A COMPARISON BETWEEN TWO ROBUST INVENTORY POLICIES

Alessandro Ceraudo
Department of Industrial Engineering Pontificia Universidade Catolica Rio de Janeiro
Rua Marques de São Vincente 225 Rio de Janeiro
sandycer79@hotmail.com

Silvio Hamacher
Department of Industrial Engineering Pontificia Universidade Catolica Rio de Janeiro
Rua Marques de São Vincente 225 Rio de Janeiro
hamacher@puc-rio.br

ABSTRACT

This article reviews a single echelon multi-period inventory problem subject to uncertainty in the demand as proposed by Bertsimas and Thiele (2006). We gave insights over the nature of the solution and in particular over its dependency on various parameters. We then tested the optimal robust policy as in Bertsimas and Thiele (2006) comparing it with another robust policy: an adaptation of the classical kσ policy; we showed how this adaptation fall within the family of robust ellipsoidal policy where uncertainty is only present in the right hand side; we then compared, along 5000 simulations, the two optimal policies against various realizations of the demand and showed that the adapted kσ policy is on average better than the robust policy: we provided insights and explanations for this phenomenon.

KEYWORDS – Robust Optimization, Inventory, Mathematical Programming

RESUMO

Este artigo apresenta o problema de estoques multi-período sujeito a incerteza na demanda, conforme formulação de Bertsimas e Thiele (2006). Inicialmente é analisada a solução do problema de programação matemática e a dependência entre os parâmetros do modelo. Em seguida é testada a política robusta ótima proposta por Bertsimas e Thiele (2006), que é comparada com outra política robusta: um adaptação da política clássica kσ. É mostrado que

1. Introduction

One of the most popular studied problems in the supply chain O.R. literature is stock management. As of state of the art, tremendous achievements were reached when no uncertainty is considered. This situation changed whenever one allows some parameters of the system to be random variables. For example, lead time, demand, replenishment time can all be considered uncertain quantities.

In this paper we considered uncertainty only in the demand process; the approach we chose to follow does not need any information about the probability density function of the demand (as opposed to the stochastic framework where this function has to be explicitly specified). We only required that demand lay within a closed and bounded interval called uncertainty set or support, where we specified a nominal value and the half-length, or radius, of the interval: for example if the nominal value is 100, in appropriate units, and the radius of the interval is 20, we require our demand to fall within the interval [80 120]. Nether less to say if we allow the nominal value of the uncertainty set to be the mean and equal probability within the set, we are implicitly considering a uniform distribution.

The main contributions of this article are three fold:

a) Gaining insights on the nature of the solution of the robust optimization algorithm presented in Bertsimas and Thiele (2006) (called from now on the budget of uncertainty algorithm) in particular on its dependency on input parameters that has not yet been stated in the literature to the best of our knowledge.

b) We tested the (poor) effectiveness of the robust policy in cases where demand is correlated over time and non-stationary (which is classical in an industry environment especially when dealing with new and /or short cycle products) and we gave explanation for that.

c) We made a comparison and then an assessment, between the robust policy subject to the budget of uncertainty and another robust policy based on hypothesis of bounded demand: a kσ policy. Although this second policy was already introduced in the literature, to the best of our knowledge, a comparison between the two policies, in both the cases of stationary and non-stationary demand, has yet to be done.

This paper is organized as follow: in section 2 we state the problem and we then follow with a brief literature review on operation research problems subject to uncertainty aiming attention to the robust framework; in section 3 we review the model of Bertsimas and Thiele (2006), we explain the choice of parameters and we gain important insights on the nature of the solution in the case of a single echelon problem; in section 4 we re-introduce a kσ policy and we explain why this methodology follows the framework of the ellipsoidal uncertainty set introduced in Ben-Tal and Nemirovski (1999); we then present extensive simulation results: a comparison between the two policies is made and we gain insights on the reason why both policies can fail in the case of non-stationary demand. Section 5 concludes.

