University of Groningen

The routed inventory pooling problem with multiple lateral transshipments

Bouma, Harmen W.; Teunter, Ruud H.

Published in:
International Journal of Production Research

DOI:
[10.1080/00207543.2015.1082668](https://doi.org/10.1080/00207543.2015.1082668)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2016

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Bouma, H. W., & Teunter, R. H. (2016). The routed inventory pooling problem with multiple lateral transshipments. *International Journal of Production Research*, 54(12), 3523-3533.
<https://doi.org/10.1080/00207543.2015.1082668>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

The routed inventory pooling problem with multiple lateral transshipments

Harmen W. Bouma and Ruud H. Teunter*

Department of Operations, University of Groningen, Groningen, The Netherlands

(Received 27 June 2014; accepted 6 August 2015)

We consider a single item, single-period inventory problem with two identical retailers who can pool stocks at multiple predetermined points in time. Since routing decisions are often taken at the multi-item level in practice, we use a predetermined route for redistribution at the single-item level. The objective is to determine order-up-to levels and transshipment decisions that minimise the sum of inventory holding, backorder and fixed as well as variable transshipment costs. We present a Dynamic Programming formulation to solve the problem and we compare the performance of our routed pooling policy to situations without pooling and with complete pooling in an extensive numerical study. The results reveal when routed pooling performs well and when multiple pooling moments provide significant benefits.

Keywords: inventory; routed pooling; stochastic demand; Dynamic Programming

1. Introduction

Combined inventory control and routing/transportation problems have recently received a lot of attention (see [Andersson et al. 2010](#), for a review). However, the existing literature on these so-called inventory routing problems is heavily biased towards deterministic problems and has not addressed inventory pooling as a mechanism for coping with demand uncertainty. Inventory pooling *has* been widely studied in the more general inventory control literature (see e.g. [Paterson et al. 2011](#)), supported by the increasing popularity of vendor managed inventory systems and technical advances in the design of electronic data interchange systems.

In a recent paper, [Bouma and Teunter \(2014\)](#) are the first to explore the interface between inventory pooling and routing. They consider a system with a pre-determining route (in order to allow coordination across items), assume that pooling is only allowed once during the replenishment cycle, and further assume that the transshipment cost is proportional in the number of units transshipped, i.e. that there is no fixed transshipment cost. Although the setting with a single pooling opportunity per replenishment cycle is a natural starting point, it does limit the potential savings from stock pooling. Consider, for instance, the situation where a number of retailers are located much closer to each other than to the supplier. This is very realistic, with e.g. the retailers situated in one or more European countries or American states and the supplier located in Asia. Given the large distance between the supplier and retailers, and the corresponding large fixed transportation (by container) cost and customs clearing effort, long replenishment cycles are typically used. In such situations, transporting stocks between retailers is much easier and less costly, and may be done on a more frequent basis. In this paper, we therefore explore what the cost benefits are of allowing multiple shipments between retailers during a supplier replenishment cycle.

Transshipping items between retailers does of course still involve certain costs, in particular for handling and moving the items. Since multiple items can often be shipped on a single pallet, part of those costs will be independent of the number of items shipped. Instead of only including a variable transport cost per item, as is done by [Bouma and Teunter \(2014\)](#), we will therefore also include a fixed transport cost per shipment. Our results will reveal the effects of the cost setting on the transshipment policy and its cost savings potential. These exact results are obtained numerically using a Dynamic Programming formulation, and are restricted to a system with two identical retailers in order to limit the state space and thereby ensure solvability within reasonable time. So, our model is not a generalisation of that in [Bouma and Teunter \(2014\)](#) in all aspects, but does shed new light on the effects of the transshipment cost structure and of allowing multiple transshipments per replenishment cycle.

The organisation of the paper is as follows. In Section 2, we review related literature and position our work. In Section 3, we state the problem definition and present a Dynamic Programming formulation. Section 4 contains an extensive numerical study in which the performance of multiple lateral transshipments and the influence of transshipment costs on optimal

*Corresponding author. Email: r.h.teunter@rug.nl

order-up-to levels and total expected costs are examined. Results are compared to the no pooling and complete pooling policies. We conclude in Section 5 and provide directions for further research.

2. Related work

This study is on the interface of inventory pooling and inventory routing. We will therefore review the two corresponding streams of literature, before discussing the few studies that combine these topics and pointing out the contribution of this paper.

