Alignment between an ABC classification and results from an optimization approach

Christian Larsen

CORAL, Department of Economics and Business Economics, BSS, Aarhus University, Fuglesangs Allé 4, DK-8210 Aarhus V, Denmark

e-mail: chl@econ.au.dk

Abstract

The paper analyzes a multi-item inventory system, using independent *(R,Q)* policies. The aim is to minimize the average stock level while keeping a weighted (across items) fill rate service level at a given target. The paper explores whether the optimal order sizes and the optimal safety factors are aligned with the ABC classification. Our results show that A items should be ordered more frequently followed by B items and C items. Concerning the optimal safety factors, our results depend on the weights applied in the specification of the weighted service level. If all items have equal weight, A items should have the lowest safety factors followed by B and C items. If the weights are the demand rates, the ranking of the safety factors follows the ratios between the demand rate and the unit purchase price, and the higher this ratio is, the higher is the optimal safety factor. However, this ranking is completely aligned with the ABC classification given that the demands are transformed into monetary units.

Keywords: Inventory control, ABC classification

1. Introduction

In most companies, the number of items is too large to implement specific inventory control policy for each item. Therefore a starting point is to group the items, and for each group to make overall guidelines for the inventory management of the items within a particular group (van Kampen et al. (2012)).

An ABC analysis is a very common starting point for making such a grouping (Silver et al. (2017)). Assume you have in all N items where c_i is the unit purchase price of item i and μ_i is the average demand per time unit of item i. The ABC classification scheme ranks the items based on their $c_i\mu_i$ values. Thus, A items are those with the highest $c_i\mu_i$ values followed by B items and C items. In general, it is acknowledged that A items should be monitored most closely followed by B items while C items can be monitored more loosely. When comparing some of the classical textbooks in the field, there seems to be some disagreement about how literal this concern must be taken. In Silver et al. specific models are recommended for A, B and C items as can be seen in the way the chapters of this textbook are entitled. A discrete critique of this viewpoint can be found in Zipkin (2009) who on p. 119 writes:

Some writers offer much more specific advice, recommending particular models for specific groups (e.g. the EOQ model for B items). This is in my opinion too specific.

It could be tempting to conclude that A items should have the highest service levels followed by B items and C items. However, this may not be the case, as this paper also shows. A reason for this is partly because that the ABC classification scheme is somewhat too crude. Say, you compare two items in the A group. The first item is rather cheap but has a very high demand, while the other item is very expensive but with a very low demand. The latter item ought not to have a too high service level, as this would result an excessive inventory investment, while it makes good sense to let the first item (a so-called high runner) have a high service level. For a somewhat similar discussion concerning the comparison of two extreme A items, see also p. 320 in Silver et al. (2017). A way to resolve this is to use a two-dimensional classification, where the first dimension applies the ABC classification while the second dimension applies some other criterion, see for example Flores and Whybark (1987). In this paper, we will not pursue this idea (as well as other approaches presented in the overview paper by van Kampen et al. (2012)) but consider an optimization model as our starting point. Then we explore whether from the derived optimality conditions we can recognize any ranking according to the ABC classification. As we are considering a multi-item case, the optimization approach can either be bottom-up or top-down. A bottom-up approach requires an item-by-item solution. For applications of this approach, see Chakravarty (1981) and Stanford and Martin (2007). A disadvantage of this approach is that all cost parameters in principle need to be estimated which can be elusive according to Gardner (1980). An alternative approach, the top-down approach, is to avoid any specifications of the cost parameters but to minimize the aggregate (across items) inventory investment subject to a workload constraint and a constraint specifying that the weighted (across items) service level should be at or above a given target. For a general recommendation of this approach, see Sections 1.5 and 2.5 in Silver et al. (2017) as well as Lenard and Roy (1995). For specific applications of this approach, see for instance Alscher and Schneider (1982), Gardner and Dannenbring (1979), Gardner (1987), Zhang et al. (2001) and Zipkin (2000). In our paper, we also apply this top-down approach. Similar to Gardner (1987) and Zhang et al. (2001) we assume all items are individually controlled by (R,Q) inventory policies. But contrary to these two references we use a two-stage decision model. In the first stage, we assume deterministic demand and decide the order sizes. As everything is deterministic, we disregard any concern about the weighted service level, which is not invoked until the second stage. Because in this second stage the objective is to minimize the aggregate investment in safety stock subject to the condition that the weighted service level should be at or above a given target, where in our case, the relevant service measure is the well-known fill rate service measure. Thus, in this second stage we use the order sizes as given values determined from the first stage. The idea behind the second stage, to let the order sizes be pre-determined from the first stage, is very much aligned with standard procedures, see Axsäter (2006) p. 129, justified by the robustness of the Economic Order Quantity, *EOQ*, first to determine Q , by *EOQ*, and subsequently, R by a service level criterion.

