

e - c o m p a n i o n

ONLY AVAILABLE IN ELECTRONIC FORM

Electronic Companion—“Approximation Algorithms for Capacitated
Stochastic Inventory Control Models” by Retsef Levi, Robin O. Roundy,
David B. Shmoys, and Van Anh Truong
Operations Research, DOI 10.1287/opre.1080.0580.

Approximation Algorithms for Capacitated Stochastic Inventory Control Models - Online Appendices

Retsef Levi

Sloan School of Management, MIT, Cambridge, MA, 02139., retsef@mit.edu,

Robin O. Roundy

School of ORIE, Cornell University, Ithaca, NY 14853, robin@orie.cornell.edu,

David B. Shmoys

School of ORIE and Dept. of Computer Science, Cornell University, Ithaca, NY 14853., shmoys@cs.cornell.edu,

Van Anh Truong

School of ORIE, Cornell University, Ithaca, NY 14853, truong@orie.cornell.edu,

Subject classifications: Stochastic Inventory Control; Heuristics; Approximation Algorithms.

Area of review: Supply Chain Management.

History: Submitted August 2005. Revised July 2006, December 2006, May 2007, September 2007, September 2007, October 10, 2007.

Appendix A: Marginal Backlogging Cost Accounting Approach - Numerical Example

To provide more intuition, we illustrate the new backlogging cost accounting through a simple example. Suppose that the order capacity is 5 in all periods, $L = 0$ and $\alpha = 1$. Assume that the inventory position at the beginning of period 3 was $x_3 = 3$, and that we have ordered $q_3 = 3$, $q_4 = 5$, $q_5 = 4$ and $q_6 = 2$ units in periods 3, 4, 5 and 6, respectively. Now say that the demands were $d_3 = 3$, $d_4 = 3$, $d_5 = 5$ and $d_6 = 11$ in periods 3, 4, 5 and 6, respectively. In particular, the accumulated demand over periods $[3, 6]$, $d_{[3,6]}$, is equal to 22. This implies that in period 6 we had a shortage of 5 units, each of which incurred a penalty cost of p_t at the end of period 6. Out of these 5 units of shortage at the end of period 6, we associate a backlogging penalty of 3 units of shortage with period 6 (the unused slack capacity in this period is 3), a penalty of 1 unit of shortage with period 5 (the unused slack capacity in this period is 1), no cost is associated with period 4 since we ordered up to capacity, and finally the penalty of 1 units of shortage is associated with period 3 ($d_{[3,6]} - (3 + 3 + 5 + 5 + 5) = 1$). In other words, $w_{36} = 1$, $w_{46} = 0$, $w_{56} = 1$ and $w_{66} = 3$. This example illustrates how we backtrack the ‘source’ of each unit of shortage and its corresponding backlogging cost incurred in period t , and associate it as forced backlogging cost to past periods. If $L > 0$, then we start the backtracking in period $t - L$, since only orders in periods earlier than $t - L + 1$ could have arrived by time t .

Appendix B: Experimental Design

In this appendix we give a detailed description of the scenarios that form the basis of the experiments done in Section 6. The space of potential parameter settings for this study is very large. In addition to parameters describing the inventory system, there are many parameters that describe the manner in which forecasts of demand evolve over time. A fully comprehensive study is beyond

the scope of this paper. Our goal is to study a broad range of potential application settings, with emphasis on the demand and forecasting processes.

The experimental design is oriented around a Base Case and six sets of *scenarios*, each of which expands the Base Case in an interesting dimension. In each set of scenarios we vary specific input parameters. The first three of these scenario sets study *first-order effects*, in this case, trends and seasonality patterns in the demand. The final three scenario sets study *second order effects* by varying the probability model that governs the variance of the demand and of the forecast errors, and the correlations that exist between them.

We begin by reviewing the structure, and some of the notation, of the MMFE model. Then we discuss the Base Case. After that we describe the manner in which the parameters of the Base Case are varied, in each of the six scenario sets.

