commit their orders with such extensive lead times. Hewlett-Packard (HP) launched, in 2000, a Procurement Risk Measurement framework that "enables the simultaneous measurement and management of multiperiod and correlated demand, cost, and availability uncertainties"; see Nagali et al. (2008). Orders are split among different numbers of suppliers, for different components. To allow for a beneficial sharing of yield risks with its suppliers, HP often engages in flexible quantity agreements, whereby it commits itself to minimum order quantities, the full delivery of which is guaranteed by the suppliers; on top of these minimum quantities, HP places short-term orders where it remains exposed to the yield uncertainties.[2] We refer to Federgruen and Yang (2008a) for several other (e.g., the vaccine, cell phone and oil refinery) industries in which multisourcing is essential to mitigate supply risks.

In this paper, we analyze a general periodic-review model, in which the firm has access to a set of potential suppliers, each with specific yield and price characteristics. Assuming that unsatisfied demand is backlogged, the firm incurs, as in standard inventory models, (i) procurement costs, (ii) inventory carrying costs for units carried over from one period to the next, and (iii) backlogging costs. However, in contrast to standard inventory models, a procurement strategy requires the specification, in each period, of (i) the *set* of suppliers to be retained, (ii) the *aggregate* order to be placed, and (iii) the suppliers' respective *shares* in this order.

The model is a generalization of the standard periodic-review model with linear procurement costs, introduced by Arrow et al. (1951) and Dvoretzky et al. (1952) and analyzed by Karlin (1958); see Zipkin (2000) for a detailed review. This standard model assumes a single, fully reliable supplier who delivers any given order either immediately or after a fixed lead time. The model also generalizes

Henig and Gerchak (1990), who assume that the single supplier faces a random yield factor; see (Zipkin 2000, §9.4.8). More specifically, we consider a planning horizon of T periods, in which the firm faces a given sequence of independent, but possibly nonstationary, continuously distributed random demands. Suppliers are differentiated by their time-dependent prices and (general) yield distributions. The firm either pays for all ordered units, regardless of whether they are delivered as usable units or not, or it is charged only for the usable ones. Most generally, the direct cost consequences of yield risks may be shared between the firm and the suppliers, in that a base cost rate is charged for *every* ordered unit and an additional charge only for those that are usable. The firm's holding and backlogging costs are proportional with the end of the period's inventory levels and backlog sizes, respectively. (More generally, all of our results continue to apply when these costs are given by convex functions of the inventory levels.) We initially assume that (the usable parts of) orders placed in any period become available in time to meet that period's demand. However, as discussed in §4, all of our results can easily be extended to allow for a fixed lead time, as long as the actual yield realizations of a given period's orders are revealed to the firm by the beginning of the next period.

We develop an efficient algorithm that identifies the optimal procurement strategy. Most importantly, we derive many structural results that generate important managerial insights. Here we summarize the main results. In contrast to the classical model with a single fully reliable supplier, it is no longer optimal to use a base-stock policy, i.e., to order up to a given base-stock level whenever the starting inventory is below this level. However, as in the classical model, there continues to be in each period, a "maximum ordering inventory level," such that orders are placed if and only if the starting inventory is below this level. (The maximum ordering inventory level can be obtained as the unique root of an analytically available increasing function.)

We show that in any given period, the set of retained suppliers is consecutive in the effective cost rates the suppliers charge in that period. (The effective cost rate is the expected procurement cost the firm incurs per useable unit.) In other words, in a given period, it is optimal to retain the k^* suppliers that are cheapest in terms of *that* period's effective cost rates, for some $1 \leqslant k^* \leqslant N$. The degree of supplier diversification, i.e., k^*, does, however, depend on the suppliers' yield characteristics and the demand distribution of this period, as well as all cost parameters and yield and demand distributions pertaining to future periods. Interestingly, these various parameters and distributions impact on k^* via a single aggregate measure, which we refer to as the benchmark cost rate. This benchmark cost rate represents the expected value of the total cost saving, associated with a marginal effective unit, delivered—for free and outside of the normal procurement process—beyond those arising from the optimal set of orders; here, the total cost saving relates to current holding and backlogging costs as well as all future costs. The benchmark cost rate is decreasing in the starting inventory. This implies, in particular, that the number of suppliers, optimally retained in each period, decreases with the starting inventory.

