THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

The use of simple calibrations of individual locations in making transshipment decisions in a multi-location inventory network

Citation for published version:

Archibald, T, Black, D & Glazebrook, K 2010, 'The use of simple calibrations of individual locations in making transshipment decisions in a multi-location inventory network', *Journal of the Operational Research Society*, vol. 61, no. 2, pp. 294-305. <https://doi.org/10.1057/jors.2008.127>

Digital Object Identifier (DOI):

[10.1057/jors.2008.127](https://doi.org/10.1057/jors.2008.127)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Journal of the Operational Research Society

Publisher Rights Statement:

This is a post-peer-review, pre-copyedited version of an article published in Journal of the Operational Research Society. The definitive publisher-authenticated version, Archibald, T., Black, D. & Glazebrook, K. Feb 2010, "The use of simple calibrations of individual locations in making transshipment decisions in a multi-location inventory network", in Journal of the Operational Research Society. 61, 2, p. 293-305; is available online at: <http://dx.doi.org/10.1057/jors.2008.127>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

The use of simple calibrations of individual locations in making transshipment decisions in a multi-location inventory network

T. W. Archibald¹, D. P. Black² and K. D. Glazebrook^{2,3}

¹University of Edinburgh Management School, Edinburgh UK;

²Lancaster University Management School, Lancaster University, Lancaster UK;

³Department of Mathematics and Statistics, Lancaster University, Lancaster UK

Abstract

Demands occur at each location in a network of stock-holding retail outlets. Should a location run out of stock between successive replenishments, then subsequent demands may be met either by transshipping from another location in the network or by an emergency supply from a central depot. We deploy an approximate stochastic dynamic programming approach to develop a class of interpretable and implementable heuristics for making transshipment decisions (whether and from where to transship) which make use of simple calibrations of the candidate locations. The calibration for a location depends upon its current stock, the time to its next replenishment and the identity of the location needing stock. A numerical investigation shows strong performance of the proposed policies in comparison with standard industry practice (complete pooling, no pooling) and a recently proposed heuristic. It points to the possibility of substantial cost savings over current practice.

Keywords: dynamic programming; inventory; stochastic process

Introduction

We consider a network of M stock holding locations (retail outlets), at each of which demands occur at random. Stock at a location is replenished independently of the other locations according to a periodic review policy. Should a stock-out occur at a location between replenishment epochs, subsequent demands must be met either by transshipping stock from another location in the network or by an emergency supply from a central depot. The paper describes how (cost) effective transshipment policies (i.e. rules for determining whether, and from where transshipments should be made in the interests of the network as a whole) may be developed. Our models have been developed in conjunction with a retailer of car parts which has a network of 50 service depots within the UK. They reflect its practices regarding periodic review and transshipment.

While it is true that many models supporting the study of transshipments have been proposed in the literature (Burton and Banerjee, 2005; Kukreja *et al*, 2001), we believe that much past research has failed to fully satisfy the demands of contemporary retailing. In our experience, transshipment in a retail network is often considered as a remedial action taken only when a stock-out occurs. When considering this type of transshipment, one should really weigh the benefit of meeting the demand using inventory held in the system against the cost of transshipment and the increased likelihood of future stock-outs at the location providing the transshipment. This requires an accurate estimate of the marginal value of inventory at each location in the network. This value is a function of the stock level and time to replenishment at each location in the system and is extremely difficult to

assess due to the uncertainty of demand and the potentially large number of locations in the network. Consequently, managers often resort to simple heuristics such as “complete pooling” or “no pooling”. While convenient to implement, these simple policies take little account of the stock levels at locations and are completely independent of the replenishment times.

Faced by the formidable technical challenge to analysis posed by the problem of making transshipment decisions which fully exploit available information at all network locations, it is perhaps not surprising that some studies have resorted to imposing unrealistic limits on network size (Archibald *et al.*, 1997; Rudi *et al.*, 2001) while others have imposed restrictions on the timing and frequency of transshipments (Jönsson and Silver 1987; Tagaras and Cohen 1992; Burton and Banerjee 2005). It has not been uncommon to assume that transshipments are made (as it were) retrospectively, once the demand within a review period has been realised. Under this assumption, the inventory network can either be centrally managed (Wee and Dada, 2005; Hu *et al.*, 2005; Herer *et al.*, 2006) or decentralised (Granot and Sosic, 2003; Hu *et al.*, 2007) and the transportation links can have limited capacity (Özdemir *et al.* 2006). In contrast to the above contributions, our model places no restriction on the number of locations in the network and allows transshipments each time a stock-out occurs.

Studies allowing transshipments in response to stock-outs have generally adopted one-for-one replenishment policies (Lee, 1987; Grahovac and Chakravarty, 2001; Wong *et al.*, 2005) or other continuous review replenishment policies (Asäter, 2003; Minner *et al.*, 2003; Kukreja and Schmidt, 2005). A major difference between these papers and our research is the assumption of a periodic review replenishment policy in our model. Zhao *et al.*, 2006 propose a model with periodic replenishment in which transshipment is considered whenever a demand occurs. However, in sharp contrast to our work, their model assumes a decentralised network in which each location has no information about the inventory level at other locations in the network. Models of retail networks with periodic replenishment often assume simultaneous replenishment of all locations (Cao and Silver, 2005; Herer *et al.*, 2006; Archibald *et al.*, 2007). This can be difficult to achieve in practice in large retail networks. Our model allows independent periodic replenishment in the network.

The model we propose is a stochastic decision process with a finite action space and uncountable state space. A standard proposal would apply the techniques of stochastic dynamic programming (DP) to a suitably designed finite state approximation. However, the high dimensionality of the state space (related to M , the network size) renders this unrealistic in problems of practical size. A recently proposed (pairwise) heuristic, derived by the authors (Archibald *et al.*, 2007), is based on the decomposition of the network into $M(M-1)/2$ two-location systems. Each two-location subproblem is modeled as a two-dimensional stochastic DP and solved using a finite state approximation. A heuristic transshipment policy for the M -location system is then constructed from the optimal value functions of the DP models. In contrast, we deploy an approximate DP approach in which a simple approximation to the DP problem’s high dimensional value function is developed by optimizing costs over a class of static (state independent) transshipment policies for the M -location system, each of which is determined by an $M \times M$ stochastic matrix P . Our proposal, a fully dynamic policy, is then obtained by applying a single DP step which utilises our approximating value function. This broad approach has been used previously in designing dynamic controls for queueing systems (Krishnan, 1987), but to our knowledge opens up a new avenue to the development of dynamic policies for transshipments. The analysis results in a policy which is expressed in terms of a collection of *calibrating indices*, one for each location, and which operates as follows: should a demand arise at location k which has no stock then for each stock-holding location j a simple quantity $I_j(i_j, t_j, k)$ is computed which depends upon j ’s current stock level (i_j) and the time to its next replenishment (t_j). This quantity will also take account of the cost

of transshipping from j to k . The policy nominates the location to supply k as that with the smallest calibrating index, unless this minimal value exceeds the cost of supplying k by emergency order which then becomes the nominated option. We would argue that the policy has the considerable virtue of interpretability in addition to its (cost) effectiveness.

The paper presents the model and develops our index-based dynamic transshipment policy in the next section. A numerical investigation is described in which the cost performance of our proposal (which we denote by SPI) is compared with the optimum for small problems ($M = 3$) for which a (close to) exact analysis via stochastic DP is possible, if expensive. SPI is then compared with that which arises from standard industry practice (complete pooling, no pooling) and from the pairwise heuristic (denoted PW) of Archibald *et al* (2007) in larger problems ($M = 10, 20$) including some in which the depot locations are assumed to be clustered in centres of population. This latter set of comparisons utilises Monte Carlo simulation.

