



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/84377/>

---

**Monograph:**

Riddalls, C.E. and Bennett, S. (2000) The Optimal Control of Inventory Systems. Research Report. ACSE Research Report 762 . Department of Automatic Control and Systems Engineering

---

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.

# THE OPTIMAL CONTROL OF INVENTORY SYSTEMS\*

C.E. Riddalls and S. Bennett

Department of Automatic Control and Systems Engineering, University of Sheffield,

S1 3JD. Tel 0114 222 5186, fax 0114 273 1729, [c.e.riddalls@sheffield.ac.uk](mailto:c.e.riddalls@sheffield.ac.uk)

RESEARCH REPORT NO. 762

*ABSTRACT: The science of Operational Research has traditionally dominated in the field of inventory theory. Yet there have been sporadic attempts to apply methods from control theory in this domain. We review these efforts, concentrating particularly on optimal control approaches. We also enumerate the generic benefits and limitations of operational research methods. Through these efforts we judge whether the supremacy of operational research techniques in the field of inventory systems is warranted. We apply a novel optimal control algorithm to a differential equation model of an inventory system. This enables us to mimic the cost structures implied by quantity discounts and approximate capacity constraints. Some examples illustrate how optimal responses to instantaneous jumps in demand are generated and how these are affected by quantity discounts.*

## 1. Introduction

Inventory theory is concerned with the decisions which affect the flow of goods through the business enterprises which connect to satisfy a consumer demand. At each enterprise these decisions are based on answers to the questions: How much should be



200452129



\*This work was supported by EPSRC under the grant number GR/M27432

ordered (or produced)? When should the orders be placed so that total inventory costs are minimised?

The economies of scale that accrue to the manufacturer or purchaser by producing or ordering in large batches can be significant. The inventory manager, however, favours policies which meet a (typically) varying and uncertain demand with the minimum of tied-up inventory. The reconciliation of these conflicting objectives is at the heart of inventory theory.

The science of Operational Research has traditionally dominated in the field of inventory theory, although there have been sporadic attempts to apply methods from control theory in this domain. In the next section we review these two distinct approaches with a view to their appraisal. Our aim is not expositional, rather to enumerate the generic merits and limitations of each approach. Through these efforts we judge whether the supremacy of operational research techniques in the field of inventory systems is warranted. To this end we cite the influential work which illustrates our theses. We conclude that optimal control methods constitute a viable alternative to the traditional OR tools for some applications. However, their somewhat esoteric nature remains an impediment to their wider adoption.

In this area optimal control methods are often criticised for their inability to model realistic cost structures and capacity constraints. In section three we introduce a relatively new and novel optimal control algorithm which will enable us to redress these deficiencies.

In section four we present cost structures arising from quantity discounts. The new-found flexibility afforded by the algorithm enables the treatment of such intrinsically discontinuous cost structures. The punitive nature of optimal control can also be exploited more fully. Therefore, by adjusting the cost structure we may avoid

transgressions beyond notional capacity constraints. By applying our algorithm to a differential-equation model of an inventory system we generate optimal responses to instantaneous changes in demand. Used in conjunction with an existing OR policy, our algorithm also furnishes the inventory manager with a new additional capability of evaluating 'what if' scenarios. A number of examples are given to illustrate the theory.

## **2. Inventory Systems**

The economies of scale achievable through relatively high levels of batched production often call for the availability of a significant and expensive storage capacity. Quantifying and reconciling this implicit trade-off constitutes a large part of production-inventory theory. Various factors serve to compound the difficulty of this calculation. Greater competition in some markets disposes a more exacting consumer intolerant to stock-outs. This shifts the balance of the trade-off in favour of larger safety stocks. Militating against this inclination are the capricious nature of consumer tastes and concomitant compression of product life-cycles which increase the risk of obsolescence to safety stocks. Further, a more discerning public, demanding a greater diversity of product ranges, discourages production in large batches and increases the cost of holding safety stocks across a product range.

The advent of computer power in the 1970s enabled companies to analyse and streamline their production processes using Materials Requirements Planning (MRP) and Manufacturing Resource Planning (MRP II). Then, in the 1980's, the emergence of intuitively attractive management philosophies like Just-In-Time (JIT) (see Cook and Roowski 1996, Ehrhardt 1998, Kim and Takeda 1996) from Japan, a country

admired for its manufacturing industry, prompted renewed scrutiny by western companies of their inventory systems. Today many companies use analytic tools to provide both structural and operation solutions for inventory systems. Most of these tools come from the discipline of operational research.

## **2.1 Operational Research Approaches**

OR theory comprises a disparate collection of mathematical techniques, such as linear programming, queuing theory, Markov chains and dynamic programming. The common theme running through all these tools is their suitability for the solution of man-made problems. Indeed, the construction of most of these approaches was motivated and consequently circumscribed by the particular application. This stands in contrast to calculus based approaches, whose foundations were conceived independently of the eventual application. We shall see that these facts have ramifications for the improvement and sophistication of OR models.

Inventory systems are a relatively mature topic in the discipline of operational research. The theory is devoted to the questions: How much should be ordered (or produced)? When should the orders be placed so that total inventory costs are minimised? In this paper we do not examine tools which are based on specific demand forecasts since mathematically that is a quite distinct problem. The most rudimentary models determine an optimal reorder quantity, the Economic Order Quantity (EOQ), given certain simple demand characteristics and cost structures. These latter are usually linear or linear-affine in nature. For example, a notional penalty cost is attached to stockouts and is usually proportional to their number. The cost of resupply (or production) may comprise fixed set-up and proportional

components. More advanced models incorporate a stochastic element in the demand or lead-time characteristics. These models fall into two categories: periodic and continuous review models, in which, respectively, inventories are monitored periodically or whenever a demand occurs. Both models can be constructed over a finite or infinite time horizon.

