



Managing transshipments in a multi-retailer system with approximate policies

Nagihan Çömez-Dolgan^{1*} and Nilgun Fescioglu-Unver²

¹*Bilkent University, Ankara, Turkey;* and ²*TOBB University of Economics and Technology, Ankara, Turkey*

Using the right transshipment policy is important when transshipments are exercised under demand uncertainty. Optimal transshipment policy can be quite complex in a multi-firm system as optimal actions depend on all system variables. Moreover, both how to select requested retailer and how to respond to requests are in question. We introduce simple, close-to-optimal heuristic transshipment policies for multiple retailers. We first show that heuristic policies may perform even better than self-optimal policy, which is explained by Braess's paradox. Then we test the performances of various heuristics with respect to centrally optimal policy. When retailers can observe others' inventory levels, more effective transshipments can be made. Otherwise, a random selection performs quite well. We also observe that although always-accept respond policy is quite close to centrally optimal in small systems, the performance of pairwise-optimal holdback levels to respond requests is more clear and consistent for larger systems.

Journal of the Operational Research Society (2015) **66**(6), 947–964. doi:10.1057/jors.2014.59

Published online 18 June 2014

Keywords: inventory sharing; decentralized system; heuristic transshipment policy; Braess's paradox

1. Introduction

The difficulty and importance of supply-demand match lead both practitioners and researchers to investigate alternative mitigation strategies to support the regular inventory management decisions. One of these well-applied mitigation methods is using transshipments, which is basically inventory sharing among the same echelon players in a supply chain. Transshipments are observed at various supply chain levels and for different product categories such as in retail level to satisfy customer demand (Rudi *et al.*, 2001; Özdemir *et al.*, 2006), among manufacturing facilities to satisfy spare part demand (Kranenburg and van Houtum, 2009), and among suppliers to swap commodity products (Kosansky and Schaefer, 2010). In this study, without loss of generality (wlog), we call the firms using transshipments as retailers. The main advantages of transshipments compared with trades between supply chain members in different echelons are shorter lead times, lower transportation costs, and rebalancing of inventories to decrease shortage and inventory costs simultaneously.

With the wider appreciation of transshipments in practice, the questions on the more effective use of transshipments are raised. The basic questions include, but not limited to, effects of using transshipments on optimal stocking levels, when to send transshipments and when not, how to share transshipment profits or how to dynamically negotiate for them, and how to

coordinate ordering levels of retailers that are using transshipments. However, most of these questions are studied on two-retailer systems relying on the easier analytical tractability of the systems compared with a multi-retailer system.

There are some aspects of the transshipment problems, which are not existing in a two-retailer setting, but relevant and crucial in a multi-retailer setting. First, in a two-retailer system, transshipments are in question when there is only and exactly one retailer with on-hand inventory, which is called the sole requested retailer. Thus, the question of from whom to ask for a transshipment is added as an additional decision in a multi-retailer system. Second, it is shown in the literature that the optimal transshipment sending decision between two-retailers requires real-time information not only on sender and receivers, but also on all members of a multi-retailer system, which complicates the definition of the transshipment policy (Archibald, 2007). Thus, both the optimal requesting and transshipment respond decisions cannot be easily stated as a closed-form policy when there are more than two retailers in the systems.

Here we study easily implementable transshipment policies for independently managed multi-retailer systems, which will increase the expected profits of retailers compared with no-transshipment policy. We study a system, in which transshipments exercised immediately after a customer arrival whose demand cannot be satisfied at the originally visited retailer. We consider the possibility of the overflow of the demand of this customer to another retailer if a transshipment cannot be obtained by the visited retailer. Such a setting helps us to model

*Correspondence: Nagihan Çömez-Dolgan, Faculty of Business Administration, Bilkent University, Bilkent, Ankara 06800, Turkey.

E-mail: comez@bilkent.edu.tr

the transshipments that are highly close to those in practice as a stocked-out retailer is allowed to do transshipment at different points in time and with different retailers.

For this purpose, we introduce eight heuristic policies. Each heuristic policy defines a heuristic transshipment requesting retailer selection decision for the stocked-out retailer and a transshipment respond policy for the requested retailer. The selection of requested retailer is based mainly on the availability of the real-time inventory state information in the system. If the real-time inventory levels of the retailers are available, the requested retailer can be selected among the retailers with on-hand inventory as a function of the inventory levels. For such a case, we introduce two different functions of inventory levels for the requested retailer selection. If only the set of retailers with on-hand inventory can be observed, but not the inventory levels, then three different requested retailer selection methods are introduced that do not account for the inventory levels.

For the requested retailers, the decision of whether to accept/reject a request, in other words respond policy, is handled by two different heuristic policies. First, it is shown in the literature (Archibald *et al.*, 1997; Çömez *et al.*, 2012b) that in a two-retailer system, when transshipments are handled one by one after each stock-out, then the optimal transshipment accept/reject decision of a requested retailer can be defined by optimal holdback inventory levels. Thus, in each small time period during a selling cycle, if the actual inventory level is greater than the optimal holdback level of the period, the request is accepted, otherwise, the request is rejected. Thus, a requested retailer may be advised to not send a transshipment even if the available inventory level is positive. Departing from these results, we introduce a heuristic transshipment respond policy for multi-retailer systems, where a requested retailer responds to a request by using her pairwise optimal holdback level from a two-retailer system that includes only herself and the requesting retailer. Second, we introduce a relatively simpler policy of always accepting transshipment requests for a requested retailer to benchmark against both the pairwise optimal holdback level-based policy and also the extreme policy of no-transshipment. In summary, by using five different requested retailer selection heuristics and two respond heuristics, we introduce and test the performance of eight heuristic transshipment policies.

We first test the performance of one of the heuristic policies with respect to (wrt) the solution in which each stocked-out retailer makes his optimal requested retailer selection decision and each requested retailer responds to a request optimally for herself, which is called self-optimal transshipment policy. The results show that using heuristic transshipment policy may even lead to an increase in total profits compared with the case where each retailer uses his/her best policy. The dominance of the heuristic policy over the self-optimal policy is surprising, but it can be explained by 'Braess's Paradox'. When self-optimal transshipment policies are used and resulting equilibrium order levels are obtained, this equilibrium is not necessarily Pareto efficient, which means that at least one retailer can be better off

without anyone else being worse off. Thus, the inefficiency of the equilibrium may result in decreases in total retail profits. To our knowledge, such an over efficiency of heuristic transshipment policies wrt the self-optimal policy within a multi-retailer system is shown for the first time in the literature.

Given that heuristic policy may perform better than self-optimal policies, next we measure the performance of the heuristic policies wrt the profits obtained by using centrally optimal transshipment policy. Besides, we also measure the profit losses due to no-sharing of inventories so that the profit gaps due to the use of heuristic policies can be better appreciated. The heuristic performances are numerically tested in systems including 3–10 retailers based on randomly generated problem parameters that follow car industry-based parameter distributions.

In the overall, the heuristic policies, in which the requested retailer selection is a function of the inventory levels of other retailers and the request respond is given according to pairwise holdback levels, perform closer to the centrally optimal policy. Although the policies using always-accept respond policy perform fairly good for small systems, when the number of retailers increases, they perform poorly. Because as each available unit becomes more valuable in a larger system due to a higher chance to send transshipment later in time, so holdback level-based transshipment policies prove their performance in these larger systems.

In this study, we also test the performance of random selection of a requested retailer. The results indicate that while the random selection is worse than the selection based on actual inventory levels, it performs quite close to those that are using partially observable inventory information such as salvage values or demand rates of retailers with on-hand inventory. In conclusion, we can claim that when the inventory levels of all retailers are not fully observable, selecting the requested retailer randomly and responding requests using pairwise optimal holdback-levels result in fairly close profits to the centrally optimal policy. For example, for a 10-retailer system, it results in around 2% profit decrease, while not using transshipments decreases the costs by more than 30% on average.

The rest of the paper is structured as follows. Section 2 reviews the literature and our contributions. The studied model is explained in Section 3. Section 4 introduces the heuristic policies investigated. The extensive numerical tests are reported and their results are discussed in Section 5. Section 6 provides some concluding remarks. Details of some numerical analyses can be found in the e-companion appendix.

2. Related literature and our contributions

Earliest studies on transshipments are Krishnan and Rao (1965) and Gross (1963). In former, all transshipments are initiated by the end of a season after all demand realization. So transshipments are allowed for all unsatisfied demand as much as the overstock is enough to cover. In latter, the ordering from the

manufacturer and transshipments, in other words, inventory redistribution among retailers are done before the demands are realized. Das (1975) extends the model of Gross (1963) by allowing transshipment in the middle of each period for the redistribution of inventories. In many follow-up studies with similar one-time transshipment opportunity models, the transshipments are analysed in terms of their effect on stocking decisions either in centrally (Robinson, 1990; Tagaras and Cohen, 1992) or decentrally managed settings (Hu *et al.*, 2007; Zhao and Atkins, 2009).

When the transshipments are considered before the total demand realization, the right transshipment policy raises as a question. Archibald *et al.* (1997) analyse such transshipment decisions for a two-retailer centralized system and prove that optimal transshipment respond policy of the stocked-out retailer can be explained by optimal holdback levels, which decrease in time. In a similar model, except that unsatisfied demand can be backordered but not lost, Çömez *et al.* (2012a) show that holdback levels increase during a replenishment cycle until the next replenishment. Çömez *et al.* (2012b) obtain the optimality of holdback level-based transshipment policy in a decentralized two-retailer system. Zhao *et al.* (2008) study optimal production and transshipment policies for centrally managed two-production facilities that make one-for-one production.

When there are more than two firms to exercise transshipments, it is quite complicated to handle the transshipment requesting and sending policies both in centralized as well as decentralized systems. Thus, there are only studies considering heuristic policies for in-season transshipment decisions. Using a centrally managed model, Archibald (2007) introduces a heuristic policy where the requested retailer is selected according to descending transshipment price and the transshipment requests are responded using pairwise optimal holdback levels. In a similar setting, Archibald *et al.* (2009) apply a decomposition method to split a multi-retailer system into two-retailer systems and select the requested retailer according to the value functions that are determined by the decomposition method. Van Wijk *et al.* (2012) analyze the transshipments among spare part depots by assuming holdback levels at each depot to respond to others' requests. The aim of this study is to develop a method to evaluate the approximate performance of the system under random holdback levels.

Study of transshipment policies for decentrally managed multi-retailer systems is quite scarce. Çömez *et al.* (2012b) introduce and test the performance of a heuristic transshipment policy, in which the requested retailer is selected according to highest ratio of inventory level to demand rate and pairwise optimal holdback levels are used for responding. Huang and Sošić (2010) work with a single season model of multiple decentralized retailers. They consider that all transshipments are handled by the end of the season, after all demand is realized and assume complete sharing of residual inventories. The decision of how to match overstocked and understocked retailers is handled by some heuristics, which are tested by using simulation.