Our findings are summarised as follows:

- SPI was found to be very close to optimal in the small problems ($M = 3$) on which it was tested;
- SPI results in cost savings over competitor heuristics in the large majority of cases studied. Typically these savings are at the level 2 – 3% over the best performing competitor. These savings are recurrent and hence can be significant;
- SPI copes easily with situations in which the locations are *not* all replenished simultaneously. This is in contrast with the heuristic PW which was developed under an assumption of simultaneous replenishment. While SPI outperforms PW (even) when locations are replenished simultaneously, cost savings tend to be greater when this is not the case;
- SPI is considerably easier to compute than PW. A simple version of SPI (called SPI **I**) is particularly convenient to compute and performs very well. SPI has a structure which is straightforward to understand and interpret.

The Model and Methodology

An inventory network consisting of M locations is required to meet demands which occur at random. The inventory at each of the locations is replenished periodically by a common supplier. Periodic replenishment is often used in large networks because of the supplier’s need to coordinate deliveries. We shall suppose that all depots are replenished at equally spaced time points with T_j the time between successive replenishments of location j . Please note that we do *not* require replenishments at all locations to occur with the same frequency or, indeed, simultaneously. The model assumes that the inventory level at location j is restored to level S_j at each replenishment. However, the very simple structure of the proposed heuristic means that it can easily be used as an approximation to inform transshipment decisions in cases where this is not always possible due, for example, to scarcity of supply or non-zero replenishment lead time. We shall suppose the vector $\mathbf{S}$ of replenishment levels to be given. In practice it will be determined by a post-hoc optimisation.

One unit of inventory is required to satisfy each instance of demand in the network. This can be supplied from stock held at the location facing the demand, from stock held at another location in the network via transshipment or by emergency order. It is assumed that the supplier is always able to provide an emergency order. Hence, demand can always be met, even when there is no stock in the inventory network. It is also assumed that the customer accepts whatever method of supply is

offered. This would be the case if, for example, the lead times for transshipment and emergency order are negligible. However, it could also be an appropriate assumption if the lead times are relatively short and the customer is compensated by a discount or free delivery.

The aim of the analysis is to determine the most cost effective way of satisfying each instance of demand in the network. It is therefore necessary to model the demand process at the network locations. If we assume a large population of independent customers each of which experiences *demand generating events* at a constant rate then it is appropriate that instances of demand in the system occur according to a Poisson process with rate λ . Further we assume that an instance of demand occurs at location j with probability ϕ_j , $1 \leq j \leq M$. Successive determinations of location are independent. It follows that location j faces a Poisson process of local demand which is of rate $\lambda\phi_j$, with demand processes for distinct locations being independent.

Costs are incurred as follows:

1. *Purchase cost.* We shall assume that all demands are indeed met. Here an assumption of fixed unit purchase cost per item (C , say) means that these costs will be incurred at rate $C\lambda$ under all policies and hence may be ignored in the analysis. Hence, we set $C = 0$.
2. *Transshipment (fixed) costs.* We suppose that certain fixed costs arise when satisfying a demand arising at location k with an item from location j and this is written R_{jk} . We write R_{Ek} for the fixed cost arising in meeting such a demand by an emergency order. Typically,

$$R_{Ek} \gg R_{jk}, j \neq k \gg R_{kk} = 0, 1 \leq k \leq M.$$

3. *Holding costs.* The cost per item per time unit of holding inventory at location j is h_j , $1 \leq j \leq M$.

The goal of analysis is the identification of a policy (rule for making decisions) for determining which location should meet each demand arising so that the total average cost incurred in operating the network per unit time is minimised. The theory of stochastic dynamic programming (see, for example, Puterman 1994) tells us that (optimal) decisions can be based on the current *system state* only. Here we take the current system state to be the vector whose components are the current inventory levels and the times to next replenishment of all locations. We can therefore restrict attention to the class of *stationary policies* in which decisions are made on the basis of the current state

$$(\mathbf{i}, \mathbf{t}) \equiv (\{i_1, i_2, \dots, i_M\}, \{t_1, t_2, \dots, t_M\}) \quad (1)$$

only. In (1), i_j is the current inventory level at depot j with t_j the time until depot j 's next replenishment, $1 \leq j \leq M$. Even under the restriction to stationary policies, determination of an ϵ -optimal policy via

- (i) the development of a finite state approximation, and
- (ii) direct application of dynamic programming (DP)

is unrealistic other than for very small problems. Hence, our search is for *heuristic policies* which come close to minimising the total system cost rate. We shall propose and implement an approach which develops and modifies a proposal of Krishnan (1987) who utilised it in the context of the dynamic routing of incoming customers to parallel queues for service.

Stage 1: Static location determinations

Our methodology first requires the development of an optimal (or good, at least) *static* policy (i.e. state independent policy) for determining which locations should satisfy each demand. Any such policy will be assumed to be determined by an $M \times M$ stochastic matrix $\mathbf{P}$. The (j, k) th component of $\mathbf{P}$, written P_{jk} , gives the probability that any demand arising at location k is met from location j . If $j \neq k$, this will be a *transshipment*. Successive determinations are made independently using $\mathbf{P}$. We assume that when the supply route indicated by $\mathbf{P}$ is not implementable because of a stock-out (zero inventory), the demand is satisfied via an emergency order to the location at which the demand originated.

Write

$$v_j^{T_j}(S_j, \mathbf{P})$$

for the total expected cost incurred under static regime $\mathbf{P}$ at depot j during time $[0, T_j)$ given that at time 0 the inventory level at j is S_j , $1 \leq j \leq M$. An *optimal static policy* $\mathbf{P}^*$ minimises the aggregate cost rate

$$\Sigma_{j=1}^M v_j^{T_j}(S_j, \mathbf{P})/T_j \quad (2)$$

over the $M(M-1)$ -dimensional space of stochastic matrices. We generalise the above notation to

$$v_j^t(i, \mathbf{P}) \quad (3)$$

in the obvious way and proceed to obtain closed-form formulae for these quantities. To do this, we shall consider costs arising from transshipments, emergency supplies and holding costs in turn and write

$$v_j^t(i, \mathbf{P}) = v_j^t(i, \mathbf{P}, \text{trans}) + v_j^t(i, \mathbf{P}, \text{emerg}) + v_j^t(i, \mathbf{P}, \text{hold}) \quad (4)$$

for the corresponding decomposition of aggregate costs.

First, observe that, under $\mathbf{P}$, requests to supply stock arrive at location j according to a Poisson process with overall rate

$$\lambda_j(\mathbf{P}) = \lambda \Sigma_{k=1}^M \phi_k P_{jk} \equiv \Sigma_{k=1}^M \lambda_{jk}(\mathbf{P}).$$

The transshipment/emergency costs in (4) are those arising when location j is nominated by $\mathbf{P}$ as supplier. We compute the three terms on the r.h.s. of (4) in turn.

