TRANSPORTATION SCIENCE

Stochastic inventory routing for perishable products

Journal:	*Transportation Science*
Manuscript ID	TS-2016-0163.R1
Manuscript Type:	Resubmitted Manuscripts
Date Submitted by the Author:	n/a
Complete List of Authors:	Crama, Yves; University of Liege, HEC-Management School Rezaei, Mahmood; University of Liege, HEC-Management School Savelsbergh, Martin; Georgia Institute of Technology, Industrial and Systems Engineering Van Woensel, Tom; Eindhoven University of Technology, Industrial Engineering
Keywords/Area of Expertise:	inventory routing problem, perishable inventory system, stochastic inventory routing

SCHOLARONE™
Manuscripts

Stochastic Inventory Routing for Perishable Products

Yves Crama

HEC Liège, University of Liège, 4000 Liège, Belgium, y.crama@ulg.ac.be

Mahmood Rezaei

HEC Liège, University of Liège, 4000 Liège, Belgium, m.rezaei@ulg.ac.be

Martin Savelsbergh

School of Industrial and Systems Engineering, Georgia Institute of Technology Atlanta, GA 30332-0205, USA,
martin.savelsbergh@isye.gatech.edu

Tom Van Woensel

School of Industrial Engineering, Eindhoven University of Technology, 5612AZ Eindhoven, Netherlands, t.v.woensel@tue.nl

Different solution methods are developed to solve an inventory routing problem for a perishable product with stochastic demands. The solution methods are compared empirically in terms of average profit, service level, and actual freshness. The benefits of explicitly considering demand uncertainty are quantified. The computational study highlights that in certain situations a simple ordering policy can already reach a very good performance, but that statistically and economically significant improvements are achieved when using more advanced solution methods. Managerial insights concerning the impact of shelf life and store capacity on profit are also obtained.

Key words: inventory routing problem; perishable inventory system; perishable inventory routing; stochastic inventory routing.

History:

1.   Introduction

Consider a retail chain whose main goal is to optimize the long-term profit of a distribution network where products are shipped from a central warehouse (depot) to several stores. The decisions to be made are: (1) how much to deliver to each store in each period, and (2) which delivery routes to use. If customer demands at the stores are deterministic and known for the entire planning horizon, then this decision problem is known as the Inventory Routing Problem (IRP).

In this paper, we consider the IRP for a (single) perishable product, e.g., a dairy product, flowers, fruits, or vegetables, with stochastic customer demands. We develop and compare four solution methods: (1) an expected-value method in which stochastic demands are replaced by their expected values, (2) a deliver-up-to-level policy with a high target service level, (3) a decomposition method that relies on independent inventory control models for each store and on estimates of the routing

2

costs that can be attributed to each store, (4) a decomposition-integration method which improves the solution obtained by the decomposition method by analyzing and reducing the routing costs.

1.1. Motivation

In an IRP, both the delivery quantities at a store and the delivery routes to serve the stores are determined by a centralized decision-maker for a given planning horizon. IRPs are very difficult to solve to optimality even when the distribution network is small, far smaller than those typically encountered in practice. Coelho and Laporte (2014) report that instances of a single-product IRP with deterministic demands can rarely be solved to optimality when the number of stores exceeds 30. In contrast, a typical retail chain in Belgium (a small country) dispatches over 18,000 (perishable and dry) products from a central warehouse to more than 800 stores.

In most real-life IRP settings, demand experienced at stores is uncertain, giving rise to the more complex Stochastic IRP (SIRP). In such a setting, information regarding inventory levels at stores is periodically transmitted to the central warehouse, where a central decision-maker uses this information, and any available information about anticipated future demand, to determine the delivery quantities and the delivery routes for the next period, or the next few periods. When the actual demand from the customers at the stores is observed and updated inventory levels are transmitted to the warehouse, new next-period (or short-term) decisions are made. Due to the complexity of such a distribution system, retail chains frequently use a two-step decision-making process, in which each store deploys its own inventory management system to place orders (ignoring any impact these order placements may have on routing costs), and the centralized decision-maker uses a vehicle routing model to determine low-cost delivery routes to serve the placed orders. Such a two-step decision-making process does not necessarily yield an optimal profit for the retail chain, but offers a pragmatic approach for managing this complex system.

In many real-life IRP settings, product shelf life is not unlimited. Perishable products constitute over 52% of sales revenue of grocery retail chains (Chao et al. 2015), and roughly 10% of product is wasted before being sold (Kouki et al. 2015), while the food retail profit margin hardly exceeds 2%; see EBRD–FAO (2009), NAICS (2012), FMI (2014). Therefore, the profitability highly depends on efficient and effective inventory routing policies. The delivery frequency plays a major role in determining the profit, the service level, and the freshness. Indeed, infrequent deliveries of large delivery quantities reduces the routing costs, but also brings loss of freshness and a commensurate reduction in customer satisfaction (and possibly even lost sales due to the stochasticity of demand). Conversely, smaller quantities delivered more frequently improves freshness of products (and, consequently, ensures customer satisfaction) and reduces the risk of lost sales, but does imply higher routing costs.

In this paper, we study the stochastic inventory routing problem for a single perishable product (PSIRP) in which we seek to maximize profit subject to a given service level requirement. Even though freshness is considered and evaluated, it is not directly taken into account in the optimization.

1.2. Problem statement

Consider a generic retail chain that aims to maximize the expected net profit generated by the sales of a single perishable product. The net profit is measured by deducting acquisition, distribution, and other miscellaneous costs from the total revenue. Acquisition costs and revenue mostly depend on the quantities delivered to the stores, whereas distribution costs are also a consequence of the way the vehicles are dispatched. Miscellaneous costs are viewed as independent of either decision and are regarded as constant costs. As a result, we do not consider them in the objective function: the net profit is simply computed as (Revenue – Acquisition costs – Distribution costs), and its expected value defines the objective function.

Assume a finite planning horizon consisting of T consecutive periods. We consider an implicit complete graph $G = (V, A)$, whose vertices represent the depot and the stores, and arcs represent the road segments between pairs of vertices. The travel cost from vertex i to vertex j is denoted as c_{ij}. Product is picked up from the depot and delivered to the stores. Each route starts at the depot, ends at the depot, and cannot exceed a predefined length. The demand in period t from end customers at a store i is an integer random variable D_{ti} (assuming independence for all periods and all stores) with a known probability distribution. We define L to be the deterministic shelf life of the product from the arrival at the store. The acquisition cost of a unit of product is a. All units delivered in period t have the same selling price p during L periods. Unsold units perish at the end of period $t + L - 1$ with no salvage value. Unmet demand leads to lost sales but does not generate any other cost. The inventory holding cost is zero. We assume that the depot has an unlimited supply of product. The capacity of store i is denoted by C_i. The retail chain deploys a sufficiently large fleet of vehicles to make deliveries, each with capacity Q. Each vehicle incurs a fixed cost K per period when it is used, and a variable cost equal to c_{ij} when traveling from vertex i to vertex j. Each store is served by at most one vehicle in each period (no split deliveries).

The retail chain uses a centralized decision-making system to determine delivery quantities and routes in each period. The inventory state in store i at the beginning of period t is denoted by $X_{ti} = (x_1, x_2, \ldots, x_{L-1})_{ti}$, where $(x_k)_{ti}$ is the inventory level of product with remaining shelf life k. At the beginning of each period t, based on the inventory states, the retail chain decides about the delivery quantities, y_{ti}, and the delivery routes, R_t, to be used in the current period. There is no time window for the delivery to stores, but in each period each vehicle performs at most one route with a predefined maximum length. The delivery lead time is zero, i.e., the delivery quantities y_{ti}

4

are available on the shelves at the beginning of period t, right after the decision is made. The real demand for period t is observed during the period and after quantity y_{ti} has been delivered. We assume that the oldest units of product are sold first (FIFO issuing policy), i.e., $(x_k)_{ti}$ is stored until $(x_{k-1})_{ti}$ is used up or perished.

A predefined Target Service Level (TSL) must be respected in every period and in every store. More precisely, the total inventory available in store i in period t must be such that the probability of not incurring a stockout in period t is at least equal to TSL. In practice, the average service level of perishables is estimated to be around 92% in Europe and the USA (Minner and Transchel 2010).

All notation is summarized in Table 1.

Table 1   Indices, parameters, and decision variables

Indices:	
i, j	indices for vertices (depot and stores)
k	index for remaining shelf life
r	index for routes
t	index for periods
Parameters:	
T	length of the planning horizon
N	number of stores
L	shelf life of the product
TSL	target service level to be respected in every period and in each store
a	acquisition price of each unit of the product
p	selling price of each unit of the product
C_i	capacity of store i
$(x_k)_{ti}$	inventory level with remaining shelf life k in period t in store i before delivery
$(x_1, x_2, \ldots, x_{L-1})_{ti}$	state of the system in period t in store i before delivery
I_{ti}	total inventory level in period t in store i before delivery
D_{ti}	random demand of end customers in period t in store i (integer-valued)
$Pr(D_{ti} = d)$	probability function of demand in period t in store i
Q	capacity of each vehicle
c_{ij}	distance and travel cost from vertex i to vertex j
K	fixed cost of using a vehicle
F_{ti}	estimated cost-to-serve assigned in period t to store i
Decision variables:	
y_{ti}	delivery quantity in period t to store i (integer values)
R_t	set of routes used in period t (index r)
π	expected profit generated by all stores over the planning horizon

1.3.   Scientific contributions

As outlined above, we investigate a stochastic inventory routing problem for a single perishable product. Our main contributions are:

- The development of four different solution methods for the PSIRP. The methods differ in sophistication and emphasis placed on features of the problem, e.g., perishability of the product, stochasticity of the demand, or target service level. The decomposition method and the decomposition-integration method, presented in Sections 5–7, rely on an original combination of stochastic dynamic programming and combinatorial optimization techniques.

- The analysis of the results of an extensive computational study, which establishes the relative performance of the solution methods (and the influence of their parameter settings) for different classes of instances and on different metrics of interest, e.g., total profit, average freshness, and simplicity of implementation. Among others, we demonstrate that a simple, easy-to-implement replenishment policy, derived from one of the more sophisticated solution methods, is highly effective in a variety of settings.

- The creation of managerial insights related to the impact of store capacity and shelf life on the expected profit.

The remainder of the paper is organized as follows. In Section 2, an extensive literature review is presented. Sections 3–6 introduce the four solution methods, two elementary methods and two sophisticated methods. In Section 7, we develop a matheuristic to solve an optimization subproblem arising in our most sophisticated solution method. Section 8 presents a heuristic algorithm for the case in which full information is available regarding future demand is available (which will be used for comparison purposes in the computation study). Section 9 discusses the setup of our comprehensive computational study. Results of the computational experiments are presented and analyzed in Section 10. Finally, concluding remarks are given in Section 11.

2.   Literature review

Inventory-routing models naturally relate to various management practices, and in particular, to vendor-managed inventory (VMI) and to retailer-managed inventory (RMI) systems. The VMI approach relies on cooperation and on information sharing between a supplier and its customers. When VMI is implemented, the supplier takes over the responsibility of managing the inventory of the customers by deciding on replenishment quantities and delivery periods. The consequences can be beneficial for both parties: customers need to employ less resources for controlling their inventory, and the supplier has more flexibility for integrating the replenishment quantities and periods to different customers (Desaulniers et al. 2016). In contrast, in an RMI system, the customers decide when and how much to order, independently of each other, so that the ability of the supplier to optimize its transportation costs is strongly restricted by the customers' decisions (Archetti and Speranza 2016, Bertazzi and Speranza 2012).

6

2.1. Inventory-routing

The (classical) IRP deals with deterministic demands and is concerned with the distribution of a single product from a single depot to a set of customers with deterministic demands over a given planning horizon. The objective is to minimize the distribution and inventory costs during the planning period without causing stock-outs at any of the customers. As in a VMI system, the main decisions in an IRP are: (a) when to serve each customer, (b) how much to deliver to a customer when it is visited, and (c) which routes to use. The IRP has a wide range of applications including the distribution of gas (Campbell and Savelsbergh 2004a, Gronhaug et al. 2010), fuel (Popović et al. 2012), automobile components (Alegre et al. 2007, Stacey et al. 2007), perishable products (Federgruen and Zipkin 1984, Federgruen et al. 1986), groceries products (Gaur and Ficher 2004), cement (Christiansen et al. 2011), and blood products (Hemmelmayr et al. 2009).

Since the IRP has the flexibility to decide how much to deliver to each customer and which routes to use in each period, the decision space becomes enormously large even when compared, for instance, with classical vehicle routing problems (VRP). Exact approaches are typically based on mixed integer programming (MIP) formulations using arc-flow decision variables, although route-based formulations have also been used (Desaulniers et al. 2016). But finding the optimal solution of such models is quite challenging, even for very small instances of the IRP (Campbell et al. 1998). Therefore, many early algorithms proposed in the literature decompose the IRP into two stages: (a) inventory control (determining the delivery amounts), and (b) vehicle routing (Campbell and Savelsbergh 2004b, Federgruen and Zipkin 1984, Qu et al. 1999). In some cases, an overall solution is found by iterating between these two problems (Federgruen and Zipkin 1984, Qu et al. 1999). Exact algorithms for the IRP are more recent. They include the branch-and-cut algorithms of Coelho and Laporte (2013a), Coelho and Laporte (2013b), and Coelho and Laporte (2014d), and the branch-and-price-and-cut algorithm of Desaulniers et al. (2016).