2. The Uncertainty Nature of the Problem

As stated before, we considered a stock management problem where uncertainty is present in the demand process. To solve this problem we use the robust framework approach; an extensive overview of the robust framework literature up to 2007 is made in Bertsimas et Al. (2007). It is important and helpful in this context, to point out that an optimal solution is said robust because it remains feasible against all the possible realizations of the uncertainty parameters within the uncertainty set; there is clearly a tradeoff between optimality and robustness: the bigger the uncertainty set, the worse the solution in terms of optimality. To deal with this tradeoff, Ben - Tal and Nemirosky (1999) proposed a model where uncertainty falls within an ellipsoidal set, moving away from the very first robust methodology introduced by Soyster (1973), the so called worst case scenario (this name comes from the fact that every uncertainty variables are set to their worst case value). Because introducing an ellipsoidal uncertainty set means adding to the problem some conic quadratic constrains, Bertsimas and Sim (2004) introduced a framework (usually called the budget of uncertainty), where they narrowed the uncertainty set symmetrically around the nominal value. In this approach, there is no need to introduce quadratic constraints. Bertsimas and Thiele (2006) applied successfully this methodology in the case of a stock management problem. We would like to point out that all these methods still lay in the framework of “worst case”, they just change the place (i.e. the uncertainty set, or support) where this worst case is chosen to lay.

3. Problem formulation

Because dealing with uncertainty is far more difficult than solving the deterministic problem, we relied our analysis on a stylized and as simple as possible representation of the reality: our aim was not providing a tool of implementation for a real industry but rather gaining insights that can help managers to deal with demand uncertainty. This is the reason why we applied the framework to a single echelon, one product supply chain.

Consider a single firm producing a single good over a finite period of time. This firm faces external demand for this good from a client. The aim of the firm is to produce (order) an optimal quantity over time in order to minimize its cost function. The cost incurred by the firm has two components: a cost of production and a cost incurred if the firm does not produce the exact quantity requested by its clients, i.e. a holding cost if it produces in excess and a backlogging cost if it produces less. We are implying that demand is fully backlogged over time (i.e. whenever demand is not satisfied at any given period of time, it will be satisfied over the next period(s)). The production cost is made of two components: a linear one with respect to the production quantity, we express it with C [$/produced units], and a fixed cost K [$] incurred whenever the firm produces something. The holding and backlogging cost are linear with respect to the current inventory over time, respectively as H and P [$/units]. Note that P is conceptually different from a lost sale cost: it just represents the cost the firm is incurring whenever it delays its production.

These considerations allow us to use the following mathematical program introduced in Bertsimas and Thiele (2006)

Set

Time (k)

Parameters

K fixed cost
H holding cost
P backlogging cost
C unitary production cost
M big m
T number of periods in horizon
X_0 initial stock level

Variables

u_k order quantity
v_k binary variable that takes value 1 if a order is made in the k period, 0 otherwise
y_k back-logging vs. holding cost variable
w_k customer demand

Mathematical Program (M.P.)

$$min_{u,v,y} \sum_{k=1}^{T} (Cu_k + Kv_k + y_k) \qquad (1)$$

s.t.

$$y_k \geq h\left(x_0 + \sum_{k=1}^{T} (u_k - w_k)\right) \qquad \forall\, k = 1,2 \ldots T \qquad (2.a)$$

$$y_k \geq -p\left(x_0 + \sum_{k=1}^{T} (u_k - w_k)\right) \qquad \forall\, k = 1,2 \ldots T \qquad (2.b)$$

$$0 \leq u_k \leq Mv_k \qquad \forall k = 1,2, \ldots T \qquad (3)$$

$$v_k \in \{0,1\} \qquad \forall k = 1,2 \ldots T \qquad (4)$$

We now explain how this formulation works; note that the current inventory is our state variable, called it x, and it is at any given time k, the sum between the inventory level at the previous period and the difference between the produced quantity and the realized demand in the current period:

$$x_{k=} x_{k-1} + u_k - w_k \qquad (5)$$

We can thus obtain a recursive formula for the state variable

$$x_k = x_0 + \sum_{i<k} (u_i - w_i) \qquad (6)$$

At each time period k = 1,2,3…T, the firm will incur a cost

$$Z_k = K + C * u_k + \max\{ -P * x_k \, , \; H * x_k \} \tag{7}$$

We will then try to minimize the sum over k of this last equation; using the state equation (6) in (7) we can express the holding vs. backlogging cost component as a function of demand and produced quantity. Note that we need an initial stock level, i.e. the positioning of the stocks before the horizon begins, x_0.

