

University of Groningen

Transshipment of cross-channel returned products

Dijkstra, Arjan S.; van der Heide, Gerlach; Roodbergen, Kees Jan

Published in:
International Journal of Production Economics

DOI:
[10.1016/j.ijpe.2017.09.001](https://doi.org/10.1016/j.ijpe.2017.09.001)

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:
2019

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

Citation for published version (APA):
Dijkstra, A. S., van der Heide, G., & Roodbergen, K. J. (2019). Transshipment of cross-channel returned products. *International Journal of Production Economics*, 209(3), 70-77.
<https://doi.org/10.1016/j.ijpe.2017.09.001>

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.



Transshipments of cross-channel returned products



Arjan S. Dijkstra, Gerlach Van der Heide^{*}, Kees Jan Roodbergen

University of Groningen, P.O. Box 800, 9700 AV, Groningen, The Netherlands

ARTICLE INFO

Keywords:
E-commerce
Multi-channeling
Returns
Lateral transshipments
Markov decision processes
Heuristics

ABSTRACT

Companies increasingly employ dual-channeling strategies with online and offline channels to reach customers. The combination of high return rates in e-commerce and the possibility for customers to return products ordered online at any offline store may result in unbalanced inventories. Transshipments can be used to deal with these unbalanced inventories. In this paper we study dynamic policies for transshipment of products that are returned cross-channel from online to offline stores. At the end of each period in a finite sales season, cross-channel returned products can be transshipped back to the online store or kept on-hand at the offline store. Optimal transshipment policies are obtained using a Markov decision process. We introduce a well-performing heuristic based on the expected costs during the sales season, with a maximum deviation of 1.59% from the optimal costs in experiments. Furthermore, we show that in all instances our heuristic outperforms static policies in which products are either always or never shipped back to the online store. We observe that dynamic transshipment policies are more effective than static policies in dealing with imbalances in the initial stock. Dynamic transshipment of cross-channel returns seems to open up possibilities for more effective demand fulfillment of dual-channel companies.

1. Introduction

Dual-channeling is a distribution strategy increasingly applied by business-to-consumer companies in practice (Agatz et al., 2008). A common configuration for dual-channeling uses separate inventories from offline stores and online stores to meet customer demands for products. Products demanded from the online store are sent to the customer from a distribution center. After buying, customers can often return products to the company. Return percentages of as much as 75% have been reported for some product categories in fashion (Mostard and Teunter, 2006). Products are predominantly returned due to either buyer's remorse or an unclear motivation not related to the state of the product (Lawton, 2008). Some companies selling consumer electronics or clothing provide customers the opportunity to return products at any store, regardless of where they were bought originally. These products can be resold in the store they are returned at. The vast majority of cross-store returns are products ordered from the online channel and returned by a customer to a nearby store of the offline channel. In practice, typically all cross-channel returns are shipped back to the distribution center of the online channel, potentially incurring more transportation costs than necessary. On the other extreme, if no products are shipped back, imbalances in the inventories of the two channels may occur.

By carefully coordinating the transshipment of cross-channel returns, companies can increase the availability of products during the sales season. For both types of channels, demand is typically lost to competitors if a customer encounters an out-of-stock situation. Some stock may be unsold at the end of the sales season, incurring costs because products have to be disposed of or sold at a discount. Efficient transshipment policies should determine when the transportation cost weigh up against the costs of unsold products.

In this paper, we study the transshipment of returned products in a dual-channel supply chain for a product that is sold during a single sales season consisting of multiple periods. Sold products return to the store they are sold from with some probability, i.e., returns depend endogenously on fulfilled demand. Moreover, products sold in the online channel return cross-channel to stores of the offline channel with a certain probability. At the end of every period, these cross-channel returns can either be added to the inventory of the offline store, or sent back to the distribution center. Returned products are assumed to be as good as new and can be resold at full price. The goal is to minimize costs during the sales season, which comprise of costs for holding stock, carrying out transshipments, and having unsold stock at the end of the sales season. Using Markov decision processes, we study optimal transshipment policies during the sales season. Furthermore, we formulate a

^{*} Corresponding author.

E-mail addresses: a.s.dijkstra@rug.nl (A.S. Dijkstra), g.van.der.heide@rug.nl (G. Van der Heide), k.j.roodbergen@rug.nl (K.J. Roodbergen).

transshipment heuristic, which we compare to the optimal policy and static policies typically used in practice. Our heuristic is shown to outperform these static policies considerably, showing the potential of dynamically determining the transshipment of returns.

Lateral transshipments can either take place at predetermined moments in time or in reaction to stock-outs. The former are called proactive, whereas the latter are called reactive (Paterson et al., 2011). Recent examples of papers studying reactive transshipments are Axsäter et al. (2013), Howard et al. (2015), and Olsson (2015). Hybrid lateral transshipments, which combine proactive and reactive lateral transshipments are studied by Paterson et al. (2012) and Glazebrook et al. (2015). Since the primary purpose of the transshipment of cross-channel returned products is preventing stock-outs, they are proactive lateral transshipments. Furthermore, as products are transshipped from offline stores to the online store, the transshipments are unidirectional (Axsäter, 2003).

Models studying lateral transshipments consider either a finite or an infinite horizon. Policies for models with a finite horizon mainly focus on situations with a single transshipment opportunity. A number of heuristics have been proposed to determine transshipment quantities in such situations (see, e.g., Jönsson and Silver, 1987; Bertrand and Bookbinder, 1998; Agrawal et al., 2004). Optimal transshipment quantities can be determined for models consisting of a single period with a single transshipment decision (see, e.g., Noham and Tzur, 2014). Our model differs in that we study a situation in which transshipment is possible in every period during the finite horizon. This implies that we cannot determine transshipments by considering the remaining periods after the transshipment in isolation, which is a key characteristic of the previously studied models. In a finite model, allowing multiple transshipment opportunities leads to an optimal policy with a distinct structure (Abouee-Mehrzi et al., 2015). However, it is unclear whether this structure still holds when cross-channel returns are possible and only a part of the inventory can be transshipped. A simulation-optimization approach to obtain a transshipment policy with a fixed threshold levels is proposed by Hochmuth and Köchel (2012). Fixed threshold levels are unlikely to work for our situation, as the number of remaining periods is an important factor in determining whether or not to ship a cross-channel returned product (Abouee-Mehrzi et al., 2015).