Researchers have also striven to develop solution methods for simplified versions of IRP models, rather than the classical IRP model. Examples of simplifying assumptions include:

- deliver-up-to-level (UL) replenishment policies; Bertazzi et al. (2002), Archetti et al. (2007), Solyali and Süral (2011);

- direct delivery policies; Bertazzi and Speranza (2012);

- deliveries occurring only when inventory levels are down to zero; Chan et al. (1998) and Jaillet et al. (2002);

- single vehicle; see, e.g., Archetti et al. (2007) and Solyali and Süral (2011);

- constant demand rate over time; Raa and Aghezaaf (2008, 2009), Ekici et al. (2015);

- periodic deliveries; Bertazzi and Speranza (2012), Campbell and Wilson (2014).

Andersson et al. (2010), Bertazzi and Speranza (2011, 2012, 2013) and Coelho et al. (2014) provide excellent reviews on IRPs, both from the application and the methodological points of view.

2.2. Stochastic inventory-routing

In the stochastic inventory routing-problem (SIRP), future customer demands are uncertain, and are given by their probability distributions. While the majority of papers on the SIRP assume that demands are fully realized at the end of each period, some models assume that demands are realized upon arrival of the vehicle at a customer (Berman and Larson 2001, Huang and Lin 2010).

Demand stochasticity implies that shortages may occur, and there is often a positive probability that a customer runs out of stock. To restrict shortages, a penalty is imposed whenever a customer runs out of stock, and this penalty is usually taken to be proportional to the unsatisfied demand. Unsatisfied demand is typically considered to be lost sale (Minkoff 1993, Kleywegt et al. 2004), and is rarely dealt with as backlogging, as in the work by Yu et al. (2012). In either case, penalties may apply. A pre-defined service level may apply too, which imposes a minimum inventory level at each customer in each period (Yu et al. 2012). The objective is to choose a delivery policy that minimizes the expected total (inventory, distribution and penalty) costs per period. Due to the complexity of the SIRP, simplifying assumptions are frequently made, as in the IRP. These assumptions may include considering a single capacitated vehicle (Coelho and Laporte 2014, Reinman et al. 1999, Schwartz et al. 2006), a single uncapacitated vehicle (Qu et al. 1999), or direct deliveries (Barnes-Schuster and Bassok 1997, Kleywegt et al. 2002, Reinman et al. 1999).

Markov decision processes (MDP) can be used to model SIRPs over an infinite planning horizon. MDP models formulate a value function which depends on inventory levels. When the demand probability distribution is stationary, a deterministic optimal policy can be calculated for each state, and the value function can be optimized by standard techniques such as policy iteration, value iteration, or successive approximation. These algorithms are practical only if the state space is small and if the optimization problem can be solved efficiently. None of these requirements are satisfied by real-world instances of the SIRP, as the state space is usually extremely large, even if inventories are discretized, and the optimization problem includes a vehicle routing problem as a special case (Campbell et al. 1998). Due to the curse of dimensionality, researchers often use approximations of the value function (Minkoff 1993, Adelman 2004) or decompose it (Kleywegt et al. 2004). Based on the LP model proposed by Puterman (1994), Adelman (2004) formulates and interprets two primal-dual approximations. Such approximations relax the feasible region of the dual problem, and so provide upper bounds on the original dual maximization problem. Kleywegt et al. (2004) formulate an MDP model of the SIRP and propose approximation methods to find good solutions within reasonable time. This is the extension of an earlier paper (Kleywegt et al. 2002) in which the authors formulated an SIRP with direct deliveries as an MDP and proposed an approximate dynamic programming approach for its solution.

8

Solution methods other than MDP have also been used for solving SIRPs. In Jaillet et al. (2002), for instance, long-term delivery costs are incorporated into shorter planning horizons. In Yu et al. (2012), rather than dealing with an exact stochastic model, an approximate SIRP model is proposed and transformed into a simplified deterministic one. Then, Lagrangian relaxation is used to decompose the model into an inventory problem and a vehicle routing problem.

2.3. Inventory-routing for perishables

The IRP for perishables (PIRP) is identical to the classical IRP, except that products have a limited shelf life after which they lose their value. In the deterministic case, shortages are not allowed. Moreover, thanks to the complete knowledge of demand, no products need to deteriorate. This implies that the cost components in the PIRP are the same as for non-perishables, i.e, total inventory holding costs and routing costs over the planning horizon. The main difference with the classical IRP is that the deliveries to customers are now restricted by the maximum shelf life of the product. As a consequence, delivery frequency plays an important role. Less frequent deliveries reduce the routing costs, but result in more units with short remaining shelf lives in the following periods; these units are subject to holding cost and to deterioration, and their sales may be reversely affected by their lack of freshness. Therefore, finding a right trade-off between costs and freshness is crucial. The main objective in most applications is to minimize costs (or to maximize profit), while freshness is controlled by imposing additional side constraints on delivery quantities. Hemmelmayr et al. (2009) investigate the delivery of blood products to hospitals, where the tour length is restricted, but vehicle capacity is ignored in view of the small size of blood bags, and no inventory cost is imposed. The objective is to minimize travel costs over a finite horizon. The authors develop and evaluate two delivery strategies, namely: (1) using fixed routes but deciding about delivery days, and (2) repeating delivery patterns for each hospital. Coelho and Laporte (2014) consider age-dependent holding costs and selling prices in a PIRP, where the supplier has the choice to deliver either fresh or aged products, and each case yields different holding costs and revenues. The objective function maximizes the total sales revenue, minus inventory and routing costs. Le et al. (2013) propose a mathematical model based on feasible routes that start from the depot, visit a subset of stores at most one time, and then return to the depot, without the necessity of respecting the vehicle capacity. The objective function represents the sum of transportation costs and inventory costs. They use a column-generation based algorithm to solve the problem. Their work is extended by Al Shamsi et al. (2014) to include the cost of CO_2 emissions, based on the vehicle load and distance. The resultant model is an MINLP problem which is solved using a commercial solver. In Mirzaei and Seifi (2015), dependency of the end-customers demand on the age of the inventory is formulated as a PIRP so that a portion of the demand is considered as lost sale if inventory is not as fresh as it could be. The objective function is the total cost of transportation, lost sale, and holding inventories. The authors develop a hybrid simulated annealing and tabu search algorithm for solving the problem.

2.4. Inventory control of perishables in an RMI system

Recall that in an RMI system, the stores decide when and how much to order, independently of each other. Therefore, the main decision variables in an RMI perishable inventory system are the order time and the order quantity. To place an order, the current inventory level and age of the stocked products (state of the system) are observed. In most problem settings, an MDP provides an exact solution approach. However, the computation of the optimal order for every state of the system using classical techniques is in general intractable because of the curse of dimensionality. Thus, many researchers turned to effective heuristic policies to handle these problems, mostly for the case of independent and identically distributed demands (Chao et al. 2015). The most widely used periodic-review ordering policies are (R, S) (Chiu 1995, Cooper 2001, Deniz et al. 2010) and (R, s, S) (Broekmeulen and Van Donselaar 2009, Lian and Liu 1999), where R refers to the number of periods between two consecutive reviews of the inventory system, s denotes the inventory level below which an order is triggered, and S is the order-up-to level value. When demands are stochastic, obtaining optimal parameters in periodic-review policies even for a single perishable product with deterministic shelf life is notoriously complicated. The fixed shelf life perishability problem remains complex when the product lifetime is longer than two units of time in a periodic review system (Kouki and Jouini 2015). Hence, researchers have worked on approximating outdate costs (Broekmeulen and Van Donselaar 2009, Chiu 1995) or calculating upper and lower bounds on the number of outdates (Chiu 1995, Cooper 2001). Some models deal with batch demands (Lian and Liu 1999) or batch orders (Broekmeulen and Van Donselaar 2009). Finally, service level is regarded as a constraint in some papers including Adachi et al. (1999), Broekmeulen and Van Donselaar (2009), and Minner and Transchel (2010). The interested reader is referred to Goyal and Giri (2001), Karaesmen et al. (2011), and Nahmias (2011), for the review works on perishables.

In conclusion, the IRP and its variants are complex problems. Considering perishable products and stochastic demands complicates the model even further. As a result, applying existing solution methods is mathematically very difficult and computationally very inefficient.

3. The expected value method

A classical way to reduce the complexity of stochastic models is to replace random variables by their expected values. Unfortunately, even deterministic IRPs are extremely difficult to solve (see, e.g., Coelho and Laporte (2014)), which is why we settle here for a simple heuristic solution. The *expected value method* (EV) provides a first benchmark.

In EV, all demands are deterministic, given by $E(D_{ti})$, for all i, t. In a first step, each store independently determines its delivery quantities. Under these assumptions, the inventory control problem of a store can be viewed as a lot sizing problem where routing induces an implicit ordering

10

cost. The exact routing costs are determined in a second step. Hence, it is optimal for each store to place as few orders as possible while satisfying demand.

More specifically, assume that the inventory levels in period t in store i before delivery are given as $(x_1, x_2, \ldots, x_{L-1})_{ti}$, so that the total inventory level is

$$I_{ti} = \sum_{k=1}^{L-1} (x_k)_{ti}. \tag{1}$$

At the beginning of period t, for each store i, EV determines the delivery quantity as follows: if the current inventory level I_{ti} is larger than or equal to the expected demand $E(D_{ti})$, the delivery quantity is zero. Otherwise, EV delivers enough to satisfy the expected demands of L periods, including the current period, provided that store capacity C_i is respected. Then, a VRP based on these delivery quantities is solved and implemented, real demands are observed, new inventory levels are calculated for the next period, and the same decision making process is repeated in period $t + 1$.

All necessary computations can be carried out as follows. Assume that in period t, y_{ti} units are delivered to store i and the actual demand observed in store i is d_{ti}. Then, the inventory level of the store in period $t + 1$ is determined by relation (3), where $(z)^+ = \max(z, 0)$ and $(x_L)_{ti} = y_{ti}$ by convention:

$$((x_1, x_2, \ldots, x_{L-1})_{ti}, y_{ti}) \xrightarrow{d_{ti}} (x_1, x_2, \ldots, x_{L-1})_{t+1,i} \tag{2}$$

$$(x_k)_{t+1,i} = ((x_{k+1})_{ti} - (d_{ti} - \sum_{l=1}^{k} (x_l)_{ti})^+)^+, \quad \text{for } k = 1, \ldots, L-1. \tag{3}$$

The algorithm is summarized as follows.

The expected value algorithm (EV)

Begin

Step 0. Set $t = 1$.

Step 1. For each store i, if $I_{ti} \geq E(D_{ti})$, set $y_{ti} = 0$; otherwise, set $y_{ti} = \min\{C_i - I_{ti}, \lfloor E(D_{ti}) + \ldots + E(D_{t+L-1,i}) \rfloor - I_{ti}\}$.

Step 2. Solve a VRP for the delivery quantities y_{ti}'s and serve the stores through the optimal VRP routes.

Step 3. For each store i, observe the actual demand in period t, say d_{ti}. Calculate the state of the system in period $t + 1$, i.e., $X_{t+1,i}$ by Relations (3). Set $t = t + 1$ and go to Step 1.

End

In each period t and store i, the expected value method can be viewed as an (R, s, S) policy, where $R = 1$, $s = E(D_{ti})$, and $S = \min\{C_i, \lfloor E(D_{ti}) + \ldots + E(D_{t+L-1,i}) \rfloor\}$.

4. A deliver-up-to-level method

Another simple heuristic is obtained by replacing Step 1 of the *EV* algorithm with a replenishment rule that explicitly takes into account the stochasticity of demand. More specifically, for any $\lambda \leq L$, let us denote by $q_{ti}^{(\lambda)}$ the smallest integer quantity which suffices to meet the demand at store i during λ consecutive periods $\{t, \ldots, t + \lambda - 1\}$ with probability TSL:

$$Pr(D_{ti} + \cdots + D_{t+\lambda-1,i} \leq q_{ti}^{(\lambda)}) \geq TSL, \tag{4}$$

and consider the *deliver-up-to-level* method UL_λ.

The deliver-up-to-level algorithm (UL_λ)

Begin

  Step 0. Set $t = 1$.

  Step 1. For each store i, if $Pr(D_{ti} \leq I_{ti}) \geq TSL$, set $y_{ti} = 0$; otherwise, set $y_{ti} = \min\{C_i - I_{ti}, q_{ti}^{(\lambda)} - I_{ti}\}$.

  Step 2. Solve a VRP for the delivery quantities y_{ti} and serve the stores through the optimal VRP routes.

  Step 3. For each store i, observe the actual demand in period t, say d_{ti}. Calculate the state of the system in period $t + 1$, i.e., $X_{t+1,i}$, by Relations (3). Set $t = t + 1$ and go to Step 1.

End

The algorithm UL_λ (greedily) sets $y_{ti} = 0$ whenever I_{ti} suffices to satisfy TSL in period t, since a positive delivery quantity would increase the routing costs in period t. If the inventory does not suffice to satisfy TSL in period t, UL_λ delivers a quantity y_{ti} which should be sufficient to satisfy the demand in the next λ periods, unless the required quantity would exceed the store capacity C_i. So, UL_λ acts like an (R, s, S) policy where $s_{UL} = q_{ti}^{(1)}$, $S_{UL} = \min\{C_i, q_{ti}^{(\lambda)}\}$, and $R = 1$ so as to enforce TSL in every period. Note that $\lambda = 1$ tends to provide the freshest products on shelf, thanks to daily deliveries, whereas bigger values of λ yield lower routing costs and possibly higher profit.

5. A decomposition method

The deliver-up-to-level method UL_λ focuses on the target service level in order to determine the delivery quantities, and mostly ignores the importance of revenues and routing costs. Our next methods rely on a Stochastic Dynamic Programming (SDP) model, which explicitly account for these aspects.

