Distributionally robust solution to the unreliable multisourcing Newsvendor problem with probabilistic constraints

Christian van Delft* and Jean-Philippe Vial†

November 2013

Abstract

In the unreliable multi-sourcing Newsvendor problem the decision maker wants to meet an uncertain demand for a single product by ordering from heterogeneous unreliable suppliers. Each supplier is characterized by a cost and a random yield factor. Because of the uncertainty in the demand and the yield factors, the profit is itself uncertain. The problem is to select how much to order from each supplier, so as to maximize the expected profit, possibly under some additional constraints such as a target service level or a limiting budget on the procurement cost. We consider a version of this problem in which the probability distributions of the demand and the yield factors are imperfectly known and described by their means and covariance matrix only. We formulate the problem as a maximin expected profit model, where the objective function is the worst-case expected profit over the set of probability distributions having the given mean and covariance. The optimal orders are solutions of a tractable conic quadratic programming approach. Via the optimality conditions, we give managerial insight concerning the relative importance of supplier purchase cost and reliability. The model is extended in order to include constraints on the service level and the procurement budget.

1 Introduction

1.1 Overview

The multi-sourcing Newsvendor problem with heterogeneous uncertain suppliers is an extension of the classical Newsboy in which the buyer can split his order among several suppliers. The suppliers are uncertain and may deliver quantities that do not match the orders. Moreover the supply cost varies from one supplier to another. In the present study we perform a quantitative analysis for the maximization of the Newsboy expected profit, when the knowledge on the uncertain parameters is limited to the first and second moments of their distribution. The solution we propose maximizes the expected profit under the worst distribution in the class of distributions having the given first and second moments.

The main issue in the above problem is one of mitigating the impact of uncertainty by proper diversification among several suppliers. Its formulation involves the expectation of the positive and negative parts of the shortfall, which is the actual demand minus the actual total delivery. Even when the probability distribution functions of the uncertain parameters—demand and reliabilities of the suppliers—are known, the computation of these expectations is difficult

*HEC School of Management, Paris (GREGHEC), 78351 Jouy-en-Josas Cedex, France, vandelft@hec.fr.

†Prof. Emeritus, University of Geneva and ORDECSYS, 4 place de l'Etrier, 1224 Chêne-Bougeries, Switzerland, jphvial@gmail.com.

and does not lead to closed form formulas. The maximization of the expected profit is thus a challenging problem. Quite a few papers in the literature deal with it, but, to our knowledge, no general quantitative treatment has been proposed so far, especially under some additional realistic requirements such as probabilistic bounds on the service level or on the procurement budget, or under the assumption of partial information on the probability distributions. In the class of problems under consideration, the profit function can be reformulated as the sum of a linear function of the uncertain parameters (demand and supplier reliabilities) and a term involving the shortfall, that is the positive part of the difference between the demand D and the total supply S. In the simplest version, namely the traditional Newsboy problem, the demand alone is uncertain and the distribution of the difference between demand and supply is inferred from the distribution of the demand. Simple analysis shows then that the optimal order, which is the same as the effective supply, can be obtained by matching the inverse distribution of the demand and a threshold percentile called *critical ratio*. In the version with uncertain suppliers the distribution of the difference $D - S$ is a linear combination of several uncertain factors, the demand and the supplier yields. Except for special distributions, such as the normal, the distribution function of $D-S$ cannot be given in closed form and the computation of the expected shortfall must be performed through complex numerical integration. The maximization of the expected profit is thus challenging and no general solution is proposed.

[Scarf (1958)] proposed an alternative approach to the simplest Newsvendor problem in the context of inventory theory. Assuming that the information on the demand is limited to the knowledge of the mean and the variance, but that the distribution function of the demand itself is unknown, he considered the alternative problem of maximizing the expected profit against the worst demand distribution among all distributions having the given mean and variance. Doing so, he could exhibit a closed form expression for the solution. In view of more recent contributions in robust optimization, Scarf's maximin solution can be said to be *distributionally robust* with respect to the class of distributions having the assigned mean and variance. The argument can be summarized as follows. The profit function of the Newsvendor can reformulated as the sum of linear function of the demand D and the order q minus a constant times the shortage $\max\{D-q, 0\}$. Scarf found an upper bound for the expected shortfall $E[\max\{D-q, 0\}]$ in terms of the mean and the variance. He showed that this bound is attained by a two-point distribution belonging to the class under consideration. A simple presentation of Scarf's ideas appeared years later, together with some extensions, in [Gallego and Moon (1993)].

The extension of Scarf's result to multiple uncertain suppliers is not immediate. In that configuration the supply is not anymore the deterministic total order $\sum q_i$, but a scalar equal to linear combination of the uncertain yields. The mean and variance of $D - S$, where S is the effective (random) supply, can still be computed (for given orders q_i) and Scarf's bound on the expected shortfall is valid. Because the mean and variance of a linear combination of random variables is easily computed from the means and covariance of the uncertain factor, the bound is a simple function of the first and second moments. Moreover, the bound is tight for a two-point distribution of the linear combination. It is not obvious that this distribution can be generated by appropriate distributions on the demand and the reliabilities, but a recent theoretical contribution, [Popescu (2007)], provide a positive answer. This result makes it possible to extend the distributionally robust approach of Scarf to our problem of interest. It follows that the original maximization problem can be replaced by the maximization of a lower bound of the expected profit. The solution is distributionally robust because the bound is attained by a permissible distribution. The optimization problem turns out to be a simple conic quadratic problem. This is the main contribution of our paper. As shown in [Chen et al. (2011)], the same approach can be applied to the computation of the conditional value at risk (CVaR), which is defined as the mean of the tail distribution exceeding VaR, which in turn is the percentile value of the distribution.

By the same token as for the expected profit, one can provide a tight bound for the CVaR in terms of the mean and covariance of the random factors. When applied to the quantity $D - S$, which is used to characterize the service quality, the CVaR value gives the expected service. In the present paper we show how this property can be used to handle constraints on the service level or on the procurement budget.

To conclude this overview, let us summarize the main merits of a distributionally robust approach to the Newsvendor with multiple unreliable suppliers. Firstly, it enables the formulation of the optimization problem as a conic quadratic problem which is easily solvable by widely available tools. the analytic formulation of the objective makes it possible to perform sensitivity analysis at the optimum. In our experiments we use [cvx (2012)] in the MATLAB environment. Secondly, it makes it possible to handle probabilistic constraints of particular interest for the problem. Lastly, but not the least, it deals with the very realistic feature of most practical problems in this area, the one of an incomplete knowledge on the probability distribution functions of the demand and the supplier yields.