As to what drives the suppliers' *market shares* in any period, clear insights can be obtained when all yield and demand distributions are Normal.[3] Here, a supplier's market share is given by the relative value of a specific supplier score, which is the product of a reliability and a cost score: the former is this period's reciprocal of the squared coefficient of variation of the supplier's yield distribution, and the latter is given by the amount by which the current supplier's effective cost rate falls below the benchmark cost rate. The market share of each selected supplier is given by his overall score relative to the sum of the suppliers' scores.[4] As with the supplier set itself, the market shares in any given period depend on all future costs, yield and demand distributions only via a single measure, i.e., the benchmark cost rate. We also make systematic comparisons with the classical inventory model, in which only demand risks prevail, focusing on optimal safety stocks and order-up-to levels, and their dependence on the starting inventory level.

The remainder of this paper is organized as follows: §2 provides a review of the relevant literature. Our results are obtained in §3, whereas §4 discusses how the base model and its results can be extended to account for fixed supplier costs, positive lead times, capacity limits, and correlated yield and demand distributions. All proofs are deferred to the electronic companion to this paper, which is available as part of the online version at http://or.journal.informs.org/.

2. Literature Review

Karlin (1958) already recognized the need to address settings where replenishment quantities are random. In §§4–8 of his paper, he considered the case where every period's order quantity, from a single supplier, is exogenously given, but the actual supply depends on a random yield factor. Subsequently, a large literature on inventory systems with random yields has developed; see Yano and Lee (1995) and Grosfeld-Nir and Gerchak (2004) for surveys. Almost all studies assume a *single* supplier. Only Henig and Gerchak (1990), Anupindi and Akella (1993), and Swaminathan and Shanthikumar (1999) considered *multiperiod* planning models.

Henig and Gerchak (1990), assuming a single supplier, demonstrated the existence of a time-dependent critical inventory level such that an order is placed if and only if the period's starting inventory is below this level. See also Zipkin (2000) for a treatment of this model, as well as the development of a linear inflation heuristic under the long-run average cost criterion.

Anupindi and Akella (1993) were the first to analyze a multiperiod procurement problem with *two* suppliers, however, assuming that any order is either delivered *completely* without any yield loss, or not at all. For this special case, these authors established that the optimal supplier set is consecutive in the effective cost rates, and the number of suppliers decreasing with the starting inventory, properties demonstrated in full generality in this paper. Swaminathan and Shanthikumar (1999) showed, both in a single- and in a two-period model, that the optimal set of suppliers may fail to be consecutive when the demand distributions are discrete.

Systems with an arbitrary set of suppliers have been analyzed only by Agrawal and Nahmias (1997), Dada et al. (2007), and Federgruen and Yang (2009), and only in a single-period setting. Agrawal and Nahmias (1997) consider the special case of constant demand, Normal yield distributions, and zero starting inventory. For a *given* set of suppliers, the paper shows that the optimal order sizes satisfy a set of nonlinear equations, without providing a method to solve them. When $N = 2$, the authors prove that this system of equations has a unique solution. As to identifying the optimal supplier base, they suggest enumerating all possible sets. Dada et al. (2007) prove the above consecutiveness property of the optimal supplier base in a single-period model with zero starting inventory. (The authors study a yield model more general than the multiplicative structure we consider.) Federgruen and Yang (2009) derive an efficient algorithm to determine the optimal set of suppliers and order quantities in a single-period setting.

This paper generalizes all of the above references by characterizing the structure of the optimal procurement strategy in an arbitrary nonstationary periodic-review setting, where the buyer has access to an arbitrary set of suppliers.

Finally, our paper is related to other recent papers that establish important structural results of the value function in dynamic programs in which the action space is given by a subset of some Euclidean space and the expected immediate and future cost function by a nonlinear function of the action vector; see Zhu and Thonemann (2009) and Frank et al. (2003).

3. Model and Structural Results

We consider a periodic-review procurement-planning model consisting of T periods, numbered forwards towards the end of the planning horizon. The firm has access to N suppliers. The suppliers differentiate themselves from each other in terms of their yield distributions and unit prices. Let:

$c_{it}^o =$ the price charged by supplier i in period t for every *ordered* unit.

$c_{it}^e =$ the additional price charged by supplier i in period t for every *effective* unit delivered.

$X_{it} =$ the random yield factor of supplier i, with cdf $G_{it}(\cdot)$, mean p_{it}, variance $\varsigma_{it}^2 > 0$, and coefficient of variation $\gamma_{it} = \varsigma_{it}/p_{it}$; $i = 1, \ldots, N$; $t = 1, \ldots, T$.

