Operations Research
Vol. 57, No. 6, November–December 2009, pp. 1451–1468
ISSN 0030-364X | EISSN 1526-5463 | 09 | 5706 | 1451

DOI 10.1287/opre.1080.0667

Optimal Supply Diversification Under General Supply Risks

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 analyze a planning model for a firm or public organization that needs to cover uncertain demand for a given item by procuring supplies from multiple sources. The necessity to employ multiple suppliers arises from the fact that when an order is placed with any of the suppliers, only a random fraction of the order size is usable. The model considers a single demand season with a given demand distribution, where all supplies need to be ordered simultaneously before the start of the season. The suppliers differ from one another in terms of their yield distributions, their procurement costs, and capacity levels.

The planning model determines which of the potential suppliers are to be retained and what size order is to be placed with each. We consider two versions of the planning model: in the first, the service constraint model (SCM), the orders must be such that the available supply of usable units covers the random demand during the season with (at least) a given probability. In the second version of the model, the total cost model (TCM), the orders are determined so as to minimize the aggregate of procurement costs and end-of-the-season inventory and shortage costs. In the classical inventory model with a single, fully reliable supplier, these two models are known to be equivalent, but the equivalency breaks down under multiple suppliers with unreliable yields.

For both the service constraint and total cost models, we develop a highly efficient procedure that generates the optimal set of suppliers as well as the optimal orders to be assigned to each. Most importantly, these procedures generate a variety of important qualitative insights, for example, regarding which sets of suppliers allow for a feasible solution, both when they have ample supply and when they are capacitated, and how various model parameters influence the selected set of suppliers, the aggregate order size, and the optimal cost values.

Subject classifications: inventory/production: approximations/heuristics; operating characteristics; reliability: failure models.
Area of review: Manufacturing, Service, and Supply Chain Operations.
History: Received January 2007; revisions received September 2007, January 2008; accepted March 2008. Published online in *Articles in Advance* June 25, 2009.

1. Introduction and Summary

We analyze a planning model for a firm or public organization that needs to cover uncertain demand for a given item by procuring supplies from multiple sources. The necessity or desirability of employing multiple suppliers arises from the fact that when an order is placed with any of the suppliers, only a random fraction of the order size is useable or materializes. This random fraction is referred to as the yield factor and it follows a general (supplier-specific) probability distribution. An important special case arises when the yield distribution has a positive mass at zero, representing the possibility of a complete breakdown due to natural or man-made disruptions or the firm's bankruptcy. The model considers a single demand season with a given demand distribution, where all supplies need to be ordered simultaneously before the start of the season. The suppliers differ from each other in terms of their yield distributions as well as their per-unit procurement costs. The yield factors at different suppliers are independent of each other as well as the season's demand. (See, however, §7 for a treatment of correlated yield and demand distributions.)

The planning model determines which of the potential suppliers are to be retained and what size order is to be placed with each. We consider two versions of the planning model: In the first, orders are to be chosen to minimize procurement costs, while ensuring that the available supply of useable units covers the random demand during the season with (at least) a given probability. In the second version of the model, the orders are determined so as to minimize the aggregate of procurement costs and end-of-the season inventory and shortage costs. We refer to these two versions as the service constraint model (SCM) and the total cost model (TCM), respectively. In classical inventory theory with fully reliable suppliers, assigning a direct stockout penalty for each unsatisfied unit of demand and employing a constraint on the probability of a stockout represent the two common approaches to control the stockout phenomenon. Much has been written about the relative merits of both modeling approaches, see, e.g., Zipkin (2000). Both approaches continue to be pursued in parallel, even though in classical inventory models, the two approaches are known to be equivalent: an instance of (TCM) with a given backlogging rate induces the same optimal inventory strategy as an instance of (SCM) with a corresponding permitted shortfall probability, and vice versa. See Boyaci and Gallego (2001) for a recent discussion of this equivalency in classical inventory models. The equivalency breaks down under multiple suppliers with unreliable yields, adding to the need to pursue both planning approaches in parallel. First, whereas a feasible solution always exists in the (TCM), in the (SCM), feasibility requires a minimum number of sufficiently reliable suppliers, as specified below. A key concept in both models is the so-called expected effective supply, i.e., the expected total number of usable units obtained from the various suppliers. We show that, for a given expected effective supply level, the optimal set of suppliers and the orders can be obtained in closed form, after determining the root of a single-variable function. Also, the total cost is a strictly convex function of the expected effective supply with a unique minimum. In the (SCM), a larger expected effective supply is optimally assigned to the same number, or fewer, suppliers, i.e., if one is willing to place a larger aggregate order, there is less need to diversify among suppliers; in the (TCM), this monotonicity property may fail to hold. In the (SCM), the safety stock (= expected inventory after ordering – expected demand) is always larger than in the classical model without supply risks. Once again, this is not always the case in the (TCM).

Recently, much attention has been given to the need to diversify the supplier pool, so as to provide adequate protection against the possibility of uncertain yields or complete disruptions due to natural causes, such as fires or hurricanes, man-made breakdowns (e.g., sabotage or terrorist attacks), as well as bankruptcy. Most recently, the nation has focused on this challenge after hurricane Katrina demolished almost 10% of the U.S. refinery capacity, driving the price of gasoline and other refined oil products through the roof. The need to "plan for disaster" and to adequately diversify the pool of supply options has been recognized as one of the premier challenges in supply chain management in the twenty-first century. See, for example, Longitudes (2004), the proceedings of a conference bringing together government, business, and academic leaders, in which this theme was highlighted with a case study of the cellular phone industry. This study contrasted the supply strategies of Ericsson and Nokia. Both used a Philips chip supplier in New Mexico as the primary source for one of the key electronic chips. However, whereas Ericsson relied entirely on this supplier, Nokia had put in place a variety of alternative supply options. When the Philips plant had to be shut down for an extended period of time due to a major fire, Ericsson suffered major losses in its sales volumes and profits, as well as a large reduction of its market share, for years to come.

The failure to satisfy quality standards or regulations represents another potential source for major disruptions in the supply process. In the fall of 2004, the United States saw half of its flu vaccine supply disappear when one of its two suppliers had to bow out after the Food and Drug Administration and its British counterpart closed the Chiron plant in Britain. Similar supply shortages have occurred repeatedly with this and other vaccines. Finally, it is generally recognized that future terrorist attacks are likely to target the supply process of vital commodities or food products. We refer to Federgruen and Yang (2008) for a more extensive discussion of these supply disruptions and their impact on the economy and general welfare.

Even companies that were able to develop tight partnerships with their suppliers have come to realize that single sourcing is far too risky. A prime example is Toyota, whose assembly plants were forced to shut down in 1997 after a fire at Aisin. (Prior to 1997, Aisin provided 90% of all brake components and practically all brake valves for Toyota; see Nishiguchi and Beaudet 1998.) Thereafter, Toyota sought multiple parallel suppliers for each part; see Treece (1997). Hewlett-Packard's Procurement Risk Management group launched, in 2000, a multisourcing strategy for its components, the ultimate profit contribution of which is estimated to amount to \$1 billion; see Nagali et al. (2008, p. 51).[1]

Two recent developments have acted as additional catalysts for the multisourcing movement. First, electronic commerce provides a platform with far lower overheads for becoming a supplier or for splitting orders. As Ketchpel and Garcia-Molina (1998, p. 603) write: "Customers can have many suppliers offer bids for a contract, or split a 'bundled' order across multiple suppliers more easily." Second, modern management information systems allow companies to track the performance of many suppliers along multiple dimensions, including costs and yield distributions. To this end, many firms maintain scorecards for their suppliers, who are graded for each of a list of criteria. Based on a relative weighting of the criteria, an overall score for each supplier is derived. These aggregate scores are used to select the supplier set and to determine their shares in the procurement orders. Developing scorecard systems has become one of the consulting world's premier supply chain services; see, e.g., Oracle's Peoplesoft Manufacturing Scorecard at http://www.oracle.com/media/peoplesoft/en/pdf/datasheets/e_epm_ds_manufacture_42005.pdf. Several cases have been written on the use of scorecard systems at specific firms, e.g., Holloway et al. (1996) regarding Sun Microsystems, and Kulp and Naravanan (2002) regarding Metal Craft, a \$13 billion supplier of automotive parts. (See also Pyke and Johnson 2003 for a discussion of their use at Air Products.)

Hitherto, the scoring methods for the individual criteria, as well as the aggregate scheme that generates the supplier's overall score, fail to be based on specific planning models. The results in this paper provide insights for the design of scorecard-based supplier selection and allocation systems. For both the (SCM) and (TCM), we develop a highly efficient procedure that generates the optimal set of suppliers as well as the optimal order for each. We also derive the following important qualitative insights: first, when ranking the suppliers in ascending order of their effective cost rates, the optimal set is, in both the (SCM) and (TCM), *consecutive*, i.e., it consists of the first k suppliers for some $k = 1, \ldots, N$. This result generalizes that of Anupindi and Akella (1993), obtained for the case of $N = 2$ and a general continuous-demand distribution. (For the same special case with $N = 2$ suppliers, Swaminathan and Shanthikumar 1999 showed that the optimal set of suppliers may fail to be consecutive, i.e., only the most expensive supplier is used for certain discrete-demand distributions.) In both the (SCM) and (TCM), each selected supplier is assigned an overall score, given by the *product* of a reliability and a cost score: The former is the mean-to-variance ratio of the supplier's yield distribution, and the latter is given by the amount by which the supplier's effective cost rate falls below a specific threshold value. The market share of each selected supplier is given by his overall score relative to the sum of the suppliers' scores.[2]

We systematically characterize the ramifications for (i) the supplier base, (ii) the expected effective supply (and hence the safety stock), and (iii) the optimal cost value resulting from changes in the supply risks, the demand magnitude and risks, and the amount of initial inventory one possesses. For example, we analyze what impact an increase in either the mean or the standard deviation on any of the yield distributions has on any of the above-mentioned characteristics of the solution. We do the same with respect to the mean and standard deviation of the demand distribution. We also show in the (SCM) that whether a supplier achieves a positive market share or not depends only on his own effective cost rate and those of his less expensive competitors, along with the coefficients of variation of their yield distributions, as the sole characteristic of these distributions. The supplier's own yield distribution is immaterial to ensuring membership of the patronized supplier base. Alternatively, if it is optimal for the buyer to patronize the k^* cheapest suppliers, the only way for any of the other suppliers to become part of the supplier base is to reduce their effective cost rate to a given maximum value.

When the suppliers charge identical effective cost rates, the optimal expected effective supply depends on the suppliers' yield distributions via a single aggregate measure. More specifically, the optimal supply quantity decreases convexly with this aggregate reliability measure. Thus, the cost reduction a new supplier realizes by joining a given industry of suppliers is larger than if he joins an industry with additional suppliers. However, surprisingly, when the suppliers' effective cost rates are different, any given supplier's market share as well as the marginal benefit he provides to the buyer may be larger when competing with additional suppliers.

Our base model assumes that the procurement costs associated with any supplier are proportional with his order size, that all suppliers have ample capacity, that the initial inventory is perfectly known, and that the supply risks of different suppliers are independent of each other and of the demand risk. In §6, we extend each of these restrictions. In particular, we show that capacity limits, fixed costs, uncertainty surrounding the initial inventory, and correlated yield and demand distributions can be incorporated into the model.

We refer to Yano and Lee (1995) and Grosfeld-Nir and Gerchak (2004) for surveys of a large literature on inventory systems with random yields. Almost all studies assume a *single* supplier. Gerchak and Parlar (1990), Yano (1991), and Parlar and Wang (1993) were among the first to demonstrate the benefits of *dual* sourcing in the presence of supply uncertainty. As mentioned, Anupindi and Akella (1993) and Swaminathan and Shanthikumar (1999) considered the (TCM) with $N = 2$ suppliers. These authors also generalize some of their results to allow for multiple periods. Ilan and Yadin (1985) appear to have been the first to address a model with an arbitrary set of potential suppliers. Yano and Lee (1995, p. 329) explained, in their survey paper, that the complexity of dealing with a general set of suppliers is extreme and "hence it is difficult to obtain structural results."

Agrawal and Nahmias (1997) address the (TCM) with an arbitrary number of suppliers and the special case of constant demand, Normal yield distributions, and zero starting inventory. (As mentioned and discussed throughout our paper, many of the structural properties of the optimal solution depend on the value of the 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 for identifying the optimal supplier base, they suggest enumerating all possible sets. After completing our paper, we became aware of Dada et al. (2007), who, for the (TCM) with zero starting inventory, established the above consecutiveness property of the optimal supplier base. This paper derives the consecutiveness property, along with a few other structural properties, without developing a solution method. (The authors study a yield model, more general than the multiplicative structure we consider.) Burke et al. (2007) consider the special case of the (TCM) in which all suppliers have fully predictable yields. The challenge to optimally diversify among unreliable suppliers is discussed in Chapter 9 of Van Mieghem's (2007) textbook on operations strategy.

Our work is also related to the literature on multidimensional newsvendor models; see Harrison and Van Mieghem (1999), Van Mieghem (1998), Rudi and Zheng (1997), and Van Mieghem and Rudi (2002). In the latter, multiple suppliers provide *complementary* products for the delivery of one or more final consumer goods with uncertain demand, whereas in this paper, the suppliers provide *substitutes*. We refer to Federgruen and Yang (2007) for a more detailed literature review. As mentioned, that paper confines itself to the special case of the (SCM) where all suppliers have identical variable procurement cost rates, but allows for general demand distributions. The authors develop and characterize the properties of the CLT-based approximation along with those of an alternative approximation based on a large deviation technique.

The remainder of this paper is organized as follows. The (SCM) and (TCM) models are specified in §2 and analyzed in §§3 and 4, respectively. Section 5 discusses the impact of reliability improvement, additional suppliers, the initial stock, and the magnitude of the demand and supply risks on the optimal supplier base and effective supply. Section 6 addresses the capacitated version of the models and discusses how fixed costs, uncertainty surrounding the initial inventory, and correlated yield and demand distributions can be incorporated into the analysis. Section 7 contains concluding remarks. An electronic companion to this paper is available as part of the online version that can be found at http://or.journal.informs.org/. The electronic companion consists of three parts: part A contains all proofs; part B shows the lack of monotonicity of the number of suppliers with respect to the expected effective supply in the (TCM) model with an example, and part C studies two special cases of capacitated suppliers.