1.2 Literature Review

The practical importance of unreliable suppliers in production and inventory systems is attested by an abundant literature. A survey of publications prior to 1995 on the problem of inventory systems with single supplier having random yields can be found in [Yano and Lee (1995)]. The quantitative analyses of the problem turn out to be much more involved than with classical Newsvendor problem. Under general assumptions, no closed form expression can be given to the solution (see for example [Bollapragada and Morton (1999)]). However, several papers have developed numerical solution methods for such problems, often under specific assumptions, and have shown potential benefits of dual, or multiple sourcing in the presence of supply uncertainty. [Rekik et al. (2007)] provide a comprehensive study of the problem with a single unreliable supplier. They develop a methodology to compute the optimal order policy and the optimal expected cost and apply it for two distributions of the yield: uniform and Gaussian. [Agrawal and Nahmias (1997)] gave optimality conditions for the order sizes in the form of sets of nonlinear equations considered under the assumption of Gaussian distribution of the yields. [Anupindi and Akella (1993)] analyze the single-period dual sourcing setting under a Bernoulli probability model: either the whole order quantity is available (the yield is 1) or nothing is delivered (the yield is zero). Optimality conditions for the allocation among the pair of suppliers are derived. The results are extended to the multi-periodic case. [Burke et al. (2009)], updated by [van Delft and Vial (2013)], consider a class of single product sourcing problems with a uniform demand and multiple uncertain suppliers. Assuming that the supplier reliabilities are independent of the demand, they can approximate the expected profit by a quadratic function. They derive a closed form expression for the orders that maximize the approximation of the expected profit. [Chopra et al. (2007)] consider a dual sourcing model with an expensive, but perfectly reliable supplier and a cheap, but unreliable supplier who is subject to both recurrent and disruption uncertainties. In this setting, the reliable supplier is used as a recourse in case of significant disruption occurrences. They show that using the cheaper supplier is optimal in case of recurrent supply uncertainty, while it is more efficient to compensate against major disruptions through recourse to the reliable supplier. [Dada et al. (2007)] proposed an enumerative algorithm solving successive sets of nonlinear equations based on the specific demand and reliability distributions. Furthermore, these authors showed that the cost and the reliability impact the optimal ordering quantities in different manners: costs can be viewed as order qualifiers, while reliabilities can be interpreted as order winners. Suppliers with excessive costs are potentially be left without any order, no matter the reliability level. On the contrary, suppliers with low costs will have some order to deliver, but the size of the order depends on their reliability.

[Federgruen and Yang (2008)] analyze a multi-sourcing model with unreliable suppliers and a service level constraint. The objective is to minimize total procurement costs, including fixed and variable costs, under the constraint that the uncertain demand is met with a given probability. The optimal orders are computed via new approximations for the shortfall probability. The same authors, in [Federgruen and Yang (2009)], extend the results of their previous paper under the assumption that randomness is Gaussian. They are able to characterize the impact of risk magnitude on the ordering policy. [Gurnani et al. (2013)] exhibit a case with a deterministic demand and two suppliers who differ in their costs and reliabilities. For this particular case, they show that single sourcing is optimal, selecting either from the more reliable and more costly supplier or the more risky but cheaper supplier, depending on the cost and reliability parameters. Furthermore, the authors showed, through experimental investigation, that contrary to the theoretical results, the decision-makers favor diversification when facing unreliable supply. [Wang et al. (2010)] explore a dual sourcing model in which the suppliers can exert effort to improve reliability. The authors characterize the optimal procurement quantities and improvement efforts and give managerial insights.

1.3 Organization of the paper

In Section 2 we present the base model for the multisourcing with unreliable suppliers and give a maximin formulation for the maximization of the expected profit. We develop the lower bound for the expected profit and give the distribution under which it is tight. We conclude formulate the distributionally robust equivalent of the maximin optimization problem as a conic quadratic problem. We conclude this section by managerial insights derived through sensitivity analysis at the optimum solution. In Section 3 we introduce probabilistic constraints on the service level and/or the budget allocated for the procurement cost. We substitute to such intractable a constraint, the tractable approximation based on the conditional value at risk (CVaR) concept. We apply to this new constraint the tools developed in Section 2 to provide a tight bound. We formulate a distributionally robust version of the problem with a constraint. We conclude the section with a scheme to enforce a solution (possibly sub-optimal) for the original chance-constraint formulation. In Section 4 we report results on various parameter configurations of the model. We separately study the cases of a deterministic demand and an uncertain one. In each of the two cases, we analyze the impact of a service level constraint. We also show the impact of a budget constraint in the case with a deterministic demand. To illustrate the fact that a two-point distribution of the shortfall can be generated by a permissible distribution of the yields, we perform a simulation and display the empirical distributions for each yield. Finally we illustrate the negative impact of a positive correlation and in contrast, the positive impact of a negative correlation. Section 5 gives a conclusion with suggestions for further research.

2 The multisourcing problem with unreliable suppliers

2.1 The model and the maximin formulation

We consider the single period variant of the Newsboy problem with multiple suppliers. The newsboy aims at serving an uncertain demand D for a single product by placing orders to n uncertain suppliers. The components of the profit function are the revenue, the purchase cost, the shortage cost and the salvage value of leftover units. We use the following notation for the cost and revenue parameters: p the unit selling price, c_i the unit purchase cost from supplier i, u the unit shortage cost (of unmet demand) and s the unit salvage value (of unsold items).

The uncertain parameters are the demand $D \in \mathbb{R}$ and the suppliers reliability vector $R = (R_i)_{i=1,\dots,n} \in \mathbb{R}^n$. If an order q_i is placed to supplier i, the actual delivery will be $q_i r_i$, where r_i is the realized value of uncertain reliability R_i. Following [Burke et al. (2009)] or [Federgruen and Yang (2008)] we assume that the purchase cost applies to the deliveries and not to the orders.

The profit associated with the realization $(d, r_1, \ldots, r_n)$ of the uncertain parameters is

$$\begin{aligned}\pi(q; d, r) &= p \min\{d, \sum_{i=1}^{n} q_i r_i\} - \sum_{i=1}^{n} c_i q_i r_i + s(\sum_{i=1}^{n} q_i r_i - d)_+ - u(d - \sum_{i=1}^{n} q_i r_i)_+ \\ &= (p-s)d - \sum_{i=1}^{n}(c_i - s) q_i r_i - (p+u-s)(d - \sum_{i=1}^{n} q_i r_i)_+. \end{aligned} \tag{1}$$

The profit is a function is separately concave in q and in the uncertain parameters and is of course itself uncertain. The classical objective is to maximize its expected value (see for example [Burke et al. (2009)] or [Federgruen and Yang (2008)]) through the classical stochastic optimization problem

$$\max_q \{E_{D,R}[\pi(q; D, R)]\}. \tag{2}$$

Since the expectation preserves concavity, the objective in concave in q. However, the objective is not well defined if we assume, as it is done in this paper, that the probability distributions of the uncertain parameters are unknown, but the means μ_D and $\mu_R \in \mathbb{R}^n$, the variance σ_D^2 and the covariance are given $\mathrm{Cov}(R) = \Gamma_R \in \mathbb{R}^{n \times n}$. We assume that the reliabilities can be correlated, but the demand is uncorrelated to the reliabilities. The last assumption is not necessary; we make it for the sake of more transparent notation only. The partial information on the uncertain parameters is not sufficient to compute the expectation (2). However, one can define a class of distributions for D and R that we denote $D \sim (\mu_D, \sigma_D^2)$ and $R \sim (\mu_R, \Gamma_R)$. The corresponding distributionally robust formulation is

$$\max_q \min_{\substack{D \sim (\mu_D, \sigma_D^2) \\ R \sim (\mu_R, \Gamma_R)}} \{E_{D,R}[\pi(q; D, R)]\}, \tag{3}$$

which can be rewritten as

$$\max_q \left\{ (p-s)\mu_D - \sum_{i=1}^{n}(c_i - s) q_i \mu_{R_i} - (p+u-s) \max_{\substack{D \sim (\mu_D, \sigma_D^2) \\ R \sim (\mu_R, \Gamma_R)}} \{E_{D,R}[(D - q^\intercal R)_+]\} \right\}. \tag{4}$$

Note that the inner optimization problem (over the set of distributions) is the maximization of a convex function. It is potentially intractable, but the tight bound we shall develop for it is simple enough.

2.2 Equivalence with a unidimensional maximin

In order to solve (4) for fixed q efficiently, we need a closed form expression for the inner maximization problem in the multivariate distributions $D \sim (\mu_D, \sigma_D^2), R \sim (\mu_R, \Gamma_R)$

$$\max_{\substack{D \sim (\mu_D, \sigma_D^2) \\ R \sim (\mu_R, \Gamma_R)}} \{E_{D,R}[(D - q^\intercal R)_+]\}. \tag{5}$$

We exploit here results from [Chen et al. (2011)] and [Popescu (2007)] concerning the expectation of the positive part of a linear combination of random variables. The key point consists of a general projection property for multivariate distributions with given means and covariances, which reduces an optimization problem under multivariate distribution to optimizing a univariate stochastic program, allowing the use of known univariate results. Combining Theorem 1 and Proposition 1 of [Popescu (2007)] we state the following key proposition.