$D_t =$ the uncertain demand in period t, with a general continuous cdf $F_t(\cdot)$, pdf $f_t(\cdot)$, ccdf $\bar{F}_t(\cdot)$, mean μ_t, standard deviation σ_t, where the pdf is assumed to be continuously differentiable on the interior of the distribution's support, $t = 1, \ldots, T$.[5]

$h_t =$ the inventory cost per unit carried in inventory at the end of period t, $t = 1, \ldots, T$.

$b_t =$ the backlogging cost rate per unit backlogged at the end of period t, $t = 1, \ldots, T$.

$\alpha =$ the discount factor $(0 < \alpha \leqslant 1)$.

Note that the expected total price paid to supplier i in period t, per ordered unit, is given by:

$$c_{it} = c_{it}^o + c_{it}^e p_{it}; \quad i = 1, \ldots, N; \; t = 1, \ldots, T. \tag{1}$$

The above two-part fee structure includes, as special cases, settings where the firm only pays for effective units ($c_{it}^o = 0$) or where it pays exclusively for every ordered unit ($c_{it}^e = 0$). In general, different suppliers ask the firm to assume a different fraction c_{it}^o/c_{it} of the cost risks associated with yield losses.

The yield factors $\{X_{it}\}$ and the sequence of demand volumes $\{D_t\}$ are assumed to be independent. We initially assume that orders placed in a given period become available in time to meet that period's demand. See §4 for relaxations of both assumptions. To formulate the planning problem as a dynamic program, we need the following variables and value functions:

$I_t =$ the inventory level at the beginning of period t, $t = 1, \ldots, T$.

$y_{it} =$ the order placed with supplier i in period t.

$Y_t = \sum_{i=1}^N y_{it} =$ the aggregate order in period t.

$Y_t^e = \sum_{i=1}^N p_{it} y_{it} =$ the expected *effective* supply in period t, i.e., the expected amount of usable supply resulting from the various orders.

$v_t(I_t) =$ the minimum expected discounted costs in periods $t, t+1, \ldots, T$, when starting period t with an inventory level I_t. (2)

The inventory dynamics are described by the following recursive scheme:

$$I_{t+1} = I_t + \sum_{i=1}^N X_{it} y_{it} - D_t, \quad t = 1, \ldots, T, \tag{3}$$

and the value functions therefore satisfy the following recursive equations: (Let $x^+ \overset{\text{def}}{=} \max\{x, 0\}$.)

$$v_t(I_t) = \min_{y_t \geqslant 0} H_t(\mathbf{y}_t, I_t), \quad \text{where} \tag{4}$$

$$\begin{aligned} &H_t(\mathbf{y}_t, I_t) \\ &= \sum_{i=1}^N c_{it} y_{it} + h_t E\left[I_t + \sum_{i=1}^N X_{it} y_{it} - D_t\right]^+ \\ &\quad + b_t E\left[D_t - I_t - \sum_{i=1}^N X_{it} y_{it}\right]^+ + \alpha E v_{t+1}\left(I_t + \sum_{i=1}^N X_{it} y_{it} - D_t\right), \\ &\qquad t = 1, \ldots, T, \quad \text{and} \end{aligned} \tag{5}$$

$$v_{T+1} \overset{\text{def}}{=} 0. \tag{6}$$

We first show that the value functions are strictly convex and that a unique optimal order vector exists for every starting inventory level. In the appendix, we show that

$$H_t(\mathbf{y}_t, I_t) = \sum_{l=1}^{N} c_{lt} y_{lt} + h_t\left(I_t + \sum_{l=1}^{N} p_{lt} y_{lt} - \mu_t\right) + (h_t + b_t) E_{\{X_{lt}\}} \int_{I_t+\sum_{l=1}^{N} X_{lt} y_{lt}}^{+\infty} \bar{F}_t(u)\, du + \alpha E_{\{X_{lt}\}} \int_{-\infty}^{+\infty} v_{t+1}\left(I_t + \sum_{l=1}^{N} X_{lt} y_{lt} - u\right) dF_t(u). \tag{7}$$

Theorem 1. *Fix* $t = 1, \ldots, T$.

(a) *For all* $-\infty < I_t < +\infty$, $v_t(I_t) < \infty$ *and* $\lim_{I_t \to -\infty} v_t(I_t) = \lim_{I_t \to +\infty} v_t(I_t) = \infty$.

(b) *The value function* $v_t(\cdot)$ *is strictly convex, and the function* $H_t(\mathbf{y}_t, I_t)$ *is jointly strictly convex and supermodular in* $(\mathbf{y}_t, I_t)$.

(c) *For each starting inventory level* I_t, *there exists a unique optimal order vector* $\mathbf{y}_t^*(I_t)$.