One other stream of research on decentralized transshipment systems investigates coordinating transshipment policies. They aim to seek transshipment mechanisms that can lead independent retailers to make system-optimal ordering decisions from the supplier at the beginning of selling season. In Anupindi *et al.* (2001), transshipments are cooperatively made at the end of the season and end-of-season profit allocations to coordinate orders are found. Extending this model, Sošić (2006) considers partial sharing of inventories. Using a cooperative game framework, Huang and Sošić (2010) compare the performances of profit allocations and constant transshipment prices to coordinate retailers' decisions. Hu *et al.* (2007) derive conditions for the existence of coordinating transshipment prices. Rong *et al.* (2010) conclude that with the preventive transshipments, there are no coordinating transshipment prices. Hanany *et al.* (2010) introduce a retailer-paid fund to subsidize the cost of transshipments. Yan and Zhao (2011) propose a mechanism to collect participation fees from retailers before the season to subsidize end-of-season transshipments. Li *et al.* (2012) study a supplier-facilitated transshipment mechanism to achieve coordination, in which retailers are offered to buy more inventory or sell back excess in the middle of the selling season.

This study contributes to the literature mainly by studying various heuristic transshipment policies for multi-retailer systems with the objective of highlighting easy-to-implement policies for independent retailers. Different than most of the past studies, first, transshipments are exercised one by one individually for each stock-out, not at once after all demand realization. Second, we show that interestingly heuristic policies can perform better than the self-optimal heuristic policies, which can be explained by Braess's paradox. Third, we report that when inventory level information of retailers are not fully observable in the system, selecting the requested randomly performs considerably well compared with a more educated selection method relying on retailer characteristics. Fourth, our numerical results are based on expected profit calculation, which include up to 10 retailers. Past studies mostly rely on simulation when the number of retailers in the investigated system exceeds 5.

3. Transshipment model for multiple retailers

We study a decentralized system of M retailers, who receive inventory from a manufacturer at the beginning of a sales season. Each retailer determines its order quantity at the beginning of the sales season. Retailer i purchases Q_i units from the manufacturer, each costing c . We develop a discrete time model by dividing the sales season into N short decision periods such that transshipment per need can be studied. The objective of each retailer i is to maximize his/her total expected profit, which is the expected operating profit obtained throughout the season from the sales $V_i^N(Q)$ minus the cost of inventory purchased cQ_i , where $Q = (Q_1, Q_2, \dots, Q_M)$ is the vector of

beginning of season inventory levels at all retailers. The expected net profit of retailer i is denoted as $J_i(Q)$, which is optimized as follows.

$$\max_{Q_i} J_i(Q) = V_i^N(Q) - cQ_i. \tag{1}$$

It is clear that as the total expected profit of each retailer depends on the stocking levels of others, the solution to obtain the initial stocking levels of retailers requires solving a game on stocking levels.

The number of decision periods remaining until the end of the sales season is denoted by n , $0 \leq n \leq N$. The periods are short enough so that there can be at most one unit demand in each period; with probability p_i at retailer i or at none of them with probability p_0 , where $\sum_{i=0}^M p_i = 1$. As N increases by a factor, and p_i decreases by the same factor, the demand at retailer i converges to independent Poisson process with mean Np_i . Time discretization is common in the literature to study system responses to an individual customer demand, for example, Talluri and van Ryzin (2004), Agrawal *et al* (2004), Wee and Dada (2005), and Çömez *et al* (2012a).

Let x_i denote the inventory level of retailer i in a period during the season and $x = (x_1, x_2, \dots, x_M)$ be the vector of inventory levels. Let $\mathcal{P}(x)$ denote the set of indices for retailers with positive inventory levels, where $|\mathcal{P}(x)|$ is the number of retailers with positive inventory. In period n , if all retailers have available on-hand (positive) inventory and if retailer i receives a demand, he immediately satisfies his demand and earns a unit sales price r . If retailer i does not receive a customer demand, he has no cost or revenue. Therefore, when all retailers have positive inventory, the total expected operating profit of retailer i in period n and all remaining periods until the end of the season can be written as follows, for $i \in \{1, \dots, M\}$.

$$V_i^n(x) = p_0 V_i^{n-1}(x) + p_i [r + V_i^{n-1}(x - e_i)] + \sum_{m=1, m \neq i}^M p_m V_i^{n-1}(x - e_m), \quad |\mathcal{P}(x)| = M. \tag{2}$$

Above e_i denotes a unit vector whose i th element is one while all others are zero.

When all retailers are stocked-out ($x = \{0\}$), each received demand is lost. With zero inventory in stock, there is no change in a retailer's operating profit from period n to $n - 1$.

$$V_i^n(x) = V_i^{n-1}(x) = \dots = V_i^0(x) = 0, \quad |\mathcal{P}(x)| = 0. \tag{3}$$

At the end of the sales season, if there is any left-over inventory unsold at retailer i , it is sold at unit salvage price s_i to obtain the profit

$$V_i^0(x) = s_i x_i. \tag{4}$$

Transshipments may occur only when a retailer is stocked-out. If retailer i has no inventory to satisfy his

demand ($x_i = 0$) in period n , he (*requesting retailer*) sends a transshipment request to another (*requested*) retailer to transship a unit inventory. If the request is accepted, the requested retailer charges transshipment price t to the requesting retailer. No economies of scale are considered in transshipment costs. Wlog, the cost of transportation τ is paid by the requesting retailer. If the stocked-out retailer cannot receive a transshipment to satisfy his waiting customer, the unsatisfied customer visits retailer m with overflow probability θ_m . With probability $1 - \sum_{m=1, m \neq i}^M \theta_m$, the customer leaves the system of retailers without visiting any other retailer. In fact, an unsatisfied customer of the requesting retailer may visit several other retailers until she finds a retailer with on-hand inventory, but because of difficulty of tracking all possible customer overflows, here we restrict the overflow of an unsatisfied customer to only one retailer, who may be out of stock or has on-hand inventory. In Section 5.3, we discuss this assumption further and report numerical analyses on allowing a second request option.

Let \mathcal{R}_{mi}^n be the indicator variable regarding whether retailer i asks for a transshipment from retailer m . It takes a value of 1, if a transshipment request is made to retailer m and 0 otherwise. The response of the requested retailer m is also denoted by an indicator variable called \mathcal{A}_{mi}^n . \mathcal{A}_{mi}^n takes a value of 1, if retailer m accepts the transshipment request of retailer i and a unit is sent from retailer m to retailer i . It takes value 0 otherwise. Similarly, $\bar{\mathcal{A}}_{mi}^n$ is the complementary variable such that $1 - \bar{\mathcal{A}}_{mi}^n = \mathcal{A}_{mi}^n$. Then, the expected operating profit of retailer i when he is stocked-out is as follows.

$$V_i^n(x) = p_0 V_i^{n-1}(x) + \sum_{m \in \mathcal{P}(x)} p_m V_i^{n-1}(x - e_m) + p_i \sum_{m \in \mathcal{P}(x)} \mathcal{R}_{mi}^n \left[\mathcal{A}_{mi}^n [r - t - \tau + V_i^{n-1}(x - e_m)] + \bar{\mathcal{A}}_{mi}^n \left[\sum_{k \in \mathcal{P}(x)} \theta_k V_i^{n-1}(x - e_k) + \left(1 - \sum_{k \in \mathcal{P}(x)} \theta_k \right) V_i^{n-1}(x) \right] \right] + \sum_{\substack{j \notin \mathcal{P}(x) \\ j \neq i}} p_j \left[\sum_{m \in \mathcal{P}(x)} \mathcal{R}_{mj}^n \left[\mathcal{A}_{mj}^n V_i^{n-1}(x - e_m) + \bar{\mathcal{A}}_{mj}^n \left[\sum_{k \in \mathcal{P}(x)} \theta_k V_i^{n-1}(x - e_k) + \left(1 - \sum_{k \in \mathcal{P}(x)} \theta_k \right) V_i^{n-1}(x) \right] \right] \right], \tag{5}$$

If $|\mathcal{P}(x)| = 0$, (5) reduces to (3).

When retailer i has inventory on-hand, but there is at least one other retailer, which is out of stock, then retailer i has a chance to receive a revenue of t by sending a transshipment or r

by receiving an overflowed customer from a stocked-out retailer.

$$\begin{aligned}
 V_i^n(x) = & p_0 V_i^{n-1}(x) + \sum_{\substack{m \in \mathcal{P}(x) \\ m \neq i}} p_m V_i^{n-1}(x - e_m) \\
 & + p_i [r + V_i^{n-1}(x - e_i)] \\
 & + \sum_{j \notin \mathcal{P}(x)} p_j \left[\mathcal{R}_{ij}^n \mathcal{A}_{ij}^n (t + V_i^{n-1}(x - e_i)) \right. \\
 & + \sum_{\substack{m \in \mathcal{P}(x) \\ m \neq i}} \mathcal{R}_{mj}^n \mathcal{A}_{mj}^n V_i^{n-1}(x - e_m) \\
 & + \left. \left(\mathcal{R}_{ij}^n \bar{\mathcal{A}}_{ij}^n + \sum_{\substack{m \in \mathcal{P}(x) \\ m \neq i}} \mathcal{R}_{mj}^n \bar{\mathcal{A}}_{mj}^n \right) \left(\theta_i (r + V_i^{n-1}(x - e_i)) \right. \right. \\
 & + \left. \left. \sum_{\substack{k \in \mathcal{P}(x) \\ k \neq i}} \theta_k V_i^{n-1}(x - e_k) + \left(1 - \sum_{k \in \mathcal{P}(x)} \theta_k \right) V_i^{n-1}(x) \right) \right], \\
 & i \in \mathcal{P}(x). \tag{6}
 \end{aligned}$$

If $|\mathcal{P}(x)| = M$, (6) reduces to (2). Note that both in (5) and (6), $\sum_{m \in \mathcal{P}(x)} \mathcal{R}_{mi}^n = 1$, that is, the stocked-out retailer i requests a transshipment from exactly one of the retailers with on-hand inventory.

For a stocked-out retailer j , the profit from satisfying a demand with a transshipment from retailer i is $r - t - \tau$. When $r - \tau > t$, receiving a transshipment is profitable. When $r - \tau = t$, despite zero profit of satisfying the demand at the requesting retailer, the transshipment can be executed to increase the customer service level. On the other hand, the requested retailer i may send a transshipment to earn t or reject, expecting to sell a unit to the overflowing customer at r with probability θ_i . With probability $\bar{\theta}_i = 1 - \theta_i$, the unsatisfied customer does not visit retailer i , then in the worst case, the unit inventory may be kept on-hand until the end of the season and salvaged at s_i . So retailer i does not accept any transshipment request if $t < \theta_i r + \bar{\theta}_i s_i$. For transshipments to be feasible, parameters should satisfy $\theta_i r + \bar{\theta}_i s_i \leq r - \tau$ for $i \in \{1, \dots, M\}$. We assume that all retailers know the parameters of other retailers, which is similarly modelled in Rudi *et al* (2001) and Hu *et al* (2007).

