Computers & Operations Research 140 (2022) 105647

Contents lists available at ScienceDirect

Computers and Operations Research

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

Inventory rationing and replenishment for an omni-channel retailer

Joost Goedhart *, René Haijema, Renzo Akkerman

Operations Research and Logistics, Wageningen University, Hollandseweg 1, 6706 KN Wageningen, The Netherlands

ARTICLE INFO

Keywords:
Dynamic rationing
Bricks-and-clicks
Markov decision process
Lost sales
Lead time
Retail operations

ABSTRACT

The growth of omni-channel retailing resulted in many new challenges for retailers, especially in relation to the replenishment and allocation of inventories for the different channels. In this paper we consider a brick-and-mortar store that uses its inventory to fulfil both in-store demand as well as online orders. In addition to deciding on replenishment quantities, such a retailer also has to decide how to ration its inventory across the channels. Practically, rationing inventory relates to storing part of the inventory in the backroom to satisfy online demand. The rationing process occurs regularly (e.g. daily) whereas inventory replenishment typically occurs less regular. To analyse this decision process, we model the rationing and ordering decisions as a Markov Decision Problem that maximises the expected profit. Based on the structure of the optimal policies, we determine heuristics that near-optimal results and scale well to retailers with many products.

1. Introduction

The increased competition of online shopping has put pressure on brick-and-mortar stores. Online shopping has made it effortless for consumers to satisfy their demands from home. To compete against online retailers, physical store retailers are increasingly adopting online shopping channels in their channel portfolio. This integration of different sales channels is referred to as omni-channel retailing (Verhoef et al., 2015). The goal of omni-channel retailing is to provide customers with a seamless shopping experience and enhance customer loyalty and satisfaction.

When an omni-channel retailer adopts new channels, it needs to reconsider its inventory policies (Jalilipour Alishah et al., 2015). Retailers for instance have to decide whether or not to integrate inventories of different channels. Retailers that add an online channel often choose to use store inventory to satisfy online demand (ENC, 2016), as this has the lowest initial investment (Fernie and Sparks, 2004). The integration of online sales with offline sales is an ongoing discussion in the field of retail operations and is referred to as bricks-and-clicks (Agatz et al., 2008). Especially smaller retailers are adopting this concept, in which their store essentially become a distribution centre for their online order fulfilment (Mou et al., 2018). In this way, retailers are able to leverage their brick-and-mortar store with online sales.

By using store inventory for the fulfilment of online orders, several advantages can be achieved. First, since offline inventory is also used to serve online customers, uncertainty in demand can be reduced, lowering the total inventory. Second, higher inventory turnover rates result in a decrease in left-overs at the end-of-sales periods (Bendoly, 2004; Bayram and Cesaret, 2021). Furthermore, stores are often located closer to customers than distribution centres, thus environmental and economic benefits for the delivery of online orders can be achieved (Jalilipour Alishah et al., 2015).

However, using stores to fulfil the demand of online customers also has disadvantages. Using store personnel for picking orders in-store might influence the customer shopping experience and can be a high expense for the retailer (Baird and Kilcourse, 2011; Ishfaq and Bajwa, 2019). Furthermore, store inventory needs to be monitored more closely to ensure online orders can be satisfied. If online demand occurs, this means picking the order from the store shelves and updating the inventory level on the website. If a customer buys the product in-store, this also needs to be coordinated with the online channel to ensure that there is no conflict in which an online customer tries to buy the same product digitally. To mitigate these problems retailers often choose to satisfy online demand from the remaining store inventory at the end of the day when customers have left. However, stock-outs might occur resulting in decreased customer loyalty and satisfaction as online demand cannot be satisfied (Nguyen et al., 2018).

To mitigate the negative effects of using a store as a fulfilment centre for online orders, managerial studies (e.g. Hobkirk, 2015; ENC, 2016) have suggested to reserve a certain amount of in-store inventory to satisfy online demand. In practice, this relates to storing the inventory for the online demand in the backroom (Aastrup and Kotzab, 2010). By doing so, the retailer does not need to continuously coordinate the two sales channels, preventing sales of products that turn out to actually be unavailable. Furthermore, there is no need to pick online orders at the moment they arrive. The retailer can inform their

* Corresponding author.
E-mail address: joost.goedhart@wur.nl (J. Goedhart).

https://doi.org/10.1016/j.cor.2021.105647
Received 24 October 2020; Received in revised form 13 September 2021; Accepted 19 November 2021
Available online 11 December 2021
0305-0548/

online customers immediately whether a product is available in their respective channel. Also for the in-store channel, the retailer needs to know how many items to display at the store, as a product that is no longer on display may generate no demand. Managing the store inventory happens at a daily basis, thus unbalanced inventory among channels can be reduced by regular rebalancing of the inventory. When retailers ration their inventory across their physical and digital channels in this way, they can balance the satisfaction of demand from both channels. This avoids the disadvantages outlined above, while creating a new one: when a stock-out happens in one of the channels, a lost sale occurs even though the other channel might have inventory left.

If a retailer rations its inventory, a trade-off therefore exist between serving in-store customers and serving online customers. Although serving an online customer from a store normally results in a lower net revenue due to shipping and handling costs (e.g. Bayram and Cesaret, 2021), the inventory cost are normally also lower since products do not need to be displayed on store shelves (Xu and Cao, 2019). Little research has been conducted on how retailers should adjust to this new store operation and its impact on the retailers' performance (Mou et al., 2018). Finding an optimal rationing policy for retailers using in-store inventory for fulfilling online and offline demand is thus an important challenge. Additionally, the retailers inventory replenishment decisions are linked to this rationing policy. Therefore, the objective of this study is to identify an optimal replenishment and rationing policy for an omni-channel retailer.

Current rationing policies are not sufficient for the abovementioned retailer settings, as they typically consider a negligible lead time and the possibility to backorder. As a positive lead time and lost sales are common in retail practice (e.g. Jalilipour Alishah et al., 2015; Bayram and Cesaret, 2021), we focus in this paper on a retailer who faces a deterministic non-zero lead time, and who has to deal with lost sales in case of stock-outs. Deterministic, non-zero lead times are necessary to account for handling time and transportation efforts. Furthermore, lost sales are relevant because, when a stock-out occurs, demand cannot be backordered, as customers will often satisfy their demand elsewhere. In-store demand may only happen when products are displayed. Online customers nowadays expect immediate fulfilment of their demand with next day delivery, and they do normally not wish to backorder their demand.

Including non-zero lead times and lost sales does however increase the complexity of the decision problem, and makes an analytical approach cumbersome. Exact solution can be obtained numerically by formulating and solving the model as a Markov Decision Problem (MDP). Solving the MDP is time consuming due to the curse of dimensionality, where the amount of states for which an optimal action needs to be determined can grow exponentially. Therefore, a faster solution would be preferred for practical implementation. In this paper, we therefore also derive a heuristic based on the structure of the optimal policy of the MDP.

We contribute to the academic literature in several ways. First, we provide an exact method for replenishment and rationing in an omni-channel bricks-and-clicks context. Second, based on our results, we are able to derive heuristics that omni-channel retailers could easily apply in practice. Third, with an extensive numerical study, we compare the heuristics with the optimal policy to find their performances.

The remainder of this paper is structured as follows. Section 2 presents related research on order fulfilment in an omni-channel setting and on inventory rationing. In Section 3, we outline the decision problem and formulate it as a MDP. In Section 4, the MDP model is numerically investigated to identify the structure of the optimal policy. In Section 5, we derive heuristics for the ordering and rationing decision based on the structure of the optimal policies found numerically. In Section 6, we compare the performance of the derived heuristics in relation to the optimal policy for a wide range of instances. Section 7 concludes the paper and discusses future research directions.

2. Literature review

Our work is related to the literature on ship-from-store strategies in omni-channel retailing and to the literature on inventory rationing. Below, we first briefly address the work on ship-from-store strategies, followed by a more comprehensive overview of the relevant inventory rationing literature.

2.1. Ship-from-store strategies

The rich body of literature on online order fulfilment discusses the different strategies retailers can adopt to fulfil online demand. In this research, we are interested in ship-from-store strategies, a concept that uses store inventory to satisfy online demand.

One of the earliest works related to online fulfilment from store inventory is the work by Bendoly (2004). This work concluded that using store inventory for online fulfilment decreases the inventory cost for the online channel, however, satisfying in-store demand decreased due to higher stock-outs. Therefore, a trade-off between lower inventory cost and satisfying demand occurs. Further research by Bendoly et al. (2007) discussed when using store inventory is beneficial for retailers, finding that with lower percentages of online sales, store fulfilment is preferable. Hübner et al. (2016) also mention that using stores for online fulfilment is preferable for retailers who are aiming to adopt an online channel in a fast and inexpensive manner.

Early work on using stores for online fulfilment is mainly focused on how to allocate an online sale to different online fulfilment locations. For instance, Mahar and Wright (2009) and Mahar et al. (2009) studied a case in which the online order could be fulfilled from either a store or an online fulfilment centre. Here, the decision on where to fulfil an online order from is chosen centrally, taking into account the location of the online fulfilment centre. Bretthauer et al. (2010) and Mahar et al. (2012) extend this research by additionally taking the decision of whether a store should be included in the allocation of online orders. By not incorporating all stores for online fulfilment, their inventory can be protected from stock-outs. Similarly, Aksen and Altinkemer (2008) study store fulfilment in settings with multiple stores, deciding on which store should fulfil an online order, based on the distance to the customer and the related cost.

More recent papers focus on the order fulfilment and the operational costs. Ishfaq and Bajwa (2019) gives insight on the profitability of online fulfilment from stores when other choices of fulfilment such as vendors, distribution centres or online fulfilment centres are available. Bayram and Cesaret (2021) researches a similar setting however, includes the demand generated in-store. Difrancesco et al. (2021) specifically study the ship-from-store strategy where online orders are picked from store shelves and then shipped to the consumer. Through simulation they determine the number of pickers and packers that ensures a good balance between service levels and costs.

Most literature on ship-from-store strategies is thus focused on the strategic decision on whether to adopt a ship-from-store strategy instead of adopting alternatives such as online fulfilment centres. On the operational side, the focus is often on which location should satisfy an online order. The impact of integration of channels on store operations is however not much addressed in the literature, even though an increasing number of stores are involved in online order fulfilment (Mou et al., 2018). Optimal implementation of the ship-from-store strategy is essential, as failure can lead to stock-outs, higher costs, and higher customer dissatisfaction (Difrancesco et al., 2021). Therefore, it is important for omni-channel retailers to have good inventory management and properly manage their inventory in relation to the different sales channels. In this paper, we therefore focus on inventory management related to our study which will be elaborated in the following section.

Table 1
Literature review of dynamic inventory rationing.