As is immediate from its proof, all of the structural results in Theorem 1 continue to apply when [some of] the demand distributions are discrete or mixed. Additional structural results depend on the value functions $v_t(\cdot)$ being continuously differentiable and twice differentiable almost everywhere, as shown in Lemma 1 below. These analytical properties of the value functions rely on our assumption that the demand functions have continuous distributions.

Lemma 1. *Fix* $t = 1, \ldots, T$.

(a) *The function* $H_t(\mathbf{y}_t, I_t)$ *is twice continuously differentiable.*

(b) *The value function* $v_t(\cdot)$ *is continuously differentiable with*

$$v_t'(I_t) = \frac{\partial H_t(\mathbf{y}_t^*, I_t)}{\partial I_t} \quad \text{and} \quad -\sum_{s=t}^{T} \alpha^{s-t} b_s \leqslant v_t'(I_t) \leqslant \sum_{s=t}^{T} \alpha^{s-t} h_s. \tag{8}$$

(c) *The value function* $v_t(\cdot)$ *is twice continuously differentiable except for at most N points.*

In view of the joint convexity of the function $H_t(\mathbf{y}_t, I_t)$, by (7), the optimal set of orders to be used in period t is the unique solution of the following system of equations and inequalities:

$$\partial H_t(\mathbf{y}_t^*, I_t)/\partial y_{it} = c_{it} + h_t p_{it} - (b_t + h_t) E_{t\{X_{lt}\}}\left[X_{it} \bar{F}_t\left(I_t + \sum_{l=1}^{N} X_{lt} y_{lt}^*\right)\right] + \alpha E_{\{X_{lt}\}}\left[X_{it} \int_{-\infty}^{+\infty} v_{t+1}'\left(I_t + \sum_{l=1}^{N} X_{lt} y_{lt}^* - u\right) dF_t(u)\right] \begin{cases} = 0, & \text{if } y_{it}^* > 0 \\ \geqslant 0, & \text{if } y_{it}^* = 0. \end{cases} \tag{9}$$

This characterization permits us to show that the optimal set of suppliers to be retained in any period t is consecutive in the effective cost rates $\{c_{it}/p_{it}\}$. In fact, in Theorem 2, below, we show that the optimal set of suppliers in period t consists of those whose effective cost rate falls below a benchmark cost rate:

$$\lambda_t^E(I_t) \overset{\text{def}}{=} -v_t'(I_t) = -h_t + (b_t + h_t) E_{\{X_{lt}\}}\left[\bar{F}_t\left(I_t + \sum_{l=1}^{N} X_{lt} y_{lt}^*\right)\right] - \alpha E_{\{X_{lt}\}}\left[\int_{-\infty}^{+\infty} v_{t+1}'\left(I_t + \sum_{l=1}^{N} X_{lt} y_{lt}^* - u\right) dF_t(u)\right] \tag{10}$$

$$\leqslant \bar{\lambda}_t(I_t) \overset{\text{def}}{=} -h_t + (b_t + h_t)\bar{F}_t(I_t) - \alpha \int_{-\infty}^{+\infty} v_{t+1}'(I_t - u)\, dF_t(u). \tag{11}$$

(The second equality follows from (8), see (20) in the appendix; the inequality follows from the fact that its left-hand side is decreasing in each of the order quantities, itself a consequence of Theorem 1(b).) The benchmark cost rate represents the expected value of the total cost saving, associated with a marginal effective unit, delivered—for free and outside of the normal procurement process—beyond those arising from the optimal set of orders; here, the total cost saving relates to current holding and backlogging costs, as well as all future costs.

Theorem 2. *Fix period* $t = 1, \ldots, T$. *Renumber the suppliers such that* $c_{1t}/p_{1t} \leqslant c_{2t}/p_{2t} \leqslant \cdots \leqslant c_{Nt}/p_{Nt}$.

(a) *The optimal set of suppliers is consecutive in the sequence of effective cost rates* $\{c_{it}/p_{it}\}$.

(b) *The benchmark cost rate* $\lambda_t^E(I_t)$ *is a strictly decreasing, continuous function of* I_t.

(c) $k^*(I_t)$, *the optimal number of suppliers, is decreasing in* I_t, *i.e., the optimal set of suppliers shrinks as a function of* I_t. *In other words, additional units of starting inventory may result in the elimination of one or more of the most expensive suppliers in the retained supplier set.*

(d) $\mathbf{y}_t^*$ *is a continuous function that is differentiable everywhere except for at most N points, where this vector function has left and right derivatives.*

(e) *There exists an inventory level* S_t, *such that it is optimal to place orders if and only if the starting inventory* $I_t < S_t$.