The MMFE Model Hurley et al. (2006) described the MMFE model of forecast evolution. In the multiplicative version of the MMFE, for every pair of times s, t , $0 \leq s \leq t \leq T$, $1 \leq t$, there is a forecast D_{st} of the demand that will occur in period t , which was generated at the end of period s , i.e., at the beginning of period $s + 1$. The actual demand is $D_t = D_{tt}$, observed at the end of period t . We assume that forecasts are unbiased, so that $D_{st} = E[D_t | \mathcal{F}_{s+1}]$. There is a forecast horizon $H \leq T$. The corporate forecasting process generates forecasts that extend H time periods into the future. Therefore, $D_{st} \neq D_{s-1,t}$ if $t < s + H$, because in that case D_{st} was effected by the forecasting process that occurred at the end of period s . However, if $t \geq s + H$ then the end-of-period- s forecasting process did not consider the period- t demand, and $D_{s,t} = D_{s-1,t}$. At the beginning of the time horizon we are given the initial set of forecasts, $\mathbf{d}_0 = (d_{0,t} : 1 \leq t \leq T)$. (In this case we use lower case because these forecasts have already been observed). Seasonality and trend are introduced into the model by choosing the vector \mathbf{d}_0 appropriately.

We model the process by which forecasts are created as follows. The period- s update vector is $\gamma_s = (\gamma_{st} : s \leq t < s + H)$. At the end of time period s the update vector γ_s is observed, and the multiplicative MMFE model updates forecasts using the formula $d_{st} = \gamma_{st} d_{s-1,t}$ for $t = s, s + 1, \dots, s + H - 1$, and by $d_{st} = d_{s-1,t}$ for $t \geq s + H$. In our experiments $\gamma_s = e^{\epsilon_s}$, where the H -dimensional random vector ϵ_s is normally distributed with mean $-\frac{1}{2} \text{diag}(\Sigma_s)$ and variance-covariance matrix Σ_s , and γ_s has a multivariate lognormal distribution whose mean is a vector of ones. $\Sigma_s \sim \Sigma$ and $\gamma_s \sim \gamma$ are both stationary over time.

In the multiplicative MMFE model, it is not hard to show that at the end of period s , given the current information set f_{s+1} and forecast vector $\mathbf{d}_s = (d_{s,t} : s \leq t \leq T)$, the future demands ($D_{tt} : s < t \leq T$) have a conditional distribution that is multivariate lognormal, with easily-computable parameters. Three of our six scenario sets study second order effects, which we create by using different variance-covariance matrices Σ .

The Base Case In the Base Case our holding and backorder costs per unit per period are stationary, equal to $h_t = 1$ and $p_t = 10$. All experiments are conducted for two different lead times: $L = 0$ and $L = 4$. Therefore, to facilitate comparisons between different scenarios, costs are not counted during the first four time periods. Note that when $L = 4$, in the first four time periods the costs incurred are determined by decisions made in the past, and are not influenced by our choice of policy. There is neither trend nor seasonality in the Base Case, so the initial demand forecast is flat, with $\mathbf{d}_0 = (400, 400, \dots, 400)$. The time horizon has length $T = 40$, and the horizon over which the user actively generates forecasts has length $H = 12$. This implies that at all times s , the first 13 elements of the forecast vector \mathbf{d}_s will be different from each other, but the 13-th element and every subsequent element will be equal to 400.

In the Base Case, we have *constant learning*, meaning that all of the entries on the diagonal of Σ are equal. The diagonal elements are selected so that for $t \geq 12$, the coefficient of variation of the period- t demand D_{tt} , seen from the beginning of time period 1, is 0.75.