2. The Service Constraint Model and the Total Cost Model

In this section, we formulate the (SCM) and (TCM) and develop some preliminary results that are common to both models. We first need the following notation for the primitives of the models:

$c_i =$ procurement cost, paid for every unit ordered from supplier i, $i = 1, \dots, N$;
$I^0 =$ initial inventory, before ordering;
$X_i =$ random yield factor of supplier i, with c.d.f. $G_i(\cdot)$, mean p_i, variance $\varsigma_i^2 > 0$, and coefficient of variation $\gamma_i = \varsigma_i/p_i$, $i = 1, \dots, N$;
$D =$ uncertain demand during the season, assumed to be Normal with mean μ, standard deviation σ, and coefficient of variation $\gamma_D = \sigma/\mu$;
$\epsilon =$ a standard Normal random variable with c.d.f. $\Phi(\cdot)$ and complementary c.d.f. $\bar{\Phi}(\cdot)$.

Without loss of generality, we rank the set of suppliers $S = \{1, \dots, N\}$ in ascending order of their effective cost rate, i.e., the expected cost value of an effectively produced unit, c_i/p_i, i.e., $c_1/p_1 \leqslant c_2/p_2 \leqslant \cdots \leqslant c_N/p_N$. Although we initially assume that the firm pays for every unit ordered, irrespective of whether it is delivered as a usable unit or not, we discuss alternative scenarios below. Let $S^0 = \{i\colon c_i/p_i = c_1/p_1\}$ denote the set of the cheapest suppliers.

When the planning model is driven by a service constraint, we have:

$\alpha =$ maximum permitted probability of a shortfall ($\leqslant 0.5$) and $z_\alpha = \Phi^{-1}(1-\alpha) \geqslant 0$.

Alternatively, in the (TCM), we have the following pair of cost parameters:

$h =$ cost of carrying an unsold unit at the end of the season;
$b =$ cost associated with any unit of unsatisfied demand.

We need the following additional notation:

$y_i =$ order placed with supplier i, $i = 1, \dots, N$;
$Y = \sum_{i=1}^N y_i =$ aggregate order;
$w_i = y_i/Y =$ fraction of the total order assigned to supplier i, $i = 1, \dots, N$;
$Y_E = \sum_{i=1}^N p_i y_i =$ expected effective supply, i.e., the expected amount of usable supply resulting from the various orders;
$\Psi^S(Y_E)[\Psi^T(Y_E)] =$ minimum cost value in the (SCM) [(TCM)] when selecting an expected effective supply Y_E; and
$I = I^0 + \sum_{i=1}^N y_i X_i - D =$ end-of-the season inventory level.
(If $I < 0$, $-I$ represents the end-of-the season shortfall.)

The CLT-based approximation replaces I by a Normally distributed random variable

$$\tilde{I} = I^0 + Y_E - \left(\epsilon\sqrt{\sum_{i=1}^N \varsigma_i^2 y_i^2 + D}\right), \tag{1}$$

with the same mean and standard deviation as I, i.e.,

$$\Pr[\tilde{I} \leqslant x] = \Phi\left(\frac{x - I^0 - \sum_{i=1}^N p_i y_i + \mu}{\sqrt{\sigma^2 + \sum_{i=1}^N y_i^2 \varsigma_i^2}}\right). \tag{2}$$

The following substantiation of the CLT-based approximation has been proven in Federgruen and Yang (2008); see Corollary 6.2 there.

Lemma 1. *Assuming all orders are placed with suppliers $1, \dots, n$, consider an arbitrary allocation scheme $\{w_{i,n}\}$ such that*

$$\frac{\max_{1\leqslant i\leqslant n} w_{i,n}}{\min_{1\leqslant i\leqslant n} w_{i,n}} \leqslant A \quad \textit{for some constant } A, \tag{3}$$

i.e., the ratio of the largest and the smallest order size remains bounded, as $n \to \infty$. For any $\eta > 0$, there exists a constant C_η such that for all x,

$$\left|\Pr[I \leqslant x] - \Pr[\tilde{I} \leqslant x]\right| \leqslant C_\eta n^{-\eta}.$$

In particular, $\lim_{n\to\infty} |\Pr[I \leqslant x] - \Pr[\tilde{I} \leqslant x]| = 0$.

Thus, approximating the distribution of I by that of $\tilde{I}$, (SCM) can be formulated as

$$\text{(SCM)} \qquad \min_{Y_E} \Psi^S(Y_E)$$

$$\text{s.t. } Y_E \geqslant \mu - I^0 + z_\alpha \sigma, \quad \text{where} \tag{4}$$

$$\Psi^S(Y_E) \stackrel{\text{def}}{=} \min \sum_{i=1}^N c_i y_i \tag{5}$$

$$\text{s.t. } \sum_{i=1}^N p_i y_i = Y_E, \tag{6}$$

$$(Y_E-\mu+I^0)^2-z_\alpha^2\left(\sum_{i=1}^N s_i^2y_i^2\right)$$
$$-z_\alpha^2\sigma^2\geqslant 0, \tag{7}$$
$$y_i\geqslant 0,\quad i=1,\ldots,N. \tag{8}$$

To explain (7) and the lower bound for Y_E in (4), note that the service constraint

$$\alpha\geqslant \Pr[\tilde I\leqslant 0]=\Phi\left(\frac{\mu-I^0-Y_E}{\sqrt{\sigma^2+\sum_{i=1}^N y_i^2s_i^2}}\right)$$
$$=1-\Phi\left(\frac{Y_E+I^0-\mu}{\sqrt{\sigma^2+\sum_{i=1}^N y_i^2s_i^2}}\right)$$
$$\Leftrightarrow\quad \frac{Y_E+I^0-\mu}{\sqrt{\sigma^2+\sum_{i=1}^N y_i^2s_i^2}}\geqslant z_\alpha$$
$$\Leftrightarrow\quad \{(7), Y_E\geqslant \mu-I^0+z_\alpha\sigma\}.$$

(To verify the last equivalence, the $\Rightarrow$ part is immediate by squaring both sides of the inequality $Y_E+I^0-\mu\geqslant z_\alpha\sqrt{\sigma^2+\sum_{i=1}^N y_i^2s_i^2}$ and from $\sqrt{\sigma^2+\sum_{i=1}^N y_i^2s_i^2}\geqslant\sigma$. To verify the $\Leftarrow$ part, (7) $\Rightarrow\{Y_E+I^0-\mu\geqslant z_\alpha\sqrt{\sigma^2+\sum_{i=1}^N y_i^2s_i^2}$ or $Y_E+I^0-\mu\leqslant -z_\alpha\sqrt{\sigma^2+\sum_{i=1}^N y_i^2s_i^2}\}$. However, because $\alpha\leqslant 0.5$, i.e., $z_\alpha\geqslant 0$, the second inequality can be ignored, in view of the constraint $Y_E\geqslant\mu-I^0+z_\alpha\sigma$.) Again, because $z_\alpha\geqslant 0$, when $I^0\geqslant\mu+z_\alpha\sigma$, the service constraint is met without any procurement whatsoever, and $\mathbf{y}^*=\mathbf{0}$ is optimal. We therefore confine ourselves to the case $I^0<\mu+z_\alpha\sigma$.

(1) shows that the end-of-the season inventory level $\tilde I$ has the same distribution as that in an ordinary inventory system that starts the season with an inventory of (Y_E+I^0) units and faces a demand distribution $D'=D+\epsilon\sqrt{\sum_{i=1}^N y_i^2s_i^2}$. Using a standard derivation in inventory theory (see, e.g., Zipkin 2000, §6.2), we obtain the following formulation for (TCM):

$$\text{(TCM)}\quad \min\quad \Psi^T(Y_E)$$
$$\text{s.t.}\quad Y_E\geqslant 0,\quad\text{where} \tag{9}$$
$$\Psi^T(Y_E)\stackrel{\text{def}}{=}\min\sum_{i=1}^N c_iy_i+h(Y_E+I^0-\mu)$$
$$+(b+h)\int_{Y_E+I^0}^{\infty}\bar\Phi\left(\frac{u-\mu}{\sqrt{\sigma^2+\sum_{i=1}^N y_i^2s_i^2}}\right)du$$
$$\text{s.t.}\quad \sum_{i=1}^N p_iy_i=Y_E, \tag{10}$$
$$y_i\geqslant 0,\quad i=1,\ldots,N. \tag{11}$$

The objective functions (5) and (10) for the (SCM) and (TCM), respectively, are based on the assumption that the firm pays for all ordered units, regardless of whether or not it is delivered as a useable unit. The alternative case, where it only pays for effectively delivered units, can be accommodated merely by replacing the coefficients $\{c_i\}$ in (5) and (10) by $\{p_ic_i\}$. Most generally, the risk is shared between the suppliers and the purchasing firm, i.e., a fraction f_i of the per-unit cost c_i applies to every ordered unit and the remaining fraction $(1-f_i)$ only to those that are effectively produced. In this case, the coefficients c_i in (5) and (10) are to be replaced by $\{[f_i+(1-f_i)p_i]c_i\}$.

The key to solving both the (SCM) and (TCM) is the determination of Y_E^*, the optimal level of the expected effective supply. We will show that for any given choice of Y_E, the corresponding set of suppliers and their market shares are easily obtained. This permits one to project the decision variables in the full-optimization models onto the single aggregate supply measure Y_E. Both the (SCM) and (TCM) are then solved by showing that the functions $\Psi^S(\cdot)$ and $\Psi^T(\cdot)$ are strictly convex and differentiable—with a unique optimal effective supply level, Y_E^*—and that these functions and their derivatives are easily evaluated.

3. The Service Constraint Model

Federgruen and Yang (2008, Theorem 6.3) derives the necessary and sufficient condition for the existence of a feasible solution in the (SCM).

LEMMA 2. (a) *A feasible solution exists if and only if it exists under the allocation scheme* $\mathbf{w}^*$, *with*

$$w_{i,n}^*=\frac{p_i/s_i^2}{(\sum_{l=1}^n p_l/s_l^2)},\quad i=1,\ldots,n.$$

(b) *A feasible solution exists if and only if condition* (F) *is satisfied*:

$$\text{(F)}\quad \text{(i) if } I^0\leqslant\mu,\quad \sum_{i=1}^N\gamma_i^{-2}>z_\alpha^2;$$
$$\text{(ii) if } I^0>\mu,\quad \sum_{i=1}^N\gamma_i^{-2}\geqslant z_\alpha^2-(I^0-\mu)^2/\sigma^2.$$

Thus, a set of suppliers is feasible if and only if it contains a sufficiently large number of sufficiently reliable suppliers. More specifically, defining a (hypothetical) supplier with a yield distribution c.v. value of *one* as a "Base Supplier," a supplier with c.v. $=\gamma$ represents γ^{-2} Base Supplier Equivalents (BSE). In the absence of capacity constraints, if the initial inventory is lower than the mean demand, a set of suppliers is feasible if and only if its total number of BSEs is in excess of a critical number, given by a simple function of the permitted shortfall probability only. In particular, feasibility of a set of suppliers depends on the characteristics of the yield distribution via a *single* measure only, i.e., the number of BSEs the set provides. The same observations apply when the initial inventory exceeds the mean demand by, say, s standard deviations of demand, except that the minimum threshold for the number of BSEs is reduced by s^2.

When all suppliers are equally expensive, i.e., the expected cost of an effective unit is identical across all suppliers ($c_1/p_1 = c_2/p_2 = \cdots = c_N/p_N$), $\sum_{i=1}^N c_i y_i = (c_1/p_1)(\sum_{i=1}^N p_i y_i) = (c_1/p_1)Y_E$, i.e., it is optimal to choose $\underline{Y}_E$, the smallest feasible value of Y_E, where $\underline{Y}_E = H(r^*)$, $r^* \stackrel{\text{def}}{=} \sum_{i=1}^N \gamma_i^{-2}$, and

$$H(r) = \begin{cases} \left(1-\dfrac{z_\alpha^2}{r}\right)^{-1}\left[(\mu-I^0)+z_\alpha\sqrt{\dfrac{(\mu-I^0)^2}{r}+\sigma^2\left(1-\dfrac{z_\alpha^2}{r}\right)}\right] & \text{if } r\neq z_\alpha^2, \\ [z_\alpha^2\sigma^2-(I^0-\mu)^2]/[2(I^0-\mu)] & \text{if } r=z_\alpha^2 \text{ and } I^0>\mu; \end{cases} \tag{12}$$

see Federgruen and Yang (2008, Theorem 6.3). Note that $\underline{Y}_E$ depends on the suppliers' characteristics only via a single measure, i.e., the number of BSEs the set of the suppliers represents.

When the effective cost rates $\{c_i/p_i\}$ fail to be identical, it is not necessarily optimal to choose the *smallest* possible value of Y_E: As we will show, this smallest feasible value requires the participation of all N suppliers, and a cheaper solution may be obtained by enlarging the effective supply Y_E while allocating the aggregate order only to some of the less expensive suppliers. As mentioned in §2, we first show how, for any given value of Y_E, the corresponding optimal set of suppliers and their orders can be evaluated effectively. Next, we prove that the function $\Psi^S(\cdot)$ is strictly convex and differentiable with a unique minimum.