2.1 Inventory routing

Inventory routing models trade off transportation costs from one or multiple depots to retailers against the inventory costs at those retailers. The objective is to determine cost-minimising routes and order/transport frequencies, given possible restrictions on vehicle capacity and maximum driving distance. We next highlight key contributions and refer interested readers to [Andersson et al. \(2010\)](#) for an extensive literature review.

Most early studies, such as that by [Federgruen and Zipkin \(1984\)](#), consider a single item, single period, single depot setting. Some contributions in the 80s do already generalise to multi-period situations, e.g. [Blumenfeld et al. \(1985\)](#) and [Miller \(1987\)](#). More recently, multiple products and/or multiple depot models have been analysed. [Bertazzi, Paletta, and Speranza \(2002\)](#) introduced a finite horizon model with multiple items. They develop a heuristic for finding a near-optimal policy of the *must-can order* type, where some item must be replenished when the inventory drops to or below its must order level, and can be replenished opportunistically when the inventory drops to or below the larger can order level and some other item must be replenished, thereby reducing fixed replenishment costs. [Savelsbergh and Song \(2007\)](#) consider a network with multiple depots and multiple demand locations for which they develop a heuristic that finds high-quality solutions for realistic size instances.

2.2 Inventory pooling

Studies on inventory pooling can be divided into two classes. The first contains models with proactive lateral transshipments in which stock levels are balanced at predetermined moments in time, while the second contains models with reactive lateral transshipments in which stock movements only occur in case of stock-outs. In this paper, we consider proactive lateral transshipments, and we will therefore focus on contributions to this particular research area in this section. For a broader overview of developments in both proactive and reactive transshipment models we refer to [Paterson et al. \(2011\)](#).

In the literature, different options for the timing of redistribution are considered. Some pioneering papers ([Allen, 1958, 1961, 1962](#)) deal with redistribution at the beginning of a period. Early papers that deal with redistribution during the period, as we do in this study, are those by [Gross \(1963\)](#) and [Karmarkar and Patel \(1977\)](#). Finally, [Agrawal, Chao, and Seshadri \(2004\)](#) take a step further by determining the moment of redistribution dynamically, based on the actual demand observed. In these studies, order-up-to levels are either assumed to be known or not taken into consideration.

Many papers (for instance [Tagaras and Vlachos 2002](#); [Jönsson and Silver 1987](#)) deal with the determination of order-up-to levels in combination with proactive transshipments. [Tagaras and Vlachos \(2002\)](#) conclude that the performance of pooling is influenced by the moment of redistribution, which in turn depends on the type and variability of demand distributions. [Jönsson and Silver \(1987\)](#) consider redistribution towards the end of the order cycle since backorders are most likely to occur at that stage.

2.3 Inventory routing and pooling

To the best of our knowledge, only two studies consider the combination of inventory routing and pooling. [Coelho, Cordeau, and Laporte \(2012\)](#) analyse a rolling horizon framework where decisions are based on forecasts that may change over time. They allow pooling of stocks, in reaction to updated forecasts, to occur simultaneously with regular replenishments. A heuristic solution procedure is developed for two ordering policies, and pooling is shown to reduce costs. [Bouma and Teunter \(2014\)](#) show that further savings can be obtained by allowing pooling to take place in between replenishments. For reasons of tractability, they restrict their attention to policies with a single opportunity to pool stocks per replenishment cycle, which obviously limits the savings potential. For the same reason, they assume that transportation costs are proportional to the number of items transshipped, i.e. that the fixed transshipment cost is zero. This does not apply to most practical situations, as discussed in the introduction, since multiple items can often be jointly picked from a warehouse and shipped on a single

pallet. Under these simplifying assumptions, Bouma and Teunter (2014) are able to perform an exact analysis for systems with up to three non-identical retailers.

In this paper, we build on the work of Bouma and Teunter (2014) by (a) including both fixed and variable transshipment costs and (b) allowing multiple pooling opportunities per replenishment cycle. The downside of considering this more realistic cost and policy setting is that the system is no longer analytically tractable. We will therefore apply a Dynamic Programming approach and develop insights in an extensive numerical study. As the computational time of DP algorithms increases rapidly with the state space (the so-called curse of dimensionality), our study is limited to systems with two identical retailers.