To complete the literature review on interconnections between classification schemes and economic optimization, we would also like to mention Teunter et al. (2010), (2017). They do not directly apply a top-down approach as described above, but use well-known optimality conditions for the economic optimization of a single item problem to gain insights into ranking criterions depending on which service measure is most relevant. Also Millstein et al. (2014) and Yang et al. (2017) use optimization models for inventory classification. However, their orientation is slightly different from ours and besides they use mixed integer programming models.

The main objective of our paper is to explore whether the optimal order sizes (and thereby the order frequencies) and the optimal safety factors show characteristics that are aligned with the ABC classification. With regard to the order frequencies, our results show that A items should be ordered more frequently followed by B items and C items. With regard to the optimal safety factors, the results depend on the weights applied in the specification of the weighted service level. If all items have equal weight, then it turns out that A items should have the lowest safety factors followed by B and C items. This result might seem contrary to popular beliefs, but can be shown mathematically (in fact this result can also be seen in the numerical results of Zhang et al. (2001)). If the weights are the demand rates as in Zhang et al (2001), then the ranking of the safety factors follows the ratios of μ_i/c_i, that is, the higher this ratio is, the higher is the optimal safety factor. This result is in alignment with the example we gave in the beginning of this section, about the expensive and inexpensive item, both being classified as A items, where we claimed the expensive one should have a lower service level. This ranking is slightly different from the criterion in Zhang et al. (2001) who rank items after $\frac{\mu_i}{L_i c_i^2}$, where L_i is the lead-time of item i.

The structure of the paper is as follows. Section 2 contains the mathematical analysis of our inventory control model that has been sketched in this section. Finally, Section 3 contains concluding remarks.

2. Analysis

Suppose that we have N items, all controlled by independent (R,Q) policies. In an (R,Q) policy (Axsäter, 2006, p. 48), when the inventory position (on-hand minus backordered plus on-order) drops at or below the re-order point R, a batch quantity of size Q is ordered (in principle, however rather unlikely, multiple batches may be ordered in order to bring the inventory position above R). Our notation (where some of the symbols have already been introduced in the previous section) is as follows:

c_i: unit purchase price of item i

L_i: lead-time of item i

μ_i: average demand per time unit of item i

σ_i: the standard deviation of the demand per time unit of item i

Q_i: order size of item i

R_i: re-order point of item i

k_i: safety factor of item i

α_i: workload weight of item i

W: workload constraint

FR(R,Q): the fill rate (fraction of demand fulfilled immediately from stock) when using an *(R,Q)* policy

β_i: the weight put item *i*'s individual fill rate

FR_{Target}: a target for the weighted fill rate

We assume demands are stationary and normally distributed. For simplicity, we assume all replenishment lead-times are constant.

2.1 Stage 1

In the first stage, we decide the order sizes Q_i such that the value of the cycle stock is minimized subject to a workload constraint. The workload constraint is measured as the weighted sum of the number of replenishment orders issued per time unit. That is, we consider the following optimization problem:

$$Min \quad \sum_{i=1}^{N} c_i \frac{Q_i}{2} \tag{1.1}$$

s.t.

$$\sum_{i=1}^{N} \alpha_i \frac{\mu_i}{Q_i} \leq W \tag{1.2}$$

The workload weights α_i are pre-determined and the higher the weight is, the more work is involved when issuing and receiving an order.