1. *Transshipment costs.* We use the fact, based on properties of the Poisson process, that *given* a demand of size d *overall* at location j (i.e. d requests for supply arriving at j in a time period of length t), then the demand arising from requests to transship from j to k has a conditional binomial

$$\text{Bin}\{d, \lambda_{jk}(\mathbf{P})/\lambda_j(\mathbf{P})\}$$

distribution. We then obtain

$$v_j^t(i, \mathbf{P}, \text{trans}) = \left[\exp\{-\lambda_j(\mathbf{P})t\} \sum_{k=1}^M R_{jk} \lambda_{jk}(\mathbf{P})/\lambda_j(\mathbf{P}) \right] \left[\sum_{n=0}^{i-1} n \frac{\{\lambda_j(\mathbf{P})t\}^n}{n!} + i \sum_{n=i}^{\infty} \frac{\{\lambda_j(\mathbf{P})t\}^n}{n!} \right]. \quad (5)$$

2. *Emergency costs.* Similarly, the expected emergency costs incurred on occasions when location j is nominated supplier are given by

$$v_j^t(i, \mathbf{P}, \text{emerg}) = \left[\exp\{-\lambda_j(\mathbf{P})t\} \sum_{k=1}^M R_{Ek} \lambda_{jk}(\mathbf{P}) / \lambda_j(\mathbf{P}) \right] \left[\sum_{n=i+1}^{\infty} (n-i) \frac{\{\lambda_j(\mathbf{P})t\}^n}{n!} \right]. \quad (6)$$

3. *Holding Costs.* In order to compute the holding cost contribution to (4), we consider an initial inventory of i at location j and a time horizon t . Policy $\mathbf{P}$ is in operation. We first condition on the event that the number of requests for supply arriving at location j during $[0, t]$ is $n \leq i$. The consequential holding cost rate at time $s \in [0, t]$ will be $h_j(i-m)$ with probability

$$\binom{n}{m} \left(\frac{s}{t}\right)^m \left(1 - \frac{s}{t}\right)^{n-m}, \quad 0 \leq m \leq n. \quad (7)$$

From (7) we infer that the conditional expected holding cost incurred over the period $[0, t]$ is given by

$$\begin{aligned} \int_0^t \sum_{m=0}^n h_j(i-m) \binom{n}{m} \left(\frac{s}{t}\right)^m \left(1 - \frac{s}{t}\right)^{n-m} ds &= h_j \int_0^t \left(i - \frac{ns}{t}\right) ds \\ &= h_j \left(i - \frac{n}{2}\right) t, \quad n \leq i, \end{aligned} \quad (8)$$

where in the first equality in (8) we use properties of the binomial distribution. If we now condition on the event that the number of requests for supply arriving at location j during $[0, t]$ is $i+n$, $n > 0$, utilisation of standard integral identities yields that the conditional expected holding cost over period $[0, t]$ is given by

$$\begin{aligned} \int_0^t \sum_{m=0}^{i-1} h_j(i-m) \binom{i+n}{m} \left(\frac{s}{t}\right)^m \left(1 - \frac{s}{t}\right)^{i+n-m} ds \\ = h_j t \sum_{m=0}^{i-1} (i-m) \binom{i+n}{m} \frac{m!(i+n-m)!}{(i+n+1)!} \\ = h_j t \sum_{m=0}^{i-1} \frac{(i-m)}{(i+n+1)} = \frac{h_j i(i+1)t}{2(i+n+1)}, \quad n > 0. \end{aligned} \quad (9)$$

From (8) and (9) we infer that

$$v_j^t(i, \mathbf{P}, \text{hold}) = h_j t \exp\{-\lambda_j(\mathbf{P})t\} \left[\sum_{n=0}^i \left(i - \frac{n}{2}\right) \frac{\{\lambda_j(\mathbf{P})t\}^n}{n!} + \sum_{n=i+1}^{\infty} \frac{i(i+1)}{2(n+1)} \frac{\{\lambda_j(\mathbf{P})t\}^n}{n!} \right]. \quad (10)$$

We now obtain $v_j^t(i, \mathbf{P})$ by aggregating the quantities in (5), (6) and (10). We are now equipped with explicit formulae for expected costs incurred under $\mathbf{P}$ which facilitate the optimization in (2). Details of how we approach the optimization numerically are given in the next section describing our numerical investigation.

Stage 2: Policy Improvement Step

We now take the static policy $\mathbf{P}^*$ which optimizes (2) and apply a single dynamic programming (DP) policy improvement step to it. The result is a dynamic policy which takes account of current inventory levels and time to next replenishment at each location before determining how (i.e. from where) any particular demand should be met. This policy will be constructed in such a way that it enjoys the following property: suppose that a demand arises which cannot be met while the system is in some state $(\mathbf{i}, \mathbf{t})$, defined in (1). The policy will take the decision (either to transship from a specified location or to make an emergency supply) such that expected costs incurred over *any* horizon $\tau \geq \max_j t_j$ are minimised under an assumption that all future decisions are made according to optimal static policy $\mathbf{P}^*$. We now show how to design a policy to achieve this.

Consider the situation described in the preceding paragraph, namely that a demand has arisen at some location k which cannot be met with the system in state $(\mathbf{i}, \mathbf{t})$. Hence we must have $i_k = 0$. Call the current time 0. Now consider some time horizon $\tau \geq \max_j t_j$ and write

$$\tau = t_j + N_j T_j + r(t_j, \tau), 1 \leq j \leq m, \quad (11)$$

where the N_j are integers and $0 \leq r(t_j, \tau) < T_j$. The expression in (11) disaggregates horizon τ into (a) time until the next replenishment at j (t_j), (b) a whole number of replenishment cycles for j ($N_j T_j$) and (c) a remainder ($r(t_j, \tau)$). Observe that $N_j + 1$ is the number of replenishments at location j over the time period $[0, \tau)$.

Suppose now that the demand arising at time 0 at location k is met by transshipment from location j (where $i_j > 0$) and that all subsequent decisions up to τ are made according to $\mathbf{P}^*$. Deploying the notation in (3) we can use (11) to write the expected cost arising *at location j* during $[0, \tau)$ as

$$R_{jk} + v_j^{t_j}(i_j - 1, \mathbf{P}^*) + N_j v_j^{T_j}(S_j, \mathbf{P}^*) + v_j^{r(t_j, \tau)}(S_j, \mathbf{P}^*). \quad (12)$$

The first term in (12) is the cost of the transshipment at time 0, the second term is the expected cost until the next replenishment once the transshipment is made, the third term is the expected cost of N_j complete replenishment cycles, and the fourth term is the expected cost incurred during the final time $r_j(t_j, \tau)$ of the horizon. Similarly, the expected cost arising *at location $l \neq j$* during $[0, \tau)$ is

$$v_l^{t_l}(i_l, \mathbf{P}^*) + N_l v_l^{T_l}(S_l, \mathbf{P}^*) + v_l^{r(t_l, \tau)}(S_l, \mathbf{P}^*). \quad (13)$$

Combining the expressions in (12) and (13), we can write the aggregate system cost incurred over $[0, \tau)$ when the decision at time 0 is a transshipment from location j to location k as

$$\begin{aligned} & R_{jk} + v_j^{t_j}(i_j - 1, \mathbf{P}^*) + \sum_{l \neq j} v_l^{t_l}(i_l, \mathbf{P}^*) + \sum_{l=1}^M N_l v_l^{T_l}(S_l, \mathbf{P}^*) + \sum_{l=1}^M v_l^{r(t_l, \tau)}(S_l, \mathbf{P}^*) \\ &= R_{jk} + v_j^{t_j}(i_j - 1, \mathbf{P}^*) - v_j^{t_j}(i_j, \mathbf{P}^*) \\ & \quad + \sum_{l=1}^M \left\{ v_l^{t_l}(i_l, \mathbf{P}^*) + N_l v_l^{T_l}(S_l, \mathbf{P}^*) + v_l^{r(t_l, \tau)}(S_l, \mathbf{P}^*) \right\}, 1 \leq j, k \leq M. \end{aligned} \quad (14)$$