Most noteworthy are the  $(Q,r)$  [Gallagher et al., 1959] and  $(s,S)$  policies [Gallego, 1998]. In the former, (a continuous review model) whenever the inventory level drops to  $r$ , a fixed order of size  $Q$  is made. In the latter, which can operate through both continuous and periodic review, an order of size  $S-s$  is made when the inventory drops to  $s$  or below. Obviously,  $(Q,r)$  policies agree with  $(s,S)$  policies when there is no possibility of overshooting the re-order point  $r$ . For production systems,  $(Q,r)$  policies have obvious logistical advantages over  $(s,S)$  policies.

The standard theory can be generalised to the case of multiple products with capacity constraints. There are many introductory books presenting this theory, for instance [Naddor 1966].

We now turn to our critical appraisal of the OR approaches to inventory systems. Although not our intention to wantonly denigrate OR methods, by highlighting their failings we aim to justify the investigation of other methodologies in this area. Through these efforts we may also judge whether the primacy of OR approaches in industry belies their efficacy.

OR methods derive paradigmatic strategies for inventory systems. Their calculation is off-line (typically in determining  $Q$  and  $r$ ) and the results determine a policy which is implemented on-line. By this we mean that they provide simple rules for an inventory manager to follow given certain conditions. The most widely adopted inventory rule is the simple EOQ model. For the purposes of the rule derivation, the demand process

is treated as a constant stream. This engenders simple closed form solutions for optimal  $Q$  and  $r$  and the total costs incurred, enabling a rigorous sensitivity analysis. Indeed, it is well documented that the total cost incurred per unit time, as measured by the model, is fairly insensitive to variations in  $Q$  about its optimum [Dobson, 1988]. However, since this sensitivity is calibrated by the same (perhaps inaccurate) model and, in any case, the specific calculation of the optimal  $Q$  is not problematic, we would regard with dubiety any appeal to these facts in advocacy of the EOQ approach. More useful is the sensitivity analysis of model parameters carried out by [Dobson, 1988]. Again, the cost is found to be relatively insensitive to spurious parameter estimates. This is gratifying since the cost parameters (especially the holding costs) can be somewhat notional and subjective. Again, however, these results are generated by the same grossly simplified model and so should be interpreted with circumspection.

The genesis of the OR approach was motivated and influenced by expediency and pragmatism. These imperatives neglected the development of an independently viable theoretical framework. The ad hoc nature of this approach, although initially facilitating its swift adoption, has hampered efforts towards its greater sophistication. For example, the static nature of the basic EOQ models (characterised by the absence of any driving dynamics) has required the encapsulation of the demand characteristics in a simple, tractable form (usually constant demands). The resulting solution is consequently only valid for that qualitative type of demand. More importantly, the extent to which varying demands impair the optimality of the solution is unclear. In an attempt to counteract this deficiency, researchers considered stochastic demand parameters within  $(Q,r)$  and  $(s,S)$  policies. The requisite estimation of the distribution parameters by the inventory manager constitutes a considerable effort and potential

source of error . Further, the inclusion of a stochastic component in demand might still not sufficiently capture its full quiddity. Another consideration is that the demand characteristics might be non-stationary over the time horizon of the model, thus vitiating the solution. Some authors have considered non-stationary demand distributions [Veinott, 1965]. However, their generation presupposes an implausible prescience on the part of the inventory manager since the accurate estimation of future demand characteristics is an even more unlikely prospect than that of the present demand distribution.

The complications resulting from stochastic components render closed form solutions elusive. The determination of optimal pairs has typically been accomplished by the 'back box' methods of numerical analysis. This, and the absence of an intrinsic theoretical framework implied that for many years there existed no capability for a rigorous analytical sensitivity analysis of these models. These shortcomings meant that no great insight into the underlying process was afforded by these models. Hence the change in solution engendered, for example, by small variations in parameters could only be ascertained by onerous calculation or heuristic methods [Archibald and Silver 1978]. In this paper it was found that when demand is erratic, the solution is sensitive to the demand distribution.

Recently there have been efforts to plug this analytical deficit. Zheng [1992] was one of the first authors to address these issues. He quantified the differences between the stochastic  $(Q,r)$  policy and its deterministic first order approximation, the EOQ model. He concluded that, as a function of  $Q$ , the stochastic cost curve is even flatter than the deterministic. Further, the relative increase of the costs incurred by using the EOQ  $Q$  instead of the stochastic one is no more than  $1/8$ . Gallego [1998] extends these results to clarify the cost impact of demand variability.

Inventory systems rarely exist in isolation. Many production processes comprise parallel flows of goods controlled by a network of inventory systems. For these reasons, researchers quickly apprehended the potential utility of multi-echelon approaches. A seminal paper in this area [Clarke and Scarf, 1960] developed a procedure which obviated the recursive computation of  $N$  individual optimal policies in a multi-echelon periodic review  $(s,S)$  system. However, this natural extension to the standard theory serves to obfuscate any insight afforded by the single echelon sensitivity analysis, when extrapolated to these systems.

The same can be said for capacitated systems. In reality both inventory and production (or even purchasing) capacities are finite. Federgruen and Zipkin [1986a,b] show that a base-stock policy is optimal in a single stage production-capacitated system. But for multi-echelon or inventory-capacitated systems such definitive analytical results remain elusive. As in other fields, such analytical shortcomings can be partially circumvented by simulation methods. Glasserman and Tayur [1995] employ simulation-based techniques to tune the parameters in a periodic review multi-echelon  $(s,S)$  system in order to improve the performance of an analytically determined policy. The sub-optimality of these policies is offset by the ersatz dynamic nature of the simulation procedure which affords some understanding of the governing elements in the system.

