International Journal of Production Research, 2017
Vol. 55, No. 21, 6470–6484, http://dx.doi.org/10.1080/00207543.2016.1249429

Evaluating warehouse strategies for two-product class distribution planning

Bradley Guthrie[a], Pratik J. Parikh[a]* and Nan Kong[b]

[a]*Department of Biomedical, Industrial and Human Factors Engineering, Wright State University, Dayton, OH, USA;* [b]*Weldon School of Biomedical Engineering, Purdue University, West Lafayette, IN, USA*

(*Received 12 October 2015; accepted 10 October 2016*)

Distribution networks often manage products with varying life cycles, where demand for some products is relatively stable throughout the year (basic) and the demand for others is short-lived (fashion). Beyond the coordination of inventory and transportation decisions, decisions at the warehouse must be considered as its resources are frequently shared by both product classes simultaneously. For this two-product class distribution planning problem, we focus on characterising three real-world distribution strategies observed in industry and evaluating them based on total distribution cost and warehouse measures (e.g. workforce plan and workload variation) against a benchmark ILS-based heuristic. Experimental results suggest that there are in fact strategies in industry that under specific system configurations may provide competitive solutions compared to the benchmark heuristic on large problem instances (e.g. 200 stores, 1000 products, and 28 days). Several managerial insights are derived to compare such distinct warehouse strategies and the corresponding impact on the network.

Keywords: distribution planning; warehouse; fashion products; industry strategy

1. Introduction

Real-world distribution networks frequently manage distributing not just one, but multiple products, each with a distinct life cycle. While the demand for some set of products may be steady throughout the year, the demand for others may only exist for a short amount of time (e.g. months, weeks). Based on their life cycles, such products can be categorised into two main product classes: long (often referred to as basic products) and short (also known as fashion products) (USOTA 1987), with each product class comprising a fairly large set of stock keeping units (SKUs); e.g. cell phones and trend apparel (short life cycle), and sugar and jeans (long life cycle).

For basic products, relatively low variance in demand makes it amenable for the distribution network to focus on developing cost-effective plans. In contrast, for fashion products, distribution plans are often aimed at providing maximum availability of products to consumers with focus on due date. In general, no matter which product class, it is vital that the supply chain fits the product (Fisher 1997).

This research was motivated by the operations at our industry partner, a US-based apparel distributor. They sell a wide variety of apparel products (basic and fashion) through their network of retail outlets and e-commerce channels. Their senior director of distribution informed us that the simultaneous flows of both basic and fashion products causes frequent operational challenges at their warehouse as both product classes utilise similar resources (e.g. workers, loading docks, material handling equipment). The inherent nature of the two products classes (i.e. one class exhibits a relatively stable flow, while the other arrives in bursts) thus leads to substantial workload variation throughout the planning horizon. This has been a big concern when hiring temporary workers on a daily basis (over the already hired permanent workers) and dealing with their low productivity, and sometimes even adopting expensive transportation modes to avoid missing the due date. These have negatively impacted the warehouse operating costs.

To address such challenges, many distribution networks (including our industry partner) have developed their own solution strategies that tend to be simple, quick, implementable and robust, all of which have been implied in the literature as important criteria for successful industry implementation (e.g. Barr et al. 1995; Viana, Sousa, and Matos 2005; Bartz-Beielstein and Preuß 2014). However, there has not been much consensus on which strategy is better and under what circumstances.

---

*Corresponding author. Email: pratik.parikh@wright.edu

This real-world observation motivated the following questions*: Which industry-deployed warehouse strategy is the best to solve the two-product class distribution problem? How close is it to the known best solution for this problem?* We attempt to address these questions by making the following contributions:

- characterise, for the first time in academic literature, several real-world distribution strategies that attempt to solve the two-product class distribution problem in a multi-period, multi-echelon setting;
- evaluate their performance against a benchmark metaheuristic on several measures (e.g. distribution cost, warehouse workforce plan and workload variation);
- identify problem attributes, which may suggest if and when one strategy is better than the others under specific conditions; and
- provide managerial insights to assist industry practitioners to appreciate the complex relationship between warehouse, inventory and transportation decisions.

We strongly believe our research will help bridge a perceived gap between academic research and industry practice.

2. Relevant literature

Although the role of warehouses in overall distribution planning has been acknowledged (Lambert, Stock, and Ellram 1998; De Koster, Le-Duc, and Roodbergen 2007), they have often not been integrated in the corresponding decision-making models. Most distribution models assume the warehouse as a node; e.g. the *inventory-routing problem* (Campbell et al. 1998), the *integrated inventory-distribution problem* (Abdelmaguid and Dessouky 2006), when considering inventory and transportation decisions (Çetinkaya, Tekin, and Lee 2006; Lei et al. 2006; Zhao, Chen, and Zang 2008). Only recently have models been proposed that integrate warehouse decisions alongside inventory and transportation; e.g. the *warehouse-inventory-transportation problem* (Sainathuni et al. 2014). Such integrated approaches have shown substantial benefits in not only total distribution cost savings, but also obtaining a relatively balanced workload at the warehouse allowing warehouse managers to plan and manage their workforce effectively. Note that almost all of the above distribution planning models considered only one class of products (typically, basic).

A parallel stream of research has emerged that primarily focuses on fashion products. Fisher, Rajaram, and Raman (2001) offer a heuristic that focuses on minimising lost sales, backorders and out-of-date inventory by determining fashion product replenishment order quantities. Weng and McClurg (2003) discuss the effect of coordination between suppliers and buyers when considering uncertain demand and delivery time. Patil, Avittathur, and Shah (2010) examine quantity discounts and transportation costs with respect to procurement, pricing and transportation decision-making. Caro and Gallien (2010) design a mixed-integer programme considering inventory and transportation decisions to maximise overall predicted sales across all Zara (Spain-based leader in fast fashion) stores. Recently, Mehrjoo and Pasek (2016) discuss risks in a quantitative manner that come natural to fast fashion distribution networks.

While it is implied in the literature that distribution networks must accommodate the unique life cycles of their products (Aitken, Childerhouse, and Towill 2003), we have not found any research that considers flows of both product classes, as well as the resulting complexities that emerge at the warehouse level. Further, while real-world distribution networks we have interacted with have employed various strategies at their warehouses based on insights from a combination of their prior experience and internal analysis, it remains unclear to decision-makers in industry the quality of solutions from such strategies. It then becomes critical to benchmark these industry strategies against optimal or near-optimal approaches.

Realising this gap between the academic literature and industry practice, we reiterate the focus of this research as follows: characterise real-world warehouse strategies for a two-product class distribution problem, evaluate them against a benchmark heuristic, identify network attributes for which a certain strategy performs better than the others, and derive managerial insights.

3. Distribution network representation

We now provide details of how we model the features and decisions in the two-product class distribution planning problem discussed above.

3.1 *Product class and planning horizon*

Let P and Q be the sets of basic and fashion products (each composed of distinct quantities), respectively. Each product ($p \in P$ or $q \in Q$) is delivered by a specific vendor, and shipped to a store $s \in S$. The planning horizon during which all processes occur at the warehouse is denoted as $t = 1, 2, \ldots T$, where T is the length of the planning horizon in days. We refer to the portion of the planning horizon during which fashion products arrive, are handled and eventually shipped out of the warehouse as the *fashion horizon*. The beginning of this horizon is designated as t_q^b, and the end is represented by the due date (t_{sq}^y) of fashion product q at store s. We denote the final day for which fashion products can *arrive* at the warehouse as t_q^e, which is calculated by the due date (t_{sq}^y) minus the maximum lead time across all stores (Max$\{L_s\}$), minus the processing time of fashion products in the warehouse (pt). Within the fashion horizon, there are distinct windows that designate the feasible range for inbound and outbound shipments. We will discuss these with respect to each strategy individually later in this section. Further, the expected demand for basic and fashion products at each store during each day is assumed to be known in advance.