In the SDP model, the state of the system in period t is defined by the inventory levels in all stores, i.e., $((x_1, \ldots, x_{L-1})_{t1}, \ldots, (x_1, \ldots, x_{L-1})_{tN})$. The decision variables are the delivery quantities

12

in period t, i.e., $(y_{t1}, \ldots, y_{tN})$, and the routing decisions. Given a decision on delivery quantities in period t, the direct costs (acquisition and routing) as well as the expected revenue in period t can be formulated, as well as the potential states of the system in period $t+1$ and the transition probabilities. Theoretically, one can set SDP relations to determine the optimal delivery quantities in each period based on the state of the system. However, this can only be applied to very small-size instances. Considering N stores, a maximum shelf life L, and inventory levels to be integers in the interval $[0, C]$, there are $(C+1)^{N(L-1)}$ potential states in each period. Therefore, it is necessary to resort to heuristic methods to solve even small instances through SDP.

We solve an independent SDP for each store, aiming to optimize an estimate of the store's expected revenue over the planning horizon. Although such a decomposition yields sub-optimal solutions, the complexity no longer depends exponentially on the number of stores. The number of states in each period for each store is $(C+1)^{(L-1)}$, which is computationally tractable for small values of L. In each period, the SDP relations allow us to determine a delivery quantity to each store, based on its current inventory level, while neglecting the routing costs. Then, we solve a VRP to obtain optimal routes for these delivery quantities. We call this the *decomposition method* (DE).

Since the SDP model considers each store independently, it does not properly account for the routing costs. Therefore, in the model associated with store i, we charge a fixed *cost-to-serve* F_{ti} if the store is visited in period t. This cost-to-serve acts as a proxy for the routing cost for delivery to store i. The choice of F_{ti} is discussed later.

Define $X_{ti} = ((x_1)_{ti}, \ldots, (x_{L-1})_{ti})$, $X_t = (X_{t1}, \ldots, X_{tN})$, and $Y_t = (y_{t1}, \ldots, y_{tN})$, so that (X_t, Y_t) denotes the complete state of the system at time t after the quantities $y_{t1}, \ldots, y_{tN}$ have been delivered. Define $f_{ti}(X_{ti}, y_{ti})$ as the total expected profit for store i from period t until the end of the planning horizon when the state of the store is X_{ti} and the delivery quantity is y_{ti}. The function f_{ti} includes total revenue, acquisition costs, and the cost-to-serve. The optimal expected profit generated by store i from period t to the end of the horizon is denoted $f_{ti}^*(X_{ti})$, that is,

$$f_{ti}^*(X_{ti}) = \max_{y_{ti}^{(1)} \leq y_{ti} \leq C_i - I_{ti}} f_{ti}(X_{ti}, y_{ti}), \tag{5}$$

where $y_{ti}^{(1)}$ is the smallest integer satisfying Inequality (4). The optimal delivery quantity is specified by Equation (6):

$$y_{ti}^* = y_{ti}^*(X_{ti}) = \arg \max_{y_{ti}^{(1)} \leq y_{ti} \leq C_i - I_{ti}} f_{ti}(X_{ti}, y_{ti}). \tag{6}$$

13

In order to determine the optimal delivery quantity y_{ti}^*, we solve the recursive Equations (7) by backward induction:

$$
\begin{aligned}
f_{ti}(X_{ti}, y_{ti}) =\ & -F_{ti} \cdot \mathbb{1}\,(y_{ti} > 0) - a y_{ti} \\
& + Pr(D_{ti} > I_{ti} + y_{ti})\big(p(I_{ti} + y_{ti}) + f_{t+1,i}^*(0,\dots,0,0)\big) \\
& + \sum_{d=0}^{I_{ti}+y_{ti}} Pr(D_{ti} = d)\big(pd + f_{t+1,i}^*(X_{t+1,i})\big)
\end{aligned}
\tag{7}
$$

where $X_{t+1,i} = (x_1,\dots,x_{L-1})_{t+1,i}$ is defined by Equation (3). The first term in Equation (7) is an estimate of the routing cost incurred to serve store i; the second term is the acquisition cost of y_{ti} units; the third term accounts for the expected revenue collected from store i when the demand in period t is larger than the inventory available at i in period t, and for the expected profit in periods $t+1,\dots,T$; similarly, the last term expresses the expected revenue in period t and the expected profit in periods $t+1,\dots,T$ when the demand does not exceed the available inventory. In order to solve (7), we use the boundary condition:

$$
f_{T+1,i}^*((x_1, x_2,\dots,x_{L-1})_{T+1,i}) = \frac{a}{2} I_{T+1,i},
\tag{8}
$$

where the right-hand side of (8) is an estimate of the profit generated by the inventory left over at the end of the horizon.

The decomposition algorithm is as follows.

The decomposition algorithm (DE)
Begin

Step 0. Set a cost-to-serve, F_{ti}, for each store i and each period t based on one of the algorithms described in Appendix A. Set $t = 1$.

Step 1. Use Equations (6)–(7) to determine a delivery quantity to each store i in period t, say y_{ti}^*, given the state of the system $X_{ti} = (x_1,\dots,x_{L-1})_{ti}$.

Step 2. Solve a VRP for the delivery quantities y_{ti}^* and serve the stores through the optimal VRP routes.

Step 3. For each store i, observe the actual demand in period t, say d_{ti}. Calculate the state of the system in period $t+1$, i.e., $X_{t+1,i}$ by Relations (3). Set $t = t+1$ and go to Step 1.

End

The *DE* algorithm can yield adequate solutions for the PSIRP provided that the costs-to-serve F_{ti} are reliable estimates of the actual routing costs. For symmetric travel costs satisfying the triangle inequality, a natural range for the cost-to-serve of store i is $[0, K + 2c_{i0}]$, where the upper bound

represents the cost of a direct delivery to store i. When we set $F_{ti} = 0$ for all stores, we obtain an algorithm that we call DE_0. This leads to a high delivery frequency. It provides very fresh products but ignores, and, implicitly increases, the routing costs. When the delivery quantity is close to the vehicle capacity Q, no other store can be served on the same route, and $F_{ti} = K + 2c_{i0}$ is the correct routing cost. In this case, each store can be dealt with independently.

When more than one store is served by each route, neither $F_{ti} = 0$ nor $F_{ti} = K + 2c_{i0}$ proves to be a good setting. In Appendix A, we introduce two methods to calculate an intermediate cost-to-serve to be assigned to each store. The first approach yields a *distance-based cost-to-serve* F_{ti}^d which focuses on the average distance between each store and its closest neighbors. The second approach produces a *route-based cost-to-serve* F_{ti}^r which allocates the total cost of a route to the stores it includes.

6.   A decomposition-integration method

In this section, we improve our estimate of the expected profit given in Equation (7), by taking into account the actual routing costs in period t and by refining the approximation of the routing costs in period $t+1$ (as compared to the costs-to-serve F_{ti}). The values F_{ti} are still used from period $t+2$ onward. Note that given the state of the system at time t, say X_t, and the vector of delivery quantities denoted by Y_t, a first estimate of the total profit for periods t to T is simply obtained as:

$$\pi_t^1(X_t, Y_t) = \sum_{i=1}^{N} f_{ti}(X_{ti}, y_{ti}) = \sum_{i=1}^{N} f_{ti}((x_1, \dots, x_{L-1})_{ti}, y_{ti}). \tag{9}$$

Now, let $R(y_1, \dots, y_N)$ represent the optimal routing cost with delivery quantities $(y_1, \dots, y_N)$. Then, Equation (9) can be improved by replacing the fixed costs-to-serve by the true routing cost in period t. This leads to:

$$\pi_t^2(X_t, Y_t) = \sum_{i=1}^{N} f_{ti}(X_{ti}, y_{ti}) + \sum_{i=1}^{N} F_{ti} \cdot \mathbb{1}(y_{ti} > 0) - R(y_{t1}, \dots, y_{tN}). \tag{10}$$

In order to apply a similar correction to the routing costs for period $t+1$, denote by $y_{t+1,1}^+, \dots, y_{t+1,N}^+$ the optimal delivery quantities in period $t+1$. Note that these quantities depend in a complex way on (X_t, Y_t) and are actually random variables, since they also depend on the realization of the demands $D_{t1}, \dots, D_{tN}$ in period t. With these notations, another estimate of the total expected profit can be derived from Equation (10), as follows:

$$\pi_t^3(X_t, Y_t) = \sum_{i=1}^{N} f_{ti}((x_1, \dots, x_{L-1})_{ti}, y_{ti})$$
$$+ \sum_{i=1}^{N} F_{ti} \cdot \mathbb{1}(y_{ti} > 0) - R(y_{t1}, \dots, y_{tN})$$
$$+ \sum_{i=1}^{N} F_{ti} \cdot Pr(y_{t+1,i}^+ > 0 \,|\, Y_t) \tag{11}$$
$$- \sum_{(y_1, \dots, y_N)} Pr((y_{t+1,1}^+, \dots, y_{t+1,N}^+) = (y_1, \dots, y_N)|Y_t) \times R(y_1, \dots, y_N).$$

In this expression, the fourth term corrects the expected value of the cost-to-serve in period $t+1$, and the last term represents the expected value of the routing cost in period $t+1$, given the delivery decisions Y_t.

The optimal delivery quantity $y_{t+1,i}^+$ can be approximated by the expected value of $y_{t+1,i}^*$. Based on Equation (6), this can be estimated as follows (compare with Equation (7)):

$$
\begin{aligned}
E(y_{t+1,i}^* \,|\, y_{ti}) = {}& Pr(D_{ti} > I_{ti} + y_{ti})\, y_{t+1,i}^*(0,\ldots,0,0) \\
& + \sum_{d=0}^{I_{ti}+y_{ti}} Pr(D_{ti}=d)\, y_{t+1,i}^*(X_{t+1,i})
\end{aligned}
\tag{12}
$$

where $X_{t+1,i}$ is defined by Equation (3).

Replacing the random quantities $y_{t+1,i}^+$ by $\lfloor E(y_{t+1,i}^* \,|\, y_{ti}) \rfloor$ turns (11) into a deterministic problem where the delivery cost in period $t+1$ can be approximated by solving a single VRP. This approach has the drawback, however, of yielding strictly positive values $\lfloor E(y_{t+1,i}^* \,|\, y_{ti}) \rfloor$ for almost all stores i, which is unlikely to happen for the optimal delivery quantities $y_{t+1,i}^*$ because this would result in high routing costs. Therefore, we further modify our approximation by considering delivery quantities $(\tilde{y}_{t+1,i} \,|\, y_{ti})$ defined by Equation (13) hereunder, where ϵ_i is a user-parameter whose value depends on the magnitude of the demand ($\epsilon_i = \frac{1}{2}E(D_{it})$ proved suitable in our numerical experiments):

$$
(\tilde{y}_{t+1,i}|y_{ti}) = \begin{cases} \lfloor E(y_{t+1,i}^*|y_{ti}) \rfloor & \text{if } \lfloor E(y_{t+1,i}^*|y_{ti}) \rfloor > \epsilon_i, \\ 0 & \text{otherwise.} \end{cases}
\tag{13}
$$

The optimal expected total profit is then approximated by solving the optimization problem (14).

$$
\begin{aligned}
\max_{(y_{t1},\ldots,y_{tN})} \tilde{\pi}(y_{t1},\ldots,y_{tN}) = {}& \sum_{i=1}^{N} f_{ti}((x_1,x_2,\ldots,x_{L-1})_{ti},y_{ti}) \\
& + \sum_{i=1}^{N} F_{ti} \cdot (\mathbb{1}(y_{ti}>0) + \mathbb{1}((\tilde{y}_{t+1,i}|y_{ti})>0)) \\
& - R(y_{t1},\ldots,y_{tN}) - R((\tilde{y}_{t+1,1}|y_{t1}),\ldots,(\tilde{y}_{t+1,N}|y_{tN})) \\
\text{subject to} \quad & y_{ti}^{(1)} \leq y_{ti} \leq C_i - I_{ti}, \quad i=1,\ldots,N.
\end{aligned}
\tag{14}
$$

This formulation takes into account the routing costs in period t, the approximated expected routing costs in period $t+1$, and costs-to-serve for the following periods. The corresponding algorithm is summarized as follows:

The decomposition-integration algorithm (DI)

Begin

Step 0. Set a cost-to-serve, F_{ti}, for each store i and each period t based on one of the algorithms described in Appendix (A). Set $t=1$.

Step 1. Solve Problem (14) to obtain the delivery quantities y_{ti} and the corresponding routes in period t.

Step 2. Serve the stores with the delivery quantities y_{ti} through the routes obtained in Step 1.

Step 3. For each store i, observe the actual demand in period t, say d_{ti}. Calculate the state of the system in period $t+1$, i.e., $X_{t+1,i}$ by Relations (3). Set $t = t+1$ and go to Step 1.

End

In the next section, we propose a matheuristic algorithm to solve Problem (14), as required by Step 1 of Algorithm DI.

7. A matheuristic for Problem (14)

In this section, we propose a matheuristic for Problem (14).

Starting from the initial solution $Y_t = (y_{t1}^*, \ldots, y_{tN}^*)$ proposed by Equation (6), we generate a new feasible solution $Y_t' = (y_{t1}', \cdots, y_{tN}')$ as explained below, and we explore whether $\tilde{\pi}(Y_t')$ is larger than $\tilde{\pi}(Y_t)$. If so, we move to the new solution; otherwise, we generate another solution. The local improvement algorithm stops when a pre-defined number of consecutively generated new solutions are rejected due to either lack of improvement or infeasibility. Calculating $\tilde{\pi}$ in Problem (14) for any new solution involves solving two VRPs. In order to avoid these expensive computations, each new solution is not generated randomly but in a systematic way which allows us to recompute $\tilde{\pi}(Y_t')$ incrementally, by difference with $\tilde{\pi}(Y_t)$.

