Omega 141 (2026) 103510

Contents lists available at ScienceDirect

Omega

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

Incorporating stochastic optional pickup demand in routing operations with divisible services for hub-and-spoke e-commerce returns management systems[☆]

Alessandro Gobbi [a],*, Daniele Manerba [b], Francesca Vocaturo [c]

[a] Dipartimento di Ingegneria Meccanica, Energetica e Gestionale, via Pietro Bucci, Cubo 41/C, Arcavacata di Rende, 87036, CS, Italy
[b] Department of Information Engineering, Università degli Studi di Brescia, via Branze 38, Brescia, 25123, BS, Italy
[c] Dipartimento di Economia, Statistica e Finanza "Giovanni Anania", Università della Calabria, via Pietro Bucci, Cubo 0/C, Arcavacata di Rende, 87036, CS, Italy

ARTICLE INFO

Keywords:
Sustainable logistics
Hub-and-spoke return management
Vehicle routing problem with deliveries and pickups
Stochastic optional pickup demand
Stochastic programming with recourse
Detour operations
Spot-market services

ABSTRACT

Nowadays, e-commerce is associated with many returns due to emotional consumption, information asymmetry, factory defects, or, more generally, customer dissatisfaction. However, little attention has been paid to reverse logistics in the e-commerce industry, although it has been proven crucial to improving the perceived quality of service and profit revenue. Depending on the nature of the goods, one successful option is to design combined *forward-and-reverse* logistics systems, where the collection of returns is ensured along with the traditional distribution of products, together with *hub-and-spoke* networks in which both distribution and collection demand from many spokes are aggregated into a few hubs. In this context, we study a variant of the vehicle routing problem with divisible deliveries and pickups, in which each hub may be associated with a mandatory delivery demand and a mandatory return pickup demand, and it may be visited more than once within the same or different routes. To address realistic scenarios, and given the large fluctuation of demand within the aggregating hubs, we also assume that an uncertain optional pickup quantity may arise and formulate the problem through two-stage Stochastic Programming, proposing and modeling ad-hoc recourse actions. Moreover, an integer L-shaped method enhanced with ad-hoc valid inequalities is developed for solving the resulting problem. Managerial insights on the underlying tactical and operational policies are inferred from extensive computational experiments on a case study and on realistic artificial instances.

1. Introduction

In the last two decades, the fast development of information and communication technologies, online services, and mobile applications has resulted in significant growth of the e-commerce segment and the proliferation of virtual stores (*e-retailers*). A further acceleration was pushed by the COVID-19 pandemic, during which e-commerce played an essential role in the economy and society by ensuring continued access to producers and services to consumers. A recent *European E-commerce Report* [1] exhibits surprisingly high numbers of internet purchases in recent years. For instance, the 86% of citizens in Northern Europe bought goods or services online (these citizens can be defined as *e-shoppers*). The same report also discloses projections that expect online sales to make up an average of 30% of retail turnover by 2030.

Despite being a virtual activity, e-commerce strongly influences freight flows, especially in urban areas, since its fundamental characteristic is the physical distribution of goods to the consumers' home or to appropriate sites (lockers, hubs, etc.). In this last-mile context, many independent players operate with poor vehicle utilization for low-profit margins in a customer-oriented delivery business, thus generating several nuisances such as air polluting emissions, congestion, noise, and infrastructural damages [2,3]. Unsurprisingly, many scientific articles have addressed the management and optimization of such last-mile distribution operations, also focusing on sustainability issues [4]. Instead, little attention has been paid to the so-called *reverse logistics* in the e-commerce industry, although it has been proven crucial to improving the perceived quality of service and profit revenue. In fact, nowadays, a considerable number of products are returned by e-shoppers due to many reasons, e.g., information asymmetry (there always exists a gap between real products and virtual pictures), factory defects, emotional consumption pushed by many ubiquitous promotions [5], or in general by the customer dissatisfaction. For example, only in 2023 in the U.S.A.,

☆ Area: Transportation and Logistics. This manuscript was processed by Associate Editor Merve Bodur.

* Correspondence to: Dipartimento di Ingegneria Meccanica, Energetica e Gestionale, Università della Calabria, Via Pietro Bucci 41/C, Rende, Italy.
E-mail address: alessandro.gobbi@unical.it (A. Gobbi).

https://doi.org/10.1016/j.omega.2025.103510
Received 31 December 2024; Accepted 27 December 2025
Available online 29 December 2025

e-shoppers generated a total amount of retail returns of 17.6%, for a total value of about 250 billion dollars [6].

We focus on returns management solutions that can be incorporated into the existing last-mile distribution business. A successful option for e-retailers is to design a combined distribution-collection logistics system in which the collection of unsatisfactory items is ensured along with the traditional distribution of products to e-shoppers. Many studies (e.g., [7]) show that vehicle routing used in freight flow optimization becomes more effective and balanced in a combined *forward-and-reverse* logistics network. This practice is also suggested by company managers and business consultants. For example, Jenkins [8] confirms that combining returns with deliveries is a practice implemented by many e-commerce companies that makes returns management a low-cost transaction. Concurrently, logistics companies are increasingly adopting *hub-and-spoke* last-mile solutions, where goods are temporarily stored in intermediate hubs before reaching the final customers [9]. This approach optimizes vehicle usage and inventory management, reduces the environmental impact of deliveries, and effectively handles unforeseen issues such as the unavailability of the final customer [10]. If a hub-and-spoke scheme is combined with the aforementioned forward-and-reverse logistics vision, hubs aggregate not only the delivery but also the return demand from many private customers. This allows e-commerce companies to speed up the process of collection, inspection, refurbishment, and resale of returned goods [11].

From a practical perspective, decision-makers must solve a specific Vehicle Routing Problem (VRP) known in the scientific literature as the VRP with Deliveries and Pickups (VRPDP). In the VRPDP, some locations receive goods (*delivery locations*) and some others return goods (*pickup locations*), and a fleet of vehicles is available to perform both the delivery and pickup service at these locations with the aim of reducing logistics costs. Another important feature that could improve the process efficiency is the possibility that a location requiring both pickup and delivery service is fulfilled, if beneficial, in two separate visits (one for each service). With this vision, Nagy et al. [12] proposed a VRPDP version called *Vehicle Routing Problem with Divisible Deliveries and Pickups* (VRPDDP). In their problem, all delivery goods come from one depot and all return goods are brought to the same depot, excluding the possibility of goods traveling directly from one location to another. A vehicle starts from the depot carrying only goods to deliver and goes back to the depot carrying only returns; at an intermediate stage, it may carry delivery goods, return goods, or a mixture of them. At each location, the total load on the vehicle may increase or decrease.

In this work, we extend the just-described VRPDDP by assuming that only a part of the pickup demand is deterministic and mandatory, i.e., that there exists a portion of the pickup demand not deterministically known in advance and optional. In addition, we impose a minimum threshold on the optional pickup demand to be collected during the current day. These extensions allow us to address a more realistic problem for several reasons. First, return requests can be considered uncertain since, upon the quantity already declared by a location at the cut-off time, additional demand may appear while the vehicle is doing its route. Differently from delivery requests, which can be fulfilled only if the vehicle has loaded the relative goods at the depot, unexpected return requests could be satisfied provided that there is enough residual capacity on the visiting vehicle. Clearly, since such additional pickup demand has been received after the cut-off time, its collection is optional for the current day. Second, the deadline for retrieving merchandise is typically set in a broad time window, making most of the pickup services non-mandatory for many days. Hence, the definition of a minimum threshold on the optional demand helps the company not to accumulate too many collection services that will become mandatory soon. Anticipating workload in an uncertain environment has indeed, in many industrial contexts, a positive impact on the capacity of the system under study to satisfy future needs [13]. According to McCue [14], returns should be handled even immediately. In fact, another reason to avoid accumulation is the will to put returned goods on the market as soon as possible to recover their value.

We name the resulting problem the *Stochastic VRPDDP with optional pickup demand* (SVRPDDP$_{op}$) and we model it through a two-stage Stochastic Programming (SP) formulation [15]. In particular, our second-stage problem involves the possibility of redirecting vehicles to the depot (i.e., performing a *detour*) to unload already picked-up items (and then continue to fulfill the remaining requests) as well as the possibility of activating spot-market pickup services, thus outsourcing some operations.[1] The obtained stochastic model is eventually approximated through a deterministic Mixed-Integer Linear Programming (MILP) formulation considering a finite number of scenarios. To solve the problem, we develop an exact algorithm, namely, an Integer L-shaped method enhanced with ad-hoc valid inequalities, and we run a large set of experiments (both on case-study inspired instances and on artificial ones) to obtain statistically solid economic and managerial insights on the optimization setting.

Our contribution is threefold. First, to the best of our knowledge, our work represents the first study concerning the SVRPDDP$_{op}$, i.e., a more sophisticated extension of the classical VRPDDP that includes both deterministic mandatory and stochastic optional pickup demands, together with deterministic mandatory delivery demands. This enriches the literature concerning e-commerce returns management, which is still poor with respect to routing activities under uncertainty, with mathematical modeling and properties of the basic operational setting, as well as of the recourse actions available to hedge against the uncertainty. Second, we develop and validate a two-stage SP formulation and an exact algorithm able to solve to optimality SVRPDDP$_{op}$ instances with up to 20 hubs, by leveraging specific properties of the model (such as valid inequalities and decomposability). Third, the results obtained by our economic analysis yield a considerable list of managerial insights for route planners in the e-commerce industry. In particular, they will be able to compare and assess different strategies to produce more robust day-ahead plans and to lead their actual operations into a more flexible state, thus facilitating adaptation if necessary. Note that, to provide tailored insights on the reverse logistics side of the network, we decided to deal solely with uncertain pickups. However, it is not complicated to extend our study to scenarios in which uncertainty also affects deliveries.

The rest of this article is organized as follows. The related scientific literature is discussed in Section 2. The formal description and mathematical formulation of the problem are given in Section 3, while Section 4 illustrates the exact solution methodology developed. After presenting in Section 5 the experimental setting used to validate our approach, Section 6 discusses the results of an extensive computational phase and derives useful managerial insights, both on artificial but realistic instances and on a specific case study. Conclusions are drawn in Section 7.