3.2 *Warehouse activities*

While the activities at the warehouse are quite complex and interleaved, we take an aggregate approach given the relative focus of the distribution planning we consider; more towards tactical and less operational. With this idea, we use ‘putaway’ to represent all the activities that are involved from unloading, staging, and eventual putaway to the appropriate storage location. Similarly, we use ‘picking’ to include actual picking, sorting, staging and loading. Figure 1 illustrates a simplistic representation of a product’s flow through the warehouse over time.

Further, we assume inbound (outbound) trailers can be gradually unloaded (loaded) as practised by our industry partner. Because, as we have observed, it is difficult to frequently reschedule inbound shipments from vendors, gradually unloading trailers is one way our partner’s warehouse has been able to delay decisions to avoid excessive costs. Based on the receipt of inbound shipments on a given day, the warehouse unloads inbound trailers gradually over several days (typically, 3–4 days) in an effort to spread the workload and balance warehouse workforce to some extent. A similar situation occurs on the outbound side where the committed due date to the retail store must not be altered. Consequently, the picking and loading of the outbound trailers are often performed gradually over a few days to balance workload at the warehouse. Not treating the warehouse as a node in the distribution network (similar to previous work) and accounting for several operational details (albeit at some aggregate level) enables us to evaluate the interaction between warehouse, transportation and inventory decisions, simultaneously for both basic and fashion products.

3.3 *Warehouse workforce*

Based on the common theme derived from our interactions with several warehouse managers, two worker types (at the minimum) can be noticed; permanent and temporary. Permanent workers are similar to full-time workers with a 40-h work week, fairly skilled at performing various activities at the warehouse, and are often salaried (with or without additional benefits). Temporary workers are typically hired on a daily basis for no more than 8 h, may not have the same productivity as permanent workers (as they often are unfamiliar with the layout, processes, and material handling equipment), and are paid hourly. These workers are frequently utilised as a way to manage variations in worker hours at the warehouse, more often during the fashion horizon.

We refer to α and β_t to denote the number of permanent (required throughout the planning horizon) and temporary workers (required during a specific day, t), respectively. The associated costs are denoted as C^{α} and C^{β}, in \$/hr. Worker rates in items/hr are specified based on warehouse activity, with Λ_{put} referring to an aggregate rate for all putaway activities, and Λ_{pick} for picking. Because in real-world settings, most warehouse managers prefer to use a mix of permanent (higher skilled, but higher cost) and temporary (lower skill, but lower cost) workers, we use $\gamma \geq 0$ to represent the maximum allowable proportion of permanent workers that can be employed as temporary workers. Further, because the

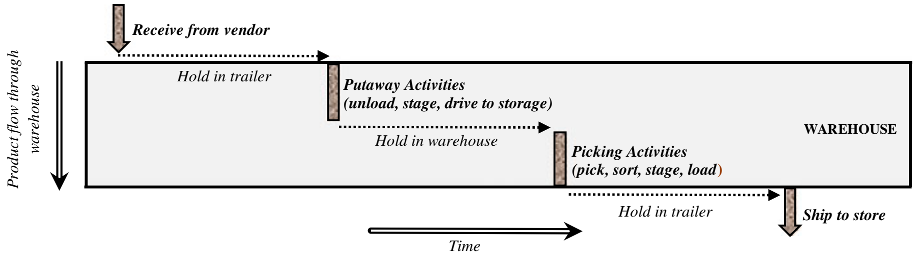

Figure 1. Pictorial representation of product flow through the warehouse.

hired temporary workers may not be identical across days when needed (depending on the availability at the third party that provides such workers), we can attach a productivity rate ($0 \leq \phi \leq 1$) for the temporary workers.

3.4 *Inventory and transportation*

Holding time at the warehouse begins once a product arrives on an inbound shipment at the warehouse until that product departs on an outbound shipment. Holding time at a store begins once an outbound shipment arrives at a store, and ends on the day of demand for that product (basic) or the due date (fashion). We let C_p^{hw} (C_{sp}^{hs}) represent the holding cost of basic product p at the warehouse (store s) in \$/item/day, and similarly C_q^{hw} (C_{sq}^{hs}) for fashion product q.

To model inbound transportation, we consider an inbound network with designated vendors supplying either basic or fashion products. On the outbound side, however, each store is expected to have demand for both product classes, which may cause consolidation of the two product classes on a shipment to a specific store when possible. We utilise a cost structure proposed in Sainathuni et al. (2014), which is composed of a fixed cost of a shipment and variable costs based on both distance (between source and destination) and weight of the shipment.

3.5 *Decisions, objective function and constraints*

The joint decisions to be made across both product classes include warehouse workforce (permanent and temporary), inbound shipment schedule from vendors, outbound shipment schedule to stores, and inventory levels at the warehouse and stores. The objective function is to minimise total distribution cost, which includes warehouse (picking + putaway), inventory (stores + warehouse) and transportation (outbound + inbound). Key constraints include meeting the store demand for basic products each day, not violating the fashion product due date, ensuring temporary workers are no more than the allowable limit, and non-splitting of fashion shipments.

4. Characterising real-world warehouse operational strategies

Having discussed how we characterise the distribution network and model the warehouse operations, we now briefly describe the four strategies to be compared in our experiments. More details about the algorithm and the moves are presented in Appendix 1.

4.1 *Fashion Release*

A specific version of the Fashion Release (FR) strategy is currently in practice at our industry partner. Essentially, their FR strategy is set up such that as soon as a fashion product q arrives from the vendor to the warehouse, and after incurring some processing time (pt) at the warehouse, all quantities of q will be shipped to each store with associated demand in the following day. While the company leadership projected that in so doing they may possibly avoid spikes in workload at the stores which otherwise would arrive just before the due date, it has possible drawbacks such as increased inventory cost at the stores and possible opportunity loss during outbound shipment consolidation to stores.

We generalise their strategy by introducing a parameter R_q, which denotes the number of days (beyond processing time) after which fashion product q must be shipped from the warehouse to the corresponding stores; $0 \leq R_q \leq (t_{sq}^y - t_q^b - pt - L_s)$, the length of the Fashion Window (FW). Note t_{sq}^y is the due date for product q at store s, t_q^b is the beginning of the fashion horizon, pt is the processing time at the warehouse, and L_s is the lead time to store s. The outbound day of product q can then be determined by $t_q^O = t_q^I + pt + R_q$, with the inbound day (t_q^I) constrained as $t_q^b \leq t_q^I \leq t_q^e - R$. If the warehouse manager chooses to use a single value for each fashion product q, then $R = R_q \forall q$. See Figure 2 for an illustration of this strategy.

In short, the FR strategy schedules the *outbound day* of fashion products based on their *inbound day* at the warehouse; the smaller value of R, the quicker the release from the warehouse (accounting for processing time).