(f) S_t *is the unique root of the strictly increasing function* $c_{1t} + h_t p_{1t} - (b_t + h_t) p_{1t} \bar{F}_t(I_t) + \alpha p_{1t} \int_{-\infty}^{+\infty} v_{t+1}'(I_t - u) dF_t(u)$.

Theorem 2(a) shows that the optimal set of suppliers in any given period is consecutive in *that* period's effective cost rates. The *degree* of supplier diversification, i.e., how *many* suppliers are to be retained depends, of course, on all current and future yield and demand distributions, as well as all current and future cost rates. Strikingly, the dependence on these various distributions and cost parameters occurs via a single aggregate measure, i.e., the benchmark cost rate $\lambda_t^E(I_t)$: the retained suppliers are precisely those whose effective cost rate is strictly below this benchmark. This characterization also implies that if a supplier fails to be part of the set of retained suppliers, he cannot become competitive by improving his yield distribution alone. In the words of Hill (2000), cost can be thought of as an "order qualifier," whereas reliability acts as an "order winner." The consecutiveness property of the optimal supplier set was first obtained by Anupindi and Akella (1993) in a two-supplier, but multiperiod, setting with Bernoulli yield factors, and by Dada et al. (2007) and Federgruen and Yang (2009) for a single-period setting with an arbitrary number of suppliers. Our assumption that the demand distributions are continuous is essential for the consecutiveness property. Under discrete or mixed demand distributions, any of the functions H_t, although convex, may fail to be differentiable at the optimal solution $\mathbf{y}_t^*$. Similarly, the value function v_{t+1} may not be differentiable in countably many points, so that the (last term in the) benchmark cost rate $\lambda_t^E(I_t)$, itself, is ill-defined when the demand distribution has a positive mass in any of the points where v_{t+1} fails to be differentiable. Indeed, Swaminathan and Shanthikumar (1999) show that under discrete distributions, the optimal supplier set needs not be consecutive, even in a single-period setting with $N=2$ suppliers; i.e., with a discrete demand distribution, it may be optimal to order from the more expensive supplier exclusively.

Similarly, the monotonicity of the optimal supplier set as a function of the starting inventory generalizes the same result obtained by Anupindi and Akella (1993) in their two-supplier model. In the case of a *single* supplier, as first shown by Henig and Gerchak (1990), the supermodularity of the value function v_t implies that the optimal order quantity is decreasing in the starting inventory level; see Theorem 1(b). Based on extensive numerical results, we conjecture that this property applies under an arbitrary number of suppliers.[6]

Our model with multiple unreliable suppliers inherits the well-known property in the classical model, that, in each period t, orders are placed if and only if the starting inventory is below a given threshold S_t. However, the threshold S_t is no longer the order-up-to level for all the inventory levels below it. The single-period example in Figure 2 of Federgruen and Yang (2009) exhibits the following phenomena: (i) the expected order-up-to level usually decreases, but it sometimes fails to be monotone, and may even be increasing in the starting inventory; and (ii) the expected order-up-to level is sometimes larger, but sometimes smaller, than the level in the corresponding classical model. Observation (ii) is somewhat surprising, because one might conjecture that the need for safety stocks would be larger when supply risks compound on demand risks. Actually, it is the *relative* magnitude of the supply risks compared to the demand risks that determines whether the order-up-to level is larger or smaller than what is optimal in the absence of supply risks. In particular, when the cost consequences of a shortage are relatively low, the additional supply risks may render it optimal to target a lower, rather than a higher, expected inventory level after ordering.

The remaining question is how the different cost parameters and yield and demand distributions impact on the *allocation* of the aggregate orders among the retained suppliers. A very insightful characterization of these market shares can be obtained when all distributions are Normal; see Theorems 4(b) and 5(c) in Federgruen and Yang (2008b):

$$p_{it}y_{it}^* = \frac{[(\lambda_t^E(I_t) - c_{it}/p_{it})^+]\gamma_{it}^{-2}}{\sum_{l=1}^N [(\lambda_t^E(I_t) - c_{lt}/p_{lt})^+]\gamma_{lt}^{-2}} Y_t^{e*}, \qquad i = 1, \ldots, N. \tag{12}$$

In other words, the share of each supplier in the expected effective total supply Y_t^{e*} is given by a supplier score, itself the product of a reliability score γ_{lt}^{-2}, and a cost-saving score that measures the saving, relative to the benchmark cost rate $\lambda_t^E(I_t)$, per effective unit, of using this supplier. By (12), the optimal market shares of the suppliers are simply proportional to their supplier scores and can be obtained as a simple closed-form expression in terms of the model parameters once the (single) benchmark cost rate has been computed. The dependence of the optimal supplier set and their market shares on all current and future cost, demand, and yield considerations arises via a single aggregate quantity, i.e., the benchmark cost rate $\lambda_t^E(I_t)$. Federgruen and Yang (2008b) also show that the market share formula (12) can be used as a close approximation when (some of) the demand and yield distributions fail to be Normal.