When retailers use their self-optimal transshipment policies, then stocked-out retailer j determines from whom to ask for a transshipment, \mathcal{R}_{ij}^n and the requested retailer i determines whether to accept a request or not, \mathcal{A}_{ij}^n by maximizing their respective expected profits, for $i, j \in \{1, \dots, M\}$, $i \neq j$, $i \in \mathcal{P}(x)$, and $j \notin \mathcal{P}(x)$. When the system inventory level is x , then the

requested retailer i determines \mathcal{A}_{ij}^n in period n as follows.

$$\mathcal{A}_{ij}^n = \begin{cases} 1 & \text{if } t + V_i^{n-1}(x - e_i) \geq \theta_i (r + V_i^{n-1}(x - e_i)) + \\ & \sum_{\substack{m \in \mathcal{P}(x) \\ m \neq i}} \theta_m V_i^{n-1}(x - e_m) + \left(1 - \sum_{m \in \mathcal{P}(x)} \theta_m \right) V_i^{n-1}(x); \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

Given the information symmetry in the system, then a stocked-out retailer j determines from which retailer to request a transshipment as follows.

$$\mathcal{R}_{ij}^n = \begin{cases} 1 & \text{if } i = \arg \max_{m \in \mathcal{P}(x)} \mathcal{A}_{mj}^n (r - t - \tau + V_j^{n-1}(x - e_m)) \\ & + \bar{\mathcal{A}}_{mj}^n \left(\sum_{k \in \mathcal{P}(x)} \theta_k V_j^{n-1}(x - e_k) + \left(1 - \sum_{k \in \mathcal{P}(x)} \theta_k \right) V_j^{n-1}(x) \right); \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

For a two-retailer system, \mathcal{R}_{ij}^n becomes irrelevant as there can be only one other retailer to ask for a transshipment for the stocked-out retailer. Çömez *et al* (2012a) show that in a two-retailer system, a requested retailer's optimal accept/reject decision can be defined by hold-back levels. For a requested retailer there exists a single value of optimal hold-back level in each period. Then the requested retailer accepts a transshipment request in period n if his actual inventory level is greater than his optimal hold-back level for the period. Otherwise, the request is rejected. However, for a centralized system, Archibald (2007) shows that when there are more than two retailers in the system, the requested retailer's optimal accept/reject decision depends on his inventory level as well as the inventory levels of the other (other than the requested retailer) retailers. Thus, the optimal transshipment request response decision cannot be defined by simple hold-back level variables, but can be complicated functions of the system inventory level. In our multi-retailer decentralized transshipment system, we can also show that both \mathcal{R}_{ij}^n and \mathcal{A}_{ij}^n are in fact functions of inventory levels x , so should more rigorously be denoted as $\mathcal{R}_{ij}^n(x)$ and $\mathcal{A}_{ij}^n(x)$.

Although the self-optimal requested retailer selection and request accept/reject decisions of retailers cannot be defined as closed-form policies, the optimal expected profits of retailers and equilibrium stocking levels can be calculated numerically for given parameter values. The pseudocode for calculating expected profits and equilibrium ordering levels under self-optimal transshipment policies is provided in the online companion. Although theoretically it is possible to make the computation for systems with arbitrary number of retailers, the computation time significantly increases with the number of retailers in the system. Because, as M increases, the number of \mathcal{R}_{ij} and \mathcal{A}_{ij} to compute in each period n increases with the second power of M . Second, to find the equilibrium ordering levels of retailers, for every possible pre-season order quantity at every retailer, the expected payoffs at all retailers should be computed. For an N -period problem, the maximum possible

order quantity at a retailer can be N . Thus, for a single problem instance, this leads to N^M initial order quantity states, which is a very large number for example for $M=10$ retailers and $N=30$ periods. Thus, it does not seem practical for each retailer to analyse her self-optimal transshipment policy and find equilibrium order quantity. Therefore, in the following section, we introduce various easy-to-implement heuristic transshipment policies for independent retailers, some of which require real-time information on the inventory states of other retailers, while the rest can be used even under limited inventory information.

4. Approximate heuristic policies

In a transshipment process, there are two decisions that should be made by the retailers. Retailer with stock-out should select which retailer to request the transshipment from and the requested retailer should decide whether to accept the transshipment request or not. Under self-optimal transshipment policy, each stocked-out retailer makes the requested retailer selection to optimize his own profit, while assuming that a requested retailer is also rational and would make her accept/reject decision to maximize her own operating profits. However this policy is a computationally complex procedure and cannot be executed in a short time either by a requested retailer or a requested one without using computers. To encourage the use of transshipments to benefit both the stocked-out retailers and retailers with on-hand inventory, it would be beneficial to introduce policies that can be determined in-advance of transshipments, so that they can be quickly utilized when needed.

We introduce and investigate the performance of heuristic policies both to select the requested retailer and also to accept/reject requests for the requested retailer. For the requested retailer's accept/reject decision, we introduce two heuristic policies. First is based on the optimal policy for a two-retailer system. In a two-retailer system, $\mathcal{A}_{ij}^n(x)$ defined in (7) can be simplified such that

$$\mathcal{A}_{ij}^n = \begin{cases} 1 & \text{if } x_i > \bar{x}_i^n; \\ 0 & \text{otherwise.} \end{cases}$$

Here \bar{x}_i^n is the optimal holdback level for retailer i in period n , which can be obtained as $\bar{x}_i^n := \max\{x_i \geq 0 : t - \theta_i r < (1 - \theta_i)(V_i^{n-1}(x) - V_i^{n-1}(x - e_i))\}$. Thus, if the other retailer asks for a transshipment from retailer i in period n , then retailer i accepts the request if her current inventory level x_i is greater than her holdback level ($x_i > \bar{x}_i^n$). Otherwise, retailer i rejects the request. Optimal holdback levels \bar{x}_i^n can be computed in-advance by considering the expected operating profit functions and can be used during a season when the actual transshipment request is received and by comparing with the actual inventory level of the requested retailer. Thus, in a two-retailer system, N optimal holdback levels are enough to define the self-optimal accept/reject policy of a retailer. Departing from this simple policy, we introduce the use of optimal holdback levels from two-retailer system as a heuristic

accept/reject policy for the requested retailers in a multi-retailer system. For this purpose, for each retailer i to respond the transshipment requests of retailer j , the holdback levels \bar{x}_{ij}^n are calculated as if retailers i and j operate in a two-retailer system. So for a problem instance with M retailers and N periods, $M(M-1)N$ holdback levels are obtained by solving $M(M-1)/2$ two-retailer problems and calculating $2N$ optimal holdback levels for each two-retailer problem.

A second heuristic policy for the requested retailer's accept/reject decision is to always accept the requests. In the literature, such a policy is generally called as complete pooling or complete sharing. Complete sharing is optimal in a decentralized setting when all transshipments are done by the end of a season after all demand is realized and fixed transshipment prices are used (Hanany *et al*, 2010) or used as assumption to simplify the transshipment sending decision (Huang and Sočić, 2010). In our problem setting, as the transshipments are initiated throughout the season while the full demand realization is not resolved, a retailer with on-hand inventory has the incentive to reject requests expecting her inventory to be sold to either the overflowing unsatisfied customer or a new customer of her own in remaining periods. When the optimal accept/reject policy is not attainable to use, the suboptimal pairwise holdback levels may lead to an over-conservative transshipment policy for the requested retailers. Our preliminary numerical analyses indicate that pairwise holdback levels may lead to losses due to infrequent use of transshipments compared with centrally optimal transshipments. Therefore, as a second heuristic accept/reject policy, we define an always-accept transshipment request response behaviour.

In addition to accept/reject decisions, a transshipment policy in a multi-retailer system should specify a list of possible requested retailers for each requesting retailer in each period. The requesting retailer can traverse an ordered list from top to bottom and continue to ask for a transshipment until a transshipment request is accepted. Sequencing retailers on the list is a combinatorial problem even for a static list. Here, we introduce five different policies for selecting one retailer to request from for each transshipment. The heuristic retailer selection policies we consider make use of information on the following variables/parameters: x_i —the current inventory level of the retailers, p_i —the probability of a customer demand at a retailer, and s_i —the salvage price per unit at retailer i at the end of sales season. While selecting these variables/parameters, we benefit from the findings of Çömez *et al* (2012b), which study a two-retailer system of decentralized retailers. They show that the optimal holdback levels of a retailer are nondecreasing in her salvage price and retailers' demand probabilities.

The resulting transshipment policies can be classified in two groups—heuristics with fully observable and partially observable requested retailer inventory information. Three heuristics we introduce belong to the former group and five of them to the second one. The policies with fully observable retailer inventories assume that the requesting retailer has knowledge about the current inventory levels of the other retailers and

makes its retailer-to-request decision accordingly. In the second group, retailer inventories are only partially observable, the retailers only know whether the other retailer is stocked-out or not. Below is a brief explanation of each policy and the possible reasoning of the requesting retailer while choosing this policy.

First, heuristic policies using pairwise optimal holdback level for the requested retailer's respond decision are defined. For the requested retailer selection, if requested retailers' inventory information are fully observable, these information can be used as follows.

[Max x_i |holdback]: Request from the retailer i with highest inventory level x_i . Higher inventory level at the requested retailer may help to increase the probability of acceptance as the requested retailer would send a transshipment if her inventory level is above her holdback level.

[Max x_i/p_i |holdback]: Request from the retailer i with highest inventory-demand probability ratio x_i/p_i . The retailer with highest inventory wrt the demand rate may not be able to sell all the inventory and will have lower holdback levels due to lower demand probability. Therefore she will be more willing to accept the transshipment request.

Policies with partially observable requested retailer inventory information and the pairwise optimal holdback level using requested retailers are defined as follows.

[Min p_i |holdback]: Request from the retailer which has the lowest demand probability p_i . The requested retailer will be willing to transship if there is less probability of selling the product to the customers, so the optimal holdback levels of the requested retailer may be lower with low demand probability.

[Min s_i |holdback]: Request from the retailer i with smallest salvage price s_i . The retailer, which has the lowest salvage price, can have lower holdback levels, so can be willing to transship the product as this retailer will earn the least amount of money from the salvage sales at the end of the season.

[Random i |holdback]: Request from any one of the retailers with on-hand inventory by choosing arbitrarily. When the retailers in the system are not much different from each other and/or the stock-out retailer does not have much information on other retailers, then he can ask for a transshipment by randomly choosing a retailer with on-hand inventory. Here, the requesting retailer does not even need to know which retailers have on-hand inventory, but he can continue contacting other retailers in a random order until he finds a retailer who is not out of stock to make an official request.

The holdback level-based transshipment request responding policy is replaced by the always-accept policy in some of the heuristics above to see whether changing the respond policy may improve the performance of the heuristic. Thus, we defined and tested three additional heuristic policies.