Using similar methods the same authors conduct a stability analysis of the same system [Glasserman and Tayur, 1994]. By stability one simply means, given a particular policy, does the systems meet demands, or does it become increasingly backlogged? The system is stable under the obvious condition that the mean demand per-period be smaller than the per-period production capacity at every echelon. The

authors also develop more useful results on how often the inventories return to their target levels.

Typically couched in the language of purchasing systems, the emphasis of most inventory theory is on the demand side. However, these methods can provide limited production solutions with careful interpretation. For example, a well known refinement to the standard model consists of incorporating quantity discounts into the cost structure on the purchasing (supply) side [Gallagher et al., 1959]. This can equally be construed as representing the cost of batched production. Notwithstanding this flexibility, these tools are mainly used by inventory managers who order goods from suppliers rather than produce themselves. This is because the complexities of many production operations demand the development of dedicated scheduling and job sequencing tools.

To this end, aggregate production planning techniques have been developed [Elsayad & Boucher 1985] in order to generate a master schedule over a planning horizon. Typically formulated as optimisation problems, they can be solved using mathematical programming techniques. These methods are again off-line and require demand projections over the planning horizon. The specific timing of the flows in the production system generated by these production planning techniques can be handled by MRP systems [Elsayad & Boucher 1985].

## **2.2 Control-Theoretic Approaches**

Modelling inventory systems using differential or difference equations holds great appeal for the control theorist. This is because many of the influential characteristics of the problem can be succinctly expressed in equation form. Then a vast array of

tools and methodologies from control theory can be invoked to gain insight into the system dynamics. The rationale for this approach is that models of modest complexity, which are therefore amenable to analytical study, can provide an insight into the factors which are common to much larger 'live' systems. Unfortunately, such analytical utility has hitherto not spawned tools which challenge the sovereignty of OR methods in real applications.

Since differential equations produce 'smooth' outputs, they are not universally suited to the modelling of all production-inventory system. The system must be considered at an aggregate level, in which individual entities in the system (products) are not considered. Rather, they are aggregated into levels and flow rates, which vary smoothly with time. So these methods are unsuited to production processes in which each individual entity has an impact on the fundamental state of the system. For the same reasons differential equation approaches cannot solve lot sizing and job sequencing problems. The same limitations apply to discrete time difference equation models which can be more computationally burdensome to implement. However, since the dynamics of production-inventory systems are fundamentally discrete, they possess an innate advantage over continuous time models. They enable the invocation of classical control theory within a discrete framework, which facilitates the inclusion of pure time delays and obviates the (sometimes) artificial smoothing of the differential equation approach. The choice of differential over difference equation methods should therefore amount to the subordination of pure time delays to the possible greater complexity of these latter methods.

We start with the simplest linear deterministic systems. Simon was one of the first to use classical Laplace transform techniques to analyse simple continuous-time production-inventory systems [Simon 1952]. In essence, he solved the system

equations analytically and investigated how the solution characteristics are determined by the system parameters. In the cited paper he highlighted a fact which persists in influencing research today, namely that pure delays are hard to deal with within this framework. Instead he used exponential delays, which approximate pure delays by smoothing the output signal over time. The extent to which this substitution compromises the true reproduction of the system characteristics must be judged for each individual system.

Porter and Taylor dealt with the set of equations:

$$\begin{aligned}\frac{di}{dt} &= p_a(t) - d(t) \\ \frac{dp_a}{dt} &= \alpha(p_d(t) - p_a(t))\end{aligned}\tag{2.2.1}$$

where  $p_d(t)$  is the desired production rate,  $p_a(t)$  is the actual production rate,  $d(t)$  is the demand and  $i(t)$  is the inventory level [Porter and Taylor 1972]. By specifying a constant desired inventory level, these equations can be rearranged and the control inputs chosen to be the desired production rate and its derivative. Simple state feedback techniques are then developed to stabilise the system. Bradshaw and Porter use similar modal control techniques in a slightly more complex environment in which advertising affects demand [Bradshaw and Porter 1975]. By discretizing the continuous time systems found in [Porter and Taylor 1972], Porter and Bradshaw [Porter and Bradshaw 1974] again use simple state feedback design to generate piecewise constant controllers for inventory systems.

A similar approach is adopted in [Vassian, 1955] who applied discrete variable servomechanism theory to a periodic review system. The aim of this approach is to

minimise fluctuations about a desired safety stock level through periodic review. The most significant difference to an archetypal OR approach is that the reorder quantities are calculated on-line and the actual safety stock level is not determined by the policy. However, out of all the control-theoretic approaches, these are the most redolent of OR tools we have examined.

An important limitation of these approaches is their failure to incorporate the capacity constraints that invariably exist in the real world. In [Bradshaw and Erol 1980] they address this issue by using bounded control. However, more importantly, this framework fails to take into account the cost implications of such control strategies. A more natural setting in which to tackle this problem and attach costs to different policies is provided by optimal control.

Using standard optimal control theory over finite time intervals the various costs of the production strategy can be accounted for and traded-off against one another. The optimal control of such systems is tackled in the papers [Bensoussan and Proth 1982], [Abad 1985], [Lieber 1973] and the books [Bensoussan et al 1974] and [Arrow et al 1958]. A common attribute of this work is that the faithful replication of real-world cost structures has been sacrificed in order to achieve tractable solutions. For example, [Bensoussan et al 1974], [Tzafestas and Kapsiotis, 1994] and [Abad 1985] all assume quadratic production and inventory holding costs, enabling them to invoke the relatively mature subject of linear-quadratic optimal control. Using these methods over a limited operating range does constitute a plausible approximation to linear cost structures. Further, away from the equilibrium, the accelerating cost penalties can legitimately be used to represent capacity constraints. However, large fluctuations in demand may require calculations over more extended cost ranges whose differing qualitative nature cannot so easily be defended. [Arrow et al. 1958] use linear cost

structures to investigate smoothed production plans when demand is known over a fixed planning horizon. The simplicity of linear costs enables significant analytical progress in this problem.