Paper	Replenishment		Demand			Optimisation focus			Method		
	Order policy	Lead time[a]	Backordering	Lost sales	Classes	Cost	Profit	Service level	MDP	Simulation	Mathematical analysis
Kaplan (1969)			×		2	×					×
Haynsworth and Price (1989)	(s,Q)[b]	D	×		2			×		×	
Ha (1997b)			×		2	×			×		
Carr and Duenyas (2000)				×	2		×		×		
Melchiors et al. (2000)	(s,Q)[b]	D		×	2	×					×
Melchiors (2001)	(s,Q)[b]	S		×	N	×			×		
Deshpande et al. (2003)	(s,Q)	D	×		2	×					×
Frank et al. (2003)	(s,S)	Z		×	2	×					×
Melchiors (2003)	(s,Q)[b]	D		×	N	×			×		
Teunter and Klein Haneveld (2008)	(s,Q)	Z	×		2	×					×
Gayon et al. (2009)		S	×		N	×					×
Benjaafar et al. (2010)			×	×	2	×			×		
Fadıloğlu and Bulut (2010)	(s,Q)	D	×	×	N	×				×	
Zhao and Lian (2011)	(s,Q)	S	×		2	×					×
Chen et al. (2012)	(s,Q)	S		×	2	×					×
Hung et al. (2012)	(R,S)	D	×		N	×					×
Chew et al. (2013)	Dynamic	Z & D	×		N	×			×		
Hung and Hsiao (2013)	(s,S)[b]	S		×	N			×		×	
Wang et al. (2013a)	(s,Q)	D	×		2	×		×			×
Wang et al. (2013b)		S	×		2	×			×		
Wang and Tang (2014)	(R,S)	Z	×	×	N	×			×		
Liu and Zhang (2015)			×		2	×			×		
Liu et al. (2015)	(R,S)	Z	×		N	×					×
Turgay et al. (2015)				×	N		×				×
Alfieri et al. (2017)	(R,S)[b]	D	×		2	×				×	
Bao et al. (2018)	Dynamic	Z	×		N	×					×
Xu and Cao (2019)	Dynamic	Z	×		2		×				×

[a]D = Deterministic, S = Stochastic, Z = Zero.
[b]Fixed parameters for replenishment policy.

2.2. Inventory rationing

Ensuring profitability for the retailer using in-store inventory for online fulfilment is an import topic in omni-channel retailing. To enable this, retailers need to develop a strategy for reserving part of their inventory specifically for online orders. Even though such a strategy is important, there is limited research on the topic. As mentioned by Jalilipour Alishah et al. (2015), Ma and Jemai (2019), and Xu and Cao (2019) the problem setting resembles inventory rationing, a topic that has received significant attention in the literature since the seminal work by Topkis (1968). Rationing is often used to connect different types of demand with different fulfilment options. These types of demand are mostly referred to as classes, which are similar to sales channels in our research. The initial rationing work by Topkis (1968) is a static strategy. Since then, research has moved from static to dynamic rationing strategies, which have been shown to be superior (Teunter and Klein Haneveld, 2008). In the remainder of this section, we therefore only focus on dynamic rationing strategies.

The dynamic rationing literature can be classified along several dimensions: replenishment policy, lead time, the number of demand classes, how shortage is dealt with, the objective of the study, and the method that is applied. The replenishment policy considers how much should be ordered and when an order should take place. Most papers consider simple policies while some consider the policy to be dynamic. A few papers do not consider replenishment in their study at all. The lead time in the studies can be either zero, a deterministic value, or a stochastic value. The shortage treatment of demand can be divided into backordering or lost sales. The number of demand classes considered is either two or a more generic N. The objective of the studies is either to reduce cost, increase profit, or improve service level. The method the authors applied can be differentiated in MDP, simulation, or mathematical analysis. Table 1 presents a chronologically ordered overview of the literature on dynamic rationing, including the classification on the six dimensions.

2.2.1. Replenishment policy and lead time

From Table 1, it is observed that the first papers investigating dynamic rationing only considered an (s,Q) replenishment policy, which is the optimal replenishment strategy for static rationing according to Ha (1997a). Not all papers consider a replenishment policy. Some papers only investigate a single period with a fixed initial inventory level, while other consider a manufacturing system in which some production planning decisions are considered instead (Carr and Duenyas, 2000; Turgay et al., 2015).

Wang and Tang (2014) concluded that, for periodic review, a base-stock replenishment policy is the optimal ordering policy in situations with zero lead time. Chew et al. (2013) were the first to study a dynamic replenishment policy and concluded that the optimal dynamic replenishment policy resembles a base-stock replenishment policy, which is confirmed in later studies (e.g. Bao et al., 2018; Xu and Cao, 2019). However, optimal replenishment policies in situations with non-zero lead times could not be derived due to the complexity of the problem.

The replenishment policy is relevant for the rationing decision, as the rationing decision is based on current and future inventory levels. Thus outstanding replenishment orders influence the rationing decision. It can be concluded that the optimal dynamic replenishment policy for situations with deterministic lead times has yet to be addressed in the literature.

2.2.2. Demand

Wang and Tang (2014) were the first to compare the backordering and lost-sales setting. From the comparison, it was found that in a backordering setting the rationing level increases occasionally to ensure inventory for future demand. With lost sales, the rationing level decreases when coming closer to replenishment. In general, considering lost sales increases the complexity of the system, as the inventory level does not change with unmet demand (Frank et al., 2003; Fadıloğlu and Bulut, 2010; Chen et al., 2012; Hung and Hsiao, 2013).

Melchiors (2001) was the first to propose a model with N demand classes, although the numerical results are only presented for 2 demand classes. Most papers propose a model for N demand classes, but are not able to derive exact solutions if the demand classes exceeds 2 due to the complexity (Melchiors, 2003; Chew et al., 2013). Bao et al. (2018) is able to derive an exact solution for N demand classes but mentions introducing non-zero lead time makes finding the optimal policy much more complicated.

2.2.3. Optimisation focus

The rationing decision is often based on minimising costs. This is also related to how shortages are considered: if demand is always satisfied (through backordering), the rationing decision only influences cost and not profit. Studies that consider lost sales and base their rationing on minimising cost, consider a shortage cost for the lost sales. By applying different shortage costs, they are able to achieve different service levels for the demand classes (Chew et al., 2013; Liu et al., 2015).

Carr and Duenyas (2000) and Turgay et al. (2015) were the first to base their rationing decisions on profit maximisation for a make-to-stock production system, concluding that the structure of the optimal policy is complex, and therefore exploring the solution with a simpler policy. Xu and Cao (2019) used profit to seek a balance between fulfilling online demand and the related high fulfilment cost. In general, using profit as optimisation focus increases the complexity of the structure of the optimal policy. However, it better captures the omni-channel fulfilment strategies as it allows the balancing between satisfying in-store and online demand.

2.2.4. Method

In most papers, dynamic programming is used to find the optimal solution and find the structure of the rationing policies. However, as mentioned previously, the structure of the system can become complex for certain system characteristics. As the problem easily become multi-dimensional, the curse of dimensionality results in complex and not insightful formulas (Teunter and Klein Haneveld, 2008). Therefore, research often resorts to finding near-optimal policies for the problem (Liu and Zhang, 2015).

By applying MDP to solve the dynamic rationing model, an exact solution can be derived. Most literature using MDP to find the optimal solution also uses the structure of the solution to derive simple heuristics (e.g. Wang et al., 2013a; Chew et al., 2013) as MDP can require high computational time.

2.3. Knowledge gap

The dynamic rationing problem resulting from the ship-from-store context is characterised by deterministic lead times and lost sales. Because online demand cannot be backlogged as the consumer requires same day shipment, shortages lead to lost sales. As the in-store and online channel have different cost and benefits, only the profit is able to encompass the trade-off between the different channels.

In the literature review we can see that the paper by Chew et al. (2013) is closely related as it applies a dynamic order policy with deterministic lead time. They do however not include lost sales, but mentions its importance for further research. Melchiors (2001, 2003) and Wang and Tang (2014) are also related as they investigate the lost sales problem using MDP. Wang and Tang (2014) specifically mention the investigating of different order policies and the inclusion of lead time as potential future research. Although Melchiors (2001, 2003) does add lead time in their research, they limit themselves by fixing the order policy. Finally, Xu and Cao (2019) study an omni-channel retailer in a similar setting, but acknowledge that their assumption of zero lead time limits their findings.

Based on these findings, we conclude that no lost sales model exits that integrates the ordering and rationing decision for a profit maximising retailer who serves both an online and an offline demand from a stock point with non-zero replenishment lead time. We contribute to the existing literature by deriving heuristics and compare them with the exact solution obtained from the MDP. In the next section we present the problem and a model to derive an exact solution.

3. Markov decision problem

3.1. Problem definition

The problem that is studied is determining an optimal rationing and replenishment policy. A typical setting for a retailer is that they can order new products every week but can ration their products on a daily basis. Managing the in-store inventory thus happens quite regularly, while ordering and replenishment are more dependent on fixed delivery schedules. The rationing decision at the start of every day is motivated by practices in which the online sales will be packed and shipped at the end of a day. Which is very common to happen after some cut-off time for online orders. The rationing decision is needed to ensure that enough products are left to fulfil the online orders placed during the day. These two decision problems can be formulated as a hierarchical decision problem. The problem has two levels, where at level I the replenishment decision is made and at level II the rationing decision is made. The time between two ordering decisions (at level I) is called a period. As the rationing decision at level II is taken more frequent, level II is split into $R = 7$ sub-periods.

The objective of the problem is to maximise the profit resulting from sales revenues on the one hand and holding costs, shipment costs, and procurement costs on the other hand. We differentiate two types of customers: those who visit the offline channel (the physical store) and those who use the online channel. It is assumed that there is no channel substitution and that the cost of the online and offline channels differ. Shipment costs only applies to online sales. Nowadays, customers expect free delivery thus the costs are carried by the retailer. The sales price in the two channels are the same in an omni-channel setting.

In Fig. 1 the problem is presented with its two levels, where at level I the retailer makes the replenishment decision is based on the actual inventory level. At level II, the rationing decision is set based on the actual inventory level and the outstanding order (which was set at level I). Every week (M) the replenishment decision is made at the start of the first day, which will be delivered after a fixed lead time ℓ. The replenishment occurs at the start of the delivery day ℓ. At the start of every day, the inventory is rationed. The rationing decision a_t sets how much products are made available for the offline channel on day t, with the remainder made available for the online channel. We do not assume that throughout the day the retailer checks for excess stock in one channel if the other channel has a stockout. A product sold through the online channel is not directly removed from stock (as shipment occurs at the end of the day), thus the retailer needs to be certain the product is not already sold. Additionally, online customers might have their product in their online basket, thinking the product is available while the retailer might be using their product to fulfil in-store demand. Furthermore, it is difficult to capture lost demand, as consumers who face empty shelves often walk out of the store. Hence to avoid lost sales in the offline channel by taking a product from the online channel requires keeping track of the in-store inventory level. But as the retailer often has many products, continuously replenishing the store from the back storage is cumbersome and might also influence the shopping experience of consumers. For the replenishment decision we assume the retailer orders every $R = 7$ days and faces a lead time (ℓ) of at most 7 days. After $R = 7$ days, the process repeats with the inventory level at the start of week $m+1$, equal to the closing inventory of week m.

3.2. MDP model

The problem described above can be formulated as a MDP, where the demand in the channels introduces uncertainty in the MDP. With the ordering and rationing decision the retailer can control its inventory levels in both channels. As the MDP consists of two different decisions on different time periods, it can be integrated in a Hierarchical Markov Process as described in Kristensen (1988).