Theorem 1 *Define $X = a^\intercal Z$, with a and Z in $\mathbb{R}^m$, and Z is a random variable with mean $\mu \in \mathbb{R}^m$ and covariance $\Gamma \in \mathbb{R}^{m\times m}$. Denote $\mu_x = E(X) = a^\intercal\mu$ and $\sigma_x^2 = Var(X) = a^\intercal\Gamma a$. Suppose $\Gamma \succ 0$. For any a in $\mathbb{R}^m$, one has*

$$\sup_{Z\sim(\mu,\Gamma)} E((a^\intercal Z)_+) = \sup_{X\sim(a^\intercal\mu, a^\intercal\Gamma a)} E(X_+).$$

Furthermore every distribution $X \sim (a^\intercal\mu, a^\intercal\Gamma a)$ can be obtained by $X = a^\intercal Z$ from some distributions $Z \sim (\mu, \Gamma)$ whose explicit form is

$$Z = \frac{C^{\frac{1}{2}}\beta}{\sqrt{a^\intercal\Gamma a}} + \frac{(X - a^\intercal\mu)\Gamma a}{a^\intercal\Gamma a} + \mu, \tag{6}$$

with $C = (a^\intercal\Gamma a)\Gamma - \Gamma a a^\intercal\Gamma$ and $\beta \sim (0, I_m)$, i.e. , a m-dimensional random variable with $E[\beta_i] = 0$, $E[\beta_i^2] = 1$ and $Cov[\beta_i, \beta_j] = 0$ if $i \neq j$.

The central idea behind this proposition is that for any nonzero vector a, the optimization over random vectors $X = a^\intercal Z$, with $E[Z] = \mu \in \mathbb{R}^m$ and covariance of Z given by $\Gamma \in \mathbb{R}^{m\times m}$, is equivalent to solving the corresponding projected problem over the class of univariate distributions with mean $a^\intercal\mu$ and variance $a^\intercal\Gamma a$. Furthermore, for any univariate distribution with mean $a^\intercal\mu$ and variance $a^\intercal\Gamma a$ it is possible to find the corresponding multivariate distribution in the original setting. A possible choice for the distribution of the β_i is the two-point distribution on $\{-1, 1\}$ with equal probability on the two points.

2.3 Unidimensional maximin tight bound

The consequence of Theorem 1 is that the worst case value of the expected shortage term for multivariate distributions (5) reduces to an optimization problem over the class of univariate distributions X_q with given mean $\mu_D - q^\intercal\mu_R$ and variance $\sigma_D^2 + q^\intercal\Gamma q$, namely

$$\max_{X_q\sim(\mu_D - q^\intercal\mu_R, \sigma_D^2 + q^\intercal\Gamma q)} \{E[X_q]_+\}. \tag{7}$$

The following proposition, proved in [Chen et al. (2011), Popescu (2007)], shows that problem (7) has a simple answer.

Proposition 1 *Let X be a random variable in $\mathbb{R}$, with mean $\mu_1 = E(X)$ and second moment $\mu_2 = E(X^2) < +\infty$. The following inequality holds*

$$E(X_+) \leq \frac{\mu_1 + \sqrt{\mu_2}}{2}$$

and is tight for the two point distribution

$$X = \begin{cases} -\sqrt{\mu_2}, & \textit{with probability } \frac{-\mu_1+\sqrt{\mu_2}}{2\sqrt{\mu_2}}, \\ \\ \sqrt{\mu_2}, & \textit{with probability } 1 - \frac{-\mu_1+\sqrt{\mu_2}}{2\sqrt{\mu_2}}. \end{cases}$$

Applying Proposition 1 to $X = D - q^\intercal R$, we have $E(D - q^\intercal R) = \mu_D - q^\intercal \mu_R$ and $E((D - q^\intercal R)^2) = \sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R)^2$, yielding the bound

$$E[(D - q^\intercal R)_+] \leq \frac{1}{2}\left(\mu_D - q^\intercal \mu_R + \sqrt{\sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R)^2}\right)$$

which is tight for the two-point (worst case) distribution

$$D - q^\intercal R = \begin{cases} -\sqrt{\sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R)^2}, & \text{with probability} \\ & \frac{-\mu_D - q^\intercal \mu_R + \sqrt{\sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R)^2}}{2\sqrt{\sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R)^2}}, \\ \sqrt{\sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R)^2}, & \text{with probability} \\ & 1 - \frac{-\mu_D - q^\intercal \mu_R + \sqrt{\sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R)^2}}{2\sqrt{\sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R)^2}}. \end{cases} \tag{8}$$

The equivalent of the distributionally robust problem (4) is thus

$$\begin{aligned} \max_q \quad & \varpi(q) := (p - s)\mu_D - \sum_{i=1}^{n}(c_i - s)q_i\mu_{R_i} \\ & -\frac{(p + u - s)}{2}\left(\mu_D - q^\intercal \mu_R + \sqrt{\sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R)^2}\right). \end{aligned} \tag{9}$$

2.4 Range, mean and coefficient of variation

We exhibit here a property of the model under examination. The ratio between the mean and the standard deviation of the reliability (the so-called coefficient of variation) is the key parameter for the reliabilities. Suppose the mean and standard deviation of the reliability of supplier i is multiplied by some arbitrary constant $\theta > 0$, letting the coefficient of variation unchanged. If q^* solves the original problem, it is easy to see that the order $\tilde{q}_j = q_j^*$ for $j \neq i$ and $\tilde{q}_i = q_i^*/\theta$ solves the problem with the new mean and standard deviation. Hence, the critical parameter is the coefficient of variation of R and not its mean and standard deviation.

2.5 Managerial insights via sensitivity analysis

In view of the analytic representation of the equivalent distributionally robust expected profit, one can perform sensitivity analysis at the optimal solution. Following the discussion in the preceding subsection, we assume, without loss of generality, $\mu_{R_i} = 1$ for all i. For the ease of notation, but again without loss of generality, we take $u = 0 = s$. We also rank the suppliers by increasing procurement cost: $c_1 \leq c_2 \leq \cdots \leq c_n$. Finally we assume that the covariance matrix Γ is diagonal: all covariances are zero. We denote q^* the optimal solution.

The sensitivity analysis is based on the derivative with respect to q_i of the expected profit for the worst case distribution

$$\frac{d}{dq_i}\varpi(q) = -c_i - \frac{\sigma_i^2 q_i}{K} + L, \tag{10}$$

where $K = \frac{2}{p}\sqrt{\sigma_D^2 + \sum_j q_j\sigma_j^2 + (\mu_D - \sum_j q_j)^2}$ and $L = \frac{p}{2} + \frac{\mu_D - \sum_j q_j}{K}$. We shall use this expression to compare the situation of two suppliers. We classify suppliers into two categories according to their participation to the total order size at the optimum. We say that supplier i belongs to the bundle of active suppliers if $q_i^* > 0$; then $\frac{d}{dq_i}\varpi(q^*) = 0$. A supplier j with $q_j^* = 0$ is said to be inactive; then $\frac{d}{dq_j}\varpi(q^*) \leq 0$.