In an infinite horizon setting, papers considering multiple transshipment opportunities typically use balancing heuristics, in which stock levels are compared to future demand in some way (Banerjee et al., 2003; Lee et al., 2007). These balancing policies typically do not depend on cost parameters, which can influence their performance (Lee et al., 2007). Liu et al. (2016) show that a myopic rebalancing policy is optimal for a pooled virtual stockpile. However, such a policy is unlikely to be optimal for other transshipment problems (Abouee-Mehrzi et al., 2015). Firouz et al. (2016) use simulation-optimization to solve a stochastic MILP to determine transshipment quantities. None of the above finite and infinite horizon articles consider returns, and extending these models to accommodate for returns is not straightforward.

Returns can be in an as-good-as-new condition, meaning that they are resalable, or they can be damaged, requiring an extensive refurbishing or remanufacturing process. The latter is studied in reverse logistic models (Fleischmann et al., 1997; Tai and Ching, 2014). In our setting the predominant reason for returning are not defects. Therefore, we study resalable returns. Resalable return models have been studied in settings with a single location and multiple locations. Returns are modelled either as an independent exogenous process, or as an endogenous process depending on fulfilled demand. Single location settings with dependent returns include Kelle and Silver (1989), Buchanan and Abad (1998), and Mostard and Teunter (2006). Kiesmüller and Van der Laan (2001) show that the stock processes under independent and depend returns differ substantially, especially in case of high return rates. To the best of our knowledge, papers studying multiple locations with resalable returns only consider independent return processes (see, e.g., Ching et al., 2003; Mitra, 2009). As high return rates are common in practice, in this paper we study a setting with multiple stock locations and a dependent return

process. Since future returns depend on the availability of stock, transshipment decisions should account for this.

The remainder of the paper is organized as follows. In §2, we introduce the model and assumptions. In §3, we formulate an MDP for obtaining the optimal policy. We develop a heuristic in §4 and compare it with the optimal policy and heuristics from practice in §5. Finally, in §6 we provide conclusions and directions for future research.

2. Problem definition

We consider the inventory control of a single product for a dual-channel company which sells through online and offline channels. The online channel consists of one online store (distribution center), indexed $i = 0$, and the offline channel consists of n offline stores, indexed $i = 1, \dots, n$. The product is sold during a sales season with duration T . At the beginning of the sales season (period 1), the stores have initial inventory $I^1 = (I_1^1, \dots, I_n^1)$. Each period t , $t = 1, \dots, T$, store i faces generally distributed non-negative demand D_t^i with mean λ_t^i . Demand in excess of the on-hand inventory is lost.

Each item sold during a period has a probability of being returned in that same period, analogous to Mostard and Teunter (2006). Products returned in a period are resalable in the next period. There are regular returns and cross-channel returns. Regular returns return to the online or offline store from which they were sold, with a probability $0 \leq p_{ii} < 1$ for each sold item at store i , $i = 0, \dots, n$. Cross-channel returns are items sold in the online store and returned to one of the offline stores. A sold item at the online store returns to offline store i , $i = 1, \dots, n$ with probability p_{0i} . Clearly, we require $0 \leq p_{00} + \sum_{i=1}^n p_{0i} < 1$.

Stock levels are reviewed at the beginning of each period t , and are denoted I^t . After each review, transshipments can be carried out. We are allowed to transship (part of) the cross-channel returns at offline store i from the previous period back to the online store. As transshipments are typically carried out overnight, the lead time of transshipments is assumed to be negligible. At the end of the period, demand is observed and fulfilled to the extent possible from on-hand stock. Finally, inventory costs are incurred at the end of the period.

The costs are as follows. Transshipment between store i and store j costs c_{ij} per unit. Clearly, we have $c_{ii} = 0$. Moreover, we have $c_{ij} = \infty$ if $i \neq j$ and $j \neq 0$, which implies that cross-channel returns can only be transshipped to the online store. In a model extension we later relax this assumption and allow transshipments between offline stores. A holding cost h is incurred for each unit on-hand at the end of the period. We do not consider a direct penalty cost for lost demand. Since the goal of the company is to sell as much as possible of the remaining inventory during a finite sales season, instead a penalty s is incurred for each unsold unit of stock by the end of period T . For our purpose of optimizing transshipment policies, unsold inventory and lost demand costs are functionally equivalent, because each extra unit of lost demand prevented by a certain policy results in one less unsold unit at the end of the sales season. Hence, one can take a similar approach to setting s for a practical setting as in standard lost-sales models, see e.g., Zipkin (2008).

We aim to find a transshipment policy that minimizes costs during the sales season. Even though we consider a single sales season, our model extends to the case with replenishments when the replenishment policy and transshipment policy are set independent from each other, as in, e.g., Banerjee et al. (2003) and Lee et al. (2007). Nonetheless, a finite model without replenishments is realistic when fashion companies are considered. In that case, long lead-times lead to single batches being ordered for the entire sales season (Mantrala and Raman, 1999; Mostard and Teunter, 2006; Caro and Gallien, 2012).

3. Markov decision process

In order to solve the problem to optimality, we formulate a Markov decision process (MDP). In what follows we provide the state space, the

admissible actions in each state, and the state transitions.

Since we have to decide which of the cross-channel returns to ship back to the online store, our state variable has to include the on-hand inventory levels as well as the cross-channel returns that are eligible for transshipment. We therefore define the discrete state variable $x^t \in \mathbb{N}^{2n+1}$ in period $t = 1, \dots, T$ as

$$x^t = (I_0^t, I^t, R^t),$$

where $I_0^t \in \mathbb{N}$ is the on-hand stock at the online store, $I^t = (I_1^t, \dots, I_n^t) \in \mathbb{N}^n$ are the on-hand stock levels at the offline stores, and $R^t = (R_1^t, \dots, R_n^t) \in \mathbb{N}^n$ are the cross-channel returns from the previous period available at the offline stores at the start of period t .

In principle, we can start from a given initial state x^1 in period 1 and determine the state space from this given state. If at the start of period 1 there are K items in total, the state space \mathcal{S} consists of all combinations of on-hand stock levels and cross-channel returns summing up to K or less in total, i.e.

$$\mathcal{S} = \left\{ (I_0, I, R) : \sum_{i=0}^n I_i + \sum_{i=1}^n R_i \leq K \right\}.$$

In the special case where the demand distribution at the online store has a bounded support with an upper limit lower than K , fewer states can be considered since $\sum_{i=1}^n R_i$ is bounded by the maximum demand at the online store.