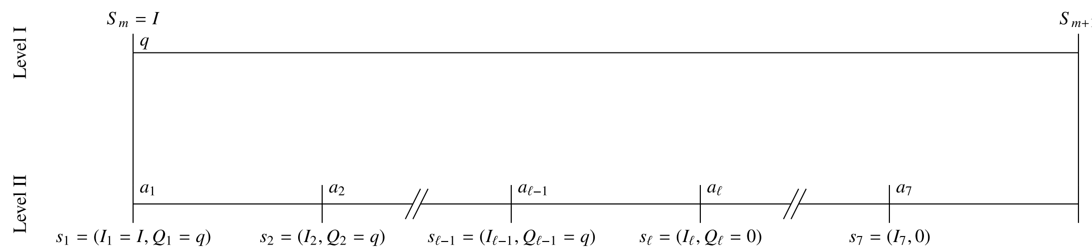

Fig. 1. Visualisation of the actions for the hierarchical problem.

3.2.1. States and state spaces

The state at level I of the MDP consists of the inventory level I at week m and is given as $S_m = I$. It is assumed that the retailer will not have more products in store than the maximum demand until the next replenishment. As holding cost are positive and no fixed ordering cost applies, there is no motivation for the retailer to order more than the maximum demand as excessive stock would negatively influences the profit. Therefore, the state space at level I is $I \in \mathcal{I}$ with $\mathcal{I} = \{0, 1, \ldots, (7 + \ell) \cdot D\}$, in which the maximum demand of a day is $D = \sum_i D_i$. D_i indicates the maximum possible demand of the individual channels with $i \in \{1 = \text{offline}, 2 = \text{online}\}$. The demand in both channels during each sub-period is modelled by Poisson distributions, $P_i(d_i)$ is the probability that the demand of channel i is d_i where $d_i \in \{0, 1 \ldots, D_i\}$.

The state at level II consist of the inventory level and the outstanding replenishment quantity and is given as $s = (I_t, Q_t)$. The state space of the inventory at level II is the same as at level I, $I_t \in \mathcal{I}$. The state space of the replenishment quantity depends on expected demand and the current inventory level, $Q_t \in \mathcal{Q}_t(I_t)$ with $\mathcal{Q}_t(I_t) = \{0, 1, \ldots, 7 \cdot D - I_t\}$. The upper bound of the state space is set under the assumption that the retailer will not replenish more products than total maximum expected demand per 7 days minus the current inventory.

3.2.2. Actions and action spaces

At level I of the MDP, a replenishment quantity q is set at the beginning of the period. The action space is thus equal to the state space of the replenishment quantity at level II thus $q \in \mathcal{Q}_t(I_t) = \{0, 1, \ldots, 7 \cdot D - I_t\}$.

At level II of the MDP, the rationing across the inventory a_t is set. It indicates how many products of the total inventory is withheld from the online channel and placed in the offline channel. The action space of a_t is dependent on the inventory I_t at sub-period t since the rationing quantity is clearly limited to the current inventory. Thus, a_t can be defined as $a_t \in \mathcal{A}_t(I_t)$ with $\mathcal{A}_t(I_t) = \{0, 1, \ldots, I_t\}$. We assume that the rationing decision is made at the beginning of the sub-period and will not be revised during the sub-period.

3.2.3. Transitions

At level I of the MDP, the state only transitions at the end of the week, while at level II the states transition every day. At level II of the MDP, we model the transition from state $s_t = (I_t, Q_t)$ to $s_{t+1} = (I_{t+1}, Q_{t+1})$. The transition of I_t to I_{t+1} depends on the current inventory, the rationing decision, demand of the individual channels, and the inventory replenishment:

$$I_{t+1} = (a_t - d_1)^+ + (I_t - a_t - d_2)^+ + \delta(\ell = t) \cdot Q_t \tag{1}$$

where $x^+ = \max(x, 0)$ and $\delta(x)$ denotes the Kronecker delta, which returns the value 1 if $x = \text{True}$, otherwise 0. Eq. (1) refers to the inventory being rationed across the two channels and the replenishment, with a_t and $I_t - a_t$ being the inventory level of the individual channels at weekday t from which the fulfilled demand of the individual channels is subtracted from. The replenishment takes place at $t = \ell$, which is at the beginning of day ℓ.

The transition of Q_t to Q_{t+1} depends on the weekday. At the first weekday the state Q_1 takes the value of replenishment quantity q. Q_{t+1} remains Q_t until the replenishment is added to the stock at time ℓ after which it is set to zero:

$$Q_{t+1} = \begin{cases} Q_t & \text{if } 1 < t \leq \ell \\ 0 & \text{else} \end{cases} \tag{2}$$

At level I of the MDP, the next state is noted by S_{m+1} which equals is the closing inventory of the week m, which is obtained from the last state transition on level II of the MDP:

$$S_{m+1} = (a_7 - d_1)^+ + (I_7 - a_7 - d_2)^+ + \delta(\ell = 7) \cdot Q_7 \tag{3}$$

3.2.4. Expected profit

The objective of the MDP is to maximise the expected profit (over an infinite horizon) which consists of revenue generated from selling products, and the cost of shipment, holding, and replenishment. At level I of the MDP the cost only consists of the replenishment cost:

$$\mathbb{E}C_{\mathrm{I}}(q) = -c \cdot q \tag{4}$$

The replenishment cost is based on a variable replenishment cost c and the replenishment quantity q.

For level II of the MDP the expected profit depends on the revenue from selling products, cost of shipment and holding cost of the individual channels. The expected profit depends on the action a_t taken in state $s_t = (I_t, Q_t)$:

$$\mathbb{E}C_{\mathrm{II}}(s_t, a_t) = \begin{cases} p\left(\sum\limits_{d_1 < a_t} d_1 \cdot P_1(d_1) + \sum\limits_{d_1 \geq a_t} a_t \cdot P_1(d_1)\right) \\ +(p-u)\left(\sum\limits_{d_2 < I_t - a_t} d_2 \cdot P_2(d_2) + \sum\limits_{d_2 \geq I_t - a_t} (I_t - a_t) \cdot P_2(d_2)\right) \\ -\left(h_1 \cdot a_t + h_2 \cdot (I_t - a_t)\right) \end{cases} \tag{5}$$

The first term is the revenue from the offline channel for price p and the second term the revenue from the online channel. The same sales price is applied but when satisfying an online demand from the store a unit shipment cost u is incurred. The quantity sold through one channel depends on the demand and rationing. Shortage cost is not included into this model, if shortage occurs the retailer faces a lost sales causing in a penalty by losing profit margins. Next to shipment costs, the model accounts for the holding cost in the last term.

3.3. Value iteration

The aim of the MDP is to maximise the long-term weekly expected profit. With value iteration, the problem is solved iteratively backwards, where one iteration relates to a single sub period, e.g. a day. Hence a sequence of $R = 7$ iterations relate to one week. To ease the explanation of the algorithm, the two levels of the MDP are integrated into one level.

Let n be the iteration counter. The maximum expected profit over n consecutive days when starting in state s_t is defined as $v_n(s_t)$. The long run weekly profit is thus $g = \lim_{n\to\infty}(v_n - v_{n-7})$, which does not depend on the initial state s, and is thus the same for all s (the value of g is called the gain of the underlying Markov chain). To determine the value of g by value iteration one starts setting $v_0(s) = 0$ for all states s. Next, one computes for all s : $v_1(s_t) = \max_{a_t\in\mathcal{A}_t(I_t)}\{\mathbb{E}C_{\mathrm{II}}(s_t, a_t)\}$, and continues by computing v_2, v_3, etc. using the so called recursive Bellman equations in (6) and (7). The day t can be calculated from the iteration n as follow: $t = (7 - (n-1) \mod 7)$. The distinction of day t is relevant to keep track of whether next to the rationing decision an order decision should be taken or an outstanding order has arrived.

For $t = 1$, the Bellman equation incorporates ordering and rationing decision and the respective costs and revenues to maximise the expected profit:

$$v_n(s_1) = \max_{q\in\mathcal{Q}_1(I_1)}\Bigg\{\mathbb{E}C_{\mathrm{I}}(s_t, q) + \max_{a_1\in\mathcal{A}_1(I_1)}\Bigg\{\mathbb{E}C_{\mathrm{II}}(s_1, a_1) + \sum_{d_1=0}^{D_1}\sum_{d_2=0}^{D_2} P_1(d_1)\,P_2(d_2)\,v_{n-1}(s_2)\Bigg\}\Bigg\} \tag{6}$$

For $t = 2, 3, \ldots, 7$ the Bellman equation has to consider the rationing decisions:

$$v_n(s_t) = \max_{a_t\in\mathcal{A}_t(I_t)}\Bigg\{\mathbb{E}C_{\mathrm{II}}(s_t, a_t) + \sum_{d_1=0}^{D_1}\sum_{d_2=0}^{D_2} P_1(d_1)\,P_2(d_2)\,v_{n-1}(s_{t+1})\Bigg\} \tag{7}$$

As $n\to\infty$, the span of the average weekly expected profit decreases, thus $\|v_n - v_{n-7}\|$ converges to 0, which implies that the difference between the largest and the smallest element of $v_n - v_{n-7}$ becomes zero, and all elements equal the maximum expected long run weekly profit. If $\|v_n - v_{n-7}\|$ is smaller than ϵ the value iteration stops, we specify $\epsilon = 0.001$. Algorithm 1 formalises the algorithm and shows how value iteration can be implemented. In particular it demonstrates that the two maximisation actions in Eq. (6), can be implemented in serial to reduce the computation complexity of value iteration. The algorithm applies backward induction. The first for loop deals with the rationing decisions at level II. Next the ordering decision is optimised and the procurement costs are added to v_n. As the average weekly expected profit converges to an vector consisting of equal values, the value iteration is stopped.