[Max x_i |always]: Request from the retailer i with highest inventory x_i .

[Min s_i |always]: Request from the retailer i with smallest salvage price s_i .

[Random i |always]: Request from any one of the retailers with on-hand inventory by choosing arbitrarily.

In each of the heuristic policies introduced above, \mathcal{R}_{ij}^n defined in (8), can be redefined accordingly. For example, for the heuristic [Max x_i/p_i |holdback],

$$\mathcal{R}_{ij}^n = \begin{cases} 1 & \text{if } i = \arg \max_{m \in \mathcal{P}(x)} x_m/p_m; \\ 0 & \text{otherwise.} \end{cases}$$

A heuristic policy can be executed as follows. If the retailers have fully observable inventory information about the other retailers, they use one of the policies in the first group. Otherwise they use a policy from the second group. Stocked-out retailer j selects retailer i from a list of retailers with inventory, according to the selection criteria of the chosen policy. Depending on the request respond policy, retailer i accepts the request either always or if her current inventory level x_i is above her pairwise optimal holdback level against retailer j in that period. If the requested retailer rejects the request, the customer demand may overflow to another retailer (including the requested retailer i) or may get lost. The requesting retailer can make only one transshipment request in a single period. The relaxation of this assumption to a second request before the demand overflow is discussed as an extension in Section 5.3. Section 5 analyses the performance of the heuristic policies introduced above. Note that we also introduced and tested several other heuristic policies such as [Max x_i/θ_i |holdback], [Max x_i/s_i |holdback], [Max p_i |holdback], and [Min θ_i |holdback], but the performances of all of these policies are significantly dominated by the policies defined above. Thus, we neglected reporting the results of these alternative policies.

5. Numerical tests of the approximate heuristic policy performances

In this section, we first test the performance of the first heuristic [Max x_i |holdback] wrt the self-optimal solution in Section 5.1. We show that the equilibrium solution with the heuristic transshipment policy may result in a higher total profit than the equilibrium solution with the self-optimal transshipments, which is quite counter-intuitive. Thus, in Section 5.2, we test the performance of eight different heuristic transshipment policies wrt the centrally optimal transshipments. Section 5.3 includes discussions on two different extensions of the previously introduced heuristic policies. Throughout all numerical tests, the number of periods in a season is set as $N = 30$.

5.1. Self-optimal versus approximate heuristic policies

We first compare the performance of one of the heuristic policies [Max x_i |holdback] to the self-optimal transshipment policy under their respective equilibrium ordering levels. We select this heuristic for comparison, because the use of self-optimal transshipment policy implicitly assumes that retailer

inventories are fully observed, which is also the case in [Max x_i |holdback] heuristic. As detailed in Section 3, the computation time for obtaining the expected profits under self-optimal transshipments increases significantly with the number of retailers in the system. This is because first, for each stocked-out retailer's requested retailer choice all possible retailers as well as their rational accept/reject decisions should be considered. Moreover, to obtain the equilibrium ordering quantities, the expected profits under each possible ordering quantity set should be computed and best response functions should be obtained. We make the comparison in a 5-retailer system as the computation time for a larger system increases extensively.

Table 1 contains comparisons between self-optimal and heuristic solutions for 32 problem instances. The base parameter setting is such that $p_i=0.15$, $s_i=6$, $c=6.5$, $t=7.5$, $r=8$, $\tau=0.15$, and $\theta_i=0.2$, for $i=\{1, 2, \dots, 5\}$. For each problem instance in the table, only the parameter that has changed from the base setting is reported. For example, in the first setting, all parameters are at their base levels, except $p_1=0.05$.

Q^* and Q^h denote the vectors of equilibrium beginning of season ordering levels at five retailers under self-optimal and heuristic transshipment policies, respectively. There may exist multiple equilibria in some of the problem instances, but we report one of these equilibria solutions in Table 1. The table also reports the optimal order levels of retailers by ignoring transshipments and demand overflow, which reduces the problem of each retailer to a newsboy setting. The vector of these order levels is denoted by Q^0 . Our purpose by calculating and reporting these optimal newsboy order levels is to conclude

whether considering transshipments and also the transshipment policy selected change the order level decisions of retailers. J_i^* and J_i^h denote the expected profits of retailer i at the equilibrium order quantity under the self-optimal and the [Max x_i |holdback] heuristic transshipment policies, respectively. The per cent decrease in the total expected profits of retailers by using the heuristic policy wrt self-optimal policy is denoted by ΔTJ^{*h} such that

$$\Delta TJ^{*h} = \frac{\left(\sum_{i=1}^5 J_i^* - \sum_{i=1}^5 J_i^h\right)}{\sum_{i=1}^5 J_i^* * 100}$$

According to Table 1, the heuristic transshipment policy for multiple retailers provides equilibrium profits up to 5.16% lower than those under self-optimal transshipments when $s_i=5.6$ and all other parameters are at their base levels. At this setting, the equilibrium ordering levels with self-optimal and heuristic transshipment policies are $Q^*=(6, 5, 5, 5, 5)$ and $Q^h=(6, 5, 5, 6, 6)$, respectively. The optimal ordering levels for the no-transshipment, no-overflow setting are $Q^0=(5, 5, 5, 5, 5)$.

Interestingly, in some problem instances, total profits obtained with the heuristic are *higher* than total self-optimal profits, which are reported in bold letters in Table 1. Although being interesting, this result is not unexpected for a game setting. Optimal transshipment decisions help a retailer benefit optimally from other retailers. The other retailers may also optimize their transshipment decisions. Thus when all retailers use their self-optimal transshipment policies, the net benefit of

Table 1 Performance of heuristic policy [Max x_i | holdback] over self-optimal policy

	Q^*	Q^h	Q^0	$\Delta TJ^{*h}(\%)$		Q^*	Q^h	Q^0	$\Delta TJ^{*h}(\%)$
p_1					t				
0.05	(2, 6, 6, 6, 6)	(2, 6, 6, 6, 6)	(2, 6, 6, 6, 6)	-0.07	6.9	(6, 5, 5, 5, 5)	(6, 5, 5, 5, 6)	(6, 6, 6, 6, 6)	1.02
0.15	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.05	7.2	(6, 5, 5, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	2.87
0.25	(9, 6, 6, 6, 6)	(10, 6, 6, 6, 6)	(9, 6, 6, 6, 6)	1.33	7.5	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.05
0.35	(12, 6, 6, 6, 6)	(13, 6, 6, 6, 6)	(13, 6, 6, 6, 6)	1.17	7.8	(6, 6, 6, 6, 6)	(7, 6, 6, 6, 7)	(6, 6, 6, 6, 6)	3.15
s_1					r				
5.2	(5, 6, 6, 6, 6)	(4, 6, 7, 7, 7)	(5, 6, 6, 6, 6)	2.42	7.65	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(5, 5, 5, 5, 5)	-0.07
5.6	(5, 6, 6, 6, 6)	(5, 6, 6, 6, 7)	(5, 6, 6, 6, 6)	1.67	8	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.05
6	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.05	8.35	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.04
6.4	(8, 6, 6, 6, 6)	(10, 6, 6, 6, 6)	(8, 6, 6, 6, 6)	1.38	8.7	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.03
s_i					τ				
5.2	(5, 5, 5, 5, 5)	(6, 5, 5, 5, 5)	(5, 5, 5, 5, 5)	3.15	0.05	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.02
5.6	(6, 5, 5, 5, 5)	(6, 5, 5, 6, 6)	(5, 5, 5, 5, 5)	5.16	0.15	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.05
6	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.05	0.25	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.08
6.4	(7, 7, 7, 7, 7)	(8, 8, 8, 8, 8)	(8, 8, 8, 8, 8)	1.37	0.35	(6, 6, 6, 6, 6)	(7, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	1.40
c					θ_i				
6.1	(8, 8, 8, 8, 8)	(8, 8, 8, 8, 8)	(8, 8, 8, 8, 8)	0.00	0	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.05
6.5	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.05	0.05	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.05
6.9	(5, 5, 5, 5, 5)	(6, 5, 5, 5, 5)	(5, 5, 5, 5, 5)	2.86	0.1	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.05
7.4	(5, 4, 4, 4, 5)	(5, 4, 4, 4, 5)	(3, 3, 3, 3, 3)	-0.34	0.2	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	(6, 6, 6, 6, 6)	-0.05

Base setting has $p_i=0.15$, $s_i=6$, $c=6.5$, $t=7.5$, $r=8$, $\tau=0.15$, and $\theta_i=0.2$, for $i=\{1, 2, \dots, 5\}$

transshipments for a retailer can be low and can even be less than the case where all retailers use suboptimal transshipment decisions. This phenomenon is a realization of the *Braess's paradox*. Non-cooperative equilibria in a game are generally Pareto inefficient, that is, one player could be better off without anyone else being worse off. When more resources (eg, optimal transshipment policies) are provided to players of the game, the inefficiency of the equilibrium results in decreases in performance (eg, equilibrium profit); see Korilis *et al* (1999). In Table 1, for 18 problem instances total profit of retailers is higher under heuristic transshipments. Moreover, in 16 of these instances, each of five retailers is better-off when retailers use heuristic transshipment policy.

Table 1 shows that in general, the equilibrium ordering levels with heuristic transshipment policy [Max x_i |holdback] and self-optimal transshipments and optimal order levels for no-transshipment, no-demand overflow system are similar, despite sometimes higher equilibrium order levels resulting in heuristic transshipment system. Similarity among Q^0 , Q^h , and Q^* indicates that use of transshipment and/or the transshipment policy used do not have significant impact on the equilibrium inventory levels of retailers. The slight impact of transshipments on retailers' choice of ordering levels is also indicated by several past studies such as Grahovac and Chakravarty (2001), Zhao *et al* (2005), Zhang (2005), and Çömez *et al* (2012a). This conclusion is in fact important in terms of our computational analyses. As stated above, computation time requirement to obtain equilibrium ordering levels increases extensively with the number of retailers in the system. The low impact of transshipment usage on the choice of ordering levels rationalizes the use of optimal order levels from a no-transshipment system in a transshipment system without much harm, which is the case in Section 5.2.

5.2. System-optimal versus approximate heuristic policies

Shown that the profits obtained through the use of a heuristic transshipment policy may exceed those obtained by using self-optimal policy, in following we evaluate the performance of the various heuristic transshipment policies we introduce by comparing to the total expected system profits under centrally optimal transshipments. Although the structure of the optimal transshipments policy even for a centrally managed system is not easy to define because of the dependence of a transshipment decision on complete system state (Archibald, 2007; Çömez *et al*, 2012a), we can calculate the resulting expected profits for a given problem instance. Moreover, we also evaluate the expected profits under no transshipments, which is also called no pooling, so that decrease in payoffs due to the use of heuristic transshipment policy can be compared with the loss due to no use of transshipments that can be regarded as an extreme transshipment policy.

