Stochastic inventory routing for perishable products

Citation for published version (APA):

Crama, Y., Rezaei, M., Savelsbergh, M. W. P., & van Woensel, T. (2018). Stochastic inventory routing for perishable products. *Transportation Science*, 52(3), 526-546. <https://doi.org/10.1287/trsc.2017.0799>

Document license:
TAVERNE

DOI:
[10.1287/trsc.2017.0799](https://doi.org/10.1287/trsc.2017.0799)

Document status and date:
Published: 01/05/2018

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Transportation Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Stochastic Inventory Routing for Perishable Products

Yves Crama, Mahmood Rezaei, Martin Savelsbergh, Tom Van Woensel

To cite this article:

Yves Crama, Mahmood Rezaei, Martin Savelsbergh, Tom Van Woensel (2018) Stochastic Inventory Routing for Perishable Products. *Transportation Science* 52(3):526-546. <https://doi.org/10.1287/trsc.2017.0799>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2018, INFORMS

Please scroll down for article—it is on subsequent pages

INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Stochastic Inventory Routing for Perishable Products

Yves Crama,^a Mahmood Rezaei,^a Martin Savelsbergh,^b Tom Van Woensel^c

^a HEC Management School, University of Liège, 4000 Liège, Belgium; ^b School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332; ^c School of Industrial Engineering, Eindhoven University of Technology, 5612 AZ Eindhoven, Netherlands

Contact: yves.crama@uliege.be, <http://orcid.org/0000-0002-7470-4743> (YC); m.rezaei@uliege.be (MR); martin.savelsbergh@isye.gatech.edu (MS); t.v.woensel@tue.nl (TVW)

Received: June 28, 2016

Revised: July 7, 2017

Accepted: August 5, 2017

Published Online in Articles in Advance:
February 6, 2018

<https://doi.org/10.1287/trsc.2017.0799>

Copyright: © 2018 INFORMS

Abstract. Different solution methods are developed to solve an inventory routing problem for a perishable product with stochastic demands. The solution methods are empirically compared 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 although a simple ordering policy can achieve very good performance, 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.

History: Michel Gendreau served as the guest editor-in-chief for this article.

Funding: The first two authors have been partially supported in the course of this research by the Interuniversity Attraction Poles Program initiated by the Belgian Science Policy Office [Grant COMEX P7/36].

Keywords: inventory routing problem • perishable inventory system • perishable inventory routing • stochastic inventory routing

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

1.1. Motivation

In an IRP, 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 (2014a) report that instances of a single-product IRP with deterministic demands can rarely be solved to optimality when the number of stores exceeds 30. By 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 about 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. Because of 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); roughly 10% of product is wasted before being sold (Kouki, Jemai, and Minner 2015), while the food retail profit margin hardly exceeds 2%; see EBRD–FAO (2009), Government of Canada (2012), and FMI (2014). Therefore, 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. 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 implies 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 a consequence of the way the vehicles are dispatched. Miscellaneous costs are viewed as independent of either decision and are regarded as constant. As a result, we do not consider them in the objective function: The net profit is simply computed as (Revenue, Acquisition costs, and 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 (i.e., 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 $X_{ti} = (x_1, x_2, \dots, x_{L-1})_{ti}$, 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} 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 (first-in-first-out (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 stock-out 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 United States (Minner and Transchel 2010).

All notation is summarized in Table 1.

1.3. Scientific Contributions

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

- 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.

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
$(x_1, x_2, \dots, x_{L-1})_{ti}$	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

• 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.

• 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 about future demand is available; this information will

be used for comparison purposes in the computation study. Section 9 discusses the set-up of our comprehensive computational study. Results of the computational experiments are presented and analyzed in Section 10. 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 information sharing between a supplier and its customers. When VMI is implemented, the supplier takes over the responsibility of managing the customers' inventory by deciding on replenishment quantities and delivery periods. The consequences can be beneficial for both parties: Customers can use fewer resources to control their inventory, and the supplier has more flexibility for integrating the replenishment quantities and periods to different customers (Desaulniers, Rakke, and Coelho 2016). By contrast, in an RMI system, the customers decide when and how much to order, independent 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).

2.1. Inventory-Routing

The (classic) 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, Vidović, and Radivojević 2012), automobile components (Alegre, Laguna, and Pacheco 2007; Stacey, Natarajarathinam, and Sox 2007), perishable products (Federgruen and Zipkin 1984; Federgruen, Prastacos, and Zipkin 1986), grocery products (Gaur and Ficher 2004), cement (Christiansen et al. 2011), and blood products (Hemmelmayr et al. 2009).