The optimality conditions (9), along with the consecutiveness property of the optimal set of suppliers, suggest the following algorithm (GA) to find the optimal order vector $\mathbf{y}_t^*$ in any given period t. (As in Theorem 2, the suppliers are numbered such that $c_{1t}/p_{1t} \leqslant c_{2t}/p_{2t} \leqslant \cdots \leqslant c_{Nt}/p_{Nt}$.)

Algorithm 1 (General Algorithm (GA))

FOR $k := 1$ TO N DO
BEGIN
Step 1. Find the unique nonnegative solution $\{y_{1t}^*, \ldots, y_{kt}^*\}$ of the following system of k equations in k unknowns $\{y_{1t}, \ldots, y_{kt}\}$:

$$c_{it} + h_t p_{it} - (b_t + h_t) E_{\{X_{lt}\}}\left[X_{it}\bar{F}_t\left(I_t + \sum_{l=1}^k X_{lt}y_{lt}\right)\right] + \alpha E_{\{X_{lt}\}}\left[X_{it}\int_{-\infty}^{+\infty} v'_{t+1}\left(I_t + \sum_{l=1}^k X_{lt}y_{lt} - u\right) dF_t(u)\right] = 0, \qquad i = 1, \ldots, k$$

Step 2. Set $y_{it}^* = 0$ for $i > k$. Calculate $\lambda_t^E(I_t)$ from (10).
IF $\lambda_t^E(I_t) \leqslant c_{k+1}/p_{k+1}$, THEN
exit with $\mathbf{y}_t^*$ as the optimal order vector
ENDIF
END (13)

(If the test in Step 2 is met, the right-hand side of (9) is nonnegative for all $i = k+1, \ldots, N$, whereas in view of Step 1, it equals zero for $i = 1, \ldots, k$, thus verifying that $\mathbf{y}_t^*$ is the optimal order vector. Also, the test in Step 2 is met for exactly *one* value of k, permitting an exit as soon as it is met: if it were met for two different values of k, there would be two distinct optimal solutions, whereas by Theorem 2(a), (9) would fail to have a solution if the test in Step 2 was never met.) (GA) amounts to solving at most N systems of well-behaved equations, each of which has a unique solution. The main difficulty arises from the evaluation of the multivariate expectations to the right of (13). For general yield and demand distributions, this is best done by a simulation technique.

4. Extensions

We conclude this paper with a discussion of several important generalizations of our model.

4.1. Fixed Supplier Cost; Price Benefits Associated Multisourcing

Our model ignores any fixed cost K_i incurred for each supplier i that is added to the pool of potential suppliers $\mathcal{P}$ (for example, costs associated with buyers and information systems). Such fixed costs provide an incentive to limit the degree of supplier diversification. They can, of course, easily be incorporated when comparing aggregate expected costs under two or more sets of suppliers. Such comparisons may also allow us to model a second benefit of supplier diversification, i.e., the ability to negotiate better prices, when dealing with a larger pool of qualified suppliers. Representing the cost rates $\{c_{it}\}$ as decreasing functions of the number of suppliers in the pool $\mathcal{P}$, i.e., $c_{it} \stackrel{\text{def}}{=} c_{it}(|\mathcal{P}|)$, the overall cost of a pool $\mathcal{P}$ is given by $C(\mathcal{P}) \stackrel{\text{def}}{=} \sum_{i\in\mathcal{P}} K_i + v_1(I_1 \mid c_{it} = c_{it}(|\mathcal{P}|))$. Identifying the optimal pool of potential suppliers $\mathcal{P}^* = \arg\min_{\mathcal{P}} C(\mathcal{P})$ is in general a complex combinatorial problem, which is NP-hard even in single-period settings, see Proposition 1 in Federgruen and Yang (2008a).

