Omega 40 (2012) 660–671

Contents lists available at SciVerse ScienceDirect

Omega

journal homepage: www.elsevier.com/locate/omega

Computing globally optimal (s,S,T) inventory policies

A.G. Lagodimos [a,*], I.T. Christou [b], K. Skouri [c]

[a] University of Piraeus, Department of Business Administration, GR-185 34 Piraeus, Greece
[b] Athens Information Technology, GR-190 02 Paiania, Greece
[c] University of Ioannina, Department of Mathematics, GR-45 110 Ioannina, Greece

ARTICLE INFO

Article history:
Received 6 April 2011
Accepted 16 December 2011
This manuscript was processed by Associate Editor Lodi
Available online 28 December 2011

Keywords:
Supply chain
Newsvendor
Optimal review period
Stochastic demand

ABSTRACT

We consider a single-echelon inventory installation under the (*s*,*S*,*T*) periodic review ordering policy. Demand is stationary random and, when unsatisfied, is backordered. Under a standard cost structure, we seek to minimize total average cost in all three policy variables; namely, the reorder level *s*, the order-up-to level *S* and the review interval *T*. Considering time to be continuous, we first model average total cost per unit time in terms of the decision variables. We then show that the problem can be decomposed into two simpler sub-problems; namely, the determination of locally optimal solutions in *s* and *S* (for any *T*) and the determination of the optimal *T*. We establish simple bounds and properties that allow solving both these sub-problems and propose a procedure that guarantees global optimum determination in all policy variables via finite search. Computational results reveal that the usual practice of not treating the review interval as a decision variable may carry severe cost penalties. Moreover, cost differences between (*s*,*S*,*T*) and other standard periodic review policies, including the simple base stock policy, are rather marginal (or even zero), when all policies are globally optimized. We provide a physical interpretation of this behavior and discuss its practical implications.

1. Introduction

Periodic review inventory policies, where ordering decisions are taken at regular time intervals, are very popular in industry (e.g. [23]). While theoretically inferior from their continuous review counterparts in terms of cost performance (e.g. [14]), their operational benefits seem to outweigh this deficiency for practical applications. The worldwide acceptance of MRP, a control system known to implement periodic review policies (e.g. [11,2]), provides a good example of such applications. Considering now the existing periodic review policies, there are three different standard types: the base stock policy, denoted as (*S*,*T*) or (*R*,*T*); the batch ordering policy, denoted as (*s*,*nQ*,*T*) or (*r*,*nQ*,*T*); and the (*s*,*S*,*T*) policy. Focusing on the latter policy, this paper studies the optimal policy variables determination and proposes a procedure that ensures globally optimal control.

To our knowledge, previous related research has almost exclusively studied the (*s*,*S*) policy, a reduced version of (*s*,*S*,*T*), having a fixed review interval *T* (not specified in absolute terms). For this restricted policy, by considering average total cost *per period* (review interval) as the objective function, the respective optimization problem is two-dimensional, with only the reorder level *s* and the order-up-to level *S* being decision variables. In fact, under this traditional formulation, the review interval *T* does not even appear in the objective function. In contrast, considering average total cost *per unit time* as the objective, in this paper all three (*s*,*S*,*T*) policy variables are simultaneously optimized, leading to solutions that may offer substantial cost savings over those of the reduced problem.

Research on the (*s*,*S*,*T*) policy, mostly on its restricted (*s*,*S*) variant originally introduced and modeled by Arrow et al. [1], is extensive. This reflects the existence of early results establishing the (*s*,*S*) policy optimality over any other periodic review policy for finite horizon [19] and infinite horizon [8] problems under general conditions (see [29] for early work on relaxing these conditions). Due to the optimality of its form, the (*s*,*S*) policy often serves as the benchmark for assessing other inventory policies. One policy often considered is (*s*,*nQ*,*T*), in its restricted (*s*,*nQ*) variant. Implicitly assuming a common review interval for both policies, comparative studies concluded that *optimal cost* (per period) difference between (*s*,*S*) and (*s*,*nQ*) is not large (e.g. [15,31,33]). Observable differences, however, do exist. For example, Zheng and Chen [33] found the optimal (*s*,*S*) to outperform (*s*,*nQ*) by circa 6% in several cases and even to exceed 10% in one case.

Focussing on the (*s*,*S*) policy, several exact and heuristic algorithms for solving the respective two-dimensional optimization

* Corresponding author. Tel.: +30 210 4142197.
E-mail addresses: alagod@unipi.gr (A.G. Lagodimos), ichr@ait.edu.gr (I.T. Christou), kskouri@uoi.gr (K. Skouri).

0305-0483/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.omega.2011.12.008

problem have been proposed. This research uses average total cost per period as the objective function and apply the end-of-period costing scheme (i.e. cost evaluation is based on end-of-period inventory information) introduced by Arrow et al. [1] and traditionally underlying the analysis of periodic review policies. In this context, early exact algorithms are essentially grid search procedures, using bounds to confine the state-space. Examples include: Veinott and Wagner [30], Johnson [10], Stidham [24], Federgruen and Zpkin [5]. In contrast, Zheng and Federgruen [34] proposed a directed enumeration scheme resulting in increased computation speed (later modified by Feng and Xiao [6]). With this algorithm, search is embedded in an iterative procedure (based on cost properties including a Newsvendor relation) gradually converging to the optimum. Finally, of the numerous heuristic solution approaches that exist, the following are just a few representative examples: Naddor [15], Schneider [21], Ehrhardt [4], Tijms and Groenevelt [26], Porteus [16].