Because 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 enormous even when compared, for instance, with classic vehicle routing problems (VRPs). 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, Rakke, and Coelho 2016). Yet 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, i.e., (a) inventory control (determining the delivery amounts), and (b) vehicle routing (Campbell and Savelsbergh 2004b; Federgruen and Zipkin 1984; Qu, Bookbinder, and Iyogun 1999). In some cases, an overall solution is found by iterating between these two problems (Federgruen and Zipkin 1984; Qu, Bookbinder, and Iyogun 1999). Exact algorithms for the IRP are more recent. They include the branch-and-cut algorithms of Coelho and Laporte (2013a, b, 2014d), and the branch-and-price-and-cut algorithm of Desaulniers, Rakke, and Coelho (2016).

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

- deliver-up-to-level (UL) replenishment policies; see, e.g., Bertazzi, Paletta, and Speranza (2002); Archetti et al. (2007); and Solyali and Süral (2011);
- direct delivery policies; see, e.g., Bertazzi and Speranza (2012);
- deliveries occurring only when inventory levels are down to zero; see, e.g., Chan, Federgruen, and Simchi-Levi (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; see, e.g., Raa and Aghezaaf (2008, 2009) and Ekici, Ozener, and Kuyzu (2015);
- periodic deliveries; see, e.g., Bertazzi and Speranza (2012) and Campbell and Wilson (2014).

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

2.2. Stochastic Inventory Routing

In the 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, Nori, and Savelsbergh 2004), and is rarely

dealt with as backlogging, as in the work by Yu et al. (2012). In either case, penalties may apply. A predefined 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. Because of 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 2014a; Reinman, Rubio, and Wein 1999; Schwartz, Ward, and Zhai 2006), a single uncapacitated vehicle (Qu, Bookbinder, and Iyogun 1999) or direct deliveries (Barnes-Schuster and Bassok 1997; Kleywegt, Nori, and Savelsbergh 2002; Reinman, Rubio, and Wein 1999).

Markov decision processes (MDP) can be used to model SIRPs over an infinite planning horizon. MDP models formulate a value function that 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 the optimization problem can be efficiently solved. 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 VRP as a special case (Campbell et al. 1998). Because of the curse of dimensionality, researchers often use approximations of the value function (Minkoff 1993, Adelman 2004) or decompose it (Kleywegt, Nori, and Savelsbergh 2004). Based on the linear programming (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 provide upper bounds on the original dual maximization problem. Kleywegt, Nori, and Savelsbergh (2004) formulate an MDP model of the SIRP and propose approximation methods to find good solutions in reasonable time. This is the extension of an earlier paper (Kleywegt, Nori, and Savelsbergh 2002) in which the authors formulated an SIRP with direct deliveries as an MDP and proposed an approximate dynamic programming approach for its solution.

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 VRP.

2.3. Inventory-Routing for Perishables

The IRP for perishables (PIRP) is identical to the classic 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 classic IRP is that the deliveries to customers are now restricted by the maximum shelf life of the product. Consequently, 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 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 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, i.e., (1) using fixed routes but deciding about delivery days, and (2) repeating delivery patterns for each hospital. Coelho and Laporte (2014a) consider age-dependent holding costs and selling prices in a PIRP, where the supplier has the choice to deliver 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, Al Raisi, and Aftab (2014) to include the cost of CO₂ emissions, based on the vehicle load and distance. The resultant model is a mixed integer linear programming problem that 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, independent 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 classic techniques is in general intractable because of the curse of dimensionality. Thus, many researchers turned to effective heuristic policies to address 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, Karaesmen, and Scheller-Wolf 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, Nose, and Kuriyama (1999); Broekmeulen and Van Donselaar (2009); and Minner and Transchel (2010). See Goyal and Giri (2001); Karaesmen, Scheller-Wolf, and Deniz (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 further complicates the model. As a result, applying existing solution methods is mathematically very difficult and computationally very inefficient.

3. The Expected Value Method

A classic 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 2014a), 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 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 $(x_1, x_2, \dots, x_{L-1})_{ti}$, so that the total inventory level is

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

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

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

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

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 $y_{ti} = \min\{C_i - I_{ti}, [E(D_{ti}) + \dots + E(D_{t+L-1,i})] - I_{ti}\}$.

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 , i.e., 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 a (R, s, S) policy, where $R = 1$, $s = E(D_{ti})$, and $S = \min\{C_i, [E(D_{ti}) + \dots + E(D_{t+L-1,i})]\}$.

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 that suffices to meet the demand at store i during λ consecutive periods $\{t, \dots, t + \lambda - 1\}$ with probability TSL

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

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 $\Pr(D_{ti} \leq I_{ti}) \geq \text{TSL}$, 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 , i.e., 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} that 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 a (R, s, S) policy where $s_{\text{UL}} = q_{ti}^{(\lambda)}$, $S_{\text{UL}} = \min\{C_i, q_{ti}^{(\lambda)}\}$, 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 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 accounts 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., $((x_1, \dots, x_{L-1})_{t1}, \dots, (x_1, \dots, x_{L-1})_{tN})$. The decision variables are the delivery quantities in period t , i.e., $(y_{t1}, \dots, y_{tN})$, and the routing decisions. Given a

decision on delivery quantities in period t , the direct costs (acquisition and routing) and 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 suboptimal 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).