3. Model

We consider an inventory system with two identical retailers, numbered 1 and 2, that are (re)supplied from an external supplier at specific coordinated times. In most real-life settings, firms limit the frequency of external replenishments to their retailers in order to reduce fixed (external) ordering costs, obtain quantity discounts and/or satisfy minimum order quantity restrictions. The obvious downside of doing so is that retailers have to cope with demand uncertainty over the period between regular replenishments, i.e. during a replenishment cycle, and this is where pooling stocks between retailers offers potential benefits. By pooling, retailers can rebalance stocks in between regular replenishments.

We therefore focus on one replenishment cycle, which we refer to as the period. The stock of retailer i , $i = 1, 2$, is replenished up to level S_i at the start of the period and no more external resupply is allowed during the period. The period is divided into M stages of equal length. During each stage t , $t = 1 \dots, M$, the retailers face identically and independently distributed demands D_{it} . At the end of each stage, holding costs h for all on hand items and backorder costs b for all demands backordered are incurred at the retailers. Furthermore, (after holding and backorder costs are determined) at the end of each stage $t = 1, 2, \dots, M - 1$, there is a potential opportunity to balance stock levels of the retailers by a transshipment. For each lateral transshipment, a fixed cost F and variable costs T per item transshipped are incurred. Backorders are satisfied as soon as on hand stock becomes available (at any stage). We assume that transshipment times are negligible compared to the length of the period, i.e. they take place overnight or take less than a day, whereas regular orders are placed weekly or monthly. The timing of events in a replenishment cycle is depicted in Figure 1.

The main focus of this paper is on the routed pooling policy, for which stock can only be moved from retailer 1 to 2. Recall from the introduction that retailers typically stock many items (stock keeping units) and a predetermined directed route allows coordination of shipments. However, we will also look into the complete (either direction) and no pooling policies, since they serve as benchmarks to investigate the performance of routed pooling. Furthermore, we will compare the performance of allowing multiple lateral transshipments to a single transshipment opportunity in a numerical study in Section 4.

In order to formally state the problem described above, we will make use of the notation in Table 1. As each transshipment considered is a lateral transshipment, we will drop the term 'lateral' in the remainder of the paper.

3.1 No pooling

If pooling stocks is not allowed, the problem simplifies to a standard newsvendor problem (see for instance Axsäter 2003). If we define $D_{it}^c = \sum_{s=1}^t D_{is}$, i.e. the cumulative demand at retailer i during stages $1, \dots, t$, then the inventory level at the end of stage t , $t = 1, \dots, M$, is given by $S_i - D_{it}^c$. Holding costs are incurred over the positive part at the end of each stage and backorder costs are incurred over the negative part. Therefore, the expected cost function for retailer i is given by

Figure 1. Graphical illustration of the time-line with costs and events in a period.

Table 1. Notation used in Section 3.

i, j	retailer indices, $i, j = 1, 2$
M	number of stages in a period
t	stage index, $t = 1, \dots, M$
h	holding cost for each item that remains unused at the end of stage t , $t = 1, \dots, M$
b	backorder cost for each demand that cannot be satisfied immediately
F	fixed cost of transshipping $x \neq 0$ units
T	unit transshipment cost
x_t	number of items transshipped from retailer 1 to 2 ($x_t \geq 0$) or from 2 to 1 ($x_t < 0$) at the end of stage t , $t = 1, \dots, M - 1$
S_i	order-up-to level of retailer i
D_{it}	demand at retailer i during stage t
D_{it}^c	cumulative demand at retailer i during stages $1, \dots, t$
R_{it}	residual stock level at the end of stage t , $t = 1, \dots, M$, just before redistribution
S_{it}	stock level at beginning of stage t , $t = 1, \dots, M$, right after redistribution
$E(Y)$	expected value of random variable Y
$f_Y(\cdot)$	probability density function of random variable Y
$F_Y(\cdot)$	cumulative distribution function of random variable Y
$(a)^+$	$\max\{a, 0\}$
$(a)^-$	$\max\{-a, 0\}$

$$\begin{aligned}
C_i(S_i) &= \sum_{t=1}^M [h E(S_i - D_{it}^c)^+ + b E(S_i - D_{it}^c)^-] \\
&= \sum_{t=1}^M \left[h \sum_{y=0}^{S_i} (S_i - y) f_{D_{it}^c}(y) + b \sum_{y=S_i+1}^{\infty} (y - S_i) f_{D_{it}^c}(y) \right].
\end{aligned}$$

The objective is to determine the order-up-to levels S_i that minimise total expected costs. Hence, we need to solve

$$\min_{S_i \geq 0} \sum_{i=1}^2 C_i(S_i),$$

which can easily be done numerically.