The off-diagonal entries of the covariance matrix Σ determine the degree of correlation between the updates that are observed in a given time period, say, time period s . The Base Case assumes

that there is some correlation between these updates, modeled by having non-zero, positive values in the first off-diagonal of Σ . Consequently, in the Base Case, if the forecast for the demand in month t will go up in period s (i.e., if $D_{st} > D_{s-1,t}$), then the forecast for demand in month $t + 1$ is likely to increase in period s as well (i.e., $P(D_{s,t+1} > D_{s-1,t+1}) > 0.5$). However, $D_{st} > D_{s-1,t}$ does not tell us anything about the forecast $D_{s,t+2}$ for demand in month $t + 2$. The values of the non-zero off-diagonal elements are chosen to give a correlation coefficient of 0.5 for each pair of adjacent forecast updates. That is, for each s and each t , $s \leq t \leq s + H - 2$, the update factors γ_{st} and $\gamma_{s,t+1}$ observed in period s have correlation coefficient 0.5, but γ_{st} and $\gamma_{s,t+2}$ are stochastically independent.

Product Launch Scenarios In this set of scenarios we study the effect of rising demand, as might be encountered at a product launch. Again, only the initial forecast vector \mathbf{d}_0 is varied. For comparison with the base case, we ensure that the mean of the values in \mathbf{d}_0 is 400. We consider upward demand trends of +5, +10 and +20 per period. In addition, we consider two examples in which the demand rises in a steeper, non-linear manner, mid-way through the horizon; these are generated using an appropriately scaled normal CDF curve. The five initial forecast vectors are plotted in Figure 1.




Figure 1 Initial forecast vectors used in Product Launch Scenarios.

End-of-Life Scenarios Here, we study scenarios associated with products that are in an end of life situation, namely those with decreasing initial forecast vectors. Essentially, these are the reverse of the Product Launch scenarios; we have initial forecast vectors with forecasted demand decreasing by 5, 10 and 20 per period. We also consider two products whose demands have steeper drop-off curves, generated using the normal complementary CDF curve. In addition, we study a total demand crash, in which the demand is forecast to crash to 0 midway through the time horizon.

Seasonality Scenarios In the seasonality study, we use the common base-values described above for all parameters except for the initial forecast vector \mathbf{d}_0 . We conduct experiments with two forms of seasonality, one defined via a sinusoidal function and the other via a step function. In both cases, the maximum value attained is 700 and the minimum is 100. This allows us to compare results more easily with the base case, because the mean of the entries in the initial forecast vector is 400 in all cases.

By the *cycle length*, we mean the number of time periods between two consecutive high-points. We consider cycle lengths with values 2, 4 and 8. For example, for the step-function with period 4, we have $\mathbf{d}_0 = (700, 700, 100, 100, 700, 700, 100, 100, \dots)$.

The above scenario sets test the effect of varying \mathbf{d}_0 , the initial forecast vector. In the final three scenario sets, we focus instead on varying Σ . In all of these, we take $\mathbf{d}_0 = (400, 400, \dots, 400)$.

Coefficient of Variation Scenarios In this scenario set, we study the effect of varying the magnitude of the variance in the demands and the forecasts. Note that for $t \geq H = 12$, at the end of time period $t - H$, we have $D_{tt} = \Gamma_t d_{t-H,t}$, where Γ_t is random and has the same distribution as

$$\Gamma = \Gamma_H = \prod_{i=1}^H \gamma_{H+1-i,H} = \exp \left(\sum_{i=1}^H \epsilon_{H+1-i,H} \right). \quad (1)$$

The $\epsilon_{H+1-i,H}$'s are independent normal random variables, with mean such that $E[e^{\epsilon_{H+1-i,H}}] = 1$, and with variance σ_{ii} , the i -th diagonal element of Σ , our forecast update matrix. (Note that σ_{ii} is a variance, not a standard deviation). Thus, Γ is log-normal, with mean one and variance $e^{(\sum_{i=1}^H \sigma_{ii})} - 1$. The coefficient of variation of Γ , and of D_{tt} for $t \geq H$, is $[e^{(\sum_{i=1}^H \sigma_{ii})} - 1]^{1/2}$. In the Base Case this number is 0.75. In the scenarios where we investigate the effect of variance, we scale the entries of Σ such that the coefficient of variation of Γ takes specific values, namely 0.5, 0.7, 1, 2, 4, and 8. This corresponds to different levels of demand variability.