Because 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 at the end of this section.

Define $X_{ti} = ((x_1)_{ti}, \dots, (x_{L-1})_{ti})$, $X_t = (X_{t1}, \dots, X_{tN})$, and $Y_t = (y_{t1}, \dots, y_{tN})$, so that (X_t, Y_t) denotes the complete state of the system at time t after the quantities $y_{t1}, \dots, y_{tN}$ 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,

$$f_{ti}^*(X_{ti}) = \max_{y_{ti}^{(0)} \leq y_{ti} \leq C_i - I_{ti}} f_{ti}(X_{ti}, y_{ti}), \quad (5)$$

where $y_{ti}^{(0)}$ is the smallest integer satisfying Inequality (4). The optimal delivery quantity is specified by Equation (6)

$$y_{ti}^* = y_{ti}^*(X_{ti}) = \arg \max_{y_{ti}^{(0)} \leq y_{ti} \leq C_i - I_{ti}} f_{ti}(X_{ti}, y_{ti}). \quad (6)$$

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

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

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. To solve (7), we use the boundary condition

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

where the right-hand side (RHS) of (8) is an estimate of the profit generated by the inventory left 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 , i.e., 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 , i.e., 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 DE₀. 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 that 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 that 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 , i.e., 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

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

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

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

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:

$$\begin{aligned}\pi_t^3(X_t, Y_t) &= \sum_{i=1}^N f_{ti}((x_1, \dots, x_{L-1})_{ti}, y_{ti}) \\ &\quad + \sum_{i=1}^N F_{ti} \cdot \mathbb{I}(y_{ti} > 0) - R(y_{t1}, \dots, y_{tN}) \\ &\quad + \sum_{i=1}^N F_{ti} \cdot \Pr(y_{t+1,i}^* > 0 | Y_t)\end{aligned}$$

$$\begin{aligned}&- \sum_{(y_1, \dots, y_N)} \Pr((y_{t+1,1}^*, \dots, y_{t+1,N}^*) \\ &= (y_1, \dots, y_N) | Y_t) \times R(y_1, \dots, y_N).\end{aligned}\quad (11)$$

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)):

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

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

Replace here the random quantities $y_{t+1,i}^*$ by $[E(y_{t+1,i}^* | y_{ti})]$ 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 a drawback, however, of yielding strictly positive values $[E(y_{t+1,i}^* | y_{ti})]$ 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 ($\epsilon_i = \frac{1}{2}E(D_{ti})$ proved suitable in our numerical experiments)

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

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

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

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 , i.e., 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 Section 7, 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 $Y_t = (y_{t1}^*, \dots, y_{tN}^*)$ proposed by Equation (6), we generate a new feasible solution $Y'_t = (y'_{t1}, \dots, y'_{tN})$ 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 predefined number of consecutively generated new solutions are rejected due to lack of improvement or infeasibility. Calculating $\tilde{\pi}$ in Problem (14) for any new solution involves solving two VRPs. To avoid these expensive computations, each new solution is not randomly generated but in a systematic way that 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)

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

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 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

$$\delta = R(y_{t1}, \dots, y_{tN}) - R(y'_{t1}, \dots, y'_{tN}). \quad (16)$$

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

$$\begin{aligned} \Delta &= \sum_{i=1}^N (f_{ti}((x_1, x_2, \dots, x_{L-1})_{ti}, y'_{ti}) \\ &\quad - f_{ti}((x_1, x_2, \dots, x_{L-1})_{ti}, y_{ti})) \\ &\quad + \delta + \sum_{i=1}^N F_{ti} \cdot (\mathbb{I}(y'_{ti} > 0) - \mathbb{I}(y_{ti} > 0)). \end{aligned} \quad (17)$$

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. In general, however, since the routes were optimally selected for the delivery quantities Y_t , route r' will be infeasible with respect to its maximum allowed length or 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 $(\tilde{Y}_{t+1} | Y_t)$ that are likely to increase). To keep some control over this effect, therefore, we restrict ourselves to certain modifications of Y_t that 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, e.g., $R = \{r\}$, 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 to restore feasibility of route r' (D is for “decrease”);

- $\bar{R}_{t+1} = \{r \in R_{t+1} | D \cap N(r) \neq \emptyset\}$ is the set of routes in period $t + 1$ that contain at least one store in D ; these are the routes in period $t + 1$ that 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. 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 $\bar{R}_t$ and I ;