We consider first the case of two active suppliers i and j. In view of (10) and $\frac{d}{dq_i}\varpi(q^*) = 0 = \frac{d}{dq_j}\varpi(q^*)$, we have

$$\sigma_j^2 q_j^* - \sigma_i^2 q_i^* = -K(c_j - c_i). \tag{11}$$

If the two suppliers have the same procurement cost, then $\sigma_j^2 q_j^* = \sigma_i^2 q_i^*$. The order to the more risky supplier, say j (i.e., $\sigma_j > \sigma_i$), satisfy

$$q_i^* = \frac{\sigma_j^2}{\sigma_i^2} q_j^* > q_j^* = \frac{\sigma_i^2}{\sigma_j^2} q_i^* > 0.$$

The more risky supplier strictly contributes to the total supply, but for a fraction of the less risky supplier. By splitting the orders between the two (and allocating less to the more risky supplier), the optimal solution decreases the global variability of the total supply. If i is more expensive ($c_j < c_i$), then the right-hand side in (11) is positive and the order to the more expensive supplier cannot be too large: $q_i^* < (\sigma_j^2/\sigma_i^2)q_j^*$. The ratio between the standard deviations may smaller or greater than one. In the latter case, where the more expensive supplier is less risky, it is conceivable that $q_j^* < q_i^*$.

We consider now the situation where some suppliers are inactive at the optimal solution. Suppose $q_i^* = 0$ and $q_j^* > 0$ at the optimum. Then $\frac{d}{dq_i}\varpi(q^*) \leq 0$ and $\frac{d}{dq_j}\varpi(q^*) = 0$. Putting the inequalities together, we get

$$c_j - c_i + \frac{\sigma_j^2 q_j^*}{K} \leq 0. \tag{12}$$

The above inequality implies $c_j < c_i$. The suppliers who are part of the active bundle have procurement costs that are strictly less than those of the inactive suppliers. We thus proved the following theorem:

Theorem 2 *The optimal distributionally robust solution has the property that a supplier with a higher procurement cost cannot enter the solution unless all suppliers with lesser costs are part of the solution.*

Note that the exclusion condition depends on the procurement cost alone, and not on the reliability of the excluded supplier. The property holds for all inactive suppliers, including the fully reliable one. The striking result is that the exclusion of a supplier from the set of active suppliers depends on his procurement cost alone, and not on his reliability. The above property of distributionally robust solutions was proved for other distributions by [Dada et al. (2007)].

Using (13) and the above theorem, we can elaborate on the barrier to entry for a supplier. We already know that a new supplier will enter the bundle of active suppliers if his cost is strictly less than the cost of one of the active suppliers. But, even if his cost is higher, he can still be part of the group, provided his cost is not too large. The amount in excess of the costs of the active suppliers can be computed. Let k be the index of the inactive supplier with least procurement cost among all inactive suppliers. Because the derivatives with respect to the order are zero for all suppliers in the set J of active suppliers we have

$$c_j + \frac{\sigma_j^2 q_j^*}{K} = L, \text{ for any } \in J.$$

The condition for supplier k to join the bundle of active suppliers is to have a procurement cost

$$c_k \leq c_j + \frac{\sigma_j^2 q_j^*}{K}, \text{ for any } j \in J. \tag{13}$$

Theorem 3 *Let J is the set of indices of the active suppliers. The inactive supplier k cannot become active unless*

$$c_k < L = \frac{p}{2} + \frac{\mu_D - \sum_{j\in J} q_j^*}{K} = c_j + \frac{\sigma_j^2 q_j^*}{K}, \text{ for any } j \in J.$$

Again, the condition is independent of the reliability σ_k of the candidate supplier. We conclude that a supplier with a procurement cost not too large can become active, even though his reliability is poor. The procurement cost is the only criterion. Of course, the amount of the total order that will be assigned to the entrant will be strongly affected by the reliability.

Finally, we derive a condition on the total order size from the status of supplier k. Because k is not selected $\frac{d}{dq_k}\varpi(q^*) \leq 0$. From (10), $-c_k + L \leq 0$, i.e.,

$$\sum_{i<k} q_i^* \geq \mu_D + (\frac{p}{2} - c_k)K.$$

If $p \geq 2c_k$, the total order is greater than the mean demand. Many papers focus on the case of 2 suppliers, an unreliable one and a reliable one. Assume $c_1 < c_2$, $\sigma_1 > 0$ and $\sigma_2 = 0$. Suppose $q_2^* > 0$. From Theorem 2, $q_2^* > 0$ implies $q_1^* > 0$. Since $q_2^* > 0$, equation (10) hods and $\sigma_2 = 0$ imply $L = c_2$. We conclude

$$q_1^* + q_2^* = \mu_D + (\frac{p}{2} - c_2)K.$$

This property can be extended to more than two suppliers, provided all suppliers are active and there exists a fully reliable supplier in the bundle.

3 Distributionally robust formulation for probabilistic constraints

3.1 Constraints on the service level and the procurement budget

The model for the Newsvendor problem of the previous section is adequate for a risk-neutral decision-maker. It may be desirable to incorporate to the base formulation managerial concerns for some critical criteria. For instance, a manager may want to restrict shortfalls so as to ensure a satisfactory service level. The total procurement cost may be another concern if a budget limit exists. Because the two quantities, shortfall and procurement cost, are essentially random, constraints on them must be formulated in probabilistic terms. The problem is then a chance-constraint programming. The two quantities of interest in the present analysis are the end inventory function $I(q; D, R) = D - q^\intercal R$ and the budget function $B(q; D, R) = b - q^\intercal R$, where b is the target budget. The corresponding probabilistic constraints are thus

$$\text{Prob}(I(q; D, R) \geq 0) = \text{Prob}(D - q^\intercal R \geq 0) \leq \beta, \tag{14}$$

and

$$\text{Prob}(B(q; D, R) \geq 0) = \text{Prob}(b - \sum_i c_i q_i R_i \geq 0) \leq \beta. \tag{15}$$

In these formulations β is a probabilistic threshold assigned by the manager. The two constraints are very similar. We shall discuss the first one only.

An alternative interpretation of the chance-constraint (14) is the constraint on the Value at Risk (VaR) on the end inventory for the percentile $1 - \beta$

$$\text{VaR}_\beta(I(q; D, R)) \leq 0. \tag{16}$$

It is well-known that a chance constraint (14) or (15) is non-convex and in general numerically intractable. However the conditional value at risk (CVaR) is also known to be the best convex upper approximation of VaR, (see [Ben-Tal et al. (2009)]). In view of this, the chance-constraint, or constraint on $\mathrm{VaR}_\beta(I(q;D,R))$, is often replaced by the constraint on the conditional value at risk $\mathrm{CVaR}_\beta(I(q;D,R)) \leq 0$.

3.2 Distributionally robust formulation of CVaR constraints

Let $f(x,\xi)$ be the loss function associated with a decision problem in the variable $x \in \mathbb{R}^m$ and the random variable $\Xi \in \mathbb{R}^p$. Let $\pi(.)$ be the cumulative distribution function of Ξ. We assume $E_\Xi(|f(x,\xi)|) < \infty$ for all x. For a given confidence level β and a fixed x, the value at risk is

$$\mathrm{VaR}_\beta(x) := \min\{\alpha \in \mathbb{R} \mid \int_{f(x,\xi)\leq\alpha} d\pi(\xi) \geq \beta\}. \tag{17}$$

The corresponding conditional value at risk is the expected value of loss exceeding $\mathrm{VaR}_\beta(x)$ and is given by

$$\mathrm{CVaR}_\beta(x) := \frac{1}{1-\beta}\int_{f(x,\xi)\geq \mathrm{VaR}_{\beta(x)}} f(x,\xi)d\pi(\xi). \tag{18}$$

By [Rockafellar and Uryasev (2000)], it also satisfies the following equality

$$\mathrm{CVaR}_\beta(x) = \min_\alpha F_\beta(x,\alpha) \tag{19}$$

where

$$F_\beta(x,\alpha) := \alpha + \frac{1}{1-\beta}\int [f(x,\xi)-\alpha]_+ d\pi(\xi) \tag{20}$$

and $\mathrm{VaR}_\beta(x)$ satisfies

$$\mathrm{VaR}_\beta(x) = \arg\min_\alpha F_\beta(x,\alpha). \tag{21}$$

Equations (20) and (21) imply that $\mathrm{VaR}_\beta(x) \leq \mathrm{CVaR}_\beta(x)$. Therefore, the constraint $\mathrm{CVaR}_\beta(x) \leq 0$ is a surrogate for the constraint on $\mathrm{VaR}_\beta(x) \leq 0$.