Time of Learning Scenarios If $s \leq t \leq s + H$ then the logic behind equation (1) above indicates that at the end of period s , the random variable $(D_{tt}|d_{st})$ has mean 1 and variance $e^{(\sum_{i=1}^{t-s} \sigma_{ii})} - 1$. Therefore we use $(\sum_{i=1}^{t-s} \sigma_{ii})$, which in the Base Case ranges from 0 to 0.446, to measure the portion of the total variability in D_{tt} that is still *unresolved* in period s . In Figure 2 we plot $(\sum_{i=1}^{t-s} \sigma_{ii})$ as a function of $t - s$, for $0 \leq t - s \leq H$. The different curves represent four different possibilities for the way in which variability is resolved. In the Base Case we have *constant learning*, meaning that all of the entries in $\text{diag}(\Sigma)$ are equal, and the curve is a straight line. When the diagonal of Σ has relatively large values in the lower right portion of the matrix, the plot is convex, and the unresolved uncertainty is low when s is close to t . This corresponds to *early learning*. Conversely, when the values in the diagonal of Σ are weighted towards the upper right corner of the matrix we have *late learning*, the plot is concave, and most of the uncertainty about the true value of D_{tt} is resolved in periods s that are close to t . We also consider the setting in which there is more weight in the center of the diagonal of Σ than at the extremes. In this case most of the learning takes place near the middle of the forecast horizon.

We construct variance-covariance matrices Σ to correspond with these four cases: constant, early, late and mid-horizon learning. In all cases, the values of Σ are scaled to ensure that the coefficient of variation of Γ , and of D_{tt} for $t \geq 12$, remain constant at 0.75.

Correlation Scenarios In this scenario set we test the effect of different types of correlation between the updates. We vary correlation in two ways. First, we set the number of non-zero off-diagonals of our 12x12 matrix, Σ , to 0 (which corresponds to no correlation), 1, 4 and 8. Secondly, the sign of the off-diagonal elements can be all positive, all negative, or entries alternating between positive and negative. (The base case corresponds to 1 off-diagonal with non-zero elements which are all positive.) As in the base case, the diagonal elements of Σ are all equal (the constant learning case), and the coefficient of variation of Γ is 0.75.




Figure 2 Cumulative sum of the diagonal elements of Σ for constant, early, late and mid-horizon learning.

Table 1 summarizes the scenarios we study. The number of scenarios for each set is given in parentheses after the set name; we see that there are 38 in total. We run each of these with lead times $L = 0, 4$. In addition, the 6 seasonality-based scenarios were run with $L = 8$. That makes a total of $38 \times 2 + 6 = 82$ scenario - lead time pairs. For each of the pairs, we ran $N = 1,000$ independent trials.

Appendix C: Extensions of the Dual-Balancing Policy

Integer-Valued Demands We now discuss the case in which the demands are integer-valued random variables, and the order in each period is also restricted to an integer. A simple, illustrative example of the Dual-Balancing policy with integer-valued demands is found in Appendix A. In the integer-valued demand case, in each period s , the functions $H_s^B(q_s^B)$ and $\tilde{\Pi}_s^B(q_s^B)$ are originally defined only for integer values of q_s^B . We now define these functions for any value of q_s^B by interpolating piecewise linear extensions of the integer values. It is clear that these extended functions preserve the convexity and monotonicity properties discussed in the previous (continuous) case. However, it is still possible (and even likely) that the value q'_s that balances the functions l_s^B and $\tilde{\pi}_s^B$ is not an integer. Instead we consider the two consecutive integers q_s^1 and $q_s^2 := q_s^1 + 1$ such that $q_s^1 < q'_s < q_s^2$. In particular, $q'_s := \lambda q_s^1 + (1 - \lambda)q_s^2$ for some $0 < \lambda < 1$. We now order q_s^1 units with probability λ and q_s^2 units with probability $1 - \lambda$. This constructs what we call a *randomized Dual-Balancing policy*.

