On the value of multi-echelon inventory management strategies for perishable items with on-/off-line channels

Transportation Research Part E 180 (2023) 103354

Contents lists available at ScienceDirect

Transportation Research Part E

journal homepage: www.elsevier.com/locate/tre

On the value of multi-echelon inventory management strategies for perishable items with on-/off-line channels

Daniele Giovanni Gioia [a,b,*], Stefan Minner [b,c]

[a] Department of Mathematical Sciences, Politecnico di Torino, Turin, 10129, Italy
[b] Logistics & Supply Chain Management, School of Management, Technical University of Munich, Munich, 80333, Germany
[c] Munich Data Science Institute (MDSI), Technical University of Munich, Garching, 85748, Germany

ARTICLE INFO

Dataset link: https://github.com/DanieleGioia/PerishableMEC

Keywords:
Perishable inventory
Simulation-based optimization
Sales channels
Multi-echelon inventory management

ABSTRACT

The progress of digitization makes the integration of online and offline sales channels increasingly necessary for retailers. Multichannel and omnichannel multi-echelon networks are gradually more common in responding to customer demands, but their complexity makes the optimization of replenishment and item allocation policies among different channels challenging, especially if products have a short shelf life, as in the case of food retailers, where customer behavior (e.g., first-/last- in-first-out selection) also plays a role. It is not always possible to solve this problem exactly and heuristics are required. We propose a dynamic model and jointly optimize allocation and replenishment policies in the case of perishable goods with stochastic demand, uncertainty in customer selection preferences, and fixed lead times. We study complexity and structure of optimal policies. Furthermore, we explore several intuitive generalizations of base-stock policies over multi-echelon networks, analyzing the effect that potential correlations and imbalances in demand volumes across channels generate on the heuristics and identifying the pros and cons of such solutions. Results show that inventory-pooling effects in multi-echelon models for perishable items are often combined with the allocation of fresher products to offline channels. Generalizations of the well-known constant-order or base-stock policies can be a viable solution that generates benefits and increases system flexibility. They advantageously leverage negative channel correlation, but in the case of unbalanced demand distributions, increased offline demand can impoverish the quality of some heuristics.

1. Introduction

The retail industry is rapidly evolving towards a combination of online and offline channels. Digitization is proceeding apace, and the number of customers buying online is growing. Moreover, the recent Covid-19 pandemic has increased product requests through online channels (UNCTAD, 2022). One exception, however, is the grocery sector. Until a few years ago, it did not receive much customer pressure and only recently has been moving towards multichannel and omnichannel experiences (Eriksson et al., 2019b).

The growing importance of online channels generates new challenges in meeting diverse customer needs, requiring new logistics models and supply chain integration (Bell et al., 2014). For example, Hübner et al. (2016) gather empirical data of over 60 international executives from retail and logistics enterprises (with 16% from the grocery sector), analyzing the characteristics of

* Corresponding author at: Department of Mathematical Sciences, Politecnico di Torino, Turin, 10129, Italy.
E-mail addresses: daniele.gioia@polito.it (D.G. Gioia), stefan.minner@tum.de (S. Minner).

https://doi.org/10.1016/j.tre.2023.103354
Received 17 March 2023; Received in revised form 22 October 2023; Accepted 2 November 2023
Available online 7 November 2023
1366-5545/

retailers who deal with multichannel configurations. They find that retailers with multiple channels integrate inventory systems to enhance flexibility. However, grocery retail requires further reflection. Indeed, while in non-food logistics the integration of different channels is generally considered an advantage, in the case of grocery this is not evident (Wollenburg et al., 2018). Several requirements need special attention. In particular, unlike technology and fashion online markets, orders generally consist of a large number of items, and perishability of goods plays a key role. Nevertheless, some particular advantages make perishables suitable and profitable in omnichannel operations. In fact, Song et al. (2021) point out how fresh products in physical stores are subject to accelerated spoilage due to selection by customers and how the advantage given by ad-hoc product selection through online channels can reduce extra losses and increase profits.

In the case of perishables in food retailing, several solutions have been implemented to handle the new opportunities of the online market. When dealing with small portions of demand, some retailers fulfill all orders through the inventory of the physical store. However, this solution creates cannibalization and low service levels due to timing misalignments for demands for the same items through different channels and differences in customers' behavioral patterns (Wollenburg et al., 2018). This solution is therefore feasible only in the case of low online demand. Other approaches use distribution centers dedicated to home delivery, where there is no interaction between different outlets. In particular, several grocery retailers, in order to cope with differences in demand volume and interest more in picking efficiency than market considerations, have decided to handle online orders through special types of distribution centers called *Online Fulfillment Centre* (OFCs) (Eriksson et al., 2019b).

The aging of perishables in physical inventories depends on the employed order policy and, in the case of distribution centers, on dispatch policies (Akkaş and Honhon, 2022). Moreover, food produced for human consumption is lost or wasted when appropriate stock management at the retail level is not in place (Rezaei and Liu, 2017). For this reason and with a view to new methodologies for managing different channels in food retailing, we model, simulate and analyze inventory control for a single perishable product of integrated multichannel networks with both online and offline customers, comparing them with separate single-echelon models in which there is no interaction between the inventories of different retailers. We investigate the benefits of pooling inventory of various sales channels from the retailer's perspective, optimizing their joint management, and studying their complexities. Initially, simple structures from which optimal allocation and order policies can be derived by value iteration are considered. Then, we present models with larger demand volumes, more complex demand distribution types, and longer deterministic lead times, which require heuristic approaches. In particular, we use parametric strategies that are well-known in the perishable goods industry, both in single- (Haijema and Minner, 2016, 2019; Gioia et al., 2023) and multi- (Van der Vorst et al., 2000; Van der Heijden et al., 1997; Noordhoek et al., 2018) echelon networks, also known to be asymptotically optimal on certain configurations (Bu et al., 2023). Demand is assumed stochastic and we consider different characteristics (i.e., coefficient of variation, correlation). As the system becomes more complicated, the role of the flexibility of heuristics is examined. We evaluate the proposed policies in terms of profits, waste, and number of stock-out events. We place special emphasis on cases where orders from suppliers are constrained by a constant and fixed flow of items, where only the presence of a shared distribution center allows a flexible allocation between channels.

Despite ongoing developments of multichannel structures for joint online and offline sales of perishable food, a study that extends and investigates well-known approaches with lookup tables and parametric policies of classic inventory management literature is lacking. This study aims at answering the following questions:

- What benefits does pooling inventories for managing different online and offline channels produce with perishable items?
- Since multiple channels may entail correlated demand distributions, what effects are observed when online and offline channels share the inventory in presence of correlation?
- How does the subdivision of demand among the channels (i.e., more online or more offline) affect a model in which a central distribution center is shared, and what precautions are necessary for generating heuristics?
- Sometimes, regardless of the existence and computation of optimal replenishment policies, it is only possible to set up contracts with constant quantities with producers. In this case, a distribution center would give the opportunity to implement flexible allocation strategies after receiving a constant amount of supplies. How beneficial is this flexibility?

The paper is organized as follows. Section 2 reviews the related literature of inventory management for perishable items. Section 3 introduces the assumption and the dynamics of the model. Moreover, it defines the policies to solve the joint order and allocation problem. Section 4 reports the main experimental results, whereas Section 5 discusses conclusions and possible future work.

2. Related literature

We address two problems that, in the literature of inventory management for perishable items, are treated both jointly and independently. Specifically, regardless of the presence of a distribution center, we have to decide about the quantity of perishable goods to order from the supplier. On the other hand, the possible presence of a distribution center and a multi-echelon structure implies decisions about the number of items to dispatch in the internal shipping process.

The literature on inventory systems for perishable items is quite extensive, and several literature reviews dealing jointly with perishability, obsolescence, and deterioration are available. Starting from the first comprehensive review on the problem of determining suitable ordering policies for both fixed lifetime perishable items and inventory subject to continuous decay made by Nahmias (1982), Goyal and Giri (2001) review studies from the early 1990s to 2000. Bakker et al. (2012) update the state of the art with studies from 2001 to 2011, followed by Janssen et al. (2016) from 2012 to 2016. Chaudhary et al. (2018) present a literature review on inventory models for perishable items from 1990 to 2016, focusing on assumptions and specifications of various studies, noticing the growing importance of multi-echelon structure with centralized information. A more specific review is proposed by Karaesmen et al. (2011). They focus on models in which the quantity of perishables in inventory must be controlled throughout a horizon, taking into account the demand and shelf life of items. In particular, our work falls into the class of problems with stochastic demand with known and *fixed* shelf life, lead time greater than zero and periodic inventory review. Furthermore, given the analytical difficulty, we proceed with simulation-based optimization, a widely used methodology for stock replenishment problems (Jalali and Nieuwenhuyse, 2015; Deng et al., 2023) that allows tackling the complexities of real-world multi-echelon structures, which are often intractable even for small instances (De Kok et al., 2018). Such techniques are already availed in food contexts with multi-echelon models and heuristic solutions like Noordhoek et al. (2018), where the authors optimize (s,S) inventory policies, applying them to different echelons of a supermarket supply chain.