Let an action be a vector a with elements $a_i, i = 1, \dots, n$ denoting the number of products that are transshipped from offline store i to the online store. The set of admissible actions in each state $x^t \in \mathcal{S}$ are

$$A(x^t) = \{a : 0 \leq a_i \leq R_i \text{ for } i = 1, \dots, n\}.$$

The cross-channel returns not transshipped are kept on-hand. The post-action state \hat{x}^t , starting from state x^t and taking action a , is therefore given by

$$\hat{x}^t \equiv \hat{x}^t(x^t, a) = \left(I_0^t + \sum_{i=1}^n a_i, I^t + R^t - a, 0 \right).$$

Letting $\hat{\mathcal{S}}$ be the set of possible post-action states, it is evident that $\hat{\mathcal{S}} \subset \mathcal{S}$. This can be exploited to reduce the number of computations in numerical procedures.

After the action, stochastic transitions occur due to demands and returns. The system makes a transition from a post-action state $\hat{x}^t \in \hat{\mathcal{S}}$ in period t to a state $x^{t+1} \in \mathcal{S}$ in period $t + 1$.

Now we formalize the transitions. Let d_0^t and $d^t = (d_1^t, \dots, d_n^t)$ be the observed demands of the online and offline stores. Let $r^t = (r_{11}^t, \dots, r_{nn}^t)$ be the observed regular returns of the offline stores. Finally, let the regular return of the online store be r_{00}^t and the cross-channel returns to the offline stores $r_0^t = (r_{01}^t, \dots, r_{0n}^t)$. The new state after the transition is

$$x^{t+1} = \left(\hat{I}_0^t - d_0^t + r_{00}^t, \hat{I}^t - d^t + r^t, r_0^t \right).$$

For completeness, below we specify the probabilities of observing a particular outcome of the convoluted demand and return distributions, when starting from state $\hat{x}^t \in \hat{\mathcal{S}}$. The most difficult aspect here is that demand is limited by the on-hand stock, and that returns are limited by the demand.

$$\begin{aligned} P(\min\{D_0^t, \hat{I}_0^t\} = d_0^t, \dots, \min\{D_n^t, \hat{I}_n^t\} = d_n^t, R_{11}^t = r_{11}^t, \dots, R_{nn}^t = r_{nn}^t, R_{00}^t = r_{00}^t, R_{01}^t = r_{01}^t, \dots, R_{0n}^t = r_{0n}^t) \\ = P(R_{00}^t = r_{00}^t, R_{01}^t = r_{01}^t, \dots, R_{0n}^t = r_{0n}^t) \times \prod_{i=1}^n P(R_{ii}^t = r_{ii}^t) \\ \times \prod_{i=0}^n \left(\mathbb{1}_{\{d_i = \hat{I}_i^t\}} P(D_i^t \geq \hat{I}_i^t) + \mathbb{1}_{\{d_i < \hat{I}_i^t\}} P(D_i^t = d_i) \right) \end{aligned}$$

Since each sold product has a fixed probability of returning, the returns of the offline store i , R_{ii}^t , are Binomial(d_i^t, p_{ii}) distributed, $i = 1, \dots, n$. The returns of the online store, $R_{00}^t, R_{01}^t, \dots, R_{0n}^t$, are Multinomial($d_0^t, p_{00}, p_{01}, \dots, p_{0n}$) distributed. Calculating the convolution of these $3n + 2$ random variables is challenging when n increases, limiting the size of possible instances that we can solve.

As terminal cost, we set

$$V^{T+1}(x^{T+1}) = s x^{T+1} \cdot \mathbf{1}$$

as the penalty cost for unsold inventory, where $\mathbf{1}$ is a $(2n + 1)$ vector of ones. Then we can iteratively solve

$$V^t(x^t) = \min_{a \in A(x^t)} \{h x^t \cdot \mathbf{1} + c \cdot a + \mathbb{E}[V^{t+1}(x^{t+1})]\}$$

for all $x^t \in \mathcal{S}$ by backward programming to obtain optimal costs for the MDP.

The MDP has to be adjusted for the model extension with lateral transshipment. In that case an action $a_{ij}, i = 1, \dots, n, j = 0, \dots, n$, denotes the number of cross-channel returns transshipped from store i to store j . The set of admissible actions in $x^t \in \mathcal{S}$ then becomes

$$A(x^t) = \left\{ a : \sum_{j=0}^n a_{ij} = R_i^t \text{ for } i = 1, \dots, n \right\},$$

with action cost $c \sum_{i=1}^n \sum_{j=0}^n a_{ij}$. The post-action state \hat{x}^t follows from transshipment accounting.

In addition, the MDP can be adjusted for situations with non-zero transshipment lead times. Suppose L_{ij} is the transshipment lead time from node i to j . For each store i we can introduce a pipeline in the state variable tracking the aggregate number of transshipments arriving at that store τ periods from now, $\tau = 1, \dots, \max_j \{L_{ji}\}$. Solving this extension is challenging because the state space size increases quickly in the lead times and the number of stores.

4. Transshipment heuristic

Since the MDP from §3 is difficult to solve for realistic instances with more than three stores, we propose a heuristic that applies to instances of any size. The idea for this heuristic is as follows. At the start of a period t , we observe the cross-channel returns at the offline stores. From now on we drop time indices to reduce notational clutter and include them as function arguments where needed. We define $C_i(I_i + 1, I_0, t)$ to be the expected cost of adding a product to the on-hand stock at offline store i for the remainder of the sales season, including possible penalty costs for being unsold. We then select store i with unassigned cross-channel returns that has the lowest value of $C_i(I_i + 1, I_0, t)$. Subsequently, we assign a cross-channel return from this store i to the store j that minimizes the direct shipment costs plus expected future costs. Hence, if $i \neq j$ the cross-channel return is transshipped from store i to store j . The heuristic is summarized formally in Heuristic 1.

When transshipment in the offline channel is allowed, any store can be selected to transship the return to and we substitute $j := \min_{k \in \{0, 1, \dots, n\}} C_k(I_k + 1, I_0, t) + c_{ik}$. The heuristic does not preclude positive lead times between stores. In this case, one can use the long-run costs starting from the moment of delivery of the product at location j , instead of t .