To isolate the effects of different initial inventory levels, we use the same beginning of season ordering levels with heuristic and centrally optimal transshipment policies for a given

problem instance. For this purpose, we obtain the optimal ordering level of each retailer by ignoring transshipments and demand overflow, which reduces the problem of a retailer to a newsboy setting as defined in Section 5.1. Besides, by setting the ordering levels, computation time to obtain the expected profits under any transshipment policy reduces to a reasonable amount. Thus, setting the beginning of season ordering levels according to optimal newsboy quantities helps us both to focus on transshipment policy performances for comparison and also to conduct more numerical exercises by shortening the computation time significantly.

We conduct tests for multi-retailer systems from 3 retailers up to 10 retailers, $M = \{3, 4, \dots, 10\}$. For each retailer system with M retailers, 350 problem instances are run. For each problem instance of an M -retailer system, first, the beginning of season ordering levels are calculated by regarding each retailer as a newsboy and the optimal newsboy order quantities are obtained. Then, the total expected profit of retailers is obtained in each of the cases when centrally optimal transshipments, no transshipments, or each one of the eight heuristic transshipment policies defined in Section 4 are used. The total expected profit with centrally optimal transshipments is denoted by TJ^C . The total expected profit of retailers when transshipments are managed by a heuristic policy is denoted by TJ^h . Then the per cent decrease in total expected profit of a centrally managed system by the use of heuristic policy h is denoted by ΔTJ^h such that

$$\Delta TJ^h = \frac{(TJ^C - TJ^h)}{TJ^C} * 100.$$

In addition to the total expected profit, the distribution of the expected demand satisfaction method is calculated for each problem instance and transshipment method used. The total expected demand of the retailers over a sales season is $N^* \sum_{i=1}^M p_i$. Depending on the transshipment policy used, some part of this total expected demand will be satisfied from direct retailer stock (S), some by transshipments (T), some by retailer stock after demand overflow (F), and the rest will be lost (L), such that $S + T + F + L = N^* \sum_{i=1}^M p_i$. Then per cent of expected demand to be satisfied directly from stock, by transshipments, from stock after overflowed, and lost are denoted, respectively by S , T , F , and $L\%$, where they sum up to 100%.

For the parameter selection, we inspired from the car retailing industry considering the well-known application of transshipments among car dealers. We take the selling price of a car as the base parameter and set the purchase price c , salvage price s , and transshipment price t relative to r . First, we set the relation between the selling price and purchase price by $r = \alpha c$, where α denotes the profit margin added to the purchase price. According to Autos.com (2013) the average profit margin for a vehicle is between 6 and 9%, which can increase to 15% for luxury cars. During numerical analyses α is selected from a range of [1.05, 1.25] indicating profit margins between 5 and 25%. For the salvage price, we define $s_i = \delta_i r$, where $1 - \delta_i$ denotes the discount rate on the selling price r to salvage the

end of season inventory. Smartmoney.com (2013) reports various discount rates on ten 2010 car models according to the time that a car sits on the lot. The discount rate ranges from 5% off up to 21%. They add that end of year discounts can go up to 35%. In following, we let δ_i to change within $[0.65, 0.8]$, which indicates a discount rate between 20 and 35% off. Next, we assume that the transshipment price is also determined as a discounted function of the selling price such that $t = \beta r$ and $1 - \beta$ is the discount rate to send a transship. By definition $t \geq \theta_i r + \bar{\theta}_i s_i$ should satisfy, otherwise retailer i does not accept any transshipment request. The discount rate applied for transshipping should not be greater than the salvage discount, so that it can be profitable for the requested retailer to send her unit inventory as a transshipment instead of salvaging. Thus, β is selected from the range $[0.9, 0.95]$ that indicates a transshipment price obtained by a discount rate between 5 and 10% on the selling price.

Regarding the transportation cost τ , we contacted a few car shipping companies to get quotes for transshipping a mid-size car among two cities approximately 200 miles apart. The average quote was around \$300. If we take the average selling price r of a new mid-size economy level car to be $r = \$20,000$, we arrive at a ratio of $r/\tau = 200/3$. During the numerical analyses, wlog, we scale r and change it within $[6, 10]$ with an average of 8. We change τ between 0.05 and 0.2 with the average value of 0.125, which is around $8 \times (3/200) = 0.12$.

To measure the sensitivity of results to system parameters, for each problem instance, we set some of the parameters to some certain values and generate the rest of the parameters randomly from continuous Uniform distributions with the ranges explained above and summarized in Table 2. For the same fixed parameters, we generate 10 problem instances, in each of which the remaining parameters are randomly generated. Thus, by taking the average of performance measures over each of 10 problem instances, we indicate the average performance under these fixed parameters. By this way, the effect of the fixed parameters on heuristic performances can be evaluated more rigorously by eliminating the effects of other parameters.

We first compare the distributions of how expected demand is satisfied among different transshipment policies. For this purpose, in an M -retailer system, the average per cent of expected demand to be satisfied from retailers' stock (for direct and overflowed customer demand) $(S + F)\%$, by transshipments $T\%$, and lost $L\%$ are calculated for 350 problem instances under each of the different transshipment policies. Detailed numerical results on $(S + F)$, T , and $L\%$ for all transshipment policies (for M equal to 3 to 10) are available in the online companion.

Figures 1 and 2 show that as the number of retailers increases, the percentage of transshipments in total demand

Table 2 Distributions of randomly generated parameters.
 $p_0 = 1 - \sum_{i=1}^M p_i$ and $i = \{1, 2, \dots, M\}$

p_i	$U(0, 1/M)$	δ_i	$U(0.65, 0.80)$	α	$U(1.05, 1.25)$	β	$U(0.90, 0.95)$
r	$U(6, 10)$	τ	$U(0.05, 0.20)$	θ_i	$U(0, 1/M)$		




Figure 1 The change in average $T\%$ versus M .




Figure 2 The change in average $(S + F)\%$ versus M .

satisfaction increases and sales from stock percentage decreases. This shows that retailers rely more on transshipments in larger systems. The result is intuitive as the system gets larger and the total average expected demand is assumed to be constant, the expected demand per retailer decreases. This causes the stocking levels at each retailer to decrease as well. Then the demand and supply match for each retailer becomes more challenging. Thus, this results in more often use of transshipments as an emergency demand satisfaction activity. Moreover, the lost sales percentage $L\%$ also increases with the number of retailers as the difficulty in demand satisfaction increases in the overall as shown by Figure 3. Note that in all these figures, always-transship heuristics are represented by dashed lines.

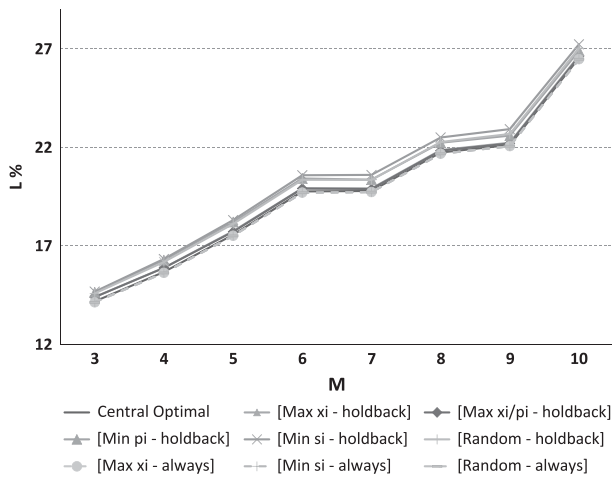


Figure 3 The change in average $L\%$ versus M .

Policies using holdback levels show slightly higher lost sales than the always-tranship heuristics. Thus, always-tranship policies result in closer lost sales percentages to centrally optimal policies than the holdback level policies. Also note that all heuristics having always-transship respond policy lead to the same lost sales percentages, because independent of how to choose the requested retailer with inventory on-hand, a requested retailer will always send transshipments. Therefore, as long as there is inventory in the system, there will not be any lost sales regardless of requesting decision. The lower lost sales amounts with always-tranship policies are expected as when a transshipment request is rejected due to the holdback level, there is a possibility that the demand will be lost.

Heuristics that always tranship have higher transshipment percentages $T\%$ as expected, but holdback level policies have closer $T\%$ to centrally optimal transshipments. As shown by Figures 1 and 2, in the overall, while the holdback level-based heuristics lead to lower transshipment usage than the centrally managed policy for smaller retailer systems, in larger systems, they utilize transshipments at least as much as the centralized one. The reason is that as the number of retailers increases in the system, there is a higher chance to obtain transshipments for retailers using heuristic policies. This is mainly due to lower pair-wise optimal holdback levels in larger systems where each retailer has a lower expected demand. In cases with 3–9 retailers, heuristics [Max x_i]holdback] and [Max x_i/p_i]holdback] have the closest transshipment percentages to that of centrally optimal solution. In the 10-retailer case, [Min s_i]holdback] results in the closest transshipment amount. On the other hand, compared with centrally optimal transshipments, with holdback heuristic policies a larger portion of the total demand is directly satisfied from stock $(S+F)$ in small systems (3–5 retailers), but this portion is below that of centrally optimal policy in larger systems. This result directly follows as a consequence of higher utilization of transshipments with holdback level policies compared with the centralized one in larger systems.

Next, we report the average, maximum, and standard deviation of the per cent gap in the expected profits between the centrally optimal transshipment policy and each of the eight heuristic policies, as well as the no-transshipment policy denoted by ΔTJ^h through Table 3 and Figure 4. In Table 3, we also report the range of the total expected profit with centrally optimal transshipments TJ^C for each M -retailer system over 350 problem instances. Note that this range is very similar for different number of retailers. The total expected demand for each M -retailer system in each setting is $\sum_{i=1}^M E(p_i) = \sum_{i=1}^M 1/(2M) = 1/2$. Here the expected demand for each retailer in an M -retailer system is $E(p_i) = 1/(2M)$, which follows from the Uniform distribution of p_i in the range $U(0, 1/M)$ as reported in Table 2. Thus, as the total expected demand potentials in different size retail systems are the same, the ranges of the total profits with centrally optimal transshipments TJ^C are similar.

In addition to comparing the average, maximum, and standard deviation of the gap between the centrally optimal transshipment policy and the eight heuristic policies, for each M -retailer system, paired t -tests have been performed using 350 problem instances to compare the gaps ΔTJ^h provided by various heuristics. For each heuristic policy pair of h_i and h_j , the null hypothesis is $H_0: \Delta TJ^{h_i} - \Delta TJ^{h_j} = 0$. When the null hypothesis is rejected, the alternative hypothesis $H_1: \Delta TJ^{h_i} - \Delta TJ^{h_j} \neq 0$ is accepted. Thus the hypothesis tests whether the difference in mean gaps, thus performances, of heuristics i and j are statistically significant or not. The p -value of a t -test resulted lower than 0.05 indicates that the null hypothesis is rejected, that is, compared heuristics' performances are significantly different. p -values for all comparisons in each M -retailer system are reported in Table 4 for $M = \{3, \dots, 10\}$.