Related to the amount of goods that channels separately or jointly have to order, studies that assume the aforementioned demand and shelf life characteristics are Haijema and Minner (2019, 2016). However, they involve single-echelon problems with a single retailer. The authors suggest several parametric heuristics, studying both the possible benefits of more detailed inventory information in decision-making and several variations on *Base-Stock-Policies* (BSP). Under similar assumptions, Minner and Transchel (2010) focus on BSP and *Constant-Order-Policies* (COP) replenishment strategies with analysis related to multiple service level constraints. Such parametric heuristics are also common in multi-echelon approaches for perishable items (Van der Vorst et al., 2000), along with periodic review assumptions that are widely practiced in agri-food inventory systems (Kanchanasuntorn and Techanitisawad, 2006; Broekmeulen and Van Donselaar, 2019), in both macro- and micro-period planning (Janssen et al., 2018). We adapt such heuristics to the multi-echelon online/offline setting in the case of large volumes of demand and complex distributions that make exact or full lookup tables derived from dynamic programming approaches intractable. We compare the performance of connected multi-echelon models with separated single-echelon networks where each channel is treated independently, following the original versions of such heuristic policies. Regarding simpler instances where approximations are not necessary, similar to Hendrix et al. (2019) and Haijema (2013), we optimize the policies by value iteration (Powell, 2022; Brandimarte, 2021).

After determining how to order products from the supplier, the focus of multi-echelon models shifts to the allocation of products made by the distribution center (DC). Focusing on perishable products, much work has been done in the area of blood products by Prastacos (1978), Federgruen et al. (1986). We assume a *retention system*, in which each retailer *maintains* the inventory assigned, without returning the goods at the end of each period (i.e., rotation systems). However, we assume that the distribution center is an OFC and has an inventory that, while interacting with offline (physical) stores, actively serves the online channel of demand, thus not merely shipping and storing goods, but participating in the sales process. Inventory allocation in multi-echelon structures with an active role of the distribution center is studied in systems with direct upstream demand. Axsäter et al. (2007) use a two-echelon system in which customer demand occurs at both retailers and the DC, developing heuristics that decide how much inventory to hold in an active distribution center. Berling et al. (2023) study multi-echelon networks with direct upstream demand from an omnichannel and multichannel perspective, using order-up-to policies. However, these studies do not consider perishability. Moreover, they assume independent demand and continuous review of inventory.

For products with a fixed shelf life, Akkaş and Honhon (2022), similar to Prastacos (1978), analyze allocation policies with a focus on the remaining age of products shipped and they maximize the average long-run profit. In particular, their work deals with *consumer packaged goods* (CPGs) with a specific expiration date, following a specific type of vendor-managed inventory named *Direct Store Delivery* (DSD) and operating on pallets. They consider different patterns of behavior of end customers (*First In First Out* (FIFO) or *Last In First out* (LIFO)) and different internal strategies regarding both quantity and age of products. However, the scope of their study diverges notably due to several factors. Firstly, their work cannot be characterized as omnichannel and multichannel application but only as a multi-echelon study. In comparison, our research straddles both fields and we tackle issues with an active role of the distribution center as an online retailer that significantly shifts the problem, the application domain, and complexity. Furthermore, the inactive role of the depot in their study precludes the possibility of an analysis of correlation and pooling effects across different channels. We incorporate replenishment policies alongside allocation policies and consider an OFC to unpack and break down pallets, thus operating with customer units (Wollenburg et al., 2018). This assumption affects the type of demand distribution and increases its cardinality with a considerably larger number of items.

When products are non-perishable, management of items flow in distribution centers with online and offline channels is studied by Alawneh and Zhang (2018). They consider different levels of storage, in a two-echelon serial inventory control system, assuming continuous review and backorders. The optimal allocation problem for integrated online/offline channels is addressed by Goedhart et al. (2022). They approach the problem with value iteration, or by heuristics when the difficulty of the problem increases due to the curses of dimensionality. However, products are not perishable and orders from the supplier occur with longer time windows than the frequency of internal shipments. Conversely, perishability assumptions suggest more (daily) frequent orders (Broekmeulen and Van Donselaar, 2019). Handling multichannel and omnichannel retail infrastructure challenges for items with a short shelf life is of practical interest to some of the largest e-commerce platforms of the world, such as Alibaba. Deng et al. (2023) present practical implementations and solutions based on simulation-based approaches, using heuristics of (s, S) type in single- and multi-echelon structures.

3. Assumptions and dynamics of the problem

We model a network with an *Online Fulfillment Center* (OFC) and physical offline retailers, where these channels differ in the following characteristics:

- **Offline**: Offline retailers *always* have their own inventory where, even in the case of interconnected inventory across channels, it is not possible to return goods backward from the *offline* physical store to the distribution center. Moreover, it is not possible to control the issuing policy when dealing with physical customers that make their own choices according to a First-/Last-In-First-Out fashion (LIFO/FIFO), thus employing a stochastic percentage of LIFO and FIFO customers
- **Online**: For online orders, the OFC manages the issuing policy and sells products according to an optimal FIFO logic, thus minimizing waste. Furthermore, if there is a multi-echelon structure, the OFC jointly fulfills online orders and manages the dispatch to offline stores, using the center itself as both distributor and online supplier. Conversely, when the two channels operate as single-echelons, it only serves online customers.

The way consumers select the age of products to purchase is a complicating factor in perishable goods inventory management. Although from the retailer's perspective a FIFO issuing policy, in which older items are purchased first, would be optimal, this is not possible when the customer makes the choice offline. LIFO/FIFO mixtures are widely applied in the literature (Haijema and Minner, 2019, 2016) and observed empirically (Akkaş and Honhon, 2022; Deng et al., 2023). However, depending on the type of perishable good, the proportion between these two extreme solutions (pure LIFO or pure FIFO) may vary. In general, for goods where the expiration date is well defined like packaged products, it is reasonable to assume that LIFO sales are predominantly observed in physical retailers. However, in other cases such as fruits and vegetables, the phenomena of random picking, awareness of more sustainable behaviors, and preferences for specific degrees of product ripeness may lead customers to not always lean towards the freshest product, even in the absence of a dedicated discount strategy.

3.1. Dynamics of the distribution center

Optimizing order policies and internal dispatching strategies requires a model that deals with sequential decisions under conditions of uncertainty. Following Powell (2022), we describe the dynamics of the problem by defining the state and decision (actions) variables and introduce the respective reward and transition functions. We first consider the multi-echelon network, in which the OFC is the sole decision maker for the size of the orders from the supplier and acts as a distribution center for both channels. Table 7 summarizes the notation.

State of the system and actions. Suppose that the scheduled orders arrive before sales, with their maximum *shelf life* (SL), after a deterministic positive *Lead Time* (LT), and we immediately decide how many units of the same single product to send to the offline retailers, also distinguishing with regards to the available shelf life in the inventory of the OFC. The dynamics of our model start after the end of sales of the period and after the expired products are removed. At that time, we observe the following status of the OFC:

$$\mathsf{DC}_t = \left[O_d^{\mathsf{LT}-1}, \dots, O_d^0 | I_d^{\mathsf{SL}-1}, \dots, I_d^1\right]. \tag{1}$$

We label:

- O_d^l: Ordered items that will arrive in l periods at the DC (d). Therefore, O_d^0 represents the order delivered to the DC (with full shelf life SL) when the store reopens.
- I_d^r: Current physical inventory of the DC (d) with residual shelf life r.

We do not allow for back orders but assume *lost sales* in case of stock-outs. This assumption well represents retail applications, where we typically cannot even observe the number of customers who found the shelf empty. In addition, we consider information about the offline retailers $k \in \{1, \dots, K\}$. We assume that it is possible to distribute products with different residual shelf life from the OFC inventory as soon as the freshest ones (O_d^0) are delivered, with their maximum shelf life. Thus, we have:

$$\mathsf{RT}_t^k = \left[O_k^{\mathsf{RLT}_k-1,\,\mathsf{SL}-\mathsf{RLT}_k}, \dots, O_k^{\mathsf{RLT}_k-1,\,1}, \dots, O_k^{0,1} | I_k^{\mathsf{SL}-\mathsf{RLT}_k-1}, \dots, I_k^1\right], \tag{2}$$

where we define:

- RLT_k: *Retailer Lead Time* is the *deterministic* delay between product delivery from the distribution center and arrival at the physical store.
- $O_k^{l,r}$: Dispatched items that will arrive in l periods at retailer k with residual shelf life r.
- I_k^r: Current physical inventory of retailer k with residual shelf life r.

The complete state of the system after sales and after expired products are removed, before placing new orders, is:

$$S_t = \left[\mathsf{DC}_t | \mathsf{RT}_t^1, \dots, \mathsf{RT}_t^K\right]. \tag{3}$$

If we consider daily orders with weekly seasonality, the day of the week should be taken into account in the state of the system.