4.2. Lead Times

Positive lead times of $L \geqslant 1$ periods can be handled in a similar manner as in the classical model with a single, fully reliable supplier, provided the actual yields for the orders placed at the beginning of a given period become known before the start of the next period: only orders in periods $t = 1, 2, \ldots, T-L$ are relevant, thereafter, the orders fail to be received during the considered planning horizon. We now use as the state variable: $\mathrm{IP}_t = I_t + \sum_{\tau=t-L}^{t-1} y_{i\tau} X_{i\tau} =$ *inventory position* at the beginning of period $t =$ the inventory on hand, plus the effective supply in process. Because all unsatisfied demand is backlogged, the inventory position satisfies the same recursive scheme as (3): $\mathrm{IP}_{t+1} = \mathrm{IP}_t + \sum_{i=1}^N X_{it} y_{it} - D_t$, $t = 1, \ldots, T$, whereas $I_{t+L} = \mathrm{IP}_t + \sum_{i=1}^N X_{it} y_{it} - (D_t + D_{t+1} + \cdots + D_{t+L})$. Recognizing that the expected end-of-the-period holding and backlogging costs in periods $t = 1, \ldots, L$ cannot be affected by any of the procurement decisions, these cost terms may therefore be eliminated from the dynamic program. Charging to period t the expected inventory and backlogging costs that incur at the end of period $t+L$, we obtain the modified dynamic program:

$$v_t(\mathrm{IP}_t) = \min_{y_t \geqslant 0} H_t(\mathbf{y}_t, \mathrm{IP}_t), \quad t = 1, \ldots, T-L, \quad \text{where}$$

$$\begin{aligned} H_t(\mathbf{y}_t, \mathrm{IP}_t) &= \sum_{i=1}^N c_{it} y_{it} + \alpha^L h_t E\Bigg[\mathrm{IP}_t + \sum_{i=1}^N X_{it} y_{it} \\ &\qquad - (D_t + D_{t+1} + \cdots + D_{t+L})\Bigg]^+ \\ &\quad + \alpha^L b_t E\left[D_t + D_{t+1} + \cdots + D_{t+L} - \mathrm{IP}_t - \sum_{i=1}^N X_{it} y_{it}\right]^+ \\ &\quad + \alpha E v_{t+1}\left(\mathrm{IP}_t + \sum_{i=1}^N X_{it} y_{it} - D_t\right), \end{aligned}$$

$$v_{T-L+1}(\cdot) = 0.$$

It is easily verified that *all* of the results in this paper continue to apply. The problem is considerably more complex when the yields of a given period's orders do not become known to the purchasing firm before the start of the next period. As an extreme case, assume that the actual yields are not revealed until the orders are delivered. In this case, the inventory position at the beginning of period t is itself unknown and only partially observable. To model this situation, we need to keep track of the L order vectors in process, requiring a dynamic program with a state space of dimension $LN + 1$.[7]

4.3. Capacities

Capacity bounds represent an additional complication in many applications. Thus, assume that a capacity limit M_{it} prevails for any order placed with supplier i in period t, i.e., $y_{it} \leqslant M_{it}$ ($i = 1, \ldots, N$; $t = 1, \ldots, T$). One easily verifies that all of the results in Lemma 1, Theorems 1 and 2, continue to apply for arbitrary capacities; see Theorem 3 in the appendix. In particular, in any period t, the optimal set of retained suppliers continues to be the consecutive set $\{i\colon c_{it}/p_{it} < \lambda_t^E(I_t)\}$. However, the simple market share formula (12) no longer applies because the market shares are now affected by the capacity levels in addition to the yield reliabilities and cost differentials.

4.4. Correlated Yields and Demands

In some settings, supply risks may be correlated, for example, when natural disasters (storms, floods) or sabotage by terrorists are likely to hit multiple facilities in a given geographic region, or when the suppliers depend on common second-tier suppliers. Similarly, the yields and demand distributions in a given period may be correlated, for example, when both are dependent on weather-related factors or common economic variables; see, for example, Babich et al. (2007) for a procurement model with multiple suppliers subject to correlated yield risks.

It can be verified that all of the results in Lemma 1 and Theorem 1 continue to apply. However, it is no longer true that the optimal set of suppliers, in any given period, is consecutive in the effective cost rates, i.e., consists of those whose effective cost rate is below a given benchmark rate.

More specifically, the second equality in (22) in the appendix (electronic companion) breaks down under correlated yields, so that a supplier i may not be selected even if its effective cost rate falls below $\lambda_t^E(I_t)$, the benchmark cost rate. It is therefore easy to construct examples where the optimal supplier set fails to be consecutive.