3.2 Complete and routed pooling

We first present the complete pooling model. The routed pooling model is obtained after a minor modification.

3.2.1 Complete pooling

Let S_{it} , $t = 1, \dots, M$, denote the stock level of retailer i at the beginning of stage t . Note that $S_{i1} = S_i$. Furthermore, let $R_{it} = S_{it} - D_{it}$, $t = 1, \dots, M$, denote the residual stock level of retailer i at the end of stage t just before redistribution. Note that we can only move stock from retailer i to j , $i \neq j$, at the end of stage t , if $R_{it} > 0$. Finally, let x_t denote the size of the transshipment from retailer 1 to 2 at the end of stage t , $t = 1, \dots, M - 1$ (a negative transshipment from 1 to 2 implies a transshipment of size $-x_t$ from retailer 2 to 1). Holding cost and backorder costs are incurred for the positive and the negative part, respectively, of the inventory level ($S_{it} - D_{it}$) at the end of stage t ($t = 1, \dots, M$). The unit transshipment cost at the end of stage t ($t = 1, \dots, M - 1$) is proportional to the number of items transshipped, $|x_t|$, and the fixed transshipment cost is incurred if $x_t \neq 0$. So we need to minimize

$$\min_{S_{it}, x_t} E \left[\sum_{i=1}^2 \sum_{t=1}^M (h (S_{it} - D_{it})^+ + b (S_{it} - D_{it})^-) + \sum_{t=1}^{M-1} (\mathbf{1}_{x_t \neq 0} F + |x_t| T) \right] \quad (1)$$

$$\text{s.t. } S_{1,t+1} = (S_{1t} - D_{1t})^+ - x_t \quad t = 1, \dots, M-1, \quad (2)$$

$$S_{2,t+1} = (S_{2t} - D_{2t})^+ + x_t \quad t = 1, \dots, M-1, \quad (3)$$

$$S_{it} \geq 0 \quad i = 1, 2, \quad t = 1, \dots, M, \quad (4)$$

$$x_t \in \mathbb{Z} \quad t = 1, \dots, M-1, \quad (5)$$

where

$$\mathbf{1}_{x_t \neq 0} = \begin{cases} 1 & \text{if } x_t \neq 0, \\ 0 & \text{else.} \end{cases}$$

The objective function (1) consists of the sum of expected holding, backorder and transshipment costs. Constraints (2) and (3) ensure that the inventory level at the beginning of stage $t+1$ at retailer 1 (retailer 2) equals the residual stock level of stage t minus the amount that is transshipped to retailer 2 (plus the amount transshipped from retailer 1). Constraints (5) state that the inventory level at the start of each period is non-negative and Constraints (4) that we can only transship an integer number of items. Note that from Constraints (2)–(4) it directly follows that $-(S_{2t} - D_{2t})^+ \leq x_t \leq (S_{1t} - D_{1t})^+$.

We will now present a Dynamic Programming formulation for solving (1)–(5). Let $\mathcal{X}_t = \{x \mid -(S_{2t} - D_{2t})^+ \leq x \leq (S_{1t} - D_{1t})^+, x \in \mathbb{Z}\}$ be the set of feasible transshipments at the end of stage t . Let $C_t(R_{1t}, R_{2t})$, $t = 1, \dots, M-1$, denote the expected cost at the end of stage t , just before transshipping x_t items, when the residual stock levels of retailers 1 and 2 are R_{1t} and R_{2t} , respectively. Then the inventory level that determines the holding and backorder costs at the end of stage $t+1$ (so before a possible redistribution at the end of stage $t+1$) is equal to the residual stock level at the end of stage t (before redistribution), plus/minus an incoming/outgoing transshipment at the end of stage t , minus the demand during stage $t+1$. Hence,