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

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

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

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

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

$$\text{subject to } \sum_{h=0}^{\underline{m}_i} v_{ih} = 1, \quad \forall i \in D, \quad (19)$$

$$\sum_{h=0}^{\bar{m}_i} v_{ih} = 1, \quad \forall i \in I, \quad (20)$$

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

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

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

$$v_{ih} \in \{0, 1\}, \quad \forall i \in D, h \in [0, \underline{m}_i] \text{ and } \forall i \in I, h \in [0, \bar{m}_i]. \quad (24)$$

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 stores in D and I are calculated using Equations (25) and (26), respectively. Delivery quantities to other stores do not change

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

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

Note that only the routes belonging to $\bar{R}_t$ or $\bar{R}_{t+1}$ appear in the IP formulation. Moreover, 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 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 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, in summary, our local search approach to Problem (14) acts as a “large neighborhood search” framework that explores the neighborhood of the current solution by solving the IP subproblem (18)–(24). The solution of (18)–(24) hopefully yields new delivery quantities that 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 $(\tilde{y}_{t+1,i} \mid y_{ti})$ 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' that 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.

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 that 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, denoted 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, 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.

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 $y_{ti} = d_{ti} + \dots + d_{t+\lambda-1,i}$, 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 $y_{t+1,i} = \dots = y_{t+\lambda-1,i} = 0$. 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 1.8 GHz CPU and 8 GB RAM. No time limit is imposed on 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., $D_{ti} \sim \text{Bin}(200, 0.1)$. The average demand is $E(D_{ti}) = 20$ for each period and for each store. We consider three shelf lives, i.e., $L \in \{2, 3, 4\}$. 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). In Section 10.5, we will 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

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 consists of initial inventory as well as demands of the stores over the planning horizon. We use the same scenarios for all solution methods and 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. We 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)

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

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

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

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)

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

Second, 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

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

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 $t + 1$, i.e., R_{t+1} .

We define the degree of similarity between these sets by Equation (31)

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

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 life $L \in \{2, 3, 4\}$. The second column indicates the solution methods applied to determine delivery quantities and routes for each scenario, i.e., 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 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.

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 in 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).

10. Discussion

In this section, 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

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.

We 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 to DI. Figure 1(a) 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 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 1(a) also shows the profit gained by applying UL_{L-1} . Interestingly, the

Figure 1. (Color online) Contribution of Different Solution Methods to Increasing Profit

performance of this simple deliver-up-to-level policy is very close to the performance of the more sophisticated method DE. Figure 1(b) can be interpreted in the same way as Figure 1(a), 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, i.e., the quantity given by Equation (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 on 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 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.

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 rigor-

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 (%)
2	EV	4	1.000 (66,542)	3.598	2.331	0.267	0.149	8.5	3.2	1.6	1.6	1.4	91	98
	UL ₁	1,626	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}	1,626	1.126	3.596	2.175	0.295	0.008	7.7	5.6	1.0	1.7	1.7	93	99
	DE ₀	1,802	1.129	3.623	2.196	0.297	0.011	7.8	5.5	1.0	1.7	1.6	97	99
	DE _d	1,484	1.130	3.623	2.196	0.297	0.011	7.8	5.5	1.0	1.7	1.6	97	99
	DE _r	1,482	1.130	3.623	2.196	0.297	0.011	7.8	5.3	1.0	1.7	1.6	97	99
	DI ₀	8,783	1.133	3.622	2.196	0.293	0.011	7.5	5.5	1.0	1.7	1.6	97	99
	DI _d	7,003	1.133	3.621	2.196	0.293	0.011	7.6	5.7	1.0	1.7	1.6	97	99
	DI _r	7,188	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
3	EV	3	1.000 (72,320)	3.376	2.164	0.212	0.103	8.5	2.0	2.6	2.3	1.9	95	99
	UL ₁	1,673	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
	DE ₀	1,575	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	7,697	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
4	EV	0	1.000 (72,182)	3.418	2.182	0.236	0.084	11.7	1.0	3.6	3.0	2.5	97	99
	UL ₁	1,672	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
	DE ₀	2,568	1.054	3.364	2.036	0.275	0.000	7.7	5.5	1.0	3.6	3.3	99	99
	DE _d	1,086	1.111	3.425	2.109	0.205	0.012	8.2	1.9	2.8	3.1	2.5	99	99
	DE _r	1,107	1.106	3.431	2.121	0.204	0.018	8.2	1.4	2.9	3.0	2.4	99	99
	DI ₀	7,716	1.070	3.362	2.039	0.253	0.005	7.4	5.3	1.1	3.5	3.2	99	99
	DI _d	1,365	1.111	3.424	2.104	0.210	0.007	7.9	2.6	2.2	3.2	2.7	99	99
	DI _r	1,597	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

ous optimality gap, since FI assumes perfect information and relies on a heuristic algorithm. Yet the value of the gap suggests that DI is performing reasonably well.