Suppose now that the demand arising at time 0 at location k is met by an emergency order and that all subsequent decisions are made according to $\mathbf{P}^*$. We see from a similar calculation to the above that the resulting aggregate system cost over $[0, \tau)$ is then given by

$$R_{Ek} + \sum_{l=1}^M \{v_l^{t_l}(i_l, \mathbf{P}^*) + N_l v_l^{T_l}(S_l, \mathbf{P}^*) + v_l^{r(t_l, \tau)}(S_l, \mathbf{P}^*)\}, 1 \leq k \leq M. \quad (15)$$

Now compare the expressions in (14) and (15). After some straightforward algebraic manipulation we deduce that the cost minimising choice at time 0 (and hence the choice made by our dynamic policy for transshipments) will be to meet the demand arising at location k by a transshipment from location j if $i_j > 0$ and

$$R_{jk} + v_j^{t_j}(i_j - 1, \mathbf{P}^*) - v_j^{t_j}(i_j, \mathbf{P}^*) = \min_l \{R_{lk} + v_l^{t_l}(i_l - 1, \mathbf{P}^*) - v_l^{t_l}(i_l, \mathbf{P}^*)\} \leq R_{Ek}. \quad (16)$$

If the inequality in (16) is *not* satisfied then the cost minimising choice at time 0 will be to make an emergency order.

We draw together this discussion in the statement of Theorem 1. From (16), we introduce the *location index* $I_j(i_j, t_j, k)$, used by our policy to assess the cost implications of a transshipment from j to k when (i_j, t_j) summarises the current state of location j , as

$$\begin{aligned} I_j(i_j, t_j, k) &= R_{jk} + v_j^{t_j}(i_j - 1, \mathbf{P}^*) - v_j^{t_j}(i_j, \mathbf{P}^*) \\ &= R_{jk} + \exp\{-\lambda_j(\mathbf{P}^*)t_j\} \left(\left\{ \sum_{l=1}^M (R_{El} - R_{jl}) \lambda_{jl}(\mathbf{P}) / \lambda_j(\mathbf{P}^*) \right\} \left[\sum_{n=i_j}^{\infty} \frac{\{\lambda_j(\mathbf{P}^*)t_j\}^n}{n!} \right] \right. \\ &\quad \left. - \sum_{n=0}^{i_j-1} h_j t_j \frac{\{\lambda_j(\mathbf{P}^*)t_j\}^n}{n!} - \sum_{n=i_j}^{\infty} \frac{h_j i_j t_j}{(n+1)} \frac{\{\lambda_j(\mathbf{P}^*)t_j\}^n}{n!} \right). \end{aligned} \quad (17)$$

Note that we obtain expression (17) by substitution from (4)-(6) and (10).

Theorem 1 (DP policy improvement from static policy $\mathbf{P}^*$).

The dynamic policy which results when a single DP policy improvement step is applied to the optimal static policy $\mathbf{P}^$ is constructed as follows: Suppose that a demand arises at location k when the system is in state $(\mathbf{i}, \mathbf{t})$, where $i_k = 0$. For each location j such that $i_j > 0$ compute the index $I_j(i_j, t_j, k)$ in (17). If*

$$\min_j \{I_j(i_j, t_j, k)\} \leq R_{Ek} \quad (18)$$

where the minimum in (18) is taken over all stock holding locations then the dynamic policy chooses to supply location k from any j achieving the minimum. If

$$R_{Ek} < \min_j \{I_j(i_j, t_j, k)\}$$

the dynamic policy chooses to supply location k via an emergency order.

Proof.

The proof is given in the discussion to (17) above. $\square$

Comments

1. Note that the quantity $v_l^{t_l}(i_l, \mathbf{P}^*) - v_l^{t_l}(i_l - 1, \mathbf{P}^*)$ can be interpreted as the marginal value of a unit of inventory (from a total of i_l) at location l when time t_l remains until the next replenishment. From (17), the index $I_j(i_j, t_j, k)$ assesses the cost impact of transshipping from j to k by combining the direct transshipment cost R_{jk} with the loss in value of the inventory at j occasioned by the surrender of a single unit.
2. We have described above the construction of a dynamic policy developed by applying a single policy improvement step to the optimal static policy $\mathbf{P}^*$. We do *not* apply a second (or subsequent) policy improvement step for (at least) three reasons. First, the simplicity and interpretability of the policy structure described in (18) will be lost under further steps. Second, in solutions to dynamic programs developed via policy improvement, much the biggest cost improvement is invariably achieved by the first step. See, for example, Puterman (1994). Third, the implementation of further policy improvement steps will be computationally prohibitive other than for very small networks.
3. We have found that implementation of the above heuristic can be achieved comfortably even with modest computing resources. As is reported below, a competitive $\mathbf{P}^*$ can be found quickly on a standard PC for problems of reasonable size. An *on-line* implementation of the heuristic will then call for (at most) M calculations of the form in (17) whenever a stock-out occurs. Alternatively, a library of values of $v_l^{t_l}(i_l, \mathbf{P}^*) - v_l^{t_l}(i_l - 1, \mathbf{P}^*)$ may be constructed where $1 \leq i_l \leq S_l$ and $t_l = r\alpha$, $1 \leq r \leq T_l/\alpha$ for some discrete time quantum. Good approximations to index values may then be easily inferred from look ups in the library, which is of size $\sum_{l=1}^M S_l T_l / \alpha$.

Numerical Investigation

An extensive numerical investigation has been undertaken to explore the quality of the allocation heuristic developed in the previous section in relation to policies either in current use or proposed in the literature. These are

- *Complete Pooling* (CP): If a depot cannot meet a demand, a transshipment is made from the nearest location in the network with available stock. If there are none, an emergency supply is mandated.
- *No Pooling* (NP): No transshipments are used. If a depot cannot meet a demand, an emergency supply is mandated.
- *Pairwise Heuristic* (PW): This is a heuristic which is based upon the solution of $M(M - 1)/2$ dynamic programs, one for each pair of locations in an M -location network. Suppose that depots j, k have current stock levels i_j, i_k and that both will be replenished in t time units. We write

$$v_{jk}^t(i_j, i_k) \tag{19}$$

for the expected cost incurred *at depots j and k alone* during the time up to the next replenishment when an optimal allocation policy is used between them. The quantity $v_{jk}^t(i_j, i_k)$ is determined by a dynamic programming recursion. Consider now a situation in which a demand occurs at depot k , which has no stock and time t to go until its next replenishment. PW was developed under an assumption of simultaneous replenishment across all locations so, to apply the heuristic,

it is necessary to assume that there is t to go until the next replenishment at each location. If depot j has stock, $i_j > 0$, then the index

$$R_{jk} + v_{jk}^t(i_j - 1, 0) - v_{jk}^t(i_j, 0) \quad (20)$$

is computed and has an interpretation as an indifference emergency cost. If emergency cost R_{Ek} were equal to the quantity in (20) then one would be indifferent (when considering depots j and k alone) between sourcing the demand from location j and via an emergency order. The heuristic PW meets the ~~k -demand~~ from whichever stock-holding location has the smallest value of the index in (20) unless this smallest value exceeds R_{Ek} . In the latter case PW mandates an emergency order.