It is worth noting that the firm will never pay both the backlogging and holding cost in the same period of time, but just one of those: the latter if the inventory x is positive and the former if x is negative in any given period; an important issue is that this formulation is not convex given the presence of a max operator in the objective function. To avoid this problem we will then introduce an auxiliary non negative variable y (for each time step instance) subject to the two antithetic constrains (2.a) (2.b): at optimality only one of the two constrains is binding (representing the actual cost incurred by the firm), while the other one will be trivially satisfied; this observation allows us to obtain a convex representation of the optimization problem as in Bertsimas and Thiele (2006).

The presence of a fixed cost in the objective function leads clearly to a mixed integer problem and one has to deal with binary variables: constrains (3) and (4) are introduced for this
.

3.1 The Robust Optimization Approach

We present here the framework to deal with uncertainty; we stated that the only random variable is the demand over time (i.e. demand is a random walk and can be seen as a trajectory moving as time passes). The basic assumption is that final demand is i.i.d, uncorrelated and bounded; we imply as stated before that demand at any time step k falls within a limited and bounded interval; the center of this interval is the nominal value of the demand; the radius of the interval is proportional to the variance of the demand. Consider the original deterministic formulation of the problem and replace the demand w_k by its random counterpart; it is clear now that we are dealing with infinite constrains because infinite are the possible realization of w_k within its uncertainty interval for each time step k; the aim is finding a robust solution, that is, it remains feasible (i.e. satisfies all the constrains) whatever the realization of w_k will be. Before entering into details of the robust counterpart formulation of the problem, it is worth highlighting that intuitively for a solution to be robust feasible, it has to satisfy the worst case realization of the uncertain parameters because in that case, it will satisfy any other realizations; there is a sort of "link" between the uncertainty set and the feasible set, the bigger the former the smaller the latter: in our problem we have a number of uncertain parameters equal to k (the realization of demand at every time step), and worst case scenario in our case will be when all the realization of our k random variables will be at the worst value (at one of the extreme value of the uncertainty set for each time step t): this is true due to the piecewise linear structure of the objective function. It is also intuitive that the more robust a solution will be (in the sense that it is giving protection against a very large uncertainty set), the worst the value of the objective function (in this case, higher costs). Note that whenever we are asking protection against the worst case, we are implying that all our k random variables realizations are at the worst case.

To cope with this overprotection phenomenon, Bertsimas and Thiele (2006) introduced a parameter, Γ, called budget of uncertainty: this parameter indicates (in a sense that it will be clear later) the maximum number of parameters that can be at their worst case realization.
Let´s then introduce the standardize random variable z_k:

$$z_k = \frac{w_k - \mu_k}{\sigma_k} \tag{8}$$

where μ_k is the center of the support or nominal value of the random variable w_k and σ_k is the half-length or radius. For each k, z_k is then a number in [0,1] and it represents, in a sense, how far w_k is from its nominal value, relatively to the half-length of the symmetric support (see equation (10) below); we then impose a further constrain: the sum of all the deviations from the nominal value has to be less than the budget of uncertainty for any given time step.

$$\sum_{k=1}^{t} |z_k| \leq \Gamma_k \tag{9}$$

Note that we can write

$$w_k = \mu_k + z_k * \sigma_k \tag{10}$$

Remembering that we are dealing in a "against worst case scenario protection" formulation (where the worst case is now limited by the budget of uncertainty, and it lays at the end of the budgeted uncertainty set), we solve for each pair of shortage-holding constraints (details provided in Bertsimas and Thiele (2006).

$$max \sum_{k=1}^{t} \sigma_k \; * z_k \tag{11}$$

s.t.

$$\sum_{k=1}^{t} z_k \leq \; \Gamma_k \qquad \forall k \tag{12}$$

$$0 \leq |z_k| \leq 1 \qquad \forall k \tag{13}$$

These problems can be solved all at once using strong duality (see Bertsimas and Thiele (2006) for details). Note that following this approach we do not change the nature of the deterministic optimization problem, i.e. the robust counterpart remains linear.

3.2 Numerical Example: Data

This section presents the data used in our experiments. This is particularly important because for some choices of parameters the solution is trivial, for instance if fixed cost is too high or production or holding costs are bigger than backlogging cost , the optimal solution can be the trivial, i.e. doing nothing. To avoid this problem, we choose parameters that won´t make the problem trivial, i.e, C=1, P=6, H=4, and K =100.