2. Literature review

In this section, we first review some studies concerning the VRPDPs, also presenting their classification (Section 2.1). Then, we illustrate some fundamental concepts of two-stage SP and present some problems tackled via this paradigm, with a focus on transportation with delivery and pickup services (Section 2.2).

2.1. Vehicle routing problems with deliveries and pickups

The VRPDP is an extension of the classical VRP in which goods are not only delivered to customers (*linehauls*) but also picked up

[1] On the spot market, freight and released trucks are offered daily. Through digital freight exchange platforms, a company indicates which goods need to be transported and various transportation service providers submit their offers. The transaction carried out by the highest bidder is then completed within a very short time.

from them (*backhauls*). The scientific literature divides the VRPDPs associated with reverse logistics applications into four classes [16]: besides the VRPDDP, the others are the *VRP with Backhauls* (VRPB), the *VRP with Mixed Linehauls and Backhauls* (VRPMLB), and the *VRP with Simultaneous Pickups and Deliveries* (VRPSPD). These problems are referred to as "one-to-many-to-one" since they consider transportation requests for loads that originate at one depot and must be delivered to locations and loads that originate at locations and must be delivered to the depot. Other problems involving pickup and delivery services exist in the literature. For instance, Jargalsaikhan et al. [17] focus on the "one-to-one" pickup and delivery problem with divisible pickups and deliveries in which transportation requests between pairs of locations are considered, while the vehicle performing the service starts and ends empty at the depot. Instead, Li et al. [18] address a "many-to-many" pickup and delivery problem in which multiple wholesale markets for agricultural products collaborate horizontally to distribute food materials to surrounding customers.

In the VRPB and VRPMLB, referred to as *single-demands* problems by Wassan and Nagy [19], customers require either deliveries or pickups. In the VRPB, deliveries must be performed before pickups. This implies that the vehicles first deliver all goods to linehauls, then pickup goods from backhauls. The VRPB is typical of those settings where pickup loads cannot be accommodated in the vehicle together with the delivery loads, e.g., when there is a risk of cross-contamination. Mingozzi et al. [20] propose an exact approach to the VRPB, while a heuristic algorithm for the problem is described by Ropke and Pisinger [21]. Instead, in the VRPMLB, the vehicles can perform the services in any sequence, implying that linehauls and backhauls can be fully mixed within a route. The VRPMLB has not received as much attention in the literature and some studies are discussed in the review of Koç and Laporte [22]. It may be useful to produce a practical compromise between the VRPB and the classical version of the VRPMLB in which the mixture is totally free. In other words, allowing mixed linehauls and backhauls may be appropriate while satisfying some restrictions. For instance, it could be possible to visit backhauls before all linehauls have been served while adding a restriction for which a certain number of linehauls are served before the first backhaul customer can be considered [23]. Finally, the VRPSPD is referred to as a *combined-demand* problem in Wassan and Nagy [19] and involves customers that are both linehauls and backhauls. In this case, a single visit is allowed to each customer, i.e., both delivery and pickup services must occur simultaneously. The literature on VRPSPD is rich and active. To solve the problem, Tasan and Gen [24] describe a genetic algorithm-based approach, while Subramanian et al. [25] propose a branch-cut-and-price (also capable of solving the VRPMLB). More recently, Koç et al. [26] have discussed the progress made on the traditional version of the VRPSPD and its variants/extensions studied over time.

When customers with combined demand are not committed to being satisfied in a single visit, better vehicle routes can be designed. The term "divisible" is used to refer to this relaxed version of the VRPSPD. In the VRPDDP (the problem studied in this paper), the pickup and the delivery quantities may be served, if beneficial, in two separate visits at the same customer location by the same vehicle or even by two different ones. A customer with two visits is denoted as a *split customer*. Gribkovskaia et al. [27] propose a MILP model for the single-vehicle case and several heuristics. Instead, Nagy et al. [12] deal with the multi-vehicle case and formally analyze the significant savings that can be achieved by allowing the pickup and delivery quantities to be served separately with respect to the simultaneous service. An important aspect generally studied in the VRPDDP is the shape of the obtained vehicle routes. In the case of breweries, e.g., bottles are first delivered until the vehicle is partly unloaded, then both pickup and delivery are performed at the remaining customers, and empty bottles are eventually picked up from the first visited customers. These customers are revisited in reverse order, thus giving rise to *lasso-shaped* solutions [28]. More recently, Polat [29] has proposed a parallel variable neighborhood search approach to the VRPDDP, in which asynchronous cooperation with a centralized information exchange strategy is used to efficiently explore different regions of the solution space and reduce computational times.

Finally, several VRPDPs with optional pickup demands have been presented in the scientific literature. For example, Bruck and Iori [30] eliminate the mandatory collection of pickup demands in the single-vehicle case and propose non-elementary formulations to address the benchmark instances proposed by Gribkovskaia et al. [27]. Moreover, concerning the multi-vehicle case, in the studies of Gribkovskaia et al. [31], Gutiérrez-Jarpa et al. [32], and Bruck and dos Santos [33], each customer has only an optional pickup demand, and no demand needs to be necessarily satisfied. The only work that simultaneously studies mandatory and optional pickups is by Santos et al. [34], where the aim is to minimize the fuel cost minus the revenue obtained by visiting the customers. However, as far as we know, no stochastic VRPDP with both mandatory and optional pickup demands has been analyzed yet. From this perspective, our study enriches the literature with a practical variant in which stochastic optional pickup is incorporated into the service. This option allows e-retailers (and eventually other companies managing returned products) to anticipate workload and obtain the corresponding advantages.

2.2. Two-stage SP models

SP represents a common way to deal with uncertain data in a decision-making process. It is particularly used when the decision-maker is interested in building solutions that are good on average, and when the problem's decisions can be made over time and decomposed in stages, according to the realizations of some random variables. When all the realizations of the random variables take place at a single point in time, we have a two-stage SP. In practice, the first stage consists of here-and-now decisions before uncertainty is revealed, while the second stage consists of wait-and-see decisions after realization of uncertainty. Generally, the aim is to minimize or maximize the total expected costs or revenues by considering here-and-now decisions and possible future recourse [35]. Here, we restrict the attention on routing problems or integrated problems in which routing represents a fundamental part. The reader is referred to the surveys of Oyola et al. [36], Oyola et al. [37], and De Maio et al. [38] for a broader discussion about the topic.

Laporte et al. [39] study a waste collection routing problem in which stochastic demands are associated with the arcs of the graph representing the street network. In this problem, modeled through two-stage SP with recourse, a route failure occurs whenever the realized demand exceeds the vehicle capacity. The vehicle then drives to a dump site to unload and returns to the point of failure to resume its route. Zhao et al. [40] address a heterogeneous VRP with time windows and stochastic demand, formulating it as a two-stage SP model. A no-split detour to depot policy is used as recourse action. In a real-world case study, the solutions obtained through an approach based on this model achieve an average cost reduction of 2.35% compared to an alternative approach. Beraldi et al. [41] analyze a one-to-one pickup and delivery problem in which each transportation request is associated with a Bernoulli random variable taking value 1 if the pair "pickup-delivery" associated with that request requires service, and 0 otherwise. To solve the problem, the authors propose a two-stage SP framework where a set of feasible routes is determined at the first stage, while in the second stage the customers requiring no service are simply skipped while performing the planned routes. Zhang et al. [42] deal with a production inventory routing problem with simultaneous pickups and deliveries. In the system under study, a set of manufacturing and re-manufacturing facilities are taken into account to reduce the environmental effects through the re-manufacturing of worn-out products. The authors introduce a two-stage SP formulation to consider the uncertain nature of customers' demand and a real-world variant of the problem that

Table 1
Overview of main notation.

Sets and parameters	
H	Set of hubs
$N_P = \{1, \dots, \|H\|\}$	Set of pickup requests (or pickup points)
$N_D = \{\|H\|+1, \dots, 2\|H\|\}$	Set of delivery requests (or delivery points)
$N = N_P \cup N_D$	Set of service requests (or service points)
$V = N \cup \{0\}$	Set of nodes, where 0 represents the depot
A	Set of arcs
$\mathcal{S}$	Set of scenarios
c_{ij}	Travel cost for $(i,j) \in A$
Q	Vehicle capacity
q_j	Mandatory demand for $j \in N$
o_j	Nominal optional pickup demand for $j \in N_P$
$\tilde{o}_j$	Stochastic non-negative variation of optional pickup demand for $j \in N_P$ over o_j
α	Minimum percentage of optional demand to be fulfilled $(0 \leq \alpha \leq 1)$
κ_j	Cost of a spot-market service for $j \in N_P$
$\tilde{o}^s_j$	Non-negative variation of optional pickup demand for $j \in N_P$ in $s \in \mathcal{S}$ over o_j
π^s	Probability associated with the occurrence of $s \in \mathcal{S}$
Decision variables	
x_{ij}	1 if a vehicle is expected to travel on $(i,j) \in A$, 0 otherwise
y_j	1 if a vehicle visits $j \in N_P$, 0 otherwise
D_{ij}	Delivery quantity carried on $(i,j) \in A$
P_{ij}	Expected pickup quantity carried on $(i,j) \in A$
$\tilde{x}_{ij}$	1 if a vehicle keeps traveling on $(i,j) \in A$ as planned by x_{ij}, 0 if a detour is performed between i and j
$\tilde{P}_{ij}$	Variation of the pickup quantity carried on $(i,j) \in A$ with respect to P_{ij}
$\tilde{z}_j$	1 if a vehicle fulfills the optional demand of $j \in N_P$, 0 otherwise
$\tilde{\lambda}$	Total optional demand which is not satisfied
$\tilde{r}_j$	1 if $j \in N_P$ is served through a spot-market service, 0 otherwise
$\tilde{x}^s_{ij}$	1 if a vehicle keeps traveling on $(i,j) \in A$ in $s \in \mathcal{S}$ as planned by x_{ij}, 0 if a detour is performed between i and j
$\tilde{P}^s_{ij}$	Variation of the pickup quantity carried on $(i,j) \in A$ with respect to P_{ij} in $s \in \mathcal{S}$
$\tilde{z}^s_j$	1 if a vehicle fulfills the optional demand of $j \in N_P$ in $s \in \mathcal{S}$, 0 otherwise
$\tilde{\lambda}^s$	Total optional demand which is not satisfied in $s \in \mathcal{S}$
$\tilde{r}^s_j$	1 if $j \in N_P$ is served through a spot-market service in $s \in \mathcal{S}$, 0 otherwise

accounts for carbon emission regulations, assuming a carbon cap-and-trade policy. A stochastic reverse logistics production routing problem with re-manufacturing and simultaneous pickups and deliveries is also studied by Shuang et al. [43]. Their problem, formulated as a two-stage SP model, considers different carbon emission control policies with heterogeneous transportation fleets and allows for lost sales.