When moving from the current solution (the current set of delivery quantities and their optimal routes) to a new solution, the difference in the approximated expected total profit is calculated by Equation (15):

$$
\begin{aligned}
\Delta = \widetilde{\pi}(y_{t1}', y_{t2}', \ldots, y_{tN}') - \widetilde{\pi}(y_{t1}, y_{t2}, \ldots, y_{tN}) = \\
\sum_{i=1}^{N} (f_{ti}((x_1, x_2, \ldots, x_{L-1})_{ti}, y_{ti}') - f_{ti}((x_1, x_2, \ldots, x_{L-1})_{ti}, y_{ti})) \\
- (R(y_{t1}', \ldots, y_{tN}') - R(y_{t1}, \ldots, y_{tN})) \\
- (R((\tilde{y}_{t+1,1}'|y_{t1}'), \ldots, (\tilde{y}_{t+1,N}'|y_{tN}')) - R((\tilde{y}_{t+1,1}|y_{t1}), \ldots, (\tilde{y}_{t+1,N}|y_{tN}))) \\
+ \sum_{i=1}^{N} F_{ti} \cdot (\mathbb{1}(y_{ti}' > 0) - \mathbb{1}(y_{ti} > 0)) \\
+ \sum_{i=1}^{N} F_{ti} \cdot (\mathbb{1}((\tilde{y}_{t+1,i}'|y_{ti}') > 0) - \mathbb{1}((\tilde{y}_{t+1,i}|y_{ti}) > 0)).
\end{aligned}
\tag{15}
$$

Assume that in every move from the current solution to a new solution, we change the delivery quantities in such a way that the routes, and so the routing costs, in period $t+1$ do not change. Moreover, let us indicate the decrease in the routing costs in period t by δ:

$$
\delta = R(y_{t1}, \ldots, y_{tN}) - R(y_{t1}', \ldots, y_{tN}').
\tag{16}
$$

17

Then, one can rewrite Equation (15) as follows:

$$\Delta = \sum_{i=1}^{N} (f_{ti}((x_1, x_2, \ldots, x_{L-1})_{ti}, y'_{ti}) - f_{ti}((x_1, x_2, \ldots, x_{L-1})_{ti}, y_{ti}))$$

$$+ \delta + \sum_{i=1}^{N} F_{ti} \cdot (\mathbb{1}(y'_{ti} > 0) - \mathbb{1}(y_{ti} > 0)). \tag{17}$$

In the proposed matheuristic, a new solution is generated in such a way that the routing cost in period t decreases while the expected routing cost in period $t+1$ does not change. We explain the main idea here.

Assume that in period t, a store j is ejected from its current route and is inserted into another route, say route r^*, without modification of the delivery quantities Y_t. Denote by r' the route in period t derived from route r^* after inserting store j into it. If route r' is feasible for the delivery quantities Y_t and if the routing cost in period t decreases as a result of this ejection-insertion step, then, clearly, the new VRP solution is preferred to the previous one. Generally speaking, however, since the routes were optimally selected for the delivery quantities Y_t, route r' will be infeasible either with respect to its maximum allowed length or with respect to the capacity of the vehicle. In the first case, we simply reject the new solution. In the second case, we try to determine whether the delivery quantities Y_t can be adapted (presumably, decreased) in such a way that r' becomes feasible. However, modifying Y_t also induces an effect on period $t+1$ (more precisely, on the quantities $(\tilde{Y}_{t+1}|Y_t)$ which are likely to increase). In order to keep some control over this effect, therefore, we restrict ourselves to certain modifications of Y_t which do not affect the feasibility of the current routes in period t and period $t+1$.

For an arbitrary set of routes R, we denote by $N(R)$ the set of stores contained in some route of R (excluding the depot); when R contains a single route, say, $R = \{r\}$, we simply write $N(r)$ instead of $N(R)$. Then, we define:

- R_t, R_{t+1} are the sets of routes in period t and $t+1$, respectively, after the ejection-insertion step has been performed on store j;

- $D = N(r')$ is the set of stores visited on route r'; we allow their delivery quantities to decrease in period t so as to restore feasibility of route r' (D is for "decrease");

- $\overline{R}_{t+1} = \{r \in R_{t+1} \mid D \cap N(r) \neq \emptyset\}$ is the set of routes in period $t+1$ that contain at least one store in D; these are the routes in period $t+1$ which may be affected when we decrease a delivery quantity to a store in D in period t; we need to make sure that these routes remain feasible, and this can be achieved by decreasing the expected delivery quantities to some of the corresponding stores in period $t+1$; or indirectly, by increasing the deliveries to these stores in period t; we model this through the introduction of the sets $\overline{R}_t$ and I;

18

- $\overline{R}_t = \{r \in R_t \mid N(\overline{R}_{t+1}) \cap N(r) \neq \emptyset\}$ is the set of routes in period t that contain at least one store in $N(\overline{R}_{t+1})$; the routes in $\overline{R}_t$ are considered as being potentially affected in period t;

- $I = (N(\overline{R}_{t+1}) \cap N(\overline{R}_t)) \setminus D$ is the set of stores (excluding stores in D) in the affected routes in both periods t and $t+1$; we allow their delivery quantities to increase in period t so as to maintain the feasibility of the routes in $\overline{R}_{t+1}$ (I is for "increase").

Moreover, define the following binary decision variables:

- for each store $i \in D$, $v_{ih} = 1$ if the delivery quantity to store i decreases by h units; else, $v_{ih} = 0$;

- for each store $i \in I$, $v_{ih} = 1$ if the delivery quantity to store i increases by h units; else, $v_{ih} = 0$.

Assume that $\underline{m}_i$ (resp., $\overline{m}_i$) is an upper bound on the largest possible decrease (resp., increase) of the delivery quantity y_{ti} in period t. We will explain in Appendix B how such bounds can be computed. Then, an IP model to determine new delivery quantities in period t is set as follows.

$$\max \quad \sum_{i \in D} \sum_{h=0}^{\underline{m}_i} f_{ti}(X_{ti}, y_{ti} - h) \cdot v_{ih} + \sum_{i \in I} \sum_{h=0}^{\overline{m}_i} f_{ti}(X_{ti}, y_{ti} + h) \cdot v_{ih} \tag{18}$$

subject to

$$\sum_{h=0}^{\underline{m}_i} v_{ih} = 1 \qquad \forall i \in D \tag{19}$$

$$\sum_{h=0}^{\overline{m}_i} v_{ih} = 1 \qquad \forall i \in I \tag{20}$$

$$\sum_{i \in D} \sum_{h=0}^{\underline{m}_i} (y_{ti} - h) \cdot v_{ih} \leq Q \tag{21}$$

$$\sum_{i \in N(r) \cap I} \sum_{h=0}^{\overline{m}_i} (y_{ti} + h) \cdot v_{ih} + \sum_{i \in N(r) \setminus I} y_{ti} \leq Q \qquad \forall r \in \bar{R}_t \setminus r' \tag{22}$$

$$\sum_{i \in N(r) \cap D} \sum_{h=0}^{\underline{m}_i} (\tilde{y}_{t+1,i} | y_{ti} - h) \cdot v_{ih} + \sum_{i \in N(r) \cap I} \sum_{h=0}^{\overline{m}_i} (\tilde{y}_{t+1,i} | y_{ti} + h) \cdot v_{ih}$$
$$+ \sum_{i \in N(r) \setminus (D \cup I)} (\tilde{y}_{t+1,i} | y_{ti}) \leq Q \qquad \forall r \in \bar{R}_{t+1} \tag{23}$$

$$v_{ih} \in \{0,1\} \qquad \forall i \in D, h \in [0, \underline{m}_i] \ \text{ and } \ \forall i \in I, h \in [0, \overline{m}_i]. \tag{24}$$

The objective function (18) maximizes the total expected profit obtained by the new delivery quantities to the stores in sets D and I, i.e., the stores whose delivery quantities may change. Constraints (19) and (20) along with Constraints (24) imply that exactly one of the decision variables v_{ih} takes value 1 for each store $i \in D \cup I$. Constraint (21) indicates that the new delivery quantities to the stores in the expanded route r' must respect the vehicle capacity. Constraints (22)–(23) guarantee

that for every affected route in period t or $t+1$, the sum of the new delivery quantities does not exceed the vehicle capacity.

If the IP has a feasible solution, the new delivery quantities to the stores in D and I are calculated by using Equations (25) and (26), respectively. Delivery quantities to other stores do not change.

$$y'_{ti} = \sum_{h=0}^{m_i} (y_{ti} - h) \cdot v_{ih} \quad \forall i \in D \tag{25}$$

$$y'_{ti} = \sum_{h=0}^{\overline{m}_i} (y_{ti} + h) \cdot v_{ih} \quad \forall i \in I. \tag{26}$$

Note that only the routes belonging to either $\bar{R}_t$ or $\bar{R}_{t+1}$ appear in the IP formulation. Moreover, in order to decrease the current excess load on route r', we only consider in (18)–(24) a subset of promising stores (those in $D \cup I$) for which the current delivery quantities can either increase or decrease. Thus, we cannot claim that the optimal solution of Problem (18)–(24) provides the optimal adjustment of delivery quantities to restore the capacity constraint in route r'. In particular, Problem (18)–(24) may be infeasible, while there actually exists an adjustment of delivery quantities such that the capacity of route r' is not exceeded and all other routes in periods t and $t+1$ remain feasible.

Thus, summing up, our local search approach to Problem (14) acts as a "large neighborhood search" framework, which explores the neighborhood of the current solution by solving the IP sub-problem (18)–(24). The solution of (18)–(24) hopefully yields new delivery quantities which increase the expected total profit.

The following algorithm must be embedded in Step 1 of algorithm DI:

The Matheuristic

Begin

Step 0. Initial solution: Solve two independent VRPs for periods t and $t+1$ where the delivery quantities are respectively $y_{ti} = y^*_{ti}$ and $(\tilde{y}_{t+1,i}|y_{ti})$ calculated by Equations (6) and (13).

Step 1. Termination: If Steps 2-5 have been repeated for a predetermined number of iterations, then stop.

Step 2. Ejection-insertion: Choose two random stores j and j' which are served in period t but are not included in the same route. Assume that j is ejected from its current route and is inserted immediately before or after j', whichever leads to a lower cost for the expanded route r'. If the expanded route r' is infeasible in terms of the route length, go to Step 1.

Step 3. Saving: Calculate δ as the decrease in the routing costs in period t resulting from the ejection-insertion in Step 1. If $\delta \leq 0$, go to Step 1.

20

Step 4. New deliveries: If the sum of the current delivery quantities on r' does not exceed Q, go to Step 5; otherwise, solve Problem (18)–(24). If the problem does not have a feasible solution go to Step 1; otherwise, calculate the new delivery quantities by Equations (25)–(26).

Step 5. Move: Use Equation (17) to calculate Δ, i.e., the difference between the expected total profit for the new solution and the current solution. If $\Delta > 0$ move to the new solution. Go to Step 1.

End

The Matheuristic proposed to solve Problem (14) relies on decreasing the routing costs in period t while keeping the expected routes in period $t + 1$ unchanged. We tested the reverse as well, i.e., modifying the delivery quantities in period t so that the routes in period t do not change while the expected routing costs in period $t + 1$ decrease. Our results show that this strategy does not perform well. This may be due to the fact that the second strategy tries to decrease the expected costs of routes which may not be realized at all in period $t + 1$.

8. Full information

When assessing the performance of the above algorithms, it is interesting to consider the *value of full information*, that is, the additional expected profit that could be reaped if full information about future demands was available to the decision-maker. In that case, the PSIRP simplifies to a deterministic PIRP. In this section, we develop a simple heuristic, denote by FI, for the resulting PIRP.

From an inventory perspective, FI delivers these quantities leading to no waste and no lost sales. There is thus no difference between delivering the demand of the current period only, or the demands of two periods ahead, or the demands of λ periods ahead as long as $\lambda \leq L$, since no inventory holding cost is charged. From the routing point of view, however, it can be beneficial to serve all stores in the same periods and with larger delivery quantities. In this sense, when demands are deterministic, serving all stores every $\lambda = L$ periods sounds like an effective strategy. However, $\lambda < L$ might be a better choice than $\lambda = L$ because the average filling rate of vehicles could be higher. Therefore, in our experiments, we tested other λ values too.

The algorithm FI is as follows.

The full information algorithm (FI)

Begin

Step 0. Set $t = 1$ and λ.

Step 1. For each store i, set $y_{ti} = d_{ti} + \ldots + d_{t+\lambda-1,i}$, where d_{ti} is the deterministic demand in period t in store i.

Step 2. Solve a VRP for period t by considering delivery quantities y_{ti}, and serve the stores with these delivery quantities through the optimal routes.

Step 3. For each store i, set $y_{t+1,i} = \ldots = y_{t+\lambda-1,i} = 0$. Set $t = t + \lambda$ and go to Step 1.

End

9. Computational study

All algorithms are coded in Java and the instances are run on an Intel Core i7 processor with 1.8GHz CPU and 8GB RAM. No time limit is imposed to any of the algorithms.

To solve the VRP models, we use a fast but effective heuristic. The heuristic first solves the LP relaxation of a route-based formulation by column generation (Righini and Salani 2006). Then, the restricted master problem obtained at the end of the column generation process is solved to optimality as an integer programming problem by calling ILOG CPLEX 12.4. Testing this heuristic on the original random instances created by Solomon (1987) showed an average optimality gap of 0.6% with respect to the exact optimal values. CPLEX is also used to solve the integer programming problems (18)-(24) described in Section 7.

9.1. Instances

For the computational experiments, the first $N = 40$ stores in the R-series random instances created by Solomon (1987) are considered with some modifications. Each route length remains limited to 230 time units, but we remove time window constraints. The vehicle capacity is $Q = 120$. Demands are randomly generated during a planning horizon of $T = 30$ periods.