Algorithm 1: Value iteration of the MDP

```
initialisation: n = 0; v_0 = 0;
repeat
    for t = 7 : −1 : 1 do
        n = n + 1
        for s_t = (I_t, Q_t) ∈ {(I_t, Q_t) | I_t ∈ 𝓘, Q_t ∈ 𝓠_t(I_t)} do
            v_n(s_t) =
              max_{a∈A_t(I_t)} { 𝔼C_II(s_t, a) + Σ_{d_1=0}^{D_1} Σ_{d_2=0}^{D_2} P_1(d_1) P_2(d_2) v_{n−1}(s_{t+1}) }
            where: s_{t+1} is from equation (1) and (2)
        end
    end
    for I_1 ∈ 𝓘 do
        w(I_1) = max_{q∈𝓠(I_t)} {𝔼C_I(q) + v_n(s_1)} where: s_1 = (I_1, q)
    end
    for I ∈ 𝓘 do
        v_n(I, 0) = w(I)
    end
until ‖v_n − v_{n−7}‖ < ε;
```

The (nearly) optimal strategy for the two levels of the MDP can be obtained from the results of the value matrices. First the optimal replenishment policy $\pi^{\mathrm{I}}(s_1)$ at $t = 1$ for level I is found by Eq. (8):

$$\pi^{\mathrm{I}}(s_1) = \arg\max_{q\in\mathcal{Q}_1(I_1)}\Bigg\{\mathbb{E}C_{\mathrm{I}}(s_t, q) + \max_{a_1\in\mathcal{A}_1(I_1)}\Bigg\{\mathbb{E}C_{\mathrm{II}}(s_1, a_1) + \sum_{d_1=0}^{D_1}\sum_{d_2=0}^{D_2} P_1(d_1)\,P_2(d_2)\,v_{n-1}(s_2)\Bigg\}\Bigg\} \tag{8}$$

Second, the optimal rationing decision $\pi_t^{\mathrm{II}}(s_t)$ for all states in $t = 1, 2, \ldots, 7$ for level II is found by Eq. (9):

$$\pi_t^{\mathrm{II}}(s_t) = \arg\max_{a_t\in\mathcal{A}_t(I_t)}\Bigg\{\mathbb{E}C_{\mathrm{II}}(s_t, a_t) + \sum_{d_1=0}^{D_1}\sum_{d_2=0}^{D_2} P_1(d_1)\,P_2(d_2)\,v_{n-1}(s_{t+1})\Bigg\} \tag{9}$$

Table 2
Parameter values.

Parameter	Values
Review period	2, 3, 4, 5, 6, 7
Lead time	2, 3, 4, 5, 6, 7
μ_1	2, 4, 6
μ_2	2, 4, 6
p	100
c	20, 30, 40
u	0, 5, 20
h_1	1, 2, 3
h_2	0.5

4. MDP computational complexity and results

4.1. Data and design of experiments

The optimal policy is investigated for a range of instances, consisting of a base test case and various alterations. The data set used is based on recent literature on similar omnichannel retail environments (Jalilipour Alishah et al., 2015; Li et al., 2015; Ovezmyradov and Kurata, 2019; Xu and Cao, 2019; Bayram and Cesaret, 2021). We alter different parameters to investigate the performance of the MDP. In Table 2, the parameter values are found for the review periods, lead time, mean daily demand, cost of product, fulfilment cost, and holding cost of the channels are given. Without loss of generality, we set $p = 100$ and $h_2 = 0.5$, and vary the other cost parameters in relation to these fixed parameters. The different instances are all created by varying specific subsets of parameters.

One typical instance is used as a base test case. The demand is independently Poisson distributed with a mean daily demand of $\mu_1 = 6$ and $\mu_2 = 2$, with a maximum demand of $D_1 = 12$ and $D_2 = 6$. In line with related papers, such as Jalilipour Alishah et al. (2015), Xu and Cao (2019) and Bayram and Cesaret (2021), we assume no substitution between offline and online demand. The cost of the product $c = 30$, the handling cost incurred for satisfying an online order with store inventory is $u = 5$. The holding cost of the offline is $h_1 = 1$. Replenishment orders are placed on Monday ($t = 1$) and delivered on Wednesday ($t = 3$), where both events occur at the beginning of the day. In the remainder of this paper, if an instance does not specify a certain parameter value, it will be equal to their base test case setting.

4.2. Computational implementation

The number of transitions from one state to the other depends on the action space and demand distribution. The action space is already limited by inventory and demand level. As is common in numerical analysis of MDPs, we truncate the demand distribution to cover 99%

Table 3
Computational complexity of the MDP for different instances.

Instances	States	CPU time (s)	RAM (MB)
$R=7, l=1$	127 890	180	52.56
$R=7, l=2$[a]	143 766	448	74.14
$R=7, l=3$	159 642	810	101.00
$R=7, l=4$	175 518	1570	133.72
$R=7, l=5$	191 394	2710	172.89
$R=7, l=6$	207 270	3780	219.10
$R=7, l=7$	223 146	5110	270.41
$R=2, l=2$	5 256	33	6.33
$R=3, l=2$	14 742	76	12.75
$R=4, l=2$	31 392	144	22.00
$R=5, l=2$	57 150	204	34.91
$R=6, l=2$	93 960	321	52.09
$R=3, l=3$	17 658	163	21.33
$R=4, l=4$	41 760	505	50.52
$R=5, l=5$	81 450	1430	98.61
$R=6, l=6$	140 616	2970	170.33
$\mu_1=2, \mu_2=2$	64 092	173	25.43
$\mu_1=4, \mu_2=4$	143 766	478	65.01
$\mu_1=2, \mu_2=6$	143 766	478	74.14

[a]Base test case.

of the cumulative distribution function to limit the state space. The demand distribution is reshaped as a right-truncated Poisson distribution as described in Cohen (1954).

Additionally, to ensure that enough memory is available to compute all the expected reward matrices, not all matrices are stored. Matrices at level I older than $n-1$ are deleted. At level II, matrices older than $t+1$ are not stored. However, to find the optimal rationing policy, matrices older than $t+1$ are needed to derive the optimal policy. This is solved by storing the action a_t in $\pi^{\mathrm{II}}(s_t)$ after the second Bellman equation. In the last, iteration the optimal rationing strategy is approximated by the action $\pi^{\mathrm{II}}(s_t)$.

Since the transition probabilities require high computation time, they are calculated beforehand and stored in a matrix. This ensures that all transition probabilities are only calculated once. As the amount of transition probabilities is only $4.3 \cdot 10^6$, it is possible to store them.

For solving the MDP, the policy needs to be calculated for a large number of states. The results were obtained by implementing the MDP in Python version 3.7.2. The model was run on a Personal Computer with Intel Xeon W-2133 CPU @ 3.60 GHz and 16 GB of RAM. Table 3 shows for all instances the total number of states, the CPU time in seconds, and RAM usage in megabyte (MB) for the MDP. The instances with different cost structures are not presented as they do not increase the complexity of the model.

From Table 3 it is observed that an increase in lead time or review period increases the computational complexity of the MDP. A longer review period increases the inventory and ordering state space. A longer lead time means that the ordering state is included in more states, increasing the total amount of states. The total amount of states grows linearly with the lead time, but exponentially with the review period. The time to solve the MDP grows with the number of states; as the problem becomes bigger, more RAM is needed to solve an instance. A higher average daily demand also increases the complexity of the MDP, as the maximum demand increases with a higher average demand. As the required number of states for the inventory and ordering decisions is related to the maximum demand, the total amount of states increases with a higher average daily demand.

Table 3 shows that solving the MDP can take more than one hour, which is impractical for retailers. To avoid long computation time heuristics are preferred. By analysing the structure of the optimal policy, we can derive effective rules that are able to encompass the complexity of the MDP.

4.3. Structure of the optimal policy

We numerically investigate the optimal policy and derive general rules from the structure of the policy so that the heuristics are applicable for all instances presented in Table 2. The structure of the optimal replenishment policy and rationing policy is identified for the base test case.

4.3.1. Replenishment policy

Fig. 2(a) shows the order-up-to level at the day when orders are placed for different inventory levels. The order-up-to level is given by adding the ordered quantity to the inventory level. From the figure it can be concluded that at lower inventory levels the optimal replenishment policy resembles an (R, Q) replenishment policy, as the replenishment quantity remains (almost) the same. At higher inventory levels, it resembles a base stock policy as observed by the maximum order-up-to level at inventory levels above 22.

In order to get better insight in which order-up-to levels occur more frequently the optimal policy is simulated for 100 000 weeks. From the simulation the frequency of different inventory levels can be found, as seen in Fig. 2(b). It is seen that both replenishment policies occur with the base stock policy being more common. As the average demand is $\mu_1 + \mu_2 = 8$, and the inventory being below 8 is 0.2%, the chance of a lost sale is small.

By using an approach similar to Haijema et al. (2007), a frequency table is used to find the structure of the optimal policy for different inventory and order-up-to levels. Table 4 shows the simulation results of the frequency of each order-up-to level with inventory level. The total frequency of each order-up-to level is presented in the last column. From this table, the optimal order quantities can also be derived. For instance, when looking at the columns for inventory level 17 and 18, we can see that in 100 000 simulations, there were 1 523 occasions in which the retailer had an inventory level of 17 upon ordering, and 1 570 occasions in which the inventory level was 18. For both these cases, the optimal order-up-to level is 83 (represented in the row).

From Table 4 it is observed that for inventory levels 0 to 10 no fixed order-up-to level but a fixed order quantity of 68 applies. This indicates that the optimal policy expects and anticipates that the available inventory will be sold during the procurement lead time. At higher inventory levels there is a chance that leftovers occur upon replenishment, thus lower quantities are ordered. When plenty of inventory is left, an uncertain amount is carried over to the next week. For stock levels higher than 22 a base stock policy applies, which occurs around 85% of the time. At most 17 products are reserved for demand during lead time, as the base stock is 85 and the optimal inventory level is 68.

The behaviour of the optimal replenishment policy can be explained by the lead time. Current orders that are placed will be replenishment after the lead time has exceeded. The optimal inventory level when orders are replenished is 68, thus the replenishment policy tries to achieve this. At low inventory levels it is expected that the stock is depleted at $t = \ell$, thus 68 products are ordered, following an (R, Q) order policy. At higher inventory levels, a certain amount of left-over products is expected at $t = \ell$, and to compensate for these left-overs they are subtracted from the order quantity, resembling an (R, S) order policy.

4.3.2. Rationing policy

Fig. 3(a) displays the optimal rationing policy at day 2, which is the day before replenishment. Day 2 is chosen for this analysis, as the day before replenishment inventory levels are lowest, thus rationing is most important. The order quantity 62 is chosen, as it is the most frequent order quantity. At high inventory levels, the amount of in-store inventory is a fixed amount as additional in-store inventory would not be worth the additional costs. This base stock level can also be proven analytically, which can be seen in Appendix. At low inventory levels, there is a choice between whether to store products in-store or offline.

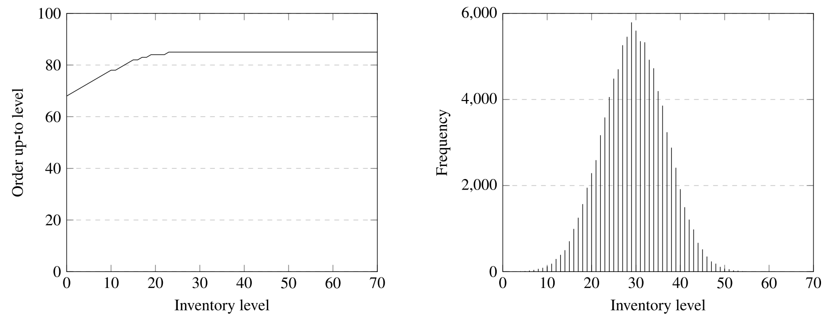

Fig. 2. Order-up-to level (a) and inventory frequency (b) for different inventory levels at the beginning of day 1.