We are interested in a distributionally robust version of the CVaR constraint. Consider the family of random variables $\Xi \sim (\mu,\Gamma)$, with $\Gamma \succ 0$. The robust CVaR with respect to this class of distributions is

$$\mathrm{RCVaR}_\beta(x) := \sup_{\Xi\sim(\mu,\Gamma)} \mathrm{CVaR}_\beta(x) = \sup_{\Xi\sim(\mu,\Gamma)} \min_\alpha F_\beta(x,\alpha). \tag{22}$$

The following proposition, borrowed from [Chen et al. (2011)] (Theorem 2.9), establishes that interchange of the sup and min operator is allowed.

Proposition 2 *When the loss function $f(x,\xi)$ is linear in ξ, one has*

$$\mathrm{RCVaR}_\beta(x) = \sup_{\Xi\sim(\mu,\Gamma)} \min_\alpha F_\beta(x,\alpha) = \min_\alpha \sup_{\Xi\sim(\mu,\Gamma)} F_\beta(x,\alpha). \tag{23}$$

It is thus possible, when the loss function f is linear in ξ, to explicit the term $\sup_{\Xi\sim(\mu,\Gamma)} F_\beta(x,\alpha)$ using the bound exhibited in Proposition 1.

In general $\mathrm{VaR}_\beta(x) < \mathrm{CVaR}_\beta(x)$, with equality holding but on exceptional degenerate cases. The constraint $\mathrm{CVaR}_\beta(x) \leq 0$ is too restrictive in general. A common practice is to relax it to $\mathrm{CVaR}_\beta(x) \leq c$, with $c > 0$ and adjust c by some numerical scheme to achieve $\mathrm{VaR}_\beta(x) = 0$. It is important to note that solving an optimization with the constraint $\mathrm{CVaR}_\beta(x) \leq 0$ and a c ensuring $\mathrm{VaR}_\beta(x) = 0$ is not the same as solving the problem with the VaR constraint $\mathrm{VaR}_\beta(x) \leq 0$. The latter constraint is less restrictive.

3.3 The distributionnally robust equivalent optimization problem

For known probability distributions for (D, R) and for a given $c > 0$, the stochastic optimization problem with a service level constraint is

$$\max_{\alpha, q \geq 0} \quad E_{D,R}[\pi(q; D, R)] \tag{24a}$$

$$\text{CVaR}_\beta(q) = \min_\alpha \left\{ \alpha + \frac{1}{(1-\beta)} E[(D - q^\intercal R - \alpha)_+] \right\} \leq c, \tag{24b}$$

The corresponding distributionally robust equivalent is

$$\max_q \quad \min_{\substack{D \sim (\mu_D, \sigma_D^2) \\ R \sim (\mu_R, \Gamma_R)}} E_{D,R}[\pi(q; D, R)] \tag{25}$$

$$\text{s.t} \quad \min_\alpha \left\{ \alpha + \min_{\substack{D \sim (\mu_D, \sigma_D^2) \\ R \sim (\mu_R, \Gamma_R)}} \frac{1}{1-\beta} E[(D - q^\intercal R - \alpha)_+] \right\} \leq c. \tag{26}$$

Note that the worst case distributions for (25) and (26) need not be the same. A stronger, but intractable formulation of the distributionally robust equivalent would impose that the worst case distributions be the same in the objective and the constraint. We shall see later a case where the stronger property holds.

Proposition 3 *The distributionnally robust formulation of the maximization of the expected profit under a CVaR constraint on the shortfall given by the following conic quadratic programming problem*

$$\max_{\alpha, t, q \geq 0} \quad (p - s)\mu_D - \sum_{i=1}^n (c_i - s) q_i \mu_{R,i} - \frac{(p + u - s)}{2} (\mu_D - q^\intercal \mu_R + t_1) \tag{27a}$$

$$\alpha + \frac{1}{2(1-\beta)} (\mu_D - q^\intercal \mu_R - \alpha + t_2) \leq c, \tag{27b}$$

$$\sqrt{\sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R)^2} \leq t_1, \tag{27c}$$

$$\sqrt{\sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R - \alpha)^2} \leq t_2. \tag{27d}$$

Proof: According to Theorem 1 and Proposition 1, the objective (25) and the constraint (26) can be replaced by their conic quadratic bound. The minimization in α need not be made explicit because the constraint is embedded into an optimization problem. ■

We now discuss a scheme to achieve $\alpha = 0$ at the optimum of (27). This will ensure that the value at risk constraint (14) holds at the optimum. Let us sketch here the algorithmic procedure to enforce $\text{VaR}_\beta(q) = 0$. Recall first that if (27b) is tight, the minimum in α is achieved in (26) for the worst-case distribution. Consequently, $\alpha = \text{VaR}_\beta(q)$. The basic step of the procedure is as follows. Let $0 < c_1 < c_2$ and let (α_i^*, q_i^*), $i = 1, 2$, be the corresponding optimal solutions associated with c_i in (27). Assume (27b) is tight in both cases and that $\alpha_1^* = \text{VaR}_\beta(q_1^*) < 0$ and $\alpha_2^* = \text{VaR}_\beta(q_2^*) > 0$. Let $c' = (c_1 + c_2)/2$ and solve (27) for this new right-hand side in (27b) with $c = c'$. Let (α^*, q^*) be the optimal solution. If $\alpha^* = \text{VaR}_\beta(q^*) < 0$ repeat the procedure with c' and c_2. Otherwise, repeat the procedure with c_1 and c'.

As a final remark, let us examine the two expectations in problem (24), one in the objective and one in the constraint, that are replaced by upper bounds. We know that each upper bound is tight for a certain worst case distribution. If the two worst case distributions are not the same, the solution of the problem might be conservative. We give here a sufficient condition to ensure that the worst case distributions are the same.

Proposition 4 *Assume that at the optimal solution of problem (27) the service level constraint (27b) with right-hand side c is tight and $\alpha = 0$ at this solution. Then the bound*

$$E[(D - q^\intercal R)_+] \leq (1 - \beta)c$$

is tight for the same distribution as the one which makes the CVaR constraint tight.

Proof: If the constraint (27b) is active at the optimum, the variable α minimizes the left-hand side of the constraint. If $\alpha = 0$ at the optimum, the constraints (27c) and (27d) are identical, and the same two-point distribution makes the bound on the profit and the bound on the CVaR simultaneously tight. ∎

4 Numerical experiments

The numerical data are similar as those in [Burke et al. (2009)]. We use the same selling price, salvage value and shortage cost. In the present paper, the number of suppliers is larger and the costs are different. The suppliers are characterized by their cost and their reliability. The latter is measured by the mean and the coefficient of variation (standard deviation over the mean). The reliabilities are between zero and one, though it can be argued that larger values than one

supplier	1	2	3	4	5	6
cost	621	624.5	628	631.5	635	638.5
mean	0.75	0.8	0.8	0.85	0.9	0.9
standard deviation	0.0825	0.072	0.056	0.0425	0.027	0.009
coefficient of variation	0.11	0.09	0.07	0.05	0.03	0.01

	deterministic	uncertain
mean demand	7500	7500
coefficient of variation	0	0.04
standard deviation	0	300

Table 1: Demand, supplier costs and supplier reliabilities

are meaningful (see for instance [Chopra et al. (2007)]). Note that the purchasing price does not vary much with the supplier (less than 0.5% between two neighboring suppliers and less than 3% between suppliers 1 and 6). For the demand, we consider two cases, a deterministic demand and an uncertain one. The deterministic demand case can be viewed as the one faced by a producer who needs a supply of exactly 7500 units to meet his commitments. The uncertain demand case will be referred to as the one of a reseller. To solve problem (27) we use the tool `cvx` for convex optimization in Matlab [cvx (2012)].