Demands from the end customers to the stores, D_{ti}, are i.i.d. random variables following a binomial distribution with parameters $n = 200$ and $p = 0.1$, i.e., $D_{ti} \sim Bin(200, 0.1)$. The average demand is $E(D_{ti}) = 20$ for each period and for each store. We consider three shelf lives, namely, $L \in \{2, 3, 4\}$. As in the original instances, the fixed cost of using each vehicle is $K = 0$, and Euclidean distances represent the cost c_{ij} of traveling from store i to store j. The acquisition price and selling price per unit are respectively $a = 6$ and $p = 10$. A target service level of $TSL = 90\%$ is to be respected in every period and every store. We set store capacities such that they are not restrictive in any solution method. When algorithms UL_λ, DE, and DI are applied, $C_i = 40$ (resp., 60, 80) is large enough for $L = 2$ (resp., 3, 4). When algorithm FI is applied, $C_i = 80$ (resp., 100, 120) is considered for $L = 2$ (resp., 3, 4). We will later analyze the impact of limited store capacity on profit, freshness, and actual service level. In the next subsections, we discuss some of the performance measures that we have collected.

9.2. Simulation

In order to evaluate the performance of different solution methods, we use random scenarios to simulate the sequence of decisions made by each method over a rolling horizon of $T = 30$ periods. We generate a set of 30 scenarios, where each scenario consists of initial inventory as well as demands

22

of the stores over the planning horizon. We use the same scenarios for all solution methods and for all shelf lives. The initial inventory of each store is a uniform random number in the interval $[0, 30]$ (resp., $[0, 50]$, $[0, 70]$) for $L = 2$ (resp., 3, 4), and is considered to have shelf life $L - 1$.

For each solution method, the expected profit is estimated by averaging the total profit over the 30 random scenarios. dWe also collect other useful information such as average actual service level and freshness as additional criteria to measure the performance of each method.

9.3. Actual service level

For each run of the simulation, we calculate the average actual service level, based on (1) the number of stock-outs (ξ_s) and (2) the fill rate (ξ_f). Recall that, according to Equation (1), I_{ti} indicates the inventory level at the beginning of period t in store i, i.e., the inventory level before delivery. The quantity $\mathbb{1}(d_{ti} \leq I_{ti} + y_{ti})$ is 1 if no stock-out happens in period t in store i, and 0 otherwise. Hence, in Equation (27) hereunder, ξ_s is the proportion of observations where no stock-outs occurred, over all stores and all periods. This metric is consistent with our initial definition of TSL in Equation (4).

$$\xi_s = \frac{\sum_t \sum_i \mathbb{1}(d_{ti} \leq I_{ti} + y_{ti})}{T N}. \tag{27}$$

Our second definition of service level considers the fill rate of demands. In this case, $\min\{d_{ti}, I_{ti} + y_{ti}\}$ shows the demand satisfied in period t in store i. Thus, Equation (28) calculates the average fill rate of demand in all stores over the planning horizon.

$$\xi_f = \frac{\sum_t \sum_i \min\{d_{ti}, I_{ti} + y_{ti}\}}{\sum_t \sum_i d_{ti}}. \tag{28}$$

9.4. Actual freshness

For each run of the simulation, the actual freshness of products is calculated in two ways. First, the average actual freshness on shelf ϕ_s is calculated by Equation (29).

$$\phi_s = \frac{\sum_t \sum_i (1x_1 + 2x_2 + \cdots + (L-1)x_{L-1})_{ti} + Ly_{ti}}{\sum_t \sum_i (I_{ti} + y_{ti})}. \tag{29}$$

Secondly, Equation (30) is used to calculate ϕ_c, the average actual freshness from a customer's perspective. In this definition, $(s_k)_{ti}$ is the number of units with remaining shelf life k sold in period t in store i.

$$\phi_c = \frac{\sum_t \sum_i (1s_1 + 2s_2 + \cdots + (L-1)s_{L-1} + Ls_L)_{ti}}{\sum_t \sum_i (s_1 + s_2 + \cdots + s_L)_{ti}}. \tag{30}$$

23

9.5. Verifying route estimations in period $t+1$

In the decomposition-integration method DI, we use Equation (13) to approximate the expected deliveries in period $t+1$. The main purpose of this approximation is to estimate the routing costs in period $t+1$; see Equation (14). Therefore, the accuracy of the approximation is evaluated for each scenario by measuring the similarity between the set of routes forecasted when using Equation (13), denoted here by $E(R_{t+1})$, and the set of routes actually used in period $t+1$, namely, R_{t+1}. We define the degree of similarity between these sets by Equation (31):

$$Similarity = \frac{\sum_{(i,j)\in(R_{t+1}\cap E(R_{t+1}))} c_{ij}}{\sum_{(i,j)\in(R_{t+1}\cup E(R_{t+1}))} c_{ij}}. \tag{31}$$

9.6. Results

For each scenario, each solution method is applied over a rolling horizon of $T=30$ periods. Table 2 summarizes the results. The first column denotes the maximum shelf live $L \in \{2,3,4\}$. The second column indicates the solution methods applied to determine delivery quantities and routes for each scenario, namely: the expected value method (EV), deliver-up-to-level with daily deliveries (UL_1), deliver-up-to-level with large delivery quantities to satisfy TSL for $\lambda = L-1$ periods (UL_{L-1}), decomposition without costs-to-serve (DE_0), decomposition with distance-based costs-to-serve (DE_d), decomposition with route-based costs-to-serve (DE_r), decomposition-integration without costs-to-serve (DI_0), decomposition-integration with distance-based costs-to-serve (DI_d), decomposition-integration with route-based costs-to-serve (DI_r), and the full information method (FI).

Column 3 displays the average computation times over 30 scenarios for each instance. When $L=2$, most of the computation time is spent in solving the VRPs, in that all $N=40$ stores are served in every period when applying UL_λ, DE, or DI. When $L=4$, however, most of the computation time is devoted to solving the expensive SDP relations (7). In the latter case, solving the VRPs takes almost no time because the average number of stores served in each period is around 15.

The next columns report, respectively, average values over 30 scenarios of the profit, revenue, acquisition cost, routing cost, waste cost, average number of vehicles per period, average number of stores per route, average time between two consecutive visits to stores, freshness on shelf, freshness from customers' perspective, service level based on the number of stock-outs, and service level based on the filling rate.

In order to increase the readability of the table, the values in Columns 4–8 are normalized with respect to the profit obtained by EV for each shelf life (the absolute value of the profit obtained by EV is shown between parentheses). For UL_λ, we obtained the highest profit by setting $\lambda = L-1$. For FI, we obtained the highest profit by setting $\lambda = 2$ (resp., 3, 3) for $L=2$ (resp., 3, 4).

24

Table 2   Comparing different solution methods

L	Method	Time(sec)	Pro.=	Rev.	−Acq.	−Rou.	Waste	Veh.	Cus.	Bet.	ϕ_s	ϕ_c	ξ_s	ξ_f
	EV	4	1.000 (66542)	3.598	2.331	0.267	0.149	8.5	3.2	1.6	1.6	1.4	91%	98%
	UL_1	1626	1.126	3.596	2.175	0.295	0.008	7.7	5.6	1.0	1.7	1.7	93%	99%
	UL_{L-1}	1626	1.126	3.596	2.175	0.295	0.008	7.7	5.6	1.0	1.7	1.7	93%	99%
2	DE_0	1802	1.129	3.623	2.196	0.297	0.011	7.8	5.5	1.0	1.7	1.6	97%	99%
	DE_d	1484	1.130	3.623	2.196	0.297	0.011	7.8	5.5	1.0	1.7	1.6	97%	99%
	DE_r	1482	1.130	3.623	2.196	0.297	0.011	7.8	5.3	1.0	1.7	1.6	97%	99%
	DI_0	8783	1.133	3.622	2.196	0.293	0.011	7.5	5.5	1.0	1.7	1.6	97%	99%
	DI_d	7003	1.133	3.621	2.196	0.293	0.011	7.6	5.7	1.0	1.7	1.6	97%	99%
	DI_r	7188	1.133	3.621	2.196	0.293	0.011	7.5	5.5	1.0	1.7	1.6	97%	99%
	FI	7	1.210	3.604	2.170	0.224	0.007	7.8	2.6	2.0	1.6	1.5	100%	100%
	EV	3	1.000 (72320)	3.376	2.164	0.212	0.103	8.5	2.0	2.6	2.3	1.9	95%	99%
	UL_1	1673	1.045	3.310	1.994	0.271	0.000	7.7	5.6	1.0	2.7	2.7	93%	99%
	UL_{L-1}	2	1.082	3.379	2.065	0.232	0.011	8.1	3.1	1.8	2.4	2.0	99%	99%
3	DE_0	1575	1.052	3.359	2.032	0.274	0.000	7.7	5.5	1.0	2.6	2.3	99%	99%
	DE_d	49	1.082	3.382	2.073	0.227	0.016	8.2	2.8	1.9	2.3	1.9	99%	99%
	DE_r	50	1.078	3.384	2.079	0.227	0.020	8.2	2.0	2.0	2.3	1.9	99%	99%
	DI_0	7697	1.057	3.356	2.036	0.263	0.005	7.5	5.5	1.0	2.6	2.3	99%	99%
	DI_d	98	1.088	3.384	2.073	0.223	0.015	7.8	3.0	1.9	2.4	2.0	99%	99%
	DI_r	98	1.085	3.383	2.077	0.221	0.019	7.8	2.0	2.0	2.3	1.9	99%	99%
	FI	1	1.132	3.317	1.995	0.189	0.005	8.0	1.7	3.0	2.3	2.0	100%	100%
	EV	0	1.000 (72182)	3.418	2.182	0.236	0.084	11.7	1.0	3.6	3.0	2.5	97%	99%
	UL_1	1672	1.046	3.315	1.997	0.272	0.000	7.7	5.6	1.0	3.7	3.7	93%	99%
	UL_{L-1}	0	1.109	3.423	2.110	0.205	0.014	8.2	1.0	2.8	3.1	2.5	99%	99%
4	DE_0	2568	1.054	3.364	2.036	0.275	0.000	7.7	5.5	1.0	3.6	3.3	99%	99%
	DE_d	1086	1.111	3.425	2.109	0.205	0.012	8.2	1.9	2.8	3.1	2.5	99%	99%
	DE_r	1107	1.106	3.431	2.121	0.204	0.018	8.2	1.4	2.9	3.0	2.4	99%	99%
	DI_0	7716	1.070	3.362	2.039	0.253	0.005	7.4	5.3	1.1	3.5	3.2	99%	99%
	DI_d	1365	1.111	3.424	2.104	0.210	0.007	7.9	2.6	2.2	3.2	2.7	99%	99%
	DI_r	1597	1.117	3.426	2.113	0.197	0.014	7.7	2.1	2.6	3.1	2.5	99%	99%
	FI	1	1.138	3.322	1.998	0.186	0.005	8.0	1.7	3.0	3.3	2.9	100%	100%

10. Discussion

We analyze and discuss the results of the computational study. We demonstrate that the differences between the profits generated by different solution methods are statistically and economically significant. We also draw some additional managerial insights.

10.1. Comparing the solution methods

Regarding freshness, Table 2 shows that method UL_1 provides the freshest products, while the other solution methods strive to reap a higher profit. All methods lead to extremely high service levels, especially when measured by the fill rate of demand ξ_f. These high service levels are obtained even for EV (which does not explicitly enforce the target service level TSL), especially when L is large. As indicated by ξ_s, the proportion of stock-outs is higher with EV and UL_1 than with the other methods, but is still reasonably low. Finally, all solution methods but EV yield extremely low waste costs.

Let us now take a closer look at the profit. Recall that all values are normalized so that the expected profit generated by EV is 1 in all scenarios. As expected, the average profit tends to increase when we move from method EV to UL, to DE, and to DI. Figure 1a illustrates the average additional profit when each of the methods is applied. When the demand distribution is known, DE_0 can be applied, whereby the profit increases on average by 12.9% (resp., 5.2%, 5.4%) for $L = 2$ (resp., 3, 4). This increase can be interpreted as the value of accessing the probability distribution and of explicitly accounting for the uncertainty of demand. The additional gap filled by DE_d shows the value of considering some aspect of routing when we determine delivery quantities. On average, it amounts to an additional increase in profit of 0.1% (resp., 3.0%, 5.7%) for $L = 2$ (resp., 3, 4). Then, by using DI_d the profit increases again on average by 0.3% (resp., 0.6%, 0.0%) for $L = 2$ (resp., 3, 4): this measures the value of further integrating inventory and routing-related decisions. Finally, accessing full information and applying FI provides some 7.7% (resp., 4.4%, 2.7%) average additional profit for $L = 2$ (resp., 3, 4). This can be interpreted as an estimate of the *value of full information*. Figure 1a also shows the profit gained by applying UL_{L-1}. Interestingly, the performance of this simple deliver-up-to-level policy is very close to the performance of the more sophisticated method DE. Figure 1b can be interpreted in the same way as the previous figure, when the cost-to-serve F^d is replaced by F^r.

We have also tested the performance of DI if the profit function in Problem (14) is replaced by a simpler estimate, say, (10). In the latter estimation, costs-to-serve $F_{t+1,i}$ are considered for period $t + 1$, whereas Problem (14) calculates a VRP routing cost for period $t + 1$ based upon the expected delivery quantities. Solving (10) is easier than (14), but our computational results show that the profit generated by (10) falls between the profits generated by DE and by DI.