$$\begin{aligned} C_t(R_{1t}, R_{2t}) = & \min_{x \in \mathcal{X}_t} \sum_{i=1}^2 \left(h E \left[(R_{it})^+ + (-1)^i x - D_{i,t+1} \right]^+ + b E \left[(R_{it})^+ + (-1)^i x - D_{i,t+1} \right]^- \right) \\ & + \mathbf{1}_{x \neq 0} F + |x|T + E \left[C_{t+1} \left((R_{1t})^+ - x - D_{1,t+1}, (R_{2t})^+ + x - D_{2,t+1} \right) \right], \end{aligned} \quad (6)$$

where

$$C_M(R_{1M}, R_{2M}) = 0. \quad (7)$$

Equation (6) states that we minimise expected future costs over transshipment quantities $x \in \mathcal{X}_t$ at the end of stage t , $t = 1, \dots, M-1$. The expected future costs at the end of the final stage M are zero (Equation (7)). Finally, let $C_0(S_1, S_2)$ denote the expected cost at the beginning of the period when we determine the optimal order-up-to levels S_i . We have

$$\begin{aligned} C_0(S_1, S_2) = & \min_{S_i \geq 0} \sum_{i=1}^2 \left(h E[S_i - D_{i1}]^+ + b E[S_i - D_{i1}]^- \right) \\ & + E \left[C_1(S_1 - D_{11}, S_2 - D_{21}) \right]. \end{aligned} \quad (8)$$

3.2.2 Routed pooling

In the routed pooling problem, stock can only be transshipped from retailer 1 to 2 at the end of stages $1, \dots, M-1$. Hence, we have

$$x_t \in \mathbb{N}_0 \quad t = 1, \dots, M-1, \quad (9)$$

The routed pooling problem is equivalent to (1)–(4) and (9). The corresponding Dynamic Programming formulation is given by (6)–(8), in which $\mathcal{X}_t = \{x \mid 0 \leq x \leq (R_{1t})^+, x \in \mathbb{Z}\}$.

The Dynamic Programming formulation (6)–(8) can be solved numerically using backward induction. Interested readers are referred to Bertsekas (2012) for a step-by-step explanation of solving Dynamic Programming formulations.

4. Computational experiments

We implemented the Dynamic Programming formulations for complete and routed pooling as well as the no pooling model in software package R 2.15.1 and solved it for a variety of instances. We normalised the number of stages M at 4, the holding cost h at 1, and the backorder cost b at 50. Four stages relates to the practical setting with monthly regular replenishments and

Table 2. Demand distributions considered in Section 4.

		Probability that a demand (at any stage) is of size																
Mean	Variance	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
2	Medium	$\frac{1}{5}$	$\frac{1}{5}$	$\frac{1}{5}$	$\frac{1}{5}$	$\frac{1}{5}$												
	Low	$\frac{1}{46}$	$\frac{2}{46}$	$\frac{4}{46}$	$\frac{8}{46}$	$\frac{16}{46}$	$\frac{8}{46}$	$\frac{4}{46}$	$\frac{2}{46}$	$\frac{1}{46}$								
4	Medium	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$								
	High	$\frac{1}{3}$	0	0	0	$\frac{1}{3}$	0	0	0	$\frac{1}{3}$								
8	Medium	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$	$\frac{1}{17}$

weekly pooling opportunities. It allows us to study the benefits of allowing multiple pooling opportunities whilst avoiding excessive computation times. As for the cost settings, it is well known (see e.g. Axsäter 2006) that the optimum service level for a single-period problem (without pooling) is given by the ratio of the underage cost and overage cost. These are, respectively, $b = 50$ and $b + h = 51$ for the selected cost parameters, implying a ratio of 0.98 and therefore an expected optimal service level of around 98%, which is in the 95–98% range that firms typically aspire to in our experience.

As one of our main contributions is to study the effect of the transshipment cost structure on the optimal transshipment policy, we consider a broad range of values for the fixed and variable transshipment cost.

- $F \in \{0, 1, \dots, 20\}$;
- $T \in \{0, 1, \dots, 50\}$.

Preliminary testing revealed that allowing even larger values does not provide further insights.

We considered five different discrete demand distributions for demand D_{it} at retailer i during stage t . These are chosen such that we can investigate the effect of the size of both the mean demand rate and variability of demand on optimal order-up-to levels, cost benefits of pooling, pooling probabilities and average number of items transshipped. The discrete uniform distribution on the interval $[0, 8]$, with mean demand of four units, serves as the base distribution for all stages. However, we also consider situations where the expected demand is reduced to 2 or increased to 8 for all stages. Furthermore, we consider situations with mean demand of four units, as in the base distribution, but with lower and higher variance of demand respectively. All distributions are presented in Table 2.