Wollenberg [44] proposes a two-stage SP approach to the single-vehicle VRPSPD, in which the quantities to be picked up follow a finite discrete probability distribution. Zhu and Sheu [45] generate a priori routes for a fleet of vehicles that pick up and deliver items with uncertain demands through SP. Under uncertainty, a vehicle may reach a customer location without enough capacity to pick up returned items or enough inventory to deliver new products (or both), causing a route failure that requires a recourse action. The authors propose a failure-specific cooperative recourse strategy to explore a risk pooling mechanism for routing in the context of simultaneous pickups and deliveries, i.e., when a route failure is encountered, the lead vehicle is asked to satisfy the unmet demand after returning to the depot or to leave the unmet demand to a partner vehicle. More realistic assumptions are made by Zhang et al. [46] that also propose a failure-specific cooperative recourse strategy, labeled as *dynamic multi-stage*, for the VRPSPD. Our study extends these works on the stochastic VRPSPD in terms of flexibility and efficiency. Indeed significant savings can be achieved by allowing the pickup and delivery quantities to be served separately with respect to the simultaneous service.

3. Problem description and mathematical formulation

In this section, we first formally define the stochastic problem under study (Section 3.1). Then, we propose a two-stage SP model (Section 3.2) and the corresponding deterministic scenario-based problem (Section 3.3). An overview of all notation is given in Table 1.

3.1. Problem definition

In the following, we describe in detail the optimization problem addressed, highlighting the information known and the possible actions available at each decisional level.

3.1.1. Day-ahead decisions

Let us consider a fleet of homogeneous vehicles, each with capacity Q, located at a single depot 0 and available for a joint delivery and pickup service to a set H of hubs, aggregating the pickup and delivery demands from many customers. Let $N = \{1, 2, \dots, 2|H|\}$ be the set of service requests, in which j and $j + |H|$ represent the pickup and the delivery request of the same hub $j \in H$, respectively. Therefore, N can be partitioned into two subsets, $N_P = \{1, 2, \dots, |H|\}$ and $N_D = \{|H|+1, |H|+2, \dots, 2|H|\}$, where the former includes all the pickup requests and the latter includes all the delivery requests. Each request $j \in N$ has a *mandatory* demand $q_j \geq 0$. Clearly, when $j \in N_P$, then q_j represents the mandatory demand to pickup, while when $j \in N_D$, then q_j represents the mandatory demand to deliver. Moreover, only pickup requests also show an *optional* demand, which may or may not be collected. More precisely, each $j \in N_P$ is also associated with an optional demand $o_j + \tilde{o}_j$, where $o_j \geq 0$ is a deterministic nominal value, while $\tilde{o}_j \geq 0$ is a stochastic variable representing a non-negative variation over the nominal value. Finally, let α be the minimum percentage of optional pickup demand that the company is expected to collect.

In the above-described setting, the decision maker aims to plan in advance, specifically a day ahead, a significant portion of their operations, particularly those related to assigning hubs to vehicle routes and sequencing their visits. This may be required due to other collateral operations for the hubs (such as synchronization for last-mile delivery or palletizing returns) and for the vehicles (loading and packing parcels), or because solving the entire problem at operational time could be impractical. Such decisions are subject to uncertainty, since the actual optional pickup demand that will emerge the next day is not deterministically known.

It is important to clarify that, in day-by-day processes like the one under study, mandatory demands (both delivery and pickup) are those services that must be performed by the end of the incoming working day (e.g., because of the presence of strict deadlines on them), while considering an optional demand means having the flexibility to postpone some pickup services to the following days. In this regard, the parameter α allows the company to control the minimum percentage of optional pickup demand to be fulfilled with the aim of *(i)* avoiding the accumulation of too many pickup requests for the upcoming days and *(ii)* recovering in short time the value of the returned products and put them back on the market. Finally, the definition of a partially unknown optional pickup demand is coherent with the possibility of accumulating some of it and, at the same time, with the dynamism of the information coming. In fact, the deterministic part corresponds to the demand that emerged in the previous days and is still not satisfied, while the optional part corresponds to that appearing during the daily process.

3.1.2. Operational setting and recourse policies

At the operational level, all the demands become known, and the decision maker adjusts and completes their plan by implementing *recourse actions*. Note that, even if the actual optional pickup demand could be revealed at different times during the operations (e.g., when the vehicle arrives at a hub), we assume a single moment at which it is possible to make new decisions and implement the recourse action against the full knowledge of the problem data. This allows the planned decisions to be conservative enough with respect to the possible future scenarios, despite the actual dynamism of the recourse actions.

The first basic recourse we want to implement in our model is the chance to review the decisions on whether to collect or not an optional demand from a hub, and to accordingly redistribute the pickup quantities carried by vehicles to face the unexpected appearance of optional demand. However, despite the above flexibility, it is easy to see that the day-ahead planning may turn out to be infeasible. In particular, some of the planned routes may show insufficient capacity for pickup demands when they are revealed, thus failing to satisfy the minimum percentage of optional demand to collect. Hence, to always guarantee a feasible solution against any day-ahead choice (i.e., to achieve the so-called *relatively complete recourse*), two different practically used strategies are applied concurrently:

- **detour policy**: vehicles are provided with the possibility to take a *detour*, i.e., of redirecting to the depot to unload the returns already picked up and then fulfilling the remaining hubs' requests. Note that, since vehicles are not allowed to load additional delivery goods at the depot during a detour, a solution including a vehicle detouring during its expected route is definitely different from a solution including two separate vehicles serving the two sub-routes before and after the detour point. Moreover, a detour may occur at any point within the planned route and not necessarily from those hubs for which the residual vehicle capacity is not enough to accommodate the pickups. A similar idea was proposed by Yang et al. [47] for the identification of *restocking points* in a pure delivery context.
- **spot-market policy**: *spot-market* pickup services can be activated, thus outsourcing part of the operations to an external transportation company. The spot-market cost reasonably depends on the hub involved in the service, in particular on its distance and its dimension in terms of demand, hence, we define a fixed cost $\kappa_j > 0$ when a spot-market service is activated for any $j \in N_P$. A hub for which a spot-market service is activated is called *critical* since it is necessary to collect its optional pickup to satisfy the minimum required percentage of optional demand.

A graphical exemplification of the decisions' dynamic of the problem is reported in Fig. 1. In particular, given the day-ahead planned route of Fig. 1(a), Fig. 1(b) shows a detour implementation, while Fig. 1(c) shows the use of a spot-market service.

3.1.3. Overall problem and assumptions

We name the resulting optimization setting as the *Stochastic Vehicle Routing Problem with Divisible Deliveries and Pickups with optional pickup demand* (SVRPDDP$_{op}$). We define a directed graph $G = (V, A)$, with node set $V = N \cup \{0\}$ and arc set $A = \{(i,j) : i,j \in V, i \neq j\}$. A travel cost $c_{ij} > 0$ is associated with each arc $(i,j) \in A : |i-j| \neq |H|$, while $c_{ij} = 0$ for any arc $(i,j) \in A : |i-j| = |H|$ (since i and $i+|H|$ represent the same hub). Then, the SVRPDDP$_{op}$ seeks to find a set of vehicle routes, i.e., directed cycles over G starting and ending at the depot, able to fulfill the mandatory delivery and pickup demands of the hubs and to satisfy the minimum percentage of the optional pickup demand, while minimizing the expected total cost given by the routing operations (comprises of the detours) and the spot-market costs.

In this problem, we stick with the following assumptions: (*i*) all deliveries come from the depot, and all pickups need to be transported to the same depot; (*ii*) each vehicle can transport both pickup and delivery goods, as long as the total quantity carried does not exceed the vehicle's capacity; (*iii*) each request (pickup or delivery) must be satisfied by a single visit. However, since we adopt a divisible service perspective, hubs requiring both a pickup and delivery service can be served by two different visits; (*iv*) to avoid unnecessary detours, we make the triangle inequality hold for the arc costs, so that coming back to the depot between two hubs is always more costly than moving directly between them.

3.2. A two-stage SP formulation

In the following, we provide a two-stage SP formulation for the SVRPDDP$_{op}$, which is naturally tailored for modeling the day-ahead decisions (at the first stage) and the recourse actions (at the second stage).