Because the mathematical program (5)–(8) is a convex program, the following Karush-Kuhn-Tucker (KKT) conditions are both necessary and sufficient for an optimal solution $\mathbf{y}^*(Y_E)$, for any fixed value of $Y_E > 0$:

$$c_i - \lambda_1 p_i + 2\lambda_2 z_\alpha^2 s_i^2 y_i \geqslant 0, \quad i=1,\ldots,N, \tag{13}$$

$$y_i[c_i - \lambda_1 p_i + 2\lambda_2 z_\alpha^2 s_i^2 y_i] = 0, \quad i=1,\ldots,N, \tag{14}$$

$$\sum_{i=1}^N p_i y_i = Y_E, \tag{15}$$

$$(Y_E-\mu+I^0)^2 - z_\alpha^2\left(\sum_{i=1}^N s_i^2 y_i^2\right) - z_\alpha^2\sigma^2 \geqslant 0, \tag{16}$$

$$\lambda_2\left[(Y_E-\mu+I^0)^2 - z_\alpha^2\left(\sum_{i=1}^N s_i^2 y_i^2\right) - z_\alpha^2\sigma^2\right] = 0, \tag{17}$$

$$y_i \geqslant 0, \quad -\infty < \lambda_1 < +\infty, \quad \lambda_2 \geqslant 0. \tag{18}$$

We call a solution $\{\mathbf{y}, Y_E\}$ *undominated* if it is feasible and satisfies the service constraint as an equality. Note that any optimal solution $\{\mathbf{y}^*, Y_E^*\}$ of (SCM) is undominated: If it satisfies the service constraint as a strict inequality, it follows from the continuity of the function to the left of (16) that any solution $\mathbf{y}' = \beta\mathbf{y}^*$, $Y_E' = \beta Y_E^*$ for $0 < \beta < 1$, with β sufficiently large, is feasible; clearly, the solution $\{\mathbf{y}', Y_E'\}$ is cheaper than $\{\mathbf{y}^*, Y_E^*\}$, contradicting the optimality of the latter. Let $(c/p)^{(2)} = \min\{c_i/p_i\colon c_i/p_i > c_1/p_1\}$ denote the second-lowest effective cost rate, which is well defined, because the effective cost rates fail to be identical, i.e., $S^0 = \{i\colon c_i/p_i = c_1/p_1\} \subsetneq \{1,\ldots,N\}$. We will show that in an optimal undominated solution, $\lambda_1(Y_E)$ and $\lambda_2(Y_E)$ are uniquely determined. $\lambda_1(Y_E)$ may then be interpreted as the marginal cost saving that can be obtained if a marginal unit of expected effective supply could be procured risk and cost free, i.e., without placing orders to any of the (unreliable) suppliers.[3] We refer to it as the *benchmark cost rate*. We will show that in an optimal undominated solution, this *benchmark cost rate* is at least as large as $(c/p)^{(2)}$, and that the optimal supplier base consists of the suppliers whose effective cost rate is strictly below the benchmark cost rate. (If Y_E is part of an undominated solution, its reduction by one unit requires the retention of a supplier whose effective cost rate is larger than the cheapest cost rate; hence $\lambda_1 \geqslant (c/p)^{(2)}$.) For $\lambda \geqslant (c/p)^{(2)}$, define $k^*(\lambda)$ as the number of suppliers whose effective cost rate falls below λ: $k^*(\lambda) = \max\{k\colon c_k/p_k < \lambda\}$.

As mentioned in §2, Theorem 1 below shows that for any value of the expected effective supply Y_E, which is part of an undominated solution, the associated optimal set of orders $\mathbf{y}$ and the benchmark cost rate λ_1 can be determined in closed form, after computing the root of an increasing single variable function. In particular, $\lambda_1(Y_E)$ is the unique root on $[(c/p)^{(2)}, \infty)$ of the equation

$$r(\lambda_1) = \frac{z_\alpha^2 Y_E^2}{(Y_E-\mu+I^0)^2 - z_\alpha^2\sigma^2}, \quad \text{where} \tag{19}$$

$$r(\lambda_1) \stackrel{\text{def}}{=} \frac{[\sum_{i=1}^{k^*(\lambda_1)}(\lambda_1 - c_i/p_i)\gamma_i^{-2}]^2}{[\sum_{i=1}^{k^*(\lambda_1)}(\lambda_1 - c_i/p_i)^2\gamma_i^{-2}]} \quad \text{for } \lambda_1 \geqslant \frac{c}{p}^{(2)}. \tag{20}$$

We first need the following lemma, which shows that $r(\cdot)$ increases from the number of BSEs represented by the cheapest suppliers to the number of BSEs represented by all. In view of this property, $r(\lambda_1)$ may be interpreted as the modified number of BSEs, given the benchmark cost rate λ_1.

Lemma 3. *On its domain $[(c/p)^{(2)}, \infty)$, $r(\cdot)$ is continuously strictly increasing from $\sum_{i\in S^0}\gamma_i^{-2}$ to $\sum_{i=1}^N \gamma_i^{-2}$.*

Theorem 1. *Let Y_E be part of an undominated solution.*

(a) *The mathematical program (5)–(8) has a unique optimal solution, given by*

$$p_i y_i^* = \begin{cases} \dfrac{(\lambda_1 - c_i/p_i)\gamma_i^{-2}}{\sum_{l=1}^{k^*(\lambda_1)}(\lambda_1 - c_l/p_l)\gamma_l^{-2}} Y_E, & i=1,\ldots,k^*(\lambda_1), \\ 0, & i=k^*(\lambda_1)+1,\ldots,N, \end{cases} \tag{21}$$

where λ_1 is the unique root of Equation (19) on $[(c/p)^{(2)}, \infty)$.

(b) *The set of retained suppliers in the optimal solution is consecutive and*

$$w_i^*(\lambda_1) \stackrel{\text{def}}{=} y_i^* \Big/ \sum_{l=1}^N y_l^* = \frac{[(\lambda_1 - c_i/p_i)^+]p_i/\varsigma_i^2}{\sum_{l=1}^N [(\lambda_1 - c_l/p_l)^+]p_l/\varsigma_l^2} \quad \text{for } \lambda_1 \geqslant \frac{c}{p}^{(2)}. \tag{22}$$

Thus, the optimal supplier base is always consecutive, and consists of all suppliers whose effective cost rate is below the above-defined *benchmark cost rate*, $\lambda_1(Y_E)$. Also, as mentioned in the introduction, each supplier in the selected supplier base is assigned an overall score, given by the product of two factors: The first factor is the mean-to-variance ratio of the supplier's yield distribution; the second factor is the net cost saving, relative to the benchmark cost rate. A supplier's market share is given by his overall score relative to the sum of the selected suppliers' scores.

The next lemma shows that the search for an optimal value Y_E^* can be restricted to an interval $(\underline{Y}_E, \overline{Y}_E]$, with $\overline{Y}_E \leqslant \infty$, such that every $Y_E \in (\underline{Y}_E, \overline{Y}_E)$ is part of an undominated solution, and $\lambda_1(Y_E)$ is a decreasing function on this interval.

LEMMA 4. *Assume that* (F) *holds.* (a) *Let* $\underline{\lambda}_1 \stackrel{\text{def}}{=} \min\{\lambda_1 \geqslant (c/p)^{(2)}: r(\lambda_1) \geqslant z_\alpha^2 - [(I^0 - \mu)^+]^2/\sigma^2\} < \infty$. *For* $\lambda_1 > \underline{\lambda}_1$, *let* $Y_E(\lambda_1)$ *denote the smallest root of the quadratic Equation* (19) (*in* Y_E) *which satisfies* (4), *i.e.,* $Y_E(\lambda_1) = H(r(\lambda_1))$, *with* $H(\cdot)$ *defined by* (12). $Y_E(\lambda_1)$ *is continuous and strictly decreasing on* $(\underline{\lambda}_1, \infty)$. *For* $\lambda_1 > \underline{\lambda}_1$, *if* $Y_E = Y_E(\lambda_1)$, *the mathmatical program* (5)–(8) *has an optimal solution with* λ_1 *as the Lagrangean multiplier of constraint* (7).

(b) $\overline{Y}_E = \lim_{\lambda_1 \downarrow \underline{\lambda}_1} H(r(\lambda_1))$, *where*

$$\lim_{\lambda_1 \downarrow \underline{\lambda}_1} H(r(\lambda_1)) = \begin{cases} \infty & \text{if } \underline{r} = z_\alpha^2 \text{ and } I^0 \leqslant \mu, \\ [z_\alpha^2\sigma^2 - (I^0-\mu)^2]/[2(I^0-\mu)] & \text{if } \underline{r} = z_\alpha^2 \text{ and } I^0 > \mu, \\ \left(1 - \dfrac{z_\alpha^2}{\underline{r}}\right)^{-1}\left[(\mu - I^0) + z_\alpha\sqrt{\dfrac{(\mu-I^0)^2}{\underline{r}} + \sigma^2\left(1 - \dfrac{z_\alpha^2}{\underline{r}}\right)}\right] & \text{if } \underline{r} \neq z_\alpha^2, \end{cases} \tag{23}$$

and

$$\underline{r} = \lim_{\lambda_1 \downarrow \underline{\lambda}_1} r(\lambda_1) = \max\left\{ r\left(\frac{c}{p}^{(2)}\right); z_\alpha^2 - \frac{[(I^0-\mu)^+]^2}{\sigma^2}\right\}. \tag{24}$$

The range of $Y_E(\cdot)$: $(\underline{\lambda}_1, \infty) \to \mathbb{R}$ *is given by* $\mathbb{Y} = (\underline{Y}_E, \overline{Y}_E)$. *An optimal solution for the* (SCM) *exists for some* $Y_E \in (\underline{Y}_E, \overline{Y}_E]$, *and every value in this interval is part of an undominated solution.*

The next theorem shows that the unique value λ_1 associated with any $Y_E \in \mathbb{Y} \cup \{\overline{Y}_E\}$, and hence $\Psi^S(Y_E) = \sum_{i=1}^N c_i y_i^*$, can, in fact, be obtained in closed form. This closed-form expression is obtained by showing that $\mathbb{Y}$ can be partitioned into at most N consecutive intervals, such that the same value k (and hence the same set of) suppliers is optimal for all Y_E in the same interval. Finally, the theorem shows that the function $\Psi^S(Y_E)$ is differentiable and strictly convex.

THEOREM 2. (a) *Let* $\underline{k} = k^*(\underline{\lambda}_1)$, $Y_E^{\underline{k}-1} = \overline{Y}_E$, $Y_E^N = \underline{Y}_E$, *and for* $k = \underline{k}, \ldots, N-1$: $Y_E^k = Y_E(c_{k+1}/p_{k+1})$. *Let* $\hat{k}(Y_E)$ *denote the optimal number of suppliers to be used for a given expected effective supply* Y_E:

$$\hat{k}(Y_E) = \begin{cases} k & \text{if } Y_E^k \leqslant Y_E < Y_E^{k-1},\ k = \underline{k}, \ldots, N, \\ \underline{k} & \text{if } Y_E = \overline{Y}_E. \end{cases} \tag{25}$$

If $Y_E = \overline{Y}_E$, $\lambda_1(Y_E) = \underline{\lambda}_1$, *and for* $Y_E \in \mathbb{Y}$,

$$\lambda_1 = \frac{\sum_{i=1}^{\hat{k}}(c_i/p_i)/\gamma_i^2}{(\sum_{i=1}^{\hat{k}} 1/\gamma_i^2)} + \left(\sum_{i=1}^{\hat{k}} 1/\gamma_i^2\right)^{-1} z_\alpha Y_E \cdot \sqrt{\frac{(\sum_{i=1}^{\hat{k}}(c_i/p_i)^2/\gamma_i^2)(\sum_{i=1}^{\hat{k}}1/\gamma_i^2) - (\sum_{i=1}^{\hat{k}}(c_i/p_i)/\gamma_i^2)^2}{[(Y_E-\mu+I^0)^2 - z_\alpha^2\sigma^2](\sum_{i=1}^{\hat{k}}1/\gamma_i^2) - z_\alpha^2 Y_E^2}}. \tag{26}$$

(b) *The function* $\Psi^S(Y_E)$ *is strictly convex and differentiable on* $\mathbb{Y} \cup \{\overline{Y}_E\}$, *with*

$$\Psi^S(Y_E) = \frac{\sum_{i=1}^{\hat{k}}(c_i/p_i)/\gamma_i^2}{(\sum_{i=1}^{\hat{k}} 1/\gamma_i^2)} Y_E - \left(\sum_{i=1}^{\hat{k}} \frac{1}{\gamma_i^2}\right)^{-1} z_\alpha^{-1} \cdot \left(\left[\left(\sum_{i=1}^{\hat{k}} \frac{(c_i/p_i)^2}{\gamma_i^2}\right)\left(\sum_{i=1}^{\hat{k}} \frac{1}{\gamma_i^2}\right) - \left(\sum_{i=1}^{\hat{k}} \frac{c_i/p_i}{\gamma_i^2}\right)^2\right] \cdot \left[[(Y_E-\mu+I^0)^2 - z_\alpha^2\sigma^2]\left(\sum_{i=1}^{\hat{k}} \frac{1}{\gamma_i^2}\right) - z_\alpha^2 Y_E^2\right]\right)^{1/2}, \tag{27}$$

$$\Psi^{S\prime}(Y_E) = \lambda_1(Y_E) - 2(Y_E - \mu + I^0)\lambda_2(Y_E). \tag{28}$$

($\lambda_2(Y_E)$ *is defined by* (43); *if* $Y_E = \overline{Y}_E$, $\Psi^{S\prime}(Y_E)$ *denotes the left-hand derivative.*)

(c) *If* $\underline{Y}_E < Y_E^* < \overline{Y}_E$, *then* $Y_E^* =$

$$\begin{cases} \dfrac{(\sum_{i=1}^{k^*} 1/\gamma_i^2)(\mu - I^0)}{(\sum_{i=1}^{k^*} 1/\gamma_i^2 - z_\alpha^2)} + \dfrac{z_\alpha^2(\sum_{i=1}^{k^*}(c_i/p_i)/\gamma_i^2)}{(\sum_{i=1}^{k^*} 1/\gamma_i^2 - z_\alpha^2)} \cdot \sqrt{\dfrac{\sigma^2(\sum_{i=1}^{k^*}1/\gamma_i^2 - z_\alpha^2) + (\mu - I^0)^2}{(\sum_{i=1}^{k^*}(c_i/p_i)/\gamma_i^2)^2 - (\sum_{i=1}^{k^*}(c_i/p_i)^2/\gamma_i^2)(\sum_{i=1}^{k^*}1/\gamma_i^2 - z_\alpha^2)}} & \text{if } \sum_{i=1}^{k^*} 1/\gamma_i^2 \neq z_\alpha^2, \\ \left[z_\alpha^2\sigma^2 - \left(2 - \dfrac{(\sum_{i=1}^{k^*}(c_i/p_i)^2/\gamma_i^2)(\sum_{i=1}^{k^*}1/\gamma_i^2)}{(\sum_{i=1}^{k^*}(c_i/p_i)/\gamma_i^2)^2}\right)(I^0 - \mu)^2\right] \Big/ [2(I^0 - \mu)] & \text{if } \sum_{i=1}^{k^*} 1/\gamma_i^2 = z_\alpha^2 \text{ and } I^0 > \mu. \end{cases} \tag{29}$$

Algorithm 1 (Algorithm SCM).