- *Static Policy Improvement Heuristic (SPI)*: This is the heuristic developed in the preceding section. Note that, while it has a similar index-based structure to PW, it is more flexible in application in that it easily accommodates quite general delivery (replenishment) patterns through the network. The index is also simpler in nature and exhibits an interpretable closed form. See (17) above.

One issue which arises in the design and deployment of SPI is the determination of optimal $\mathbf{P}^*$. See (2) and following. We have adopted a standard approach to this optimization problem in the form of hill-climbing from a large number of initial points in $\mathbf{P}$ -space, with our estimate of $\mathbf{P}^*$ chosen as the resulting local minimum with the smallest associated cost rate. The search for an optimal $\mathbf{P}^*$ is not computationally expensive. For a 10 location problem a competitive $\mathbf{P}^*$ can be found within 2 seconds using a standard PC with a 2.2 GHz processor. We have found, unsurprisingly, that when cost parameters are realistically chosen, our estimates of $\mathbf{P}^*$ are reasonably close to the identity $\mathbf{I}$. Moreover, extensive numerical investigation has demonstrated that the performance of SPI is affected very little if (an estimate of) $\mathbf{P}^*$ is replaced by the identity $\mathbf{I}$ in the formula for the indices $I_j(i_j, t_j, k)$ in (16). See, for example, Table 1 in which average costs per unit time from an optimal transshipment policy (OPT) are compared with those of versions of SPI utilising the identity matrix (SPI $\mathbf{I}$) and the above estimate of $\mathbf{P}^*$ (SPI $\mathbf{P}^*$). The results given are all for three depot problems and hence are small enough for average costs (including optimal) to be computed by means of DP value iteration. Each entry in the table is a cost rate averaged over ten configurations for transshipment costs among the depots (all in the range (10, 30)) with emergency costs ranging from 20 to 100. Please note that the theory does *not* guarantee that SPI $\mathbf{P}^*$ will always yield lower costs than SPI $\mathbf{I}$. Indeed, there are cases in Table 1 where this is not the case.

Note from Table 1 that the cost rates for the two versions of SPI are virtually indistinguishable. They are also never more than 0.56% in excess of the optimal rate. In the bulk of the remaining report of numerical results (Tables 2-7 and 9-11) we shall assume a version of SPI which uses the identity $M \times M$ matrix $\mathbf{I}$.

In part to facilitate easy comparison between PW and SPI, the first phase of the numerical investigation (Tables 2-4) concerns set-ups in which deliveries at all locations are made simultaneously and at equally spaced intervals. We shall take the interval between successive delivery epochs to be the unit of time ($T = 1$). Tables 5-7 concern set-ups in which deliveries are staggered. All cases reported in Tables 2-7 and 9-11 are for networks with ten locations ($M = 10$). The resulting allocation problems are considerably beyond the scope of exact (or even ϵ -approximate) solutions via dynamic programming. Further problem details for the results in Tables 2-7 are as follows:

R_E	OPT	SPI I	SPI P*
20	69.2407	69.2628	69.2663
30	77.0127	77.1206	77.1296
40	79.6126	79.841	79.8471
50	81.8432	82.1387	82.1425
60	83.9798	84.3342	84.3341
70	86.0734	86.4767	86.4753
80	88.1429	88.5892	88.5855
90	90.1972	90.6781	90.6753
100	92.2413	92.7532	92.7497

Table 1. Average costs per unit time incurred by an optimal transshipment policy and two versions of SPI (**I** and **P***) for a range of three depot problems.

- *Demand patterns.* We write $d^i = \lambda\phi_i$ for the demand rate (equivalently, mean demand between replenishments) at location i , $1 \leq i \leq 10$, and $\mathbf{d}$ for the corresponding 10-vector. In all problems $\mathbf{d}$ has the form

$$\mathbf{d} = (d_1, d_1, d_1, d_2, d_2, d_2, d_2, d_3, d_3, d_3),$$

where possible choices of d_1 , d_2 and d_3 are

$$d_1 = d_2 = d_3 = 20,$$

$$d_1 = 25, d_2 = 20, d_3 = 15 \text{ and}$$

$$d_1 = 30, d_2 = 20, d_3 = 10.$$

- *Replenishment levels.* The inventory levels after replenishment are taken to have the form

$$S_i = \lfloor \lambda\phi_i + \sqrt{\lambda\phi_i} \rfloor, 1 \leq i \leq 10,$$

where $\lfloor u \rfloor$ is the largest non-negative integer less than or equal to u . Hence replenishment levels are approximately one standard deviation above the mean demand. This means that, if left to it's own devices, each location would have a roughly 16% chance of exhausting it's stock between replenishments. The reader is referred to comments at the end of this section regarding choice of replenishment levels.

- *Emergency costs.* These are assumed to be common to all locations ($R_{Ei} = R_E$, $1 \leq i \leq 10$) and in the range $[20, 100]$.
- *Transshipment costs.* These are assumed to have a “fixed plus variable” structure, written as

$$B + [0, D] \tag{21}$$

in what follows. In (21), B is a base (fixed) cost, taken to be 10 throughout. The variable element of the cost of transshipping is distance related and is represented $[0, D]$ in (21) and, for any problem instance, is structured as follows: A map of 10 locations is obtained by sampling uniformly within a square grid. The pair of locations at greatest distance from each other (distance x_{max} say) are given a variable transshipment cost of D . Other variable transshipment costs are

Figure 1. Plots of average cost per unit time against R_E for four allocation heuristics for the case $d_1 = d_2 = d_3 = 20$.

proportional to distance. Hence two locations distance x apart have total transshipment cost given by

$$B + Dx(x_{max})^{-1}.$$

In the tables following, we take $D = 40$ throughout.

- *Holding costs.* The cost of holding one unit of inventory for one unit of time is taken to be one, namely the unit of cost.

In Tables 2-4 below find the values of the average cost per unit time in operating the inventory network under the four allocation heuristics CP, NP, PW and SPI. In each case, the estimate of average cost (AVE) is based upon 1,000 simulation runs. Each simulation consists of a burn-in period of 20 cycles before a further 50 cycles are observed. In all cases, the standard errors of the estimates of average cost are placed alongside (SE). Each row of each table corresponds to a value of the emergency cost R_E which increases from 20 (top row) to 100 (bottom row). The optimising cost rate is highlighted in bold for every R_E value.

The primary feature of the numerical results are consistent across Table 2-4. In order to make them more transparent, plots of average cost per unit time against assumed emergency order cost for each allocation heuristic are given in Figure 1 for the Table 2 results.