Let us define the following first-stage variables: x_{ij}, taking value 1 if a vehicle is expected to travel on arc $(i,j) \in A$, and 0 otherwise; y_j, taking value 1 if a vehicle visits node $j \in N_P$, and 0 otherwise; D_{ij} and P_{ij}, measuring the amount of delivery quantities and the expected amount of pickup quantities (mandatory or optional), respectively, carried on arc $(i,j) \in A$. Then, our SVRPDDP$_{op}$ can be modeled as follows[2]:

$$\min_{\mathbf{x},\mathbf{y},\mathbf{D},\mathbf{P}} \quad \sum_{(i,j)\in A} c_{ij}x_{ij} + \mathbb{E}[\delta(\mathbf{x},\mathbf{y},\mathbf{D},\mathbf{P},\tilde{\mathbf{o}})] \tag{1a}$$

$$\text{s.t.} \quad \sum_{(i,j)\in A} x_{ij} = 1 \qquad j \in N | q_j > 0 \tag{1b}$$

$$\sum_{(i,j)\in A} x_{ij} = y_j \qquad j \in N_P \tag{1c}$$

$$\sum_{(j,i)\in A} x_{ji} = \sum_{(i,j)\in A} x_{ij} \qquad j \in N \tag{1d}$$

$$\sum_{(i,j+|H|)\in A} x_{i,j+|H|} \leq y_j \qquad j \in N_P \tag{1e}$$

$$\sum_{(i,j)\in A} D_{ij} - q_j = \sum_{(j,i)\in A} D_{ji} \qquad j \in N_D \tag{1f}$$

$$\sum_{(i,j)\in A} D_{ij} = \sum_{(j,i)\in A} D_{ji} \qquad j \in N_P \tag{1g}$$

$$\sum_{(i,j)\in A} P_{ij} + q_j = \sum_{(j,i)\in A} P_{ji} \qquad j \in N_P \tag{1h}$$

[2] As commonly done, we will use bold font notation to indicate the entire vectors of variables defined with the same letter (e.g., $\mathbf{x}$ is a vector including all variables $x_{ij}, \forall (i,j) \in A$).

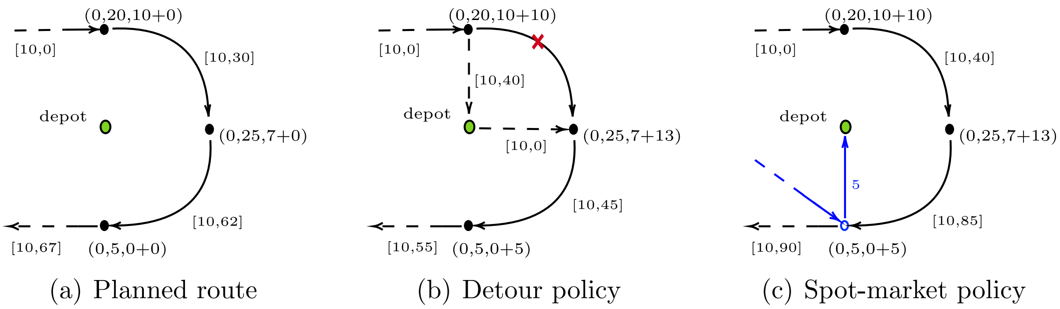

Fig. 1. Different recourse strategies in a route where a vehicle with $Q = 100$ visits three hubs. For each hub j, we indicate in round brackets the tuple $(q_{j+|H|}, q_j, o_j + \tilde{o}_j)$. Next to each arc (i,j), we report in square brackets the delivery and pickup freight transported. In Fig. 1(a), there is no uncertainty and the vehicle can satisfy all the optional demands (hubs denoted by a black filled dot). In Fig. 1(b), due to uncertainty, the vehicle satisfies the demand of the first hub and then decides to return to the depot to unload all collected items (dotted arcs). This way, it gains enough residual capacity to fulfill the optional demand of both the remaining two hubs. In Fig. 1(c), given the same uncertainty, the vehicle decides not to fulfill the optional demand (5 units) of the third hub. To complete the collecting service, however, a spot-market policy is activated (blue arcs). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

$$\sum_{(i,j)\in A} P_{ij} = \sum_{(j,i)\in A} P_{ji} \qquad j \in N_D \qquad (1i)$$

$$\sum_{(0,j)\in A} P_{0j} = 0 \qquad (1j)$$

$$\sum_{(i,0)\in A} D_{i0} = 0 \qquad (1k)$$

$$D_{ij} + P_{ij} \le Q x_{ij} \qquad (i,j) \in A \qquad (1l)$$

$$\sum_{j\in N_P} o_j y_j \ge \alpha \sum_{j\in N_P} o_j \qquad (1m)$$

$$x_{ij} \in \{0,1\} \qquad (i,j) \in A \qquad (1n)$$

$$y_j \in \{0,1\} \qquad j \in N_P \qquad (1o)$$

$$D_{ij}, P_{ij} \ge 0 \qquad (i,j) \in A. \qquad (1p)$$

Eqs. (1b) and (1c) establish that each visited node is served exactly once, whereas Eqs. (1d) are pairing constraints stating that if a vehicle enters a node, it has to leave it. Constraints (1e) impose that the pickup node corresponding to a hub visited for delivery has to be visited even if its mandatory demand is null; since the traveling cost between nodes associated with a single location is null, this requirement does not impact on the objective function but allows the optional pickup to be collected at the operational time even if the mandatory pickup was null. Eqs. (1f)–(1i) represent flow conservation constraints, while constraints (1j)–(1k) guarantee that vehicles start with zero pickup load and return to the depot with zero delivery load, respectively. Constraints (1l) ensure that the vehicle capacity Q is never exceeded and link logically the arc selection to the flow carried by the vehicles. Constraint (1m) guarantees that the vehicles visit a subset of hubs in which it is possible to collect at least α percent of the total deterministic optional pickup demand. Finally, constraints (1n)–(1p) define variables' nature.

Note that, in line with the described operational setting, first-stage solutions fix at day-ahead time the sequence of hubs visited by a vehicle and the amount of delivery quantity carried on each link of the corresponding route. Instead, the optional pickup demand is considered only to fulfill the relative minimum percentage required, whereas it is ignored by flow conservation constraints (1h) and, in turn, does not impact the vehicle capacity at this planning stage.

The objective function (1a) establishes the minimization of the overall operating costs. More precisely, the first term represents the cost of the planned routes while the second term represents the expectation of the costs deriving by applying recourse actions in the future scenarios (namely, the routing cost related to route modifications and spot-market costs). The latter are encapsulated in a function $\delta(\cdot)$ representing the optimal solution value of another problem (the *second-stage problem*), which depends on all the first-stage decisions $(\mathbf{x}, \mathbf{y}, \mathbf{D}, \mathbf{P})$ and the random variable $\tilde{\mathbf{o}}$ and models all the recourse actions already presented. To this aim, let us introduce the following second-stage variables: $\tilde{x}_{ij}$, taking value 1 if a vehicle travels on arc $(i,j) \in A$ after the resolution of the uncertainty (i.e., if it keeps traveling on arc (i,j) as planned by the corresponding decision x_{ij} or if it performs a detour from node i), and 0 otherwise; $\tilde{P}_{ij}$, measuring the (possibly negative) variation of pickup quantity carried on arc $(i,j) \in A$ with respect to the corresponding first-stage decision P_{ij}; $\tilde{z}_j$, taking value 1 if a vehicle fulfills the optional pickup demand of hub $j \in N_P$, and 0 otherwise; $\tilde{\lambda}$, measuring the total optional pickup demand not satisfied; $\tilde{r}_j$, taking value 1 if a node $j \in N_P$ is served through a spot-market service, and 0 otherwise. Then, the second-stage problem can be formulated as follows:

$$\delta(\mathbf{x},\mathbf{y},\mathbf{D},\mathbf{P},\tilde{\mathbf{o}}) := \min_{\tilde{\mathbf{x}},\tilde{\mathbf{P}},\tilde{\mathbf{z}},\tilde{\lambda},\tilde{\mathbf{r}}} \sum_{j\in N_P} \kappa_j \tilde{r}_j + \sum_{(i,j)\in A | i,j\neq 0} (c_{i0} + c_{0j} - c_{ij})(x_{ij} - \tilde{x}_{ij}) \qquad (2a)$$

$$\text{s.t.} \quad \tilde{x}_{ij} \le x_{ij} \qquad (i,j) \in A | i,j \neq 0 \quad (2b)$$

$$\sum_{(i,j)\in A} \tilde{x}_{ij} = y_j \qquad j \in N_P \quad (2c)$$

$$\sum_{(i,j)\in A} \tilde{x}_{ij} = 1 \qquad j \in N | q_j > 0 \quad (2d)$$

$$\sum_{(i,j)\in A} \tilde{x}_{ij} = \sum_{(j,i)\in A} \tilde{x}_{ji} \qquad j \in N \quad (2e)$$

$$\tilde{z}_j \le y_j \qquad j \in N_P \quad (2f)$$

$$\sum_{(i,j)\in A} (P_{ij} + \tilde{P}_{ij}) + q_j + (o_j + \tilde{o}_j)\tilde{z}_j = \sum_{(j,i)\in A} (P_{ji} + \tilde{P}_{ji}) \qquad j \in N_P \quad (2g)$$

$$\sum_{(i,j)\in A} (P_{ij} + \tilde{P}_{ij}) = \sum_{(j,i)\in A} (P_{ji} + \tilde{P}_{ji}) \qquad j \in N_D \quad (2h)$$

$$\sum_{(0,j)\in A} (P_{0j} + \tilde{P}_{0j}) = 0 \qquad (2i)$$

$$\sum_{(j,i)\in A} D_{ji} + P_{j0} + \tilde{P}_{j0} \le Q \qquad j \in N \quad (2j)$$

$$D_{ij} + P_{ij} + \tilde{P}_{ij} \le Q \qquad (i,j) \in A | j \neq 0 \quad (2k)$$

$$P_{ij} + \tilde{P}_{ij} \le Q \tilde{x}_{ij} \qquad (i,j) \in A | i,j \neq 0 \quad (2l)$$

$$\tilde{P}_{i0} \le Q \tilde{x}_{i0} \qquad (i,0) \in A \quad (2m)$$

$$P_{ij} + \tilde{P}_{ij} \ge 0 \qquad (i,j) \in A \quad (2n)$$

$$\sum_{j\in N_P} (o_j + \tilde{o}_j)\tilde{z}_j + \tilde{\lambda} \ge \alpha \sum_{j\in N_P} (o_j + \tilde{o}_j) \qquad (2o)$$

$$\sum_{j\in N_P} (o_j + \tilde{o}_j)\tilde{r}_j \ge \tilde{\lambda} \qquad (2p)$$