4.2 *Fashion Holding*

A specific version of the Fashion Holding (FH) strategy was in practice prior to the industry partner adopting the FR strategy. In their version of FH, the arriving fashion products are held at the warehouse until the latest possible day of shipment to their corresponding stores (accounting for the lead time). The core concept was to change the look of the stores overnight, in sync with the beginning of a new fashion season. The implications of this was that holding fashion products until close to the due date increased the number of opportunities for consolidation on outbound shipments, as

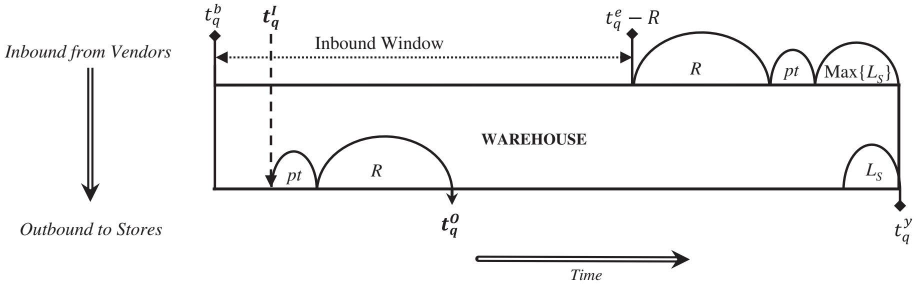

Figure 2. Illustration of the FR strategy.

well as lowered inventory levels at stores. However, sudden spikes in scheduled outbound shipments from the warehouse led to greater difficulties in managing workforce, and higher costs at the stores in terms of receiving all fashion shipments on one day and remodelling the store overnight.

We generalise this strategy by incorporating a parameter H_q, which represents the number of days prior to its due date (considering lead time) that product q must be shipped outbound; $0 \leq H_q \leq (t_{sq}^y - t_q^b - pt - L_s)$ and $H = H_q$, $\forall q$ if need be. The outbound day of product q going to store s is determined by $t_{sq}^O = t_{sq}^y - L_s - H$. Similar to FR, $t_q^b \leq t_q^I \leq t_q^e - H$. See Figure 3 for an illustration of this strategy.

In short, the FH strategy schedules the *outbound day* of fashion products based on their *due date* at stores; small values of H schedule the outbound day closer to the due date (minus lead time).

4.3 *Basic–Fashion Split*

Another strategy that was discussed with our industry partner, but not implemented yet, was to separate the flow of basic and fashion products along the horizon; we refer to the strategy as the Basic–Fashion Split (BFS) strategy. That is, depending on the anticipated number of fashion products (and the associated product quantities), the warehouse may manage the flow of all the basic products earlier in the horizon and dedicate its efforts towards managing the flow of fashion products later along the horizon. The idea is that if an appropriate split of the horizon were found, then the warehouse workload variation could be mitigated to some extent, possibly reducing warehouse worker cost. However, this reduction was thought to have negative implications via increased inventory levels of basic products at stores.

To evaluate the BFS strategy, we let t_{split} represent the day that splits the planning horizon into two sub-horizons, one for basic and one for fashion. We determine the value of t_{split} using the demand proportion of the basic products to the total demand; i.e. $t_{\text{split}} = \frac{\sum_{s\in S,p\in P,t\in T} D_{spt}}{\sum_{s\in S,p\in P,t\in T} D_{spt} + \sum_{s\in S,q\in Q,t\in T} D_{sqt}} * T$, where D_{spt} (D_{sqt}) refers to the demand of basic (fashion) product p (q) at store s in time t, and $t_q^b \leq t_{\text{split}} \leq t_q^e$ (split day cannot be outside of the fashion horizon). Figure 4 graphically shows how t_{split} determines the division of basic and fashion horizons.

In short, as opposed to structuring the flow of fashion products as in FR and FH, the BFS strategy focuses on *modifying the feasible windows for basic and fashion products* with respect to associated *product quantities*.

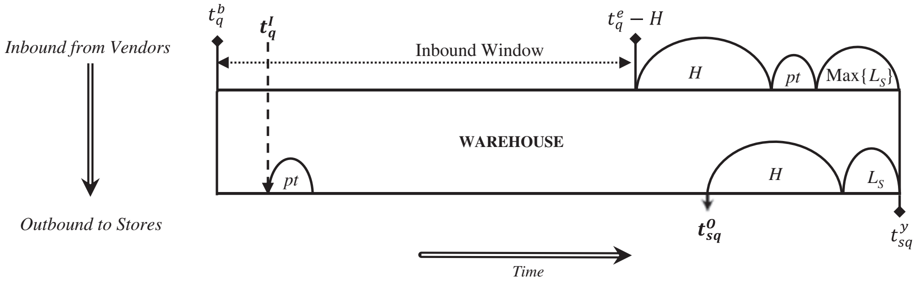

Figure 3. Illustration of the FH strategy.

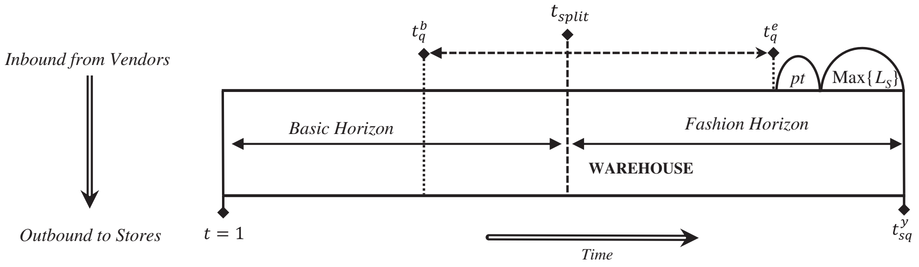

Figure 4. Illustration of the Basic–Fashion Split strategy.

4.4 *Fashion Window*

In order to adequately benchmark the three strategies discussed above, we devised the FW strategy at the warehouse. This strategy is based on an enhanced version of the ILS-based heuristic proposed recently in Sainathuni et al. (2014). Key enhancements include (i) generalisation via two parameters (R and H) and (ii) additional internal variables to allow flexibility in loading (unloading) trailers.

Figure 5 shows a schematic of the FW strategy. Let t_q^{EO} and t_{sq}^{LO} represent the earliest and latest outbound days for product q. Accordingly, $t_q^{EO} = (t_q^I + pt) + R$ and $t_{sq}^{LO} = (t_{sq}^y - L_s) - H$; the outbound window for fashion products then becomes $[t_q^{EO}, t_{sq}^{LO}]$. Similar to FR and FH, the fashion inbound window must adjust accordingly with R and H, where $t_q^b \le t_q^I \le t_q^e - H - R$. Individually, and as a sum, both R and H are both bounded by $t_{sq}^y - t_q^b - pt - L_{\text{Max}}$. Intuitively, $R = H = 0$ would provide the widest window for fashion outbound, and is expected to return lowest cost solutions. In short, the FW strategy structures the flow of fashion products through the warehouse by *modifying their feasible inbound and outbound windows,* based on product *inbound days* and *due dates* at stores.

5. Experimental evaluation of the strategies

We conducted a detailed experimental study to compare the three strategies against the benchmark FW strategy. We do not claim that the benchmark strategy provides an optimal strategy as it is a heuristic-based solution. However, a performance study in Sainathuni et al. (2016) revealed that our heuristic solutions were within 3% of the optimal solutions on small problem instances. Each of the above four strategies was simulated using a unique set of local search and perturbation moves that modify the quantity of product p (q) on day t scheduled for inbound, putaway, picking, and outbound. We used C# programming language for this using an Intel-i7 Quad-Core Desktop with 16 GB of RAM; please see Appendix 2 for run times.

All problem instances are industry-sized, comprising 20 vendors (10 for fashion and 10 for basic), 200 stores, 1000 products and 28-day planning horizon. Table 1 presents parameters that are fixed throughout our experiments.