As discussed in section 2.4 the solution is insensitive to a joint change of standard deviation and mean, provided their ratio, the coefficient of variation, remains unchanged. In other words, for a different pair of mean and standard deviation, but constant coefficient of variation, the order may change but not the expected delivery. To simplify the presentation of the numerical results we shall omit the order and only report the expected (or mean) deliveries.

4.1 Producer with a deterministic target production plan

We study the case of a single producer who must release the deterministic quantity D. We study first the case with no constraint on the service level. Table 2 shows that the producer decreases the uncertainty of the total supply by mixing low cost suppliers with important uncertainty to more certain suppliers with higher costs.

supplier	1	2	3	4	5	6	total
expected deliveries	1050	1226	1462	1759	1811	0	7309
deliveries in %	14%	17%	20%	24%	27%	0	100%

Expected profit
472,047

$\text{Prob}(D - q^\intercal R \geq \text{VaR})$	0.1	0.05	0.01
VaR value	925	1233	2435

Table 2: Distributionnaly robust solution for the unconstrained base case with a deterministic $D = 7500$.

The VaR value displays the amount that the shortage will not exceed with the probability given above. The VaR and CVaR values are obtained by solving the auxiliary problem in α, where q is the optimal solution displayed in the table.

$$\min_{\alpha} \left\{ \alpha + \frac{1}{2(1-\beta)} \left(\mu_D - q^\intercal \mu_R - \alpha + \sqrt{\sigma_D^2 + q^\intercal \Gamma q + (\mu_D - q^\intercal \mu_R - \alpha)^2} \right) \right\}.$$

The same base case is treated now with a constraint on the service level. We recall that we use the probability of shortage $\text{Prob}(D - q^\intercal R \geq 0)$ as a measure of the service level. Given a target probability β we implement the binary search procedure described in Section 3.1 to achieve $\text{VaR}_\beta = 0$. The results are given in Table 3. Producer is facing the deterministic target production 7500 and in the 3 instances the constraint on the VaR at the given levels of quantile $(0.10, 0.05, 0.05)$ to be less than zero is enforced.

supplier	1	2	3	4	5	6	total
Shortage probability ≤ 0.10; Expected profit: 406,633; CVaR $= 117.5$							
expected deliveries	104	139	203	345	812	5991	7594
Shortage probability ≤ 0.05; Expected profit: 376,911; CVaR $= 161.6$							
expected deliveries	88	120	180	318	783	6156	7645
Shortage probability ≤ 0.01; Expected profit: 251,723; CVaR $= 361.9$							
expected deliveries	69	99	155	290	762	6480	7855

Table 3: Distributionnaly robust solution with the service level constraint.

Table 4 shows the impact of a budget constraint. The producer has a deterministic target production 7500 and the allocated budget is the same as the mean cost of supply in the unconstrained case $b = 877,026$. In the first instance, the budget constraint is just inactive. The probability of not meeting the budget is then 0.5. (i.e., $\text{VaR}_{0.5} = 0$). In the next 3 instances the constraint on meeting the budget at the given levels of quantile $(0.10, 0.05, 0.05)$ to be less than zero is enforced.

The optimal solution when the budget is reduced favors diversification towards the more reliable, but more expensive, suppliers in a manner that resembles to the impact of a more restrictive service level as displayed in Table 3. However, in the case of the service level the total expected deliveries increases in order to reduce the expected shortfall, whereas in the case of the budget constraint the total expected deliveries decreases so as to meet more severe restriction on the budget constraint.

supplier	1	2	3	4	5	6	total
Probability of not meeting the budget ≤ 0.5							
Expected profit: 472047				CVaR on budget = 30,000			
expected deliveries	1 050	1226	1462	1759	1811	0	7309
Probability of not meeting the budget ≤ 0.1							
Expected profit: 445,882				CVaR on budget = 19,226			
expected deliveries	363	443	568	792	1306	3709	7180
Probability of not meeting the budget ≤ 0.05							
Expected profit: 435,517				CVaR on budget = 21,546			
expected deliveries	241	300	395	577	1056	4579	7147
Probability of not meeting the budget ≤ 0.01							
Expected profit: 409,196				CVaR on budget = 40,030			
expected deliveries	128	166	233	374	813	5268	6982

Table 4: Distributionnaly robust solution with a probabilistic budget constraint.

4.2 Reseller facing an uncertain demand

When the demand is uncertain, the shortage inherits a larger variability (its variance is larger). The impact of diversification among suppliers is relatively weaker. As a result the robust solution puts more emphasis on the cost and less on the variability. This can be seen on Table 5 in the absence of a constraint on the service level and on Table 6.

supplier	1	2	3	4	5	6	total
expected deliveries	1671	1837	1947	1681	0	0	7135
deliveries in %	23%	26%	28%	24%	0	0	100%

Expected profit
434,076

Prob($D - q^\intercal R \geq$ VaR)	0.1	0.05	0.01
VaR value	925	1233	2435

Table 5: Distributionnaly robust solution for the unconstrained base case with uncertain demand (mean 7500, coefficient of variation 0.04 and standard deviation 300).

In the following case, reseller faces uncertain demand (mean 7500, coefficient of variation 0.04 and standard deviation 300). In the 3 instances the VaR at the given level of quantile is enforced to be 0.

supplier	1	2	3	4	5	6	total
Shortage probability ≤ 0.10							
Expected profit: 193,080 CVaR $= 538.4$							
expected deliveries	287	356	466	675	1211	4923	7917
Shortage probability ≤ 0.05							
Expected profit: 63,284 CVaR $= 713.9$							
expected deliveries	215	271	365	552	1081	5658	8142
Shortage probability ≤ 0.01							
Expected profit: -470,106 CVaR $= 1565.7$							
expected deliveries	131	174	252	424	981	7072	9034

Table 6: Distributionnaly robust solution with the service level constraint.

4.3 Complementary analyses on the producer case

For the sake of shortness, we perform a few complementary analyses on the producer case only, but similar observations can be made on the reseller case.

4.3.1 Worst case distributions of the demand and the reliabilities

Let q^* be the optimal solution of problem (27), without the CVar constraints (27c) and (27d). We know that the objective is attained for a 2-point distribution of the scalar random variable $X(q) = D - q^\intercal R$. We raise the issue of finding a distribution for D and R with the appropriate properties (mean and variance, and hopefully range) yielding the two point distribution for X at q^*. Using (6) we can produce the appropriate multidimensional random variable (D, R) with the appropriate mean and variance. The only check to be performed is to see whether the range of components is admissible. We shall do it under the assumption that the variable β used in the construction has the smallest possible range. This is the case if each β_i takes values -1 and 1 with equal probabilities.

The worst case distribution corresponding to the orders solution is the two-point distribution given by Proposition 1. Formula (6) in Theorem 1 provides an easy way to construct a set of supplier distributions that generates the two-point distribution of $q^\intercal R$ and have the appropriate means and standard deviation. The distribution is not unique because the only requirement on the auxiliary random variable β is to be centered with the covariance matrix equal to the identity. For the display of Figure 1 we chose the β_i to be independent binary variables with equal probabilities on -1 and 1. This ensures the least possible range. This last feature also ensures, via formula (6), the least possible range for the reliability distributions. Those ranges are displayed in Table 7. They are all between 0 and 1. To obtain this result we chose lower reliability means for suppliers 1–5. Higher means would not have changed the solution as far as profit and expected deliveries are concerned, but they imply that some range upper bounds larger than 1. In [Chopra et al. (2007)] the authors argue that in the pharmaceutical industry (vaccines) or in the semiconductor industry the deliveries may exceed the orders. Therefore values of R greater than 1 may be acceptable.