The Lagrange multiplier of the constraint (1.2) is

$$\lambda^* = \left(\frac{1}{W}\sum_{i=1}^{N}\sqrt{\frac{\alpha_i c_i \mu_i}{2}}\right)^2 \tag{2}$$

and the optimal order sizes are

$$Q_i^* = \sqrt{\frac{2\alpha_i \mu_i \lambda^*}{c_i}} \qquad i=1,..,N \tag{3}$$

For the derivations of (2) and (3), see Appendix A (the derivations can also be seen in Alscher and Schneider (1982)). From (3) we get the optimal order interval is

$$T_i^* = \frac{Q_i^*}{\mu_i} = \sqrt{\frac{2\alpha_i \lambda^*}{c_i \mu_i}} \qquad i=1,..,N \tag{4}$$

Thereby, we get:

Theorem 1

If all items have the same workload weight, the A items will have the shortest order interval followed by B items and C items.

The result of Theorem 1 is aligned with the guideline seen in many textbooks on supply chain management for example in Simchi-Levi et. al (2008) at p. 56, in Bowersox and Close at p. 300 and at p. 47 in Wild (2002) . All items may not have the same workload weight, though. If items first are classified according to their workload weight, then the determination of the optimal order interval within each class of workload weight will be consistent with an ABC classification. Note that (3) is just a variant of the *EOQ* (Economic Order Quantity) formula where the Lagrange multiplier can be interpreted as the ratio between the order cost and the inventory carrying charge (and this ratio is then the same for all items). That is we could write $\lambda^* = A/r$ where A is the (common to all items) fixed order cost and r is the (common to all items) inventory caring charge, measured as the cost of having one dollar tied up in inventory per time unit. So under the assumption that these two cost components are the same for all items, we can derive Theorem 1 by the bottom-up approach we referred to in Section 1. This can also be derived from the results in Chakravarty (1981) and Stanford and Martin (2007). However, these two references do not all state Theorem 1 explicitly, because in both of these two papers, the focus is on grouping items such that in each group the same re-order interval is applied.

2.2 Stage 2

With the order sizes determined in Stage 1, the next stage concerns the minimization of the value of the safety stocks subject to a weighted service level constraint. Introducing safety factors k_i, $i = 1,..,N$ defined as

$$k_i = \frac{R_i - \mu_i L_i}{\sigma_i \sqrt{L_i}} \qquad i = 1,..,N \tag{5}$$

the value of safety stock is, see Axsäter (2006) and Silver et al. (2017)

$$\sum_{i=1}^{N} c_i k_i \sigma_i \sqrt{L_i} \tag{6}$$

We assume that the relevant service measure is the fill rate, and if order sizes are reasonably large (meaning covering the expected lead-time demand), the fill rate in an (R,Q) policy can be specified as (for the sake of clarity we omit the item index i here):

$$FR(R,Q) = 1 - \frac{\sigma\sqrt{L}G(k)}{Q} \tag{7}$$

The function *G(k)* is the unit loss function (Axsäter 2006, p. 91). Another popular service measure is the cycle service level measure *CSL*, denoted S_1 (Axsäter 2006) and P_1 (Silver et al. 2017), which is

$$CSL(R, Q) = \Phi(k) \tag{8}$$

Here $\Phi(k)$ is the cumulative distribution function of the standard normal distribution (mean *0* and standard deviation *1*).

Assume that we have a target for the weighted fill rate, FR_{target}, and the individual fill rates are weighted by positive numbers β_i with, $\sum_{i=1}^{N} \beta_i = 1$, then our second-stage optimization problem is:

$$\min \sum_{i=1}^{N} c_i k_i \sigma_i \sqrt{L_i} \tag{9.1}$$

s.t.

$$\sum_{i=1}^{N} \beta_i \left[1 - \frac{\sigma_i \sqrt{L_i} G(k_i)}{Q_i^*}\right] \geq FR_{Target} \tag{9.2}$$