The paucity of papers treating more sophisticated cost structures bears witness to the difficulty of the problem. However, two papers [Bensoussan and Proth 1982], [Lieber 1973] have considered finite-time interval problems with simple convexity requirements on the costs. These may not always be fulfilled for particular processes and can be difficult to check, but seem to be the extent to which general conditions can be relaxed. Lieber uses Pontryagin's Maximum Principle to derive some planning horizon results. He requires the production cost function to be twice continuously differentiable, which may not be the case in reality as these costs tend to jump up with the number of production runs, not the number of goods produced. However, continuous approximations can be made to arbitrarily closely mimic these functions. A particular innovation by [Bensoussan and Proth 1982] is to take into account capacity constraints. They use backwards dynamic programming to calculate optimum production schedules. This method can be computationally burdensome and is only as good as the chosen algorithm. Neither of these papers gives examples which replicate the costs implied by batched production or, for purchasing systems, quantity discounts. A review of some of the literature mentioned here can be found in [Axsäter 1985].

In reality demand is stochastic in nature. Yet the added complexity of stochastic systems may obscure our view of the essential dynamics and hamper the derivation of simple control strategies.

Schneeweiss uses the Weiner-Hopf technique to find a simple control strategy in the presence of stochastic demand [Schneeweiss 1971]. In [Bensoussan et al 1975]

Bensoussan et al use partial differential equations to model the inventory levels of perishable products. In his book [Bensoussan et al 1974] he studies the application of linear quadratic cost stochastic control theory (again ignoring time delays). It could be observed that since the linear optimal control problem incorporating time delays has not been adequately tackled in this field, it is rather premature to jump to stochastic systems.

### **2.3 Conclusions**

Having briefly examined two sets of tools for production-inventory systems, we now weigh up their relative merits and deficiencies. Mindful of the pre-eminence of OR techniques in industry, we judge whether optimal control methods possess the potential to compliment or replace them in any niche of application. We recognise however, that no single technique is likely to prove a panacea in this field.

OR tools rightfully dominate in the design of production-inventory systems at a strategic level. Their greatest strength is in resolving the many trade-offs inherent in the design of such systems. Indeed the standard theory presented above has been amended to tackle more subtle trade-offs. For example, in a series of papers [Liao and Shyu 1991a, Liao and Shyu 1991b, Ben-daya and Raouf 1994] the authors quantify certain trade-offs in inventory systems. By treating lead time as a decision variable, they balance the increased costs of reducing lead times (called crashing costs) against the concomitant decrease in inventory costs, whilst maintaining customer service levels.

The simplicity of OR models, whilst contributing to their currency, simultaneously determines their limitations. Their failure to elucidate the dynamics of the process is

telling. Only through knowledge of the dynamics of a system can we gain a full appreciation and understanding of the factors which affect its performance. The one-off nature of the calculations involved means that valuable dynamic real-time information is ignored in the manufacture of these solutions. To compensate for this dynamic deficit certain possibly factitious demand distributions must be postulated and their parameters estimated. The consequent potential for error is compounded by the non-stationarity of the demand characteristics, itself rendered imponderable by an increasingly capricious and demanding consumer.

Optimal control approaches do not assume any particular demand pattern, rather they react to each demand in real time. Thus their responses to a particular demand can be viewed as more appropriate than the prescriptive OR rule. The absence of a dynamic capability also restricts the applicability of OR tools to systems in steady state. In contrast, optimal control approaches can deal with transient periods and 'one off' shocks which might not be embodied in the demand probability density function of an OR system. They also provide the capability to evaluate 'what if' scenarios, thus furnishing the inventory manager with a greater understanding of the system. However, this transient superiority is not matched by a structural one. Optimal control approaches usually work with a safety stock whose level they do not calculate.

The issue of sensitivity analysis has been tackled for the most simple systems by OR methods. However, for the more complicated multi-echelon and capacitated systems these questions remain unanswered. Indeed, proven optimal policies have not been derived for some of these more complicated systems.

Indeed, generalisations of the most simple deterministic EOQ models to include multiple products and capacity constraints requires the solution of a potentially

difficult optimisation problem. This again militates against an appreciable understanding of the system characteristics.

Notwithstanding these criticisms, the supremacy of OR tools for the reconciliation of structural trade-off problems remains unimpeachable. Further, their utility in the solution of batch sizing and job sequencing problems is unassailable [Chandra 1993].

Optimal control methods must therefore challenge in different areas of application if they are to be taken up by any practitioners. Traditionally these have been in the field of production planning. Given a demand forecast, they can determine the optimal aggregate production schedule. However, the intricacies of modern production systems demand the observance of constraints and cost structures which cannot be incorporated into differential equation systems, due to the rigorous assumptions imposed by their classical foundations. For instance, the inability to deal satisfactorily with inherently discrete and discontinuous quantities such as batches and hard capacity constraints, has traditionally rendered these methods somewhat impotent.

However, for purchasing systems, we contend that the merits of optimal control techniques do warrant further investigation into possible new rôles for them. The remainder of this paper is devoted to the appointment of such a rôle which compliments existing OR techniques. Recent advances in optimal control theory enable us to make tangible progress in the hitherto intractable (for optimal control methods) treatment of the quantity discounts which arise from batched production. We introduce an algorithm which also enables the approximation of inventory capacity constraints.