5.1 *Warehouse cost contribution*

We first evaluated the warehouse cost as a percentage of the total distribution cost for each policy in 10 test instances, by varying the product mix ratio (proportion of basic to fashion products) and worker cost (permanent and temporary); see Figure 6. The fashion horizon was assumed to be two weeks for a four-week planning horizon.

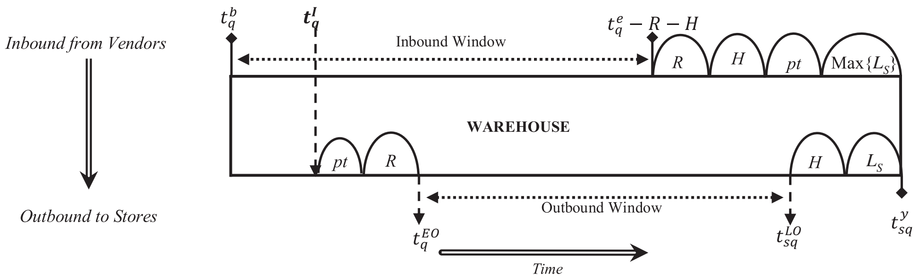

Figure 5. Illustration of the FW strategy.

Table 1. Parameters used in the experiments.

Parameter	Value
pt	1 day
L_s	0–2 days
t^y_{sq}	Day 28
t^e_q	Day 25
γ	1
ϕ	.8
Λ_{put}	1200 items/hr
Λ_{pick}	200 items/hr
C^{hw}_p and C^{hw}_q	.01 \$/item/day $\forall p, q$
C^{hs}_{ps} and C^{hs}_{qs}	.05 \$/item/day $\forall p, q, s$
R, H (FR,FH,FW)	0

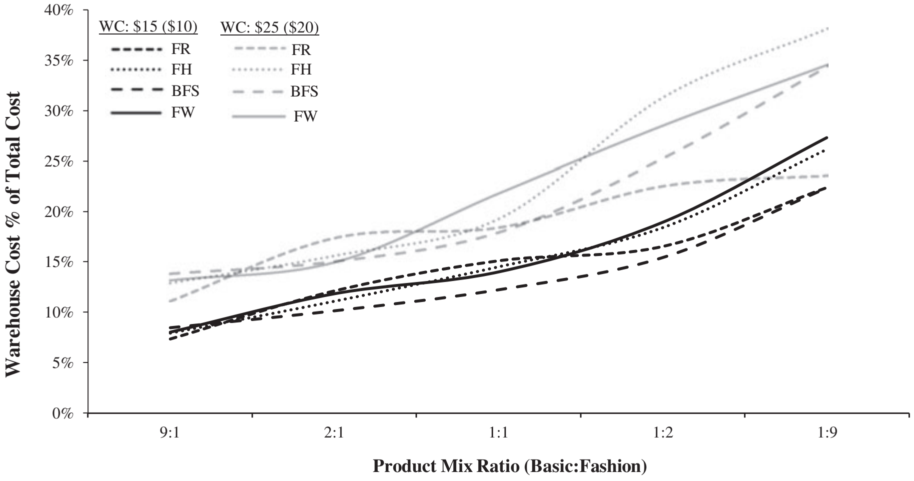

Figure 6. Comparison of warehouse cost contribution to total cost.

The two key takeaways here are the range and magnitude of warehouse cost contribution values. For each policy, there is a general trend of increasing warehouse contribution as the product mix ratio shifts toward more fashion products. This can be partially attributed to the changing proportion of fashion products that are required to be processed during the fashion horizon. For instance, with a product ratio of 1:9, the flow of 90% of the total product quantities (all fashion) through the warehouse during only 50% (two out of four weeks) of the planning horizon causes a relatively high variation in the workload, subsequently increasing the warehouse costs. Interestingly, with such a low product mix ratio, there are fewer basic products in the distribution network, resulting in low inventory levels at the warehouse and stores throughout the planning horizon, and reduced number of corresponding outbound shipments. The result is that with FW, the warehousing costs are 8–27% and 13–35% of total distribution cost given worker costs of \$15 and \$25, respectively, which are substantial.

5.2 *Warehouse workforce plan*

We then focused on analysing the workforce plan generated by the strategies. For this, we considered a product mix ratio of 1:1, two-week fashion horizon, and cost of permanent (temporary) workers as \$15/hr (\$10/hr) to compare warehouse workforce scheduling; see Figure 7. 'Permanent' and 'temporary' refer to the numbers of workers scheduled by type. 'Required' represents the actual number of workers (equivalent of permanent workers) required during that day; we further break it down by activity type ('putaway' and 'picking').

From the figure, it is evident that FW, FH and FR exhibit similar behaviour in workforce planning (i.e. temporary workers only in the fashion horizon), while the structure of BFS noticeably differs (i.e. temporary workers throughout

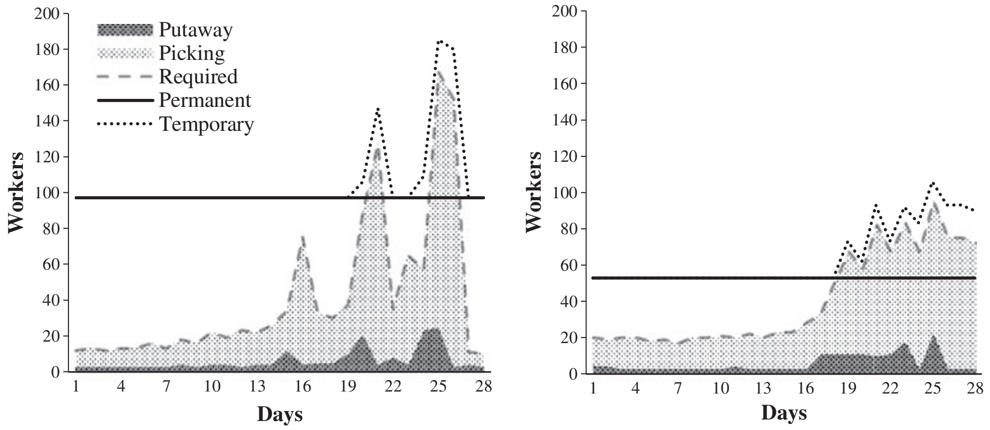

(a) Fashion Release (b) Fashion Holding

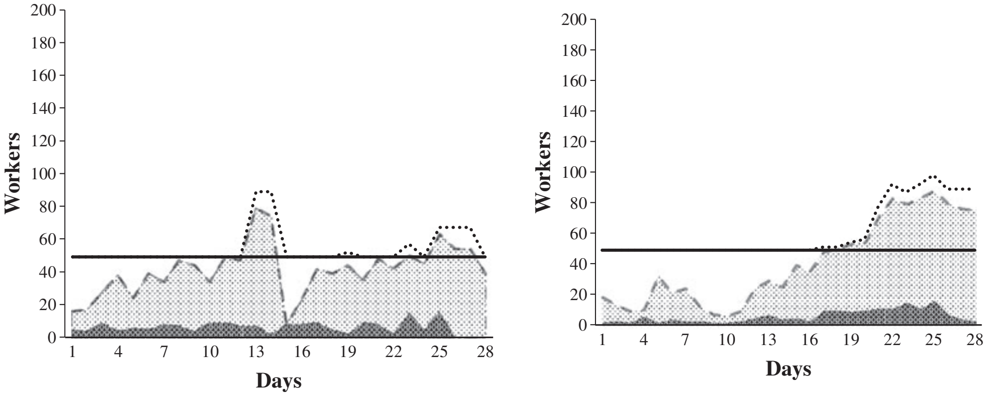

(c) Basic-Fashion Split (d) Fashion Window

Figure 7. Workforce plans generated by the four strategies.

the planning horizon). As predicted, BFS results in a relatively balanced required workload while maintaining an identical level of permanent workers (49) compared to FW. While FH is close behind (55), FR produces a considerably higher level of permanent workers (97) due to its poor ability to balance worker hours. With respect to the utilisation of temporary workers, we note that all four strategies take full advantage of temporary workers during peak days, limited only by parameter γ, as these workers are relatively inexpensive.