The decision variables are the safety factors k_i, $i=1,..,N$. Let γ denote the Lagrange multiplier of the weighted fill rate constraint. Now the optimality conditions concern solving the following equation system, for derivations, see Appendix B.

$$\left\{\begin{array}{cc} \Phi(k_i) = 1 - \frac{c_i Q_i^*}{\gamma \beta_i} & i = 1, .., N \\ \sum_{i=1}^{N} \beta_i \left[1 - \frac{\sigma_i \sqrt{L_i} G(k_i)}{Q_i^*}\right] = FR_{Target} & \end{array}\right\} \tag{10}$$

Note the resemblance to the cost optimization criterion stated as eq. (2) in Teunter et al. (2010), because the Lagrange multiplier γ is in fact an imputed backorder cost (of the type cost per unit). Inserting (3) into the upper equation of (10) gives:

$$\Phi(k_i) = 1 - \frac{\sqrt{2\alpha_i \mu_i c_i \lambda^*}}{\gamma \beta_i} \qquad i=1,..,N \tag{11}$$

With regard to specifying the fill rate weights, β_i, $i = 1,\ldots,N$, there can of course be different choices. We consider the following three possibilities:

<u>The fill rate weights are chosen in alignment with the ABC classification</u>

$$\beta_i = \frac{\mu_i c_i}{\sum_{j=1}^{N} \mu_j c_j} \qquad i=1,..,N \tag{12}$$

<u>The fill rate weights are chosen equal</u>

$$\beta_i = \frac{1}{N} \qquad i=1,..,N \tag{13}$$

<u>The fill rate weights are chosen as the demand proportions</u>

$$\beta_i = \frac{\mu_i}{\sum_{j=1}^{N} \mu_j} \qquad i=1,..,N \tag{14}$$

We find all three choices could be reasonable. The first is of course in complete alignment with the ABC criterion. The second choice could be motivated by simplicity and equality. It is used in the analysis of Alscher and Schneider (1982), who assume all items are governed by periodic review *(s,S)* policies (where *s* is the re-order point and *S* is the order-up-to level). The third choice is applied by Zhang et al. (2001).

We can now deduce different results depending on the three choices with regard to the specifications of the fill rate weights β_i, $i = 1,..,N$.

Theorem 2

If all items have the same workload weight α and the fill rate weights are set as in (12), then solving (9.1)-(9.2) results in A items are having the highest safety factors followed by B items and C items.

Proof

When inserting (12) into (11), we get (after some reductions)

$$\Phi(k_i) = 1 - \frac{\sum_{j=1}^{N} \mu_j c_j}{\gamma} \sqrt{\frac{2\alpha\lambda^*}{\mu_i c_i}} \qquad i=1,..,N \tag{15}$$

Note that $\sum_{j=1}^{N} \mu_j c_j$ is a constant. Thus the higher $\mu_i c_i$ is, the higher is $\Phi(k_i)$ and thereby k_i.

■

This result is not so surprising and maybe more a sort a validation. If the fill rate weights are set such that the fill rates of the A items should have the higher weight then naturally the safety parameters of the A items should be high, and similar for the B and C items. More interesting are the following two theorems, Theorems 3 and 4.

Theorem 3

If all items have the same workload weight α and the fill rate weights are as in (13), then solving (9.1)-(9.2) results in A items are having the lowest safety factors followed by B items and C items.

Proof

When inserting (13) into (11) we get

$$\Phi(k_i) = 1 - \frac{N}{\gamma}\sqrt{2\alpha\mu_i c_i \lambda^*} \tag{16}$$

We can now see that the higher $\mu_i c_i$ is, the lower is $\Phi(k_i)$ and thereby k_i.