### **3. Optimal Control**

The algorithm used in this paper is based on work first published in [Banks and Mhana 1992]. They control systems of the form

$$\dot{x} = A(x)x + B(x)u, \quad (3.1)$$

using the infinite-time cost functional

$$J = \int_0^{\infty} (x^T Qx + u^T Ru) dt. \quad (3.2)$$

Here  $x$  and  $u$  are the  $n$ - and  $m$ -dimensional state and control vectors, respectively.  $A, B, Q$  and  $R$  are of the appropriate dimension and  $A$  and  $B$  are analytic. They show that under some bounded growth conditions on the derivatives of  $A$  and  $B$ , their algorithm produces a stabilizing control. The Banks and Mhana approach takes its inspiration from linear quadratic control [Ogata 1990], for which  $A$  and  $B$  are constant matrices in (3.1). Given an initial condition  $x_0$ , the algorithm consists of freezing the state at any given point,  $\bar{x}$ , on the optimal trajectory and solving the corresponding linear problem:

$$\dot{x} = A(\bar{x})x + B(\bar{x})u, \quad x(0) = x_0 \quad (3.3)$$

$$J = \int_0^{\infty} (x^T Qx + u^T Ru) dt \quad (3.4)$$

Finding a positive-definite symmetric  $P$  which solves the algebraic Riccati equation:

$$P(\bar{x})A(\bar{x}) + A^T(\bar{x})P(\bar{x}) - P(\bar{x})B(\bar{x})R^{-1}B^T(\bar{x})P(\bar{x}) + Q = 0 \quad , \quad (3.5)$$

at every point on an optimal trajectory constitutes an optimal stabilizing control. The condition for existence of such a solution is that the frozen linear system  $(A(\bar{x}), B(\bar{x}), Q)$  be stabilizable and observable for all  $\bar{x}$  on the stable trajectory. This condition is difficult to prove *a priori*, but, for polynomial functions, its failure will only be ephemeral. The utility of this algorithm stems from its simplicity and the resulting ease of application. At any  $\bar{x}$  on a trajectory one is only required to solve the algebraic Riccati equation.

We must point out that after Banks and Mhana published their work there appeared a paper [Gong and Thompson 1995] questioning the reasoning behind one of the proofs in the multivariable case. These matters were resolved in [McCaffrey and Banks 1998], where the algorithm was shown to be asymptotically optimal (that is, converges to the optimal control as  $x \rightarrow 0$ ) in this case.

Both these last two references extend the problem to consider analytic state dependent cost matrices ( $Q$  and  $R$  in (3.4)). It is this innovation that allows us to model quantity discounts in the cost structure and capacity constraints. In reality the resulting cost structures are discontinuous but, to comply with the requirements of the algorithm, we have smoothed them using trigonometric functions about points of discontinuity. In this way it is possible to arbitrarily closely approximate the discontinuous cost structures. In the paper [Harrison and Banks 1998] the authors use this algorithm to control a vehicle suspension system. They use discontinuous cost matrices in order to better approximate the engineering objectives of the system. The approximation of hard constraints by the appropriate modification of the cost functional is discussed in more detail in [Russell 1965].

We shall see in the next section that our production-inventory system model has an external forcing function (the demand) as well as a feedback input. It is well known that, for constant inputs such systems reduce to the familiar regulator problem (with no external input). For the purposes of this study we shall only consider constant external inputs by assuming that fluctuations in demand are initially satisfied by a fall in the level of safety stock.

#### 4. Implementation

Consider an inventory system modelled by the following set of equations:

$$\frac{di}{dt} = R_a(t) - d(t), \quad \frac{dR_a}{dt} = \alpha[R_d(t) - R_a(t)], \quad (4.1)$$

where  $d(t)$  is the demand received from the retailer and  $i(t)$ ,  $R_d(t)$ ,  $R_a(t)$  denote the inventory level, the desired resupply rate (the control variable) and the actual resupply rate at time  $t$ , respectively. The latter two are related by a simple exponential time delay with parameter  $\alpha$ . This means that demands made of the next echelon in the supply chain are met gradually over time. This stands in contrast to the archetypal OR model, in which orders are completely fulfilled after a (possibly varying) leadtime. That said, by increasing  $\alpha$ , a pure time delay can be approximated. From the notation it is apparent that we are modelling purchasing rather than production systems.

The continuity properties required for the reliable solution of differential equations imposes similar restrictions on the type of system as a whole. Essentially, each quantity defined above (except demand) must vary smoothly over time. Many real

systems meet this requirement. For example, deliveries of high volume products sold in supermarkets are made every day and can be assumed to be continuously replenished. We shall see that intrinsically discrete phenomena like instantaneous jumps in demand can also be treated through this approach by defining the initial inventory level to be below some desired level. By instantaneous we implicitly refer to the chosen time frame. Thus in the above supermarket example, although demand during a day is 'smooth', taken over the whole day an aberrationally large demand should be regarded as instantaneous.

The inventory policy is determined by minimising a cost functional of the form

$$J = \frac{1}{2} \int_0^{\infty} (Q_{11}(i) \times (i)^2 + Q_{22}(R_a) \times (R_a)^2 + R \times (R_d)^2) dt, \quad (4.2)$$

for  $Q_{11}, Q_{22} \geq 0, R > 0$  are the components of the matrices  $Q$  and  $R$  in (3.4). The resupply cost structure comprises the following quantity discounts:

$$\text{Unit purchasing cost} = \begin{cases} c_1, & R_a < q_1 \\ c_2, & q_1 < R_a < q_2 \\ c_3, & q_2 < R_a < q_3 \\ c_4, & R_a > q_3 \end{cases}$$