For each of the four strategies, the permanent worker line passes above the required worker line, which suggests idle time for such workers. In our experiments, this effect is attributed to setting $\gamma = 1.0$, which means that whenever temporary workers are required during a given day, their number is bounded by the level of the permanent workers (for the entire planning horizon). A higher value of γ, which may come at higher cost of training temporary workers and increased errors, would lead to lowering the permanent worker line making it closer to the 'required' line (which represents permanent equivalent workers) and alleviating some idle time; a lower value of γ has the opposite effect (results not shown). Note that idle time of workers at a warehouse is often used to accomplish tasks such as cleaning of work areas, reorganisation, and other supporting activities. Finally, we observed that the 'required' line is always below the 'temporary' worker line as the productivity $\phi = 0.8$; the two lines would align when $\phi = 1.0$.

5.3 *Variation in total worker hours*

We next analysed the variation in the worker hours, from which the above workforce plans were derived, to further understand how the three strategies compared to the benchmark FW strategy. We represent workload variation by displaying the range of %-difference of worker hours required during each day from the average warehouse hours (across the planning horizon).

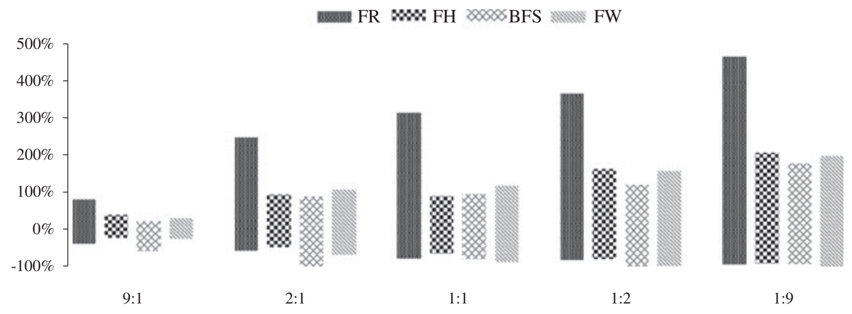

Figure 8. Variation in total warehouse hours across various product mix ratios.

From Figure 8, we noticed that all strategies demonstrate an increasing trend of variation with a decreasing product mix ratio. This is intuitive as the more the basic products, the more opportunities there are to effectively balance the workload at the warehouse; an observation similar to Sainathuni et al. (2014). However, with the inclusion of fashion products, which impose a strict due date and must be handled in a reasonable time frame, the variation in the workload is expected to increase.

We noticed further that FR demonstrates considerably higher sensitivity to product mix ratio than the other strategies, again due to fashion warehouse hours being constrained to their respective inbound and outbound days. With the ratio of 1:9, FR results in an absolute difference of 267% from FW, compared to a difference of 63% with the 9:1 ratio. Also, we note that none of FW, FH and BFS consistently provides the smallest range among the four strategies across the considered product mix ratios. As stated earlier, a higher degree of workload variation at the warehouse greatly deters the planning of workforce at the warehouse, often leading to underutilised worker hours or not enough workers, both leading to higher costs and/or affecting service downstream.

6. Sensitivity to system parameters

We now discuss how sensitive the solutions generated in the above experiments are to changes in system parameters. We first analysed the impact of *varying product mix ratio*, *fashion horizon length*, and *worker costs* for all strategies. In addition to the below, the sensitivity of FW, FR and FH strategies to values of *R* and *H* is presented in Appendices 3 and 4.

Table 2 presents the parameters, along with their associated number of levels and values. This resulted in 20 experiments we evaluated using each of the four strategies; i.e. a total of 80 experiments. Figure 9 displays for each combination the ratio of each strategy's total cost to that of FW. Also displayed is the total cost (in $) of FW for each parameter combination.

Focusing on the cost ratios for each strategy, we noticed that FH performs fairly well over all parameter combinations, consistently within 10% (1.1 times) of the FW strategy. As previously indicated, the approach of holding all fashion products until their last possible shipment date benefits both outbound consolidation and store inventory levels. Meanwhile, the approach contributes to relatively low variation in warehouse hours. In contrast, the FR strategy behaves much worse than FW (8–123%), more prominently as the product mix ratio reduces. This finding can be associated with our earlier observations that FR provides little opportunity for consolidation of fashion products on the outbound side, as well as a poorly balanced workload. Note that a tighter fashion horizon length (one vs. two weeks) has a positive impact on the performance with the FR strategy, which is largely because both FR and FW strategies have limited consolidation opportunities across a shorter horizon than longer.

Similar to FR, BFS also displays an increasing total cost ratio for a decreasing product mix ratio; ranging from 1.01 to 1.27 times of FW. We note that for BFS, at least two shipments must be sent from the warehouse to each store, one

Table 2. Parameters varied in experiment and their respective number of levels and values.

Parameter	Levels	Values
Product mix ratio (basic:fashion)	5	1:9, 2:1, 1:1, 1:2, 1:9
Fashion horizon length (weeks)	2	1, 2
Permanent (temporary) worker cost ($/hr)	2	$15($10), $25($20)

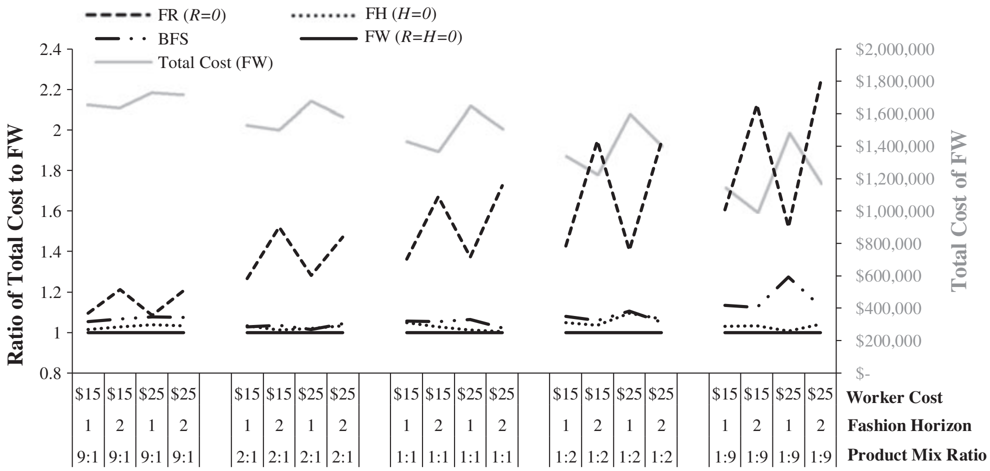

Figure 9. Sensitivity of total cost (relative to FW) to product mix ratio, length of fashion horizon, and worker cost.

carrying only basic products during the initial part of the planning horizon, and the other carrying only fashion products during the latter part. For a larger proportion of basic products (e.g. 9:1), this constraint does not separate BFS too far from FW (1.05–1.08). However, as the product mix ratio decreases, FW tends to schedule a declining number of outbound shipments (as low as one shipment per store), increasing the separation of BFS from FW up to 1.27. Further, it is also for lower product mix ratios where increasing the length of the fashion horizon improves the performance of BFS relative to FW (due to relatively lower warehouse costs). All six instances of a two-week fashion horizon with product mix ratios varying from 1:1 through 1:9, resulted in improved performance over their respective one-week counterparts, with a range 0.3–11.4%.