Step 0. (Calculation of $\underline{k}$, the *minimum* number of required suppliers and $\overline{Y}_E$, the *maximum* expected effective supply.)

$$\underline{k} := \min\left\{k \geqslant |S^0| \colon r\left(\frac{c_{k+1}}{p_{k+1}}\right) \geqslant z_\alpha^2 - \frac{[(I^0-\mu)^+]^2}{\sigma^2}\right\},$$

$$\underline{r} := \max\left\{r\left(\frac{c^{(2)}}{p}\right), z_\alpha^2 - \frac{[(I^0-\mu)^+]^2}{\sigma^2}\right\}$$

(see (24)),

$$\overline{Y}_E := \begin{cases} \infty \quad \text{if } \underline{r} = z_\alpha^2 \text{ and } I^0 \leqslant \mu, \\ [z_\alpha^2\sigma^2 - (I^0-\mu)^2]/[2(I^0-\mu)] \\ \qquad \text{if } \underline{r} = z_\alpha^2 \text{ and } I^0 > \mu, \\ \left(1-\dfrac{z_\alpha^2}{\underline{r}}\right)^{-1}\left[(\mu - I^0) + z_\alpha\sqrt{\dfrac{(\mu-I^0)^2}{\underline{r}} + \sigma^2\left(1-\dfrac{z_\alpha^2}{\underline{r}}\right)}\right] \\ \qquad \text{if } \underline{r} \neq z_\alpha^2 \end{cases}$$

(see (23)),

Step 1. FOR $k := N-1$ DOWNTO $\underline{k}$ DO
BEGIN

$$Y_E^k = Y_E(c_{k+1}/p_{k+1})$$

$$= \begin{cases} \left(1-\dfrac{z_\alpha^2}{r(c_{k+1}/p_{k+1})}\right)^{-1}\left[(\mu - I^0) + z_\alpha\left(\dfrac{(\mu-I^0)^2}{r(c_{k+1}/p_{k+1})} \right.\right. \\ \qquad\qquad \left.\left. + \sigma^2\left(1-\dfrac{z_\alpha^2}{r(c_{k+1}/p_{k+1})}\right)\right)^{1/2}\right] \\ \quad \text{if } r\left(\dfrac{c_{k+1}}{p_{k+1}}\right) \neq z_\alpha^2, \\ [z_\alpha^2\sigma^2 - (I^0-\mu)^2]/[2(I^0-\mu)] \\ \quad \text{if } r\left(\dfrac{c_{k+1}}{p_{k+1}}\right) = z_\alpha^2 \text{ and } I^0 > \mu \end{cases}$$

(see (12)),

$$\lambda_1^k = c_{k+1}/p_{k+1}, \quad \lambda_2^k = \frac{\sum_{i=1}^k (\lambda_1^k - c_i/p_i)\gamma_i^{-2}}{2z_\alpha^2 Y_E^k} \quad \text{(see (43))},$$

$$\Psi^{S'}(Y_E^k) = \lambda_1^k - 2(Y_E^k - \mu + I^0)\lambda_2^k \quad \text{(see (28))}.$$

IF $\Psi^{S'}(Y_E^k) \geqslant 0$ THEN GO TO Step 3
END

Step 2. $k := \underline{k} - 1$; IF $\Psi^{S'}(\overline{Y}_E) > 0$ THEN GO TO Step 3; $Y_E^* := \overline{Y}_E$; $\hat{k}(Y_E^*) := \underline{k}$; GO TO Step 4.

Step 3. $\hat{k}(Y_E^*) := k+1$;

$$Y_E^* = \begin{cases} \dfrac{(\sum_{i=1}^{\hat{k}} 1/\gamma_i^2)(\mu - I^0)}{(\sum_{i=1}^{\hat{k}} 1/\gamma_i^2 - z_\alpha^2)} + \dfrac{z_\alpha^2(\sum_{i=1}^{\hat{k}} (c_i/p_i)/\gamma_i^2)}{(\sum_{i=1}^{\hat{k}} 1/\gamma_i^2 - z_\alpha^2)} \\ \cdot\sqrt{\dfrac{\sigma^2(\sum_{i=1}^{\hat{k}} 1/\gamma_i^2 - z_\alpha^2) + (\mu - I^0)^2}{(\sum_{i=1}^{\hat{k}}(c_i/p_i)/\gamma_i^2)^2 - (\sum_{i=1}^{\hat{k}}(c_i/p_i)^2/\gamma_i^2)(\sum_{i=1}^{\hat{k}} 1/\gamma_i^2 - z_\alpha^2)}} \\ \quad \text{if } \displaystyle\sum_{i=1}^{\hat{k}} 1/\gamma_i^2 \neq z_\alpha^2, \\ \left[z_\alpha^2\sigma^2 - \left(2 - \dfrac{(\sum_{i=1}^{\hat{k}}(c_i/p_i)^2/\gamma_i^2)(\sum_{i=1}^{\hat{k}} 1/\gamma_i^2)}{(\sum_{i=1}^{\hat{k}}(c_i/p_i)/\gamma_i^2)^2}\right) \right. \\ \quad \left. \cdot (I^0-\mu)^2\right] \Big/ [2(I^0-\mu)] \\ \quad \text{if } \displaystyle\sum_{i=1}^{\hat{k}} 1/\gamma_i^2 = z_\alpha^2 \text{ and } I^0 > \mu \end{cases}$$

(see (29)).

Step 4. FOR $k := 1$ TO $\hat{k}$ DO $y_k^* = Y_E^*[(\lambda_1 - c_k/p_k)p_k/s_k^2][\sum_{i=1}^{\hat{k}}(\lambda_1 - c_i/p_i)\gamma_i^{-2}]$, where

$$\lambda_1 = \begin{cases} \dfrac{c^{(2)}}{p} \quad \text{if } \hat{k}(Y_E^*) = |S^0|, \\ \dfrac{\sum_{i=1}^{\hat{k}}(c_i/p_i)/\gamma_i^2}{(\sum_{i=1}^{\hat{k}} 1/\gamma_i^2)} + \left(\displaystyle\sum_{i=1}^{\hat{k}} 1/\gamma_i^2\right)^{-1} z_\alpha Y_E^* \\ \cdot\sqrt{\dfrac{(\sum_{i=1}^{\hat{k}}(c_i/p_i)^2/\gamma_i^2)(\sum_{i=1}^{\hat{k}} 1/\gamma_i^2) - (\sum_{i=1}^{\hat{k}}(c_i/p_i)/\gamma_i^2)^2}{[(Y_E^* - \mu + I^0)^2 - z_\alpha^2\sigma^2](\sum_{i=1}^{\hat{k}} 1/\gamma_i^2) - z_\alpha^2(Y_E^*)^2}}, \\ \qquad \text{otherwise.} \end{cases}$$

FOR $k := \hat{k}+1$ TO N DO $y_k^* = 0$, see (21).

Algorithm SCM is a succinct description of a highly efficient optimization algorithm for (SCM). Even when $N = 20$ and the optimal number of suppliers $k^* \geqslant 15$, a typical problem instance of (SCM) can be solved by our algorithm in a few milliseconds when implemented on a Dell Optiplex GX620 computer with Pentium D CPU of 3.00 GHz and 3.5 GB of RAM. It is easily verified that the complexity of the (SCM) Algorithm depends on k^*, and is only quadratic in the latter.

4. The Total Cost Model

As in the (SCM), we design an efficient solution method for the (TCM) as follows: We first prove that the function $\Psi^T(\cdot)$ is strictly convex and differentiable with a unique minimum Y_E^*. We next show that for any given expected effective supply level Y_E, the optimal cost value $\Psi^T(Y_E)$ as well as the associated set of suppliers and their orders can be evaluated efficiently. Note first that the nonlinear part of the objective function (10) is given by $G(Y_E, \mathbf{y}) \stackrel{\text{def}}{=} \int_{Y_E+I^0}^{\infty} \bar{\Phi}((u-\mu)/\sqrt{\sigma^2 + \sum_{i=1}^N y_i^2 s_i^2})\,du$, a jointly strictly convex function, as shown by the following lemma.

LEMMA 5. *$G(Y_E, \mathbf{y})$ is strictly convex.*

The lemma implies that the (TCM)'s objective function, when written as a function of the order vector $\mathbf{y}$ alone, is jointly strictly convex as well. Agrawal and Nahmias (1997) proved this property for the case $N = 2$ and conjectured it for general values of N.

As with its counterpart [(5)–(8)], the mathematical program (10)–(11), which defines $\Psi^T(Y_E)$, is a convex program because $G(Y_E, \mathbf{y})$ is jointly convex in $\mathbf{y}$. The KKT conditions are therefore, again, both necessary and sufficient for an optimal solution $\mathbf{y}^*(Y_E)$, for any fixed value $Y_E > 0$:

$$c_i - \lambda p_i + (b+h)\phi\left(\frac{Y_E + I^0 - \mu}{\sqrt{\sigma^2 + \sum_{l=1}^N y_l^2 s_l^2}}\right) \cdot \frac{y_i s_i^2}{\sqrt{\sigma^2 + \sum_{l=1}^N y_l^2 s_l^2}} \geqslant 0, \quad i = 1, \ldots, N, \tag{30}$$

$$y_i\left[c_i - \lambda p_i + (b+h)\phi\left(\frac{Y_E + I^0 - \mu}{\sqrt{\sigma^2 + \sum_{l=1}^N y_l^2 s_l^2}}\right) \cdot \frac{y_i s_i^2}{\sqrt{\sigma^2 + \sum_{l=1}^N y_l^2 s_l^2}}\right] = 0, \quad i = 1, \ldots, N, \tag{31}$$

$$\sum_{l=1}^N p_l y_l = Y_E, \quad \mathbf{y} \geqslant 0, \quad -\infty < \lambda < +\infty. \tag{32}$$

(To verify that the left-hand side of (30) represents the derivative of the Lagrangean with respect to y_i is analogous to the derivation of $\partial H/\partial \Sigma$ in the proof of Lemma 5.) $\lambda(Y_E)$ may, again, be interpreted as the marginal cost saving that can be obtained if the Y_Eth unit of expected effective supply could be procured risk and cost free, i.e., without placing orders with any of the (unreliable) suppliers.

THEOREM 3. *Fix $Y_E > 0$ and let $\mathbf{y}^*$ be an optimal solution to* (10)–(11).

(a) *$\lambda > c_1/p_1$ and*

$$p_i y_i^* = \begin{cases} \dfrac{(\lambda - c_i/p_i)\gamma_i^{-2}}{\sum_{l=1}^{k^*(\lambda)}(\lambda - c_l/p_l)\gamma_l^{-2}} Y_E, & i = 1, \ldots, k^*(\lambda), \\ 0, & i = k^*(\lambda)+1, \ldots, N, \end{cases} \tag{33}$$

where λ is the unique root of the equation $L(\lambda \mid Y_E) = 0$, and the function

$$L(\lambda \mid Y_E) \stackrel{\text{def}}{=} \ln\left\{\sigma^2\left[\sum_{l=1}^{k^*(\lambda)}(\lambda - c_l/p_l)\gamma_l^{-2}\right]^2 + Y_E^2\left[\sum_{l=1}^{k^*(\lambda)}(\lambda - c_l/p_l)^2\gamma_l^{-2}\right]\right\} + \ln\frac{2\pi}{(b+h)^2} - 2\ln Y_E + \frac{(Y_E + I^0 - \mu)^2[\sum_{l=1}^{k^*(\lambda)}(\lambda - c_l/p_l)\gamma_l^{-2}]^2}{\sigma^2[\sum_{l=1}^{k^*(\lambda)}(\lambda - c_l/p_l)\gamma_l^{-2}]^2 + Y_E^2[\sum_{l=1}^{k^*(\lambda)}(\lambda - c_l/p_l)^2\gamma_l^{-2}]} \tag{34}$$

is strictly increasing in λ.

(b) *There exists an optimal solution in which the set of contributing suppliers is consecutive and*

$$\frac{y_i^*}{\sum_{l=1}^N y_l^*} = w_i^*(\lambda), \quad \textit{where}$$

$$w_i^*(\lambda) = \left\{\frac{[(\lambda - c_i/p_i)^+]p_i/s_i^2}{\sum_{l=1}^N[(\lambda - c_l/p_l)^+]p_l/s_l^2}\right\}.$$

Thus, as in the (SCM), the optimal supplier base is consecutive, a property first demonstrated for the (TCM) by Dada et al. (2007). More specifically, Theorem 3 reveals that λ may again be interpreted as a benchmark cost rate, and the optimal supplier base consists of all suppliers whose effective cost rate is below the benchmark value. The market share of each supplier in the selected base is again given by the product of the same reliability and cost scores as in the (SCM): the first factor is the mean-to-variance ratio of the supplier's yield distribution; the second factor is the net cost saving, relative to the benchmark cost rate.

The next theorem identifies a simple upper bound for Y_E^*:

$$Y_E^* \leqslant \bar{Y}_E \stackrel{\text{def}}{=} \frac{p_1}{c_1}\left[h(I^0 - \mu) + (b+h)\int_{I^0}^{\infty}\bar{\Phi}\left(\frac{u-\mu}{\sigma}\right)du\right]. \tag{35}$$

Note that the function $k^*(\lambda)$ is discontinuous in any of the critical cost rates $\{c_i/p_i\}$. As a consequence, $\mathbf{y}^*(Y_E)$ and $\lambda(Y_E)$ fail to be differentiable in values of Y_E, for which $\lambda(Y_E)$ equals one of these critical cost rates. The next theorem shows, nevertheless, that the function $\Psi^T(Y_E)$ is strictly convex and differentiable and that its derivative is easily computed. Because, as we will show, $\Psi^{T\prime}(Y_E) > 0$ for some Y_E, it follows from the strict convexity of Ψ^T that the optimal expected effective supply Y_E^*, along with the corresponding optimal vector of orders $\mathbf{y}^*$, can be found simply by determining, via bisection, $Y_E^* = \min\{0 \leqslant Y_E \leqslant \bar{Y}_E \colon \Psi^{T\prime}(Y_E) \geqslant 0\}$.