Second, consider the generic expression of the expected profit as $E(\text{Profit}) = E(\text{Revenue} - \text{Acquisition}) - E(\text{Routing})$. Clearly, the revenue of the retail chain is maximized when all stores can 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. 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 little waste (around 1%–2%). This already implies that $E(\text{Revenue} - \text{Acquisition})$ is almost best possible. A more precise calculation can be carried out as follows: In each period t and for every store i , $E(\text{Revenue}_{ti} - \text{Acquisition}_{ti}) \leq E(D_{ti})(p - a)$. In our experiments, $E(D_{ti})(p - a) = 80$ so that, with 30 periods and 40 stores, $E(\text{Revenue}) - E(\text{Acquisition})$ is bounded by 96,000. For each shelf life $L = 2, 3, 4$, respectively, the average value of $(\text{Revenue} - \text{Acquisition})$ achieved by DI_r over 30 random scenarios equals 94,869,

94,465, and 94,824, respectively, which is 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

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

With our parameter settings, this leads to a lower bound of 9,040 on the expected total routing costs, as compared with 19,478, 15,998, and 14,196, respectively, for the average routing costs obtained by DI_r when $L = 2, 3, 4$, respectively.

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
	≥ 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
	≥ 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
	≥ 80	30	230	0.7	$H_0: P_{DI_d} = P_{DE_d}$	Accept	0.0

Even though this bound is weak, we conclude that the total average profit generated by DI_r is within 13%, 10%, and 7%, respectively, of the optimal expected profit when $L = 2, 3, 4$. Note that the actual optimality gap is probably much better, since our routing algorithm does 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 quickly become 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 shown in Table 3. The threshold for the t -statistic with 29° 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.

Superiority of DI over DE is not confined to improving the profit. In particular, DI uses fewer vehicles than DE , and the difference is 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 on DE .

10.4. Impact of Cost-to-Serve Values on DE and DI

Thus far, we have defined two ways to assign a positive cost-to-serve to a generic store, i.e., 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

Table 4. Statistical Tests on the Number of Vehicles

L	C	$E(V_{DE_d} - V_{DI_d})$	$Std(V_{DE_d} - V_{DI_d})$	t -statistic	Hypothesis	Result	Veh. dec. (%)
2	30	0.22	0.09	13.4	$H_0: V_{DI_d} \geq V_{DE_d}$	Reject	2.8
	≥ 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
	≥ 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
	≥ 80	0.30	0.16	10.3	$H_0: V_{DI_d} \geq V_{DE_d}$	Reject	3.6

cost-to-serve is not as clear: While F^d provides the highest profit for shelf lives $L = 2$ and $L = 3$, $F^r \approx 2.2F^d$ 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 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) and 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. Profit Obtained by DE and DI with Different Costs-to-Serve for $Q = 120$

Figure 2 also demonstrates 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 a (R, s, S) policy.

Economic interpretation of profit improvements. We argue 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, etc.). In fact, net profit in the retail food sector is of the order of 2% of revenue; see EBRD–FAO (2009), Government of Canada (2012), and FMI (2014). Miscellaneous costs thus account for about 30% of total revenue. The corresponding breakdown of the revenue is depicted in Figure 3.

Our results show an average improvement of 0.6% in profit when we exploit DI 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 provided an estimate of the value of information in Section 10.1: 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. To relax this assumption, we run a few additional experiments when the demand is binomial with parameters $(n, p) = (40, 0.5)$, compared with $(n, p) = (200, 0.1)$ in the initial instances. Observe

Figure 3. Breakdown of the Revenue

Table 5. Comparing Profits When the Variance Decreases

L	(n, p)	Profit by DI_r	Improvement
2	(200, 0.1)	75,392	0.3%
	(40, 0.5)	75,588	
3	(200, 0.1)	78,467	1.2%
	(40, 0.5)	79,407	
4	(200, 0.1)	80,628	0.5%
	(40, 0.5)	81,067	

that this modification preserves the mean value of the demand ($=20$) but decreases its variance $np(1-p)$ from 18 to 10. Thus, it allows us to investigate 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.

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, 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 influenced slightly 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)

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	75,375	1.7	1.6	97	99	70
	≥ 40	75,394	1.7	1.6	97	99	71
3	30	76,216	2.7	2.5	98	99	69
	40	76,997	2.5	2.1	99	99	50
	50	78,613	2.4	2.0	99	99	53
	≥ 60	78,628	2.4	2.0	99	99	53
4	30	76,216	3.7	3.5	98	99	69
	40	77,130	3.5	3.1	99	99	50
	50	79,790	3.3	2.9	99	99	53
	60	80,108	3.2	2.8	99	99	72
	70	80,368	3.2	2.7	99	99	70
	≥ 80	80,368	3.2	2.7	99	99	70

Figure 4. The Best Profit Obtained by DI for Different Shelf Lives

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. A further 4.3% profit loss is incurred when moving from shelf life 3 to 2. These values can 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 net profit, which are extremely significant.