We first analyzed a case in which demand is stationary over time i.e. its average and standard deviation won´t vary over time but will be kept constant at 100 and 10 respectively. The stationary hypothesis is made primarily to understand the behavior of the robust framework in an "easier" case. The extension to the case of non-stationary demand is trivial from a modeling point of view, and it was also analyzed in what follows.

It remains to fix the budget of uncertainty: we first observed that the number of uncertain parameters (demand realizations) varies with time; to make things clearer, suppose k=1, then we have only one uncertain parameter which is the realization of demand at time step 1, w_1; suppose k=2, then we do have two uncertain parameters that are demand at time step 1 w_1 and demand at time step 2 w_2; by induction this shows that the number of uncertain parameters grows linearly with time; it then makes sense to have a budget of uncertainty that is not fixed but variable over time. As we said before at each time step k, the number of uncertain parameters is exactly k (i.e. the demand realizations up to k) and so at most k of them can be at their worst case value; this implies the budget of uncertainty at each time step k is less or equal to k. Second note that between any two consecutive time instants k, k+1 the numbers of uncertain parameters have grown of one unit, i.e. the difference between the budget of uncertainty in two consecutive periods, $\Delta\Gamma$, is at most 1: any bigger growing rate would be ignored in the sense that the newest uncertainty parameter would be set up to its worst case value; we are now showing a numerical example to make things clearer: according to table 1, when k=15 G=2.04; this means that the cumulative deviation of the fifteen random demands from their original nominal value cannot exceed 2.04 or, say at most two demand realizations are allow to be at their worst case level (nominal value plus or minus their deviation). We chose to follow the framework of Bertsimas and Thiele (2006) and the value of the budget is reported in table 1.

Time	Γ	Time	Γ
1	0,721688	11	1,767767
2	0,883883	12	1,83995
3	1,020621	13	1,909407
4	1,141089	14	1,976424
5	1,25	15	2,041241
6	1,350154	16	2,104064
7	1,443376	17	2,165064
8	1,530931	18	2,224391
9	1,613743	19	2,282177
10	1,692508	20	2,338536

Table 1 : Budget of uncertainty as function of time

The reader who is interested, can found details on how this choice is made in Bertsimas and Thiele (2006)

3.3 Numerical Example: Results

We present in this section some preliminary results; the problem was implemented in AIMSS 3.11 and solved with CPLEX 12.1, in Intel® Core ™ 2 CPU 6300 @ 1.86GHz computer. The output is u_k, the ordered quantity as function of time, but it is also meaningful to report the state variable x_k which represent the inventory level at each time step k. In table 2 we present the results of applying the robust framework to a stationary demand of mean 100 and constant half length of the support 10, subject to the budget of uncertainty presented in table 1 and the choices of parameters as in section 3.2

Time	u_k	Time	u_k
1	100,46	11	104,75
2	100,88	12	105,20
3	101,30	13	105,66
4	101,70	14	106,12
5	102,13	15	106,59
6	102,55	16	107,06
7	102,98	17	107,54
8	103,41	18	108,03
9	103,85	19	108,53
10	104,30	20	109,03

Table 2: Optimal order quantity as function of time

We would like to point out the attention of the reader on the fact that the order quantity is just a slight variation of the average demand (that in this example was set to 100). The robust framework leads us to order a quantity slightly above the average mean of demand: more details are provided in Bertsimas and Thiele (2006). Intuitively the protection against the worst case is made by ordering a quantity higher than the average demand because backlogging costs are higher than holding costs and so one is better off risking some holdings rather than shortage. Finally to assess how good this policy behaves in real situations, we plot the optimal solution of the robust formulation against 5000 simulation of demand over the entire length of the time horizon; then we averaged the policy cost along all the simulations and include some descriptive statistics, like standard deviation, minimum costs maximum costs and number of stock-outs.

To make the analysis more reliable we tested the robust policy against various possible distribution of demand: we chose to use a uniform and normal distribution Results are reported in the following tables

normal (100,20)	
6708	average cost
17494	max cost
3083	min cost

uniform (80 120)	
18418,2	average cost
26737,7	max cost
11655,54	min cost

Tables 3 and 4 : average, max and min cost of optimal policy

As the reader can see, the robust policy behaves on average much better when the underlined demand process follows a normal distribution rather than a uniform one (with the same mean and standard deviation).