THEOREM 4. (a) *$0 \leqslant Y_E^* \leqslant \bar{Y}_E$.*

(b) *$\Psi^T(\cdot)$ is strictly convex.*

(c) *$\Psi^T(Y_E)$ is differentiable with $\Psi^{T\prime}(Y_E) = \lambda(Y_E) + h - (b+h)\bar{\Phi}((Y_E + I^0 - \mu)/\sqrt{\sigma^2 + \sum_{i=1}^{k^*(\lambda)} y_i^2 s_i^2})$, or*

$$\Psi^{T\prime}(Y_E) = \lambda(Y_E) + h - (b+h)\bar{\Phi}\left((Y_E + I^0 - \mu) \cdot \left(\left(\sigma^2 + Y_E^2\left[\sum_{i=1}^{k^*(\lambda)}(\lambda(Y_E) - c_i/p_i)^2\gamma_i^{-2}\right] \Bigg/ \left[\sum_{i=1}^{k^*(\lambda)}(\lambda(Y_E) - c_i/p_i)\gamma_i^{-2}\right]^2\right)^{1/2}\right)^{-1}\right). \tag{36}$$

(d) *If $(I^0 - \mu)/\sigma \geqslant \Phi^{-1}((b - c_1/p_1)/(b+h))$, $Y_E^* = 0$. If $(I^0 - \mu)/\sigma < \Phi^{-1}((b - c_1/p_1)/(b+h))$, Y_E^* is the unique root of $\Psi^{T\prime}(Y_E) = 0$.*

Based on the above theorem, Algorithm TCM below determines Y_E^* as the unique root of the function $\Psi^{T\prime}$. The algorithm has considerable similarity to its counterpart for the (SCM), with only the following two exceptions: (i) the benchmark cost rate $\lambda(Y_E)$ for any given value of Y_E cannot be obtained in closed form, but must be computed as the unique root of a nonlinear increasing function. (Recall, in the (SCM), that (26) provides a closed-form expression for $\lambda_1(Y_E)$.); (ii) the optimal number of suppliers $k^*(Y_E)$ must be recalculated for every trial value of Y_E, rather than being given directly by the position of Y_E vis-à-vis a set of up to $N-1$ breakpoints $\{Y_E^k\}$. Recall that in the (SCM), the optimal number of suppliers k^* is nonincreasing in Y_E, i.e., when considering a larger expected effective supply, the corresponding order is optimally allocated to the same or a smaller number of suppliers. It is this monotonicity property that allows for the determination of critical values $\{Y_E^k\}$ such that $k^*(Y_E)=k$ if and only if $Y_E^k \leqslant Y_E < Y_E^{k-1}$ $(k=1,\ldots,N)$; as examples in Online Appendix B exhibit, it fails, in general, to hold in the (TCM) because $\lambda(Y_E)$, the benchmark cost rate, may fail to be monotone. However, when $I^0 \leqslant \mu$ and $\mu > \sigma$, and if one restricts oneself to values of $Y_E \geqslant \mu - I^0 + \sigma$ (corresponding with an expected safety stock of at least one standard deviation of the demand distribution), the following patterns can be proven:

(i) When $I^0 = 0$ or is sufficiently small, $\lambda(Y_E)$ is decreasing.

(ii) For all $I^0 \leqslant \mu$, $\lambda(Y_E)$ is either monotonically decreasing, or it first increases until reaching a maximum and is monotonically decreasing thereafter.

To verify these observations, recall from Theorem 3(a) that $\lambda(Y_E)$ is the unique root of the equation $L(\lambda \mid Y_E)=0$ so that by the implicit function theorem, $\lambda'(Y_E)$ has the opposite sign as

$$\begin{aligned}
&\partial L(\lambda \mid Y_E)/\partial Y_E \\
&= \left(2\left[\sum_{l=1}^{k^*(\lambda)} (\lambda - c_l/p_l)\gamma_l^{-2}\right]^4 \left\{\sigma^2[(Y_E+I^0-\mu)Y_E-\sigma^2]\right.\right. \\
&\qquad \left.\left. + \left[\sum_{l=1}^{k^*(\lambda)} y_l^{*2}\varsigma_l^2\right][(Y_E+I^0-\mu)(\mu-I^0)-\sigma^2]\right\}\right) \\
&\quad \cdot \left(Y_E\left\{\sigma^2\left[\sum_{l=1}^{k^*(\lambda)} (\lambda - c_l/p_l)\gamma_l^{-2}\right]^2\right.\right. \\
&\qquad \left.\left. + Y_E^2\left[\sum_{l=1}^{k^*(\lambda)} (\lambda - c_l/p_l)^2\gamma_l^{-2}\right]\right\}^2\right)^{-1}.
\end{aligned}$$

Observations (i) and (ii) follow readily. Thus, when $I^0 \leqslant \mu$, it pays to calculate the (at most $2N$) break points $\{Y_E^k\}$ such that k^* remains constant in-between consecutive break points.

We have conducted a numerical study to investigate whether the (TCM) Algorithm, which is based on the CLT-approximation for the end-of-the period inventory-level distribution, finds a solution that is close to the exact optimum. The study employs 80 instances, all with $N=4$ potential suppliers. For $N=4$, it is still possible to find the exact optimal solution with a general purpose algorithm, for example, one based on sequential quadratic programming, where the cost associated with any given vector of orders is determined via a Monte Carlo simulation. With $N=4$, the "exact" simulation-based optimization algorithm takes approximately 10 CPU minutes per instance when implemented on a Dell Optiplex GX620 computer with Pentium D CPU of 3.00 GHz and 3.5 GB of RAM. (In contrast, Algorithm TCM takes less than 2 CPU seconds, when run on the same platform.) At the same time, determining the exact optimal solution becomes prohibitively expensive when the number of suppliers is $N \geqslant 20$, say. However, if Algorithm TCM generates very close to optimal solutions, when $N=4$, it is at least as accurate when the number of suppliers is larger. It is easily verified that the complexity of the Algorithm TCM depends on k^*, the actual number of suppliers used, and is only quadratic in the latter. (Most instances with $N=20$ suppliers can be solved in less than 2 CPU seconds, as in the case where $N=4$.)

Algorithm 2 (Algorithm TCM).

Step 0. IF $(I^0-\mu)/\sigma \geqslant \Phi^{-1}((b-c_1/p_1)/(b+h))$
  THEN BEGIN $y^* = Y_E^* = 0$; EXIT END
  $Y_1 := 0$; $Y_2 := \overline{Y}_E := (p_1/c_1)[h(I^0-\mu)+(b+h) \cdot \int_{I^0}^{\infty} \bar{\Phi}((u-\mu)/\sigma)\,du]$

Step 1. WHILE $((Y_2 - Y_1) > \epsilon_1)$ DO
  BEGIN
   $Y := (Y_1+Y_2)/2$; $k := 2$; $k^* := 1$
   WHILE $(L(c_k/p_k \mid Y) < 0$ and $k \leqslant N)$ DO
    BEGIN $k := k+1$; $k^* := k^*+1$; END
   $\lambda_1 := c_{k-1}/p_{k-1}$;
   IF $k \leqslant N$ THEN $\lambda_2 := c_k/p_k$ ELSE $\lambda_2 := \bar{\lambda}$;
   WHILE $((\lambda_2-\lambda_1) > \epsilon_2)$ DO
    BEGIN
     $\lambda := (\lambda_1+\lambda_2)/2$;
     IF $(L(\lambda \mid Y) < 0)$ THEN $\lambda_1 := \lambda$ ELSE $\lambda_2 := \lambda$
    END
   $\lambda := (\lambda_1+\lambda_2)/2$;
   FOR $i=1$ TO k^* DO
    $y_i^* = Y\{[(\lambda - c_i/p_i)]p_i/\varsigma_i^2\}/\{\sum_{l=1}^{k^*}[(\lambda - c_l/p_l)]\gamma_l^{-2}\}$
   IF $(\Psi^{T\prime}(Y) = \lambda + h - (b+h)\bar{\Phi}((Y+I^0-\mu)/\sqrt{\sigma^2+\sum_{i=1}^{k^*} y_i^{*2}\varsigma_i^2}) < 0)$
   THEN $Y_1 := Y$
   ELSE $Y_2 := Y$
  END
  $Y_E^* = Y$

The 80 instances evaluated in the study all have $N=4$, $\mu=20$, $h=1$, and uniform yield distributions on intervals $[\underline{p}_i, \bar{p}_i]$. Here, $\underline{p}_i$ and $\bar{p}_i$ are uniformly generated from the intervals $[0.6, 0.85]$ and $[0.9, 0.99]$, respectively. Finally, the effective cost rates $\{c_i/p_i\}$ are uniformly

Table 1. Comparison of Algorithm TCM and an "exact" simulation-based algorithm.

Group	I^0	σ	b	Mean(T) (TCM)	Max(T) (TCM)	Mean(T) sim	Max(T) sim	Mean ΔC (%)	Max ΔC (%)	Mean ΔY_E (%)	Max ΔY_E (%)	Mean $\Delta \mathbf{y}^*$ (%)	Max $\Delta \mathbf{y}^*$ (%)
1	0	5	6	1.86	1.92	591.06	609.57	0.38	0.76	0.58	2.46	0.75	4.31
2	0	5	10	1.95	2.00	522.66	631.72	0.50	0.81	0.47	1.58	0.66	3.31
3	0	10	6	1.90	1.99	535.82	623.21	0.27	0.42	0.74	2.15	0.80	2.39
4	0	10	10	1.89	1.95	471.46	627.67	0.30	0.43	0.35	1.69	0.37	1.69
5	10	5	6	1.94	1.99	589.09	623.43	0.28	0.59	0.23	0.90	0.27	0.90
6	10	5	10	1.92	2.00	524.63	652.08	0.23	0.62	0.21	0.82	0.19	0.63
7	10	10	6	1.90	1.95	589.02	623.86	0.25	0.36	0.18	1.06	0.23	0.86
8	10	10	10	1.89	1.96	552.79	622.99	0.23	0.46	0.67	1.84	0.67	1.76
Overall				1.91	2.00	547.07	652.08	0.31	0.81	0.43	2.46	0.49	4.31

generated from the interval $[1, 4]$. The 80 instances are partitioned into 8 groups of 10, each with a different combination of values for the three remaining parameters, I^0, σ, and b. (For each of these, two distinct values are considered.) Let $(\mathbf{y}^*, Y_E^*)$ and $(\hat{\mathbf{y}}, \hat{Y}_E)$ denote the solution generated by Algorithm TCM and the simulation-based "exact" search method, respectively. Table 1 exhibits, for each of the eight groups of instances, the mean and maximum of the CPU times (in seconds) for Algorithm TCM (Mean(T), (TCM) and Max(T), (TCM)); the mean and maximum of the CPU time of the "exact" simulation-based search method (Mean(T), Sim and Max(T), Sim); the mean and maximum of the relative difference in the expected costs of the two solutions (Mean ΔC and Max ΔC); the mean and maximum of $|Y_E^* - \hat{Y}_E|/\hat{Y}_E$ (Mean ΔY_E and Max ΔY_E); as well as the mean and maximum of the following relative distance measure between the order vectors $\mathbf{y}^*$ and $\hat{\mathbf{y}}^*$: $\sqrt{\sum_{i=1}^N (y_i^* - \hat{y}_i)^2} / \sum_{i=1}^N \hat{y}_i$. (The latter pair of measures is referred to as Mean $\Delta\mathbf{y}^*$ and Max $\Delta\mathbf{y}^*$.) We conclude that Algorithm TCM generates solutions that are very close to optimal, not just in terms of the cost value and the expected effective supply, but also in terms of the individual order sizes.

5. The Impact of Initial Inventory, Demand and Supply Risks, and the Benefit of Additional Suppliers

In this section, we discuss how the optimal cost value, the (effective) order size, and supplier base vary with the supply risks, the demand magnitude and risks, and the initial inventory. We also show that although in general the marginal benefit of an additional supplier decreases when the supplier is added to a larger potential base, this property may fail to hold in some extreme cases.

Supply Risks

The following proposition shows, both for the (SCM) and (TCM), that if a supplier is insufficiently competitive, i.e., does not share in the buyer's orders, it cannot join the supplier base merely by improving its own supply risk, i.e., the standard deviation of its yield distribution.[4] In both models, the optimal expected costs increase with any of the yield distributions' standard deviations. Also, for the (SCM) and $I^0 \leqslant \mu$, if an increase in the standard deviation of any of the suppliers' yield distributions has an impact on the supplier base, it is to include additional suppliers. (Our numerical studies indicate that the same result applies to the (SCM) with $I^0 > \mu$, as well as the (TCM).)

Proposition 1 (Impact of Supply Risks). (a) *In both the (SCM) and (TCM), if a supplier j fails to be competitive, i.e., $y_j^* = 0$, this supplier cannot become part of the supplier base by improving ς_j, the reliability of his yield distribution, alone.*

(b) *In the (SCM), the optimal cost value increases with any of the parameters $\{\varsigma_j\}$. In the (TCM), the same applies as long as $Y_E^* + I^0 \geqslant \mu$.*

(c) *Consider the (SCM) with $I^0 \leqslant \mu$. k^* is an increasing step function of any of the yield distribution standard deviations $\{\varsigma_j\}$.*

To provide better insight into the invariability result in part (a), we now derive, for the (SCM), the necessary and sufficient conditions for a given supplier j to be part of the supplier base.