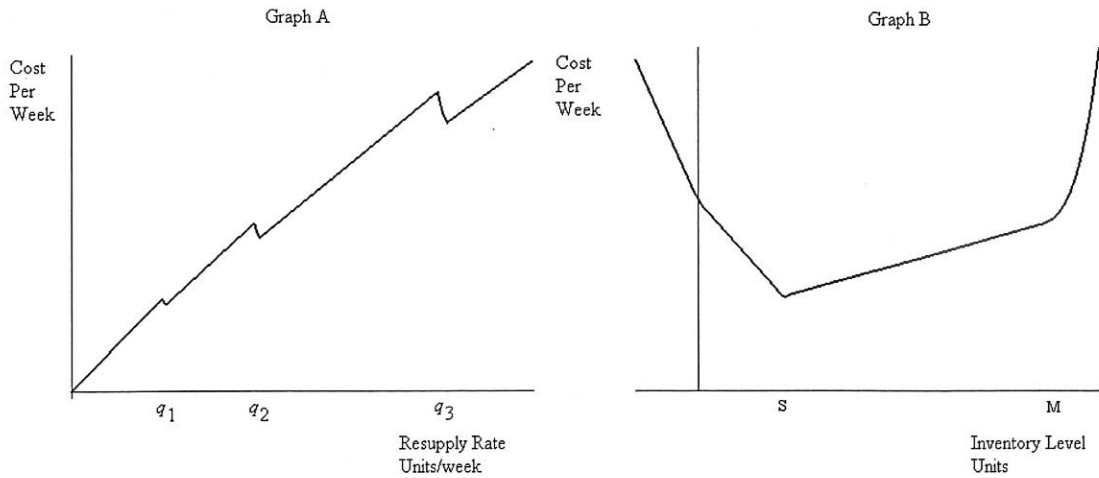


Figure 1. Cost structure for inventory system

Figure 1 shows the cost structures featuring in (4.2), which are traded off against one another by the optimal control algorithm. Graph A plots  $Q_{22}(R_a) \times (R_a)^2$ , the cost per unit time of resupply implied by the quantity discounts. The gradient of the slope in Graph A has the appropriate value  $c_1 \dots c_4$ . This approach allows for any number of quantity discounts. Mathematically it is possible for  $R_a$  to become negative. However for physically realistic systems this remains a pathological possibility. Note that both graphs have been smoothed in order to comply with the stipulations of the optimal control algorithm.

Before expounding the inventory cost structures (shown in Graph B) it is instructive to set forth the specific rôle chosen for this algorithm and how it fits into an existing OR structure. Suppose that the demand characteristics have been estimated based on past data and the average rate per unit, which we assume to be static, time is  $d$ . Based on this rate, the perceived demand variability and the inventory cost structure, an OR type model has calculated the desired safety stock level ( $S$ ). The requirement of continuous replenishment means that on an average day (if this is the smallest time unit we consider)  $d$  units are both demanded and delivered and the average inventory

level during the day is  $S + d/2$  (we assume deliveries occur en masse at the start of the day). We thus say the system is in steady state. However, during days of extra demand this fixed reorder rate would be insufficient to bring stocks back up to  $S$ . After a large instantaneous (ie during one day) demand our optimal control algorithm finds the least cost transient response to bring the system back into steady state. We can now explain the inventory cost structure.

Inventory cost structures tend to be rather notional in character since the calculation of the opportunity cost of holding inventory is based on an element of supposition [Lee and Billington 1992]. In addition to the opportunity cost of the capital and a quantification of the risk of obsolescence, this can incorporate the fixed cost of the warehousing facility. To ensure that the system remains at the steady state when  $i = S$  and  $R_a = d$ , we append an artificial fixed cost to the inventory cost structure (the plot of  $Q_{11}(i) \times (i)^2$  in Graph B above) to make the level that same as that in Graph A when  $R_a = d$ . Above this level storage costs increase linearly with their level (at a rate  $c_3$  per unit) up to the maximum capacity  $M$ , after which the cost increases precipitously to deter incursions into this 'forbidden zone'. Below  $S$  we penalise the risk of stockouts at a rate of  $-c_4$  per unit. The penalty for actual stockouts is at a rate  $-c_5$  per unit, when  $i < 0$ .

In order to avoid the function  $Q_{11}$  'blowing up' near zero we have redefined it to be constant around this point. Choosing this region to be suitably small we have found that this does not affect the optimal policy.

The parameter  $R$  is chosen to have negligible impact on the overall cost structure.

## 5. Examples

We now present two examples to illustrate the preceding theory. Both examples could find an embodiment in a high volume product at a supermarket. With this in mind we nominate one day as the smallest time unit to be considered. Assume that deliveries are made at the start of each day and that day's demand could be considered to occur instantaneously thereafter. The safety stock level could be an 'acceptably' well stocked shelf space, with some extra capacity for deliveries and some in the storeroom. Since shelf and storage space is at a premium in a supermarket we choose the per week holding cost to be relatively large (10% of the product cost). Since well stocked shelves are a primary objective of supermarkets we attach a large notional penalty to falls in safety stocks and stockouts. The first example aims to demonstrate the effect of differing quantity discounts on the transient period following an instantaneous jump in demand.

*Example 1*

Consider the following cost structure:

$$c_1 = \text{£}1.00 \text{ per unit}, c_2 = 97.5p \text{ per unit}, q_1 = 1900 \text{ units/week}, c_3 = 10p \text{ per unit per week}, \\ c_4 = \text{£}10.00 \text{ per unit per week}, c_5 = \text{£}20.00 \text{ per unit per week}, \alpha = 2, R = 1/000$$

with  $d = 1800$  units per week,  $S = 250$  and  $M = 275$ . For consistency we express  $q_1$  in units/week, but since purchasing occurs daily the same discount should be obtainable on a daily basis at the commensurate rate. Suppose that the system is in steady state and an instantaneous jump in demand causes the inventory level to fall to zero in one day. The response of the optimal control algorithm is shown in Figure 2.