Heuristic 1 Transshipment Heuristic

```

while  $\{i | R_i > 0\} \neq \emptyset$  do
     $i := \min_{\{k | R_k > 0\}} C_k(I_k + 1, I_0, t)$ 
     $j := \min_{k \in \{0, i\}} C_k(I_k + 1, I_0, t) + c_{ik}$ 
     $I_j := I_j + 1$ 
     $R_i := R_i - 1$ 
end while
    
```

The remainder of this section deals with estimating $C_i(I_i, I_0, t)$ and $C_0(I_0, I, t)$. For each cross-channel return assigned to the offline or online store, the expected costs need to include all costs until the end of the sales season, including possible penalties for unsold stock. Since returned products may return again in a future period, we account for this in the expected costs. We split the expected cost into two parts: the costs until the product is sold, and the costs if the product returns again. Since the costs if the product returns again include these same two parts, we will use a recursion to obtain the expected costs. In order to simplify the calculation, we assume that no further transshipments take place in the future.

4.1. Determining the expected costs

The expected costs $C_i(I_i, I_0, t)$ at time t for keeping the I_i 'th product at offline store $i, i = 1, \dots, n$ for the remaining sales season depend on its own stock level and that of the online store. As stated before, we split the expected costs during the $T - t$ remaining periods into the expected costs until being sold for the first time, $C_i^d(I_i, t)$, and the expected costs after being sold for the first time, $C_i^r(I_i, I_0, t)$. Hence,

$$C_i(I_i, I_0, t) = C_i^d(I_i, t) + C_i^r(I_i, I_0, t).$$

The costs $C_i^d(I_i, t)$ consists of expected holding and penalty costs, which depend on the number of periods a product is at a store. Therefore, let $W_i(I_i)$ be the stochastic variable indicating the time until product I_i is sold at store i , that is

$$W_i(I_i, t) = \min \left\{ \tau : \tau \leq T - t : \sum_{u=t+1}^{t+\tau} D_i^u \geq I_i \right\}.$$

If the minimum does not exist, we define $W_i(I_i, t) = T - t + 1$, which corresponds to the event that the product is not sold at all during the sales season. Note that $W_i(I_i, t)$ is a stopping time.

Let $D_i^{(\tau)}(t) = \sum_{u=t}^{t+\tau} D_i^u$ denote the convolution of demand at store i during τ time periods. The distribution of $D_i^{(\tau)}$ is easily obtainable for a number of common distributions, including Normal and Poisson. For other distributions it can be determined numerically. From this convolution we can obtain the distribution of $W_i(I_i, t)$ as

$$P(W_i(I_i, t) = \tau) = \begin{cases} 1 - P(D_i^{(1)} < I_i) & \text{if } \tau = 1, \\ P(D_i^{(\tau-1)} < I_i) - P(D_i^{(\tau)} < I_i) & \text{if } 1 < \tau \leq T - t, \\ P(D_i^{(\tau)} < I_i) & \text{if } \tau = T - t + 1. \end{cases}$$

The expected costs until being sold for the first time (or not at all) are then

$$C_i^d(I_i, t) = h\mathbb{E}\{W_i(I_i, t)\} + (s - h)P(W_i(I_i, t) = T - t + 1). \tag{1}$$

The first part gives the expected holding costs, the second part the penalty costs. The penalty costs include the term $s - h$ because we need to subtract the extra holding cost we counted in case $W_i(I_i, t) = T - t + 1$.

The costs $C_i^r(I_i, I_0, t)$ after a return depend on the period in which the product is returned and the stock level at that time. Conditioning on the period in which the product is returned yields the recursive relation

$$C_i^r(I_i, I_0, t) = \sum_{\tau=1}^{T-t} P(W_i(I_i, t) = \tau) p_{ii} \mathbb{E} \left\{ C_i \left(I_i - \sum_{u=t}^{t+\tau} (d_i^u - r_{ii}^u - r_{0i}^u), I_0 - \sum_{u=t}^{t+\tau} (d_0^u - r_{00}^u), t + \tau \right) \right\}. \tag{2}$$

Due to the dependence between the demand and return processes it is complex to obtain the expectation in (2) exactly.

We will approximate the expectation in (2) by replacing some

quantities with simpler distributions or their expectations. For approximating $\sum_{u=t}^{t+\tau} d_i^u$ and $\sum_{u=t}^{t+\tau} r_{ii}^u$ the following observation is important. Given that product I_i is sold, the other $I_i - 1$ products at the offline store must also have been sold. Therefore, we replace $\sum_{u=t}^{t+\tau} d_i^u$ by I_i . Since we condition on the event that product I_i returns, the remaining $\sum_{u=t}^{t+\tau} r_{ii}^u - 1$ returns correspond to the returns of the other $I_i - 1$ products. Hence, this follows a Binomial($I_i - 1, p_{ii}$) distribution. We approximate $\sum_{u=t}^{t+\tau} d_0^u$ by its expected demand $\min\{I_0, \sum_{u=t}^T \lambda_0^u\}$. As estimates for $\sum_{u=t}^{t+\tau} r_{0i}^u$ and $\sum_{u=t}^{t+\tau} r_{00}^u$ we multiply this expected demand by the return probabilities p_{0i} and p_{00} and round them to the nearest integer. Using this approximation it is possible to recursively determine $C_i(I_i, I_0, t)$ for each online store i .

The costs for products at the online store require a more involved computation, as products bought from this store can return at any online or offline store. The idea is to separately calculate expected costs for a return at each store. Define $C_{0i}^r(I_i, I_0, t)$ as the part of the expected future costs of the I_0 'th product at the online store that can be attributed to returning at store i . Furthermore, define $C_{00}(I_0, t)$ as the expected costs incurred at the online store of the I_0 'th product. Then we compute $C_0(I_0, I, t)$ as

$$C_0(I_0, I, t) = C_{00}(I_0, t) + \sum_{i=1}^n C_{0i}^r(I_i, I_0, t). \tag{3}$$

$C_{00}(I_0, t)$ can be computed analogously to the expected long run costs at the offline stores in (1) and (2).

Now we distinguish between two scenarios for the I_0 'th product returning at offline store i in the future:

1. It returns τ periods from now to offline store i with probability p_{0i} and incurs costs there.
2. It returns τ periods from now to the online store with probability p_{00} and later returns to store i when it is sold again.