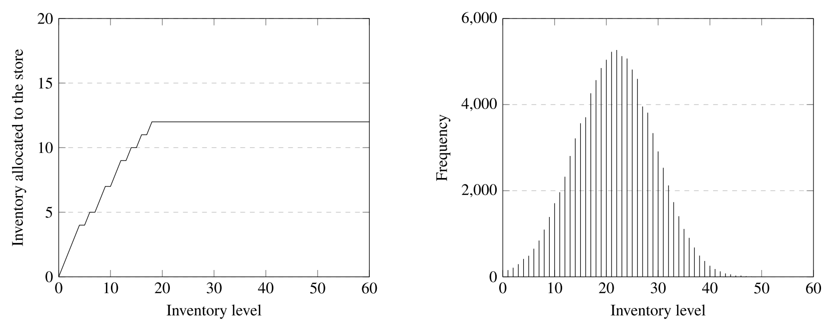

Fig. 3. Rationing policy (a) and inventory distribution (b) of different inventory levels at the beginning of day 2 if the order quantity is 62.

Table 4
Frequency table of order-up-to level with inventory level.

Inventory	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	(23,24,…,70)	Total
Up-to 85																								83 733	83 733
84																				1950	2289	2591	3170		10 000
83																		1523	1570						2823
82																707	994								1701
81															499										499
80														391											391
79													295												295
78											125	185													310
77										90															90
76									66																66
75								41																	41
74							26																		26
73						11																			11
72					8																				8
71				2																					2
70			1																						1
69		3																							3
68	0																								0
..																									..
0																									0

This trade-off is based on different drivers such as expected future sales and holding cost. It is seen that the trade-off between storing a product in-store or online is almost a fixed ratio.

From Fig. 3(b) it is seen that there is a risk of shortage, as the expected average daily demand is 8 and inventory levels below 8 occur in a significant number of instances. For 95.7% of the weeks, the inventory is able to satisfy the average daily demand at day 2. For 4.3% of the weeks, the inventory is however below 8. Comparing Figs. 2(b) to 3(b), the average inventory level shifts from 30 to 22. This shift is expected for the inventory levels of two consecutive days if the average daily demand is 8.

5. Heuristic

Solving the MDP for the base case takes around eight minutes for a single product and can take over an hour for instances with long lead times. Heuristics to reduce computational time are then preferred, especially when a retailer manages multiple products. Therefore, we develop heuristics for the ordering and rationing decision that approximate the structure of the optimal policy identified above.

From Figs. 2 and 3, it is clear that the ordering and rationing decisions turn out to have a threshold. When the inventory level is below the threshold, the optimal ordering decision can be approximated by, a fixed order quantity. At high inventory levels, a base-stock policy seems to fit. For the rationing decision, at low inventory levels a trade-off exists between the channels which is driven by expected future sales

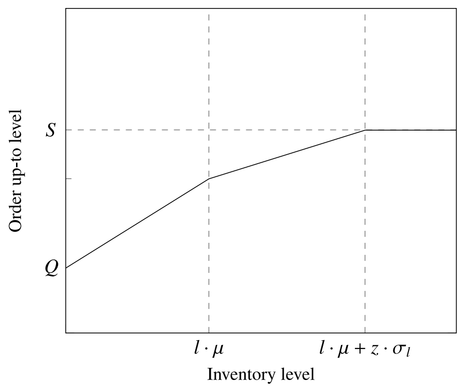

Algorithm 2: Heuristic for ordering

$z = F^{-1}\left(\frac{p-c}{p-c+R\cdot h_2}\right)$

$\mu = \sum_i \mu_i$

$Q = R \cdot \mu + z \cdot \sigma_R$

$S = (R + \ell) \cdot \mu + z \cdot \sigma_{R+\ell}$

if $I < \ell \cdot \mu$ **then**
  $q = Q$
else if $\ell \cdot \mu \leq I \leq \ell \cdot \mu + z \cdot \sigma_\ell$ **then**
  $w = (I - \ell \cdot \mu)/(z \cdot \sigma_\ell)$
  $q = \text{round}\left((1 - w) \cdot Q + w \cdot (S - I)\right)$
else
  $q = S - I$
end

Fig. 4. Visualisation of the parameters of the ordering heuristic (a) with the algorithm (b).

and holding costs. At high inventory levels, we already identified an optimal strategy which consists of a fixed amount of products stored in-store (found in Appendix). In the following, these results and insights are used to develop ordering and rationing heuristics.

5.1. Ordering heuristic

The heuristic for ordering is based on having a fixed order quantity at low inventory levels and a base-stock policy at high inventory levels, as it was found in Fig. 2 to be optimal. In between the low and high inventory levels the optimal ordering policy is a combination of a fixed order quantity and base-stock policy.

Fig. 4(a) illustrates the structure found for the optimal ordering policy. In the figure, two inventory level points are presented which are used to derive a heuristic for ordering. The first inventory level point indicates that if the inventory is below the level it is assumed there is no inventory when replenishment occurs. This inventory level point is found from the average sales procurement lead time ($\ell \cdot \sum_i \mu_i$), as this is the amount of products that likely will be sold before replenishment occurs.

The second inventory level point indicates the point in which it is certain there is left-over inventory at replenishment. This is given if the inventory level exceeds the average sales with safety stock procurement lead time $\left(\ell \cdot \sum_i \mu_i + z \cdot \sigma_\ell\right)$, it is unlikely the retailer will sell more products during lead time before replenishment occurs. The safety stock z follows the critical fractile ratio $\left((p-c)\,/\left(p-c+R\cdot h_2\right)\right)$. As the safety stock is often a formula consisting of the lost sale cost and holding cost, it is assumed in this heuristic that the lost sale cost is equal to the profit of selling a product in the offline channel. Reason is that the optimal policy is to sell the product through the offline channel. The holding cost is assumed to be from the online channel, since excess stock is preferred to be stored in the online channel. F^{-1} denotes the inverse of the Normal distribution function with mean zero and variance of one. σ_i is the standard deviation of the total demand during period i.

The region between the two inventory level points indicate there is uncertainty if there will be inventory left-over at replenishment or not. If the inventory level is closer to the average sales procurement lead time it is more likely that at replenishment the inventory level is zero. However, if the inventory is closer to the average sales with safety stock procurement lead time there is a higher change of left-overs at replenishment.

The region which assumes that no inventory is left at replenishment follows an fixed order quantity policy, where Q is the order quantity as seen in Fig. 4(a). The order quantity is derived from the average sales with safety stock during the review period and is calculated as follow: $R \cdot \sum_i \mu_i + z \cdot \sigma_R$.

If the inventory level exceeds the average sales with safety stock procurement lead time, it is certain that there is left-over inventory at replenishment thus a base-stock policy is used. The order-up-to level S is derived from the average sales during the review period and lead time. The order-up-to level uses the same formula as the fixed order quantity however, takes the average sales and safety stock over the review period plus lead time: $(R+\ell) \cdot \sum_i \mu_i + z \cdot \sigma_{R+\ell}$.

In between the fixed order quantity and base-stock policy, the optimal policy is a combination of the two. Therefore, we linearly interpolate the two policies in the region where the inventory level is between the two inventory level points. By using a weight ratio (w) we decide whether the order policy should follow more a fixed order quantity or base-stock policy. The weight ratio is a linear interpolation between the two inventory level points where the values ranges between 0 and 1. By using the weight ratio the fixed order quantity policy changes to a base-stock policy as the inventory level increases.

Fig. 4(b) gives the algorithm for ordering, where first the parameters are derived followed by checking in which region the inventory level is. As no closed form exist for the inverse of the Normal distribution function, this takes a numerical procedure. Except for this calculation, the heuristic is able to find the order quantity fast.

5.2. Rationing heuristic

The heuristic for rationing is based on operating the channels individually and trying to find the optimal quantity to allocate to each channel by considering the marginal costs and revenue for each additional product allocated to the two channels. Excess products the retailer has available are stored in the channel with the lowest inventory cost. The optimal quantity for each individual channel is found using the newsvendor model, however, the ordering cost is replaced by the holding cost. If the inventory level is below the optimal quantity of the individual channels, the decision on where to allocate a product is based on which channel gives the highest expected contribution for the product. G_i^{-1} denotes the inverse of the Poisson distribution function of either channel.

Fig. 5(a) illustrates the structure found for the optimal rationing policy. In the figure the threshold inventory level where excess products are stored in the channel with the lowest inventory cost is indicated as $r_1 + r_2$. This inventory level point is found by calculating the optimal quantity for each individual channels. At inventory levels below this threshold, a trade-off is made between storing the product in the offline or online channel. This trade-off is based on which channel gives the highest expected contribution for an extra product.

Fig. 5(b) gives the algorithm for rationing, where first the optimal quantity of the channels is derived as this takes a numerical procedure. The calculation of the highest expected contribution for the product is cumbersome however, unavoidable for the heuristic. Nevertheless, the heuristic for the rationing is able to find the rationing decision fast.

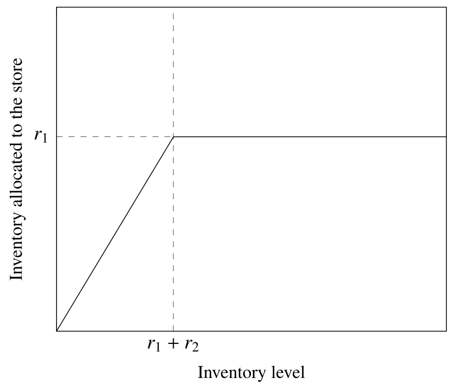

Algorithm 3: Heuristic for rationing

$r_1 = G_1^{-1}\left(\frac{p-(h_1-h_2)}{p}\right)$, $r_2 = G_2^{-1}\left(\frac{(p-u)-h_2}{(p-u)}\right)$

if $I \geq r_1 + r_2$ **then**
  $a = r_1$
else
  $i = 0, a = 0$
  while $i < I$ **do**
    $i = i + 1$
    $s = (i, 0)$
    if $\mathbb{E}C_{II}(s, a+1) > \mathbb{E}C_{II}(s, a)$ **then**
      $a = a + 1$
  end
end

Fig. 5. Visualisation of the parameters of the rationing heuristic (a) with the algorithm (b).

6. Performance of the heuristics

The result of the heuristics is evaluated for the different problem instances described in Section 4.1. We evaluate the heuristics on three different performance measures, the goodness of fit of the policy, the optimality gap and the service level. The first performance is to evaluate how much the heuristics results resembles the optimal policy, the second performance indicator gives us an indication how good the heuristics performs on the different revenue and cost parameters. The last performance measures gives an indication on the performance of the fulfilment of the individual channels.

6.1. Goodness of fit

To evaluate the goodness of fit of the heuristic, a weighted root-mean-square-error (wRMSE) is used. RMSE sums the squared differences between the optimal actions and the actions set by the heuristic over all feasible states. We chose to use RMSE because it penalises large deviations from the optimal decisions more than small deviations. Furthermore, we chose to use a weighted version of this measure because not all states are equally likely to occur, and it is important that the heuristics fits well to the optimal policy in the states that are most predominant. To achieve this, we use weights set by the probabilities that the states occurs. These probabilities can be obtained from solving the steady state distribution by Markov chain analysis, or by only considering actions that are made during a simulation period. We apply simulation, and compute the wRMSE with weights set to the relative frequency that states are visited during a simulation. Similar as described in Section 4.3, we simulate the heuristics for J periods and calculate the wRMSE as follows:

$$\text{wRMSE} = \sqrt{\frac{\sum_{j=1}^{J}\left(\pi_j - k_j\right)^2}{J}} \tag{10}$$

where π_j is the optimal action to be taken in simulation period j and k_j is the action chosen by the heuristic in the same period. For the ordering action the simulation period $J = 100\,000$ weeks and for the rationing action $J = 700\,000$ days. By using a relatively long simulation period, we implicitly include the frequency of states occurring, as the error encountered for common states will be included many times in this weighted average. In Table 5 the wRMSE of both the ordering and rationing decision is presented.