First, we observe that although the maximum is not monotone in M , the average per cent profit gap increases with the number of retailers in the system. Despite not being monotone, the standard deviation of ΔTJ^h also tends to increase with the size of the system.

Table 3 and Figure 4 show that [Max x_i]holdback] and [Max x_i/p_i]holdback] are among the best performing heuristics for 3–9 retailer systems. When the number of retailers increases to 10, [Min s_i]holdback] policy performs better than these two heuristics according to the average gap in Table 3, but its performance is not statistically different from that of [Max x_i]holdback] according to Table 4. When we consider the transshipment percentage $\%T$ and average ΔTJ^h results together, we see that heuristics which have transshipment percentages closer to central optimal perform better than other heuristics in general.

When requested retailer inventories are only partially observable, policies [Min s_i]holdback] and [Min s_i]always] perform among the best. According to the average gap, when the number of retailers in the system is between 3 and 5, [Min s_i]always] performs better than [Min s_i]holdback], which is vice versa for 6- to 10-retailer systems. When the differences in gaps are compared, it seems that although the

Table 3 Average, maximum and standard deviation of $\Delta T J^h$

Transshipment policy	Avg $\Delta T J^h$	Max $\Delta T J^h$	Std $\Delta T J^h$	Avg $\Delta T J^h$	Max $\Delta T J^h$	Std $\Delta T J^h$	Avg $\Delta T J^h$	Max $\Delta T J^h$	Std $\Delta T J^h$
<i>M</i> = 3, TJ^C = [0.8, 40]									
[Max x_i -h]	0.74	8.65	0.85	0.87	5.75	0.71	0.96	5.27	0.82
[Max x_i/p_i -h]	0.72	8.60	0.84	0.86	5.17	0.76	0.93	5.54	0.89
[Min p_i -h]	1.17	8.98	0.93	1.60	5.81	1.01	1.81	7.78	1.24
[Min s_i -h]	1.00	8.71	0.91	1.29	5.83	0.86	1.47	6.28	1.05
[Random-h]	1.05	8.84	0.88	1.38	5.98	0.84	1.58	6.19	0.99
[Max x_i -a]	0.88	8.65	1.31	1.34	13.16	1.87	1.68	18.36	2.20
[Min s_i -a]	0.76	9.27	1.47	1.16	13.83	2.15	1.46	18.45	2.51
[Random-a]	0.91	9.05	1.39	1.43	13.96	2.01	1.79	18.76	2.31
No transshipment	10.29	34.99	4.58	14.04	38.04	5.47	17.24	40.54	5.87
<i>M</i> = 4, TJ^C = [1.6, 39.9]									
<i>M</i> = 5, TJ^C = [0.4, 38.6]									
[Max x_i -h]	1.16	9.85	0.94	1.34	5.48	1.03	1.53	6.05	1.16
[Max x_i/p_i -h]	1.17	7.56	1.00	1.34	5.37	1.08	1.56	7.50	1.28
[Min p_i -h]	2.05	8.67	1.43	2.08	7.33	1.39	2.16	8.39	1.46
[Min s_i -h]	1.63	7.60	1.11	1.61	6.07	1.15	1.68	6.69	1.22
[Random-h]	1.82	9.94	1.17	1.93	6.67	1.20	2.07	6.36	1.24
[Max x_i -a]	2.22	23.28	2.73	2.70	24.82	3.29	3.15	27.86	4.01
[Min s_i -a]	1.90	23.68	3.14	2.19	26.58	3.77	2.58	27.51	4.49
[Random-a]	2.33	23.80	2.86	2.77	25.62	3.41	3.20	27.86	4.12
No transshipment	20.41	44.50	6.39	22.61	54.38	6.81	25.67	67.71	7.55
<i>M</i> = 6, TJ^C = [0.8, 33.4]									
<i>M</i> = 7, TJ^C = [0.8, 39.3]									
<i>M</i> = 8, TJ^C = [0.6, 40.5]									
[Max x_i -h]	1.66	6.61	1.22	1.81	8.92	1.32	1.81	8.92	1.32
[Max x_i/p_i -h]	1.73	8.80	1.41	1.91	6.89	1.53	1.91	6.89	1.53
[Min p_i -h]	2.29	8.85	1.67	2.31	13.15	1.74	2.31	13.15	1.74
[Min s_i -h]	1.76	7.66	1.27	1.73	9.06	1.29	1.73	9.06	1.29
[Random-h]	2.24	7.91	1.41	2.29	8.95	1.47	2.29	8.95	1.47
[Max x_i -a]	3.25	30.10	4.13	4.02	25.10	4.79	4.02	25.10	4.79
[Min s_i -a]	2.57	31.32	4.58	3.36	25.75	5.28	3.36	25.75	5.28
[Random-a]	3.28	30.10	4.19	4.06	25.10	4.83	4.06	25.10	4.83
No transshipment	28.38	55.56	7.52	31.07	56.05	8.22	31.07	56.05	8.22
<i>M</i> = 9, TJ^C = [0.4, 34.6]									
<i>M</i> = 10, TJ^C = [0.8, 36.5]									




Figure 4 The change in average $\Delta T J$ versus *M*.

performance of [Min s_i |always] is not significantly different than the performance of [Min s_i |holdback] in 3–6 retailer systems, [Min s_i |holdback] performs statistically better than all other partially observable policies in 7-, 8-, and 10-retailer systems. Relying on *p*-values in Table 4, [Random|holdback] heuristic is performing significantly better than [Min p_i |holdback] in all systems, except 9 and 10 retailers, where the performance difference is not statistically significant. Therefore we can conclude that when there is no information about x_i or s_i , selecting the retailer to ask for transshipment randomly is a better policy than using demand rates for selection.

Although the heuristics with always-accept respond policy perform even better than some of the holdback level-based policies on average when the number of retailers is small, as the number of retailers increases, they perform poorly. For example, according to Table 4, for the 3-retailer system, all always-accept respond policy heuristics perform statistically significantly better than [Min p_i |holdback] policy. On the other hand, for 10-retailer system, all holdback level-based policies' performances are statistically better than the performances of all

always-accept policies. The maximum ΔTJ^h and standard deviation of ΔTJ^h are higher in always-transship heuristics even for smaller systems. Although lost sales percentage is lower in always-transship strategies, this gain does not result in smaller ΔTJs as seen in Table 3.

Table 3 shows that ΔTJ with no-transshipment policy increases significantly as the number of retailers increases. On the other hand, the differences between the performance of no-transshipment and all other heuristic policies are statistically significantly different in all M -retailer systems. This result confirms the importance of using transshipments, especially as the number of retailers increases. We can conclude that in general it is better for the retailers to compute the pairwise optimal holdback levels and use as a transshipment respond policy. However, if they cannot, it is better to always transship rather than making no transshipments at all.

Next, we investigate the sensitivity of the profit gap wrt the system parameters individually. Recall that for a fixed parameter, we generate 10 problem instances, in each of which the remaining parameters are randomly generated according to Table 2. Then the average of profit gaps over these 10 problem instances indicates the average performance under this fixed parameter. Thus, in Table 5, each number indicates the average value over 10 instances. The conclusions regarding the sensitivity analyses are drawn from all retailer systems including 3-10 retailers, but the results are reported only for seven retailers in Table 5 due to the space limitations. The sensitivity results for 3- and 10-retailer systems are given in the online companion for a benchmark.

In general, the sensitivity results do not show strictly monotone behaviours with the change in the fixed parameter and among all retailer systems. This is mainly due to the setting that to increase rigour at each fixed value of a parameter, 10 problem instances are generated where all other parameters are randomly selected. Thus, the average gap at a fixed parameter also includes the randomness factor for other parameters. However, results indicate certain sensitivity directions for some of the parameters.

Firstly, as δ_i increases, the performances of always-transship policies improve for 3- to 9-retailer systems. Increase in δ_i means increase in salvage price s_i . With higher s_i , newsboy stocking levels are set to higher values, which results in higher from stock demand satisfaction rates as indicated by Table 5. Thus, the way that transshipments are granted get similar in always-transship and centrally managed systems.

On the other hand, when β increases, which indicates an increase in transshipment price t , holdback level-based policies perform closer to the centrally optimal policy. When transshipment price is higher, pairwise optimal holdback levels tend to decrease as transshipments become more profitable (Çómez *et al.*, 2012b). Thus, more transshipments are granted with holdback level-based policies approaching the transshipment amounts in the centrally managed system. Note that transshipment decisions are independent of transshipment price with both central optimal and always-transship policies.

The profit margin α is another variable that shows considerable effect on the performance of always-transship policies. As α increases, the purchase price c decreases, so the initial inventory levels increase and always-transship policies perform better in general. Table 5 shows that with the increase in α , the demand satisfaction directly from stock increases as a result of higher initial inventory levels. Thus, similar to the effect of increasing δ_i , the use of transshipments under always-transship policies approaches to that under central optimal policy.

The results in Table 5 indicate a few more points. One is that for $\alpha = 1.05$, always-accept policies perform poorly wrt other policies. This leads the average performance of these policies to decrease to a level notably lower than [Max x_i/p_i |holdback] and [Max x_i |holdback]. In fact, for most of the instances, [Min s_i |always] is the best performing heuristic policy. Second, policies [Min p_i |holdback], [Random|holdback], and [Random|always] do not prove to be the best performing in any of the instances. However, the average performances of these last two heuristics, in which requested retailer is selected randomly among all with on-hand inventory, are highly close to those where requested selection requires information gathering. We believe this result is notable for multi-retailer transshipment policy studies, as to our knowledge, the performance of random selection methods are not even considered for discussion in the transshipment literature before.

5.3. Further extensions of heuristic policies

As illustrated by Figure 4, although always-accept respond policies manage transshipments similar to holdback level-based respond policies up to six retailers, for seven or more retailer systems most of always-accept policies perform significantly poorer than holdback level policies. On the other hand, using the inventory level information for the requesting retailer selection results in better performance compared with the case when the inventories are partially observable and selection is based on retailer characteristics other than inventory level. Overall, [Max x_i |holdback] and [Max x_i/p_i |holdback] policies perform quite well among all tested to manage transshipments close to centrally optimal transshipment decisions. Given that using holdback level-based respond policy behave well and known that requested retailers base their respond decisions on holdback levels in this policy, it might be rational for the requesting retailer to consider holdback levels of other retailers in her requested retailer choice as well.