10.2. Optimality gaps

To assess the algorithms performance, it is necessary to compute a tight upper bound on the optimal value of the expected profit. Although difficult, we nevertheless measure the quality of our solutions in various ways.
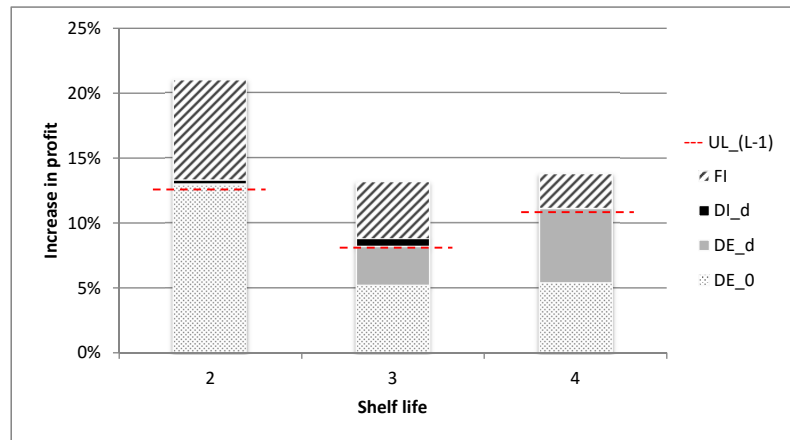
(a) distance-based cost-to-serve
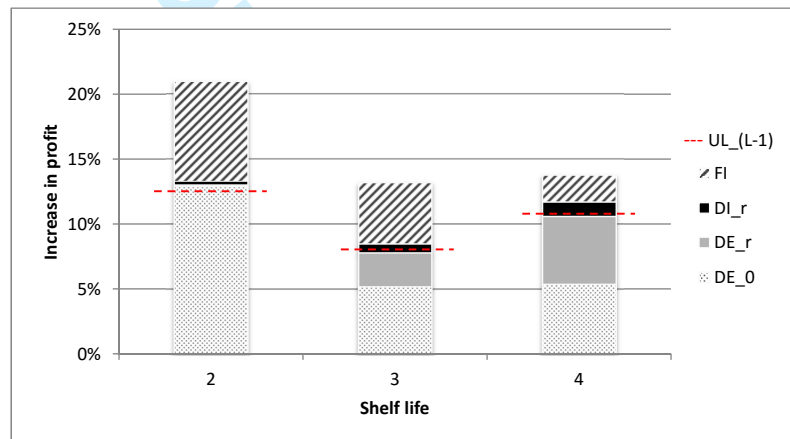
(b) route-based cost-to-serve

Figure 1   **Contribution of different solution methods to increasing profit.**

First, the gap between the best solution value (obtained by DI) and the value obtained by FI is relatively small. This gap cannot be interpreted as a rigorous optimality gap, since FI assumes perfect information and relies on a heuristic algorithm. But the value of the gap suggests, nevertheless, that DI is performing reasonably well.

Next, consider the generic expression of the expected profit as $E(Profit) = E(Revenue - Acquisition) - E(Routing)$. Clearly, the revenue of the retail chain is maximized when all stores are able to satisfy the demand of their customers, that is, when the service level reaches 100%. Under these conditions, the acquisition costs are minimized when all units bought are sold, that is, when there is no waste. Now, Table 2 indicates that the solutions obtained by DI_d or DI_r achieve a very high service level (fill rate of 99%) and produce very little waste (around 1-2%). This already implies that $E(Revenue - Acquisition)$ is almost best possible. A more precise calculation can be carried out as follows: in each period t and for every store i, $E(Revenue_{ti} - Acquisition_{ti}) \leq E(D_{ti})(p - a)$.

27

In our experiments, $E(D_{ti})(p-a) = 80$ so that, with 30 periods and 40 stores, $E(Revenue) - E(Acquisition)$ is bounded by 96000. For each shelf life $L = 2, 3, 4$, respectively, the average value of $(Revenue - Acquisition)$ achieved by DI_r over 30 random scenarios is equal to 94869, 94465, 94824, respectively, which is again within 1-2% of the upper bound.

On the other hand, it is much more difficult to bound the expected routing costs, since they depend on the delivery quantities. A rough lower bound on the expected routing costs can be computed by assuming that each store is served by direct shipments and with full vehicle capacity. When the total demand of store i over the planning horizon is d, the resulting lower bound is $2c_{0i}\lfloor d/Q \rfloor$, where c_{0i} is the travel cost from the depot to the store. Hence, a lower bound on the expected routing cost incurred by store i during the planning horizon is:

$$2c_{0i} \sum_{d=0}^{nT} Prob(D_i = d) \lfloor d/Q \rfloor.$$

With our parameter settings, this leads to a lower bound of 9040 on the expected total routing costs, as compared with 19478, 15998 and 14196, respectively, for the average routing costs obtained by DI_r when $L = 2, 3, 4$, respectively.

Even though this bound is quite weak, we conclude that the total average profit generated by DI_r is within 13%, 10%, 7%, respectively, of the optimal expected profit when $L = 2, 3, 4$. Note that the actual optimality gap is most probably much better, since our routing algorithm does quite well, as observed in Section 9.

10.3. Statistical tests

Algorithm DI takes the solution provided by DE as an initial solution and improves it by applying a matheuristic. The average improvement of profit appears to be small, but actually proves statistically significant. When estimating the profit over a sample of 30 scenarios, the standard deviation of the estimate is about 0.001 times its average value, which shows that the in-sample variation of the profit is extremely small and hence, differences between methods become quickly significant. We also checked that, when computing the profit over two independent samples, the difference between the average profit obtained for DI is one order of magnitude smaller than the difference between methods.

Additionally, we test the statistical hypothesis $H_0 : P_{DI_d} \leq P_{DE_d}$, where P_{DE_d} and P_{DI_d} indicate the total profits obtained by DE_d and DI_d, respectively. The results of the t-test for paired samples are displayed in Table 3. The threshold for the t-statistic with 29 degrees of freedom (above which the null hypothesis is rejected with confidence level 99.99%) is $t_{0.9999;29} = 4.25$. Table 3 shows that H_0 is rejected in all cases but one, which shows that DI dominates DE in terms of profit.

28

Table 3   Statistical tests on profit

L	C	$E(P_{DI_d} - P_{DE_d})$	$Std(P_{DI_d} - P_{DE_d})$	t-statistic	Hypothesis	Result	Prof. inc.
2	30	192	53	20.0	$H_0 : P_{DI_d} \leq P_{DE_d}$	Reject	0.3%
	$\geq$40	217	62	19.2	$H_0 : P_{DI_d} \leq P_{DE_d}$	Reject	0.3%
3	30	118	48	13.4	$H_0 : P_{DI_d} \leq P_{DE_d}$	Reject	0.2%
	40	379	96	21.6	$H_0 : P_{DI_d} \leq P_{DE_d}$	Reject	0.5%
	50	391	147	14.6	$H_0 : P_{DI_d} \leq P_{DE_d}$	Reject	0.5%
	$\geq$60	408	130	17.2	$H_0 : P_{DI_d} \leq P_{DE_d}$	Reject	0.5%
4	30	118	48	13.4	$H_0 : P_{DI_d} \leq P_{DE_d}$	Reject	0.2%
	40	407	85	21.6	$H_0 : P_{DI_d} \leq P_{DE_d}$	Reject	0.5%
	50	352	111	17.4	$H_0 : P_{DI_d} \leq P_{DE_d}$	Reject	0.4%
	60	911	204	24.4	$H_0 : P_{DI_d} \leq P_{DE_d}$	Reject	1.2%
	70	233	166	7.7	$H_0 : P_{DI_d} \leq P_{DE_d}$	Reject	0.3%
	$\geq$80	30	230	0.7	$H_0 : P_{DI_d} = P_{DE_d}$	Accept	0.0%

Superiority of DI over DE is not confined to improving the profit. In particular, DI uses fewer vehicles than DE, and the difference is again statistically significant, as shown in Table 4. Similar conclusions apply when F^d is replaced by F^r. These results demonstrate that it makes sense to use our Matheuristic to build upon DE.

Table 4   Statistical tests on number of vehicles

| L	C	$E(V_{DE_d} - V_{DI_d})$	$Std(V_{DE_d} - V_{DI_d})$	t-statistic	Hypothesis	Result	Vehi. dec.
2	30	0.22	0.09	13.4	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	2.8%
	$\geq$40	0.21	0.08	14.4	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	2.8%
3	30	0.13	0.11	6.5	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	1.7%
	40	0.19	0.10	10.4	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	2.5%
	50	0.32	0.10	17.5	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	3.9%
	$\geq$60	0.34	0.11	16.9	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	4.2%
4	30	0.13	0.11	6.5	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	1.7%
	40	0.22	0.10	12.0	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	2.8%
	50	0.31	0.09	18.9	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	3.8%
	60	0.91	0.12	41.5	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	10.3%
	70	0.42	0.13	17.7	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	5.1%
	$\geq$80	0.30	0.16	10.3	$H_0 : V_{DI_d} \geq V_{DE_d}$	Reject	3.6%

10.4.   Impact of cost-to-serve values on DE and DI

So far, we defined two ways to assign a positive cost-to-serve to a generic store, namely, F^d and F^r. In general, F^r is much larger than F^d. Therefore, we also tested the sensitivity of the performance of DE and DI when other values of the cost-to-serve are considered. The results are shown in Figure 2. It appears that, for all three shelf lives, DE achieves its best performance when F^d is set as cost-to-serve. (Note that the horizontal axis is normalized so that $F^d = 1$ in all cases, and the vertical axis shows the relative profit with respect to EV.) However, when DI is used, the best setting of cost-to-serve is not that clear: while F^d provides the highest profit for shelf lives $L = 2$ and $L = 3$, $F^r \approx 2.2 F^d$ results in the best profit for shelf life $L = 4$.

Our results show that if the average number of stores per route, say $\bar{n}$, is at least 3, then F^d works well for all values of L. On the other hand, $\bar{n} < 2$ implies that routes rarely include more than

29

(a) $L = 2$

(b) $L = 3$

(c) $L = 4$

Figure 2 **Profit obtained by** DE **and** DI **with different costs-to-serve for** $Q = 120$.

30

two stores. In this case, $F_i = K + 2c_{i0}$ proves a better estimation for cost-to-serve than F^d, and we can adopt a direct delivery policy for store i. This is consistent with the results in (Gallego and Simchi-Levi 1990, Bertazzi 2008), i.e., when the delivery quantity is a large fraction of the vehicle capacity, direct shipping is preferable in almost all routing strategies.

Figure 2 also demonstrate that, regardless of the cost-to-serve value, there is always profit improvement using DI.

10.5. Managerial insights

In this section, we discuss the economic significance of the profit improvements, the impact of shelf life and store capacity on the profit obtained by the best solution method, and the interpretation of DE as an (R, s, S) policy.

Economic interpretation of profit improvements. We argue here that the improvement in profit provided by DI over DE is not only statistically, but also economically significant. In our experimental setting, profit is about 32% of total revenue, but does not account for a variety of miscellaneous costs (salaries, buildings, marketing, administration, and so forth). In fact, net profit in the retail food sector is of the order of 2% of revenue; see EBRD–FAO (2009), NAICS (2012), FMI (2014). Miscellaneous costs thus account for about 30% of total revenue. The corresponding breakdown of the revenue is depicted in Figure 3.


Figure 3   **Breakdown of the revenue.**

Our results show an average improvement of 0.6% in profit when we exploit DI as compared to DE while setting an appropriate cost-to-serve (0.3% for $L = 2$ with F^d, 0.5% for $L = 3$ with F^d, and 1.0% for $L = 4$ with F^r). This translates into 0.19% of the revenue ($= 0.6\%$ of 32%), meaning about 10% of the net profit of a typical retail chain. This is economically significant.

Value of information. We already provided an estimate of the value of information in Section 10.1: namely, the results in Figure 1 clearly show that perfect information about future demand

potentially leads to significant profit improvements. Reducing demand uncertainty may be achieved, for example, by collecting and analyzing increased amounts of data relating to purchasing habits in different stores. In practice, however, it is unrealistic to assume that retail demand can be treated as completely deterministic. In order to relax this assumption, we run a few additional experiments when the demand is binomial with parameters $(n, p) = (40, 0.5)$, as compared with $(n, p) = (200, 0.1)$ in the initial instances. Observe that this modification preserves the mean value of the demand ($= 20$) but decreases its variance $np(1-p)$ from 18 to 10. So, it allows us to investigate the modification of the "optimal" profit when the uncertainty is reduced. The results obtained with the algorithm DI_r (using route-based costs-to-serve) are summarized in Table 5 (they are representative of the general trend). They show that when the variance decreases, the profit increases by roughly 0.5%, a significant economic improvement.

Table 5   Comparing profits when the variance decreases

L	(n,p)	Profit by DI_r	Improvement
2	(200,0.1)	75392	
	(40,0.5)	75588	0.3%
3	(200,0.1)	78467	
	(40,0.5)	79407	1.2%
4	(200,0.1)	80628	
	(40,0.5)	81067	0.5%

Impact of store capacity. Intuitively, one might expect that the smaller F_{ti}, the more frequently store i is visited. However, our experimental results show that the frequency of visits is relatively insensitive to F_{ti}, provided that F_{ti} is strictly positive. Therefore, selecting the value of the cost-to-serve cannot be regarded as a lever to adjust the frequency of visits and the freshness of products. On the other hand, the store capacity clearly has an effect on these performance indicators.

Table 6 shows the expected profit obtained by DI_d over $T = 30$ periods, as well as freshness and service level, when considering a limited store capacity C. We see that the service level is only slightly influenced by C. The changes in profit and freshness, however, are significant. The results suggest that providing extra store capacity beyond $(L-1)E(D) + 0.5E(D)$ does not have any major impact on profit. Observe that $(L-1)E(D)$ can be viewed as the expected required capacity between two consecutive visits (during $L-1$ periods) when the visits are maximally spread, while the quantity $0.5E(D)$ acts as buffer inventory to respect the target service level during $L-1$ periods on average.