The first decision $X_0^\pi(S_t)$ concerns the quantity of the product ordered by the DC (i.e., OFC in the multi-echelon model) with respect to a policy π and the current state of the system S_t. Then, the orders with 0 residual lead time will have arrived and the status of the distribution center will be updated as:

$$\mathsf{DC}_t = \left[X_0^\pi(S_t), O_d^{\mathsf{LT}-1}, \dots, O_d^1 | O_d^0 = I_d^{\mathsf{SL}}, I_d^{\mathsf{SL}-1}, \dots, I_d^1\right]. \tag{4}$$

The DC decides on the quantities to be sent to retailers, just after having received the supplies. The quantity of the product shipped per available age r and retailer k (i.e., $X_k^{r,\pi}$) cannot exceed the stock availability, thus the following constraint follows:

$$X_k^{r,\pi} \in \{0, \dots, I_d^r\} \quad \text{s.t.} \quad \sum_{k=1}^{K} X_k^{r,\pi} \le I_d^r \quad \forall r \in \{1, \dots, \mathsf{SL}\}. \tag{5}$$

Complexity of the state space. The state vector DC_t has LT components (from LT-1 to 0) that pertain to in-transit orders and SL-1 components that pertain to physical inventory, where discarded items have already been removed at observation time. Hence, DC_t accounts for $(\mathsf{LT}+\mathsf{SL}-1)$ components in the complete state of the system. RT_t^k has $(\mathsf{SL}-\mathsf{RLT}-1)$ components for each offline retailer k regarding physical inventory, where, differently from DC_t, RLT is subtracted to take into account the effect of transportation time on maximum available shelf life. The remaining components of RT_t^k regard in-transit dispatched items. When the freshest items are allocated, the maximum residual shelf life items can have when delivered is $\mathsf{SL}-\mathsf{RLT}_k$, due to the internal lead time. Accordingly, the minimum residual shelf life that shipped items must have to be saleable at their arrival depends on RLT_k. It follows that for each in-transit period (1 to RLT), there are $\mathsf{SL}-\mathsf{RLT}_k$ possible residual shelf lives at delivery time, overall arranged in $\mathsf{RLT}_k(\mathsf{SL}-\mathsf{RLT}_k)$ components. If we consider a maximum amount of products that can be ordered (**maxO**) and assume that all retailers have the same RLT, the number of possible states (before decisions) is:

$$\sim (\mathbf{maxO})^{(\mathsf{LT}+\mathsf{SL}-1)+K(\mathsf{RLT})(\mathsf{SL}-\mathsf{RLT})+K(\mathsf{SL}-\mathsf{RLT}-1)}. \tag{6}$$

If $\mathsf{RLT}=0$ and K=1, the exponent is $(\mathsf{SL}+\mathsf{LT}-1)+(\mathsf{SL}-1)$, where the two addends concern the dynamics of the DC and the retailer's inventory only, with no variables for the queue of internal logistics of the system. However, the dimension will become rapidly intractable with large volumes of demand.

Transition function. After ordering and dispatching, sales occur at retailers (offline) and at the distribution center (i.e., OFC). The inventory of the OFC is updated according to the assumed FIFO issuing policy, as follows:

$$I_d^r = \left[I_d^r - \sum_{k=1}^{K} X_k^{r,\pi} - [D_{d,t} - \sum_{r'=1}^{r-1} (I_d^{r'} - \sum_{k=1}^{K} X_k^{r',\pi})]^+ \right]^+ \quad \forall r \in \{1, \dots, \mathsf{SL}\}, \tag{7}$$

where $D_{d,t}$ is the demand of period t at the OFC (d). Specifically, the number of stored items with remaining shelf life r at the OFC is updated by subtracting all the items with that age shipped to retailers (i.e., $\sum_{k=1}^{K} X_k^{r,\pi}$), along with those required to fulfill the residual demand not yet met by products with shorter remaining shelf lives $[D_{d,t} - \sum_{r'=1}^{r-1}(I_d^{r'} - \sum_{k=1}^{K} X_k^{r',\pi})]^+$.

For offline retailers, FIFO and LIFO issuing policies are mixed stochastically by means of a parameter $0 \le \gamma_k \le 1$. The LIFO transition function of the inventory is

$$I_k^r = \left[I_k^r + O_k^{0,r} - [\gamma_k D_{k,t} - \sum_{r'=r+1}^{\mathsf{SL}-\mathsf{RLT}_k} (I_k^{r'} + O_k^{0,r'})]^+ \right]^+ \quad \forall r \in \{1, \dots, \mathsf{SL}-\mathsf{RLT}_k\}, \forall k \in K, \tag{8}$$

and will involve a $\gamma_k \cdot 100\%$ of the customers. The corresponding FIFO one is

$$I_k^r = \left[I_k^r + O_k^{0,r} - [(1-\gamma_k) D_{k,t} - \sum_{r'=1}^{r-1} (I_k^{r'} + O_k^{0,r'})]^+ \right]^+ \quad \forall r \in \{1, \dots, \mathsf{SL}-\mathsf{RLT}_k\}, \forall k \in K, \tag{9}$$

and involves (1 - γ_k) $\cdot$ 100% of the customers. Here $D_{k,t}$ identifies the demand of period t for each offline retailer k. This demand is divided into two parts through the value γ_k. The number of items with remaining shelf life r is updated by summing the products supplied by the distribution center and by satisfying the demand. We start fulfilling the demand from the oldest items (i.e., $r = 1$) in the FIFO case of Eq. (9), and from those with maximum possible remaining shelf life at a given retailer k (i.e., $r = \mathsf{SL}-\mathsf{RLT}_k$) when LIFO assumptions hold in Eq. (8). Finally, indexes are shifted, expired items (I_d^0, I_k^0) are scrapped and we observe the new state of the system S_{t+1}.

Reward function. We define a reward function by computing the value of the cash flow that occurred in the current period

$$C(X^\pi, S, D) = \min \left\{ \sum_{r=1}^{\mathsf{SL}} (I_d^r - \sum_{k=1}^{K} X_k^{r,\pi}), D_d \right\} p_{On} + \sum_{k=1}^{K} \min \left\{ \sum_{r=1}^{\mathsf{SL}-\mathsf{RLT}_k} (I_k^r + O_k^{0,r}), D_k \right\} p_k + \\ + g_{On} I_d^0 + \sum_{k=1}^{K} g_k I_k^0 - c X_0^\pi, \tag{10}$$

where p_k, g_k and p_{On}, g_{On} are the sales prices and the salvage values (or disposal costs if negative), c is the unit cost per ordered item and X^π contains all the actions made with respect to a decision policy π. The t index is hidden for the sake of notational clarity. Similar to Buisman et al. (2020), Haijema and Minner (2019, 2016), we assume that holding and fixed transportation costs are negligible because perishable products are usually frequently restocked along with many others, thus sharing holding and fixed transportation costs among all these products.

The aim of the optimization is to find a policy π, such that we maximize the expected cash flow.

3.2. Dynamics of the single-echelon model

When there is not a shared distribution center and the OFC acts as a single-echelon, the system status after sales and after expired products are removed is:

$$\mathsf{R}_t^k = \left[O_k^{\mathsf{LT}-1}, \dots, O_k^0 | I_k^{\mathsf{SL}-1}, \dots, I_k^1\right], \tag{11}$$

where $k \in \mathbb{K} = \{1, \dots, K\} \cup \{\mathrm{OFC}\}$, set of the offline retailers that now includes the OFC, treated as an independent retailer.

The dimensionality of the state space and the number of decision variables are reduced. Specifically, the number of possible states is $\sim (\mathbf{maxO})^{(\mathsf{LT}+\mathsf{SL}-1)}$ for each retailer $k \in \mathbb{K}$. The only decision variable will involve the order size $X_{0,k}(\mathsf{R}_t^k)$ for each retailer.

The transition function will no longer include internal shipments, thus reducing to:

$$I_k^r = \left[I_k^r - [D_{k,t} - \sum_{r'=r+1}^{\mathsf{SL}} I_k^{r'}]^+\right]^+, \tag{12}$$

for a LIFO issuing policy and

$$I_k^r = \left[I_k^r - [D_{k,t} - \sum_{r'=1}^{\mathsf{r}-1} I_k^{r'}]^+\right]^+, \tag{13}$$

for a FIFO one. The cash flow is now separable and becomes:

$$C(X^\pi, S, D) = \sum_{k \in \mathbb{K}} C(\mathsf{R}^k, D_k) = \sum_{k \in \mathbb{K}} \left(\min\left\{\sum_{r=1}^{\mathsf{SL}} I_d^r, D_k\right\} p_k - g_k I_k^0 - c X_{0,k}^\pi\right). \tag{14}$$

3.3. Decision policies

Similar to Haijema (2013), we iterate a value function evaluated on all possible states, using the transition function ω that updates the state based on the sales (which depend on overall demand D) according to (7), (8), (9),(12) and (13). The Bellman equation is:

$$V_n(S) = \max_X \left\{\mathbb{E}[C(X, S, D) + V_{n-1}(\omega(X, S, D))]\right\}.$$

from which we derive the optimal policy X^{opt} (Brandimarte, 2021).