5. Electronic Companion

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

Endnotes

1. Another major benefit of multisourcing results from increased bargaining power to reduce purchase prices. Although this benefit is not directly addressed, our model can be used to compare the overall long-term performance of the firm when having access to additional potential suppliers. In such comparisons, the suppliers' unit prices could be modeled as decreasing functions of the total number of suppliers the firm deals with; see §4.
2. Nagali et al. (2008, p. 51) state: "In periods of high demand, hi-tech suppliers place original equipment manufacturers (OEMs) such as HP under allocation whereby they supply only a fraction of the OEM's total demand. Availability uncertainty can also result from supply and delivery disruptions, such as the earthquake in Taiwan in late 1999, or supplier quality issues." HP's PRM process was a finalist in the 2007 Edelman Award Competition.
3. See, however, §3 for a discussion of how this result can be used as an approximation when the distributions are of a different type.
4. Commercially available supplier scorecard systems tend to determine aggregate scores as the sum or a weighted average scores of individual criteria, apparently without any theoretical justification; see e.g., http://www.theperformancescore.com/index.asp?pgid=21 and http://www.commercezone.co.za/CWS_CommerceZone/default.aspx?i_CategoryID=68.
5. Demand may represent net demand, net of precommitted and guaranteed deliveries, as under the flexible quantity contracts discussed in §1. In our base model, we assume that the demand distributions have the positive half line as their support. All of our results are easily extended when the support is given by a different interval, for example, the full real line, as in the case of Normal distributions.
6. This could, for example, be established by showing that the right-hand side of (25) in the appendix (electronic companion) is nonpositive.
7. See Federgruen and Yang (2008b) for a reduction to an $L+1$ dimensional state space when all distributions are Normal or when end-of-the-period inventory levels are approximated as Normals.

References

Agrawal, N., S. Nahmias. 1997. Rationalization of the supplier base in the presence of yield uncertainty. *Production Oper. Management* **6**(3) 291–308.
Anupindi, R., R. Akella. 1993. Diversification under supply uncertainty. *Management Sci.* **39**(8) 944–963.
Arrow, K., T. Harris, J. Marschak. 1951. Optimal inventory policy. *Econometrica* **19**(3) 250–272.
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.
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.
Dvoretzky, A., J. Kiefer, J. Wolfowitz. 1952. The inventory problem: I. Case of known distributions of demand. *Econometrica* **20**(2) 187–222.
Federgruen, A., N. Yang. 2008a. Selecting a portfolio of suppliers under demand and supply risks. *Oper. Res.* **56**(4) 916–936.
Federgruen, A., N. Yang. 2008b. Multi-period procurement strategies in the presence of unreliable suppliers (unabridged version). Working paper, Columbia University, New York.
Federgruen, A., N. Yang. 2009. Optimal supply diversification under general supply risks. *Oper. Res.* **57**(6) 1451–1468.
Frank, K., R. Q. Zhang, I. Duenyas. 2003. Optimal inventory policies in systems with priority demand classes. *Oper. Res.* **51**(6) 993–1002.
Grosfeld-Nir, A., Y. Gerchak. 2004. Multiple lotsizing in production to order with random yields: Review of recent advances. *Ann. Oper. Res.* **126**(1) 43–69.
Henig, M., Y. Gerchak. 1990. The structure of periodic review policies in the presence of random yield. *Oper. Res.* **38**(4) 634–643.
Hill, T. 2000. *Manufacturing Strategy*, 3rd ed. McGraw-Hill/Irwin, Boston.
Karlin, S. 1958. Optimal inventory policy for the Arrow-Harris-Marschak dynamic model. K. Arrow, S. Karlin, H. Scarf, eds. *Studies in the Mathematical Theory of Inventory and Production*, Chapter 9. Stanford University Press, Stanford, CA.
Nagali, V., J. Hwang, D. Sanghera, M. Gaskins, M. Pridgen, T. Thurston, P. Mackenroth, D. Branvold, P. Scholler, G. Shoemaker. 2008. Procurement risk management (PRM) at Hewlett-Packard Company. *Interfaces* **38**(1) 51–60.
Swaminathan, J. M., J. G. Shanthikumar. 1999. Supplier diversification: Effect of discrete demand. *Oper. Res. Lett.* **24**(5) 213–221.
Treece, J. 1997. Just-too-much single-sourcing spurs Toyota purchasing review: Maker seeks at least 2 suppliers for each part. *Automotive News* (March 3) 3.
Yano, C. A., H. L. Lee. 1995. Lot sizing with random yields: A review. *Oper. Res.* **43**(2) 311–334.
Zhu, K., U. W. Thonemann. 2009. Coordination of pricing and inventory control across products. *Naval Res. Logist.* **56**(2) 175–190.
Zipkin, P. H. 2000. *Foundations of Inventory Management*. McGraw-Hill Book Co., New York.