Translating DE into a (R, s, S) policy. 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 , independent 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 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 shows that, regardless of the state of the system, the up-to-level point determined by DE_d is very

Figure 5. The Normalized Frequency of Up-to-Level Points in DE_d for Different Shelf Lives

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

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
DE	Applicable to multiple products
	F^d performs best if $\bar{n} \geq 2$; otherwise F^r performs better
DI	Performs similarly to the deliver-up-to-level policy
	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

10.7. Extensions

All solution methods presented above can be extended to account for inventory holding costs or for decaying products, i.e., products that lose their quality gradually over their shelf life. This is straightforward 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) to enforce higher freshness. To incorporate these elements in DE and DI, we can add the term e_{ti} defined by Equation (32) hereunder to the profit function given in Equation (7)

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

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 that 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 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 that explicitly takes the target service level into account. Our numerical results show that this naïve policy performs reasonably well when the target service level is high. Next, a decomposition method is applied to independently determine delivery quantity to each store. 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 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. We show how further profit improvement is possible when accessing full information about the future demands.

Acknowledgments

The authors thank Prof. Onno Boxma for his contributions at the early stages of the work, as well as two anonymous reviewers for their constructive comments.

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 (A.1)

$$F_{ti}^d = \frac{\sum_{j \in J_i} c_{ij}}{|J_i|}. \quad (\text{A.1})$$

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. Thus, there is a smaller chance to serve store i and its nearest neighbors in the same period, and, consequently, 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, Ergun, and Savelsbergh (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 must 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 (A.2)

$$\alpha_{ti} = E(D_{ti}) + \dots + E(D_{t+\lambda-1,i}). \quad (\text{A.2})$$

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

$$\bar{\alpha}_{ti} = \frac{\alpha_{ti} + \sum_{j \in J_i} \alpha_{tj}}{1 + |J_i|}. \quad (\text{A.3})$$

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

$$\bar{n}_{ti} = \frac{Q}{\bar{\alpha}_{ti}}, \quad (\text{A.4})$$

and the cost of the route serving i as

$$R_{ti} = 2c_{i0} + (\bar{n}_{ti} - 1)F_{ti}^d. \quad (\text{A.5})$$

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

$$F_{ti}^r = \frac{\alpha_{ti}}{\alpha_{ti} + \sum_{j \in J_i} \alpha_{tj}} \cdot R_{ti}. \quad (\text{A.6})$$

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, compared to the former costs F_{ti}^d .

Appendix B. Maximum Decrease and Increase in Delivery Quantities

Following the notations given in Section 7, 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., $\bar{m}_i$ 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 that 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 $\bar{R}_{t+1}$. To analyze the latter constraint, consider two cases. In the first, store $i \in D$ is not included in any route in $\bar{R}_{t+1}$, i.e., $i \in D \setminus N(\bar{R}_{t+1})$. 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 $\underline{m}_i \leq \min_{0 \leq y \leq y_{ti}} \{y | (\tilde{y}_{t+1,i} | y) = 0\}$. Hence, for every store $i \in D \setminus N(\bar{R}_{t+1})$, the maximum decrease of the delivery quantity y_{ti} is determined as

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

The second case considers stores $i \in D$ that 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 leads 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

$$\underline{m}_i = \min \{y_{ti} - y_{ti}^{(1)}, q_{tr'} - Q\}. \quad (\text{B.2})$$

A maximum increase of delivery quantity, say $\bar{m}_i$, 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., $\bar{m}_i \leq Q - q_t(i)$ must hold. On the other hand, store capacities must be respected, i.e., $\bar{m}_i \leq C_i - I_{ti} - y_{ti}$. As a result, the maximum increase of delivery quantity to any store $i \in I$ is determined as