Heuristics. In the case of a single-echelon network, we introduce the well-known *Base-Stock Policy* (BSP) and *Constant Policy* (COP) to decide the size of the orders. Both depend on a single parameter α_k to be optimized for each retailer $k \in \mathbb{K}$. The COP represents a constant flow of items through each channel, while the BSP allows orders to vary according to the on-hand inventory and the number of products in the queue. Namely, we define:

$$X_{0,k}^{\mathsf{BSP}}(\mathsf{R}_t^k | \alpha_k) = \left[\alpha_k - \sum_{r=1}^{\mathsf{SL}-1} I_k^r - \sum_{l=0}^{\mathsf{LT}-1} O_k^l\right]^+ \quad \forall k \in \mathbb{K}, \tag{15}$$

the base stock policy, and

$$X_{0,k}^{\mathsf{COP}}(\mathsf{R}_t^k | \alpha_k) = \alpha_k \quad \forall k \in \mathbb{K}, \tag{16}$$

the COP one.

For multi-echelon systems, where the OFC dispatches items to offline retailers, we consider two different adaptations of the BSP and COP logic. Starting with the decision of the OFC on the size of orders, we define a Full Pull (FPL) order policy

$$X_0^{\mathsf{FPL}}(S_t | \alpha_0) = \left[\alpha_0 - \sum_r \left(I_d^r + \sum_k I_k^r\right) - \sum_l \left(O_d^l + \sum_{k,r} O_k^{l,r}\right)\right]^+, \tag{17}$$

and a Semi Push (SP) one

$$X_0^{\mathsf{SP}}(S_t | \alpha_0) = \alpha_0. \tag{18}$$

The values where r and l vary depend on RLT_k for each retailer k and are not made explicit to maintain readable notation. Similar to BSP in the single-echelon case, orders placed by the distribution center take into account the current stock availability of both echelons in the FPL case. Conversely, orders are made of a continuous flow of products as in COP in the SP case. We consider a flexible BSP internal dispatch policy for both ordering policies:

$$X_k^{\mathsf{FPL/SP}}(S_t | \alpha_k) = \min\left\{\left[\alpha_k - \sum_{r=1}^{\mathsf{SL}-\mathsf{RLT}_k-1} I_k^r - \sum_{l=0}^{\mathsf{RLT}_k} \sum_{r=1}^{\mathsf{SL}-\mathsf{RLT}_k} O_k^{l,r}\right]^+, \sum_{r=1}^{\mathsf{SL}} I_d^r - \sum_{k' \neq k} X_{k'}\right\} \quad \forall k \in \{1, \dots, K\}, \tag{19}$$

where retailer k makes requests according to a BSP logic ($[\alpha_k - \sum_{r=1}^{\mathsf{SL}-\mathsf{RLT}_k-1} I_k^r - \sum_{l=0}^{\mathsf{RLT}_k} \sum_{r=1}^{\mathsf{SL}-\mathsf{RLT}_k} O_k^{l,r}]^+$) and those are fulfilled taking into account the current inventory available in the OFC. If the number of retailers K is larger than one, it is necessary to decide on

Table 1
Summary of the investigated heuristics. Order and allocation parts of the policies are summarized for each heuristic.

	Single-echelon		Multi-echelon					
	COP	BSP	SP	FPL	SC	FPC	SP2K	FPL2K
Order	Constant	Base-Stock	Constant	Base-Stock	Constant	Base-Stock	Constant	Base-Stock
Allocation	–	–	Base-Stock	Base-Stock	Base-Stock + min OFC inventory		Weighted Base-Stock	
N params	$K+1$	$K+1$	$K+1$	$K+1$	$K+2$	$K+2$	$4K+1$	$4K+1$

an execution priority or rationing policy (e.g., Lagodimos, 1992; Akkaş and Honhon, 2022). Moreover, regardless of K, the dispatch can follow a FIFO (f) or LIFO (l) allocation logic with respect to the available inventory of the OFC. For this reason, we define FPL_l, SP_l, FPL_f, and SP_f, where FPL_l and SP_l allocate products starting with the freshest and FPL_f and SP_f with the oldest.

Since it is not always possible to have enough flexibility to change the size of orders on a daily basis with the supplier, the OFC allows flexible allocations through the internal BSP dispatch policy even if the size of the orders is constant (SP). When comparing the different models and their respective heuristics, it may be useful to directly compare a policy based on constant values (COP) for the single-echelon with the SP in the multi-echelon setting. The aim of this comparison is to quantify the benefit of inventory when BSP is not possible on the supplier side and is implemented only after the arrival of orders.

We also investigate versions of such heuristics that maintain a minimum inventory in the OFC, avoiding complete emptying in case of internal requests that saturate the central inventory capacity. This adjustment is inspired by Axsäter et al. (2007). Analytically we define an allocation policy:

$$X_k^{\mathsf{FPC/SC}}(S_t|\alpha_k,\alpha_C) = \min\left\{ \left[\alpha_k - \sum_{r=1}^{\mathsf{SL}-\mathsf{RLT}_k-1} I_k^r - \sum_{l=0}^{\mathsf{RLT}_k}\sum_{r=1}^{\mathsf{SL}-\mathsf{RLT}_k} O_k^{l,r}\right]^+, \sum_{r=1}^{\mathsf{SL}} I_d^r - \sum_{k'\neq k} X_{k'} - \alpha_C \right\}, \quad \forall k\in\{1,\dots,K\}. \tag{20}$$

where $\alpha_C \geq 0$ is the minimum number of items to maintain in the OFC. Consistent with the policies introduced earlier, we call these heuristics FPC and SC for a BSP and constant replenishment strategy, respectively.

Lastly, since the problem considers perishable goods with different residual shelf life, we weigh inventories by age adapting the BSP-W2Sk policy from Haijema and Minner (2019) to our Full Pull and Semi Push multi-echelon strategies. In this regard, we define FPL2K and SP2K, where the allocation policy is

$$\begin{aligned} &X_k^{\mathsf{FPL2K/SP2K}}(S_t|\alpha_k, w_k^{\mathsf{new}}, w_k^{\mathsf{old}}, v_k) = \\ &\quad \min\left\{ \left[\alpha_k - w_k^{\mathsf{old}} \sum_{r=1}^{\mathsf{SL}-\mathsf{RLT}_k-v_k} I_k^r - w_k^{\mathsf{new}} \sum_{r=\mathsf{SL}-\mathsf{RLT}_k-v_k+1}^{\mathsf{SL}-\mathsf{RLT}_k-1} I_k^r - \sum_{l=0}^{\mathsf{RLT}_k}\sum_{r=1}^{\mathsf{SL}-\mathsf{RLT}_k} O_k^{l,r}\right]^+, \sum_{r=1}^{\mathsf{SL}} I_d^r - \sum_{k'\neq k} X_{k'} \right\} \\ &\quad \forall k\in\{1,\dots,K\}. \end{aligned} \tag{21}$$

$v_k \in \{1,\dots,\mathsf{SL}-\mathsf{RLT}_k\}$ divides the inventory into old and new products, assigning a weight of $w_k^{\mathsf{new}} \in [0,1]$ and $w_k^{\mathsf{old}} \in [0,1]$ to the items on-hand for each offline retailer k.

A summary of the heuristics is provided in Table 1.

4. Numerical experiments

The complexity due to the number of decision variables and the dimensionality of the problem (6) makes solutions of real-size instances through exact approaches impractical. Therefore, in the first part of this experimental section, we analyze a small example, which is useful for understanding the structure of the optimal policy. We use the insights in the second part, where heuristics are used.

4.1. Design of experiments and optimal policy

We consider a simple case where there is only one offline retailer $K = 1$, one DC/OFC and the following parameters:

- A lead time of two periods ($\mathsf{LT}=2$) and a maximum shelf life of three periods ($\mathsf{SL}=3$).
- A zero internal lead time ($\mathsf{RLT}=0$), assuming that the facilities reside in the same urban area.
- A discrete uniform distribution of demand, identically and independent for both channels $D_{k,t} \sim U[0,1,2,3]$, $\forall k \in \mathbb{K}$, $\forall t$, following a pure FIFO assumption online and a pure LIFO offline.
- Three different newsvendor ratios, equal for both channels. Ratios are set with reference to the classical newsvendor model. Specifically, setting the salvage value to 0, we consider as critical fractile $\mathsf{newsR} = \frac{p-c}{p}$. We fix $p_k = 5\ \forall k \in \mathbb{K}$, varying the cost as $c = 3.75$, 2.5, 1.25. This results in $\mathsf{newsR} = 0.25, 0.5, 0.75$. Such values are related to empirically observed gross margin for grocery retailers. Common values for fruit and vegetables are around 25%, increasing for dairy and bakery up to 50% circa (O' Riordan, 1993). However, we cover also a higher value to provide a full picture of the effects on optimal policies.