Note that in the second scenario, a product can return more than once to the online store before being returned at store i . In the first scenario the expected future costs are $C_i(\cdot, \cdot, \tau)$ as given by 2. In the second scenario, the expected future costs are again $C_{0i}^r(\cdot, \cdot, \tau)$. Therefore, we can now write a similar recursion as above. For $i, i = 1, \dots, n$ we have

$$C_{0i}^r(I_0, I_i, t) = \sum_{\tau=1}^{T-t} P(W_0(I_0) = \tau) \left\{ p_{0i} \mathbb{E} \left\{ C_i \left(I_i + \sum_{u=t}^{t+\tau} r_{ii}^u + r_{0i}^u - d_i^u, I_0 + \sum_{u=t}^{t+\tau} r_{00}^u - d_0^u, t + \tau \right) \right\} + p_{00} \mathbb{E} \left\{ C_{0i}^r \left(I_i + \sum_{u=t}^{t+\tau} r_{ii}^u + r_{0i}^u - d_i^u, I_0 + \sum_{u=t}^{t+\tau} r_{00}^u - d_0^u, t + \tau \right) \right\} \right\}$$

The first part of this equation is the cost for returning immediately to store i , whereas the second part is the expected cost of returning to store i after returning to store 0 at time $t + \tau$ first. The latter includes multiple possible returns due to the recursive nature of the equation. The approximation of the stochastic variables in this equation is similar to the one for the stores of the offline channel. We approximate $\sum_{u=t}^{t+\tau} r_{ii}^u + r_{0i}^u - d_i^u$ as $1 + p_{0i}(I_0 - 1) - (1 - p_{ii})\lambda_i$ and $-1 + I_0 + \sum_{u=t}^{t+\tau} r_{00}^u - d_0^u$ by a Binomial($I_0 - 1, p_{00}$) distributed variable. Finally, we compute $C_0(I_0, I, t)$ as in (3). All of the functions that we derived can be precomputed to save computation time. Note that this computation is efficient only if we decompose the cost function into separate cost functions for each store as we did.

5. Experiments

To test the performance of the heuristic proposed in the previous section, we solve a number of instances to optimality with the MDP from Section 3. We compare these optimal solutions to the solutions found by our heuristic, as well as two static policies. We refer to the policy of

transshipping all cross-channel returned products to the online store as the *Ship All* policy, whereas we refer to the policy of keeping all cross-channel returned products at the offline store as the *Ship None* policy.

We have a common experimental design for $n = 1$ and $n = 2$ offline stores. The objective values of the heuristics are the mean of the simulation of 100,000 sales seasons per instance, using common random numbers. The results are shown in Tables 1–3. The number of periods T are 10 and the penalty s for unsold stock is 50. Holding costs are equal at all stores and set to 1 per unit per period. The transshipment costs c are either low (5) or high (15). The demand rate of the offline stores is kept constant at 1, whereas the online store faces a demand rate of either 1 or 4. The self-return rate p_{ii} for offline store $i, i = 1, 2$ is 0.1. The online store has a total return rate of 0.4 of which either 50% or 100% returns cross-channel to an offline store; for $n = 2$ the cross-channel returns are evenly divided over the offline stores. The cross-channel return rates are in accordance with a recent survey in the USA, in which 60% of surveyed consumers indicated to prefer returning a product bought online to an offline store when given the choice (UPS, 2016). Furthermore, a cross-channel return rate of 100% shows the performance of our heuristic most clearly, as the number of transshipment decisions in a sales season is at its maximum.

Optimal initial stock levels depend on the transshipment policy used. For example, when using the *Ship None* policy it is better to keep less stock in the online channel compared to the *Ship All* policy. Furthermore, forecasting of demand for fashion products before the sales season is complex and may lead to substantial forecast errors (Au et al., 2008). Therefore, we consider two different types of initial stock levels: one with relatively more offline stock and one with more online stock. Additionally, we consider total initial inventory equal to either 90% or 130% of total expected sales during the sales season, representing different types of forecast errors. This results in four different configurations of initial stock. See Appendix A for the exact procedure used.

5.1. Performance of the heuristic

In Table 1 we see that our heuristic has a maximum deviation of 1.29% from the optimal costs, with all other instances below 1% when we consider only one offline store. For two offline stores the results are similar as can be seen in Table 2. The maximum deviation is 1.53%, with all other instances being below 1.1%. The results for two stores and transshipments in the offline channel are shown in Table 3. Here, the maximum deviation is slightly lower at 1.42%, but overall the heuristic's performance is similar to the case without transshipments in the offline channel. In all instances our heuristic beats the static *Ship All* and *Ship None* policies.

In general, the deviations from the optimal solution for the heuristic increase in the cross-channel return rate. This seems logical, as with higher cross-channel return rates, the decision whether to transship cross-channel returns has a higher impact on total costs, inflating the effect of suboptimal decisions. Interestingly, the transshipment cost does not appear to have a large influence on the performance of the heuristic. Hence, the heuristic seems to take these costs into account in an effective way.

The relative difference with the optimal solution of both static policies increases when the demand rate of the online channel increases. Which of the static policies is best depends on the other parameter settings. Clearly, a higher transshipment cost means the cost increases of the *Ship All* policy. By comparing instances which have the same cost and demand parameters but different initial inventories, we see that the initial inventory is the most significant factor for determining the performance of the static policies. When there is relatively more stock in the online channel, *Ship None* typically performs best. In this case, the extra inventory at the online store generates cross-channel returns that can replenish the inventory at the offline stores. With relatively more stock in the offline channel, *Ship All* is better because this gives an option to replenish the online channel. Under the optimal policy, the costs for different initial inventories are typically close to each other, indicating

Table 1
Results for instances with 1 offline store. Optimal objective value and relative difference from optimal by heuristics.