The last column of Table 6 indicates the similarity values calculated by Equation (31). The values show that DI estimates reasonably well the expected routes in period $t+1$.

Impact of shelf life. Figure 4 shows the profits (in absolute value) obtained by DI_d for $L = 2, 3$ and by DI_r for $L = 4$. The numbers indicate a 2.5% decrease in profit when shelf life decreases from 4 to 3. Some further 4.3% loss in profit is incurred when moving from shelf life 3 to 2. These values can

32

Table 6   Comparing the impact of different capacities on profit, freshness, and service level

L	C	Profit by DI_d	ϕ_s	ϕ_c	ξ_s	ξ_f	Similarity
2	30	75375	1.7	1.6	97%	99%	70%
	$\geq$40	75394	1.7	1.6	97%	99%	71%
3	30	76216	2.7	2.5	98%	99%	69%
	40	76997	2.5	2.1	99%	99%	50%
	50	78613	2.4	2.0	99%	99%	53%
	$\geq$60	78628	2.4	2.0	99%	99%	53%
4	30	76216	3.7	3.5	98%	99%	69%
	40	77130	3.5	3.1	99%	99%	50%
	50	79790	3.3	2.9	99%	99%	53%
	60	80108	3.2	2.8	99%	99%	72%
	70	80368	3.2	2.7	99%	99%	70%
	$\geq$80	80368	3.2	2.7	99%	99%	70%

be interpreted as the cost of perishability. Recall that these decreases translate into $2.5 \times 16 = 40.7\%$ and $4.3 \times 16 = 68.6\%$ decreases in the net profit, which are extremely significant.
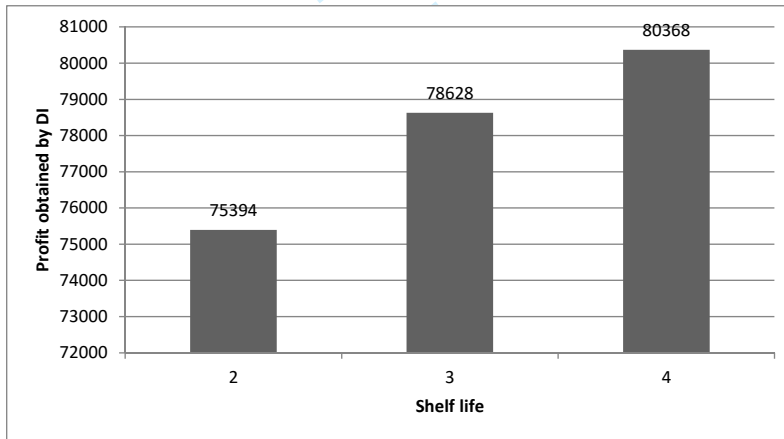

Figure 4   The best profit obtained by DI for different shelf lives.

Translating DE ***into an*** (R, s, S) ***policy.*** Interestingly, as illustrated by Figure 2, the simple algorithm UL_{L-1} is a strong competitor for DE when $L \geq 3$, regardless of the value of the cost-to-serve. For some cost-to-serve values, UL_{L-1} even outperforms DE (but not the improved solution produced by DI). In our instances, the inventory level triggering a delivery in UL_{L-1} is $s_{UL} = 1.25E(D)$ for any L, and the up-to-level point is $S_{UL} = 1.3E(D)$ (resp., $2.4E(D)$, $3.45E(D)$) for $L = 2$ (resp., 3, 4) (see Equation (4) and the end of Section 4). Our computational experiments reveal that the inventory level triggering a delivery in method DE is $s_{DE} = 1.25E(D)$ for any L, independently of the value of X_{ti}, the state of the system in period t in store i. On the other hand, the delivery quantities prescribed by DE do depend on X_{ti}, but mostly through the value of the total inventory level $I_{ti} = \sum_{k=1}^{L-1}(x_k)_{ti}$. When TSL is high, the up-to-level point in DE, i.e., $S_{DE} = (I_{ti} + y_{ti}^*|X_{ti})$, is quite close to that in UL_{L-1}, i.e., $S_{UL} = (I_{ti} + y_{ti}^*|I_{ti})$, especially when L is large. The up-to-level

point in DE slightly increases when setting a higher cost-to-serve. Figure 5 shows the normalized

frequency (over all possible states X_{ti}) of values of the up-to-level point, S_{DE}, when DE_d is applied.


Figure 5   The normalized frequency of up-to-level points in DE_d for different shelf lives.

The figure shows that, regardless of the state of the system, the up-to-level point determined by

DE_d is very likely to be $1.4E(D)$ (resp., $2.5E(D)$, $3.4E(D)$) for $L = 2$ (resp., 3, 4). For example,

when $L = 3$ and the inventory level does not satisfy TSL in the current period, DE prescribes

$S_{DE} = 2.5E(D)$ as the up-to-level point in 88% of the states, whatever the breakdown of X_{ti} is.

This implies that all complex SDP Relations (7) can be developed once offline and be translated

into a simple and easy-to-interpret (R, s, S) policy, where $R = 1$, $s = y_{ti}^{(1)}$, and $S = 1.4E(D)$ (resp.,

$2.5E(D)$, $3.4E(D)$) for $L = 2$ (resp., 3, 4), without any major impact on the performance of DE.

10.6. Summary

All solution methods have their advantages and limitations, and each has proven to perform for

certain settings. The main features of EV, UL_{L-1}, DE, and DI are summarized in Table 7.

10.7. Extensions

All solution methods presented above can be extended to account for inventory holding costs or

for decaying products, i.e., products which lose their quality gradually over their shelf life. This is

straighforward for methods EV and UL, but less so for DE and DI. Let us define h as the inventory

holding cost per unit per period. Moreover, let us assume that the value of each unit of the product

decreases by h' monetary units in each period. The parameter h' can alternatively be considered as a

self-imposed penalty with the objective to increase freshness when modeling perishable products. In

other words, even if the selling price is actually constant during the shelf life (perishable products),

the retail chain may assume that the value of the product decreases linearly over time (decaying

products) in order to enforce higher freshness. In order to incorporate these elements in DE and

Table 7   Features of the different solution methods

Method	Feature
EV	Extremely simple
	Does not take the stochasticity of demands into consideration
	Leads to the lowest profit, service level, and freshness
	Leads to the highest waste cost and to the highest number of vehicles used
UL_{L-1}	Extremely simple
	Its special case, UL_1, provides freshest products
	Useless when TSL is not defined or is low
	A strong competitor for DE, especially when TSL is high
	Applicable to multiple products
DE	F^d performs best if $\bar{n} \geq 2$; otherwise F^r performs better
	Performs similarly to the deliver-up-to-level policy
DI	F^d delivers the best results provided that $\bar{n} \geq 3$
	Superior to DE statistically and economically, in terms of profit and number of vehicles per period
	Superiority over DE applies for a range of values of store capacity and different estimates of F_{ti}
	Slightly higher freshness but the same actual service level compared to DE
	Superior to UL_{L-1} even when UL_{L-1} dominates DE

DI, we can add the term e_{ti} defined by Equation (32) hereunder to the profit function given in Equation (7):

$$e_{ti} = -(h + h')I_{ti} + h'(L-1)\sum_{d=0}^{x_{t1}-1} Pr(D_{ti} = d)(x_{t1} - d). \qquad (32)$$

The first term in (32) charges the total inventory at the beginning of period t with costs h and h', since this inventory is carried from the previous period. The second term cancels out the charged costs h' during $L-1$ periods for the units which are completely deteriorated at the end of period t.

11. Conclusions

By considering uncertainty and combining inventory with routing decisions for perishable products, retail chains can obtain a significant increase in their net profit. We have shown how such benefit can be gained and we have quantified it.

The expected value method, where only the expected demands are taken into consideration in retail chain's decisions, serves as a benchmark. We then show how the knowledge of the demand distribution can add to the profit. To this end, we first propose a simple up-to-level method which explicitly takes the target service level into account. Our numerical results show that this naive policy performs reasonably well when the target service level is high. Next, a decomposition method is applied to determine delivery quantity to each store independently. Assigning virtual costs-to-serve to stores whenever they are visited accounts for some aspects of routing in the method. This leads to a significant increase in profit. Finally, we integrate the decisions independently made by each store, and we slightly divert from the latter delivery quantities with the aim to decrease the routing costs. Though the routing costs only comprise a small portion of the total costs, we showed that the final improvement in total profit is statistically and economically significant.

Our approach considers the real (expected) routes for only two periods ahead in the decomposition-integration method. This is justifiable when routing decisions cannot be made for a large number of periods and deliveries cannot be synchronized to be carried out in the same periods. This is the case when (1) demands are highly stochastic, and (2) shelf life is short or store capacity is limited for long-term deliveries. At last, we show how further profit improvement is possible when accessing full information about the future demands.

Acknowledgments

Appendix A: Cost-to-serve estimation

Distance-based cost-to-serve. The first approach to estimate costs-to-serve looks at the average distance between each store and its "closest neighbors". Defining J_i as a set of stores near store i, a distance-based cost-to-serve for store i is calculated by Equation (33):

$$F_{ti}^d = \frac{\sum_{j \in J_i} c_{ij}}{|J_i|} \tag{33}$$

Observe that F_{ti}^d does not depend on t. Our experimental results show that the size of the set J_i should increase with the maximum shelf life L. The reason is that for large L, stores are served less frequently. So, there is a smaller chance to serve store i and its nearest neighbors in the same period, and, by way of consequence, it is more likely that store i will be served together with some of its farther neighbors. We set $|J_i| = 2L$ in our experiments.

Route-based cost-to-serve. Our second approach is inspired by the work of Özener et al. (2013). These authors introduce several methods to allocate a cost to each store in an IRP. Based on the same underlying concepts, we assign the whole cost of a route to the stores it includes. However, we have to estimate the cost of the routes before solving any IRP. This can be done by calculating and comparing the average routing cost plus the average deterioration cost for different frequencies of deliveries. Assuming that the ideal periodicity of delivery is λ periods, that store i is served in period t, and that the store capacity is large enough, the delivery quantity in period t to store i can be estimated by Equation (34):

$$\alpha_{ti} = E(D_{ti}) + \ldots + E(D_{t+\lambda-1,i}). \tag{34}$$

Then, the average delivery quantity in period t to store i and to its neighbors is

$$\bar{\alpha}_{ti} = \frac{\alpha_{ti} + \sum_{j \in J_i} \alpha_{tj}}{1 + |J_i|}. \tag{35}$$

Given the average delivery quantity $\bar{\alpha}_{ti}$ and vehicle capacity Q, we approximate the average number of stores included in the route serving store i in period t as

$$\bar{n}_{ti} = \frac{Q}{\bar{\alpha}_{ti}}, \tag{36}$$

36

and the cost of the route serving i as

$$R_{ti} = 2c_{i0} + (\bar{n}_{ti} - 1)F_{ti}^d. \tag{37}$$

Finally, the portion of the estimated cost of the route allocated to store i is

$$F_{ti}^r = \frac{\alpha_{ti}}{\alpha_{ti} + \sum_{j \in J_i} \alpha_{tj}} \cdot R_{ti}. \tag{38}$$

In our experiments, we determined that the best setting for the frequency of deliveries is $\lambda = L - 1$. We also observed that the costs F_{ti}^r computed by this approach are high as compared to the former costs F_{ti}^d.

Appendix B: Maximum decrease and increase in delivery quantities.

Following the notations given in Section 7, here we determine the maximum decrease (resp., increase) in delivery quantity to the stores in D (resp., I), i.e., we determine $\underline{m}_i$ for $i \in D$ (resp., $\overline{m}_i$ for $i \in I$). This helps us to restrict the number of decision variables in the IP formulation. Let us define q_{tr} as the current load on route r in period t, and $q_t(i)$ as the current load on the route which includes store i in period t.

The delivery quantity to store $i \in D$ can decrease as long as it respects TSL, i.e., $\underline{m}_i \leq y_{ti} - y_{ti}^{(1)}$, where $y_{ti}^{(1)}$ is, as before, the smallest integer delivery quantity satisfying Inequality (4). Moreover, there is no need to decrease the delivery quantity to store $i \in D$ by more than the excess load on route r', i.e., $\underline{m}_i \leq q_{tr'} - Q$. The decrease must not cause vehicle load violation in any of the routes in $\bar{R}_{t+1}$. In order to analyze the latter constraint, consider two cases. In the first case, store $i \in D$ is not included in any route in $\bar{R}_{t+1}$, i.e., $i \in D \backslash N(\bar{R}_{t+1})$. For such a store, the delivery quantity in period t, y_{ti}, can decrease as long as the expected delivery quantity in period $t+1$, $(\tilde{y}_{t+1,i}|y_{ti})$, remains zero; otherwise, the routing costs in period $t+1$ would increase. This translates into $\underline{m}_i \leq \min_{0 \leq y \leq y_{ti}}\{y \mid (\tilde{y}_{t+1,i}|y) = 0\}$. Hence, for every store $i \in D \backslash N(\bar{R}_{t+1})$, the maximum decrease of the delivery quantity y_{ti} is determined as:

$$\underline{m}_i = \min\{y_{ti} - y_{ti}^{(1)}, \ q_{tr'} - Q, \ y_{ti} - \min_{0 \leq y \leq y_{ti}}\{y \mid (\tilde{y}_{t+1,i}|y) = 0\}\}. \tag{39}$$

The second case considers stores $i \in D$ which are also served in period $t+1$, i.e., stores $i \in D \cap N(\bar{R}_{t+1})$. A similar reasoning about the necessity of respecting TSL and the uselessness of decreasing a delivery quantity more than the excess load on r' applies for these stores and lead to the same constraints as in the previous case. Constraints (23) in the IP formulation guarantee that a decrease of the delivery quantity to $i \in D \cap N(\bar{R}_{t+1})$ does not cause any vehicle capacity violation in period $t+1$. Therefore, for all stores $i \in D \cap N(\bar{R}_{t+1})$, the maximum decrease of delivery quantity is simply determined as

$$\underline{m}_i = \min\{y_{ti} - y_{ti}^{(1)}, \ q_{tr'} - Q\}. \tag{40}$$

A maximum increase of delivery quantity, say $\overline{m}_i$, for any store $i \in I$ can also be determined. On one hand, the increase cannot be so high as to exceed the vehicle capacity, i.e., $\overline{m}_i \leq Q - q_t(i)$ must hold. On the other hand, store capacities must be respected, i.e., $\overline{m}_i \leq C_i - I_{ti} - y_{ti}$. As a result, the maximum increase of delivery quantity to any store $i \in I$ is determined as

$$\overline{m}_i = \min\{Q - q_t(i), C_i - I_{ti} - y_{ti}\}. \tag{41}$$

37

References

Adachi Y, Nose T, Kuriyama S (1999) Optimal inventory control policy subject to different selling prices of perishable commodities. *International Journal of Production Economics*, 60:389-394.