The reader should note that the heuristic NP becomes hopelessly uncompetitive at even quite modest levels of R_E . Given that under NP each location operates autonomously, the system plainly becomes vulnerable to a serious accumulation of costs from emergency orders. CP stands at the opposite extreme, free as it is to go hunting for available stock at the nearest place which has it. The average cost rate under CP increases very little as R_E goes from 20 to 100. That said, it's operation is still

	CP		NP		PW		SPI		
R_E	AVE	SE	AVE	SE	AVE	SE	AVE	SE	%
20	247.1006	1.1784	237.9020	0.8713	226.1010	0.7953	226.2461	0.799	-0.0641
30	247.3056	1.1893	286.6520	1.3258	237.3255	0.9744	236.9229	0.9812	0.1699
40	247.5106	1.2017	335.4020	1.7807	240.7875	1.0446	239.1421	1.0420	0.6880
50	247.7156	1.2157	384.1520	2.2358	242.7654	1.0827	239.6532	1.0650	1.2986
60	247.9206	1.2311	432.9020	2.6910	244.4316	1.1157	240.0239	1.0825	1.8364
70	248.1256	1.2479	481.6520	3.1463	245.8827	1.1457	240.4275	1.1024	2.2690
80	248.3306	1.2661	530.4020	3.6015	247.2587	1.1792	240.8621	1.1237	2.6557
90	248.5356	1.2855	579.1520	4.0568	248.4768	1.2099	241.3288	1.1476	2.9619
100	248.7406	1.3062	627.9020	4.5121	249.5451	1.2386	241.7654	1.1725	2.8851

Table 2. Estimates of average cost per unit time under four allocation heuristics for the case $d_1 = d_2 = d_3 = 20$.

	CP		NP		PW		SPI		
R_E	AVE	SE	AVE	SE	AVE	SE	AVE	SE	%
20	246.2437	1.1602	237.5915	0.8675	225.6918	0.7878	225.8524	0.7923	-0.0711
30	246.4487	1.1725	286.1815	1.3201	236.6558	0.9603	236.2876	0.9661	0.1558
40	246.6537	1.1864	334.7715	1.7733	239.9969	1.0276	238.4239	1.0261	0.6597
50	246.8587	1.2018	383.3615	2.2265	241.9349	1.0672	238.9218	1.0509	1.2611
60	247.0637	1.2186	431.9515	2.6799	243.6112	1.1023	239.2834	1.0707	1.8087
70	247.2687	1.2369	480.5415	3.1333	245.0871	1.1380	239.6967	1.0937	2.2488
80	247.4737	1.2564	529.1315	3.5868	246.3675	1.1713	240.1547	1.1165	2.5870
90	247.6787	1.2772	577.7215	4.0402	247.4626	1.2025	240.5802	1.1412	2.8608
100	247.8837	1.2992	626.3115	4.4937	248.4908	1.2340	240.9976	1.1675	2.8573

Table 3. Estimates of average cost per unit time under four allocation heuristics for the case $d_1 = 25$, $d_2 = 20$, $d_3 = 15$.

	CP		NP		PW		SPI		
R_E	AVE	SE	AVE	SE	AVE	SE	AVE	SE	%
20	239.3448	1.1347	232.4121	0.8536	220.9445	0.7770	221.0861	0.7824	-0.064
30	239.5498	1.1481	278.4281	1.2991	230.6964	0.9419	230.3306	0.9492	0.1588
40	239.7548	1.1631	324.4441	1.7451	233.9382	1.0128	232.1196	1.0078	0.7835
50	239.9598	1.1796	370.4601	2.1913	237.8522	1.1499	232.5930	1.0321	2.2611
60	240.1648	1.1975	416.4761	2.6375	238.2355	1.1671	232.9565	1.0524	2.2661
70	240.3698	1.2168	462.4921	3.0838	238.6835	1.1866	233.3544	1.0747	2.2837
80	240.5748	1.2373	508.5081	3.5302	240.2858	1.1587	233.8097	1.1005	2.7698
90	240.7798	1.2592	554.5241	3.9766	241.4084	1.1909	234.2101	1.1254	2.8050
100	240.9848	1.2822	600.5401	4.4229	242.4576	1.2220	234.6009	1.1518	2.7212

Table 4. Estimates of average cost per unit time under four allocation heuristics for the case $d_1 = 30$, $d_2 = 20$, $d_3 = 10$.

more costly than either PW or SPI which take more careful account of current inventory levels when choosing whether and how to transship. When R_E is small (≤ 30 , say), there is very little difference between average cost rates incurred by PW and SPI. However, SPI's relative performance becomes increasingly strong as R_E increases above these levels.

Closer numerical investigation of the indices in (18) and (20) which underlie the operation of PW and SPI lead to the following conclusions: the indices seem to coincide when the inventory i_j held by the potential sourcing location j is either very large or very small. The exclusive attention which the PW index in (20) gives to the relation between sourcing location j and the *currently* needy location k leads to slightly strange behaviour of this index for mid-range values of i_j . Deployment of the static policy $\mathbf{P}^*$ to approximate future commitments yields a rather more balanced analysis and a more soundly based calibration from (17).

The strong relative performance of SPI is even clearer for problems in which replenishments are *not* assumed to be simultaneous. The problems summarised in Tables 5-7 have the characteristics of those above save only that in each problem the 10 locations are grouped (randomly) in 5 pairs of two. Locations in the same pair have the same delivery times chosen randomly from 5 equally spaced times within the replenishment cycle. The time between successive replenishments is common to all locations and is again taken to be the unit of time.

The qualitative properties of heuristics NP and CP are much as above. NP's vulnerability to large costs incurred from an excess of emergency orders is much as before, while staggered replenishments means that CP's earlier small exposure to emergency costs is reduced further. PW is not designed for general delivery configurations and consequently performs poorly in comparison to SPI. The latter heuristic is the clear winner in all problem instances.

We conclude the account of the numerical evidence in support of the transshipment heuristic SPI with a brief description of results from a study of larger (20 depot) problems in which the depots are *clustered* in centres of population. Each row of Table 8 has a summary (AVE, SE) of cost rates arising from the application of five transshipment heuristics (CP, NP, PW, SPI **I** and SPI **P**^{*}) to 100 problems, each with an assigned value of the emergency cost R_E . The 100 problems generated for each row combine a choice from among 10 randomly generated 20-vectors of depot demand rates (each component of which is chosen independently from the uniform $U(10, 30)$ distribution) with a choice from among 10 patterns of transshipment cost, each one arising from a clustered depot geography as follows:

The positions of five *hub* depots (which might, for example, be thought to be located in city centres) were obtained by sampling independently and uniformly within the unit square. The positions of the remaining fifteen depots were obtained by first determining the hub to which they belonged (independently, with equal probabilities) and then establishing their position relative to the chosen hub. In all cases the latter determination was achieved by sampling uniformly from within a circle centred at the hub. Finally, transshipment costs in all cases were given by

$$R_{jk} = 10 + 70d_{jk}, \quad j \neq k,$$

where d_{jk} is the Euclidean distance between depots j and k . The cost rate for each of the 100 problem instances underlying each entry in the table was obtained from a simulation involving 1,000 runs, as described above. Other assumptions concerning times between deliveries, replenishment levels and holding costs are as in the earlier study reported in Tables 2-4. The results in Table 8 confirm our earlier numerical findings in this larger and more complex context. As in Table 1, the cost performance

	CP		NP		PW		SPI		
R_E	AVE	SE	AVE	SE	AVE	SE	AVE	SE	%
20	261.9085	0.2189	237.1313	0.1196	242.4200	0.1605	229.3017	0.1166	3.4145
30	261.9085	0.2189	285.1499	0.1819	255.1086	0.1929	241.5704	0.1447	5.6042
40	261.9085	0.2189	333.1686	0.2442	256.0314	0.1945	244.1097	0.1517	4.8837
50	261.9085	0.2189	381.1873	0.3066	255.9251	0.1933	245.3469	0.1532	4.3115
60	261.9085	0.2189	429.2057	0.3690	255.9080	0.1926	246.5922	0.1546	3.7778
70	261.9085	0.2189	477.2244	0.4313	255.9232	0.1922	247.8756	0.1560	3.2466
80	261.9085	0.2189	525.2430	0.4937	255.9417	0.1920	249.1573	0.1575	2.7229
90	261.9085	0.2189	573.2617	0.5561	255.9756	0.1919	250.3957	0.1591	2.2284
100	261.9085	0.2189	621.2804	0.6185	256.0155	0.1919	251.5887	0.1604	1.7595

Table 5. Estimates of average cost per unit time when deliveries are staggered for four allocation heuristics for the case $d_1 = d_2 = d_3 = 20$.