	c	λ_0	λ_1	p_{00}	p_{01}	S_0	S_1	Optimal	Heuristic (%)	Ship None (%)	Ship All (%)
1	5	1	1	0.2	0.2	7	6	137.77	0.14	2.50	22.94
2	5	1	1	0.2	0.2	5	8	134.25	0.11	9.93	4.93
3	5	1	1	0.2	0.2	10	10	416.14	0.14	0.97	4.82
4	5	1	1	0.2	0.2	8	12	418.82	0.10	3.54	1.18
5	5	4	1	0.2	0.2	24	6	264.60	0.46	30.35	18.66
6	5	4	1	0.2	0.2	22	8	277.83	0.27	44.92	5.75
7	5	4	1	0.2	0.2	34	9	814.27	0.20	3.54	8.49
8	5	4	1	0.2	0.2	31	12	827.07	0.12	7.85	2.33
9	5	1	1	0	0.4	7	6	131.27	0.44	8.95	35.38
10	5	1	1	0	0.4	5	8	137.93	0.25	20.39	7.24
11	5	1	1	0	0.4	10	10	415.73	0.13	3.82	7.07
12	5	1	1	0	0.4	8	12	424.99	0.08	8.82	1.74
13	5	4	1	0	0.4	24	6	291.21	1.29	72.72	19.54
14	5	4	1	0	0.4	22	8	307.10	0.83	83.77	6.36
15	5	4	1	0	0.4	34	9	841.79	0.08	16.56	9.21
16	5	4	1	0	0.4	31	12	858.76	0.06	23.79	2.73
17	15	1	1	0.2	0.2	7	6	140.04	0.11	0.83	32.84
18	15	1	1	0.2	0.2	5	8	142.07	0.16	3.87	9.08
19	15	1	1	0.2	0.2	10	10	418.78	0.15	0.34	8.42
20	15	1	1	0.2	0.2	8	12	427.48	0.45	1.44	3.17
21	15	4	1	0.2	0.2	24	6	300.49	0.13	14.78	27.18
22	15	4	1	0.2	0.2	22	8	328.24	0.12	22.67	9.49
23	15	4	1	0.2	0.2	34	9	830.35	0.36	1.54	15.04
24	15	4	1	0.2	0.2	31	12	860.94	0.76	3.61	6.63
25	15	1	1	0	0.4	7	6	139.33	0.41	2.64	51.45
26	15	1	1	0	0.4	5	8	155.57	0.37	6.74	13.21
27	15	1	1	0	0.4	10	10	425.84	0.50	1.36	12.90
28	15	1	1	0	0.4	8	12	446.67	0.67	3.54	4.55
29	15	4	1	0	0.4	24	6	386.89	0.94	30.01	25.25
30	15	4	1	0	0.4	22	8	418.02	0.81	35.01	9.54
31	15	4	1	0	0.4	34	9	911.65	0.20	7.63	16.62
32	15	4	1	0	0.4	31	12	955.80	0.21	11.22	7.30

Table 2
Results for instances with 2 offline stores, without transshipments in the offline channel. Optimal objective value and relative difference from optimal by heuristics.

	c	λ_0	λ_1, λ_2	p_{00}	p_{01}	p_{02}	S_0	S_1	S_2	Optimal	Heuristic (%)	Ship None (%)	Ship All (%)
1	5	1	1, 1	0.2	0.1	0.1	7	7	7	215.23	0.15	2.18	12.41
2	5	1	1, 1	0.2	0.1	0.1	5	8	8	212.65	0.12	5.81	3.47
3	5	1	1, 1	0.2	0.1	0.1	10	11	11	659.71	0.14	0.65	2.36
4	5	1	1, 1	0.2	0.1	0.1	8	12	12	662.48	0.06	2.15	0.81
5	5	4	1, 1	0.2	0.1	0.1	25	7	6	336.92	0.86	17.13	20.43
6	5	4	1, 1	0.2	0.1	0.1	22	8	8	351.40	0.44	28.58	6.26
7	5	4	1, 1	0.2	0.1	0.1	37	9	9	1061.02	0.20	1.04	10.42
8	5	4	1, 1	0.2	0.1	0.1	31	12	12	1065.72	0.25	5.83	2.39
9	5	1	1, 1	0	0.2	0.2	7	7	7	208.80	0.44	6.65	19.87
10	5	1	1, 1	0	0.2	0.2	5	8	8	215.11	0.20	11.75	5.56
11	5	1	1, 1	0	0.2	0.2	10	11	11	659.39	0.15	2.46	3.76
12	5	1	1, 1	0	0.2	0.2	8	12	12	667.69	0.08	5.37	1.32
13	5	4	1, 1	0	0.2	0.2	25	7	6	353.51	1.53	45.68	24.57
14	5	4	1, 1	0	0.2	0.2	22	8	8	376.55	1.04	55.99	7.87
15	5	4	1, 1	0	0.2	0.2	37	9	9	1066.15	0.16	7.89	13.26
16	5	4	1, 1	0	0.2	0.2	31	12	12	1094.82	0.13	18.26	2.94
17	15	1	1, 1	0.2	0.1	0.1	7	7	7	218.20	0.10	0.79	18.53
18	15	1	1, 1	0.2	0.1	0.1	5	8	8	219.96	0.12	2.30	6.45
19	15	1	1, 1	0.2	0.1	0.1	10	11	11	662.61	0.11	0.21	4.61
20	15	1	1, 1	0.2	0.1	0.1	8	12	12	670.87	0.28	0.87	2.13
21	15	4	1, 1	0.2	0.1	0.1	25	7	6	365.25	0.38	8.04	30.05
22	15	4	1, 1	0.2	0.1	0.1	22	8	8	396.31	0.30	14.01	10.78
23	15	4	1, 1	0.2	0.1	0.1	37	9	9	1067.39	0.17	0.43	16.50
24	15	4	1, 1	0.2	0.1	0.1	31	12	12	1097.96	0.72	2.72	5.92
25	15	1	1, 1	0	0.2	0.2	7	7	7	217.83	0.31	2.23	30.20
26	15	1	1, 1	0	0.2	0.2	5	8	8	230.98	0.26	4.08	10.53
27	15	1	1, 1	0	0.2	0.2	10	11	11	669.78	0.35	0.87	7.48
28	15	1	1, 1	0	0.2	0.2	8	12	12	688.58	0.47	2.17	3.27
29	15	4	1, 1	0	0.2	0.2	25	7	6	433.87	1.08	18.70	33.40
30	15	4	1, 1	0	0.2	0.2	22	8	8	476.78	1.06	23.20	12.73
31	15	4	1, 1	0	0.2	0.2	37	9	9	1110.46	0.14	3.58	21.71
32	15	4	1, 1	0	0.2	0.2	31	12	12	1189.90	0.15	8.81	6.77

Table 3
Results for instances with 2 offline stores, with transshipments in the offline channel. Optimal objective value and relative difference from optimal by heuristics.