Adelman D (2004) A price-directed approach to stochastic inventory routing. *Operations Research*, 52(4):499-514.

Alegre J, Laguna M, Pacheco J (2007) Optimizing the periodic pickup of raw materials for a manufacturer of auto parts. *European Journal of Operational Research*, 179(3):736-746.

Al Shamsi A, Al Raisi A, Aftab M (2014) Pollution-inventory routing problem with perishable goods. Golinska P, ed., *Logistics Operations*, Supply Chain Management and Sustainability, EcoProduction, Springer, Switzerland, 585-596.

Andersson H, Hoff A, Christiansen M, Hasle G, Løkketangen A (2010) Industrial aspects and literature survey: combined inventory management and routing. *Computers & Operations Research*, 37:1515-1536.

Archetti C, Bertazzi L, Laporte G, Speranza M G (2007) A branch-and-cut algorithm for a vendor managed inventory routing problem. *Transportation Science*, 41:382-391.

Archetti C, Speranza M G (2016) The inventory routing problem: the value of integration. *International Transactions in Operational Research*, 23(3):393-407.

Barnes-Schuster D, Bassok Y (1997) Direct shipping and the dynamic single-depot multi-retailer inventory system. *European Journal of Operational Research*, 101(3):509-518.

Berk E, Gurler U (2008) Analysis of the (Q,r) inventory model for perishables with positive lead times and lost sales. *Operations Research*, 56(5):1238-1246.

Berman O, Larson R C (2001) Deliveries in an inventory routing problem using stochastic dynamic programming. *Transportation Science*, 35(2):192-213.

Bertazzi L (2008) Analysis of direct shipping policies in an inventory routing problem with discrete shipping times. *Management Science*, 54(4):748-762.

Bertazzi L, Paletta G, Speranza M G (2002) Deterministic order-up-to level policies in an inventory routing problem. *Transportation Science*, 36:119-132.

Bertazzi L, Speranza M G (2011) Matheuristics for inventory routing problems. In: Montoya-Torres J R, Juan A A, Huatuco L H, Faulin J, Rodriguez-Verjan G L. *Hybric algorithms for service, computing, and manufacturing systems: routing and scheduling solutions*, IGI Global, Hershey, 1-14.

Bertazzi L, Speranza M G (2012) Inventory routing problem: an introduction. *EURO Journal on Transportation and Logistics*, 1:307-326.

Bertazzi L, Speranza M G (2013) Inventory routing problem with multiple customers. *EURO Journal on Transportation and Logistics*, 2:255-275.

Broekmeulen R, Van Donselaar K (2009) A heuristic to manage perishable inventory with batch ordering, positive lead times, and time-varying demand. *Computers & Operations Research*, 36:3013-3018.

Campbell A M, Clarke L, Kleywegt A J, Savelsbergh M W P (1998) The inventory routing problem. Crainic T G, Laporte G, eds. *Fleet Management and Logistics*, Springer, Boston, 95-113.

Campbell A M, Savelsbergh M (2004) Delivery volume optimization. *Transportation Science*, 38(2):210-223.

Campbell A M, Savelsbergh M (2004) A decomposition approach for the inventory routing problem. *Transportation Science*, 38(4):488-502.

Campbell A M, Wilson J H (2014) Forty years of periodic vehicle routing. *Networks*, 63(1):2-15.

Chan M L A, Federgruen A, Simchi-Levi D (1998) Probabilistic analyses and practical algorithms for inventory routing models. *Operations Research*, 46(1):96-106.

Chao X, Gong X, Shi C, Zhang H (2015) Approximation algorithms for perishable inventory systems. *Operations Research*, 63(3):585-601.

Chiu H N (1995) A heuristic (R,T) periodic review perishable inventory model with lead times. *International Journal of Production Economics*, 42:1-15.

Christiansen M, Fagerholt K, Flatberg T, Haugen Ø, Kloster O, Lund EH (2011) Maritime inventory routing with multiple products: a case study from the cement industry. *European Journal of Operational Research*, 208(1):86-94.

Coelho L C, Cordeau J F, Laporte G (2012) The inventory routing problem with transshipment. *Computers & Operations Research*, 39:2537-2548.

Coelho L C, Cordeau J F, Laporte G (2014) Thirty years of inventory routing. *Transportation Science*, 48(1):1-19.

Coelho L C, Cordeau J F, Laporte G (2014(b)) Heuristics for dynamic and stochastic inventory routing. *Computers & Operations Research*, 52:55-67.

Coelho L C, Laporte G (2013) A branch-and-cut algorithm for the multi-product multi-vehicle inventory-routing problem. *International Journal of Production Research*, 51:7156-7169.

Coelho L C, Laporte G (2013) The exact solution of several classes of inventory routing problems. *Computers & Operations Research*, 40:558-565.

Coelho L C, Laporte G (2014) Optimal joint replenishment, delivery and inventory management policies for perishable products. *Computers & Operations Research*, 47:42-52.

Coelho L C, Laporte G (2014) Improved solutions for inventory routing problems through valid inequalities and input ordering. *International Journal of Production Economics*, 155:391-397.

Cooper W L (2001) Pathwise properties and performance bounds for a perishable inventory system. *Operations Research*, 49(3):455-466.

Deniz B, Karaesmen I, Scheller-Wolf A (2010) Managing perishables with substitution: Inventory issuance and replenishment heuristics. *Manufacturing & Service Operations Management*, 12(2):319-329.

Desaulniers G, Rakke J G, Coelho L C (2016) A branch-price-and-cut algorithm for the inventory routing problem. *Transportation Science*, 50(3):1060-1076.

Ekici A, Ozener O O, Kuyzu G (2015) Cyclic delivery schedules for an inventory routing problem. *Transportation Science*, 49(4):817-829.

Federgruen A, Zipkin P (1984) A combined vehicle routing and inventory allocation problem. *Operations Research*, 32(5):1019-1037.

Federgruen A, Prastacos G, Zipkin PH (1986) An allocation and distribution model for perishable products. *Operations Research*, 34(1):75-82.

Gallego G, Simchi-Levi D (1990) On the effectiveness of direct shipping strategy for the one-warehouse multi-retailer R-systems. *Management Science*, 36:240-243.

Gaur V, Fisher ML (2004) A periodic inventory routing problem at a supermarket chain. *Operations Research*, 52(6):813-822.

Goyal S K, Giri B C (2001) Recent trends in modeling of deteriorating inventory. *European Journal of Operational Research*, 134:1-16.

Gronhaug R, Christiansen M, Desaulniers G, Desrosiers J (2010) A branch-and-price method for a liquified natural gas inventory routing problem. *Transportation Science*, 44(3):400-415.

Hemmelmayr V, Doerner K F, Hartl R F, Savelsbergh M W P (2009) Delivery strategies for blood products supplies. *OR Spectrum*, 31:707-725.

Hemmelmayr V, Doerner K F, Hartl R F, Savelsbergh M W P (2010) Vendor managed inventory for environments with stochastic product usage. *European Journal of Operational Research*, 202:686-695.

Hsieh T P, Dye C Y (2010) Optimal replenishment policy for perishable items with stock-dependent selling rate and capacity constraint. *Computers & Industrial Engineering*, 59:251-258.

Hsu V N (2000) Dynamic economic lot size model with perishable inventory. *Management Science*, 46(8):1159-1169.

Huang S H, Lin P C (2010) A modified ant colony optimization algorithm for multi-item inventory routing problems with demand uncertainty. *Transportation Research Part E*, 46:598-611.

Jaillet P, Bard J F, Huang L, Dror M (2002) Delivery cost approximations for inventory routing problems in a rolling horizon framework. *Transportation Science*, 36(3):292-300.

Karaesmen I Z, Scheller-Wolf A, Deniz B (2011) Managing perishable and aging inventories: Review and future research directions. Kempf K G, Keskinocak P, Uzsoy R, eds. *Planning Production and Inventories in the Extended Enterprise*, International Series in Operations Research and Management Science. Springer, New York, 151:393-436.

Kleywegt A J, Nori V S, Savelsbergh M W P (2002) The stochastic inventory routing problem with direct deliveries. *Transportation Science*, 36(1):94-118.

40

Kleywegt A J, Nori V S, Savelsbergh M W P (2004) Dynamic programming approximations for a stochastic inventory routing problem. *Transportation Science*, 38(1):42-70.

Kouki C, Jemai Z, Minner S (2015) A lost sales (r,Q) inventory control model for perishables with fixed lifetime and lead time. *International Journal of Production Economics*, 168:143-157.

Kouki C, Jouini O (2015) On the effect of lifetime variability on the performance of inventory systems. *International Journal of Production Economics*, 167:23-34.

Kouki C, Sahin E, Jemai Z, Dallery Y (2014) Analysis of a periodic review inventory control system with perishables having random lifetime. *International Journal of Production Research*, 52(1):283-298.

Le T, Diabat A, Richard J P, Yih Y (2013) A column generation-based heuristic algorithm for an inventory routing problem with perishable goods. *Optics Letters*, 7:1481-1502.

Lian Z, Liu L (1999) A discrete-time model for perishable inventory systems. *Annals of Operations Research*, 87:103-116.

Minkoff A S (1993) A Markov decision model and decomposition heuristic for dynamic vehicle dispatching. *Operations Research*, 41(1):77-90.

Minner S, Transchel S (2010) Periodic review inventory control for perishable products under service level constraints. *OR Spectrum*, 32:979-996.

Mirzaei S, Seifi A (2015) Considering lost sale in inventory routing problems for perishable goods. *Computers & Industrial Engineering*, 87:213-227.

Nahmias S (2011) Perishable inventory systems. Hillier F S, ed. *International Series in Operations Research and Management*, (Springer, NewYork), 1-80.

Özener Ö, Ergun Ö, Savelsbergh M W P (2013) Allocating cost of service to customers in inventory routing. *Operations Research*, 60(5):1-14.

Popović D, Vidović M, Radivojević G (2012) Variable neighborhood search heuristic for the inventory routing problem in fuel delivery. *Expert Systems Applications*, 39(18):13390-13398.

Powell W E (2011) *Approximate dynamic programming: solving the curses of dimensionality*, (Wiley, New Jersey).

Puterman M L (1994) *Markov Decision Processes.* John Wiley & Sons, New York.

Qu W W, Bookbinder J H, Iyogun P (1999) An integrated inventory transportation system with modified periodic policy for multiple products. *European Journal of Operational Research*, 115:254–269.

Raa B, Aghezaaf E H (2008) Designing distribution patterns for long-term inventory routing with constant demand rates. *International Journal of Production Economics*, 112:255-263.

Raa B, Aghezaaf E H (2009) A practical solution approach for the cyclic inventory routing problem. *European Journal of Operational Research*, 192:429-441.

Reinman M I, Rubio R, Wein L M (1999) Heavy traffic analysis of the dynamic stochastic inventory routing problem. *Transportation Science*, 33(4):361-380.

Righini G, Salani M (2006) Symmetry helps: Bounded bi-directional dynamic programming for the elementary shortest path problem with resource constraints. *Discrete Optimization*, 3:255-273.

Schwartz LB, Ward JE, Zhai X (2006) On the interactions between routing and inventory management policies in a one-warehouse N-retailer distribution system. *Manufacturing & Service Operations Management*, 8:253-272.

Solomon M M (1987) Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, 35(2):254-265.

Solyali O, Süral H (2011) A branch-and-cut algorithm using a strong formulation and an a priori tour-based heuristic for an inventory routing problem. *Transportation Science*, 45(3):335-345.

Soysal M, Bloemhof-Ruwaard J M, Haijema R, Van der Vorst J (2015) Modeling an inventory routing problem for perishable products with environmental considerations and demand uncertainty. *International Journal of Production Economics*, 164:118-133.

Stacey J, Natarajarathinam M, Sox C (2007) The storage constrained inbound inventory routing problem. *International Journal of Physical Distribution & Logistics Management*, 37(6):484-500.

Yu Y, Chu C, Chen H, Chu F (2012) Large scale stochastic inventory routing problems with split delivery and service level constraints. *Annals of Operations Research*, 197(1):135-158.

EBRD–FAO (2009) *Agribusiness Handbook: Food Retail*, European Bank for Reconstruction and Development and Food and Agriculture Organization of the United Nations, *http://www.fao.org/fileadmin/user_upload/tci/docs/AH4-Food%20Retail.pdf*.

FMI (2014) *Supermarket Facts*, Food Marketing Industry, *http://www.fmi.org/research-resources/supermarket-facts*.

NAICS (2012) *Canadian Industry Statistics on Grocery Stores* (NAICS 4451), Government of Canada, *https://www.ic.gc.ca/app/scr/sbms/sbb/cis/revenues.html?code=4451&lang=eng*.