$$\tilde{r}_j \leq 1 - \tilde{z}_j \qquad j \in N_P \quad (2q)$$

$$\tilde{x}_{ij} \in \{0,1\} \qquad (i,j) \in A \quad (2r)$$

$$\tilde{P}_{ij} \in \mathbb{R} \qquad (i,j) \in A \quad (2s)$$

$$\tilde{z}_j \in \{0,1\} \qquad j \in N_P \quad (2t)$$

$$\tilde{\lambda} \geq 0 \quad (2u)$$

$$\tilde{r}_j \in \{0,1\} \qquad j \in N_P. \quad (2v)$$

The objective function (2a) minimizes the total spot-market cost to fulfill critical nodes plus the total routing cost depending on the detours implemented. Constraints (2b) impose that it is not possible to travel on an arc that was not planned to be traveled unless such an arc links a node to the depot (i.e., it can be traveled for implementing a detour). Eqs. (2c)–(2e) guarantee that the new routes visit only the hubs already visited in the planned routes. Constraints (2f) state that it is possible to satisfy the optional demand of a hub only if it is visited by a vehicle. Constraints (2g)–(2i) set the flow conservation of goods to pick up in new routes, including the detour arcs. Constraints (2j)–(2m) impose the vehicle capacity restriction on the amount of goods loaded by each vehicle and deny the flow to move along arcs not traveled by any vehicle.[3] Constraints (2n) ensure that the pickup flow does not become negative, while constraint (2o) sets the value of $\tilde{\lambda}$ to the surplus between the total optional demand fulfilled and the minimum percentage required by the company. Constraints (2p)–(2q) ensure that a node j is considered critical only if its optional demand is not fulfilled by the vehicles but is strictly necessary to reach the minimum percentage of optional demand collected. Finally, constraints (2r)–(2t) define the variables' nature.

3.3. Deterministic scenario-based problem

To practically solve the SVRPDDP$_{op}$, we approximate the two-stage SP formulation by a deterministic scenario-based problem (DSP) [48, 49]. Let us consider a finite (but sufficiently large) set $\mathcal{S}$ of scenarios. Each scenario $s \in \mathcal{S}$ is associated with a realization $\tilde{o}^s_j$ of the non-negative variation of optional pickup demand for node $j \in N_P$, and occurs with known probability π^s, such that the standard axiom $\sum_{s\in\mathcal{S}} \pi^s = 1$ is satisfied. Then, the DSP is the following MILP model[4]:

$$\min \sum_{(i,j)\in A} c_{ij}x_{ij} + \sum_{s\in\mathcal{S}} \pi^s \left[\sum_{j\in N_P} \kappa_j \tilde{r}^s_j + \sum_{(i,j)\in A|i,j\neq 0} (c_{i0} + c_{0j} - c_{ij})(x_{ij} - \tilde{x}^s_{ij}) \right] \quad (3a)$$

s.t. constraints (1b)–(1p)

$$\tilde{x}^s_{ij} \leq x_{ij} \qquad (i,j) \in A|i,j \neq 0, s \in \mathcal{S} \quad (3b)$$

$$\sum_{(i,j)\in A} \tilde{x}^s_{ij} = y_j \qquad j \in N_P, s \in \mathcal{S} \quad (3c)$$

$$\sum_{(i,j)\in A} \tilde{x}^s_{ij} = 1 \qquad j \in N|q_j > 0, s \in \mathcal{S} \quad (3d)$$

$$\sum_{(i,j)\in A} \tilde{x}^s_{ij} = \sum_{(j,i)\in A} \tilde{x}^s_{ji} \qquad j \in N, s \in \mathcal{S} \quad (3e)$$

$$\tilde{z}^s_j \leq y_j \qquad j \in N_P, s \in \mathcal{S} \quad (3f)$$

$$\sum_{(i,j)\in A} (P_{ij} + \tilde{P}^s_{ij}) + q_j + (o_j + \tilde{o}^s_j)\tilde{z}^s_j = \sum_{(j,i)\in A} (P_{ji} + \tilde{P}^s_{ji}) \qquad j \in N_P, s \in \mathcal{S} \quad (3g)$$

$$\sum_{(i,j)\in A} (P_{ij} + \tilde{P}^s_{ij}) = \sum_{(j,i)\in A} (P_{ji} + \tilde{P}^s_{ji}) \qquad j \in N_D, s \in \mathcal{S} \quad (3h)$$

$$\sum_{(0,j)\in A} (P_{0j} + \tilde{P}^s_{0j}) = 0 \qquad s \in \mathcal{S} \quad (3i)$$

$$\sum_{(j,i)\in A} D_{ji} + P_{j0} + \tilde{P}^s_{j0} \leq Q \qquad j \in N, s \in \mathcal{S} \quad (3j)$$

$$D_{ij} + P_{ij} + \tilde{P}^s_{ij} \leq Q \qquad (i,j) \in A|j \neq 0, s \in \mathcal{S} \quad (3k)$$

$$P_{ij} + \tilde{P}^s_{ij} \leq Q\tilde{x}^s_{ij} \qquad (i,j) \in A|i,j \neq 0, s \in \mathcal{S} \quad (3l)$$

$$\tilde{P}^s_{i0} \leq Q\tilde{x}^s_{i0} \qquad (i,0) \in A, s \in \mathcal{S} \quad (3m)$$

$$P_{ij} + \tilde{P}^s_{ij} \geq 0 \qquad (i,j) \in A, s \in \mathcal{S} \quad (3n)$$

$$\sum_{j\in N_P} (o_j + \tilde{o}^s_j)\tilde{z}^s_j + \tilde{\lambda}^s \geq \alpha \sum_{j\in N_P} (o_j + \tilde{o}^s_j) \qquad s \in \mathcal{S} \quad (3o)$$

$$\sum_{j\in N_P} (o_j + \tilde{o}^s_j)\tilde{r}^s_j \geq \tilde{\lambda}^s \qquad s \in \mathcal{S} \quad (3p)$$

$$\tilde{r}^s_j \leq 1 - \tilde{z}^s_j \qquad j \in N_P, s \in \mathcal{S} \quad (3q)$$

$$\tilde{x}^s_{ij} \in \{0,1\} \qquad (i,j) \in A, s \in \mathcal{S} \quad (3r)$$

$$\tilde{P}^s_{ij} \in \mathbb{R} \qquad (i,j) \in A, s \in \mathcal{S} \quad (3s)$$

$$\tilde{z}^s_j \in \{0,1\} \qquad j \in N_P, s \in \mathcal{S} \quad (3t)$$

$$\tilde{\lambda}^s \geq 0 \qquad s \in \mathcal{S} \quad (3u)$$

$$\tilde{r}^s_j \in \{0,1\} \qquad j \in N_P, s \in \mathcal{S}, \quad (3v)$$

where the variables $\tilde{x}^s_{ij}$, $\tilde{P}^s_{ij}$, $\tilde{z}^s_j$, $\tilde{\lambda}^s$, and $\tilde{r}^s_j$ have the same decisional meaning as the respective second-stage variables previously introduced, now related to the occurrence of a specific scenario $s \in \mathcal{S}$. All the constraints have already been explained; however, all the second-stage constraints now explicitly appear for each scenario. The objective function now includes the calculation of the second-stage expected value as the summation over all the scenarios of each single second-stage cost weighted by the corresponding probability.

Note that the above DSP allows the creation of subtours on the planned routes defined on variables $\mathbf{x}$ only among nodes having null delivery and pickup demand. When all the deterministic demands of a hub are null, the appearance of a non-null optional pickup demand may force the visiting of some pickup nodes among which the planned route may contain a subtour. This pathological case, however, does not impact the correctness of the formulation, since, in any scenario, the implemented routes based on variables $\tilde{\mathbf{x}}$ are surely cycle-free. Anyway, we decided to fix this inconsistency by implementing ad-hoc valid inequalities within the solution algorithm proposed (see Section 4.2).

4. Solution approach

In the following, we present the mathematical programming-based exact approach developed to solve the SVRPDDP$_{op}$, namely, an Integer L-shaped method (Section 4.1) enhanced with the introduction of several valid inequalities (Section 4.2). All such inequalities and the Integer L-shaped cuts are separated in a branch-and-cut fashion.

4.1. Integer L-shaped approach

We developed an Integer L-shaped (IL) method, leveraging the block diagonal structure of the DSP and its decomposability by scenario once first-stage variables are fixed. The IL method by Laporte and Louveaux [50] is a decomposition algorithm designed to solve mixed-integer two-stage SP problems. It extends the classical L-shaped method by Van Slyke and Wets [51], which was intended for linear programs, to accommodate the non-convex nature of the second-stage value function when integer variables are involved. In this work, we decided to directly use IL and not other MIP frameworks based on the second-stage LP relaxation and the classical L-shaped method since such a relaxation returns very bad bounds. This is mainly due to the fact that relaxing the

[3] Note that, given the first-stage requirement on the arc variables $\mathbf{x}$, also second-stage solutions cannot involve subtours among the nodes.

[4] The expected value calculation resorts to a linear form in the case of a discrete number of realizations of the uncertain demand.

binarity of the recourse decisions basically disrupts the entire solution structure.

The algorithm operates by iteratively solving a so-called *master problem* (MP), which represents a relaxation of the DSP where the second-stage cost function is estimated through the value of a bounded variable $\theta \in \mathbb{R}$. At each iteration of the method, *feasibility cuts* and *optimality cuts* are possibly generated. The former cuts exclude solutions that are feasible for the MP but infeasible for the DSP, while the latter cuts are needed to approximate the second-stage cost (i.e., to provide a strict bound on the optimal value of θ) through an evaluation of the recourse costs given a specific solution of the MP. In our implementation, to exploit the powerful techniques included in the modern solvers and to avoid the need to solve many integer MPs to optimality, we embed the IL cuts separation within a branch-and-cut algorithm. More specifically, we adopted the branch-and-bound framework available in Gurobi through its callback capabilities.