The above research paradigm has ignored the review interval T as a policy variable. In their seminal work, Hadley and Whitin [7] first presented and discussed all standard periodic policies, including (s,S,T) in its general form, treating T as a decision variable. In order to allow optimization in T, they assumed continuous time and considered average total cost per unit time as the objective function. Moreover, they introduced a new cost element (i.e. the review cost) together with a continuous costing scheme (i.e. cost evaluation is based on average inventory information), identical to that used in the analysis of continuous review and deterministic inventory policies. The main advantage of this scheme over end-of-period costing is that it provides more representative cost figures, whose accuracy does not depend on the review interval T (see also [18] for a recent discussion of these costing schemes). Under this new framework, Hadley and Whitin derived models for all three standard periodic review policies, assuming Normal and Poisson demand processes. Noticeably, for each specific demand process, most expressions in the respective models are presented in a final closed-form (having very complex forms). While invaluable for computations, a major weakness of this approach is that it effectively hinders the identification of either structural similarities existing between different periodic review policies or functional properties of the respective cost models necessary for formal optimization. This lead naturally to the conclusion that (s,S,T) (as well as all other standard periodic review policies) can only be optimized via exhaustive search. Such a procedure, however, is both computationally demanding and cannot guarantee global optimum determination (effectively constituting a heuristic approach).

It is fairly recently that, building on Hadley and Whitin [7], exact algorithms to globally optimize periodic review policies have appeared. A common feature of this research is that, instead of analyzing respective cost models in their final form, analysis is initially based on general distribution-free models. Subsequently, by focussing on specific demand distribution classes, important cost function properties are established. For the base stock (S,T) policy, Rao [17] considers T as a decision variable. Assuming continuous time and a class of demand processes which is stochastically increasing and linear in time, he applied stochastic convexity analysis to show that the (S,T) policy cost per unit time is jointly convex in both policy variables (analogous to the result for the continuous review (r,Q) policy in [32]). So, also invoking a Newsvendor property, the globally optimal base stock policy is easily determined. Under identical assumptions, Lagodimos et al. [12] proposed an algorithm to globally optimize the (s,nQ,T) policy. The algorithm makes use of new cost per unit time bounds and properties. Specifically, for any T, the bounds are jointly convex in s and Q. Since they are also increasing in Q and satisfy a Newsvendor relation, locally optimal solutions (for any T) prevail via a finite search on Q. Since such local optima are also bounded by optimal (S,T) policies, jointly convex in both variables [17], a finite search on T guarantees global optimum. Lagodimos et al. [13] later showed the (s,nQ,T) holding and backorders cost per unit time to be jointly convex in all decision variables and used this property to adapt the above procedure for the policy optimization under exogenous supply-lot constraints. Note that analogous joint convexity results were also independently proposed by Shang and Zhou [20] for multi-echelon serial systems with echelon (s,nQ,T) policies and used to ensure globally optimal control (under the simplification that ordering costs are fixed).

Under this general background, this paper establishes an exact procedure that guarantees the globally optimal (s,S,T) policy determination in finite steps. Given the known optimality of this policy over all periodic review policies, this is the main contribution of this research. In our opinion, the following important parallel contributions are also worth stating: (i) It establishes formal structural relationships between all standard periodic review policies. Other than allowing interchangeability of results between these policies it also constitutes a decisive step towards their totally unified analysis and optimization. (ii) It demonstrates how results for the (s,S) policy under traditional end-of-period costing can be used in optimizing (s,S,T) under continuous costing, a finding directly applicable to all periodic policies. (iii) It provides firm evidence that all standard periodic review policies, when globally optimized, have marginal cost differences, thus encouraging the use of the simple base stock policy in practical applications.

The remainder of this paper is organized as follows. Section 2 gives the notation and assumptions used. Section 3 presents distribution-free models for the (s,S,T) policy average total cost and service under stationary demand processes and provide comparisons that form the basis of the analysis. Section 4 presents simple cost bounds (upper and lower) and a cost-equivalence property, associating (s,S,T) with the simple base stock policy. Section 5 develops optimality properties, leading to a search procedure that ensures global optimal (s,S,T) control. Issues on the procedure implementation are also presented. Section 6, based on examples presented in other sections, provides a physical interpretation of the (s,S,T) optimal policy behavior and of its affinity with optimal base stock. Section 7 presents computations designed to investigate the importance of considering the review interval as a decision variable and its effects on the relative performance of all three standard periodic review policies. Finally, Section 8 discusses the implications of the findings and presents directions for future research.

2. Preliminaries

In this section we present the general notation used along with the assumptions underlying the operation of the inventory system considered in this paper.

2.1. Notation

While some additional notation is introduced later, the following is a list with the notation used:

$D(t)$	cumulative demand in time interval $[0, t]$
λ	demand rate (mean demand per unit time)
L	replenishment lead time
T	length of review interval
s	reorder level
S	order-up-to level
Δ	the difference $\Delta = S - s$