Figure 2. Resupply rate and inventory level after a fall in safety stocks

Notice that the resupply rate (Graph A) exploits the purchasing discount at  $R_a = 1900$  units per week. After about a week and a half, during which the purchasing rate has exceeded the quantity discount point, the inventory level returns to within a fifth of its desired level. Thereafter the rate drops off more sharply since the quantity discount is no longer being enjoyed. The inventory level at the end of each day (Graph B) returns to  $S$  after three weeks. Figure 3 shows the resupply rate for different quantity discounts given the same cost structure and jump in demand.

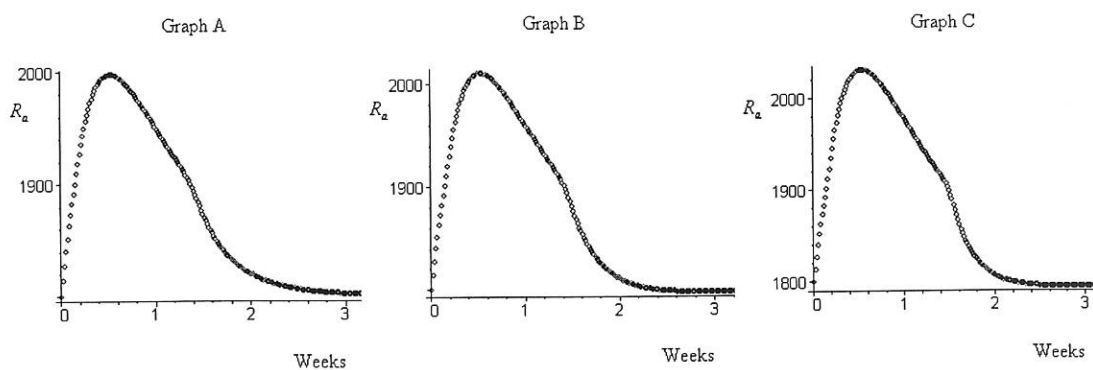


Figure 3. Resupply rates given differing quantity discounts

Graphs A, B and C show, respectively, the rates when  $c_2 = 95p, 92.5p$  and  $90p$ . Observe that as the quantity discount increases, the peak resupply rate also increases to greater exploit the savings offered.

If, during the transient period, another jump in demand caused the inventory to drop, then we could reset the initial conditions of the algorithm and carry on the optimal control. Although not now optimal with respect to the first initial condition, the resulting response constitutes a plausible near optimal policy in the absence of any forecasting capacity. Persistent increase in demand should prompt the reassessment of the medium term demand rate  $d$ . Our next example shows what happens when  $d$  is close to a quantity discount.

#### Example 2

Consider the system in example 1 with  $c_2 = 90p$  per unit but this time  $d = 1890$  units per week. Again, suppose the system is pushed away from its steady state by the same instantaneous demand. Figure 4 shows the resulting behaviour.




Figure 4. Resupply rate and inventory level for example 2.

We see that the system now enters a pattern of periodic oscillation. This is because it is now cheapest to have sustained periods of purchasing above the price break and

incur the additional inventory holding costs and then have a shorter period of much lower purchasing rates during which the inventory is depleted. This example exposes the pathology of one of the drivers of the well-known demand amplification effect [Sterman 1989]. Indeed we can see how a relatively small one-off increase in demand is transformed into wild fluctuations as it is communicated to the next echelon in the supply chain.

## **6. Conclusions**

In this paper we have contrasted the classical OR techniques for inventory systems with those based on optimal control theory. We found that the 'black box' nature of the more sophisticated OR tools hindered the acquisition of significant analytical insight into these models. We acknowledged however that the supremacy of these tools in the business world, should encourage efforts not towards their replacement, rather towards their improvement and supplementation by other methods.

We concluded that OR approaches rightfully dominate in the design of inventory systems but that the prescriptive nature of the policy rules could be improved upon by a dynamic capability. We exploited recent advances in optimal control enabling the treatment of inherently discontinuous issues such as quantity discounts and capacity constraints. The algorithm presented here was designed to compliment an existing OR type policy (having determined the structure of the system) by providing optimal responses to demand fluctuations.

Abad, P.L., Two-level algorithm for decentralized control of a serially connected dynamic system, *Int. J. Systems Sci.*, **16**, 5, 1985, 619-624.

- Archibald, B.C., Silver, E.A., (s,S) policies under continuous review and discrete compound poisson demand, *Man. Sci.*, **24**, 9, 1978, 899-909.
- Arrow, K.J., Karlin, S., Scarf, H., Studies in the mathematical theory of inventory and production, (Stanford: SUP), 1958.
- Axsäter, S., Control theory concepts in production and inventory control, *Int. J. Systems Sci.*, **16**, 2, 1985, 161-169.
- Ben-daya, M., Raouf, A., Inventory models involving lead time as a decision variable, *J. Op. Res. Soc.*, **45**, 5, 1994, 579-582.
- Banks, S.P., Mhana, K.J., Optimal control and stabilization for nonlinear systems, *IMA J. Math Con. & Inf.*, **9**, 1992, 179-196.
- Bensoussan, A., Hurst, E.G. Naslund, B., *Management applications of modern control theory*, (Oxford: North-Holland), 1974.
- Bensoussan, A., Nissen, G., Tapiero, C.S., Optimal inventory and product quality control with deterministic and stochastic deterioration-an application of distributed parameters control systems, *IEEE Trans. Aut. Con.*, June, 1975, 407-412.
- Bensoussan, A., Proth, J.M., Inventory planning in a deterministic environment continuous time model with concave costs, *Euro. Inst. Adv. Studies in Man.*, working paper, April 1982.
- Bradshaw, A., Porter, B., Synthesis of control policies for production-inventory tracking system, *Int. J. Systems Sci.*, **6**, 3, 1975, 225-232.
- Bradshaw, A., Erol, Y., Control policies for production-inventory systems with bounded input, *Int. J. Systems Sci.*, **11**, 8, 1980, 947-959.
- Chandra, P., A dynamic distribution model with warehouse and customer replenishment requirements, *J. Op. Res. Soc.*, **44**, 7, 1993, 681-692.