Our MP is as follows

$$\begin{aligned} \min \quad & \sum_{(i,j)\in A} c_{ij}x_{ij} + \theta \\ \text{s.t.} \quad & \text{constraints (1b)–(1p)} \\ & \theta \geq 0, \end{aligned}$$

i.e., a problem pursuing the minimization of an objective function that preserves the first-stage cost but substitutes the second-stage cost with a new variable $\theta \geq 0$, subject to all the first-stage constraints (1b)–(1p). Note that the non-negative requirement for θ is not restrictive since the spot-market cost is, by definition, greater than or equal to zero, while the overall cost of making detours is also always greater than or equal to zero, given the assumption of working on a graph where the triangle inequality holds (see Section 3.1.3).

First, we can observe that each feasible solution of the MP also represents a feasible solution for the DSP, since the DSP implements a relatively complete recourse. Therefore, it is never necessary to add feasibility cuts. On the contrary, to ensure optimality for the DSP, we must dynamically add optimality cuts to the MP during the exploration of the search tree whenever a new feasible solution is found. Let $(\mathbf{x}^v, \mathbf{y}^v, \mathbf{P}^v, \mathbf{D}^v, \theta^v)$ be the vth feasible solution of the MP found during the branch-and-cut, and let $\mathcal{E}^v$ be the value of the expectation of the second-stage cost given the first-stage solution $(\mathbf{x}^v, \mathbf{y}^v, \mathbf{P}^v, \mathbf{D}^v)$, i.e., $\mathcal{E}^v = \mathbb{E}[\delta(\mathbf{x}^v, \mathbf{y}^v, \mathbf{D}^v, \mathbf{P}^v, \tilde{\mathbf{o}})]$. If $\theta^v < \mathcal{E}^v$, we add to the MP the following optimality cut proposed by Laporte and Louveaux [50], which constraints θ with respect to the first-stage binary variables:

$$\theta \geq \theta^v \left(\sum_{(i,j)\in A^v} x_{ij} - \sum_{(i,j)\notin A^v} x_{ij} - |A^v| + \sum_{j\in N_P^v} y_j - \sum_{j\notin N_P^v} y_j - |N_P^v| + 1 \right), \quad (4)$$

where $A^v := \{(i,j) \in A \mid x_{ij}^v = 1\}$ and $N_P^v := \{j \in N_P \mid y_j^v = 1\}$.

In our case, the value of $\mathcal{E}^v$ can be calculated as the second-stage value of the DSP solution in which the first-stage variables are fixed to the value of $(\mathbf{x}^v, \mathbf{y}^v, \mathbf{P}^v, \mathbf{D}^v)$. Note that this value is just an approximation of the true value of $\mathcal{E}^v$ since the DSP considers a finite scenario set $\mathcal{S}$.

4.2. Valid inequalities

Despite the decomposition applied, the MP is still a challenging problem to solve. Hence, in the following, we propose several valid inequalities to enhance its LP relaxation.

First, as suggested by Nagy et al. [12], it is possible to deny visiting the pickup node of a hub before the delivery one by applying the following variable-fixing:

$$x_{j,j+|H|} = 0, \qquad j \in N_P. \quad (5)$$

Second, it is possible to impose a lower bound on the number of vehicles exiting the depot by using the following restrictions:

$$\sum_{(0,j)\in A} x_{0j} \geq \max\left\{ \left\lceil \frac{1}{Q} \sum_{j\in N_D} q_j \right\rceil, \left\lceil \frac{1}{Q} \sum_{j\in N_P} (q_j + \alpha o_j) \right\rceil \right\}. \quad (6)$$

Third, the following *generalized connectivity constraints* (GCCs) can help to improve the linear relaxation of the DSP:

$$\sum_{(i,j)\in A | i\in N', j\notin N'} x_{ij} \geq \begin{cases} y_k & \text{if } k \in N_P \\ 1 & \text{if } k \in N_D | q_k > 0, \end{cases} \qquad N' \subset N, |N'| \geq 2, k \in N'. \quad (7)$$

Given a subset N' of nodes not containing the depot, and a single node k inside the set N', inequalities (7) impose that at least one arc must exit from the set N' if $k \in N_D | q_k > 0$ or if $k \in N_P$ and it is visited (i.e., $y_k = 1$).

All the cuts (5) and (6), which are $|N_P| + 1$ in number, are simply included in the MP formulation from the beginning, while inequalities (7), which are exponentially many in the number of nodes, must be dynamically separated throughout the branch-and-bound tree. Similar to the case of basic *connectivity constraints* for the Capacitated VRP [52], such a separation is based on max-flow/min-cut resolution on an ad-hoc graph, but it also involves the variables related to the visit of pickup nodes (see, e.g., [53,54]). More precisely, let us call $\mathbf{x}^*$ and $\mathbf{y}^*$ the values of the vectors $\mathbf{x}$ and $\mathbf{y}$, respectively, in the optimal solution of the continuous relaxation of the MP or of any of its subproblems generated within the branch-and-bound tree, and let us consider a weighted graph $\bar{G} = (\bar{V}, \bar{A})$ where $\bar{V} = V$, $\bar{A} = A$, and a capacity x_{ij}^* is associated with each arc $(i,j) \in \bar{A}$. Then, given a node $k \notin N_P$ or a node $k \in N_P$ such that $y_k^* > 0$, the most violated inequality (7) corresponds to the partition $(N', V \setminus N')$ associated with a minimum-capacity cut in $\bar{G}$ separating node 0 from k, with $k \in N'$. This is equivalent to computing a maximum flow in $\bar{G}$ from node 0 to node k. The obtained inequality (7) is effective (i.e., it excludes some fractional solutions) only if the corresponding maximum flow is less than 1 when $k \in N_D | q_k > 0$ and than y_k^* when $k \in N_P$. Note that, for each partition $(N', V \setminus N')$ in which (7) is separated for a delivery node $k \in N_D | q_k > 0$, separations for other delivery nodes in N' make no sense since they would lead to redundant cuts. The specific algorithm used to solve each min-cut/max-flow problem is the one developed by Boykov and Kolmogorov [55], which is suggested by the recent literature as the most efficient, and has a polynomial worst-case complexity.

Note that the separation of inequalities (7) for any relaxed problem within the branch-and-bound tree guarantees to eliminate subtours even in the planned routes involving hubs with null deterministic demand, thus fixing the issue mentioned previously.

5. Experimental setting

As far as we know, a benchmark set of instances for the SVRPDDP$_{op}$ does not exist in the literature. Hence, starting from the observation of a case study (presented in Section 5.1), we generate artificial data (Section 5.2) to obtain a large set of realistic instances for the application at hand (Section 5.3).

5.1. Case-study presentation

To obtain realistic data for the SVRPDDP$_{op}$, in particular concerning the operational characteristics of the hubs involved, we analyzed a real-world case of a well-known e-commerce company operating in the city center of Brescia, an Italian mid-size town.

The operator works on a logistics network composed of 16 hubs and a warehouse located approximately 13 km outside the city center. Fig. 2 shows a map of the city [56] in which the 16 hubs are indicated through a location pin. Their precise geo-positions are extracted

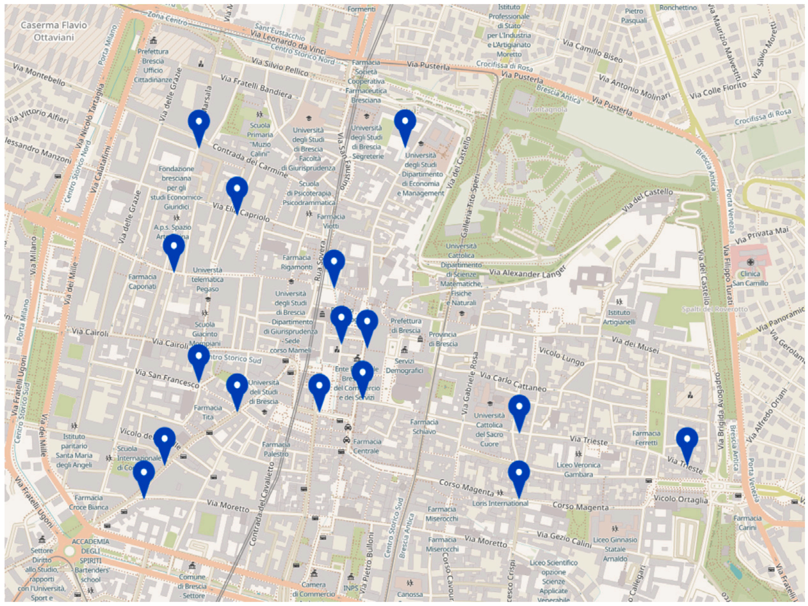

Fig. 2. Locations of hubs in the city center of Brescia.

from LockerMap [57]. At each hub, customers can either drop off a package for return or pick up their own parcel. Instead, the company uses a dedicated fleet of vehicles (standard diesel vans) serving the city center for forward and backward operations between the hubs and the base depot, where all the freight departs and, possibly, comes back. This makes their operations easily comparable to our optimization setting.

We decided to consider fashion products such as clothes, bags, and other foldable accessories (the most involved products in the return logistics); thus, we can assume that each vehicle can carry up to 300 units. Considering the density, age, and habits of the population living in the case-study area, we can also estimate a maximum daily delivery demand per hub of 100 units and a maximum mandatory pickup demand of 25 units.

5.2. Generation of artificial data

Starting from the information gathered in Section 5.1, we describe in the following how we generated diversified but realistic instances.

Concerning the topology of the network, we either consider the specific Brescia case-study network topology T_B (Fig. 2) or artificial network topologies T_A, each differing from the others for the geo-location of the depot and the hubs. In the latter case, the locations of interest are generated randomly in a square area of 20×20 km^2 and the cost c_{ij} of traversing each arc $(i,j) \in A$ is calculated as the Euclidean distance between the coordinates of the hubs associated with nodes i and j, truncated to be an integer. The cost matrix is eventually triangularized in such a way that the triangle inequality holds.