■

Thus if all items have same fill rate weight, it is the C items that have the highest safety factor followed by the B items and the A items in the bottom. This may be somewhat surprising. However, it comes straightforwardly from the mathematical analysis. The result of Theorem 3 can also be seen in the numerical experiments of Zhang et al. (2001), where they always assume all demand rates are equal, meaning $\beta_i = 1/N$ in these examples. Alscher and Schneider (1982), who use the same weights, do not offer any insights with regard to the ranking of the items, because their analysis is more concerned about reporting benefits of their algorithm. A final critical concern we could raise against the result of Theorem 3 is that it may be out of line with the general idea about an ABC classification. Because the general idea behind such a classification is to distinguish the "the trivial many" (the C (and B) items) from "the important fews", in order to better focus on the latter, and this concern tend to be blurred by the weighting scheme (13) which in effect let the "the trivial many" dominate "the important few".

Theorem 4

If all items have the same workload weight α and the fill rate weights are as in (14), then solving (9.1)-(9.2) results in that the items with the highest $\frac{\mu_i}{c_i}$ should have the highest safety factor.

Proof

When inserting (14) into (11), we get

$$\Phi(k_i) = 1 - \frac{\sum_{j=1}^{N}\mu_j}{\gamma}\sqrt{\frac{2\alpha\lambda^* c_i}{\mu_i}} \tag{15}$$

Note that $\sum_{j=1}^{N}\mu_j$ is a constant. Thus the higher $\frac{\mu_i}{c_i}$ is, the higher is $\Phi(k_i)$ and thereby k_i.

■

The result of Theorem 4, addresses the example stated in Section 1, that a high runner A item should have a high safety factor, while an expensive A item might not necessarily have a high safety factor. The fill rate weights in Theorem 4 are the same weights also assumed in Zhang et al. (2001). However, they rank items following the ratio of $\frac{\mu_i}{L_i c_i^2}$. A reason for this, is that instead of considering their weighted fill rate constraint, Zhang et al. (2001) choose to compute safety factors by instead (as a sort of simplification) considering a constraint of type CSL. A reason why Zhang et al. (2001) resort to this heuristic might be that their fill rate specification (see eq. (2.8) in their paper) is more general than that of this paper, (7). A possible critical remark to the analysis of our paper may concern that it sort of "cuts a corner" by using the more simple fill rate specification (7). It is the author's opinion, that using (7) is just as good as using the fill rate formula of Zhang et al. (2001). For a justification of that viewpoint, see Appendix C. It might also be strange that the lead-time (which does not enter into our analysis) enters into the denominator of the index of Zhang et al. (2001), as it might seem odd that a high lead-time should favor a small safety factor.

An obvious criticism to the choice of weights in (14), (and it seems to be ignored by Zhang et al. (2001)), concern that all demands are measured in units. However, in order for (14) to make sense these "units" must somehow be comparable. Say, if items are ranging from basic items like screws and nails to sophisticated items like washing machines, then (14) is less meaningful when demands are measured in units. So a sensible approach might be in that case to start, before any analysis is carried out, to transform all the values of μ_i into demands in dollars (or whatever monetary unit that is appropriate) and then set the values of c_i to *1*. Then (14) is identical to (12), and Theorem 4 is just a repetition of Theorem 2.

In finalizing this section, note that the results of Theorems 2 – 4 all are about which items should have the highest safety factors and thereby, from (8), about which items should have the highest CSL service level. Thus, the fill rate oriented analysis of Stage 2 turns out to be redirected towards the *CSL* service measure. This is also noted in Teunter et. al. (2010) who comment that it might have advantages with respect to implementation in an ERP (Enterprise Resource Planning) system like the well-known commercial system SAP.

3. Concluding remarks

In almost all textbooks on Inventory Control you can find a single chapter on ABC classifications. Often this is it, and the following mathematical analysis (most often focusing on considering just a single item) does not return to the ABC classification scheme. An exception to this is the textbook by Silver et al. (2017), who let specific chapters be devoted the study of A, B and C items. Also, in more popular books on Operations Management and Supply Chain Management there are viewpoints on how A, B and C items ought to be handled from an inventory control perspectives. However, these guidelines are not based on mathematical rigor, and as noted by Viswanathan and Bhatnagar (2005) that literature also contain conflicting viewpoints among those authors about how to set the service levels. A way to treat this in a rigorous way is to develop a mathematical multi-criteria optimization model. Instead of focusing on the algorithms, and the how to developing response surfaces to guide a decision maker to make the right trade-off, the focus in this paper is more on how we from the mathematical results can rediscover the ABC classification. In that respect, our paper is unique. Of course, there are limitations of our analysis, like we assume stationary demand and constant lead-times. Obviously relaxing some of these assumptions could open up for more interesting and possible challenging analyses. Furthermore, we only consider continuous review *(R,Q)* policies. We could consider other policies, like the periodic review *(s,S)* policy, to explore whether we will get almost the same results as presented in this paper.