For this purpose, we extend [Max x_i/p_i |holdback] policy to include potential requested retailers' holdback levels in the requested retailer selection. A new policy [Max $(x_i - \text{holdback}_{ij})/p_i$ |holdback] is introduced such that requesting retailer j selects the requested retailer i in period n , which has the highest ratio of the difference between her current inventory level (x_i) and the holdback level that she will use against retailer j (holdback_{ij}) to her demand rate p_i . Retailers that are not willing or cannot send a transshipment to the stocked-out retailer j would have non-positive ratio. Thus, by searching the highest ratio, retailer j

Table 5 Sensitivity analysis— $M = 7$ retailer system

Parameter	Value	Central optimal				$\Delta T J^h$								
		S%	T%	F%	L%	[Max x_i-h]	[Max x_i/p_i-h]	[Min p_i-h]	[Min s_i-h]	[Random-h]	[Max x_i-a]	[Min s_i-a]	[Random-a]	No transshipment
Σp	0.4	59.19	20.74	1.33	18.73	1.94	1.94	2.20	1.14	2.20	2.79	1.18	2.45	30.23
	0.5	67.47	21.10	0.69	10.74	1.69	1.69	2.50	1.77	2.50	2.33	1.06	2.15	26.01
	0.6	69.93	19.62	0.95	9.49	1.34	1.34	2.11	1.63	2.11	1.45	0.60	1.44	23.94
	0.7	61.07	11.76	3.51	23.66	1.09	1.09	1.36	1.35	1.36	2.78	2.76	2.99	21.15
	0.8	72.69	12.71	1.94	12.66	1.08	1.08	1.73	1.58	1.73	1.72	1.32	1.87	19.62
δ_i	0.60	50.69	16.50	1.48	31.32	0.61	0.59	1.06	1.08	1.08	2.68	2.88	2.88	29.81
	0.65	50.82	14.13	1.84	33.20	0.40	0.34	0.88	0.92	0.92	2.67	2.89	2.89	23.65
	0.70	55.21	15.32	1.37	28.11	0.58	0.44	1.43	1.37	1.37	2.12	2.35	2.35	27.21
	0.75	70.38	13.65	1.15	14.82	0.45	0.37	1.05	1.10	1.10	1.44	1.65	1.65	17.86
	0.80	71.80	12.32	1.38	14.50	0.68	0.66	1.75	1.60	1.60	1.18	1.46	1.46	19.29
τ	0.05	59.82	19.76	1.13	19.29	1.15	1.05	1.95	1.51	1.79	1.16	0.49	1.09	26.44
	0.10	61.96	18.64	0.93	18.46	1.08	1.09	2.37	2.13	2.06	1.37	1.06	1.51	26.19
	0.15	64.73	13.87	1.97	19.43	1.37	1.89	2.30	1.31	1.85	3.55	3.20	3.95	23.28
	0.20	54.88	18.14	1.38	25.61	1.40	1.74	2.31	1.92	2.08	5.59	5.62	6.05	25.81
	0.25	68.96	14.82	1.76	14.46	1.58	1.81	2.61	1.86	2.18	4.02	3.92	4.43	19.83
β	0.86	63.48	14.85	1.61	20.06	3.03	3.18	4.91	3.89	4.36	2.65	2.01	2.63	21.13
	0.88	52.41	13.37	2.66	31.55	1.90	1.53	3.75	3.06	3.08	3.56	3.42	3.69	24.92
	0.90	75.96	17.20	0.46	6.38	1.92	2.11	3.15	2.31	2.91	1.84	0.57	1.82	19.97
	0.92	72.32	17.42	0.46	9.81	1.58	1.73	2.66	2.09	2.46	2.40	1.34	2.32	20.81
	0.94	62.12	13.93	1.02	22.93	1.24	1.22	1.58	1.10	1.43	2.33	1.66	2.26	23.58
r	6	66.36	16.27	1.17	16.20	2.25	1.95	2.06	1.50	2.20	4.25	2.82	3.86	22.24
	7	62.25	17.02	1.25	19.49	1.16	1.32	1.94	1.27	1.73	1.75	1.16	1.89	24.02
	8	58.40	14.68	2.36	24.57	1.16	1.07	1.64	1.42	1.61	3.43	3.19	3.57	21.69
	9	66.53	17.20	0.95	15.31	1.61	1.22	2.16	2.09	2.35	2.61	1.71	2.47	23.67
	10	63.91	16.45	1.13	18.52	1.55	1.49	2.33	1.73	2.26	2.74	1.88	2.76	23.94
α	1.05	31.47	3.71	3.64	61.18	0.47	0.46	0.57	0.55	0.51	10.83	11.54	11.26	23.30
	1.10	48.25	11.51	3.28	36.96	0.63	0.59	1.02	0.94	0.88	2.53	3.08	2.88	25.15
	1.15	64.81	21.01	0.65	13.52	1.30	1.02	2.71	2.54	2.55	0.84	0.60	0.99	25.97
	1.20	72.90	22.66	0.04	4.40	1.71	1.84	2.91	2.39	2.77	1.44	0.19	1.35	24.02
	1.25	84.87	14.90	0.00	0.23	2.50	2.59	2.41	0.98	2.34	2.49	0.08	1.84	15.08
$\Sigma \theta$	0.5	58.08	13.43	2.56	25.93	1.26	1.04	1.45	1.21	1.49	3.54	3.32	3.66	21.54
	0.6	64.88	12.37	2.49	20.26	1.31	1.40	1.83	1.06	1.62	2.68	2.06	2.74	16.59
	0.7	67.67	16.01	0.81	15.51	1.34	1.64	2.31	1.50	1.97	1.68	0.71	1.67	17.21
	0.8	65.28	17.86	2.26	14.60	1.54	1.28	1.77	1.29	1.77	1.83	0.90	1.62	18.85
	0.9	60.99	14.94	3.89	20.18	0.89	1.26	1.90	1.15	1.43	2.18	2.09	2.52	17.45
				Avg.	1.34	1.34	2.08	1.61	1.93	2.70	2.19	2.77	22.61	
				Max.	3.03	3.18	4.91	3.89	4.36	10.83	11.54	11.26	30.23	
				Std.	0.58	0.63	0.83	0.66	0.73	1.73	2.03	1.82	3.62	

would guarantee to ask from a retailer who is willing to send a transshipment, if there is at least one such retailer. This property was not guaranteed in any of the previous heuristics with holdback level-based respond policies, but with always-accept respond policies. Therefore, we expect this new policy to increase the number of successful transshipment requests, so the contribution of transshipments in demand satisfaction.

To measure the performance of this extended heuristic, the heuristic policy is run for the same 350 problem instances for each M -retailer system used in Section 5.2. We report the average, maximum, and standard deviation of the per cent gap in the total expected profits between the centrally optimal transshipment policy and the $[\text{Max } (x_i - \text{holdback}_{ij})/p_i]$ holdback policy. Besides, the per cent of expected demand to be satisfied directly from stock (for direct and overflowed customers), by transshipments, and lost are also calculated. The results are given in Table 6.

As expected, using this policy utilizes transshipments slightly more than other holdback level-based respond policies. Comparing the average gaps listed in Table 6 with those in Table 3, except the 10-retailer system, $[\text{Max } (x_i - \text{holdback}_{ij})/p_i]$ holdback performs better than all other introduced heuristics. However, the magnitude of the improvement in the average performance by this heuristic may not be promising enough to compensate its additional difficulty in implementation compared with other policies introduced. A requesting retailer needs to observe the inventory levels of all other retailers with on-hand inventory as well as calculate their holdback levels to be used against himself. This high computation requirement in a decentralized system environment may not be so practical especially without the promise of a significant improvement in profit.

Next, we investigated our assumption of a requested retailer making only one request per stockout. In practice, a stocked-out retailer may communicate with several retailers one after the other until he gets a transshipment request accepted. To assess the value of considering additional requests for a rejected retailer, we reanalyse the heuristic policy $[\text{Max } x_i]$ holdback by allowing a second request option if the first one is rejected. Although allowing multiple requests until

one is accepted is the most realistic case, the additional improvement in profit that can be provided by a second request signals the importance of a single-request assumption used in earlier analyses.

To measure the performance of $[\text{Max } x_i]$ holdback policy, in which the requesting retailer has the second request option, if the first one is rejected, the heuristic policy is run for the same 350 problem instances for each M -retailer system used in Section 5.2. We report the average, maximum, and standard deviation of the gap with the centrally optimal transshipment policy in the total expected profits. Besides, the per cent of expected demand to be satisfied directly from stock, by transshipments, and lost are also calculated and the results are given in Table 7.

Comparing the profit gaps resulting from using $[\text{Max } x_i]$ holdback with two requests per stockout to those with the same heuristic, but with only one request from Table 3, it is clear that allowing a second request option is decreasing (in loose terms) the average, maximum, and standard deviation of the gaps. This profit improvement can be devoted to the slight increase in the use of transshipments and the decrease in loss sales. Besides, it is observed that the demand satisfaction rate directly from stock for retailer's own and overflowing customers decreases as well. However, we can claim that the improvement in the performance of the $[\text{Max } x_i]$ holdback heuristic after allowing a second request option may not be enough to compensate the increased difficulty in analysing the multiple transshipment requests for a stockout. Moreover, given that requested retailers are selected in decreasing order of their expected willingness to send a transshipment, allowing more than two requests in case first two are rejected may bring even a smaller marginal improvement over the assumption of maximum two requests per stockout. Therefore, allowance of a single transshipment request per stockout seems to be not restricting for the purpose of comparing performances of various easy-to-implement transshipment policies. However, still in practice, any one of the policies introduced by this study can perform slightly better when retailers spend effort to contact more than a single other retailer to get a transshipment request accepted.

Table 6 Performance of heuristic policy $[\text{Max } (x_i - \text{holdback}_{ij})/p_i]$ holdback over centrally optimal policy

# of retailers	Avg $\Delta T J^h$	Max $\Delta T J^h$	Std $\Delta T J^h$	%S + F	%T	%L
M=3	0.64	8.61	0.82	78.89	6.76	14.35
M=4	0.72	5.07	0.66	74.79	9.40	15.80
M=5	0.79	4.95	0.74	70.35	12.01	17.63
M=6	1.01	6.60	0.87	66.30	13.89	19.82
M=7	1.22	4.72	0.98	64.35	15.84	19.80
M=8	1.43	7.47	1.21	60.67	17.58	21.75
M=9	1.61	7.58	1.32	58.10	19.79	22.11
M=10	1.83	6.88	1.48	52.93	20.55	26.53

Table 7 Performance of heuristic policy $[\text{Max } x_i]$ holdback when two requests are allowed per stockout over centrally optimal policy

# of retailers	Avg $\Delta T J^h$	Max $\Delta T J^h$	Std $\Delta T J^h$	%S + F	%T	%L
M=3	0.70	8.65	0.83	78.95	6.67	14.37
M=4	0.83	5.68	0.68	74.97	9.18	15.86
M=5	0.92	5.27	0.77	70.65	11.65	17.71
M=6	1.14	9.16	0.91	66.72	13.38	19.90
M=7	1.32	5.15	1.01	64.86	15.24	19.89
M=8	1.53	6.05	1.16	61.31	16.85	21.84
M=9	1.65	6.61	1.21	58.85	18.93	22.22
M=10	1.80	7.29	1.29	53.73	19.63	26.64

6. Conclusion

This paper studies a system of multiple retailers selling the same product, where the retailer operations are decentralized. Each retailer is responsible for maximizing his own profit through the inventory stocking decision at the beginning of the selling season and demand satisfaction efforts throughout the season. In this system when a retailer gets stocked-out, he may ask for a transshipment from other retailers to satisfy a waiting customer demand. In case of a transshipment need, requesting (stocked-out) retailer should decide which retailer to request from and requested retailer should decide on whether to accept the transshipment request or not. As transshipments among independently managed retailers are considered, it becomes more important that the transshipment policy used is easy-to-exercise, as transshipments become voluntary activities that are not imposed by a central agent. In the lack of easy-to-use optimal policies, many firms may follow extreme policies such as either they do not want to commit to use transshipments (called no sharing) or commit to send transshipment as long as it is possible (called complete sharing). However, instead of extreme policies, the introduction of easy-to-use and close-to-optimal policies can encourage independent retailers to mitigate supply-demand match through the use of transshipments. Therefore, we introduce heuristic transshipment policies for selecting the retailer to request from and for making the accept/reject decision. Thus, we develop several heuristic transshipment policies and compare their performances with self-optimal and centrally optimal policy.