	CP		NP		PW		SPI		
R_E	AVE	SE	AVE	SE	AVE	SE	AVE	SE	%
20	261.1774	0.2174	236.7129	0.1191	241.1615	0.1576	228.8315	0.1156	3.4442
30	261.1774	0.2174	284.5380	0.1812	253.9632	0.1893	241.0827	0.1439	5.3428
40	261.1774	0.2174	332.3631	0.2432	254.9104	0.1908	243.5898	0.1509	4.6474
50	261.1774	0.2174	380.1884	0.3054	254.8273	0.1897	244.8014	0.1523	4.0955
60	261.1774	0.2174	428.0136	0.3675	254.8301	0.1891	246.0301	0.1535	3.5768
70	261.1774	0.2174	475.8389	0.4296	254.8657	0.1889	247.2475	0.1549	3.0812
80	261.1774	0.2174	523.6640	0.4918	254.8890	0.1887	248.4617	0.1562	2.5868
90	261.1774	0.2174	571.4891	0.5539	254.9385	0.1886	249.6773	0.1576	2.1072
100	261.1774	0.2174	619.3144	0.6161	254.9967	0.1885	250.8266	0.1590	1.6625

Table 6. Estimates of average cost per unit time when deliveries are staggered for four allocation heuristics for the case $d_1 = 25$, $d_2 = 20$, $d_3 = 15$.

	CP		NP		PW		SPI		
R_E	AVE	SE	AVE	SE	AVE	SE	AVE	SE	%
20	254.1299	0.2089	232.4410	0.1172	237.6835	0.1597	224.7297	0.1135	3.4314
30	254.1299	0.2089	278.1615	0.1783	247.6632	0.1842	235.8340	0.1406	5.0159
40	254.1299	0.2089	323.8820	0.2394	248.1058	0.1843	238.0149	0.1466	4.2396
50	254.1299	0.2089	369.6022	0.3006	248.0503	0.1833	239.2037	0.1480	3.6984
60	254.1299	0.2089	415.3226	0.3618	248.0801	0.1827	240.3930	0.1491	3.1977
70	254.1299	0.2089	461.0434	0.4230	248.1434	0.1826	241.6293	0.1506	2.6959
80	254.1299	0.2089	506.7639	0.4842	248.2265	0.1825	242.8213	0.1521	2.2260
90	254.1299	0.2089	552.4844	0.5454	248.3040	0.1825	244.0122	0.1536	1.7588
100	254.1299	0.2089	598.2046	0.6066	248.3839	0.1824	245.1590	0.1552	1.3154

Table 7. Estimates of average cost per unit time when deliveries are staggered for four allocation heuristics for the case $d_1 = 30$, $d_2 = 20$, $d_3 = 10$.

R_E	CP		NP		PW		SPI I		SPI P*	
	AVE	SE	AVE	SE	AVE	SE	AVE	SE	AVE	SE
20	470.584	0.5037	464.708	0.3821	431.687	0.3315	431.764	0.3336	432.025	0.3347
30	470.587	0.5038	557.634	0.5815	452.679	0.4051	451.246	0.4064	451.433	0.4078
40	470.589	0.5039	650.56	0.7811	459.79	0.4365	455.894	0.4338	455.937	0.4345
50	470.592	0.5039	743.486	0.9808	463.777	0.4526	457.275	0.4434	457.256	0.4439
60	470.595	0.5040	836.411	1.1805	466.796	0.4624	457.755	0.4454	457.676	0.4457
70	470.597	0.5041	929.337	1.3803	469.592	0.4711	458.104	0.4458	458.002	0.4460
80	470.6	0.5042	1022.26	1.5801	472.187	0.4797	458.478	0.4463	458.335	0.4463
90	470.602	0.5043	1115.19	1.7799	474.586	0.4878	458.854	0.4470	458.712	0.4470
100	470.605	0.5044	1208.11	1.9796	476.744	0.4952	459.275	0.4479	459.092	0.4478

Table 8. Estimates of average cost per unit time, averaged over 100 problem instances for clustered networks of depots under five allocation heuristics.

of the two versions of SPI are very close. Moreover, as in Tables 2-4 both versions of SPI outperform all other heuristics save only the single instance of the marginal superiority of PW in the case with low emergency cost $R_E = 20$.

Remark

In all of the numerical investigations reported in Tables 2-8, replenishment levels for all locations are set at approximately one standard deviation above the mean demand between replenishments. The implication of this choice is that, while individual locations have a significant chance (approximately 16%) of exhausting their stock between replenishments, the inventory network *as a whole* is very unlikely to be depleted. This is precisely the kind of set up in which transshipments can play a valuable role in meeting demands and reducing costs. In such scenarios such relatively low replenishment levels are likely to be dictated by considerations of available storage space. This is often the case, for example, when the network concerns the retail of car parts.

In order to explore this issue further, the computations of Tables 2-4 were repeated with replenishment levels of the form

$$S_i = \lfloor \lambda \phi_i + \alpha \sqrt{\lambda \phi_i} \rfloor, 1 \leq i \leq 10,$$

for a range of α between 0.5 and 2. The results are presented in Tables 9-11 for the case with emergency cost $R_E = 70$. At the bottom end of the α -range transshipments occur frequently under all of CP, PW and SPI and emergency orders very occasionally. At the top end of the α -range transshipments occur very occasionally and emergency orders almost never. See Figure 2.

In all cases the tables show that the average cost rate is minimised when the heuristic SPI is applied at replenishment levels determined by taking α to be around 1.75. Should storage space be more limited (as indicated above) a smaller α may have to be adopted thereby creating a greater role for transshipments, as demonstrated in Figure 2. However, the evidence of Tables 9-11 is that SPI outperforms the other heuristics at all α values within the range considered. SPI's performance is particularly strong in the range $0.5 \leq \alpha \leq 1.25$ for which transshipment plays a significant role.

	Policy Heuristic								
	CP		NP		PW		SPI		
α	AVE	SE	AVE	SE	AVE	SE	AVE	SE	%
0.5	409.9400	3.0746	807.1223	4.5154	391.9529	2.7234	382.4474	2.7572	2.4854
0.75	304.3043	1.9054	621.2823	3.8056	297.3307	1.7079	289.2485	1.6818	2.7942
1.0	248.1256	1.2479	481.6520	3.1463	245.8827	1.1458	240.4275	1.1024	2.2689
1.25	217.5717	0.8690	381.0690	2.5609	217.0021	0.8168	213.7524	0.7764	1.5203
1.5	202.1929	0.6249	313.0155	2.0543	202.2273	0.6017	200.4500	0.5740	0.8694
1.75	196.0133	0.4548	268.5621	1.6195	196.1639	0.4454	195.2533	0.4267	0.3892
2.0	195.8645	0.3378	241.8095	1.2594	195.9209	0.3335	195.5088	0.3220	0.1819

Table 9. Estimates of average cost per unit time under four allocation heuristics for the case $R_E = 70$, $d_1 = d_2 = d_3 = 20$