$$\bar{m}_i = \min \{Q - q_t(i), C_i - I_{ti} - y_{ti}\}. \quad (\text{B.3})$$

References

- Adachi Y, Nose T, Kuriyama S (1999) Optimal inventory control policy subject to different selling prices of perishable commodities. *Internat. J. Production Econom.* 60:389–394.
- Adelman D (2004) A price-directed approach to stochastic inventory routing. *Oper. Res.* 52(4):499–514.
- 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, Cham, Switzerland), 585–596.
- Alegre J, Laguna M, Pacheco J (2007) Optimizing the periodic pickup of raw materials for a manufacturer of auto parts. *Eur. J. Oper. Res.* 179(3):736–746.
- Andersson H, Hoff A, Christiansen M, Hasle G, Løkketangen A (2010) Industrial aspects and literature survey: Combined inventory management and routing. *Comput. Oper. Res.* 37:1515–1536.
- Archetti C, Speranza MG (2016) The inventory routing problem: The value of integration. *Internat. Trans. Oper. Res.* 23(3):393–407.
- Archetti C, Bertazzi L, Laporte G, Speranza MG (2007) A branch-and-cut algorithm for a vendor managed inventory routing problem. *Transportation Sci.* 41(3):382–391.
- Barnes-Schuster D, Bassok Y (1997) Direct shipping and the dynamic single-depot multi-retailer inventory system. *Eur. J. Oper. Res.* 101(3):509–518.
- Berman O, Larson RC (2001) Deliveries in an inventory routing problem using stochastic dynamic programming. *Transportation Sci.* 35(2):192–213.
- Bertazzi L (2008) Analysis of direct shipping policies in an inventory routing problem with discrete shipping times. *Management Sci.* 54(4):748–762.
- Bertazzi L, Speranza MG (2011) Matheuristics for inventory routing problems. Montoya-Torres JR, Juan AA, Huatuco LH, Faulin J, Rodriguez-Verjan GL, eds. *Hybrid Algorithms for Service, Computing, and Manufacturing Systems: Routing and Scheduling Solutions* (IGI Global, Hershey, PA), 1–14.
- Bertazzi L, Speranza MG (2012) Inventory routing problem: An introduction. *Eur. J. Transportation Logist.* 1:307–326.
- Bertazzi L, Speranza MG (2013) Inventory routing problem with multiple customers. *Eur. J. Transportation Logist.* 2:255–275.
- Bertazzi L, Paletta G, Speranza MG (2002) Deterministic order-up-to level policies in an inventory routing problem. *Transportation Sci.* 36(1):119–132.
- Broekmeulen R, Van Donselaar K (2009) A heuristic to manage perishable inventory with batch ordering, positive lead times, and time-varying demand. *Comput. Oper. Res.* 36:3013–3018.
- Campbell AM, Savelsbergh M (2004a) Delivery volume optimization. *Transportation Sci.* 38(2):210–223.
- Campbell AM, Savelsbergh M (2004b) A decomposition approach for the inventory routing problem. *Transportation Sci.* 38(4):488–502.
- Campbell AM, Wilson JH (2014) Forty years of periodic vehicle routing. *Networks* 63(1):2–15.
- Campbell AM, Clarke L, Kleywegt AJ, Savelsbergh MWP (1998) The inventory routing problem. Crainic TG, Laporte G, eds. *Fleet Management and Logistics* (Springer, Boston), 95–113.
- Chan MLA, Federgruen A, Simchi-Levi D (1998) Probabilistic analyses and practical algorithms for inventory routing models. *Oper. Res.* 46(1):96–106.
- Chao X, Gong X, Shi C, Zhang H (2015) Approximation algorithms for perishable inventory systems. *Oper. Res.* 63(3):585–601.
- Chiu HN (1995) A heuristic (R, T) periodic review perishable inventory model with lead times. *Internat. J. Production Econom.* 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. *Eur. J. Oper. Res.* 208(1):86–94.
- Coelho LC, Laporte G (2013a) A branch-and-cut algorithm for the multi-product multi-vehicle inventory-routing problem. *Internat. J. Production Res.* 51:7156–7169.
- Coelho LC, Laporte G (2013b) The exact solution of several classes of inventory routing problems. *Comput. Oper. Res.* 40:558–565.
- Coelho LC, Laporte G (2014a) Optimal joint replenishment, delivery and inventory management policies for perishable products. *Comput. Oper. Res.* 47:42–52.
- Coelho LC, Laporte G (2014d) Improved solutions for inventory routing problems through valid inequalities and input ordering. *Internat. J. Production Econom.* 155:391–397.
- Coelho LC, Cordeau JF, Laporte G (2014) Thirty years of inventory routing. *Transportation Sci.* 48(1):1–19.
- Cooper WL (2001) Pathwise properties and performance bounds for a perishable inventory system. *Oper. Res.* 49(3):455–466.
- Deniz B, Karaesmen I, Scheller-Wolf A (2010) Managing perishables with substitution: Inventory issuance and replenishment heuristics. *Manufacturing Service Oper. Management* 12(2):319–329.
- Desaulniers G, Rakke JG, Coelho LC (2016) A branch-price-and-cut algorithm for the inventory routing problem. *Transportation Sci.* 50(3):1060–1076.
- 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.
- Ekici A, Ozener OO, Kuyzu G (2015) Cyclic delivery schedules for an inventory routing problem. *Transportation Sci.* 49(4):817–829.
- Federgruen A, Zipkin P (1984) A combined vehicle routing and inventory allocation problem. *Oper. Res.* 32(5):1019–1037.
- Federgruen A, Prastacos G, Zipkin PH (1986) An allocation and distribution model for perishable products. *Oper. Res.* 34(1):75–82.
- FMI (2014) Supermarket facts. Food Marketing Industry, Arlington, VA, <http://www.fmi.org/research-resources/supermarket-facts>.
- Gallego G, Simchi-Levi D (1990) On the effectiveness of direct shipping strategy for the one-warehouse multi-retailer R-systems. *Management Sci.* 36(2):240–243.
- Gaur V, Fisher ML (2004) A periodic inventory routing problem at a supermarket chain. *Oper. Res.* 52(6):813–822.
- Government of Canada (2012) Canadian industry statistics on grocery stores (NAICS 4451). Ottawa, Ontario, Canada, <https://www.ic.gc.ca/app/scr/sbms/sbb/cis/revenues.html?code=4451&lang=eng>.
- Goyal SK, Giri BC (2001) Recent trends in modeling of deteriorating inventory. *Eur. J. Oper. Res.* 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 Sci.* 44(3):400–415.
- Hemmelmayr V, Doerner KF, Hartl RF, Savelsbergh MWP (2009) Delivery strategies for blood products supplies. *OR Spectrum* 31:707–725.
- Huang SH, Lin PC (2010) A modified ant colony optimization algorithm for multi-item inventory routing problems with demand uncertainty. *Transportation Res. Part E* 46:598–611.
- Jaillet P, Bard JF, Huang L, Dror M (2002) Delivery cost approximations for inventory routing problems in a rolling horizon framework. *Transportation Sci.* 36(3):292–300.
- Karaesmen IZ, Scheller-Wolf A, Deniz B (2011) Managing perishable and aging inventories: Review and future research directions. Kempf KG, Keskinocak P, Uzsoy R, eds. *Planning Production and Inventories in the Extended Enterprise*. Internat. Series Oper. Res. Management Sci., Vol. 151 (Springer, New York), 393–436.
- Kleywegt AJ, Nori VS, Savelsbergh MWP (2002) The stochastic inventory routing problem with direct deliveries. *Transportation Sci.* 36(1):94–118.
- Kleywegt AJ, Nori VS, Savelsbergh MWP (2004) Dynamic programming approximations for a stochastic inventory routing problem. *Transportation Sci.* 38(1):42–70.
- Kouki C, Jouini O (2015) On the effect of lifetime variability on the performance of inventory systems. *Internat. J. Production Econom.* 167:23–34.
- Kouki C, Jemai Z, Minner S (2015) A lost sales (r, Q) inventory control model for perishables with fixed lifetime and lead time. *Internat. J. Production Econom.* 168:143–157.
- Le T, Diabat A, Richard JP, Yih Y (2013) A column generation-based heuristic algorithm for an inventory routing problem with perishable goods. *Optics Lett.* 7:1481–1502.