From Table 5 it is observed that both the ordering and the rationing decision show low weighted RMSEs, where the maximum values are 1.699 and 1.031 respectively. For different lead times, it is observed that the weighted RMSE of ordering varies the most, where it increases with longer lead times. This can be explained by the fact that the heuristics tries to predict the quantity of products sold during lead time, and for shorter lead times this prediction is more accurate. For different review periods the weighted RMSE stays around 1.25.

Table 5
Weighted RMSE of the heuristics during review period for all instances.

Instance	Weighted RMSE	
	Ordering	Rationing
$R = 7, \ell = 1$	0.390	0.646
$R = 7, \ell = 2$[a]	0.663	0.639
$R = 7, \ell = 3$	1.066	0.681
$R = 7, \ell = 4$	0.954	0.664
$R = 7, \ell = 5$	1.272	0.645
$R = 7, \ell = 6$	1.699	0.620
$R = 7, \ell = 7$	1.548	0.658
$R = 2, \ell = 2$	1.531	0.700
$R = 3, \ell = 2$	1.678	0.761
$R = 4, \ell = 2$	0.980	0.645
$R = 5, \ell = 2$	0.998	0.677
$R = 6, \ell = 2$	1.058	0.702
$R = 3, \ell = 3$	0.893	0.589
$R = 4, \ell = 4$	0.752	0.611
$R = 5, \ell = 5$	1.198	0.614
$R = 6, \ell = 6$	1.658	0.597
$\mu_1 = 2, \mu_2 = 2$	0.259	0.131
$\mu_1 = 4, \mu_2 = 4$	0.666	0.857
$\mu_1 = 2, \mu_2 = 6$	0.592	0.122
$u = 0$	0.592	0.649
$u = 20$	0.899	0.593
$h_1 = 2$	1.237	1.031
$h_1 = 3$	1.361	0.446
$c = 20$	0.980	0.501
$c = 40$	0.499	0.970

[a]Base test case.

The goodness of fit of the rationing decision is relatively stable for different review periods and lead times. The demand and cost parameters have the largest influence on the weighted RMSE, which is expected as these parameters are used in the rationing heuristic thus have the largest influence.

6.2. Optimality gap

Table 6 shows the profit of the MDP and the heuristics. Additionally, the gap of the revenue and all individual cost factors are included. From the results of the profit, it appears that the heuristics perform quite well and is very close to the optimum in all cases. The largest gap in profit is 0.064% for the instance in which $h_1 = 2$. The heuristics performance are the best for the instance in which $\mu_1 = 2, \mu_2 = 2$ with an gap of only 0.002%. These results correspond with the weighted RMSE as in these instances they were also the smallest and largest.

From the performance results, it is observed that the heuristics are performing better for shorter lead times than for longer lead times, even though it is still close to optimal. As described above, the heuristics

Table 6
Optimality gap of the heuristics during review period for all instances.

Instance	Profit			Revenue	Costs		
	MDP	Heuristic	Gap (%)	Gap (%)	Inventory Gap (%)	Ordering Gap (%)	Shipment Gap (%)
$R = 7, \ell = 1$	3626.63	3626.30	−0.009	0.031	0.858	0.020	−0.014
$R = 7, \ell = 2$[a]	3623.84	3623.42	−0.011	0.042	1.007	0.041	0.004
$R = 7, \ell = 3$	3621.15	3620.55	−0.017	−0.026	−0.171	−0.027	−0.089
$R = 7, \ell = 4$	3618.69	3618.07	−0.017	0.057	1.339	0.057	0.028
$R = 7, \ell = 5$	3616.36	3615.56	−0.022	0.079	1.934	0.061	0.063
$R = 7, \ell = 6$	3614.18	3613.05	−0.031	0.147	3.130	0.143	0.160
$R = 7, \ell = 7$	3612.09	3611.02	−0.030	0.106	2.355	0.106	0.107
$R = 2, \ell = 2$	1057.47	1056.99	−0.045	−0.175	−3.828	−0.176	−0.348
$R = 3, \ell = 2$	1579.53	1578.81	−0.046	−0.182	−3.606	−0.183	−0.326
$R = 4, \ell = 2$	2097.13	2096.78	−0.017	−0.061	−1.060	0.058	−0.146
$R = 5, \ell = 2$	2610.35	2609.49	−0.033	−0.060	−0.752	0.038	−0.146
$R = 6, \ell = 2$	3119.23	3118.46	−0.025	−0.055	−0.678	−0.048	−0.139
$R = 3, \ell = 3$	1577.80	1577.64	−0.010	−0.030	−0.332	−0.036	−0.284
$R = 4, \ell = 4$	2093.23	2092.93	−0.014	0.045	−3.253	0.508	−0.001
$R = 5, \ell = 5$	2604.00	2603.49	−0.020	0.076	1.992	0.071	0.052
$R = 6, \ell = 6$	3110.24	3109.30	−0.030	0.146	3.695	0.144	0.159
$\mu_1 = 2, \mu_2 = 2$	1762.99	1762.95	−0.002	0.008	0.166	0.008	0.007
$\mu_1 = 4, \mu_2 = 4$	3561.35	3560.84	−0.014	0.070	1.652	0.069	0.008
$\mu_1 = 2, \mu_2 = 6$	3507.54	3507.41	−0.004	0.016	0.399	0.015	0.017
$u = 0$	3693.39	3692.97	−0.011	0.039	0.942	0.038	0.000
$u = 20$	3415.46	3414.91	−0.016	0.068	1.340	0.066	0.165
$h_1 = 2$	3542.67	3540.40	−0.064	0.312	5.043	0.311	0.041
$h_1 = 3$	3467.23	3465.88	−0.039	0.188	2.447	0.186	0.052
$c = 20$	4180.92	4180.47	−0.011	−0.055	−3.433	0.406	−0.107
$c = 40$	3067.30	3066.60	−0.023	0.080	1.663	0.080	−0.015

[a]Base test case.

try to predict the quantity of products sold during lead time, which is more accurate for short lead times. This can also be seen by the cost of ordering, which increases with the lead time.

For short review periods, the performance results are the lowest. For short review periods, the amount of days with low inventory is relatively high compared with long review periods. The rationing decision influences the revenue and inventory cost, which can be seen as highest in the cases with short review periods. The high shipment cost gap indicates that more product are sold in the online channel than offline channel, while a product in the offline channel has a larger profit margin. The gap for profit at higher review periods remains around 0.020%.

For instances in which the review period is equal to the lead time, the gap is the lowest when they are both equal to three days. When decreasing or increasing the period the gap increases. The trend in gap is similar as discussed above, where longer lead time increase the gap and longer review period decreases the gap. However, the gap resulting from the sub-optimal rationing policy is more significant than the order policy, as the gap for short review period and lead time is larger than long review period and lead time. This can also be seen in the inventory cost, which gaps are significantly larger than the ordering gaps.

For different review periods and lead times, the heuristics are performing effective. Additionally the heuristics are tested for different demand distributions. Among all instances, the gap of the heuristics are the lowest for the different demand distributions.

The heuristics are also showing near-optimal results for instances with different cost parameters. Higher costs for handling online demand or higher costs for in-store inventory do increases the gap slightly. Having higher costs means that having sub-optimal rationing decisions decreases the profit more. Increasing the ordering cost decreases the gap of the heuristics. This can be explained by the approximation of the safety stock of the ordering heuristic. It assumes that a lost sale cost only consist of the profit of selling a product in the offline channel, not taking into account the possibility of the lost sale occurring in the online channel.

6.3. Service level

Next to profit an important performance indicator is the alpha service level, that is the fraction of time one ends a day with products still in stock (i.e., the non-stock out probability). We report these for every day and for each channel. In particular one is interested in the cycle service level, that is the product availability just before a replenishment arrives.

6.3.1. Service level per weekday

Fig. 6 shows the average service level per day resulting from using the policy of the MDP and heuristics for each channel. The service level is calculated by the fraction of 100 000 simulated weeks, in which all demand can be met from stock assigned to that channel. The figure shows that demand in the online channel is fully served for the days following replenishment (day 3 till 7), and that there is some lost demand during those days in the offline channel. This is due to the rationing decision, and the holding costs being higher at the offline channel: it deems optimal not to meet the maximum possible demand of that channel. During the procurement lead time (day 1 and 2), the service level drops, as the retailer is experiencing lower inventory levels, and increased probabilities that they cannot satisfy all demand. Overall, the heuristic has a slightly higher service level than the optimal result. This is due to a slightly larger replenishment order, which happens mostly when the stock level is low upon ordering.

At high inventory levels, the two channels are not competing for products and the rationing decision can be reduced to a cost minimisation decision to determine the optimal base stock level for the offline channel. As having more products in the offline channel than the base stock level causes the expected marginal revenue to become less than the expected marginal cost, the optimal solution is to store excess products in the online channel. In extreme cases the offline channel might not have enough products to meet all demand. As the online channel holds all excessive stock it is capable to satisfy all demand in extreme cases. Both the policy of the MDP and the heuristic have this

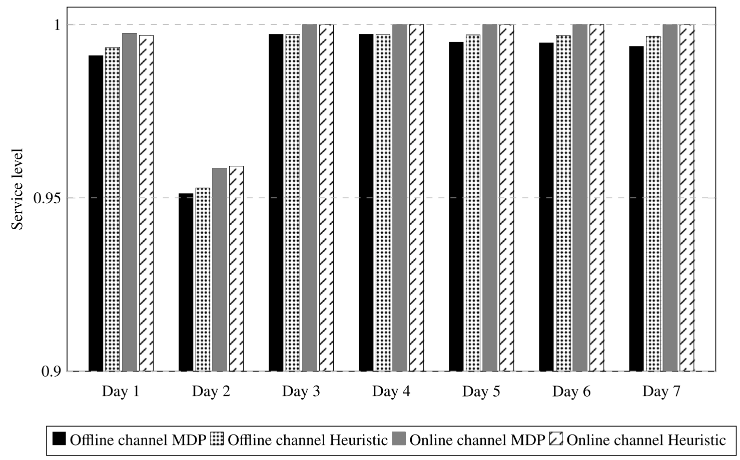

Fig. 6. Average service level per day of each channel for the MDP and heuristic for the base case.

base stock level for the offline demand, therefore the service level is identical for day 3 and 4.

The day before replenishment (day 2), the trade-off between serving the two channels is more important due to lower inventory levels. Therefore, it is important that the heuristic is able to capture the trade-off between expected future sales and holding cost as good as possible. As seen in Table 5, the rationing decision has a low RMSE, indicating that the heuristic is capable of capturing the core trade-off.