To construct an empirical distribution of the reliabilities, we construct a sample of size $N = 10,000$ of the binary variables β and ξ. In view of (6), we construct the sample of reliabilities. The empirical distributions are represented on Figure 1.

From simulation of the worst case distribution, we derive a sample of the profit function. The

	Case 1		Case 2	
supplier	mean	range	mean	range
1	0.75	[0.6155 0.9626]	0.9	[0.7426 1.1512]
2	0.8	[0.6781 0.9871]	0.9	[0.7635 1.1098]
3	0.8	[[0.6985 0.9486]	0.9	[0.7864 1.0665]
4	0.85	[0.7686 0.9621]	0.9	[0.8127 1.0198]
5	0.9	[0.8910 0.9090]	0.9	[0.8464 0.9656]
6	0.9	no order	0.9	no order

Table 7: Ranges of variation for the reliabilities when the mean reliabilities change. Base case of the producer (deterministic production plan). No constraint on the service level

profit distribution turns out to be also a 2-point one. The mean profit is 472,720, quite close to the theoretical one 472,047. The lowest profit value is at 361,000; the highest at 492,400.

4.3.2 Correlated suppliers

The impact of negatively correlated suppliers is positive: putting more emphasis on a pair of suppliers who are negatively correlated decreases the variability. The opposite holds if the suppliers are positively correlated. The results are reported in Table 8.

supplier	1	2	3	4	5	6	total
Uncorrelated suppliers							
expected deliveries	1050	1226	1462	1759	1811	0	7309
Expected profit: 472,047							
Suppliers 1 and 2 are correlated with coefficient $= -0.8$							
expected deliveries	1729	2049	1369	1464	680	0	7290
Expected profit: 482,227							
Suppliers 3 and 4 are correlated with coefficient $= 0.8$							
expected deliveries	1126	1326	1232	1298	2317	0	7300
Expected profit: 468,522							

Table 8: Producer with a deterministic plan of 7500 units facing uncertain and correlated suppliers.

4.3.3 Influence of shortage cost and salvage value

The shortage cost is a reality, but it is very difficult to quantify it. Clearly higher shortage costs should be an incentive for higher orders to achieve less shortage. In order to make comparison, we increase the shortage cost with the goal of achieving the same probability of shortage. Table 9 reveals that one can achieve the desired service level by an appropriate choice of the shortage cost. Surprisingly enough, the optimal orders are the *exactly the same* as the optimal ones for problem (27). Consequently, the expected profit is smaller, because shortage is penalized by a value that is much higher than the original value in the base case. The most striking observation

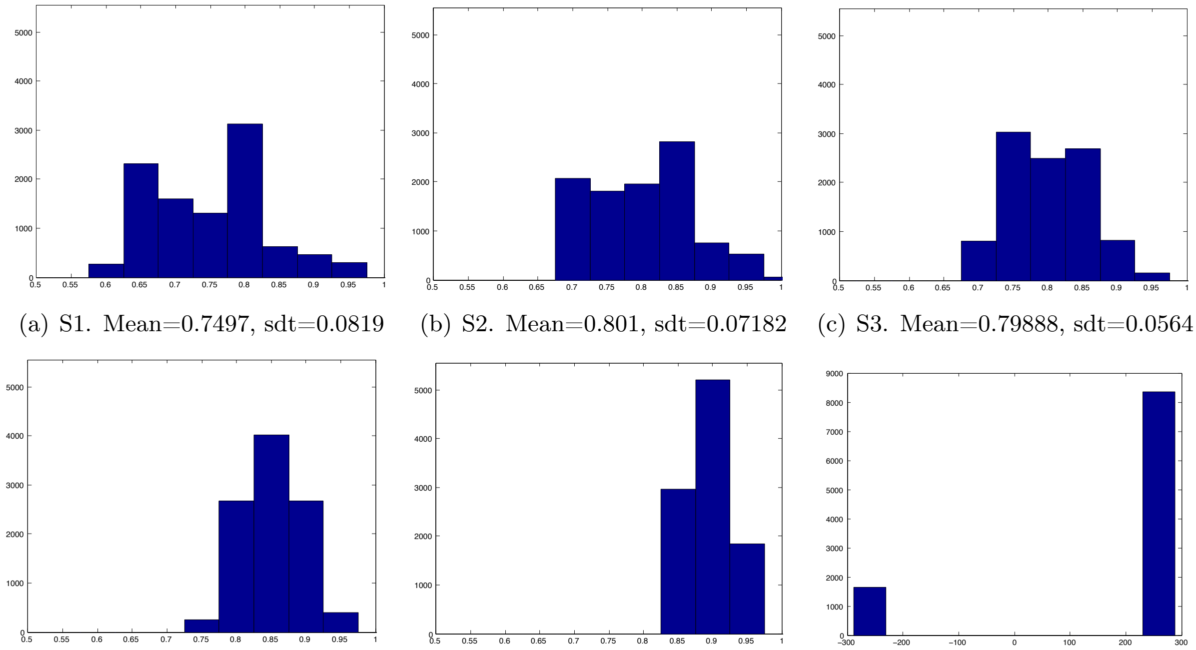

(a) S1. Mean=0.7497, sdt=0.0819 (b) S2. Mean=0.801, sdt=0.07182 (c) S3. Mean=0.79888, sdt=0.0564

(d) S4. Mean=0.8501, sdt=0.0424 (e) S5. Mean=0.8997, sdt=0.0268 (f) demand minus delivery

Figure 1: Deterministic demand. Worst case distribution of reliabilities for suppliers (a), (b), (c), (d), (e) and distribution of the demand minus the total supply (f).

Target shortage probability	0.1	0.05	0.01
Shortage cost	5,575	11,980	66,170
Expected delivery	7594	7645	7855
Expected profit	341,837	280,573	12,477

Table 9: Impact of the shortage cost for the producer case. The shortage cost is adjusted so as to meet the same probability of shortage as in 3. It appears that the computed optimal diversification among suppliers is the same as in Table Table reftab:varconst.

is the extremely high value of the shortage cost in order to achieve the required service level. The smaller the shortage probability, the lesser is the conditional expected shortage and the higher the required unit shortage cost. Note that the recommended shortage cost values are very high. Indeed, the penalty on shortage necessary to achieve the desired level of service must be high. On the other hand, the higher the service level, the lesser is the average shortage. Two factors are combined to enforce an increase of the shortage cost: the necessity to increase the global penalty of shortage and the need to compensate for a lesser number of shortage units.

The influence of the salvage value is displayed in Table 10.

4.4 Value of the information

Scarf addressed the issue of the value of information for the simple Newsboy model. What is the maximum loss in profit incurred with the distributionally robust solution when the demand distribution is known? The same question is relevant in the multi-sourcing problem. We shall not answer it in full generality. Rather, we shall analyze the behavior of the robust solution when the demand and the reliabilities follow normal laws with appropriate means and standard

Salvage	Suppliers					Total supply	Profit
	1	2	3	4	5		
0	1021	1197	1436	1749	1902	7304	470,618
100	1133	1312	1540	1788	1546	7318	475,638
200	1320	1505	1717	1851	938	7332	481,642
300	1647	1833	1993	1875	0	7348	489,514

Table 10: Impact of the salvage value in the case of a producer with deterministic plan of 7500 units.

deviations. The results are displayed on Tables 11 and 12.

supplier	1	2	3	4	5	6	
DR expected deliveries	1050	1226	1462	1759	1811	0	7309
	Distributionally robust profit						472,047
	Mean profit on the event tree						491,389
SP expected deliveries	1709	1887	1954	1664	0	0	7213
	Mean profit on the event tree						494,664

Table 11: The DR solution versus the SP solution. The case of the producer with a deterministic plan of 7500 unis. The expected profit is calculated on the event tree based on the normal distributions.

supplier	1	2	3	4	5	6	
DR expected deliveries	1671	1837	1947	1681	0	0	7135
	Distributionally robust profit						434,076
	Mean profit on the event tree						471,461
SP expected deliveries	2314	2367	1994	390	0	0	7065
	Mean profit on the event tree						474,004

Table 12: DR versus SP. The case of the reseller with a random demand. The expected profit is calculated on the event tree based on the normal distributions.