References

Alscher, J. and Schneider, H. (1982). Resolving a multi-item inventory problem with unknown costs. *Engineering Costs and Production Economics*, 6 (1), 9-15.

Axsäter, S. (2006). *Inventory Control*. 2nd. ed., Springer, New York, NY, USA.

Bowersox, D.J. and Close, D.J. (1996). *Logistics Management*. McGraw Hill. Singapore.

Chakravarty, A.K. (1981). Multi-item inventory aggregation into groups. *Journal of the Operational Research Society*, 32, 19-26.

Flores, B.E., and Whybark, D.C. (1987). Implementing multiple criteria ABC analysis. *Journal of Operations Management*, 7 (1-2), 79-85.

Gardner Jr., E. S. (1980). Inventory theory and the gods of Olympus. *Interfaces*, 10 (4), 42-45.

Gardner Jr., E.S. (1987). A top-down approach to modeling US Navy inventories. *Interfaces*, 17 (4), 1-7.

Lenard, J.D. and Roy, B. (1995). Multi-item inventory control: a multicriteria view. *European Journal of Operational Research*, 87 (3), 685-692.

Gardner Jr., E.S. and Dannenbring, D.C. (1979). Using optimal policy surfaces to analyze aggregate inventory tradeoffs. *Management Science*, 25 (8), 709-720.

Millstein, M.A., Yang, L. and Li, H. (2014). Optimizing ABC inventory grouping decisions. *International Journal of Production Economics*, 148, 71-80.

Silver, E.A., Pyke, D.F. and Thomas, D.J. (2017). *Inventory and Production Management in Supply Chains*. 4'th. ed. CRC Press, Boca Raton, FL, USA.

Simchi-Levi, D., Kaminsky, P. and Simchi-Levi, E. (2008). *Designing and managing the supply chain: concepts, strategies and case studies*. 3rd ed., McGraw Hill, Boston, MA, USA.

Stanford, R.E. and Martin, W. (2007). Towards a normative model for inventory cost management in a generalized ABC classification system. *Journal of the Operational Research Society*, 58, 922-928.

Teunter, R.H., Babai, M.Z. and Syntetos, A.A. (2010). *Production and Operations Management*, 19 (3), 343-352.

Teunter, R.H., Syntetos, A.A. and Babai M.Z. (2017). Stock keeping unit fill rate specification. *European Journal of Operational Research*, 259, 917-925.

van Kampen, T.J., Akkerman, A. and van Donk, D.P. (2012). SKU classification. A literature review and conceptual framework. *International Journal of Operations and Production Management*, 32 (7), 850-876.

Viswanathan, S., Bhatnagar, R. (2005). The application of ABC analysis in production and logistics. *International Journal of Services Operations Management*, 1 (3), 257-267.

Wild, T. (2002). *Best practice in inventory management*. 2nd ed., Butterworth and Heinemann, Oxford, UK.

Yang, L., Li, H., Campbell, J.F. and Sweeney, D.C. (2017). Integrated multi-period dynamic classification and control. *International Journal of Production Economics*, 189, 86-96.

Zhang, R.Q., Hopp, W.J., and Supatgiat, G. (2001). Spreadsheet inventory control for a distribution center. *Journal of Heuristics* 7 (2), 185-203.

Zipkin, P.H. (2000). *Foundations of inventory management*. McGraw Hill, Boston, MA, USA.