Notice the decreasing nature of the total cost for FW (grey lines in Figure 9) as product mix ratio decreases (more fashion). The intuition is that with more fashion products flowing through the warehouse, there is considerably more variation in the warehouse workload (as illustrated in Section 5.3), thus increasing warehouse costs. Though found to be true, per our discussion of warehouse cost contribution, lower product mix ratios also resulted in decreased store inventory levels, as well as decreased outbound shipments, offsetting the warehouse cost increase; in fact, lowering the total cost.

7. Conclusions and future research

Most distribution networks are challenged with effectively and efficiently managing flows of two or more classes of products with differing life cycles. For a specific two-product class distribution problem, we focused on evaluating various warehouse strategies that decision makers in the industry have implemented. We considered the case of apparel distribution given our close ties with one Midwest US distributor. For such multi-period, multi-echelon, distribution networks, we characterised three warehouse strategies, Basic-Fashion Split (BFS), Fashion Release (FR), and Fashion Holding (FH). We then compared them against our benchmark ILS-based heuristic (i.e. the FW strategy). Several measures were evaluated, including total distribution cost, warehouse workforce plan and workload variation. The following managerial insights were drawn from our study:

- The FH strategy appears to generate distribution plans very close to the benchmark FW strategy (typically <5%) even for higher percentage of fashion products.
- Both the BFS and FR strategies tend to become less competitive when the number of fashion products increase.
- The percent contribution of warehouse cost to the total distribution cost increases (up to 27–35% depending on worker cost) with a higher proportion of fashion products.
- Both FH and BFS strategies generate competitive workforce plans compared to the FW strategy.
- Variation of total warehouse hours increases with a higher proportion of fashion products.

The implications of our findings can be significant. The quantification of how such strategies compare to benchmark approaches proposed by academics (e.g. the FW strategy) can provide industry practitioners deeper understanding and insights into the impact of their chosen strategy on warehouse, inventory, and transportation decisions. Of note, our industry partner had employed the FH strategy in the past, and is currently employing the FR strategy–two entirely distinct strategies. Our findings suggest that, purely from a distribution planning perspective, what they did in the past seemed to be far better in terms of total cost than what they are doing currently. When we shared our findings with them, they were intrigued as this was not what they expected, and are in discussions with us to further understand these findings. As managing multiple product classes is becoming a norm for modern-day distribution networks, it is crucial for industry and academics to work collaboratively in not only evaluating real-world strategies against optimisation-based approaches, but also to leverage the insights generated from optimal solutions to devise simple, quick, robust strategies that have a high probability of acceptance and implementation by industry.

There are many possibilities for future research. Splitting shipments (inbound and outbound) is often practised, and thus the resulting impact on inventory levels (warehouse and stores) and corresponding workload (putaway and picking activities) is intriguing. In addition, the length of exposure of fashion products at the store and product pricing are obviously key factors if the objective is total revenue. Accounting for these would provide for an even more comprehensive understanding and comparison of these strategies.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This research was partially supported by Division of Civil, Mechanical and Manufacturing Innovation [grant number CMMI#1235061], [grant number CMMI#1235283].

References

Abdelmaguid, T. F., and M. M. Dessouky. 2006. “A Genetic Algorithm Approach to the Integrated Inventory-distribution Problem.” *International Journal of Production Research* 44 (21): 4445–4464.

Aitken, J., P. Childerhouse, and D. Towill. 2003. “The Impact of Product Life Cycle on Supply Chain Strategy.” *International Journal of Production Economics* 85 (2): 127–140.

Barr, R. S., B. L. Golden, J. P. Kelly, M. G. Resende, and W. R. Stewart Jr. 1995. “Designing and Reporting on Computational Experiments with Heuristic Methods.” *Journal of Heuristics* 1 (1): 9–32.

Bartz-Beielstein, T., and M. Preuß. 2014. “Experimental Analysis of Optimization Algorithms: Tuning and Beyond.” In *Theory and Principled Methods for the Design of Metaheuristics*, edited by Y. Borenstein and A. Moraglio, 205–245. Berlin: Springer.

Campbell, A., L. Clarke, A. Kleywegt, and M. Savelsbergh. 1998. “The Inventory Routing Problem.” In *Fleet Management and Logistics*, edited by T. G. Crainic and G. Laporte, 95–113. Norwell, MA: Kluwer.

Caro, F., and J. Gallien. 2010. “Inventory Management of a Fast-fashion Retail Network.” *Operations Research* 58 (2): 257–273.

Çetinkaya, S., E. Tekin, and C. Y. Lee. 2006. “A Stochastic Model for Joint Inventory and Outbound Shipment Decisions.” *IIE Transactions* 40 (3): 324–340.

De Koster, R., T. Le-Duc, and K. J. Roodbergen. 2007. “Design and Control of Warehouse Order Picking: A Literature Review.” *European Journal of Operational Research* 182 (2): 481–501.

Fisher, M. L. 1997. “What is the Right Supply Chain for Your Product?” *Harvard Business Review* 75: 105–117.

Fisher, M., K. Rajaram, and A. Raman. 2001. “Optimizing Inventory Replenishment of Retail Fashion Products.” *Manufacturing & Service Operations Management* 3 (3): 230–241.

Lambert, D. M., J. R. Stock, and L. M. Ellram, eds. 1998. *Fundamentals of Logistics Management*. Singapore: McGraw-Hill.

Lei, L., S. Liu, A. Ruszczynski, and S. Park. 2006. “On the Integrated Production, Inventory, and Distribution Routing Problem.” *IIE Transactions* 38 (11): 955–970.

Mehrjoo, M., and Z. J. Pasek. 2016. “Risk Assessment for the Supply Chain of Fast Fashion Apparel Industry: A System Dynamics Framework.” *International Journal of Production Research* 54 (1): 28–48.

Patil, R., B. Avittathur, and J. Shah. 2010. “Supply Chain Strategies Based on Recourse Model for Very Short Life Cycle Products.” *International Journal of Production Economics* 128 (1): 3–10.

Sainathuni, B., P. J. Parikh, X. Zhang, and N. Kong. 2014. “The Warehouse-Inventory-Transportation Problem for Supply Chains.” *European Journal of Operational Research* 237 (2): 690–700.

Sainathuni, B., B. Guthrie., P. J. Parikh, and N. Kong. 2016. “Distribution Planning for Products with Varying Life Cycles.” *Working Paper.*

USOTA. 1987. *The US Textile and Apparel Industry: A Revolution in Progress*. Washington, DC: US Office of Technology Assessment.

Viana, A., J. P. Sousa, and M. A. Matos. 2005. “Constraint Oriented Neighbourhoods – A New Search Strategy in Metaheuristics.” In *Metaheuristics: Progress as Real Problem Solvers*, edited by Y. Borenstein and A. Moraglio, 389–414. Berlin: Springer.

Weng, Z. K., and T. McClurg. 2003. “Coordinated Ordering Decisions for Short Life Cycle Products with Uncertainty in Delivery Time and Demand.” *European Journal of Operational Research* 151 (1): 12–24.