- Chung, K.J., Tsai, S.F., An algorithm to determine the EOQ for deteriorating items with shortage and a linear trend in demand, *Int. J. Prod. Econ.*, **51**, 1997, 215-221.
- Clark, A.J., Scarf, H., Optimal policies for a multi-echelon inventory problem, *Man. Sci.*, **6**, 1960, 475-490.
- Cook, R.L., Rogowski, R.A., Applying JIT principles to continuous process manufacturing supply chains, *Prod. Inv. Man. J.*, 1996 first quarter, 12-17.
- Dobson, G., Sensitivity of the EOQ model to parameter estimates, *Opns. Res.*, **36**, 4, 1988, 570-574.
- Ehrhardt, R., Finished goods management for JIT production: new models for analysis, *Int. J. Comp. Int. Manuf.*, **11**, 3, 1998, 217-225.
- Elmaghraby, S.E., *The design of production systems*, (New York: Reinhold), 1966.
- Gong, C., Thompson, S., A comment on 'stabilization and optimal control for nonlinear systems', *IMA J. Math. Con. & Inf.*, **12**, 1995, 395-398.
- Elsayed, A.E., Boucher, T.O., *Analysis and control of production systems*, (New Jersey: Prentice-Hall), 1985.
- Federgruen, A., Zipkin, P., An inventory model with limited production capacity and uncertain demands, I: The average cost criterion, *Math. Opns. Res.*, **11**, 1986a, 193-207.
- Federgruen, A., Zipkin, P., An inventory model with limited production capacity and uncertain demands, II: The discounted cost criterion, *Math. Opns. Res.*, **11**, 1986b, 208-215.
- Gallego, G., New bounds and heuristics for (Q,r) policies, *Man. Sci.*, **44**, 2, 1998, 219-233.
- Gallagher, H.P., Morse, P.M., Simond, M., Dynamics of two classes of continuous review inventory systems, *Oper. Res.*, **7**, 1959, 362-384.

- Glasserman, P., Tayur, S., The stability of a capacitated, multi-echelon production-inventory system under a base-stock policy, *Opns. Res.*, **42**, 5, 1994, 913-925.
- Glasserman, P., Tayur, S., Sensitivity analysis for base-stock levels in mutiechelon production-inventory systems, *Man. Sci.*, **41**, 2, 1995, 263-281.
- Harrison, R.F., Banks, S.P., Towards real-time, nonlinear quadratic optimal regulation, *Proc. IFAC Algorithms and Architectures for real-time control, AARC'98, Cancun, Mexico*, 1998, 417-422.
- Kapuscinski, R., Tayur, S., A capacitated production-inventory model with periodic demand, *Opns. Res.*, **46**, 6, 1998, 899-911.
- Kim, G.C., Takeda, E., The JIT Philosophy is the culture in Japan, *Prod. & Inv. Man. J.*, first quarter, 1996, 47-51.
- Lee, H.L., Billington, C., Managing the supply chain inventory: pitfalls and opportunities, *Sloan Man. Rev.*, spring, 1992, 65-73.
- Liao, C.J., Shyu, C.H., Stochastic inventory model with controllable lead time, *Int. J. Systems Sci.*, **22**, 11, 1991, 2347-2354.
- Liao, C.J., Shyu, C.H., An analytical determination of lead time with normal demand, *Int. J. Op. & Prod. Man.*, **11**, 9, 1991, 72-78.
- Lieber, Z., An extension to Modigliani and Hohn's planning horizons results, *Management Science*, **20**, 3, 1973, 319-330.
- McCaffrey, D., Banks, S.P., Lagrangian Manifolds and asymptotically optimal stabilizing feedback, Research Report No. 722, Dept Automatic Control & Systems Engineering, University of Sheffield, 1998.
- Morse, P.M., Solution of a class of discrete-time inventory problems, *Opns. Res.*, **7**, 1959, 67-78.
- Naddor, E., *Inventory systems*, (Florida: Robert E. Kreiger), 1966.



- Ogata, K., *Modern control engineering*, (Prentice-Hall: New Jersey), 1990.
- Porter, B., Taylor, F., Modal control of production-inventory systems, *Int. J. Systems Sci.*, **3**, 3, 1972, 325-331.
- Russell, D.L., Penalty functions and bounded phase coordinate control, *J. SIAM*, **2**, 3, 1965, 409-422.
- Russell, D.L., *Mathematics of finite dimensional control systems*, (Marcel Dekker, New York), 1979.
- Schneeweiss, C.A., Smoothing production by inventory-an application of the Weiner filtering theory, *Man. Sci.*, **17**, 7, 1971, 472-483.
- Simon, H.A., On the application of servomechanism theory in the study of production control, *Econometrica*, **20**, 1952, 247-268.
- Song, J., Zipkin, P., Inventory control in a fluctuating demand environment, *Opns. Res.*, **41**, 2, 1993, 351-370.
- Sterman, J.D., Modelling managerial behaviour: misperceptions of feedback in a dynamic decision making experiment, *Management Science*, **35**, 3, 1989, 321-339.
- Taha, H.A., *Operations research an introduction*, (Prentice Hall, London), 1997.
- Vassian, H. J., Application of discrete variable servo theory to inventory control, *J. Opers. Res. Soc. Am.*, **3**, 3, 1955, 272-282.
- Veinott, A.F., The optimal inventory policy for batch ordering, *Oper. Res.*, **13**, 1965, 424-432.
- Zheng, Y.S., On properties of stochastic inventory systems, *Man Sci.*, **38**, 1, 1992, 87-103.