Table 2
Optimal policy results. The expected profit and waste per period and stock-out percentage in the two channels are presented.

newsR	Echelon	Profit	Waste	Stk-out(OFC)	Stk-out(Offline)
0.25	multi-	2.1	0.2	1.5%	44.0%
	single-	2.1	0.2	19.1%	32.3%
0.5	multi-	5.8	0.6	0.2%	22.1%
	single-	5.5	0.6	8.0%	11.8%
0.75	multi-	10.2	0.6	0.0%	7.8%
	single-	9.8	0.9	2.0%	5.2%

Table 3
Orders and dispatches of the optimal policy in the multi-echelon system. The times a given quantity was ordered are shown as a percentage of the total number of decisions. As for dispatches, their percentages are categorized according to the remaining shelf life (rsl) at the time of delivery. '–' means that all values are zero.

newsR	OFC activity	n. of items					
		0	1	2	3	4	5
0.25	ordered	0.7	8.1	30.8	56.4	4.0	0
	dispatched (rsl = 1)	100	0	0	0	–	–
	dispatched (rsl = 2)	97.4	1.7	0.9	0	–	–
	dispatched (rsl = 3)	44.5	27.6	27.6	0.3	–	–
0.5	ordered	0.2	6.1	12.1	49.7	31.7	0.2
	dispatched (rsl = 1)	99.5	0.5	0	0	–	–
	dispatched (rsl = 2)	94.6	3.6	1.1	0.7	–	–
	dispatched (rsl = 3)	32.4	28.3	22.3	17.0	–	–
0.75	ordered	0	2.8	9.9	35.5	40.5	11.3
	dispatched (rsl = 1)	99.0	0.9	0.1	0	–	–
	dispatched (rsl = 2)	90.6	5.3	2.3	1.8	–	–
	dispatched (rsl = 3)	22.9	26.0	28.3	22.8	–	–

Table 4
Orders of the optimal policy in the single-echelon system. The times a given quantity of product was ordered are shown as a percentage of the total number of decisions.

newsR	Retailer	n. of items				
		0	1	2	3	4
0.25	OFC	0	64.4	35.6	0	0
	Off	66.7	0	0	33.3	0
0.5	OFC	0	41.8	58.2	0	0
	Off	33.4	0	33.3	33.3	0
0.75	OFC	2.1	24.7	66.7	6.5	0
	Off	33.4	0	33.3	0	33.3

Discussion of results. Table 2 presents the numerical results through three different performance measures. Specifically, the (*Profit*) value concerns the expected profit per period. Then, we report the scrapped items per period (*Waste*) and the (empirical) probability of stock-out (*Stk-out*) in the two channels. These values are calculated by applying the optimal policy X^{opt}, obtained through value iteration, to an out-of-sample horizon of 7000 steps, sufficient to have accurate estimates of the metrics.

The multi-echelon network generates higher or at least equal profits, because, in the worst case, it replicates single-echelon policies without actively using the inventory of the OFC. Waste follows the same pattern. For stock-outs, the OFC's shared inventory reduces stock-outs for online customers regardless of the newsvendor ratio. It also increases offline stock-outs, with smaller gaps with increasing critical ratio newsR. The complexity of a LIFO issuing policy in offline retailers makes the multi-echelon model prioritize demand saturation in the online channel, where it can be served in a simpler FIFO logic. The system gains the possibility of internally shipping products with different residual shelf life and the advantage of risk-pooling, allowing orders for different channels to be aggregated.

Tables 3 and 4 provide a statistical analysis of the optimal policy. In the multi-echelon case, the first row of Table 3 indicates the percentage of times that a specific number of items were ordered by the policy under consideration, while the other rows indicate the times that a number of items for a specific remaining shelf life were shipped to offline retailers. For example, the row "dispatched (rsl=1)" for newsR = 0.25 means that items with a remaining shelf life of 1 were never shipped to the offline retailer (100% on the zero column). On the other hand, the next row "dispatched (rsl=2)" indicates that in 97.4% of the cases no items with a remaining shelf life of 2 were shipped, in 1.7% of the shipments one item with 2 rsl periods was dispatched, and in 0.9%, 3 items were dispatched. We note that the OFC in the multi-echelon case seldom assigns expiring products to the offline retailer. This is logical because of the different behavioral patterns of customers. In case it dispatches items with low remaining life, these would not be preferred to newer items under the offline LIFO assumption, producing waste. The internal shipping policy resembles a LIFO issuing policy between the OFC and the retailer. As for the risk-pooling effects generated by the aggregation of orders, the multi-echelon structure orders about 2–3 items per period and gradually moves to 3–4 as the profit margin increases. Table 4 presents the percentage that a specific quantity is ordered by the independent single-echelon structure. Decisions fluctuate more vigorously in handling orders, especially for the offline channel, jumping between either 3 or 0 items in the $\mathsf{newsR} = 0.25$ instance. For $\mathsf{newsR} = 0.75$, the offline channel takes advantage of high margins, and the order size may exceed 3 items. With this decision, it manages LIFO demand by favoring the purchase of items with non-maximum shelf life in later periods.

4.2. Design of experiments and results of the heuristic approaches

Since we are interested in demand distributions with larger volumes and more complex interactions, we now consider parametric heuristic policies. Given a parameter $\boldsymbol{\alpha}$, we optimize the following problem:

$$\max_{\boldsymbol{\alpha}} \mathbb{E}\left[C(X^{\pi}(S|\boldsymbol{\alpha}), S, D)\right],$$

by simulation-based optimization, using an off-the-shelf surrogate optimization software (Eriksson et al., 2019a). We evaluate the objective function for an initial subset of points by sampling according to a Latin Hypercube design. We construct a surrogate model of the objective function with the evaluated points by means of radial basis functions (Regis and Shoemaker, 2007) and then use the surrogate model to decide where to evaluate the function based on a weighted-distance merit balance between the distance from previously evaluated points and the prediction of the value made by the surrogate. Since small changes in the $\boldsymbol{\alpha}$ parameter produce small changes in the inventories and do not produce large jumps in the value of the objective function, through the use of a surrogate model we can exploit previously observed values, also dealing with multiple local optima. We set the maximum number of function evaluations to $\max[200, 100 \cdot \text{Nparams}]$ and the number of points in the initial explorative subset to $\max[20, 10 \cdot \text{Nparams}]$. Other response-surface-based optimization methodologies are possible as well and applied to practical inventory management problems for perishable items (Deng et al., 2023).

Since the stochastic problem is ergodic in nature, the estimate of the expected value can be calculated by simulation over a sufficiently long horizon. To ensure that the estimate of this value is accurate, we apply two precautions:

- Since there is an initial transient where the inventory is empty because of the lead time, we do not measure any statistics for an initial time window three times as long as the sum of lead time and shelf life. $3 \cdot (\mathsf{SL} + \mathsf{LT})$.
- To decide that the estimate is sufficiently accurate, we use a 35-period sliding window. If the difference between the maximum and minimum value of the estimated expected value of profit in that window is less than 0.02% of the current estimation, we stop the simulation. To ensure multiple inventory cycles, the length of the sliding window is at least 7 times longer than the maximum shelf life of the products considered in our experiments.

Design of experiments. We assume a negative binomial distribution for demands. This distribution and its generalizations have numerous applications in marketing (Ehrenberg, 1959; Driesener and Rungie, 2022) and have proven to be a suitable solution in many real-world cases of purchasing scenarios. Furthermore, a negative binomial can be interpreted both as a Poisson distribution with parameters distributed according to gamma or as a compound Poisson process with geometrically distributed purchases quantity (Agrawal and Smith, 2015), providing a flexible choice for modeling retail sales. We parameterize the demand distribution through the mean and standard deviation (μ and σ) and we manage their relationship through the *coefficient of variation* ($\mathrm{cv} = \frac{\sigma}{\mu}$). We consider two different coefficients equal to 0.6 and 0.9. The range of these values is modeled by empirical estimations made by Broekmeulen and Van Donselaar (2019) from 3 large fresh food retailers in Europe, ruling out phasing-in, phasing-out, or promoted items, where such values may vary. Regarding the shelf life of the products, we consider $\mathsf{SL} = 3$ and 5, values common for, e.g., fruits and vegetables or bread. The choice is motivated by the fact that for inventory managers in retail, the product can be handled as a nonperishable when the shelf life is longer (Hendrix et al., 2019). Then, to better focus on the value of the multi-echelon structure rather than on rationing policies, we still consider a single offline retailer and an OFC. We simulate the behavioral pattern of offline customers through a beta distribution to represent the share of LIFO customers. Namely, we use $\gamma_{\mathrm{Off}} \sim$ Beta(2,2) and $\gamma_{\mathrm{Off}} \sim$ Beta(9,1). The first assumption represents symmetric uncertainty on the LIFO/FIFO balance, while the other option entails a LIFO-dominated issuing policy, where 90% of the offline consumers choose the freshest items on average. Other than a beta distributed LIFO/FIFO balance, it is also possible to simulate the residual shelf life of each purchased item singularly, by means of discrete choice models (Gioia et al., 2023). However, it requires a substantial number of parameters and it proves to be more adequate when discount strategies based on the age of the products are allowed in the model. The lead time is set $\mathsf{LT}= 3$, while the retailer lead time is assumed $\mathsf{RLT}= 0$, thus assuming a warehouse close to retailers. Concerning the newsvendor ratio, we explore two different cases. Namely, we define *Low Margin* and *High Margin* experiments when $\mathsf{newsR} = 0.25$ or 0.75 respectively. These configurations are applied to positively/negatively correlated, and independent scenarios. We assume symmetrical volumes of demand ($\mu = 100$) in the two channels and also unbalanced scenarios in which demand is assigned 80% to one channel and the remaining 20% to another (i.e., $\mu_{\mathrm{OFC}} = 160$ and $\mu_{\mathrm{Off}} = 40$ and vice versa).