Figure 1 represent the cost of the optimal policy for a random sample of 256 simulations extracted over the 5000 simulations in the case of uniform distribution of the underlined demand: as the reader can see, there is a reasonable gap between maximum and minimum cost.

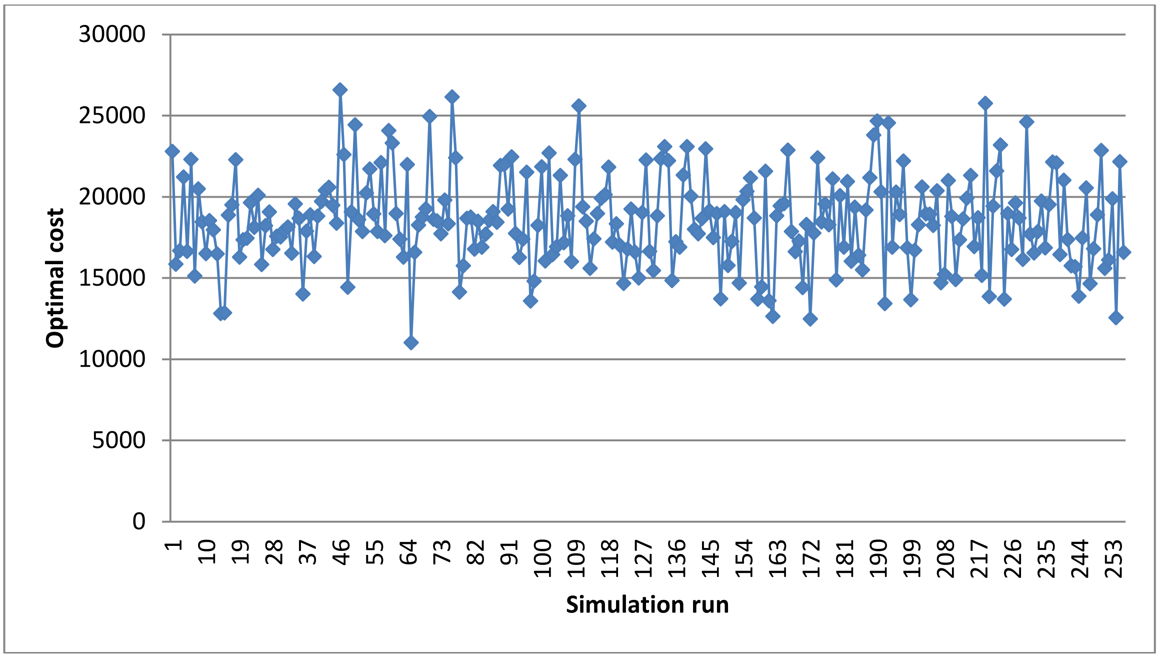

Figure 1: Optimal policy cost for normal distribution demand

It is worth noting that the robust policy behaves pretty well independently of the underlying ‘real’ demand distribution provided there is a good guess of the first moment of it, i.e. the mean (as noted before, the robust policy leads to order a quantity slightly higher than the mean).

4. Another Robust Policy

A strong hypothesis of the previous section (and in this context, of the robust framework) is that demand is bounded: we introduced the uncertainty set and we forced demand to fall within it, we then bounded the uncertainty set by the budget of uncertainty so to reduce it. This hypothesis does not follow what is normally assumed by the stochastic approach (for example, when demand is assumed to be a normal or a Poisson process, it is clearly not bounded). It is common anyway in the literature to assume – see for instance Graves (2000) and references therein - that demand is bounded by m+(θ*σ) or m-(θ*σ) where σ is the standard deviation, θ is a safety factor and m the mean of the random process. It is clear that in this case the worst case is represented by one of the two extreme points m+(θ *σ) or m-(θ *σ) at every time step k. Following the same approach as before we now impose that at each time step the pair of backlogging holding constrains are

$$y_k \geq h\left(x_0 + \sum_{i=1}^{k} u_i - w_i + \sqrt{\sum_{i=1}^{k} \sigma_i^2}\right) \qquad \forall k = 1,2 \ldots k \qquad (14)$$

$$y_k \geq -p\left(x_0 + \sum_{i=1}^{k} u_i - w_i - \sqrt{\sum_{i=1}^{k} \sigma_i^2}\right) \qquad \forall k = 1,2 \ldots k \qquad (15)$$