Observe that at the beginning of time period s the order quantity of the Dual-Balancing policy is still a random variable $Q_s^B = Q'_s$ with support $\{q_s^1, q_s^2\} = \{q_s^1(f_s), q_s^2(f_s)\}$, which is a function of the

Table 1 Scenario codes

Topic	Description
Product Launch (5)	Increment by I per period, $I \in \{5, 10, 20\}$
	Increasing scaled normal CDF curve
	Steeper increasing scaled normal CDF curve
End-of-Life (6)	Decrement by I per period, $I \in \{5, 10, 20\}$
	Decreasing scaled normal CDF curve
	Steeper decreasing scaled normal CDF curve
	Demand crash
Seasonal (7)	Initial forecast vector is flat
	Sinusoidal periodicity with cycle length n , $n \in \{2, 4, 8\}$
	Step-function periodicity with cycle length n
Coeff. of Var. (6)	Coefficient of variation equals $\beta = 0.5, 0.7, 1, 2, 4, 8$
Learning Rate (4)	Constant learning
	Late learning
	Early learning
	Mid-horizon learning
Correlation (10)	All off-diagonal elements of Σ are 0
	First n off-diagonals of Σ have positive entries, $n \in \{1, 4, 8\}$
	First n off-diagonals of Σ have negative entries
	First n off-diagonals of Σ have entries alternating positive and negative

observed information set f_s . We would like to show that this policy admits the same performance guarantee of 2. For each $t = 1, \dots, T - L$, let Z_t be again the random balanced cost of the Dual-Balancing policy in period t . Focus now on some period s . For a given observed information set $f_s \in \mathcal{F}_s$ we have for some $0 \leq \lambda = \lambda(f_s) \leq 1$,

$$z_s = E[H_s^B(Q'_s)|f_s] = \lambda E[H_s^B(q_s^1)|f_s] + (1 - \lambda)E[H_s^B(q_s^2)|f_s] := E[H_s^B(\lambda q_s^1 + (1 - \lambda)q_s^2)|f_s],$$

and

$$z_s = E[\tilde{\Pi}_s^B(Q'_s)|f_s] = \lambda E[\tilde{\Pi}_s^B(q_s^1)|f_s] + (1 - \lambda)E[\tilde{\Pi}_s^B(q_s^2)|f_s] := E[\tilde{\Pi}_s^B(\lambda q_s^1 + (1 - \lambda)q_s^2)|f_s].$$

The second equality (in each of the two expressions above) is a formal statement of the fact that we extended the domains of $H_s^B(q_s^B)$ and $\tilde{\Pi}_s^B(q_s^B)$ from integer to real values using piecewise linear interpolation. By the definition of the algorithm we have,

$$\lambda E[H_s^B(q_s^1)|f_s] + (1 - \lambda)E[H_s^B(q_s^2)|f_s] = \lambda E[\tilde{\Pi}_s^B(q_s^1)|f_s] + (1 - \lambda)E[\tilde{\Pi}_s^B(q_s^2)|f_s].$$

It is now readily seen that, for each period s and each $f_s \in \mathcal{F}_s$, we again have $E[H_s^B(Q'_s) + \tilde{\Pi}_s^B(Q'_s)|f_s] = 2z_s$, i.e., $E[H_s^B(Q'_s) + \tilde{\Pi}_s^B(Q'_s)|\mathcal{F}_s] = 2Z_s$. This also implies that Lemma ?? is still valid.