Correlation between demands is generated by bivariate Gaussian copulas (Nelsen, 2006) having negative binomial marginals with mean and variance consistent with the aforementioned configurations. Specifically, the copula is generated by a bivariate Gaussian distribution with a linear correlation coefficient ρ equal to −0.5 in the case of negative correlation and 0.5 in the case of positive correlation.

Since the optimal policy structure in the multi-echelon model analyzed in often allocates fresher products first, having a LIFO (i.e., _l) or FIFO (i.e., _f) version available for each heuristic, we do not present results with FIFO internal allocation. In our configurations, these heuristics perform worse than their FIFO counterparts and do not add managerial insights.

Discussion of results. Table 5 shows the average profit and waste divided by subsets of parameters and normalized with respect to the single-echelon constant policy (COP). Specifically, given $\Pi^s_{\mathtt{COP}}$, the profit or waste per period of COP on configuration s, for each other heuristic π, we present

$$\left(\frac{\sum_{s=1}^{S} \Pi^s_{\pi}}{\sum_{s=1}^{S} \Pi^s_{\mathtt{COP}}} - 1 \right) \times 100.$$

The average of COP is not normalized and is given to clarify the magnitude of the results. Moreover, since half of the results come from high margin scenarios and half from low margin ones, when focusing on profits, such visualization may overrepresent high margin scenarios, where profits are inherently higher, contributing more to the averaged cumulative values. For this reason, we also present the percentage of the relative improvement per scenario in Table 6. Namely,

$$\left(\frac{1}{S} \sum_{s=1}^{S} \frac{\Pi^s_{\pi}}{\Pi^s_{\mathtt{COP}}} - 1 \right) \times 100.$$

Values are estimated out-of-sample over a 7000-period-long horizon, ensuring a maximum variation lower than the 0.02% in the empirical confidence of interval over the 35-period sliding window. Single-echelon models only observe marginal distributions, not varying with respect to the copula-induced correlation, thus correspondent values are presented without distinguishing by correlation coefficient. Simulations related to a multi-echelon network with pooled inventories are identified by their policy names (e.g., SP_l and FPL_l). Similarly, simulations related to a single-echelon structure are identified by the policy names (i.e., COP and BSP).

Focusing on unbalanced demand volumes between the two channels, when demand is dominantly online, where customers follow a FIFO behavior, profits increase and waste reduces for all configurations. Differences between single- and multi-echelon models reduce compared to the symmetric case. The more demand moves toward the offline channel, the more the higher share of LIFO consumers affects the complexity of dynamics, being older items more likely to be scrapped. COP is more robust and this effect confirms the results already pointed out by Minner and Transchel (2010), where constant approaches in the case of particularly complex configurations are valuable options. Under symmetric hypotheses, all multi-echelon policies perform better than single-echelon ones, both in profits and waste. However, having a minimum inventory constraint to be maintained in the OFC helps heuristics (FPC and SC). When demand is mainly offline, multi-echelon heuristics without any minimum stock level (i.e., FPL, FPL2K, SP, and SP2K) worsen remarkably with respect to symmetric and online dominant cases. The reason for such low performance is not poor management of the offline channel, but rather the lack of products in the OFC, which assigns most items to the physical store to cope with large internal base-stock values and peaks of the offline demand, generating many stock-outs in online (OFC) orders if compared with the single-echelon heuristics. The presence of a minimum stock level at the OFC avoids excessive dispatches due to large demands from the offline retailer. However, when only a small part of the demand is online, multi-echelon structures fail to provide benefits, thus it may be appropriate to fulfill online orders directly from physical offline retailers (Wollenburg et al., 2018).

Negative correlation aids the multi-echelon structure, allowing the OFC to exploit these characteristics, improving profits and reducing waste. When orders are constant, a base-stock allocation policy managed by the OFC (SP, SP2K, and SC) improves a simple COP, where no tailored allocations are possible. On the contrary, positive correlation negatively affects multi-echelon policies. Since single-echelon structures do not observe correlation, in the case of zero correlation, values concerning single-echelon policies (COP, BSP) summarize all the configurations. The challenge of negative binomial demands, stochastic consumer behavior, and long lead times renders a more flexible single-echelon approach (BSP) unprofitable due to problem complexity. Notably, these results are closely tied to the shelf life of products. For a shelf life of five periods, BSP becomes better than COP, differing by 0.4% and significantly reducing waste by 23% when considering the percentage of relative improvement. Conversely, short shelf lives have a significant negative impact on statistics of BSP on all the subsets. Shelf life emerges as a crucial factor also for multi-echelon structures, with shorter shelf lives hindering allocation flexibility and performance of the OFC, while Full Pull policies (FPL/FPl2K/FPC) take advantage of higher shelf life. However, all the multi-echelon policies outperform single-echelon approaches when the shelf life is five periods.

When the offline LIFO share increases, waste increases as well. Moreover, policies that consider a minimum stock level at the OFC are more effective in cases of higher offline LIFO values, where demand distributions are more complex. Nevertheless, the primary determinant of waste is the newsvendor ratio. The flexibility of the multi-echelon structure helps minimize losses and enhance performance when margins are low, compared to single-echelon approaches. In scenarios with a high coefficient of variation, policies with constant order quantity yield better results (i.e., COP, SP2K, SC).

Policies with varying weights for different ages in a Full Pull strategy (FPL2K) improve performance compared to their simpler counterpart (FPL) in cases of negatively correlated demand across channels (cv = 0.6) and exhibit superior responsiveness to prevalent offline demand. On the other hand, when age-based weights are applied to the Semi Push approach (SP2K), the improvement with respect to (SP) is noticeable in all the subsets but the 80/20 unbalanced case, where the allocation policy has a lower impact due to the prevalent online demand. Lastly, even if age-weighted policies are a generalization of their non-weighted counterpart and are theoretically able to reproduce them by fixing all additional parameters to 1, their optimization is considerably more complex, and local optima are common.

Table 5
Average profit and waste (Profit|Waste) per period with respect to different subsets of parameters and policies. Values normalized w.r.t. COP (profit, higher = better|waste smaller = better). COP values presented raw.

	Subset	Single-echelon		Multi-echelon					
		COP	BSP	FPL_l	FPC_l	FPL2K_l	SP_l	SC_l	SP2K_l
On/Off	80/20	429 | 25.4	−1.7 | −1.5	0.1 | −10.2	0.2 | −10.0	0.0 | −10.8	0.7 | −7.1	0.9 | −6.9	0.7 | −14.6
	50/50	417 | 27.6	−2.2 | −7.2	0.2 | −10.2	0.9 | −16.2	0.1 | −2.0	1.4 | −9.4	1.9 | −17.1	1.5 | −4.7
	20/80	406 | 31.1	−3.8 | −0.5	−4.0 | −9.3	−1.8 | −8.2	−3.5 | −3.2	−2.1 | −5.5	−0.6 | 1.0	−0.9 | −1.7
ρ	−0.5	–	–	0 | −16.3	1.4 | −16.5	0.7 | −15.2	1.5 | −12.5	2.0 | −16.1	1.8 | −17.1
	0	417 | 28.1	−2.5 | −3.2	−1.1 | −7.4	−0.3 | −12.2	−1.2 | −4.1	0.0 | −9.1	0.8 | −5.2	0.6 | −6.3
	0.5	–	–	−2 | −6.5	−1.6 | −6.0	−2.6 | 3.4	−1.2 | −0.8	−0.3 | −1.4	−0.8 | 2.9
LIFO/FIFO	50/50	425 | 26.1	−2.2 | −7.5	−1.7 | −9.9	−0.8 | −10.4	−1.6 | −2.6	−0.6 | −4.2	0.1 | −3.8	0.0 | −7.1
	90/10	410 | 30.0	−3.1 | 0.9	−0.7 | −9.9	0.3 | −12.3	−0.7 | −7.3	0.5 | −10.0	1.4 | −10.5	0.9 | −6.3
cv	0.6	440 | 21.8	−1.3 | −14.9	−0.2 | −23.4	0.6 | −22.3	0.1 | −20.4	0.1 | −18.4	0.8 | −16.4	0.6 | −13.7
	0.9	395 | 34.4	−4.1 | 4.2	−2.2 | −5.4	−1.3 | −7.9	−2.5 | 0.6	0.1 | −8.8	0.8 | −5.7	0.4 | −4.8
SL	3	402 | 31.7	−5.2 | 10.2	−3.6 | 0.2	−2.3 | 2.0	−3.5 | 5.4	−0.4 | −1.8	0.4 | −3.6	0.1 | −2.5
	5	433 | 24.4	−0.2 | −20.3	1.1 | −23.0	1.7 | −28.8	1.0 | −18.8	0.3 | −14.4	1.0 | −12.3	0.7 | −12.0
newsR	0.75	662 | 47.3	−2.5 | −1.4	−1.6 | −7.9	−0.8 | −9.4	−1.5 | −2.2	−0.3 | −4.1	0.2 | −3.6	0.1 | −3.7
	0.25	173 | 8.9	−2.8 | −12.6	0.4 | −21.4	1.9 | −23.0	0.2 | −21.8	1.2 | −25.2	2.8 | −28.6	1.7 | −23.6