Concerning the demands, each instance (of topology T_B or T_A) is equipped with detailed demand quantities based on the case-study observation. More precisely:

- **mandatory demand**: each hub can equally likely have a positive mandatory delivery demand only, a positive mandatory pickup demand only, both of them positive, or none of them. When a positive mandatory demand q_j exists for node $j \in N$, it is randomly generated from a uniform distribution in $[0,100]$ for a delivery node $j \in N_D$, and in $[0,25]$ for a pickup node $j \in N_P$. As observed in the case study, the mandatory delivery services are likely to represent the largest part of the hubs' requests.
- **optional pickup demand**: each node $j \in N_P$ in each scenario $s \in \mathcal{S}$ also has an optional pickup demand $o_j + \tilde{o}_j^s$. The deterministic part o_j is randomly generated in $[0,225]$, assuming that the optional pickup demand is made up by daily quantities similar to mandatory pickup demand, possibly accumulated for many (9–10) days. The stochastic part $\tilde{o}_j^s$, instead, is drawn from three different Normal probability distributions, truncated to obtain only non-negative values (since we assume that returns requests are not redrawn). The three truncated Normal distributions, modeling different volumes and volatilities given by the occurrence of low, mid, and high seasons, have a mean $\mu = \{20,30,40\}$ and standard deviation $\sigma = \mu/2$, respectively. In the following, these instances are identified by their degree of uncertainty, namely, small (S), medium (M), and large (L), respectively.

Finally, the spot-market cost κ_j is set to double the back-and-forth cost $c_{0j} + c_{j0}$ between the depot and the hub $j \in H$.

5.3. Benchmark sets

In conclusion, we generated two macro sets of benchmark instances. In the first set, containing purely artificial instances, we randomly generated 10 network topologies T_A for each combination of number of hubs $|H| = \{5,10,15,20\}$, demands distributed according to $\{S,M,L\}$, and minimum threshold percentage $\alpha = \{0.5,0.6,0.7,0.8\}$, which means 480 instances in total. In the second set, always based on the case-study topology T_B involving 16 hubs, we randomly generated 5 demand occurrences for each distribution $\{S,M,L\}$ and minimum threshold percentage $\alpha = \{0.5,0.6,0.7,0.8\}$, which means 60 instances in total.

In all our experiments, we use a random sample $\bar{\mathcal{S}}$ of 50 scenarios for each instance. This choice is supported by the stability analysis performed and described in Appendix A.

6. Results and discussion

This section is devoted to analyzing, from an economic and managerial perspective, the results obtained by solving the SVRPDDP$_{op}$ both on the artificial instances generated (Section 6.1) and on the case-study data (Section 6.2). From now on, we will use the following definitions:

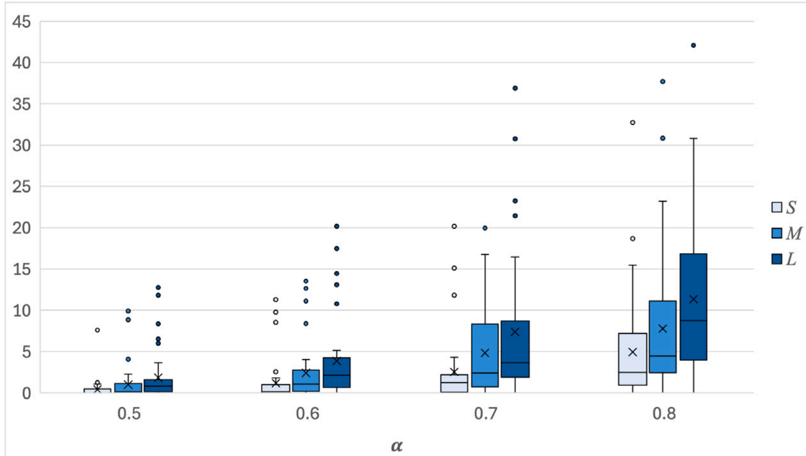

Fig. 3. $EVPI(\%)$ box-plots for different distributions and values of α.

- Recourse Problem (RP) solution: the solution of the DSP, i.e., the proposed SP model approximated through the use of scenario set $\bar{S}$. Let obj_{RP} be its value.
- Expected Value (EV) solution: the solution obtained by solving the DSP in which all the first-stage decisions are fixed to the values that are optimal for the deterministic problem that uses expected values as estimators of the random parameters. Let obj_{EV} be its value.
- Wait-and-See (WS) solution: the *ideal* solution that aggregates all the optimal solutions obtained by solving the DSP when considering every single scenario separately. Let obj_{WS} be the expected value of all these solutions.

All the algorithms are implemented in Java 9, while MILP models are solved by invoking Gurobi v12.0 through its Java APIs. All the computational tests have been run on a computer equipped with an *Intel Xeon* (*Skylake*) processor at 2.29 GHz, 32 GB of RAM, and running at 64-bit *Windows 10 pro* operating system. All the instances have been solved within a threshold time of 4 h, however, the best solutions are provided on average after 650 s and more than the 40% of this time can be saved if we allow a 1% worsening of the optimal value. The IL decomposition scheme contributes up to the 75% to the solution times and the valid inequalities introduced (especially the GCCs) improve the lower bounds by 50% on average. Further computational details can be found in Appendix B.

6.1. Results on artificial instances

In Section 6.1.1, we show the quality of the solutions obtained by our SP approach, in Section 6.1.2, we analyze the impact of the various recourse actions, while in Section 6.1.3, we derive some insights concerning the fleet dimension and the service divisibility.

6.1.1. SP indicators

The box-plots in Figs. 3 and 4 present, for the three probability distributions considered, respectively, the *Expected Value of the Perfect Information* ($EVPI$) and the *Value of the Stochastic Solution* (VSS). The $EVPI$ measures the impact of the uncertainty on our problem and is calculated as the percentage $EVPI(\%) = 100\frac{obj_{RP}-obj_{WS}}{obj_{RP}}$, while the VSS measures the advantage of using our SP model instead of simply using expected values as estimators of the random variables and is calculated as the percentage $VSS(\%) = 100\frac{obj_{EV}-obj_{RP}}{obj_{RP}}$.

The $EVPI(\%)$ naturally increases with the growth of the mean and variance and value of α. In particular, this trend can be observed for any value of α, i.e., the upper whisker of the box increases from 0.42% to 3.65% for $\alpha = 0.5$, from 1.80% to 5.14% for $\alpha = 0.6$, from 4.32% to 16.44% for $\alpha = 0.7$, and from 15.44% to 30.79% for $\alpha = 0.8$. Finally, we can notice that for low α values, the number of outliers is significantly high, pushing the average almost beyond the third quartile. For each distribution, as α increases, the number of outliers decreases substantially, and the third quartile rises. This means that, in instances where a low threshold for optional pickup is required, uncertainty has a highly variable impact, often zero or very low. Conversely, when the threshold rises, the impact of uncertainty becomes consistently greater.

Also the $VSS(\%)$ increases with the mean and variance of the distribution as α varies, although the growth pattern is not strictly monotonic, but exhibits an oscillatory behavior. The average value starts at 1.07%, when $\alpha = 0.5$, and reaches 4.87%, when $\alpha = 0.8$ for the S instances; starts at 1.96% and reaches 4.07% for the M ones; starts at 2.66% and reaches 5.34% for the L ones.

Similar to $EVPI(\%)$, $VSS(\%)$ appears unstable for low α values, gradually stabilizing as the parameter increases. As α grows, the third quartile value rises for any considered normal distribution, exceeding 5.8% when α equals 0.8. We also note that the maximum value of $VSS(\%)$, disregarding the outliers, is always equal to or greater than 6.2% when α is not 0.5. When outliers are considered too, $VSS(\%)$ sometimes reaches very high values, even greater than 25%. Finally, we notice a gradual increase in the median and an elevation of the first quartile value. In particular, when α is 0.8, for the M instances the first quartile is greater than 0. This confirms that the number of instances with very low $VSS(\%)$ values, including null values, progressively decreases.

6.1.2. Impact of the recourse actions on costs and solutions

In Fig. 5 we present, as $|H|$ increases, respectively for each distribution, the percentage weights in RP solution of the three components in the objective function of the model, i.e., the routing cost (*routing*), the spot-market cost (*spot*), and the detour cost (*detour*). There appears to be no strong correlation between the graph size and the percentage of recourse actions, suggesting that is the network topology and the quantities of demands that play a significant role on overall costs rather than dimension of the network. For instances with only 5 hubs, spot recourse emerges as the dominant recourse cost component, likely due to the limited flexibility in smaller networks, where detour possibilities are limited and often ineffective. Conversely, in instances with 10 or more hubs, detour actions become the primary source of recourse cost. This shift indicates that larger, more interconnected networks enable more structured responses to uncertainty, such as planned detours. As the degree of uncertainty increases, we observe that the total cost of recourse actions also increases, although it never exceeds 10% of the

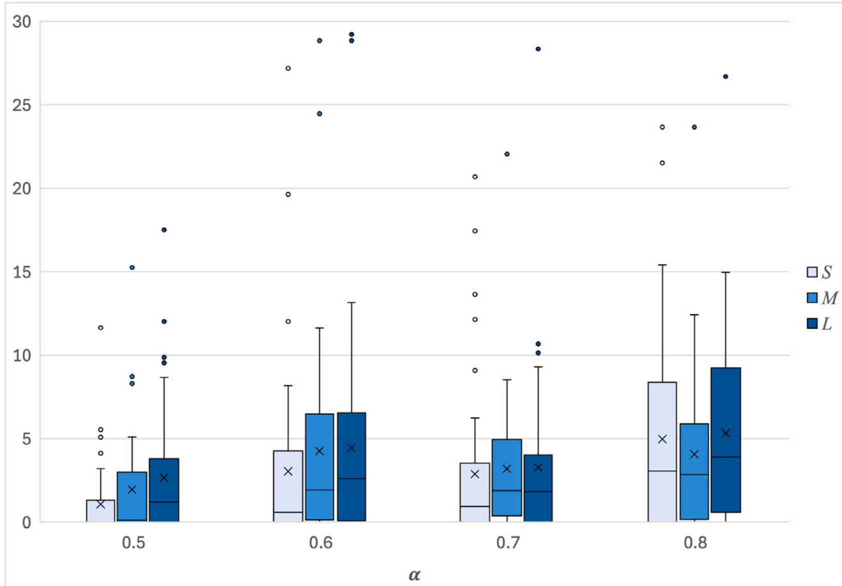

Fig. 4. $VSS(\%)$ box-plots for different distributions and values of α.