Now define the sets \mathcal{T}_H and \mathcal{T}_Π in the following way. Let $\mathcal{T}_H = \{t : X_t^B + Q_t^2 \leq Y_t^{OPT}\}$, and $\mathcal{T}_\Pi = \{t : X_t^B + Q_t^2 > Y_t^{OPT}\}$. Observe that for each period s , conditioned on some $f_s \in \mathcal{F}_s$, we know deterministically x_s^B , q_s^B and, if the optimal policy is deterministic, we also know y_s^{OPT} . Therefore, we know whether $s \in \mathcal{T}_H$ or $s \in \mathcal{T}_\Pi$. If the optimal policy is also a randomized policy, we condition not only on f_s but also on the decision made by the optimal policy in period s . Moreover, if $s \in \mathcal{T}_H$, then, with probability 1, $Y_s^B \leq Y_s^{OPT}$, and if $s \in \mathcal{T}_\Pi$, then, with probability 1, $Y_s^B \geq Y_s^{OPT}$. This implies that also Lemmas ?? and ?? are still valid. The following theorem is now established (the proof is identical to that of Theorem ?? above).

Theorem C.1 *The randomized Dual-Balancing policy has a worst-case performance guarantee of 2, i.e., for each instance of the capacitated periodic-review stochastic inventory control problem, the expected cost of the randomized Dual-Balancing policy is at most twice the expected cost of an optimal solution, i.e., $E[\mathcal{C}(B)] \leq 2E[\mathcal{C}(OPT)]$.*

Stochastic Lead Times In this section, we consider the more general model where the lead time of an order placed in period s is some nonnegative integer-valued random variable L_s . However, we assume that the random variables L_1, \dots, L_T are correlated, and in particular, that $s + L_s \leq t + L_t$ for each $s \leq t$. In other words, we assume that any order placed at time s will arrive no later than any other order placed after period s . This is a very common assumption in the inventory literature, usually described as “no order crossing”.

Similar to Levi et al. (2007), we next describe how to extend the Dual-Balancing policy and the analysis of the worst-case expected performance to this more general setting. For each $t = 1, \dots, T$, let S_t be the latest period for which an order placed in that period arrives before time t . In other words, $S_t := \max\{s : s + L_s \leq t\}$. Now modify the definition of the random variables W_{st} (for each $s \leq t$) to be

$$W_{st} := \min\{\mathbb{1}(s \leq S_t)\bar{Q}_s, \mathbb{1}(s \leq S_t) (D_{[s,t]} - (X_s + Q_s + \sum_{j \in (s, S_t]} u_j))^+\}.$$

Similar to the discussion in Section 3 above, we can write

$$W_{st} = \mathbb{1}(s \leq S_t) \left[(D_{[s,t]} - (X_s + Q_s + \sum_{j \in (s, S_t]} u_j))^+ - (D_{[s,t]} - (X_s + \sum_{j \in [s, S_t]} u_j))^+ \right],$$

and

$$W_{st} = \mathbb{1}(s \leq S_t) \left[(D_t - NI_t - \sum_{j \in (s, S_t]} \bar{Q}_j)^+ - (D_t - NI_t - \sum_{j \in [s, S_t]} \bar{Q}_j)^+ \right].$$

We again define the forced marginal backlogging cost in period s as $\tilde{\Pi}_s = \sum_{t \geq s} p_t W_{st}$. It is straightforward to check that we can still express the cost of each feasible policy P as $\mathcal{C}(P) = \sum_t (H_t + \tilde{\Pi}_t)$. In each period, we again balance the conditional expected marginal holding cost against the conditional expected forced marginal backlogging cost. It is readily verified that the same analysis described in Section 4.1 is still valid.

Theorem C.2 *The Dual-Balancing policy provides a performance guarantee of 2 for the capacitated periodic-review stochastic inventory control problem with stochastic lead times and non-crossing orders.*

References

- Hurley, Gavin, P. Jackson, R. Levi, R. O. Roundy, D. B. Shmoys. 2006. New policies for stochastic inventory control models - a theoretical and computational study. Working paper.
- Levi, R., Martin Pál, R. O. Roundy, D. B. Shmoys. 2007. Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research* **32** 284–302.