α	Policy Heuristic								%
	CP		NP		PW		SPI		
	AVE	SE	AVE	SE	AVE	SE	AVE	SE	
0.5	454.2394	3.5270	869.3237	4.7414	432.1395	3.1326	422.1207	3.1862	2.3734
0.75	326.6476	2.1685	667.7856	4.0100	317.4047	1.9416	308.4683	1.9256	2.8970
1.0	247.2687	1.2369	480.5415	3.1334	245.0871	1.1381	239.6967	1.0937	2.2488
1.25	216.5079	0.8471	380.2053	2.5420	216.2063	0.8040	212.9814	0.7610	1.5141
1.5	201.2212	0.6047	311.3294	2.0308	201.3145	0.5858	199.5993	0.5561	0.8125
1.75	195.3306	0.4413	266.9383	1.5986	195.3580	0.4299	194.5662	0.4121	0.3928
2.0	196.5525	0.3062	236.4795	1.1567	196.5421	0.2976	196.2492	0.2895	0.1492

Table 10. Estimates of average cost per unit time under four allocation heuristics for the case $R_E = 70$, $d_1 = 25$, $d_2 = 20$, $d_3 = 15$.

α	Policy Heuristic								%
	CP		NP		PW		SPI		
	AVE	SE	AVE	SE	AVE	SE	AVE	SE	
0.5	444.5052	3.5048	847.4459	4.7076	423.7336	3.1205	414.1330	3.1775	2.3182
0.75	294.5509	1.8682	598.3704	3.7314	289.0270	1.6759	280.6488	1.6454	2.9852
1.0	240.3698	1.2168	462.4921	3.0839	239.0264	1.1256	233.3544	1.0748	2.4306
1.25	219.0144	0.9180	392.1399	2.6368	218.7860	0.8647	214.9615	0.8187	1.7791
1.5	198.7155	0.5848	302.0420	1.9602	198.8731	0.5667	197.2359	0.5371	0.7501
1.75	193.2594	0.4216	259.3502	1.5417	193.4343	0.4168	192.5867	0.3959	0.3492
2.0	193.7812	0.3090	234.9930	1.1969	193.9233	0.3089	193.5003	0.2957	0.1451

Table 11. Estimates of average cost per unit time under four allocation heuristics for the case $R_E = 70$, $d_1 = 30$, $d_2 = 20$, $d_3 = 10$.

Figure 2. Plot of average number of emergency order and transshipments per unit time under four allocation heuristics for the case $R_E = 70$, $d_1 = d_2 = d_3 = 20$.

Conclusions

We have argued that much of the current literature on transshipments does not meet the needs of contemporary retailing. Our proposed model takes full account of inventory and transshipment costs and adopts a realistic view of demand uncertainty and the size of the network. The resulting stochastic decision process has a finite action space and an uncountable state space of high dimension rendering unrealistic any direct application of stochastic DP. We implement an approximate DP approach which applies a single policy improvement step to an optimal static stochastic proposal for transshipments. The result is a simple, interpretable and easily implementable class of policies which make transshipment decisions in terms of calibrations of the candidate (stock holding) locations in the network. A numerical study has shown these policies to be close to optimal for small networks and to outperform standard proposals in large ones. A further inference from the numerical work is that the first stage optimization over the static class is really unnecessary and that the simple ‘no pooling’ choice of the identity $\mathbf{I}$ at this stage yields outstanding results.

Acknowledgements

The authors acknowledge the support of the EPSRC for this work through the award of grants GR/T08562/01 and GR/S45188/01. They also would like to express their appreciation of the helpful comments of a referee.

References

Archibald, T. W., D. P. Black, K. D. Glazebrook. 2007. An index heuristic for transshipment decisions in multi-location inventory systems based on a pairwise decomposition. *European Journal of Operational Research*, to appear, DOI:10.1016/j.ejor.2007.09.019.

- Archibald, T. W., S. A. E. Sassen, L. C. Thomas. 1997. An optimal policy for a two depot inventory problem with stock transfer. *Management Science* **43**(2) 173–183.
- Axsäter, S. 2003. A new decision rule for lateral transshipments in inventory systems. *Management Science* **49**(9) 1168–1179.
- Burton, J., A. Banerjee. 2005. Cost-parametric analysis of lateral transshipment policies in two-echelon supply chains. *International Journal of Production Economics* **93-94** 189–195.
- Cao, D., E. A. Silver. 2005. A dynamic allocation heuristic for centralized safety stock. *Naval Research Logistics* **52**(6) 513–526.
- Grahovac, J., A. Chakravarty. 2001. Sharing and lateral transshipment of inventory in a supply chain with expensive low-demand items. *Management Science* **47**(4) 579–594.
- Granot, D., G. Sosic. 2003. A three-stage model for decentralized distribution system of retailers. *Operations Research* **51** 771–784.
- Herer, Y. T., M. Tzur, E. Yücesan. 2006. The multilocation transshipment problem. *IIE Transactions* **38**(3) 185–200.
- Hu, J., E. Watson, H. Schneider. 2005. Approximate solutions for multi-location inventory systems with transshipments. *International Journal of Production Economics* **97**(1) 31–43.
- Hu, X. X., I. Duenyas, R. Kapuscinski. 2007. Existence of coordinating transshipment prices in a two-location inventory model. *Management Science* **53**(8) 1289–1302.
- Jönsson, H., E. A. Silver. 1987. Analysis of a two-echelon inventory control system with complete redistribution. *Management Science* **33**(2) 215–227.
- Krishnan, K. R. 1987. Joining the right queue: a Markov decision rule. *Proc. 28th IEEE Conference on Decision Control*. 1863–1868.
- Kukreja, A., C. P. Schmidt. 2005. A model for lumpy demand parts in a multi-location inventory system with transshipments. *Computers and Operations Research* **32** 2059–2075.
- Kukreja, A., C. P. Schmidt, D. M. Miller. 2001. Stocking decisions for low-usage items in a multilocation inventory system. *Management Science* **47**(10) 1371–1383.
- Lee, H. L. 1987. A multi-echelon inventory model for repairable items with emergency lateral transshipments. *Management Science* **33**(10) 1302–1316.
- Minner, S., E. A. Silver, D. J. Robb. 2003. An improved heuristic for deciding on emergency transshipments. *European Journal of Operational Research* **148**(2) 382–400.
- Özdemir, D., E. Yücesan, Y. T. Herer. 2006. Multi-location transshipment problem with capacitated transportation. *European Journal of Operational Research* **175**(1) 602–621.
- Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, New York.
- Rudi, N., S. Kapur, D. E. Pyke. 2001. A two-location inventory model with transshipment and local decision making. *Management Science* **47**(12) 1668–1680.

- Tagaras, G., M. A. Cohen. 1992. Pooling in two-location inventory systems with non-negligible replenishment lead times. *Management Science* **38**(8) 1067–1083.
- Wee, K. E., M. Dada. 2005. Optimal policies for transshipping inventory in a retail network. *Management Science* **51**(10) 1519–1533.
- Wong, H, G. J. van Houtum, D. Cattrysse, D. Van Oudheusden. 2005. Simple, efficient heuristics for multi-item multi-location spare parts systems with lateral transshipments and waiting time constraints. *Journal of the Operational Research Society* **56**(12) 1419–1430.
- Zhao, H., V. Deshpande, J. K. Ryan. 2006. Emergency transshipment in decentralized dealer networks: When to send and accept transshipment requests. *Naval Research Logistics* **53**(6) 547–567.