The reader can note that we introduced an ellipsoidal uncertainty set as in Ben - Tal (1999); the steps to obtain formula (14) and (15) are as followed: we first defined the uncertainty set as $B = \left\{d \in R^K \middle| \sum_{i=1}^{k} \sigma_k^{-2} * |d_k - m|^2 \leq \vartheta^2 \ \forall\ k = 1..T\right\}$ where d is a particular realization of the demand random process; we then wrote the robust counterpart (also called deterministic equivalent) of our problem by applying formula 14-15-16 of Ben-Tal (1999) and as result we obtained the constrains (14) (15). It is also worth noting that our example is quite similar to that one presented in section 4 of Ben-Tal (1999) where uncertainty affected only the right hand side: this is a remarkable fact because, like the budget of uncertainty policy, the ellipsoidal policy didn´t change the nature of the problem in the sense that no nonlinear terms are introduced. In this case then the ellipsoidal robust counterpart is numerically tractable and can be solve with standard software (in general ellipsoidal constrains can lead to a conic quadratic robust counterpart that may not be solved efficiently). Cleary the optimal policy is a correction of the mean by a term that is proportional to the deviation, but this term does not increase linearly over time because of the square root effect (this effect does not necessarily happens with the budget of uncertainty policy). This reduction over time makes sense because we have a compensation effect along the horizon: deviation from the mean happens up and down so they cancel them out, or in statistical term the variance of the sum of two random variables is the sum of the variances, but the resulted standard deviation is not the sum of the standard deviations.

We made a comparison of the two robust policies: we plotted the two robust solutions against the realization of 5000 random demand where we varied the underlined probability density function of the random variable. We considered both normal and uniform distributed uncorrelated demand; for the ellipsoidal policy we considered two values of the safety factor: k = 1 and k = 2. Figures 2 and 3 report the average inventory level for each time step t of the three optimal policies when the underlined demand follows a uniform and normal distribution respectively (the inventory at each time step was averaged along the 5000 simulations).