The results will be presented in the following order. In Sections 4.1 and 4.2, we discuss the effects of changing the mean and variability of demand, respectively, on the complete/routed pooling policies and their benefits compared to no pooling. Then, in Section 4.3, we zoom in on the benefits of allowing multiple pooling opportunities per period and the influence of the transshipment cost structure, as these are the main contribution of this study.

4.1 Effects of mean and variability of demand

We first consider the effect of mean demand on the optimal policy. In Figure 2(a) the differences between total stock levels in case of no pooling and complete pooling are given for demand distributions with different mean values and $F = 0$. In Figure 2(b) results are presented for the difference between no pooling and routed pooling.

As expected, both differences increase with the unit transshipment cost. We also observe that the stock level reduction under routed pooling is not much less than under complete pooling, indicating that routed pooling performs relatively well. This is in line with the findings of Bouma and Teunter (2014). Stock reductions increase with mean demand. So in absolute terms, the stock reduction achieved by pooling increases with mean demand. It is important to realise, though, that higher mean demand also implies higher stock levels without pooling. Moreover, higher mean demand results in more units being transshipped if pooling is allowed. Therefore, larger stock reductions for higher mean demand do not automatically imply that larger relative cost savings from pooling. In fact, Figure 3(a) and (b) show that the relative cost savings are (almost) independent of mean demand if the fixed transshipment cost is zero. Apparently, the relative increase in number of units transshipped and therefore transshipment cost for higher mean demand is equal to the relative decrease in stock level. This changes if the fixed transshipment cost is positive, as is depicted in Figure 3(c) and (d). Now, the relative increase in cost from transshipping more units is smaller, and the relative savings go up as a result. Figure 3(e) and (f) further demonstrate that this effect is more pronounced if the fixed transshipment cost is larger compared to the unit transshipment cost.

Figure 2. Difference in total order-up-to levels between no pooling and pooling policies, when mean stage demand is either 2, 4 or 8, variability of demand is medium, fixed transshipment cost $F = 0$ and unit transshipment cost $T \in \{0, \dots, 25\}$. (a): Difference between no pooling and complete pooling. (b): Difference between no pooling and routed pooling.

Figure 3. Relative cost savings of complete and routed pooling for medium variability of demand. (a)–(b): Fixed transshipment cost $F = 0$ and unit transshipment cost $T \in \{0, \dots, 50\}$. (c)–(d): $F = 10$ and $T \in \{0, \dots, 50\}$. (e)–(f): $T = 5$ and $F \in \{0, \dots, 20\}$.

4.2 Effects of demand variability

To isolate the effect of demand variability, only results for settings with mean demand constant at 4 (see Table 2) are considered in this subsection. In Figure 4 costs are presented for all three policies (no, routed and complete pooling) and all demand

Figure 4. Optimal expected total cost levels for all pooling policies with mean stage demand $\mu_{it} = 4$ and with low, medium and high variability of demand. (a): Fixed transshipment cost $F = 0$ and unit transshipment cost $T \in \{0, \dots, 50\}$. (b): $T = 5$ and $F \in \{0, \dots, 20\}$.

Figure 5. Order-up-to levels and expected costs if pooling occurs once during a period (at the end of stage 1, 2 or 3) and after each stage 1, 2 and 3. Mean demand $\mu_{it} = 4$ and variability of demand is medium.

variation levels (low, medium, high). This is done for varying levels of the fixed transshipment cost in Figure 4(a) and for varying levels of the unit transshipment cost in Figure 4(b).

It appears from Figure 4 that, for both complete and routed pooling, demand variability does not seem to significantly effect the cost savings that can be achieved from pooling for a wide range of fixed and variable transshipment costs. This is in line with the results of Bouma and Teunter (2014), and therefore not discussed further here. Instead, we next explore the relative performance of routed pooling in relation to the pooling cost structure and the possibility to pool stocks multiple times.

Figure 6. Expected costs if pooling occurs once during a period (at the end of stage 1, 2 or 3) and after each stage 1, 2 and 3. Here, mean demand $\mu_{it} = 4$ and variability of demand is medium.