Theorem 2(a) and (12) give rise to a closed-form formula for Y_E^{j-1}, the maximum expected effective supply under which supplier j is patronized, i.e., for $Y_E \geqslant Y_E^{j-1}$ only less expensive suppliers in $\{1, \ldots, j-1\}$ are used:

$$
Y_E^{j-1} = \begin{cases} \dfrac{[\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)\gamma_i^{-2}]^2}{[\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)\gamma_i^{-2}]^2 - z_\alpha^2[\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)^2\gamma_i^{-2}]} \\ \quad \cdot \left\{ \mu - I^0 + z_\alpha \left(\dfrac{(\mu - I^0)^2[\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)^2\gamma_i^{-2}]}{[\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)\gamma_i^{-2}]^2} \right.\right. \\ \quad \left.\left. + \dfrac{\sigma^2([\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)\gamma_i^{-2}]^2 - z_\alpha^2[\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)^2\gamma_i^{-2}])}{[\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)\gamma_i^{-2}]^2} \right)^{1/2} \right\} \\ \qquad \text{if } \dfrac{[\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)\gamma_i^{-2}]^2}{[\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)^2\gamma_i^{-2}]} \neq z_\alpha^2, \\ [z_\alpha^2\sigma^2 - (I^0 - \mu)^2]/[2(I^0 - \mu)] \\ \qquad \text{if } \dfrac{[\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)\gamma_i^{-2}]^2}{[\sum_{i=1}^{j-1}(c_j/p_j - c_i/p_i)^2\gamma_i^{-2}]} = z_\alpha^2 \text{ and } I^0 > \mu. \end{cases}
$$

PROPOSITION 2 (COMPETITIVENESS CONDITIONS). *Consider the (SCM).*

(a) *Supplier j achieves a positive market share if and only if*

$$\frac{(Y_E^{j-1}-\mu+I^0)}{z_\alpha^2 Y_E^{j-1}}\left(\sum_{i=1}^{j-1}\left(1-\frac{c_i/p_i}{c_j/p_j}\right)\gamma_i^{-2}\right)\leqslant 1. \tag{37}$$

Thus, whether supplier j achieves a positive market share depends only on the following characteristics of the suppliers: (a) *the effective prices $\{c_1/p_1,\dots,c_{j-1}/p_{j-1}\}$ and the yield distribution coefficients of variation $\{\gamma_1,\dots,\gamma_{j-1}\}$ of all suppliers that are cheaper than supplier j;* (b) *the effective price c_j/p_j of supplier j, but not his own reliability γ_j.*

(b) *There exists a maximum price $c_{k^*+1}/p_{k^*+1} > (c/p)^* \geqslant c_{k^*}/p_{k^*}$ such that any of the currently unutilized suppliers $k^*+1,\dots,N$ receives a positive market share if and only if his effective cost rate is below $(c/p)^*$. In particular, $(c/p)^*$ is independent of the yield distributions of suppliers $k^*+1,\dots,N$.*

Although the number of suppliers increases when any of the yield distributions becomes more volatile, and although for constant input parameters, at least in the (SCM), additional suppliers allow for a reduction in the expected effective supply Y_E (see Theorem 2(a)), Y_E^* fails, in general, to decrease with any of the standard deviations $\{\varsigma_j\}$. This is illustrated by the following example:

EXAMPLE 1. Let $N=3$, $\mu=100$, $I^0=0$, $\sigma=20$; $c_1=1$, $c_2=2$, $c_3=7$; $p_1=p_2=0.65$, $p_3=0.75$, and $\varsigma_1=\varsigma_2=\varsigma_3=0.1$. Consider first the (SCM) with $\alpha=0.001$. When the standard deviations of the yield distributions $\{\varsigma_i\}$ are increased by a common factor C, Y_E^* generally increases, but it decreases from $Y_E^*=250.0143$ to $Y_E^*=243.6793$ and from $Y^*=384.6374$ to $Y^*=374.8912$ when C is increased from $C=1.15$ to $C=1.28$. Similarly, in the (TCM) with $h=1$ and $b=1{,}000$, Y_E^* generally increases, but it decreases from $Y_E^*=195.2574$ to $Y_E^*=193.4458$ and from $Y^*=300.3960$ to $Y^*=297.6089$ when C is increased from $C=1.00$ to $C=1.12$.

Example 1 shows that both the expected safety stock and the total order may fail to increase when all of the standard deviations of the yield distributions are increased, in parallel, by the same percentage. This counterexample also implies that, in both the (SCM) and (TCM), the expected safety stock, as well as the total order, may fail to be monotone in any individual supplier's yield standard deviation.

Proposition 1 shows that if a supplier improves the volatility of his yield distribution, this can only result in a contraction of the supplier base. Does the same monotonicity pattern apply when any of the suppliers improves his average yield? Example 2 disproves this conjecture:

EXAMPLE 2. Assume that the buyer has access to $N=2$ potential suppliers with Normal yield distribution. Let $I^0=0$, $\mu=48$, $\sigma=3$; $c_1=1.2$, $c_2=1.8$; $p_1=0.6$, $p_2=0.7$; $\varsigma_1^2=0.01$, and $\varsigma_2^2=0.02$. In the (SCM) with $\alpha=0.15$, orders are split between both suppliers ($k^*=2$). Supplier 2 is more expensive than supplier 1. In the (SCM), as p_2 is varied on the interval $[0.5, 0.9]$, leaving all other parameters at their base value, supplier 2 continues to be more expensive than supplier 1. He is patronized ($k^*=2$) if and only if $p_2\geqslant 0.76$, so that the number of suppliers increases with this type of improvement of the yield distribution. (Proposition 2 shows that if supplier 2 fails to be competitive, he cannot become part of the supplier base by improving the coefficient of variation of his yield distribution alone; note that an increase of p_2 has the additional effect of reducing his effective cost rate, which does enter into the acceptance condition (37).)[5] At the same time, if p_1 is varied from its base value, it is optimal to use both suppliers as long as $p_1\leqslant 0.56$, but to give supplier 1 exclusivity thereafter. In other words, k^* may fail to be monotonically increasing or monotonically decreasing. We have observed the same phenomenon in the (TCM). Similarly, as with the dependence of the expected effective supply Y_E^* with respect to any of the yield standard deviations $\{\varsigma_j\}$, Y_E^* fails, in general, to be monotone in any of the average yields $\{p_j\}$, both in the (SCM) and (TCM). (In the (SCM) and (TCM) of Example 1, Y_E^* first goes down and then goes up, as p_2 is varied on the interval $[0.51, 1)$.) Only the optimal cost value is monotonically decreasing in any of the $\{p_j\}$-values, a result paralleling Proposition 1(b). (In the (SCM), the result is, again, an immediate consequence of the feasible region expanding with any of the $\{p_j\}$-values.)[6]

Demand Magnitude and Risks

In the classical (SCM) and (TCM), it is well known that the order-up-to level increases when demands have a larger mean and become more variable. (Indeed, it is $\mu+z_\alpha\sigma$ or $\mu+\Phi^{-1}((b-c_1/p_1)/(b+h))\sigma$, simple linear functions of both μ and σ.) The following proposition shows that in the (SCM), when $I^0\leqslant\mu$, the expected order-up-to level, and hence the expected safety stock, continue to be increasing in μ and σ, albeit that the dependence is no longer linear; extensive numerical studies show that the same monotonicity results apply in the (SCM) when $I^0>\mu$ as well as in the (TCM). Similarly, the optimal number of suppliers increases with μ. Perhaps more strikingly, the Proposition also shows that, in the (SCM), when $I^0\leqslant\mu$, an increase in the demand variability can only result in the elimination of suppliers from the supplier base. In contrast, when $I^0\geqslant\mu$, increased demand variability can only result in an expansion of the supplier base. The same monotonicity patterns apply to the (TCM).

PROPOSITION 3 (IMPACT OF DEMAND MAGNITUDE AND RISK). *Consider the (SCM):*

(a) *The optimal cost value increases with μ and σ.*[7]

(b) *Both k^* and Y_E^* increase with μ.*

(c) *k^*, the optimal number of suppliers* (i) *decreases with σ, when $I^0<\mu$;* (ii) *is independent of σ when $I^0=\mu$; and* (iii) *increases with σ, when $I^0>\mu$.*

(d) *If $I^0\leqslant\mu$, Y_E^* increases with σ.*

Figure 1. The optimal number of suppliers k^* as a function of σ in the (SCM).

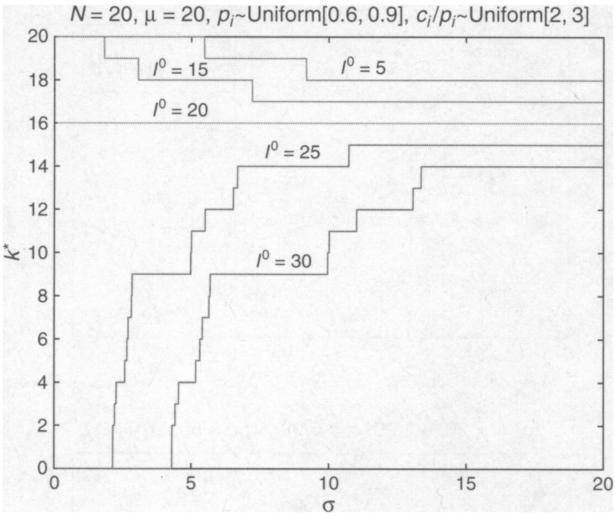

The dependence of the supplier base on the demand variability is illustrated in Figure 1, which displays the optimal number of suppliers k^* as a function of σ for five different values of the starting inventory I^0. (The figure refers to an instance with $N = 20$ suppliers, each with a two-point yield distribution, with probabilities $p_i = \Pr[X_i = 1]$ drawn uniformly from the interval $[0.6, 0.9]$ and effective cost rates c_i/p_i drawn uniformly from the interval $[2, 3]$. Also, $\mu = 20$ and $\alpha = 0.01$.) As shown in Proposition 3, the curves are decreasing step functions when $I^0 < \mu$ and increasing when $I^0 > \mu$; the curve is flat when $I^0 = \mu$.

The intuition behind the monotonicity results of k^* with respect to σ is as follows: (25) shows that it is optimal to use the cheapest k suppliers if and only if targeting an expected supply level $Y_E^k \leqslant Y_E < Y_E^{k-1}$, $k = \underline{k}, \ldots, N$, i.e., the larger an expected supply level is sought, the fewer suppliers one should use so as to control the cost. An increase in the demand volatility σ has two opposite effects: first, as in the classical model, more safety stock is needed to cover the demand, hence a larger value of Y_E is required. (This is indeed proven in part (d) for $I^0 \leqslant \mu$; our numerical study shows that the result holds throughout.) If the critical effective supply values $\{Y_E^k\}$ were invariant with respect to σ, this would imply that k^* *decreases* with σ. However, the formula for Y_E^k in Step 1 of Algorithm SCM shows that these critical values are *increasing* in σ, so that for a given value of Y_E, the *same or a larger* number of suppliers is to be used. When $I^0 < \mu$ ($I^0 > \mu$) and the gap between the initial inventory and the inventory after ordering is relatively large (small), the first (second) effect dominates.

Initial Inventory

Note that in both the (SCM) and (TCM), the optimal order quantities depend on μ and I^0 only via $(\mu - I^0)$. As a consequence, all monotonicity properties identified in Proposition 3 with respect to μ imply the reverse monotonicity pattern with respect to I^0:

PROPOSITION 4 (IMPACT OF THE INITIAL INVENTORY IN THE (SCM)). *Consider the (SCM).*
(a) *The optimal number of suppliers k^* decreases with I^0.*
(b) *Y_E^* decreases with I^0.*
(c) *the optimal cost value decreases with I^0.*[8]

It appears intuitive that the safety stock requirement, in our setting with combined demand and supply risks, should be larger than the optimal safety stock in a setting where only demand risks prevail and all suppliers are completely reliable. Indeed, in the (SCM) we showed that $I^0 + Y_E$, the expected inventory after ordering, satisfies (4), i.e., it is larger than or equal to the optimal inventory level when the suppliers are fully reliable. In contrast, the optimal safety stock in the (TCM) may be smaller than its optimal level when the suppliers are fully reliable: To allow for a meaningful comparison between the models with and without supply risks, assume that in the latter, the suppliers are completely reliable ($\hat{p}_i = 1$) and have the same effective cost rate $\hat{c}_i = \hat{c}_i/\hat{p}_i = c_i/p_i$. In the classical model, it is optimal to place a single order with a supplier i^* for which $c_i^*/p_i^* = \min_j c_j/p_j$. Whether the expected inventory after ordering is smaller or larger than the level in the classical model depends on which of the supply and demand risks dominates. This is exhibited by Figure 2, which considers four instances, again with $N = 20$ suppliers and the same yield distributions as the distributions in Figure 1.[9]

Figures 2(a)–2(d) display the expected inventory after ordering as a function of the initial inventory.[10] As one moves from instance (a) to (c), the consequence of a shortage is increasingly expensive. Moving from instance (c) to instance (d), demand is increasingly variable. In Figure 2(a), the curve is entirely *below* the classical order-up-to level. In Figure 2(b), the curve crosses this level twice, and in Figures 2(c) and 2(d), the curve is entirely *above* the classical order-up-to level.

One might also conjecture that the expected inventory *after* ordering be nonincreasing in the inventory level *before* ordering. After all, with one unit fewer in stock (before ordering), it appears desirable to increase the order sizes so as to target the same (effective) inventory after ordering with a high probability; given the supply risks, this is likely to result in an increase of the expected effective supply by more than one unit and hence in an increase in the expected inventory after ordering. Indeed, in all of our numerical experiments with (SCM) instances, the expected inventory after ordering decreases as a function of the initial stock until it hits the classical level. However, in the (TCM), the above consideration may be counterbalanced when, in the presence of relatively low stockout cost rates, the supply risks justify an expected inventory level after ordering below the classical level. Here, additional units of initial stock allow one to target a higher expected inventory level after ordering, closer to the optimal level in

Figure 2. (TCM): Expected inventory after ordering vs. initial inventory.