	c	λ_0	λ_1, λ_2	p_{00}	p_{01}	p_{02}	S_0	S_1	S_2	Optimal	Heuristic (%)	Ship None (%)	Ship All (%)
1	5	1	1, 1	0.2	0.1	0.1	7	7	7	210.90	0.07	4.41	14.65
2	5	1	1, 1	0.2	0.1	0.1	5	8	8	210.88	0.15	6.63	4.33
3	5	1	1, 1	0.2	0.1	0.1	10	11	11	656.58	0.10	1.22	2.89
4	5	1	1, 1	0.2	0.1	0.1	8	12	12	661.08	0.08	2.39	1.02
5	5	4	1, 1	0.2	0.1	0.1	25	7	6	330.68	0.59	19.27	22.63
6	5	4	1, 1	0.2	0.1	0.1	22	8	8	348.83	0.34	29.45	6.98
7	5	4	1, 1	0.2	0.1	0.1	37	9	9	1052.21	0.27	1.79	11.36
8	5	4	1, 1	0.2	0.1	0.1	31	12	12	1063.98	0.16	5.99	2.58
9	5	1	1, 1	0	0.2	0.2	7	7	7	203.45	0.40	9.46	23.06
10	5	1	1, 1	0	0.2	0.2	5	8	8	212.98	0.22	12.90	6.65
11	5	1	1, 1	0	0.2	0.2	10	11	11	655.79	0.16	3.10	4.40
12	5	1	1, 1	0	0.2	0.2	8	12	12	666.01	0.14	5.67	1.64
13	5	4	1, 1	0	0.2	0.2	25	7	6	350.16	1.42	46.94	25.73
14	5	4	1, 1	0	0.2	0.2	22	8	8	374.95	0.87	56.70	8.22
15	5	4	1, 1	0	0.2	0.2	37	9	9	1062.95	0.53	8.15	13.59
16	5	4	1, 1	0	0.2	0.2	31	12	12	1093.87	0.23	18.35	2.97
17	15	1	1, 1	0.2	0.1	0.1	7	7	7	216.05	0.10	1.75	19.69
18	15	1	1, 1	0.2	0.1	0.1	5	8	8	218.87	0.02	2.87	6.99
19	15	1	1, 1	0.2	0.1	0.1	10	11	11	661.15	0.14	0.51	4.88
20	15	1	1, 1	0.2	0.1	0.1	8	12	12	670.18	0.31	0.99	2.29
21	15	4	1, 1	0.2	0.1	0.1	25	7	6	362.02	0.28	8.98	31.18
22	15	4	1, 1	0.2	0.1	0.1	22	8	8	395.01	0.15	14.38	11.11
23	15	4	1, 1	0.2	0.1	0.1	37	9	9	1063.31	0.34	0.80	16.92
24	15	4	1, 1	0.2	0.1	0.1	31	12	12	1097.34	0.71	2.76	5.94
25	15	1	1, 1	0	0.2	0.2	7	7	7	215.03	0.25	3.56	31.95
26	15	1	1, 1	0	0.2	0.2	5	8	8	229.75	0.24	4.78	11.09
27	15	1	1, 1	0	0.2	0.2	10	11	11	668.13	0.43	1.18	7.74
28	15	1	1, 1	0	0.2	0.2	8	12	12	687.82	0.51	2.28	3.43
29	15	4	1, 1	0	0.2	0.2	25	7	6	432.47	1.01	19.03	33.76
30	15	4	1, 1	0	0.2	0.2	22	8	8	476.08	0.73	23.31	12.85
31	15	4	1, 1	0	0.2	0.2	37	9	9	1109.17	0.55	3.67	21.85
32	15	4	1, 1	0	0.2	0.2	31	12	12	1189.66	0.19	8.81	6.77

that a dynamic transshipment policy has the flexibility to cost-efficiently deal with small imbalances in initial inventories, whereas static policies generally perform poorly with small imbalances.

5.2. Transshipment frequency over time

Since our heuristic is dynamic, rather than static, the number of




Fig. 1. Fraction of cross-channel returns transshipped to the online channel each period for instance 4, 12, 20, and 28 with 1 offline store.




Fig. 2. Fraction of cross-channel returns transshipped to the online channel each period for instance 3, 11, 19, and 27 with 1 offline store.

transshipped products varies between periods. Figs. 1 and 2 show the fraction of cross-channel returns transshipped in each period of the sales season for a number of instances from Table 1. The fractions are the average of 100,000 simulation runs per instance. Fig. 1 shows instances with relatively more offline stock. We see that the transshipment frequency declines over time. On the other hand, Fig. 2 illustrates the same instances with relatively more online stock. The fraction of transshipped items is almost always lower than in comparable instances in Fig. 1. The

Appendix A. Initial stock levels

For setting initial stock levels, we used the following procedure. In case of relatively more offline stock, we use

$$I_i^{off} = \gamma \lambda_i \left(1 - \sum_{j=0}^n p_{ij} \right)$$

and round to the nearest integer. The scaling factor γ takes on values 0.9 or 1.3. The total stock resulting from this scaling is roughly 90% and 130% of the expected demand during the sales season. The above assumes that all returns are regular returns and leads to relatively much stock at the offline stores. For the case with relative more online stock, we redivide the $K = \sum_{i=0}^n I_i^{off}$ items according to the demand rates λ_i . We set

$$I_i^{on} = \frac{\lambda_i}{\sum_{j=0}^n \lambda_j} K,$$

and round to the nearest integer. If the total sum does not add up to K , we subtract/add an item from stock levels which have been rounded up/down most.

trend is different: transshipment frequencies are low at first and slightly increase before decreasing again at a later stage. The interpretation is that cross-channel returns are first used to replenish some inventory at the offline store; later the decisions become similar to Fig. 1. In both Figures we can observe that instances with higher shipment cost have lower transshipment frequencies in all time periods. The same holds for the cross-channel return rate; an increase leads to a higher transshipment frequency.

6. Conclusion and discussion

We study dynamic policies for transshipment of products that are returned cross-channel from online to offline stores. At the end of each period in a finite sales season, cross-channel returned products can be transshipped back to the online store or kept on-hand at the offline store. We derive optimal transshipment policies using Markov decision processes and propose a heuristic with a maximum deviation of 1.59% from the optimal costs in experiments. In all instances our heuristic outperforms static policies used in practice, showing that dynamic transshipment policies are more effective than static policies in dealing with imbalances in the initial stock. Dynamic transshipment of cross-channel returns seems to provide possibilities for more effective demand fulfillment for companies with online and offline channels.