4.3 Multiple pooling opportunities and the relative performance of routed pooling

We will consider potential pooling at a single moment (at the end of stage 1, 2 or 3) or at all those times. We focus on the base probability distribution with mean $\mu_{it} = 4$ and medium variability of demand. In Figure 5, optimal order-up-to levels and total expected costs are depicted for $F \in \{0, 10\}$ and $T = 0$ for both complete and routed pooling.

Let us first discuss the distribution of stocks over the retailers. For complete pooling (with identical retailers), it is optimal to distribute stocks as evenly as possible at the start, and so the difference in starting stock levels is at most one. For routed pooling, however, creating a larger initial difference by stocking more at retailer 1 increases the probability of levelling stocks, if needed, at transshipment opportunities. Of course, the initial stock at retailer 2 needs to be sufficient (with reasonable probability) to cover all demand until the first pooling opportunity. This explains why, for a single transshipment opportunity, the initial stock difference reduces as the transshipment opportunity occurs later in the replenishment period. By allowing pooling at different moments during the period, creating a large difference in initial stocks is possible without a high stock-out risk at retailer 2. Therefore, the cost savings from allowing multiple transshipment opportunities is larger for routed pooling than for complete pooling. Since these cost savings are achieved by transshipping more units and more often, the effect is dampened if the transshipment cost increases. A comparison of Figure 5(b) and (d) shows that the savings reduce from about 4 to 1% if the fixed transshipment cost is increased from 0 to 10.

Figure 6 shows that increasing the unit transshipment cost also reduces the benefit of allowing multiple transshipments opportunities, and confirms that there are no significant benefits under complete pooling.

5. Conclusion

We considered the routed inventory pooling problem with multiple lateral transshipments. In this single-period problem with two retailers, stock can be redistributed along a fixed route (to facilitate coordination of redistribution for different items) at multiple predetermined points in time. We introduced a Dynamic Programme that minimises total expected inventory and transshipment costs over the order-up-to levels of the retailers and the number of items transshipped at each pooling opportunity. In an extensive numerical study, the routed pooling policy was compared to the no pooling and complete pooling policies with multiple pooling opportunities, as well as to policies with a single opportunity for redistribution.

From the experiments we can conclude that the performance of both complete and routed pooling is strongly affected by transshipment costs. Changes in the unit transshipment cost generally have a greater impact on costs than equivalent changes in fixed transshipment costs, due to the fact that generally more than one item is transshipped for the cases that we considered, which also seems likely for real-life cases.

In comparison with complete pooling, the performance of routed pooling declines rapidly with unit transshipment costs. If transshipments are either free of charge or very cheap, the restrictions of a fixed route for redistribution can be compensated by a great imbalance in initial order-up-to levels; in these cases, routed pooling leads to almost the same amount of cost reductions as complete pooling. However, if transshipment costs increase, there is a tendency towards more balanced order-up-to levels and the performance of routed pooling deteriorates.

An interesting find is that, in case of low transshipment costs, multiple transshipment opportunities do have a positive effect on expected costs in case of routed pooling, but not in case of complete pooling. This relates to the imbalance in stock

levels that is desired in case of routed pooling. Allowing multiple transshipment opportunities makes this possible without a substantial stock-out risk at the retailer at the end of the route with low initial stocks.

A limitation of our research is that we only considered systems with two identical retailers in order to ensure solvability by Dynamic Programming in reasonable time. Previous results by Bouma and Teunter (2014), though under stricter assumptions as discussed in the introduction, do suggest that insights on the benefit of complete and routed pooling over no pooling based on such small systems carry over to somewhat larger systems with three non-identical retailers. However, more research is definitely needed in this direction. Although formulating larger systems as Dynamic Programmes is still possible, solving them is too time-consuming. Dynamic Programming suffers from the so-called curse of dimensionality, as a result of which computation times grow exponentially in for instance the number of stages per period and in the number of retailers. Simulation and approximate Dynamic Programming provide alternative research methodologies.

Other model extensions are also worthwhile to consider. One is to include truck capacities. Related transshipment policies may be considered where some but not all of the retailers are visited – balancing stocks only amongst retailers with very high or very low stocks in order to limit transshipment costs. Such extensions would put more emphasis on the routing aspects of the combined routing and pooling problem, where the main focus of this research was on pooling.