Zhao, Q. H., S. Chen, and C. X. Zang. 2008. “Model and Algorithm for Inventory/Routing Decision in a Three-Echelon Logistics System.” *European Journal of Operational Research* 191 (3): 623–635.

Appendix 1. Brief description of the algorithms that emulate the four strategies

We briefly explain how the FW strategy is simulated referencing Figure A1(d); the three other strategies are simulated in a similar way. The algorithm for the FW strategy uses an initial solution similar to FH, where the basic solution is randomised and the fashion solution schedules all fashion product SKUs to be shipped outbound to stores at the latest possible day considering lead time. The algorithm begins by holding the schedule of all fashion products static, and only modifies the schedule of basic products, through

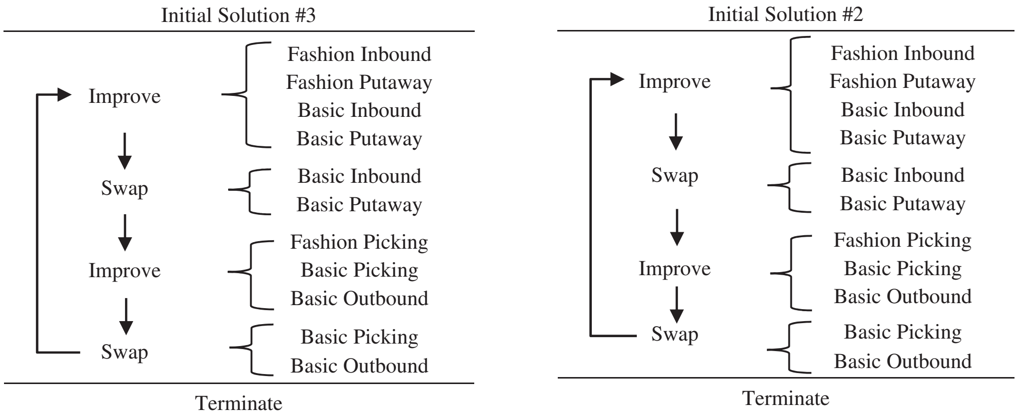

(a) Fashion Release **(b) Fashion Holding**

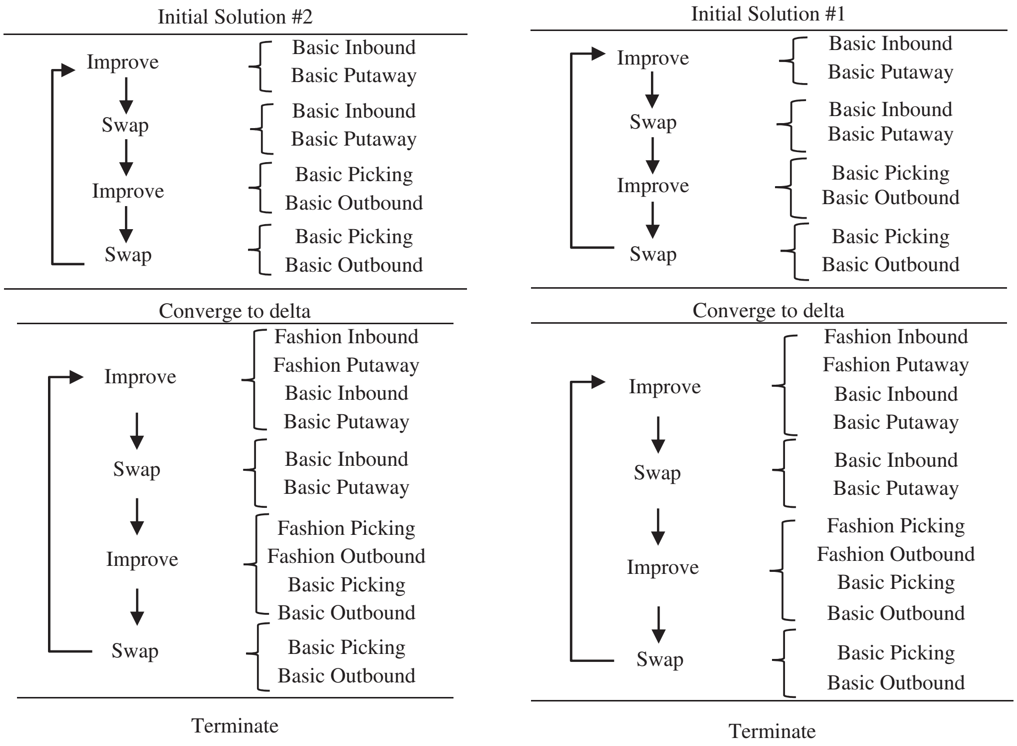

(c) Basic-Fashion Split **(d) Fashion Window (Benchmark)**

Figure A1. Complete set of moves and initial solutions required to simulation the four strategies.

local search (improve) and perturbation (swap) moves, until convergence. The algorithm then freely alters both basic and fashion solutions, and terminates upon a second convergence. The convergence factor was set at 5×10^{-9} in order to successfully provide a benchmark solution; for all other strategies a value of 0.0025 was found sufficient.

Table A1 summarises the moves used in the algorithms. Each description provides the source day (SD) of the move (day that product quantities are moved from), the destination day (DD) of the move (day that product quantities are moved to), and also the type of move employed.

Table A1. Descriptions of moves and initial solutions required to simulation the four strategies.

Initial Solution Name	Description of the Initial Solution
Initial Solution #1	Random Solution, Feasible to Basic-Fashion Split Calculation
Initial Solution #2	Random Solution, Feasible to Fashion Holding Outbound Constraint
Initial Solution #3	Random Solution, Feasible to Fashion Release Outbound Constraint
Move Name	Description of the Move
Improve Fashion Inbound	SD: Random; DD: Random; Type: Advance & delay shipments; putaway, picking, outbound moves to obtain feasibility
Improve Fashion Putaway	SD: Maximum required putaway workers; DD: Minimum required putaway workers; Type: Advance & delay putaway hours; picking move to obtain feasibility
Improve Fashion Picking	SD: Random; DD: Minimum required picking workers; Type: Advance & delay picking hours
Improve Fashion Outbound	SD: Random; DD: Random; Type: Consolidate Shipments
Improve Basic Inbound	SD: Random; DD: Random; Type: Consolidate Shipments
Improve Basic Putaway	SD: Maximum required putaway workers; DD: Minimum required putaway workers; Type: Split basic putaway hours
Improve Basic Picking	SD: Maximum required putaway workers; DD: Minimum required putaway workers; Type: Split basic picking hours
Improve Basic Outbound	SD: Random; DD: Random; Type: Consolidate Shipments
Swap Basic Inbound	SD: Random; DD: Random; Type: Swap all basic inbound shipments between days for a random number of vendors
Swap Basic Putaway	SD: Random; DD: Random; Type: Swap all basic putaway hours between days for a random number of vendors
Swap Basic Picking	SD: Random; DD: Random; Type: Swap all basic picking hours between days for a random number of stores
Swap Basic Outbound	SD: Random; DD: Random; Type: Swap all basic outbound shipments between days for a random number of stores

Appendix 2. Average run times

Run times associated with each strategy were recorded in the experiments from Section 5.1. Table A2 shows the average run times (hh:mm:ss) across four instances (varying worker cost and fashion horizon length) for the extreme product mix ratios we considered in this study (i.e. 1:9 and 9:1). Run times for all other ratios fell within these ranges.