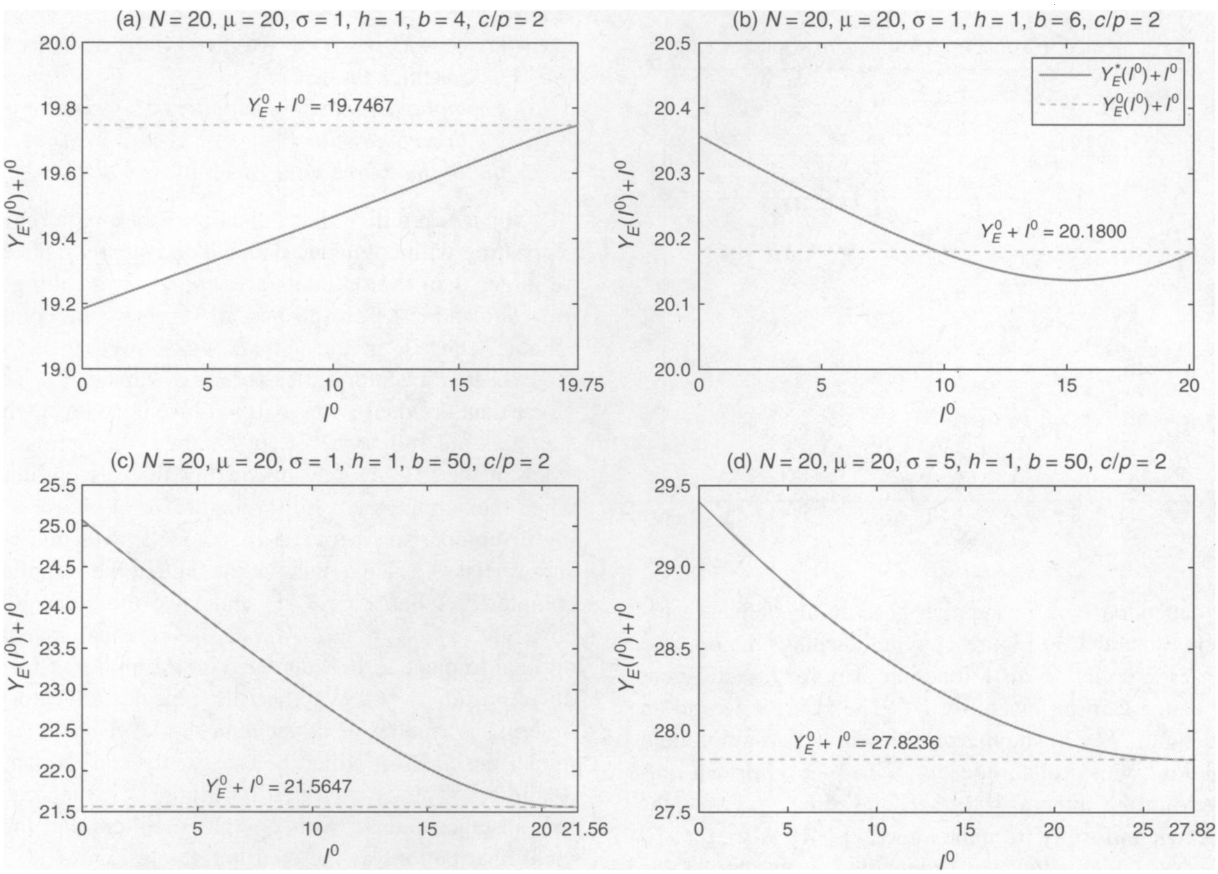

the classical model. Indeed, this situation arises in Figure 2(a), where the expected inventory-after-ordering curve is increasing throughout. In Figure 2(b), with an increased stockout cost rate, the expected inventory-after-ordering curve is first decreasing and then increasing, whereas the curve is decreasing in Figures 2(c) and 2(d), where the stockout cost rate is very high. Note that the increasing parts of the curves only arise when the curve is below the classical level, a situation never encountered in the (SCM); see (4). Note that in Figure 2(d), the deviation of the expected inventory-after-ordering curve from the optimal level in the classical newsvendor model can be as high as 16.25%. In general, large deviations of the expected inventory-after-ordering curve from the optimal level in the classical model can be expected when the suppliers' effective cost rates $\{c_i/p_i\}$ fail to be identical, the number of suppliers is relatively small, or b is relatively large. The percentage deviation also increases when supply risks measured by $\{\gamma_i\}$ are large relative to the demand risk, characterized by σ.

Benefits of Additional Suppliers

When all suppliers are equally expensive, Y_E^* depends on the suppliers' yield distributions via a single measure, i.e., the total number of BSEs, $\sum_{i=1}^N \gamma_i^{-2}$. More specifically, the optimal expected effective supply decreases convexly with this measure. Thus, the cost reduction a new supplier realizes by joining a given industry of suppliers is larger than if he joins an industry with additional suppliers. When the suppliers have different effective cost rates, the following example shows that this result may fail to apply; indeed, a given supplier may enjoy a larger market share when being part of a larger set of potential suppliers.

Example 3. Let $N = 4$; $c_1 = 1$, $c_2 = 1.1$, $c_3 = 1.2$, $c_4 = 2.65$; $p_1 = 0.6$, $p_2 = 0.61$, $p_3 = 0.62$, $p_4 = 0.99$, so that $c_1/p_1 = 1.6667$, $c_2/p_2 = 1.8033$, $c_3/p_3 = 1.9355$, $c_4/p_4 = 2.6768$. Finally, let $\varsigma_i^2 = p_i(1-p_i)/2$, $\mu = 48$, $I^0 = 0$, $\sigma = 3$, and $\alpha = 0.15$. Let $Z(S)$ denote the optimal cost value if the buyer has access to the set of suppliers $S \subseteq \{1,\ldots,4\}$ and $w_i^*(S)$ the market share of supplier $i \in S$. For $S = \{2,4\}$, $Z(S) = 138.9568$ and $w_2^*(S) = 17.22\%$, $w_4^*(S) = 82.78\%$; for $S = \{1,2,4\}$, $Z(S) = 134.6284$ with $w_1^*(S) = 23.17\%$, $w_2^*(S) = 20.74\%$, $w_4^*(S) = 56.09\%$. Thus, if supplier 1 joins the potential supplier base $\{2,4\}$, this results in a cost reduction by 4.3284 units. At the same time, for $S = \{2,3,4\}$, $Z(S) = 136.4341$ and $w_2^*(S) = 18.41\%$, $w_3^*(S) = 16.30\%$, $w_4^*(S) = 65.29\%$; whereas for $S = \{1,2,3,4\}$, $Z(S) = 130.7067$ and $w_1^*(S) = 31.08\%$, $w_2^*(S) = 27.62\%$, $w_3^*(S) = 24.13\%$, $w_4^*(S) = 17.18\%$. We conclude that if supplier 1 joins the larger supplier base $\{2,3,4\}$, his joinder results in a larger cost saving. Note that supplier 1's market share is 31.08% when retained in conjunction with $\{2,3,4\}$ and only 23.17% when retained in conjunction with $\{2, 4\}$. This is explained by the fact that supplier 4 is much more reliable (as well as much more expensive) than the others. When supplier 4 is combined with only two of the others, feasibility considerations dictate that he be given the lion's share of the orders; when supplier 4 is combined with all three of the others, his share can be reduced drastically, allowing for a higher market share for supplier 1.

6. Extensions

Capacity Limits

The (SCM) and (TCM) assume that each supplier is capable of accepting orders of any desired magnitude. However, in many applications, the supply of any given provider is bound by a capacity limit. (Recall, for example, the oil refinery industry discussed in the introduction, which has been operating at close to 100% capacity.) Thus, let

$U_i =$ capacity limit for orders placed with supplier i, $i = 1, \ldots, N$;

$u_i = p_i U_i =$ effective capacity limit of supplier i, i.e., the expected number of effective units, which can be procured from this supplier, $i = 1, \ldots, N$.

To adapt the formulations of the (SCM) and (TCM), only the constraints $y_i \leqslant U_i$, $i = 1, \ldots, N$ need to be added.

When discussing how the results for the (SCM) and (TCM) need to be modified to address the capacity limits, we confine ourselves to the most fundamental question, i.e., whether a given set of suppliers permits a feasible solution and, if so, what the range of feasible effective supply values is. (We thus omit a detailed derivation of the required adaptations of Algorithms SCM and TCM.)

Although feasibility is always guaranteed in the (TCM), we first describe how the necessary and sufficient feasibility condition (F) in the (SCM) is to be generalized. As shown in §3, in the uncapacitated model, a feasible solution exists if and only if an effective supply value Y_E exists, which under a proper allocation of the aggregate order satisfies (4) and the service constraint (7). With

$x_i = p_i y_i =$ expected effective order received from supplier i, $i = 1, \ldots, N$,

the service constraint (7) can be written as

$$(Y_E - \mu + I^0)^2 - z_\alpha^2 \left(\sum_{i=1}^{N} \gamma_i^2 x_i^2 \right) - z_\alpha^2 \sigma^2 \geqslant 0. \tag{38}$$

It follows from a simple sample path argument that if a set of orders $\mathbf{y}$ is feasible, feasibility is maintained when placing full-capacity orders, i.e., under $\hat{\mathbf{y}} = \mathbf{U} \geqslant \mathbf{y}$. (Clearly, if $\mathbf{y}$ satisfies (4), so does $\mathbf{U} \geqslant \mathbf{y}$. On every sample path, the effective supply under the larger orders is at least as large as that resulting from the orders $\mathbf{y}$; this implies that the measure of the set of sample paths for which the service constraint is satisfied under $\hat{\mathbf{y}} = \mathbf{U}$ is at least as large as the measure under $\mathbf{y}$, and hence at least equal to $1 - \alpha$.) By the same argument, if full-capacity orders do not result in a feasible solution, no set of orders does. We conclude:

Theorem 5. *A feasible solution exists in the capacitated (SCM) if and only if the following condition holds:*

(FC) (F) *holds, and*

$$\left(\sum_{i=1}^{N} u_i - \mu + I^0 \right)^2 - z_\alpha^2 \left(\sum_{i=1}^{N} \gamma_i^2 u_i^2 \right) - z_\alpha^2 \sigma^2 \geqslant 0, \tag{39}$$

$$\sum_{i=1}^{N} u_i \geqslant \mu - I^0 + z_\alpha \sigma. \tag{40}$$

Actually, (39) implies (F), and therefore (39) and (40) represent the necessary and sufficient feasibility condition, all by themselves. We nevertheless state (F) as a separate condition because it manifests that, irrespective of the capacity limits, feasibility requires the number of BSEs represented by the set of suppliers to be in excess of a given threshold, as discussed in §3.

We refer to Online Appendix B for a discussion of two important special cases, i.e., the case where all suppliers have an identical effective capacity ($u_i = u$) and that where the effective capacities are proportional to the suppliers' BSEs ($u_i = u\gamma_i^{-2}$).

Fixed Costs

Thus far, we have ignored any fixed costs associated with each individual order to a supplier. As explained, retaining a smaller set of suppliers, when feasible, has the advantage of reducing the average procurement cost per unit (even though it may come at the expense of requiring a larger aggregate order to hedge against the increased supply risks). The presence of fixed costs provides an additional incentive to pursue solutions with a smaller set of suppliers. If the same fixed cost K is incurred for every retained supplier, it is quite easy to incorporate the fixed costs into the analysis. In the (SCM), for example, let Y_E^* denote the optimal effective supply and $\hat{k}$ the number of associated suppliers, in the absence of fixed costs. Because by the above monotonicity property, $k^*(Y_E) \geqslant k^*(Y_E^*)$ for all $Y_E < Y_E^*$, Y_E^* continues to be preferred over all lower values of Y_E in the presence of fixed costs. Because $\Psi^S(Y_E)$ is increasing for $Y_E > Y_E^*$, it follows that only one of the $(N - \hat{k} + 1)$ values in $\{Y_E^*, Y_E^{\hat{k}+1}, \ldots, Y_E^N\}$ may arise as total optimal supply level $Y_E^*(K)$ and $Y_E^*(K) = \arg\min\{\Psi^S(Y_E^*) + \hat{k}K, \Psi^S(Y_E^{\hat{k}+1}) + (\hat{k}+1)K, \ldots, \Psi^S(Y_E^N) + NK\}$. This characterization also implies:

Corollary 1. *In the (SCM), the optimal effective supply $Y_E^*(K)$ is increasing in K. In particular, $Y_E^*(K) \geqslant Y_E^*(0) = Y_E^*$, the optimal level in the case without fixed costs.*

In the (TCM), even though $k^*(Y_E)$ may fail to be monotone, as explained above, it is still easy to identify all breakpoint values such that the number of suppliers k^* remains constant between consecutive breakpoints. (The breakpoint values can be found by solving (34) for $\lambda \in \{c_i/p_i\colon i = |S^0| + 1, \ldots, N\}$. As demonstrated in §4, when I^0 is sufficiently small, this results in at most N breakpoints, and at most $2N$ as long as $I^0 \leqslant \mu$.) In view of the convexity of $\Psi^T(\cdot)$, it again suffices to evaluate only these breakpoints along with Y_E^*.

When the fixed costs are supplier dependent, Federgruen and Yang (2008) already showed that the problem is NP-complete even in the special case where the suppliers have identical variable procurement cost rates $\{c_i/p_i\}$, in which case the optimal set of orders for any given selection of suppliers can be determined in closed form. Nevertheless, that paper showed that a simple greedy-type supplier selection procedure comes very close to being optimal, both empirically and in terms of a worst-case optimality gap. These results follow from the fact that the marginal benefit associated with a new supplier is smaller when the supplier is added to a larger list of potential suppliers. (This property implies that the optimal cost value, viewed as a function of the set of potential suppliers, is submodular.) We continue to advocate the use of the greedy procedure in our general setting with nonidentical cost rates $\{c_i/p_i\}$ as the same submodularity property continues to apply, except in certain extreme cases; see Example 3. Evaluating any candidate set of suppliers can, of course, be done with the (SCM) and (TCM) algorithms.

Uncertain Initial Inventory

We have assumed that the initial inventory I^0 is known precisely. Often, this is not the case. Raman and DeHoratius' (2004) field studies reveal, for example, that in the retail industry, 65% or more of the items have inaccurate inventory records. Another common problem is the inability to find items that the company's computer system claims are in inventory. See also Longitudes (2007), in which this challenge is featured prominently. Finally, inventories may be subject to theft or sabotage. In all of the above settings, it may therefore be appropriate to treat I^0 as a random variable itself. This generalization is easily accomplished, by assuming that the net demand, $D - I^0$, is normally distributed, in which case all of our results can easily be extended. Although in our base model, many of the quantitative and qualitative results depend on whether $I^0 \leqslant \mu$ or $I^0 \geqslant \mu$, in the generalized model, it is important to distinguish between the cases where $E(D-I^0) < 0$ and $E(D-I^0) \geqslant 0$.