Appendix A: Solving the minimization problem (1.1)-(1.2)

The Lagrangean function is:

$$L(\lambda; Q_1, \dots, Q_n) = \sum_{i=1}^{N} c_i \frac{Q_i}{2} + \lambda \left(\sum_{i=}^{N} \alpha_i \frac{\mu_i}{Q_i} - W \right) \qquad \text{(A1)}$$

The partial derivate wrt Q_i is:

$$\frac{\delta L}{\delta Q_i} = \frac{c_i}{2} - \frac{\lambda \alpha_i \mu_i}{(Q_i)^2} \qquad i = 1,..,N \qquad \text{(A2)}$$

Equating this to zero, and considering Q_i as a function of λ, gives:

$$Q_i(\lambda) = \sqrt{\frac{2\lambda \alpha_i \mu_i}{c_i}} \qquad i = 1,..,N \qquad \text{(A3)}$$

Because (1.2) is tight, we get from

$$\sum_{i=1}^{N} \alpha_i \frac{\mu_i}{Q_i(\lambda)} = W \qquad \text{(A4)}$$

that the optimal Lagrangean multiplier is as in (2).

Appendix B: Solving the minimization problem (9.1)-(9.2)

The Lagrangean function is:

$$L(\lambda; k_1, \ldots, k_n) = \sum_{i=1}^{N} c_i k_i \sigma_i \sqrt{L_i} - \gamma \left(\sum_{i=1}^{N} \beta_i \left[1 - \frac{\sigma_i \sqrt{L_i} G(k_i)}{Q_i^*}\right] - FR_{Target} \right) \qquad \text{(B1)}$$

Because $G'(k) = \Phi(k) - 1$, the partial derivative of the Lagrangean function wrt k_i is:

$$\frac{\delta L}{\delta k_i} = c_i \sigma_i \sqrt{L_i} + \frac{\gamma \beta_i \sigma_i \sqrt{L_i}}{Q_i^*} (\Phi(k_i) - 1) \qquad i = 1,..,N \qquad \text{(B2)}$$

Putting the partial derivative to zero, gives the upper part of (10). Because the constraint (7.2) must be tight, we have the lower part of (10).

Appendix C: A justification of the simple fill rate measure

A possible critique of this paper could be that the simple specification of the fill rate is used. That is, instead of (7) the specification if the fill rate should be (again here omitting item indices)

$$FR_{Advanced}(R, Q) = 1 - \frac{\sigma\sqrt{L}}{Q} \left[G\left(\frac{R - L\mu}{\sigma\sqrt{L}}\right) - G\left(\frac{R + Q - L\mu}{\sigma\sqrt{L}}\right) \right] \qquad \text{(C1)}$$

as done in Zhang et al (2001). Often it is claimed in the literature that is better to use (C1) than (7), because (7) is biased. It is the author's opinion that this concern often is too cautious and that the bias is in reality often opposite, if we take into consideration that the normal distribution is just an approximation of a demand process that makes aggregate demand distributions right-skewed. An example is the following. Assume customers arrive to an inventory system under a Poisson process with intensity *2* arrivals per time unit. Let any customer demand *2* units with probability *0.15*, *3* units with probability *0.2*, *4* units with probability *0.3*, *5* units with probability *0.2* and *6* units with probability *0.15*. Let the lead-time be *2* time units and let *(R,Q)* be *(35, 10)* (note, *Q* is in this numerical example set deliberately low in order to force through a difference between (5) and (C.1)). If you compute the fill rate by the exact formula (5.51) in Axsäter (2006) ("exact" means under the true assumption of a compound Poisson process), you get the fill rate is *0.9011*. Under the assumption of a compound Poisson process the lead-time demand has mean *24* and standard deviation *10.77*. By (5) you get the fill rate to be *0.9252* and by (C1) you get it to be *0.9330*. As the normal distribution is symmetric both formulas over-estimates the fil rate as they fail to recognize the right skewness. But the point is, that it is the simple formula (5) that is closest to the exact value. So in that sense (5) is more robust than (C1) with respect to non-normality of the underlying demand process. This phenomenon in addition to the fact that when *Q* is large there is no significant numerical difference between the two formulas, makes the author claim that using (5) in the analysis should not pose any major concern about the validity of the research in this paper.