Another avenue of further research is to consider storage restrictions. We ignored these restrictions as, to the best of our knowledge, did all previous contributions on inventory transshipments. In fact, the majority of the inventory literature does not take storage restrictions into account. However, they may obviously exist in practical situations and are especially worth considering if policies under consideration affect the distribution of stocks of various stock-points, as is the case for transshipment decisions.

Finally, an interesting research opportunity is to develop heuristics for finding near optimal replenishment and transshipment decisions that are easy to implement in practice. Exact solutions such as presented in this research provide valuable insights into the behaviour of the optimal transshipment policy, but that policy may be too difficult to determine and implement for some firms. Simpler heuristics may be developed that are more understandable for managers and ideally also more adaptable to specific restrictions that firms may face on e.g. truck capacity, inventory storage capacity or inventory budgets.

Disclosure statement

No potential conflict of interest was reported by the authors.

References

- Agrawal, V., X. Chao, and S. Seshadri. 2004. "Dynamic Balancing of Inventory in Supply Chains." *European Journal of Operational Research* 159 (2): 296–317.
- Allen, S. G. 1958. "Redistribution of Total Stock over Several User Locations." *Naval Research Logistics Quarterly* 5 (4): 337–345.
- Allen, S. G. 1961. "A Redistribution Model with Set-up Charge." *Management Science* 8 (1): 99–108.
- Allen, S. G. 1962. "Computation for the Redistribution Model with Set-up Charge." *Management Science* 8 (4): 482–489.
- Andersson, H., A. Hoff, M. Christiansen, G. Hasle, and A. Løkketangen. 2010. "Industrial Aspects and Literature Survey: Combined Inventory Management and Routing." *Computers & Operations Research* 37 (9): 1515–1536.
- Axsäter, S. 2003. "Evaluation of Unidirectional Lateral Transshipments and Substitutions in Inventory Systems." *European Journal of Operational Research* 149 (2): 438–447.
- Axsäter, S. 2006. *Inventory Control*. Vol. 90. New York: Springer Verlag.
- Bertazzi, L., G. Paletta, and M. G. Speranza. 2002. "Deterministic Order-up-to Level Policies in an Inventory Routing Problem." *Transportation Science* 36 (1): 119–132.
- Bertsekas, D. P. 2012. In *Dynamic Programming and Optimal Control*. 4th ed. Belmont, Massachusetts: Athena Scientific.
- Blumenfeld, D. E., L. D. Burns, J. D. Diltz, and C. F. Daganzo. 1985. "Analyzing Trade-offs between Transportation, Inventory and Production Costs on Freight Networks." *Transportation Research Part B: Methodological* 19 (5): 361–380.
- Bouma, H. W., and R. H. Teunter. 2014. "The Routed Inventory Pooling Problem with Three Non-identical Retailers." *International Journal of Production Economics* 156: 223–234.
- Coelho, L. C., J. F. Cordeau, and G. Laporte. 2012. "The Inventory-routing Problem with Transshipment." *Computers & Operations Research* 39 (11): 2537–2548.
- Federgruen, A., and P. H. Zipkin. 1984. "A Combined Vehicle Routing and Inventory Allocation Problem." *Operations Research* 32 (5): 1019–1037.
- Gross, D. 1963. In *Centralized Inventory Control in Multilocation Supply Systems*. 47–84. Stanford, CA: Stanford University Press.
- Jönsson, H., and E. A. Silver. 1987. "Analysis of a Two-echelon Inventory Control System with Complete Redistribution." *Management Science* 33 (2): 215–227.
- Karmarkar, U. S., and N. R. Patel. 1977. "The One-period, n-location Distribution Problem." *Naval Research Logistics Quarterly* 24 (4): 559–575.

- Miller, D. M. 1987. "An Interactive, Computer-aided Ship Scheduling System." *European Journal of Operational Research* 32 (3): 363–379.
- Paterson, C., G. Kiesmuller, R. Teunter, and K. Glazebrook. 2011. "Inventory Models with Lateral Transshipments: A Review." *European Journal of Operational Research* 210 (2): 125–136.
- Savelsbergh, M., and J.-H. Song. 2007. "Inventory Models with Lateral Transshipments: A Review." *Computers & Operations Research* 35 (7): 2266–2282.
- Tagaras, G., and D. Vlachos. 2002. "Effectiveness of Stock Transshipment Under Various Demand Distributions and Nonnegligible Transshipment Times." *Production and Operations Management* 11 (2): 183–198.