While the FH and FR strategies alternate in producing the lowest average times for the product mix ratios analysed, the average values for BFS are considerably longer. This makes sense considering BFS has no structured fashion solution as in FR and FH. Note that the distribution planning problem is usually tactical in nature, such run times are fairly reasonable in industry; typically, such decisions are made once every few months and the algorithms are allowed to run overnight to achieve the best possible solution.

Also note that in order to obtain the best possible (near-optimal, if not optimal) solution from the FW strategy, the stopping criterion (tolerance based) was kept very small, and so the higher run times.

Table A2. Average run times (hh:mm:ss) of strategies for product mix ratios of 1:9 and 9:1.

Strategy	Product mix ratio – 1:9	Product mix ratio – 9:1
FH	00:41:47	04:47:51
FR	00:34:55	07:34:40
BFS	00:50:36	14:56:01
FW	17:51:35	72:13:16

Appendix 3. Sensitivity of FW strategy to *R* and *H*

Since the FW strategy can result in a number of variations based on the specific values of R and H, we analysed 10 additional instances and compared them against the benchmark (i.e. $R = H = 0$). For five different product mix ratios, we calculated the ratio of

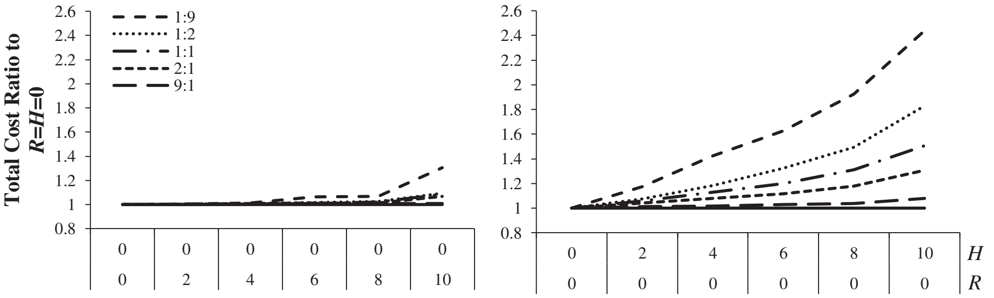

(a) Sensitivity to *R* (b) Sensitivity to *H*

Figure A2. Total cost sensitivity of FW to the two parameters: (a) *R* and (b) *H*.

the total cost for each combination of *R* and *H* to that obtained for $R = H = 0$. The length of the fashion horizon was two weeks and worker cost was set to \$15/hr for permanent and \$10/hr for temporary.

Figure A2 presents total cost ratios (five different product mix ratios) for combinations of *R* and *H* where *H* is held constant at 0 (a) and *R* is held constant at 0 (b). The solid black line (in both figures) with total cost ratio equal to 1.0 represents the ratio of $R = H = 0$ variation to itself. Two key observations should be noted. First, the FW strategy is substantially more sensitive to increases in the value of *H* (which shrinks the fashion horizon from the tail end) than increasing *R* (which shrinks it from the front end). For instance, in the case of a product mix ratio of 1:1, increasing *R* from 0 to 10 with $H = 0$ increases the cost ratio to 1.05 (5%) while increasing *H* from 0 to 10 with R = 0 generates an increase to 1.52 (52%). This finding can be explained by realising that when $H = 10$, the inbound window for fashion products is now limited to just 1 day (day 15) and the outbound window is limited to three days (days 16, 17, 18). As a result, warehouse costs are substantially affected (due to the lack of days to distribute worker hours), along with store inventory costs (i.e. fashion products arrive 10 days prior to their due date (day 28)).

Second, the magnitude of trends in total cost ratio for an increasing *R* and *H* differs for each product mix ratio. Consider the $R = 0$ and $H = 10$ combination in Figure A2(b). While the total cost for a product mix of 9:1 (more basic, less fashion) is only 1.08 times as much as that of $R = H = 0$, this ratio is 2.42 for a product mix of 1:9 (less basic, more fashion). This finding is intuitive because both *R* and *H* determine the boundaries for the fashion inbound and outbound windows. Shrinking these windows (by either increasing *R* or *H* or both) will have a substantial impact on the total cost ratio for product mixes with a larger proportion of fashion products than smaller.

Appendix 4. Sensitivity of FH and FR strategies to *R* and *H*

The Fashion Holding (FH) and Fashion Release (FR) strategies each have a single parameter that can generate multiple variations of each strategy. For FH, the parameter *H* specifies how many days before a fashion product's due date at a store it is to be shipped outbound. For FR, *R* indicates how many days following inbound a fashion product is to be shipped outbound. Although it would be possible to assign a unique parameter value to each fashion product type, we retain simplicity for our model. Thus, $R = R_q$ and $H = H_q \forall q$, and both are constrained from 0 to $t^e_q - t^b_q$. Figure A3 displays the generated total cost values of FH and FR for 5 different product mix ratios (1:9–9:1) and 6 different values of *R* and *H* (0 to 10 in steps of 2). The length of the fashion horizon was two weeks and worker cost was set to \$15/hr for permanent and \$10/hr for temporary.

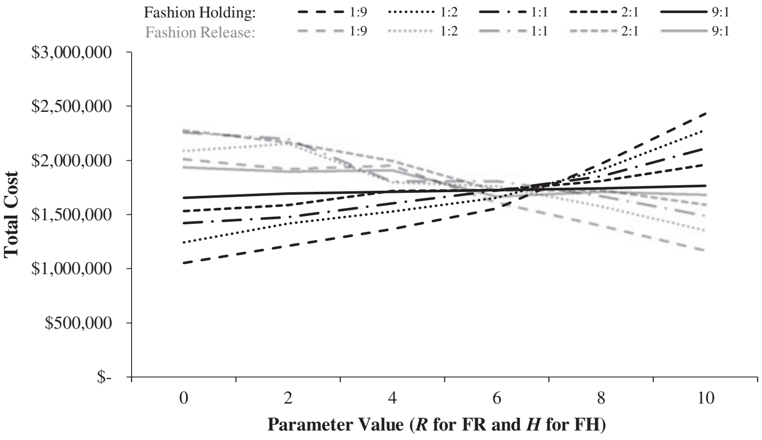

Figure A3. Sensitivity of FR strategy to *R* and FH strategy to *H*.

We first note the obvious trends for FH (increasing) and FR (decreasing) as R and H increase to 10 for their respective strategies. For FH, this finding suggests that holding fashion products at the warehouse as close to their due dates as possible (i.e. $H = 0$) is indeed the most favourable option. However, in practice, lead times to stores are by nature variable, and thus it is usually practical to provide a cushion in the outbound transportation strategy. Observing the trends of FH further, we note that for high product mixes (e.g. 9:1; more basic than fashion), the total cost is less affected by R and H values. However, for low product mixes (e.g. 1:9; more fashion than basic), the total cost increases with increasing values of H. The latter is because of the substantial increases in the cost due to sending fashion products earlier to the stores (impacting inventory holding costs; store holding is higher than warehouse holding) and the shorter time in which the warehouse has to turn around these fashion products (impacting warehouse worker costs).

A reverse trend is observed with FR where the total cost is much higher for low values of R, but gradually decreases as R increases. We also observe that all product ratio instances intersect with their respective FH instances around $R = H = 6$. In other words, after this point, it is now more favourable to schedule the outbound shipment of fashion products based on their *inbound day* rather than their *due date at stores*. That is, after $R = H = 6$, FR appears to be more favourable for all product mixes than FH.