We observe that using the DR solution instead of the optimal SP solution on the sample from the normally distributed parameters induce a minimal loss of profit. In the two cases under investigation the loss is less than 1%. One checks that the ratio between the worst case profit and the optimal profit when the distribution is known to be a multi normal is at least 0.95 in the producer case and 0.90 in the reseller case. We also observe that knowing the distribution makes it possible to lower the total mean delivery and thus decrease the purchasing cost.

It is worth mentioning that the stochastic programming solution with the uniform distribution (demand and reliabilities) is nearly the same as with the normal distribution.

5 Conclusion

The DR solution to the Newsboy problem with multiple uncertain suppliers has many advantages. It is shown to be the solution of a simple conic quadratic optimization problem (27), which is efficiently solved by interior point methods [Grant and Boyd (2008)]. The solvers that are included in the distribution of [cvx (2012)] can easily handle instances with many more suppliers than in our numerical example. The main advantage of the DR approach is that no assumption is made on the probability distributions except that the mean and covariance exist and their value is known. The DR approach avoids numerical computation of expectations which may be challenging with the multiple dimension of the uncertain parameters. Rather, it replaces these expectations by lower bounds that are numerically tractable in the optimization problem.

The DR approach can equally handle a service level constraint, a feature which is particularly relevant in a production framework. The ease of the numerical computation makes it possible to trace the influence of the parameters. For example it appears that a producer (facing a deterministic demand) can better take advantage of diversification among a panel of suppliers with varied costs and standard deviation than a reseller (facing an uncertain demand) with the same panel of suppliers. In the same logic, it appears that a service level constraint is much more demanding for the reseller than the producer. The numerical experiments also show that recent results on the value information for the simple Newsboy problem [Yue et al. (2006)] hold for the multi-sourcing case: if the probability distribution is not the worst case one, the DR solution still performs well in comparison with the optimal solution associated with this distribution.

In this paper we chose to illustrate the DR approach on a simple basic extension of the classical Newsboy problem. In view of the simplicity of problem (27) additional constraints on the orders can be incorporated. Convex constraint $q \in \mathcal{Q}$, with $\mathcal{Q}$ convex, would not add numerical difficulties. Fixed costs or thresholds on the orders would force to resort to nonlinear integer programming, a more challenging area. However, for the relatively small number of suppliers that is considered a branch and bound scheme would perform well. Another much desirable extension would be the multi-period case. It has been shown [Ben-Tal et al. (2005)] that robust optimization works pretty well on multi stage inventory problems. It would be interesting to extend the DR approach to this situation.

References

[Agrawal and Nahmias (1997)] N. Agrawal and S. Nahmias. 1997. Rationalization of the supplier base in the presence of yield uncertainty. *Production and Operations Management*, 6(3):291–308.

[Anupindi and Akella (1993)] R. Anupindi and R. Akella. 1993. Diversification under supply uncertainty diversification under supply uncertainty. *Management Science*, 39(8):944–963.

[Ben-Tal et al. (2009)] A. Ben-Tal, L.ElGhaoui, and A.Nemirovski. 2009. *Robust Optimization.* Princeton University Press.

[Ben-Tal et al. (2005)] A. Ben-Tal, B. Golany, A.Nemirovski, and J.-Ph. Vial. 2005. Supplier-retailer flexible commitments contracts: A robust optimization approach. *Manufacturing Services and Operations Management*, 7:248–271.

[Bollapragada and Morton (1999)] S. Bollapragada and T. Morton. 1999. Myopic heuristics for the random yield problem. *Operations Research*, 47(5):713–722.

[Burke et al. (2009)] G. Burke, J. Carrillo, and A. Vakharia. 2009. Sourcing decisions with stochastic supplier reliability and stochastic demand. *Production and Operations Management*, 18(4):475–484.

[Chen et al. (2011)] L. Chen, S. He, and S. Zhang. 2011 Tight bounds for some risk measures, with applications to robust portfolio selection. *Operations Research*, 29(4):847–865.

[Chopra et al. (2007)] S. Chopra, G. Reinhardt, and U. Mohan. 2007. The importance of decoupling recurrent and disruption risks in a supply chain. *Naval Research Logistics*, 54:544–555.

[cvx (2012)] CVX Research, Inc. 2012. CVX: Matlab software for disciplined convex programming, version 2.0. `http://cvxr.com/cvx`.

[Dada et al. (2007)] M. Dada, N. Petruzzi, and L. Schwarz. 2007. A newsvendor's procurement problem when suppliers are unreliable. *Manufacturing & Service Operations Management*, 9(1):9–32.

[Federgruen and Yang (2008)] A. Federgruen and N. Yang. 2008 Selecting a portfolio of suppliers under demand and supply risks. *Operations Research*, 56(4):916–936.

[Federgruen and Yang (2009)] A. Federgruen and N. Yang. 2009. Optimal supply diversification under general supply risks. *Operations Research*, 57(6):1451–1468.

[Gallego and Moon (1993)] G. Gallego and I. Moon. 1993. The distribution-free newsboy problem: Review and extensions. *Journal of the Operational Society*, 44(8):825–834.

[Grant and Boyd (2008)] M. Grant and S. Boyd. 2008. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited.

[Gurnani et al. (2013)] H. Gurnani, K. Ramachandran, S. Ray, and Y. Xia. 2013. Ordering behavior under supply risk:an experimental investigation. *Manufacturing Services and Operations Management*, 7:1–32.

[Popescu (2007)] I. Popescu. 2007. Robust mean-covariance solutions for stochastic optimization. *Operations Research*, 55(1):98–112.

[Rekik et al. (2007)] Y. Rekik, E. Sahin, and Y. Dallery. 2007. A comprehensive analysis of the newsvendor model with unreliable supply. *OR Spectrum*, 29:207–233.

[Rockafellar and Uryasev (2000)] T. Rockafellar and S. Uryasev. 2000. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–41.

[Scarf (1958)] H. Scarf. 1958. A min–max solution of an inventory problem. In *Studies in the Mathematical Theory of Inventory and Production*, pages 201–209. Arrow, K.J. and Karlin, S. and Scarf, H., (edts) Stanford University Press, Stanford.

[van Delft and Vial (2013)] C. van Delft and J.-Ph. Vial. 2013. A note on "sourcing decisions with stochastic supplier reliability and stochastic demand.". Working paper (submitted to *Production and Operations Management*), HEC-Paris.

[Wang et al. (2010)] Y. Wang, W. Gilland, and B. Tomlin. 2010. Mitigating supply risk: Dual sourcing or process improvement? *Manufacturing & Service Operations Management*, 12(3):489–510.

[Yano and Lee (1995)] C. Yano and H. Lee. 1995. Lot sizing with random yields: a review. *Operations Research*, 43(2):311–334.

[Yue et al. (2006)] J. Yue, B. Chen, and M. Wang. 2006. Expected value of distribution information for the newsvendor problem. *Operations Research*, 54(6):1126–1136.