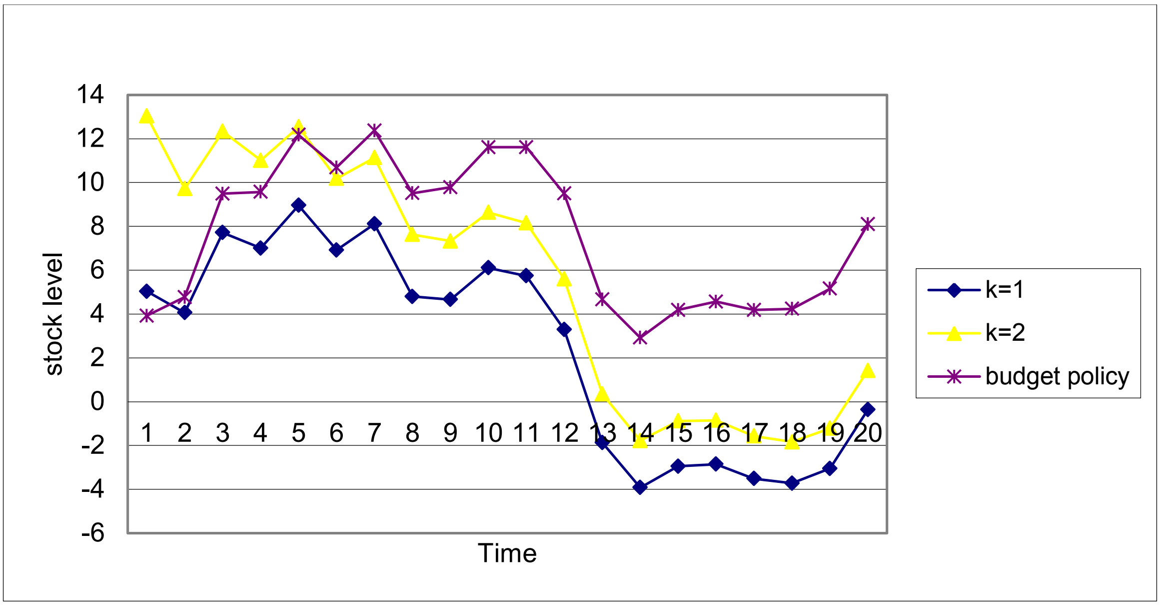

Figure 2: Average stock level for the three policies when demand is uniform

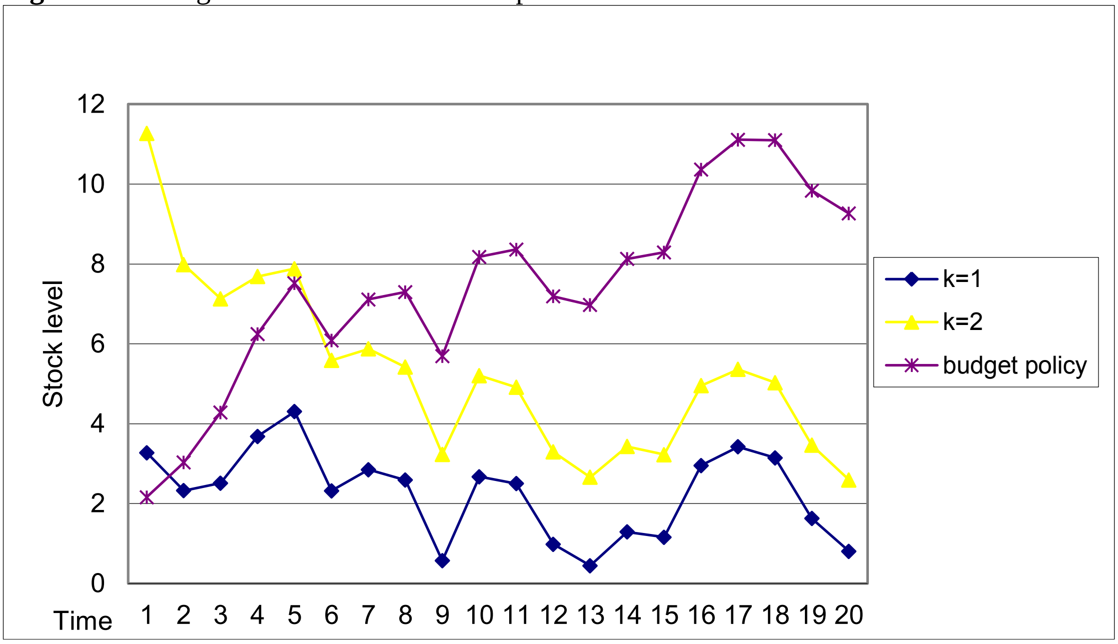

Figure 3: Average stock level for the three policies when demand is normal

As it can be seen, the policies behaved pretty well (provided we have a good guess of the first moment of demand, regardless of how good the guess is about the second moment (see previous section for details), but the ellipsoidal policy is slightly better, especially when the underlined demand is normal. This is clearly explained because the ellipsoidal policy is able to mimic pretty well a normal distribution (on average 95% of the realization falls between the mean plus minus two times the deviation).

We then proceeded to test the two policies against realizations that were not independent, i.e. time correlated or non-stationary demands; we considered for example an ARIMA 1 demand process like in Sodhi (2011): results are summarized below in figure 4 where we plotted the average stock level for each time step t along the 5000 simulations