Firstly, we show that when all retailers use heuristic transshipment policies, the total retailer pay-offs may be more than that when they use self-optimal transshipment policies. In other words, using a heuristic policy may result in higher total expected profit and better individual profits for all retailers. This is a realization of the Braess's paradox, which explains how equilibrium solution may move to a more efficient result when non-optimal actions are taken by players of a game instead of self-optimal ones. The result is important to comment on the general intuition that heuristic methods are inferior to optimal solutions. When retailers are in fact competitors due to demand overflow and their stocking decisions form an inventory game, their use of self-optimal transshipment policies may result in an equilibrium stocking solution that is inferior to the one under heuristic transshipment policies. Given that computing optimal transshipment decisions depends on the complete system state and is challenging to use in practice and shown that self-optimal transshipment policies may be even inferior to heuristic policies, then the need for the introduction of good heuristic policies and their use in practice are clearly underlined by this study.

We develop five different heuristic policies for the decision of selecting the retailer to ask for transshipment. These heuristic policies can be classified into two according to the availability of inventory level information on others: fully observable policies where the retailers know the inventory levels of others and partially observable policies where the requesting retailer

knows which retailers have stock on-hand but not the exact amounts. We use two different policies for the transshipment accept/reject decision: always-accept and pairwise optimal holdback levels. Always-accept is an extreme policy that retailers do not always prefer. Pairwise optimal holdback levels are obtained by considering a two-retailer system of only the requesting and requested retailer such that requested retailer accepts to send a transshipment if her inventory level at the time of the request is above her holdback level. We compare a total of eight policy combinations with the centrally optimal policy over 3- to 10-retailer systems.

The results show that retailers rely more on transshipment as the number of retailers in the system increases. In fewer-retailer systems, using holdback level policies results in a larger portion of the total demand directly satisfied from stock relative to central optimal policy. In larger systems, this percentage falls behind that of centrally optimal system. The average profit gap increases with the system size. However, the profit loss due to a no transshipment policy is even much higher.

In terms of transshipment request respond decision, always-accept policies result in lower lost sales and higher transshipment rate but holdback policies have closer transshipment rate to that of centrally optimal system. Heuristics with transshipment amounts closer to that of centrally optimal policy perform better in almost all cases in terms of total system profits.

In terms of transshipment requesting decision, if the inventory levels are fully observable, $[\text{Max } x_i|\text{holdback}]$ and $[\text{Max } x_i/p_i|\text{holdback}]$ are better performing policies considering total system profits. If the inventory is partially observable, $[\text{Min } s_i|\text{holdback}]$ is a good policy. One surprising result is that random selection policy performs better than some other policies

From a managerial point of view, there are several key findings of this study that can be important for the practice of transshipments in decentralized and multi-retailer systems. One is that retailers can be better off by using heuristic transshipment policies instead of self-optimal in terms of both practicality and also profits contrary to the general belief. Second, although the performance of heuristic policies is in general poorer in large systems compared with small systems, but retailers are far better off using heuristic policies in larger systems compared with the extreme policy of not using any transshipments. Third, when retailers are allowed to access others' inventory state information, it is better to select the requested retailer with higher inventory level expecting a higher chance of a request to get accepted. More interestingly, it is shown that if there is no information about inventory level or salvage prices of others, retailers are better off if they randomly select the retailer to request from. To the best of our knowledge, this is the first study to analyse the performance of random retailer selection as a transshipment requesting policy, although a random selection process is probably highly practiced in industry when there is no systematic support is available. Forth, comparison of alwaysaccept and holdback policies show that when the system is small, always-accept policies often work better. Holdback

policies perform better in larger systems. Thus, in small systems, always-accept policy can provide good enough practice instead of computing holdback levels.

In the overall, these findings, which are quite novel in the literature, mainly support many of the rational, but analytically unproved applications of the industry. They provide a promising basis for heuristic, but practical applications of transshipments in large systems when finding and applying optimal solutions have quite complex requirements.

References

- Agrawal V, Chao X and Seshadri S (2004). Dynamic balancing of inventory in supply chains. *European Journal of Operational Research* **159**(2): 296–317.
- Anupindi R, Bassok Y and Zemel E (2001). A general framework for the study of decentralized distribution systems. *Manufacturing and Service Operations Management* **3**(4): 349–368.
- Archibald TW (2007). Modelling replenishment and transshipment decisions in periodic review multilocation inventory systems. *Journal of the Operational Research Society* **58**(7): 948–956.
- Archibald TW, Sassen SAE and Thomas LC (1997). An optimal policy for a two depot inventory problem with stock transfer. *Management Science* **43**(2): 173–183.
- Archibald TW, Black DP and Glazebrook KD (2009). An index heuristic for transshipment decisions in multi-location inventory systems based on a pairwise decomposition. *European Journal of Operational Research* **192**(1): 69–78.
- Autos.com (2013). Breaking down the new car dealer cost, available at <http://www.autos.com/car-buying/breaking-down-the-new-car-dealer-cost>, accessed 17 April 2013.
- Çömez N, Stecke KE and Çakanyıldırım M (2012a). Multiple in-cycle transshipments with positive delivery times. *Production and Operations Management* **21**(2): 378–395.
- Çömez N, Stecke KE and Çakanyıldırım M (2012b). In-season transshipments among competitive retailers. *Manufacturing and Service Operations Management* **14**(2): 290–300.
- Das C (1975). Supply and redistribution rules for two-location inventory systems: One-period analysis. *Management Science* **21**(7): 765–776.
- Grahovac J and Chakravarty A (2001). Sharing and lateral transshipment of inventory in a supply chain with expensive low-demand items. *Management Science* **47**(4): 579–594.
- Gross D (1963). Centralized inventory control in multilocation supply systems. In: Scarf HE, Gilford DM and Shelly MW (eds). *Multistage Inventory Models and Techniques*. Stanford University Press: Stanford, CA, pp 47–84.
- Hanany E, Tzur M and Levran A (2010). The transshipment fund mechanism: Coordinating the decentralized multilocation transshipment problem. *Naval Research Logistics* **57**(4): 342–353.
- Hu X, Duenyas I and Kapuscinski R (2007). Existence of coordinating transshipment prices in a two-location inventory model. *Management Science* **53**(8): 1289–1302.
- Huang X and Sošić G (2010). Transshipment of inventories: Dual allocations vs. transshipment prices. *Manufacturing and Service Operations Management* **12**(2): 299–318.
- Korilis YA, Lazaar AA and Orda A (1999). Avoiding the Braess paradox in non-cooperative networks. *Journal of Applied Probability* **36**(1): 211–222.
- Kosansky A and Schaefer T (2010). Should you swap commodities with your competitors? *Supply Chain Quarterly* **2**: 42–47.
- Kranenburg AA and van Houtum GJ (2009). A new partial pooling structure for spare parts networks. *European Journal of Operational Research* **199**(3): 908–921.
- Krishnan KS and Rao VRK (1965). Inventory control in N warehouses. *Journal of Industrial Engineering* **16**(3): 212–215.
- Li R, Ryan JK and Zeng Z (2012). Supplier facilitated transshipments in a multi-retailer distribution system. Paper presented at the POMS 23rd Annual Conference, Chicago, IL, 20–23 April 2011, available at: http://www.pomsmeetings.org/confproceedings/025/FullPapers/FullPaper_files/025-0059.pdf.
- Özdemir D, Yucesan E and Herer YT (2006). Multi-location transshipment problem with capacitated transportation. *European Journal of Operational Research* **175**(1): 602–621.
- Robinson LW (1990). Optimal and approximate policies in multiperiod, multilocation inventory models with transshipments. *Operations Research* **38**(2): 278–295.
- Rong Y, Snyder LV and Sun Y (2010). Inventory sharing under decentralized preventive transshipments. *Naval Research Logistics* **57**(6): 540–562.
- Rudi N, Kapur S and Pyke DF (2001). A two-location inventory model with transshipment and local decision making. *Management Science* **47**(12): 1668–1680.
- Smartmoney.com (2013). The secret to year-end car buying, available at, http://www.smartmoney.com/spend/autos/how-to-score-discounts-at-year-end-car-sales/#article_tab_article, accessed 17 April 2013.
- Sošić G (2006). Transshipment of inventories among retailers: Myopic vs. farsighted stability. *Management Science* **52**(10): 1493–1508.
- Tagaras G and Cohen MA (1992). Pooling in two-location inventory systems with non-negligible replenishment lead times. *Management Science* **38**(8): 1067–1083.
- Talluri K and van Ryzin G (2004). Revenue management under a general discrete choice model of consumer behavior. *Management Science* **50**(1): 15–33.
- Van Wijk ACC, Adan IJBF and van Houtum GJ (2012). Approximate evaluation of multi-location inventory models with lateral transshipments and hold back levels. *European Journal of Operational Research* **218**(3): 624–635.
- Wee KE and Dada M (2005). Optimal policies for transshipping inventory in a retail network. *Management Science* **51**(10): 1519–1533.
- Yan X and Zhao H (2011). Decentralized inventory sharing with asymmetric information. *Operations Research* **59**(6): 1528–1538.
- Zhang J (2005). Transshipment and its impact on supply chain members' performance. *Management Science* **51**(10): 1534–1539.
- Zhao H, Desphande V and Ryan JK (2005). Inventory sharing and rationing in decentralized dealer networks. *Management Science* **51**(4): 531–547.
- Zhao H, Ryan JK and Desphande V (2008). Optimal dynamic production and inventory transshipment policies for a two-location make-to-stock system. *Operations Research* **56**(2): 400–410.
- Zhao X and Atkins D (2009). Transshipment between competing retailers. *IIE Transactions* **41**(8): 665–676.

Received 13 May 2013;
accepted 28 April 2014 after two revisions