Table 6
Percentage of relative improvement of profit and waste (Profit|Waste) per period with respect to different subsets of parameters and policies. Values normalized w.r.t. COP (profit, higher = better|waste smaller = better). COP values presented raw.

	Subset	Single-echelon		Multi-echelon					
		COP	BSP	FPL_l	FPC_l	FPL2K_l	SP_l	SC_l	SP2K_l
On/Off	80/20	429 | 25.4	−1.7 | −11.7	0.5 | −16.0	0.6 | −14.8	0.2 | −15.9	1.6 | −10.9	1.8 | −12.6	1.5 | −17.5
	50/50	417 | 27.6	−2.8 | −10.1	0.8 | −13.6	1.8 | −20.2	0.7 | −11.0	2.4 | −16.0	3.1 | −19.3	2.4 | −10.0
	20/80	406 | 31.1	−4.6 | −4.7	−3.8 | −13.4	−1.3 | −13.4	−3.6 | −9.7	−2.5 | −10.9	0.2 | −9.4	−1.0 | −7.7
ρ	−0.5	–	–	1 | −19.7	2.7 | −21.8	1.7 | −20.0	2.8 | −20.5	3.3 | −22.0	2.7 | −24.6
	0	417 | 28.1	−3.0 | −8.9	−0.7 | −11.2	0.3 | −16.0	−1.1 | −14.0	0.1 | −13.9	1.6 | −10.1	1.1 | −9.3
	0.5	–	–	−3 | −12.2	−1.9 | −10.5	−3.2 | −2.7	−1.4 | −3.4	0.1 | −9.2	−0.8 | −1.3
LIFO/FIFO	50/50	425 | 26.1	−2.7 | −10.7	−1.4 | −16.5	−0.2 | −16.2	−1.3 | −11.5	−0.1 | −9.4	0.7 | −9.6	0.4 | −9.6
	90/10	410 | 30.0	−3.4 | −7.0	−0.2 | −12.1	1.0 | −16.0	−0.4 | −13.0	1.1 | −15.8	2.6 | −18.0	1.6 | −13.8
cv	0.6	440 | 21.8	−1.1 | −18.1	0.4 | −18.2	1.5 | −18.4	0.6 | −16.6	0.2 | −9.1	1.3 | −12.9	0.8 | −10.7
	0.9	395 | 34.4	−4.9 | 0.4	−2.0 | −10.4	−0.7 | −13.8	−2.4 | −7.8	0.8 | −16.1	2.1 | −14.7	1.2 | −12.7
SL	3	402 | 31.7	−6.4 | 5.3	−3.9 | −1.4	−2.3 | −0.6	−3.8 | 1.8	0.3 | −6.1	1.5 | −9.7	0.9 | −6.3
	5	433 | 24.4	0.4 | −23.0	2.3 | −27.2	3.0 | −31.6	2.1 | −26.3	0.7 | −19.1	1.8 | −17.8	1.1 | −17.1
newsR	0.75	662 | 47.3	−2.7 | −4.5	−1.7 | −8.7	−0.9 | −10.4	−1.6 | −2.7	−0.4 | −2.3	0.1 | −2.8	0.1 | −3.2
	0.25	173 | 8.9	−3.4 | −13.2	0.1 | −20.0	1.7 | −21.8	−0.1 | −21.7	1.4 | −22.9	3.2 | −24.7	1.9 | −20.3

5. Conclusions and future works

We present models of multichannel structures for retailers that handle perishable goods both online and offline. We focus on the analysis of simple networks with a single offline retailer and an online fulfillment center, thus isolating the effects of the two different channels and not addressing the optimal management of rationing policies in cases where the number of retailers to be served is larger. When assumptions about demand volume, distribution type, and product characteristics allow for exact optimization, we optimize allocation and replenishment policies through value iteration. We address our first research question on the benefits of inventory pooling and we show that the OFC often allocates fresh products to offline retailers and handles online orders with older items. The multi-echelon model reduces online stock-outs because it gives priority to channels where the customer behavior is optimal (FIFO), increasing the offline ones, where customers make LIFO selections. The different customer behavior of the channels is exploited by the multi-echelon system together with the inventory-pooling effect, increasing the profit and reducing waste.

When model complexity and assumptions do not allow the problem to be solved optimally, we present generalizations of base-stock policies that are widely used in inventory management problems. Since the problem assumptions are particularly complex, a constant replenishment policy is advantageous for both multi-echelon and single-echelon models. However, although this result is already highlighted in Minner and Transchel (2010), the additional flexibility given by BSP-oriented internal allocations (SP) after constant orders (COP) often generates an advantage, making a multi-echelon allocation strategy beneficial when setting up contracts with a constant quantity from producers. As shelf life increases, the multi-echelon model and the proposed heuristics generate the highest profits, reducing waste. In fact, a longer shelf life allows a more flexible allocation policy to take advantage of products with a shorter remaining life, but enough to be shipped and sold in the offline channel. When offline and online channels have correlated demands, the proposed multi-echelon heuristics benefit from negative correlations, producing an advantage over the corresponding heuristics in single-echelon models. However, positively correlated demand is not exploited particularly well by simple extensions

Table 7
Notation.

Symbols	
SL	Full shelf life of the product
LT	Lead time from suppliers to OFC
RLT_k	Lead time from OFC to physical store k
DC_t	State variable of the OFC at time t in the multi-echelon system
S_t	Complete state vector of the multi-echelon system at time t
RT_t^k	State vector of retailer k at time t in the multi-echelon system
R_t^k	State vector of retailer k at time t in the single-echelon system
O_d^l	Ordered items that will arrive in l periods at the OFC (d) with full shelf life
$O_k^{l,r}$	Dispatched items that will arrive in l periods at retailer k with residual shelf life r
I_d^r	Current physical inventory of the OFC (d) with residual shelf life r
I_k^r	Current physical inventory of retailer k with residual shelf life r
X_0^π	Quantity ordered by the OFC according to decision policy π
$X_k^{r,\pi}$	Quantity sent to retailer k per residual shelf life r according to decision policy π
$D_{d,t}$	Demand of period t at the OFC (d)
$D_{k,t}$	Demand of period t for retailer k
γ_k	LIFO/FIFO stochastic mixing parameter for retailer k
p_k, g_k	Sales prices and salvage values (disposal costs) for retailer k per unit
c	Unit cost per ordered item
μ_k, σ_k	Mean and standard deviation of the demand distribution for retailer k
ρ	Correlation coefficient between online and offline channel demand

of base stock heuristics in the multi-echelon model because of the difficulties arising from filling base-stock-based internal requests after simultaneous peaks of demands. We show how the subdivision of demand among the channels affects model performance. If the online channel has higher volumes, the problem is simplified due to the larger share of FIFO consumers and the proposed heuristics perform adequately, but in the opposite case, with predominantly offline demands, base-stock allocation policies in the multi-echelon model generate frequent stock-outs in the OFC, allocating too many items to the offline channel, even if order quantities from the supplier are similar to the single-echelon case. One possible solution is a minimum level of stock in the distribution center, that proves to be profitable for almost all configurations. However, the implication is that different demand balances between online and offline channels require specific attention in creating heuristics for multi-echelon models. The extension of base-stock or constant-order policies in multi-echelon models with perishable goods is not trivial and requires understanding interactions and balance between online and offline channels. Policies with age-weighted inventories are investigated and they provide improvements on their non-weighted counterparts but increase optimization complexity.

Rationing policies are necessary when more offline retailers are included, increasing the complexity. Future studies should address such complications and the modeling of substitutions in the case of multiple products in the assortment. Another limitation of the present study and possible direction for future research is the non-seasonal demand. Inventory pooling under asynchronous seasonality across channels can improve the performance of multichannel networks, however, the dimensionality of the problem that will grow needs to be treated carefully. Furthermore, we assume that the frequency of orders received by the OFC and that of internal shipments within the model coincide, but the latter can also be more frequent per replenishment cycle.