The overall higher service level was also seen in Table 7, where the revenue is higher for the heuristic. The heuristic has an overall higher service level due to an average higher inventory level. The higher inventory level results from the larger replenishment orders, which was seen in Table 6 where the ordering gap is positive, indicating higher ordering costs for the heuristic. The average higher inventory level comes at the expense of higher inventory costs, which offset the higher revenue resulting in a lower profit for the heuristic.

6.3.2. Cycle service level

The service level on day l sets the cycle service level, which is an important performance indicator. Where Fig. 6, focuses on the base case, in Table 7 we present the cycle service level for all instances. The cycle service level is the service level on the day just before replenishment. For most instances, the cycle service level is above 95%, only for the instances in which the holding cost is higher or the cost of the product is higher, the service level turns out to be lower. This can be explained by the fact that the cost of trying to satisfy demand is higher in these instances, and that it therefore is more profitable to not always satisfy demand. In settings where holding costs are much smaller, the fill rates may get close to one, and we observed that the performance of the heuristic is equally close or even closer to optimal.

Table 7 shows that the cycle service level is almost always slightly higher for the online channel than the offline channel, as the holding costs in the offline channel are larger. Due to the higher holding costs, fewer products are stored in the offline channel resulting in a lower service level.

The service level remains relatively stable for different lead times, as for the offline channel it is around 95% and for the online channel between 95%–96%. Though not visible here, we have observed that for a longer lead time the optimal policy resembles less a base stock policy. This is explained by the fact the actual inventory level gives less information on the inventory level when products are delivered after a longer lead time. The proposed heuristic closely approximates the optimal ordering policy by a combination of a constant order policy and a base stock policy, as discussed in Section 6.1. Thus we achieve comparable service levels.

When increasing the review period, the cycle service level decreases. Longer review periods increase the uncertainty of demand thus the chance of hitting a stock out before replenishment increases as the retailer is less responsive to low stock levels. This is also seen in the instances in which the review period is equal to the lead time.

For most instances, the MDP results in a lower service level than the heuristic. The heuristic often has higher replenishment orders than the MDP, resulting in the higher cycle service levels at the expense of higher inventory and ordering costs (as was shown in Table 6). However, for short review periods, the heuristic orders fewer products than the MDP resulting in lower cycle service levels.

Table 7
The cycle service level for all instances.

Instances	Offline channel		Online channel	
	MDP	Heuristic	MDP	Heuristic
$R=7, \ell=1$	0.949	0.949	0.957	0.957
$R=7, \ell=2$[a]	0.951	0.953	0.959	0.959
$R=7, \ell=3$	0.951	0.946	0.960	0.954
$R=7, \ell=4$	0.947	0.951	0.956	0.958
$R=7, \ell=5$	0.949	0.955	0.956	0.961
$R=7, \ell=6$	0.945	0.956	0.955	0.965
$R=7, \ell=7$	0.946	0.954	0.954	0.961
$R=2, \ell=2$	0.985	0.977	0.990	0.983
$R=3, \ell=2$	0.979	0.968	0.985	0.974
$R=4, \ell=2$	0.972	0.966	0.979	0.973
$R=5, \ell=2$	0.965	0.959	0.972	0.965
$R=6, \ell=2$	0.957	0.950	0.966	0.958
$R=3, \ell=3$	0.978	0.976	0.985	0.981
$R=4, \ell=4$	0.969	0.971	0.976	0.976
$R=5, \ell=5$	0.962	0.966	0.969	0.973
$R=6, \ell=6$	0.953	0.964	0.961	0.969
$\mu_1=2, \mu_2=2$	0.961	0.961	0.956	0.956
$\mu_1=4, \mu_2=4$	0.955	0.958	0.952	0.952
$\mu_1=2, \mu_2=6$	0.955	0.957	0.953	0.955
$u=0$	0.950	0.952	0.961	0.960
$u=20$	0.953	0.955	0.949	0.955
$h_1=2$	0.933	0.950	0.956	0.961
$h_1=3$	0.932	0.944	0.960	0.966
$c=20$	0.959	0.953	0.966	0.959
$c=40$	0.938	0.941	0.949	0.948

[a]Base test case.

7. Discussion and conclusion

Brick-and-mortar stores are increasingly adopting online shopping channels in their portfolio, where in-store inventory is often used to fulfil both in-store demand as well as online orders. As the retailer serves demand of multiple channels by one inventory, this study proposes to ration the inventory of the retailer among the channels. The rationing is suggested as solution to the negative effects that occur when using in-store inventory for the fulfilment of online demand.

The rationing is characterised by the trade-off of serving in-store or online customers, where the profit of selling a product in-store is higher than through the online channel but the inventory cost of the online channel is lower. The proposed model encompasses this core trade-off experienced by an omni-channel retailer. The underlying cost structure and demand of both channels influence this decision. The products are present at the same location, thus the retailer can make the rationing decision on a daily basis.

This paper contributes to the academic literature on omni-channel retailing by providing a novel exact method as well as heuristics for the integration of the replenishment and rationing decision. Previous research on dynamic rationing assumes zero lead times, no lost sales, and static order policies which does not capture the characteristics of an omni-channel retailer. To capture these characteristics, we study an omni-retailer with a dynamic order policy, a deterministic lead time, and lost sales. The problem is formulated as an MDP that was solved through value iteration to get an exact solution. The structure of the optimal rationing and ordering policies are numerically investigated and captured in heuristics.

Based on our model and numerical results, we conclude that the optimal ordering policy consists of a fixed order quantity and an order-up-to level. At low inventory levels, a fixed order quantity is optimal, whereas at high inventory levels an order-up-to level is optimal, this is due to the uncertainty of demand during the procurement lead time. For the rationing policy, a maximum threshold level for high inventory levels is found, which is also proven analytically. At low inventory levels there is a trade-off on whether to store products in-store or offline. Based on these insights, heuristics are developed and tested through extensive numerical tests. The heuristics are compared with the optimal policy on its goodness of fit, optimality gap and service level. The heuristics show near-optimal results, concluding that the heuristics can be effective in encompassing the complexity of dynamic rationing. Furthermore, the heuristics is able to closely approximate the optimal ordering policy by mixing a base stock policy and a constant order policy.

Even though the model studied in this paper captures the core trade-offs presented in omni-channel retailing, several potentially complicating factors were not taken into account, and could be directions for further research. More specifically, it could be in the exploration of a network of multiple omni-channel retailer locations who can fulfil online demand, extending the rationing decision with the allocation of online orders to different locations. This research contributes to such a direction by offering heuristics that are fast and have the potential to be upgraded to these more complex problem settings. However, it should be noted that use of our heuristics in practice would require the retailer to regularly update demand parameters throughout the selling period as these might change over time. In our paper we assumed lost sales as replenishment happens only weekly and consumers may be impatient. By adopting the lost sales assumption, we prevent stock outs as much as possible as they would result in lost profit. A mixture of lost sales and backordering could be incorporated in the framework we present, e.g. if a percentage of online customers might be willing to wait for the fulfilment of online orders. Adding a backlog decision will further increase the computational complexity. Based on previous research (e.g. Benjaafar et al., 2010; Fadıloğlu and Bulut, 2010), we expect the amount of backorders to have a limit and that the decision of when to backorder will dependent on the state of the system. Additional research could be interesting to investigate the influence of product returns in the current context. Online sales have become challenging for retailers due to the large quantity of sold goods being returned. This return flow influences the inventory of the retailer, thus influencing the rationing and ordering policies. The addition of returns would however significantly increase the dimensionality of the MDP, as the state of the model would have to include the numbers of sold products (in the last days or weeks) that will be potentially returned in the coming days or weeks. The curse of dimensionality would result in unsolvable MDPs for most relevant problem sizes. The development of effective heuristic solutions would then also be highly relevant for this problem.

CRediT authorship contribution statement

Joost Goedhart: Conceptualization, Methodology, Software, Investigation, Writing – original draft. **René Haijema:** Conceptualization, Methodology, Writing – review & editing, Supervision. **Renzo Akkerman:** Conceptualization, Methodology, Writing – review & editing, Supervision.

Appendix. Derivation of inventory base stock level

We can prove that there is an optimal quantity to be stored in-store in case of excessive stock. We have the following equation to calculate the expected contribution in state s_t when taking action a:

$$\begin{aligned}\mathbb{E}C\left(s_t,a\right) = {} & p\left(\sum_{d<a} d\cdot P_1(d) + \sum_{d\geq a} a\cdot P_1(d)\right) \\ & + (p-u)\left(\sum_{d<I_t-a} d\cdot P_2(d) + \sum_{d\geq I_t-a}(I_t-a)\cdot P_2(d)\right) \\ & - \left(h_1\cdot a + h_2\cdot\left(I_t-a\right)\right)\end{aligned}$$

We add the following term:

$$p\sum_{d\geq a+1} d\cdot P_1(d) - p\sum_{d\geq a+1} d\cdot P_1(d) + (p-u)\sum_{d\geq I_t-a+1} d\cdot P_2(d) - (p-u)\sum_{d\geq I_t-a+1} d\cdot P_2(d)$$

We know that the following two conditions will always hold

$$\sum_{d<a} d\cdot P_1(d) + \sum_{d\geq a+1} d\cdot P_1(d) = \mu_1 \quad \text{and}$$

$$\sum_{d<a} d\cdot P_2(d) + \sum_{d\geq a+1} d\cdot P_2(d) = \mu_2$$

We can therefore get to the following equation after adding and subtracting:

$$\begin{aligned}\mathbb{E}C\left(s_t,a\right) = {} & p\cdot\mu_1 + p\sum_{d\geq a+1}(a-d)\cdot P_1(d) + (p-u)\cdot\mu_2 \\ & + (p-u)\sum_{d\geq I_t-a+1}\left(I_t-a-d\right)\cdot P_2(d) \\ & - \left(h_1\cdot a + h_2\cdot\left(I_t-a\right)\right)\end{aligned}$$

To find the optimal action a that maximises the expected contribution $\mathbb{E}C\left(s_t,a\right)$, we take the approximate derivative:

$$\Delta\mathbb{E}C\left(s_t,a\right) = \mathbb{E}C\left(s_t,a+1\right) - \mathbb{E}C\left(s_t,a\right)$$

After subtracting and rewriting we get the following equation:

$$\Delta\mathbb{E}C\left(s_t,a\right) = p\cdot\sum_{d\geq a+1} -P_1(d) + (p-u)\cdot\sum_{d\geq I_t-a+1} P_2(d) + \left(h_1-h_2\right)$$

We know that:

$$\sum_{d\geq a+1} P_1(d) = 1 - \sum_{d\leq a} P_1(d) \quad \text{and}$$

$$\sum_{d\geq I_t-a+1} P_2(d) = 1 - \sum_{d\leq I_t-a} P_2(d)$$

We can therefore rewrite the approximate derivative as follows:

$$\Delta\mathbb{E}C\left(s_t,a\right) = -p + p\cdot\sum_{d\leq a} P_1(d) + (p-u) - (p-u)\cdot\sum_{d\leq I_t-a} P_2(d) + \left(h_1-h_2\right)$$

If all possible demand in the online channel is satisfied due to excess inventory being available to this channel, we have $\sum_{d\leq I_t-a} P_2(d) =$

1. This assumes $h_1 > h_2$, as otherwise the excess inventory would be allocated to the offline channel. In this situation, the approximate derivative is reduced to the following.

$$\Delta\mathbb{E}C\left(s_t, a\right) = -p + p \cdot \sum_{d \le a} P_1\left(d\right) + \left(h_1 - h_2\right)$$

In order to find the maximum value of $\mathbb{E}C\left(s_t, a\right)$, we find the largest value for a for which $\Delta\mathbb{E}C\left(s_t, a\right)$ is negative:

$$\Delta\mathbb{E}C\left(s_t, a\right) \le 0 \implies -p + p \cdot \sum_{d \le a} P_1\left(d\right) + \left(h_1 - h_2\right) \le 0$$

$$\implies \sum_{d \le a} P_1\left(d\right) \le \frac{p - \left(h_1 - h_2\right)}{p}$$

Thus we get the following equation that gives us the optimum amount of products to hold in-store:

$$P_1(d \le a) \le \frac{p - \left(h_1 - h_2\right)}{p}$$

The inventory base stock level a can then be easily derived from the inverse probability function of P_1.