- Lian Z, Liu L (1999) A discrete-time model for perishable inventory systems. *Ann. Oper. Res.* 87:103–116.
- Minkoff AS (1993) A Markov decision model and decomposition heuristic for dynamic vehicle dispatching. *Oper. Res.* 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. *Comput. Indust. Engrg.* 87: 213–227.
- Nahmias S (2011) *Perishable Inventory Systems* (Springer, New York).
- Özener Ö, Ergun Ö, Savelsbergh MWP (2013) Allocating cost of service to customers in inventory routing. *Oper. Res.* 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.
- Puterman ML (1994) *Markov Decision Processes* (John Wiley & Sons, New York).
- Qu WW, Bookbinder JH, Iyogun P (1999) An integrated inventory transportation system with modified periodic policy for multiple products. *Eur. J. Oper. Res.* 115:254–269.
- Raa B, Aghezaaf EH (2008) Designing distribution patterns for long-term inventory routing with constant demand rates. *Internat. J. Production Econom.* 112:255–263.
- Raa B, Aghezaaf EH (2009) A practical solution approach for the cyclic inventory routing problem. *Eur. J. Oper. Res.* 192:429–441.
- Reinman MI, Rubio R, Wein LM (1999) Heavy traffic analysis of the dynamic stochastic inventory routing problem. *Transportation Sci.* 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 Optim.* 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 Oper. Management* 8(3):253–272.
- Solomon MM (1987) Algorithms for the vehicle routing and scheduling problems with time window constraints. *Oper. Res.* 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 Sci.* 45(3):335–345.
- Stacey J, Natarajaratnam M, Sox C (2007) The storage constrained inbound inventory routing problem. *Internat. J. Physical Distribution Logist. 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. *Ann. Oper. Res.* 197(1):135–158.