CRediT authorship contribution statement

Daniele Giovanni Gioia: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing – original draft, Visualization. **Stefan Minner:** Conceptualization, Methodology, Investigation, Writing – review & editing, Supervision, Project administration.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The code of the simulation framework is provided on https://github.com/DanieleGioia/PerishableMEC.

References

Agrawal, N., Smith, S.A., 2015. Retail Supply Chain Management. Springer.
Akkaş, A., Honhon, D., 2022. Shipment policies for products with fixed shelf lives: Impact on profits and waste. Manuf. Serv. Oper. Manag. 24, 1611–1629.
Alawneh, F., Zhang, G., 2018. Dual-channel warehouse and inventory management with stochastic demand. Transp. Res. E 112, 84–106.
Axsäter, S., Olsson, F., Tydesjö, P., 2007. Heuristics for handling direct upstream demand in two-echelon distribution inventory systems. Int. J. Prod. Econ. 108, 266–270.
Bakker, M., Riezebos, J., Teunter, R.H., 2012. Review of inventory systems with deterioration since 2001. European J. Oper. Res. 221, 275–284.
Bell, D.R., Gallino, S., Moreno, A., 2014. How to win in an omnichannel world. MIT Sloan Manage. Rev. 56, 45–53.
Berling, P., Johansson, L., Marklund, J., 2023. Controlling inventories in omni/multi-channel distribution systems with variable customer order-sizes. Omega 114, 102745.
Brandimarte, P., 2021. From Shortest Paths to Reinforcement Learning: A MATLAB-Based Tutorial on Dynamic Programming. Springer.
Broekmeulen, R.A., Van Donselaar, K.H., 2019. Quantifying the potential to improve on food waste, freshness and sales for perishables in supermarkets. Int. J. Prod. Econ. 209, 265–273.
Bu, J., Gong, X., Chao, X., 2023. Asymptotic optimality of base-stock policies for perishable inventory systems. Manage. Sci. 69, 846–864.
Buisman, M.E., Haijema, R., Hendrix, E.M., 2020. Retailer replenishment policies with one-way consumer-based substitution to increase profit and reduce food waste. Logist. Res. 13 (7), 1–15.
Chaudhary, V., Kulshrestha, R., Routroy, S., 2018. State-of-the-art literature review on inventory models for perishable products. J. Adv. Manag. Res. 15, 306–346.
De Kok, T., Grob, C., Laumanns, M., Minner, S., Rambau, J., Schade, K., 2018. A typology and literature review on stochastic multi-echelon inventory models. European J. Oper. Res. 269, 955–983.
Deng, Y., Zhang, X., Wang, T., Wang, L., Zhang, Y., Wang, X., Zhao, S., Qi, Y., Yang, G., Peng, X., 2023. Alibaba realizes millions in cost savings through integrated demand forecasting, inventory management, price optimization, and product recommendations. INFORMS J. Appl. Anal. 53, 32–46.
Driesener, C., Rungie, C., 2022. The dirichlet model in marketing. J. Consumer Behav. 21, 7–18.
Ehrenberg, A.S., 1959. The pattern of consumer purchases. J. R. Stat. Soc. Ser. C. Appl. Stat. 8, 26–41.
Eriksson, D., Bindel, D., Shoemaker, C.A., 2019a. pySOT and POAP: An event-driven asynchronous framework for surrogate optimization. http://dx.doi.org/10.48550/ARXIV.1908.00420, arXiv preprint.
Eriksson, E., Norrman, A., Kembro, J., 2019b. Contextual adaptation of omni-channel grocery retailers' online fulfilment centres. Int. J. Retail Distrib. Manag. 47, 1232–1250.
Federgruen, A., Prastacos, G., Zipkin, P.H., 1986. An allocation and distribution model for perishable products. Oper. Res. 34, 75–82.
Gioia, D.G., Felizardo, L.K., Brandimarte, P., 2023. Simulation-based inventory management of perishable products via linear discrete choice models. Comput. Oper. Res. 157, 106270.
Goedhart, J., Haijema, R., Akkerman, R., 2022. Inventory rationing and replenishment for an omni-channel retailer. Comput. Oper. Res. 140, 105647.
Goyal, S.K., Giri, B.C., 2001. Recent trends in modeling of deteriorating inventory. European J. Oper. Res. 134, 1–16.
Haijema, R., 2013. A new class of stock-level dependent ordering policies for perishables with a short maximum shelf life. Int. J. Prod. Econ. 143, 434–439.
Haijema, R., Minner, S., 2016. Stock-level dependent ordering of perishables: A comparison of hybrid base-stock and constant order policies. Int. J. Prod. Econ. 181, 215–225.
Haijema, R., Minner, S., 2019. Improved ordering of perishables: The value of stock-age information. Int. J. Prod. Econ. 209, 316–324.
Hendrix, E., Ortega Lopez, G., Haijema, R., Buisman, M., García Fernandez, I., 2019. On computing optimal policies in perishable inventory control using value iteration. Comput. Math. Methods 1 (4), 1–12.
Hübner, A., Wollenburg, J., Holzapfel, A., 2016. Retail logistics in the transition from multi-channel to omni-channel. Int. J. Phys. Distrib. Logist. Manage. 46, 562–583.
Jalali, H., Nieuwenhuyse, I.V., 2015. Simulation optimization in inventory replenishment: a classification. IIE Trans. 47, 1217–1235.
Janssen, L., Claus, T., Sauer, J., 2016. Literature review of deteriorating inventory models by key topics from 2012 to 2015. Int. J. Prod. Econ. 182, 86–112.
Janssen, L., Diabat, A., Sauer, J., Herrmann, F., 2018. A stochastic micro-periodic age-based inventory replenishment policy for perishable goods. Transp. Res. E 118, 445–465.
Kanchanasuntorn, K., Techanitisawad, A., 2006. An approximate periodic model for fixed-life perishable products in a two-echelon inventory–distribution system. Int. J. Prod. Econ. 100, 101–115.
Karaesmen, I.Z., Scheller-Wolf, A., Deniz, B., 2011. Managing perishable and aging inventories: Review and future research directions. In: Kempf, K.G., Keskinocak, P., Uzsoy, R. (Eds.), Planning Production and Inventories in the Extended Enterprise: A State of the Art Handbook, Vol. 1. Springer, New York, NY, pp. 393–436.
Lagodimos, A., 1992. Multi-echelon service models for inventory systems under different rationing policies. Int. J. Prod. Res. 30, 939–956.
Minner, S., Transchel, S., 2010. Periodic review inventory-control for perishable products under service-level constraints. OR Spectrum 32, 979–996.
Nahmias, S., 1982. Perishable inventory theory: A review. Oper. Res. 30, 680–708.
Nelsen, R.B., 2006. An Introduction to Copulas. Springer New York, NY.
Noordhoek, M., Dullaert, W., Lai, D.S., de Leeuw, S., 2018. A simulation–optimization approach for a service-constrained multi-echelon distribution network. Transp. Res. E 114, 292–311.
O' Riordan, D., 1993. Retail gross margins: Some international comparisons. Int. J. Retail Distrib. Manag. 21, 33–39.
Powell, W.B., 2022. Reinforcement Learning and Stochastic Optimization: A Unified Framework for Sequential Decisions. Wiley.
Prastacos, G.P., 1978. Optimal myopic allocation of a product with fixed lifetime. J. Oper. Res. Soc. 29, 905–913.
Regis, R.G., Shoemaker, C.A., 2007. A stochastic radial basis function method for the global optimization of expensive functions. INFORMS J. Comput. 19, 497–509.
Rezaei, M., Liu, B., 2017. Food Loss and Waste in the Food Supply Chain, Vol. 2017. International Nut and Dried Fruit Council, Reus, Spain, pp. 26–27.
Song, Y., Fan, T., Tang, Y., Xu, C., 2021. Omni-channel strategies for fresh produce with extra losses in-store. Transp. Res. E 148, 102243.
UNCTAD, 2022. Covid-19 boost to e-commerce sustained into 2021. Available online at: https://unctad.org/news/covid-19-boost-e-commerce-sustained-2021-new-unctad-figures-show. (Last accessed on 31 December 2022).
Van der Heijden, M., Diks, E., De Kok, A., 1997. Stock allocation in general multi-echelon distribution systems with (R, S) order-up-to-policies. Int. J. Prod. Econ. 49, 157–174.
Van der Vorst, J.G., Beulens, A.J., Van Beek, P., 2000. Modelling and simulating multi-echelon food systems. European J. Oper. Res. 122, 354–366.
Wollenburg, J., Hübner, A., Kuhn, H., Trautrims, A., 2018. From bricks-and-mortar to bricks-and-clicks: Logistics networks in omni-channel grocery retailing. Int. J. Phys. Distrib. Logist. Manage. 48, 415–438.