Our research indicates a number of interesting avenues for further research. We use a transshipment cost that is linear in the number of items transshipped. It would be interesting to see what transshipment policies are effective in different situations; for example, when a fixed order cost is used. Fixed order costs limit the number of periods in which transshipment is beneficial by incentivising shipment of many returns at the same time. In this case, policies should probably provide a threshold for the minimum number of returns that are transshipped, if any are transshipped at all. Another interesting addition would be the inclusion of emergency transshipments, which would enable fulfillment of demand encountering a stock-out if the other channel still has stock available at a premium. It is unclear how this would affect the lateral transshipment policy.

Acknowledgments

This research is partially funded by the Dutch Institute for Advanced Logistics, Dinalog. We thank the two anonymous referees for their helpful comments.

References

- Abouee-Mehrzi, H., Berman, O., Sharma, S., 2015. Optimal joint replenishment and transshipment policies in a multi-period inventory system with lost sales. *Oper. Res.* 63 (2), 342–350.
- Agatz, N.A.H., Fleischmann, M., Van Nunen, J.A.E.E., 2008. E-fulfillment and multi-channel distribution—a review. *Eur. J. Oper. Res.* 187 (2), 339–356.
- Agrawal, V., Chao, X., Seshadri, S., 2004. Dynamic balancing of inventory in supply chains. *Eur. J. Oper. Res.* 159 (2), 296–317.
- Au, K.F., Choi, T.M., Yu, Y., 2008. Fashion retail forecasting by evolutionary neural networks. *Int. J. Prod. Econ.* 114 (2), 615–630.
- Axsäter, S., 2003. Evaluation of unidirectional lateral transshipments and substitutions in inventory systems. *Eur. J. Oper. Res.* 149 (2), 438–447.
- Axsäter, S., Howard, C., Marklund, J., 2013. A distribution inventory model with transshipments from a support warehouse. *IIE Trans.* 45 (3), 309–322.
- Banerjee, A., Burton, J., Banerjee, S., 2003. A simulation study of lateral shipments in single supplier, multiple buyers supply chain networks. *Int. J. Prod. Econ.* 81, 103–114.
- Bertrand, L.P., Bookbinder, J.H., 1998. Stock redistribution in two-echelon logistics systems. *J. Oper. Res. Soc.* 966–975.
- Buchanan, D.J., Abad, P.L., 1998. Optimal policy for a periodic review returnable inventory system. *IIE Trans.* 30 (11), 1049–1055.
- Caro, F., Gallien, J., 2012. Clearance pricing optimization for a fast-fashion retailer. *Oper. Res.* 60 (6), 1404–1422.
- Ching, W.K., Yuen, W.O., Loh, A.W., 2003. An inventory model with returns and lateral transshipments. *J. Oper. Res. Soc.* 54 (6), 636–641.
- Firouz, M., Keskin, B.B., Sharif, H.M., 2017. An integrated supplier selection and inventory problem with multi-sourcing and lateral transshipments. *Omega* 70, 77–93.
- Fleischmann, M., Bloemhof-Ruwaard, J.M., Dekker, R., Van der Laan, E., Van Nunen, J.A.E.E., Van Wassenhove, L.N., 1997. Quantitative models for reverse logistics: a review. *Eur. J. Oper. Res.* 103 (1), 1–17.
- Glazebrook, K., Paterson, C., Rauscher, S., Archibald, T., 2015. Benefits of hybrid lateral transshipments in multi-item inventory systems under periodic replenishment. *Prod. Oper. Manag.* 24 (2), 311–324.
- Hochmuth, C.A., Köchel, P., 2012. How to order and transship in multi-location inventory systems: the simulation optimization approach. *Int. J. Prod. Econ.* 140 (2), 646–654.
- Howard, C., Marklund, J., Tan, T., Reijnen, I., 2015. Inventory control in a spare parts distribution system with emergency stocks and pipeline information. *Manuf. Serv. Oper. Manag.* 17 (2), 142–156.
- Jönsson, H., Silver, E.A., 1987. Analysis of a two-echelon inventory control system with complete redistribution. *Manag. Sci.* 33 (2), 215–227.
- Kelle, P., Silver, E.A., 1989. Forecasting the returns of reusable containers. *J. Oper. Manag.* 8 (1), 17–35.
- Kiesmüller, G., Van der Laan, E.A., 2001. An inventory model with dependent product demands and returns. *Int. J. Prod. Econ.* 72 (1), 73–87.
- Lawton, C., 2008. The war on returns. *Wall Str. J.* 8 (1).
- Lee, Y.H., Jung, J.W., Jeon, Y.S., 2007. An effective lateral transshipment policy to improve service level in the supply chain. *Int. J. Prod. Econ.* 106 (1), 115–126.
- Liu, F., Song, J.S., Tong, J.D., 2016. Building supply chain resilience through virtual stockpile pooling. *Prod. Oper. Manag.* 25 (10), 1745–1762.
- Mantrala, M.K., Raman, K., 1999. Demand uncertainty and supplier's returns policies for a multi-store style-good retailer. *Eur. J. Oper. Res.* 115 (2), 270–284.
- Mitra, S., 2009. Analysis of a two-echelon inventory system with returns. *Omega* 37 (1), 106–115.
- Mostard, J., Teunter, R., 2006. The newsboy problem with resalable returns: a single period model and case study. *Eur. J. Oper. Res.* 169 (1), 81–96.
- Noham, R., Tzur, M., 2014. The single and multi-item transshipment problem with fixed transshipment costs. *Nav. Res. Logist.* 61 (8), 637–664.
- Olsson, F., 2015. Emergency lateral transshipments in a two-location inventory system with positive transshipment leadtimes. *Eur. J. Oper. Res.* 242 (2), 424–433.
- Paterson, C., Kiesmüller, G., Teunter, R., Glazebrook, K., 2011. Inventory models with lateral transshipments: a review. *Eur. J. Oper. Res.* 210 (2), 125–136.
- Paterson, C., Teunter, R., Glazebrook, K., 2012. Enhanced lateral transshipments in a multi-location inventory system. *Eur. J. Oper. Res.* 221 (2), 317–327.
- Tai, A.H., Ching, W.K., 2014. Optimal inventory policy for a markovian two-echelon system with returns and lateral transshipment. *Int. J. Prod. Econ.* 151, 48–55.
- UPS, 2016. Pulse of the Online Shopper. Available at: <https://solvers.ups.com/ups-pulse-of-the-online-shopper/>. (Accessed 13 April 2017).
- Zipkin, P., 2008. Old and new methods for lost-sales inventory systems. *Oper. Res.* 56 (5), 1256–1263.