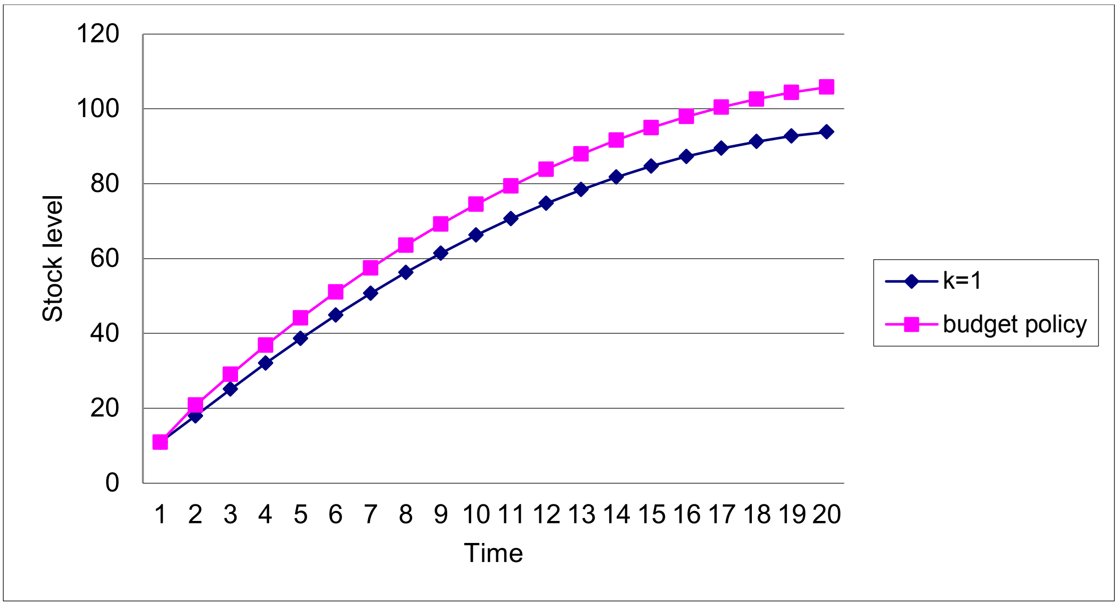

Figure 4: Average stock level for the two policies when demand follows an ARIMA process

The poor results were due to the fact that the iso-variance curves of the stochastic demand process were in this last case not ellipsoidal: both the methods failed because they were meant to mimic iso-variance curves that are ellipses. This fact was noticed by Sim (2010), but he did not provide a numerical example of this phenomenon like we did.

5. Conclusions and Future Works

In this paper we reformulated a single echelon single product stock problem when demand is uncertainty; we choose to follow the robust framework so that the main hypothesis is that demand is bounded.

We analyzed two robust formulations of the problems: the one proposed in Bertsimas and Thiele (2006) and a $\theta * \sigma$ policy; we made a comparison between the two by mean of extensive simulation results; we showed that both methods lead to an optimal solution that is a correction of the mean of the demand random process and we showed that they behave pretty well when a good guess of the mean of demand is given. Both policies fail when demand is correlated over time as we showed when we plotted those policies against a demand that follows a ARIMA 1 process.

We then gain the following novel insights:

- The failure of the two policies is due to the fact that the iso-variance curves of an ARIMA process are not ellipses: both policies are intended to reproduce the behavior of a random process whose iso-variance curves are indeed ellipses.
- In the case of i.i.d non correlated demand, the $\theta * \sigma$ policy in general behaves better because the ellipsoidal uncertainty set "matches" the ellipsoidal curve while the polyhedral uncertainty set of the budget policy doesn´t do it exactly.

A possible future extension of this work is to take advantage of the adversarial structure of the problem and study it by a game theoretical point of view: we envision being able to give a closed formula solution of the optimal policy.

6. References

Ben-Tal A., Nemirovski A. (1999), Robust Solution of Uncertainty Linear Program, *Operation Research Letter* 25, 1-13

Bertsimas D., Thiele A. (2006), A Robust Optimization Approach to Inventory Theory, *Operation Research* 54,1, 150-168

Bertsimas D., Sim M. (2004), The Price of Robustness, *Operation Research* 52, 35-53

Bertsimas D., Brown D.B., Caramanis C. (2007) Theory and Applications of Robust Optimization, *ArXiv:1010.5445v1*

Graves S.C., Willems P. (2000), Optimizing strategic safety stock placement in supply chains, *Manufacturing and Service Operation Management* Vol.2 n.1 pp. 68-83

Goh J., Sim M. (2010), Robust Optimization made easy with Rome, *Operation Research* forthcoming

Sodhi M., Tang C. (2011), Determining supply requirements in the sales-and-operation planning (S&OP) process under demand uncertainty: A stochastic programming formulation and a spreadsheet implementation, *JORS*, 62, p.526-536

Soyster A. L., (1973), Convex programming with set-inclusive constraints and applications to inexact linear programming, *Operation Research* 21, 1154-1157