References

Aastrup, J., Kotzab, H., 2010. Forty years of out-of-stock research–and shelves are still empty. Int. Rev. Retail Distrib. Consum. Res. 20 (1), 147–164.

Agatz, N.A., Fleischmann, M., Van Nunen, J.A., 2008. E-fulfillment and multi-channel distribution - A review. European J. Oper. Res. 187 (2), 339–356.

Aksen, D., Altinkemer, K., 2008. A location-routing problem for the conversion to the "click-and-mortar" retailing: The static case. European J. Oper. Res. 186 (2), 554–575.

Alfieri, A., Pastore, E., Zotteri, G., 2017. Dynamic inventory rationing: How to allocate stock according to managerial priorities. An empirical study. Int. J. Prod. Econ. 189, 14–29.

Baird, N., Kilcourse, B., 2011. Omni-Channel Fulfillment and the Future of the Retail Supply Chain. Benchmark Report, Retail Systems Research.

Bao, L., Liu, Z., Yu, Y., Zhang, W., 2018. On the decomposition property for a dynamic inventory rationing problem with multiple demand classes and backorder. European J. Oper. Res. 265 (1), 99–106.

Bayram, A., Cesaret, B., 2021. Order fulfillment policies for ship-from-store implementation in omni-channel retailing. European J. Oper. Res. 294 (3), 987–1002.

Bendoly, E., 2004. Integrated inventory pooling for firms servicing both on-line and store demand. Comput. Oper. Res. 31 (9), 1465–1480.

Bendoly, E., Blocher, D., Bretthauer, K.M., Venkataramanan, M., 2007. Service and cost benefits through clicks-and-mortar integration: Implications for the centralization/decentralization debate. European J. Oper. Res. 180 (1), 426–442.

Benjaafar, S., ElHafsi, M., Huang, T., 2010. Optimal control of a production-inventory system with both backorders and lost sales. Nav. Res. Logist. 57 (3), 252–265.

Bretthauer, K.M., Mahar, S., Venakataramanan, M., 2010. Inventory and distribution strategies for retail/e-tail organizations. Comput. Ind. Eng. 58 (1), 119–132.

Carr, S., Duenyas, I., 2000. Optimal admission control and sequencing in a make-to-stock/make-to-order production system. Oper. Res. 48 (5), 709–720.

Chen, H., Zhou, Z., Liu, M., 2012. A queuing-inventory system with two classes demand and subject to selected service. J. Inf. Comput. Sci. 9 (11), 3081–3089.

Chew, E.P., Lee, L.H., Liu, S., 2013. Dynamic rationing and ordering policies for multiple demand classes. OR Spectrum 35 (1), 127–151.

Cohen, A.C., 1954. Estimation of the Poisson parameter from truncated samples and from censored samples. J. Amer. Statist. Assoc. 49 (265), 158–168.

Deshpande, V., Cohen, M.A., Donohue, K., 2003. A threshold inventory rationing policy for service-differentiated demand classes. Manage. Sci. 49 (6), 683–703.

Difrancesco, R.M., van Schilt, I.M., Winkenbach, M., 2021. Optimal in-store fulfillment policies for online orders in an omni-channel retail environment. European J. Oper. Res. 293 (3), 1058–1076.

ENC, 2016. What Every Omni-Channel Retailer Should Know About Developing a Ship From Store Strategy. enVista Company.

Fadıloğlu, M.M., Bulut, O., 2010. A dynamic rationing policy for continuous-review inventory systems. European J. Oper. Res. 202 (3), 675–685.

Fernie, J., Sparks, L., 2004. Logistics and Retail Management: Insights into Current Practice and Trends from Leading Experts. Kogan Page Publishers.

Frank, K.C., Zhang, R.Q., Duenyas, I., 2003. Optimal policies for inventory systems with priority demand classes. Oper. Res. 51 (6), 993–1002.

Gayon, J.P., De Vericourt, F., Karaesmen, F., 2009. Stock rationing in an M/E r/1 multi-class make-to-stock queue with backorders. IIE Trans. 41 (12), 1096–1109.

Ha, A.Y., 1997a. Inventory rationing in a make-to-stock production system with several demand classes and lost sales. Manage. Sci. 43 (8), 1093–1103.

Ha, A.Y., 1997b. Stock-rationing policy for a make-to-stock production system with two priority classes and backordering. Nav. Res. Logist. 44 (5), 457–472.

Haijema, R., van der Wal, J., van Dijk, N.M., 2007. Blood platelet production: Optimization by dynamic programming and simulation. Comput. Oper. Res. 34 (3), 760–779.

Haynsworth, H.C., Price, B.A., 1989. A model for use in the rationing of inventory during lead time. Nav. Res. Logist. 36 (4), 491–506.

Hobkirk, I., 2015. Key distribution strategies of top omni-channel retailers. Logist. Manage. 55 (1), 48–53.

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 (6/7), 562–583.

Hung, H.C., Chew, E.P., Lee, L.H., Liu, S., 2012. Dynamic inventory rationing for systems with multiple demand classes and general demand processes. Int. J. Prod. Econ. 139 (1), 351–358.

Hung, Y.F., Hsiao, J.Y., 2013. Inventory rationing decision models during replenishment lead time. Int. J. Prod. Econ. 144 (1), 290–300.

Ishfaq, R., Bajwa, N., 2019. Profitability of online order fulfillment in multi-channel retailing. European J. Oper. Res. 272 (3), 1028–1040.

Jalilipour Alishah, E., Moinzadeh, K., Zhou, Y.-P., 2015. Inventory Fulfillment Strategies for an Omni-Channel Retailer. Working Paper, University of Washington, Seattle.

Kaplan, A., 1969. Stock rationing. Manage. Sci. 15 (5), 260–267.

Kristensen, A.R., 1988. Hierarchic Markov processes and their applications in replacement models. European J. Oper. Res. 35 (2), 207–215.

Li, T., Zhao, X., Xie, J., 2015. Inventory management for dual sales channels with inventory-level-dependent demand. J. Oper. Res. Soc. 66 (3), 488–499.

Liu, S., Song, M., Tan, K.C., Zhang, C., 2015. Multi-class dynamic inventory rationing with stochastic demands and backordering. European J. Oper. Res. 244 (1), 153–163.

Liu, S., Zhang, C., 2015. An efficient inventory rationing strategy for systems with two demand classes and backordering. Int. J. Prod. Res. 53 (20), 6136–6142.

Ma, S., Jemai, Z., 2019. Inventory rationing for the News-Vendor problem with a drop-shipping option. Appl. Math. Model. 71, 438–451.

Mahar, S., Bretthauer, K.M., Venkataramanan, M., 2009. The value of virtual pooling in dual sales channel supply chains. European J. Oper. Res. 192 (2), 561–575.

Mahar, S., Salzarulo, P.A., Wright, P.D., 2012. Using online pickup site inclusion policies to manage demand in retail/E-tail organizations. Comput. Oper. Res. 39 (5), 991–999.

Mahar, S., Wright, P.D., 2009. The value of postponing online fulfillment decisions in multi-channel retail/e-tail organizations. Comput. Oper. Res. 36 (11), 3061–3072.

Melchiors, P., 2001. Rationing Policies for an Inventory Model with Several Demand Classes and Stochastic Lead Times. Working Paper, Department of Operations Research, University of Aarhus.

Melchiors, P., 2003. Restricted time-remembering policies for the inventory rationing problem. Int. J. Prod. Econ. 81–82, 461–468.

Melchiors, P., Dekker, R., Kleijn, M.J., 2000. Inventory rationing in an (s, Q) inventory model with lost sales and two demand classes. J. Oper. Res. Soc. 51 (1), 111–122.

Mou, S., Robb, D.J., DeHoratius, N., 2018. Retail store operations: Literature review and research directions. European J. Oper. Res. 265 (2), 399–422.

Nguyen, D.H., de Leeuw, S., Dullaert, W.E., 2018. Consumer behaviour and order fulfilment in online retailing: a systematic review. Int. J. Manage. Rev. 20 (2), 255–276.

Ovezmyradov, B., Kurata, H., 2019. Effects of customer response to fashion product stockout on holding costs, order sizes, and profitability in omnichannel retailing. Int. Trans. Oper. Res. 26 (1), 200–222.

Teunter, R.H., Klein Haneveld, W.K., 2008. Dynamic inventory rationing strategies for inventory systems with two demand classes, Poisson demand and backordering. European J. Oper. Res. 190 (1), 156–178.

Topkis, D.M., 1968. Optimal ordering and rationing policies in a nonstationary dynamic inventory model with N demand classes. Manage. Sci. 15 (3), 160–176.

Turgay, Z., Karaesmen, F., Örmeci, E.L., 2015. A dynamic inventory rationing problem with uncertain demand and production rates. Ann. Oper. Res. 231 (1), 207–228.

Verhoef, P.C., Kannan, P.K., Inman, J.J., 2015. From multi-channel retailing to omni-channel retailing: introduction to the special issue on multi-channel retailing. J. Retail. 91 (2), 174–181.

Wang, K., Jiang, Z., Li, G., Cao, J., 2013b. Integrated capacity allocation policies for a production service system with two-class customers and items. In: 2013 IEEE International Conference on Automation Science and Engineering (CASE). IEEE, pp. 83–88.

Wang, D., Tang, O., 2014. Dynamic inventory rationing with mixed backorders and lost sales. Int. J. Prod. Econ. 149, 56–67.

Wang, D., Tang, O., Huo, J., 2013a. A heuristic for rationing inventory in two demand classes with backlog costs and a service constraint. Comput. Oper. Res. 40 (12), 2826–2835.

Xu, J., Cao, L., 2019. Optimal in-store inventory policy for omnichannel retailers in franchising networks. Int. J. Retail Distrib. Manage. 47 (12), 1251–1265.

Zhao, N., Lian, Z., 2011. A queueing-inventory system with two classes of customers. Int. J. Prod. Econ. 129 (1), 225–231.