Dependent Supply and Demand Risks

The analysis in this paper has assumed that the yield factors of the suppliers are independent. In some settings, supply risks may be correlated, for example, when natural disasters (storms, floods) or sabotage by terrorists is likely to hit multiple facilities in a given geographic region. (Recall the oil refineries example in §1.) To address these interdependencies, assume that the vector of yield factors $\{X_i: i = 1, \ldots, N\}$ has a general joint distribution, with correlation factors $\rho_{ij} = \operatorname{corr}(X_i, X_j)$: $1 \leqslant i \neq j \leqslant N$. The CLT-approximation for the end-of-the-period inventory level is easily adapted to account for any interdependence of the yield factors

$$\Pr[\tilde{I} \leqslant x] = \Phi\left(\frac{x - I^0 - \sum_{i=1}^N p_i y_i + \mu}{\sqrt{\sigma^2 + \mathbf{y}' V \mathbf{y}}}\right),$$

where the $N \times N$ variance-covariance matrix V has $v_{ii} = \varsigma_i^2$ and $v_{ij} = \rho_{ij}\varsigma_i\varsigma_j$ for $i \neq j$.

In §7 of Federgruen and Yang (2008), we have shown that the CLT-based approximation continues to be substantiated by asymptotic accuracy results, as in Lemma 1, provided the dependence of the yield factors is sufficiently "weak," a concept defined precisely there, along with a discussion of easily verified sufficient conditions. This, in turn, permits simple modifications of the model formulations of (SCM) and (TCM). For example, the latter can be formulated as

$$\Psi^T(Y_E) \stackrel{\text{def}}{=} \min_{\mathbf{y}} \left\{ \sum_{i=1}^N c_i y_i + h(Y_E + I^0 - \mu) + (b+h)\int_{Y_E+I^0}^{\infty} \bar{\Phi}\left(\frac{u-\mu}{\sqrt{\sigma^2 + \mathbf{y}'V\mathbf{y}}}\right) du : \sum_{i=1}^N p_i y_i = Y_E,\, y_i \geqslant 0 \right\}.$$

Following the proof of Lemma 5, one verifies that $\Psi^T(Y_E)$ continues to be a strictly convex function, so that its optimum value is determined as the unique effective supply Y_E for which $\Psi^{T\prime}(Y_E) = 0$. (In the proof of Lemma 5, only part (iii) requires a (minor) modification. To show that $\Sigma = \sqrt{\mathbf{y}'V\mathbf{y}}$ is a convex function of $\mathbf{y}$, note that V, as a symmetric positive definite matrix, can be factorized as $V = LL^T$. Thus, with $\mathbf{x}(\mathbf{y}) = L^T\mathbf{y}$, $\Sigma(\mathbf{y}) = \sqrt{\mathbf{x}(\mathbf{y})'\mathbf{x}(\mathbf{y})}$ is a convex function of $\mathbf{y}$, as the composition of a convex and a linear function.) Only the evaluation of the function $\Psi^T(Y_E)$ for a given value of Y_E is now more involved than in the case of independent yield factors. (The (SCM) can be generalized in similar ways.)

Finally, if the demand variable D is correlated with (some of) the yield factors, let $\rho_{Di} = \operatorname{corr}(D, X_i)$. The CLT-based approximation can again be extended along the above lines, now resulting in the following formulation of (TCM):

$$\Psi^T(Y_E) \stackrel{\text{def}}{=} \min_{\mathbf{y}} \left\{ \sum_{i=1}^N c_i y_i + h(Y_E + I^0 - \mu) + (b+h) \int_{Y_E+I^0}^{\infty} \bar{\Phi}\left(\frac{u-\mu}{\sqrt{\sigma^2 + [\mathbf{y}^T, 1]\hat{V}[1, \mathbf{y}]^T}}\right) du : \sum_{i=1}^N p_i y_i = Y_E,\, y_i \geqslant 0 \right\},$$

where the $(N+1) \times (N+1)$ symmetric positive definite matrix $\hat{V}$ is defined by $\hat{v}_{ij} = v_{ij}$ for $1 \leqslant i, j \leqslant N$, $\hat{v}_{N+1,N+1} = \sigma^2$, and $\hat{v}_{i,N+1} = \hat{v}_{N+1,i} = \sigma\varsigma_i\rho_{Di}$. The above argument establishes that the function $\Psi^T(Y_E)$ continues to be strictly convex.

7. Conclusions

We have proposed and analyzed two planning models for settings where uncertain demand for a given item is to be covered by procuring supplies from one or more suppliers. The suppliers face supply risks, in that only a random fraction of orders placed with them results in usable units. The service constraint model minimizes total procurement costs subject to a service constraint that ensures that demand is covered with a given minimum probability $1-\alpha$. In the total cost model, end-of-the-season inventory and backlogging costs are assumed and the aggregate of their expectation and the total procurement costs is minimized.

In both models, the analysis is anchored on a characterization of the functions $\Psi^S(Y_E)$ and $\Psi^T(Y_E)$, denoting the optimal cost value for a given effective supply Y_E in the (SCM) and (TCM), respectively. We have shown that both functions $\Psi^S(\cdot)$ and $\Psi^T(\cdot)$ are strictly convex and differentiable and that they have a unique minimum. This characterization permits us to obtain the optimal effective supply (and associated selection of suppliers and their respective orders) by finding the unique root of the derivative functions $\Psi^{S\prime}(\cdot)$ and $\Psi^{T\prime}(\cdot)$, respectively. We have also shown that both the function and its derivative can be evaluated very efficiently, either in closed form ((SCM)) or after computing the unique root of a nonlinear equation ((TCM)).

Much of our paper is devoted to characterizing the ramifications for (i) the supplier base, (ii) the expected effective supply (and hence the safety stock), and (iii) the optimal cost value resulting from changes in the supply risks the demand magnitude and risks, as well as the amount of initial inventory one possesses. We also show, in the (SCM), that whether a supplier achieves a positive market share or not depends only on his own effective cost rate and those of his less expensive competitors, along with the coefficients of variation of their yield distributions, as the sole characteristic of these distributions. The supplier's own yield distribution is immaterial to ensure membership of the patronized supplier base. Alternatively, if it is optimal for the buyer to patronize the k^* cheapest suppliers, the only way for any of the other suppliers to become part of the supplier base is to reduce their effective cost rate to a given maximum value. Finally, in both the (SCM) and (TCM), the optimal supplier base consists of the k^* cheapest suppliers for some $k^* = 1, \ldots, N$. Each selected supplier is assigned an overall score, given by the product of a reliability and a cost score: The former is the mean-to-variance ratio of the supplier's yield distribution, and the latter is given by the amount by which the supplier's effective cost rate falls below a specific threshold value. The market share of each selected supplier is given by his overall score relative to the sum of the suppliers' scores.

Total cost and service constraint-based models represent the two fundamental approaches in inventory theory to ensure that appropriate safety stocks are selected. In classical inventory models with a single, fully reliable supplier, the two approaches are known to be equivalent; see the discussion in §1. In the presence of supply risks compounding on demand uncertainty, we have shown that this equivalency breaks down in several ways.

First, the very existence of a feasible solution in the (SCM) is of fundamental importance, whereas it is trivially satisfied in the (TCM) as well as in classical service-constraint-based inventory models. We have obtained a very simple characterization of the necessary and sufficient condition for the existence of a feasible solution, both in the case where all suppliers have ample supply and where their supply is capacitated. In the uncapacitated case, the necessary and sufficient condition for the feasibility reduces to a comparison of a single measure characterizing how many suppliers are available and how reliable they are, with either z_α^2, if the initial inventory is below the mean demand, and $z_\alpha^2 - s^2$, if it is s standard deviations above the mean demand. The single measure characterizing the suppliers' pool is the number of BSEs they represent. In the capacitated case, this condition needs to be complemented with a quadratic and a linear inequality in the vector of effective capacities.

A second qualitative difference between the (SCM) and (TCM) is that in the former, if a larger effective supply is targeted, this larger order can be optimally assigned to a subset of the least expensive suppliers within the initial group. This monotonicity implies that a set of critical effective supply values $\{Y_E^k\colon k = \underline{k}, \ldots, N\}$ exist such that the k least expensive suppliers are used if $Y_E^k \leqslant Y_E < Y_E^{k-1}$, $k = \underline{k}, \ldots, N$. Beyond this qualitative result, the identification of these breakpoint values can speed up the search for the optimal value Y_E^*, and is therefore part of our proposed solution method (SCM). At the same time, the above monotonic relationship between the number of suppliers and the targeted effective supply may fail to hold in the (TCM).

Third, one would expect that when supply risks compound on the demand risk, a larger safety stock is required. We have shown that in the (SCM) this always holds, whereas in the (TCM) it may, for some inventory levels, be optimal to order up to a lower level than that in the classical model with a single reliable supplier. We refer the reader to §§1 and 5 for a summary of other qualitative differences in the optimal solutions of the (SCM) and (TCM).

Future work should extend our results to settings with multiple replenishment opportunities.

8. 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. Among the factors driving the need for supplier diversification, the authors mention: "Coupled with this price volatility, there is significant uncertainty about the availability of hi-tech components including memory chips and other semiconductor products. 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."

2. Commercially available supplier scorecard systems tend to determine aggregate scores as the sum or 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.

3. Similarly, $\lambda_2(Y_E)$ may be interpreted as the additional cost incurred when the square of the safety stock in the classical model increases by one unit.

4. Dada et al. (2007) show this for the case of $N = 2$ suppliers. They conclude from this and the fact that the optimal set of suppliers is consecutive in the effective cost rates, that cost rate advantages act as "order qualifiers" whereas reliability advantages serve as "order winners," i.e., to improve a supplier's market share, once qualified.

5. Paying only for the good ones, k^* decreases with p_i.

6. In the special case where all suppliers face two-point yield distributions, this result, combined with Proposition 1(b), implies that the optimal cost value decreases when any of the suppliers improves p, the likelihood of a successfully delivered order. (In this case, $\varsigma_i = \sqrt{p_i(1-p_i)}$, which decreases with p_i for $p_i \geqslant 0.5$.)

7. In the (TCM), the optimal cost value can also be proven to increase with μ, as well as with σ (the latter under the very mild restriction, $Y_E^* \geqslant \mu - I^0$); see the appendix for a proof of this endnote.

8. In the (TCM), the optimal cost value decreases with I^0 as well. This follows from Endnote 8, and the fact that the optimal cost value depends on μ and I^0, only via $(\mu - I^0)$.

9. All four instances have the same collection of two-point yield distributions, with $\Pr[X_i = 1] = p_i$ and $\Pr[X_i = 0] = 1 - p_i$. The p_i values are generated independently from a uniform distribution on [0.6, 0.9].

10. We restrict ourselves to values of $I^0 \leqslant \mu + \sigma\Phi^{-1}((b - c_1/p_1)/(b+h))$ because Theorem 4(d) shows that for a larger value of I^0, $Y_E^* = 0$, i.e., $I^0 + Y_E^*$ follows the 45°-line.

Acknowledgment

The authors thank Chung Piaw Teo for his insightful comments on an earlier version of this paper.

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.

Boyaci, T., G. Gallego. 2001. Minimizing holding and ordering costs subject to a bound on backorders is as easy as solving a single backorder cost model. *Oper. Res. Lett.* **29**(4) 187–192.

Burke, G. J., J. Carrillo, A. J. Vakharia. 2007. Single vs. multiple supplier sourcing strategies. *Eur. J. Oper. Res.* **182** 95–112.

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.

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

Gerchak, Y., M. Parlar. 1990. Yield variability, cost tradeoffs and diversification in the EOQ model. *Naval Res. Logist.* **37**(3) 341–354.

Grosfeld-Nir, A., Y. Gerchak. 2004. Multiple lotsizing in production to order with random yields: Review of recent advances. *Ann. Oper. Res.* **126** 43–69.

Harrison, J. M., J. A. Van Mieghem. 1999. Multi-resource investment strategies: Operational hedging under demand uncertainty. *Eur. J. Oper. Res.* **113** 17–29.

Holloway, C., D. Farlow, G. Schmidt, A. Tsay. 1996. Supplier management at Sun Microsystems (A): Managing the supplier relationship. Harvard Business School Case, Harvard Business School, Boston.

Ilan, Y., M. Yadin. 1985. Second sourcing and transfer of technology—Tradeoffs under uncertainty. Working paper, Department of Operations Research, Statistics, and Economics, The Technion, Haifa, Israel.

Ketchpel, S. P., H. Garcia-Molina. 1998. Competitive sourcing for Internet commerce. *Proc. 18th IEEE Internat. Conf. Distributed Comput. Systems'98*, IEEE, Piscataway, NJ, 602–611.

Kulp, S., V. G. Naravanan. 2002. Metalcraft supplier scorecard. Harvard Business School Case, Harvard Business School, Boston.

Longitudes. 2004. *Mapping the New World of Synchronized Commerce.* Harvard Business School Publishing, Boston.

Longitudes. 2007. *Competitiveness and the Global Supply Chain.* Harvard Business School Publishing, Boston.

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.

Nishiguchi, T., A. Beaudet. 1998. Case study: The Toyota Group and the Aisin fire. *Sloan Management Rev.* **40**(1) 49–59.

Parlar, M., D. Wang. 1993. Diversification under yield randomness in two simple inventory models. *Eur. J. Oper. Res.* **66** 52–64.

Pyke, D. F., M. E. Johnson. 2003. Sourcing strategy and supplier relationships: Alliances vs. eProcurement. C. Billington, H. Lee, J. Neale, T. Harrison, eds. *The Practice of Supply Chain Management.* Springer Science and Business Media, Inc., New York, 77–89.

Raman, A., N. DeHoratius. 2004. Inventory record inaccuracy: An empirical analysis. Working paper, Harvard Business School, Harvard University, Boston.

Rudi, N., Y. S. Zheng. 1997. Multi-item newsboy model with partial variety postponement. Presentation, INFORMS Dallas Meeting. INFORMS, Hanover, MD.

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.

Van Mieghem, J. A. 1998. Investment strategies for flexible resources. *Management Sci.* **44**(8) 1071–1078.

Van Mieghem, J. A. 2007. *Operations Strategy: Principles and Practice.* Dynamic Ideas LLC, Charlestown, MA.

Van Mieghem, J. A., N. Rudi. 2002. Newsvendor networks: Inventory management and capacity investment with discretionary activities. *Manufacturing Service Oper. Management* **4**(4) 313–335.

Yano, C. A. 1991. Impact of quality and pricing on the market shares of two competing suppliers in a simple procurement model. Technical report, Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor.

Yano, C. A., H. L. Lee. 1995. Lot sizing with random yields: A review. *Oper. Res.* **43**(2) 311–334.

Zipkin, P. H. 2000. *Foundations of Inventory Management.* McGraw-Hill, New York.