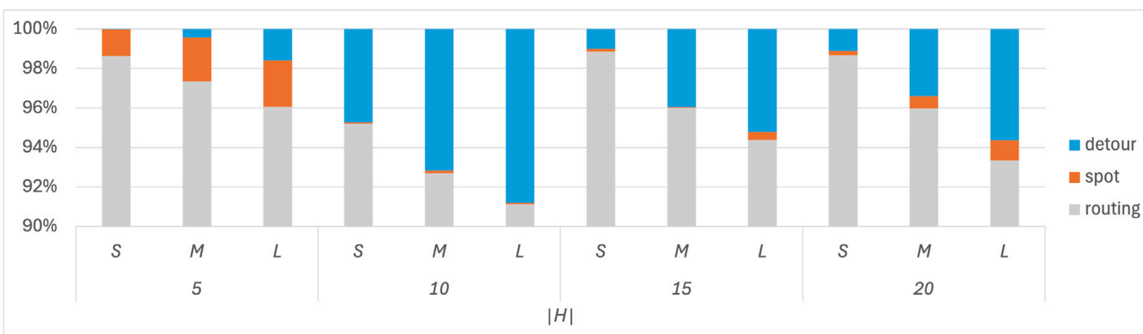

Fig. 5. Cost items (%) in RP solution per number of hubs and distribution type.

total. This highlights the significant impact of the uncertainty on the problem (as already shown in Fig. 3).

In Table 2, we present, as α increases, the percentage distribution of recourse costs in the RP solutions and the corresponding variation in the EV solutions, for each normal distribution. Regarding RP solutions, as uncertainty increases, we observe a progressive decrease in *routing* weight and a corresponding increase in *detour* weight. The former starts with an average of 97.84% and decreases to 93.73%, while the latter begins at 1.72% for S instances and reaches 5.31% for L ones. The *spot* weight, on the other hand, is significantly lower than the other two components and exhibits an average increasing trend, reaching a percentage just below 1% for the L instances. The increasing *detour* weight indicates that as uncertainty grows, the model adjusts the route by increasingly integrating detours, resulting in higher overall costs. The very modest weight of the *spot* component, with relatively insignificant variations, suggests that the stochastic model almost always resorts to detours rather than relying on the extreme option of spot-market services. In EV solutions, we generally observe that routing costs have a slightly lower impact compared to the RP solutions. On the contrary, the percentage weight of the spot market cost increases by up to +3.44% in L instances. The variations of the percentage weight of detour costs, on the other hand, are fluctuating and not particularly significant. This analysis indicates that when uncertainty is not taken into account, there is a risk of relying on the most expensive recourse action, which can lead to a significant increase in total costs.

In the following, we provide additional insights by deepening the analysis of the obtained results from the perspective of the solutions' structure. In Figs. 6 and 7, the histograms show the percentage of instances with spot-market costs greater than 0 and the percentage of instances with detour costs greater than 0, respectively, as α varies and for the three distributions. We can notice that the percentage of instances using the spot-market service in the RP solution is substantially below 30%. On the contrary, in the EV solutions, this percentage increases significantly. When α is 0.7 and 0.8, the value reaches 50% and 70%, respectively. The percentage of RP solutions resorting to detours grows as uncertainty increases, exceeding 90% for L instances with $\alpha = 0.8$. In the EV solutions, there is a slight increase in this percentage but not so pronounced as for spot-market recourse. This likely indicates that the EV solution must directly rely on spot-market actions rather than detours because the latter alone would not make the solution feasible in scenarios with large optional demand variations.

Figs. 8 and 9 display the average number of detours per scenario performed by the RP and EV solutions, respectively, using the three distributions. In Fig. 8, we observe that, for S instances in which the degree of uncertainty is low, the majority of the instances have a detour count well below 0.5. For each level of uncertainty, the average number of detours exhibits a non-decreasing trend as the α parameter increases. Additionally, the upper whisker of the box plots tends to grow, indicating a wider spread of values depending on the topology of instances. The only exception to this pattern is observed in the M instances when $\alpha = 0.8$. This probably suggests that, for particular

Table 2
Percentage of the three cost components in the RP and EV solutions.

	α	*S*			*M*			*L*		
		routing	*spot*	*detour*	*routing*	*spot*	*detour*	*routing*	*spot*	*detour*
RP solutions	0.5	99.72	0.04	0.24	97.81	0.18	2.01	96.58	0.27	3.15
	0.6	97.94	0.60	1.46	96.53	0.81	2.66	95.16	0.78	4.06
	0.7	97.89	0.14	1.97	94.74	0.43	4.83	92.80	0.79	6.41
	0.8	95.83	0.95	3.22	92.96	1.59	5.45	90.36	2.02	7.62
	avg	97.84	0.43	1.72	95.51	0.75	3.74	93.73	0.96	5.31
EV solutions	0.5	−1.39	0.57	0.82	−1.98	1.11	0.87	−1.96	1.69	0.27
	0.6	−2.36	2.67	−0.31	−3.50	2.67	0.83	−3.56	3.01	0.54
	0.7	−3.79	2.80	0.99	−1.50	2.56	−1.06	−3.23	3.08	0.15
	0.8	−5.35	4.84	0.51	−3.35	4.28	−0.93	−5.30	5.99	−0.68
	avg	−3.22	2.72	0.50	−2.58	2.65	−0.07	−3.51	3.44	0.07

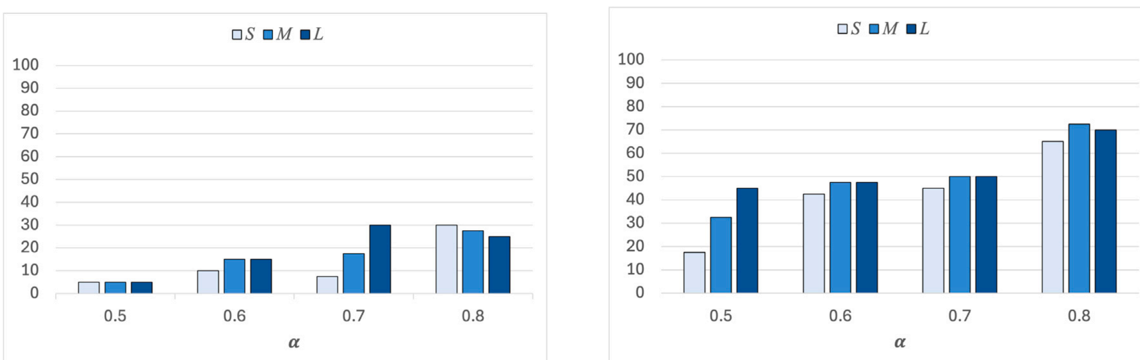

(a) RP solution (b) EV solution

Fig. 6. Number of instances (%) with solutions resorting to spot-market.

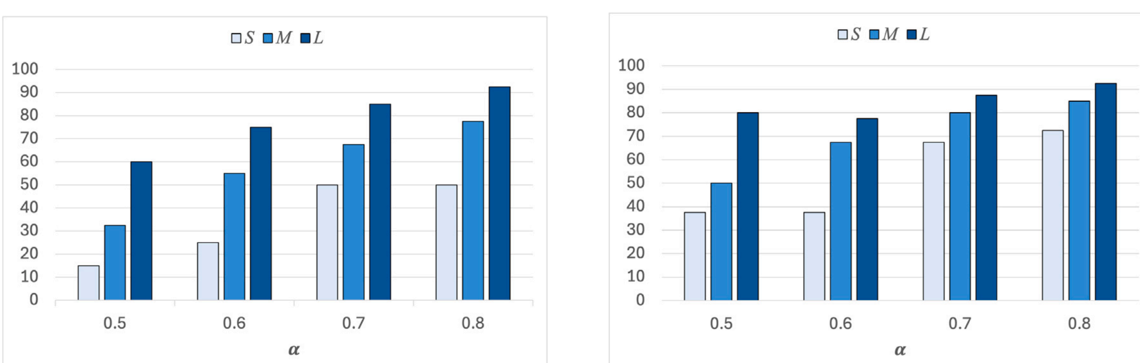

(a) RP solution (b) EV solution

Fig. 7. Number of instances (%) with solutions resorting to detour.

instances, it is possible to reach a saturation point in terms of detour flexibility. In general, as α increases, the planned route in the first stage becomes progressively more subject to modifications, allowing for a greater number of detours to accommodate the required α. Notably, in the L instances for α values of 0.7 and 0.8, the first quartile is strictly greater than zero. This indicates that, in these settings, the vast majority of instances require at least one detour, highlighting the operational impact of high uncertainty combined with high values of optional pickup. In Fig. 9, which illustrates the behavior of EV solutions, we observe a similar trend to that seen in RP solutions, albeit on a reduced scale. When α takes on higher values, particularly in the M and L instances, the box plots become significantly more compact compared to their counterparts in Fig. 8. This compactness reflects a lower variability in the number of detours, suggesting that in many cases, the EV solutions do not rely heavily on detours. Instead, the spot market is often preferred as a recourse action. In conclusion, we can say that, when uncertainty is only roughly estimated, the planning process may fail to fully exploit the potential of detours. This can occur either because detours become too costly under such approximations, or because they are ineffective in meeting the minimum α requirement.

Finally, we propose a further comparison between the RP and EV solutions. Let us define $\mathcal{M}_{RP}$ and $\mathcal{D}_{RP}$ as the sets of instances for which RP solutions rely on the spot-market service and the detour service, respectively. Similarly, we define $\mathcal{M}_{EV}$ and $\mathcal{D}_{EV}$ for EV solution. In Tables 3, 4, and 5, we report the cardinalities (expressed as percentages of the total instances) of certain sets derived from the relationship between the four mentioned above. In columns $|\mathcal{M}_{RP} \cup \mathcal{D}_{RP}|$ and $|\mathcal{M}_{EV} \cup \mathcal{D}_{EV}|$, we respectively indicate the percentage of instances for which RP and EV solutions involve the use of at least one of the two recourse actions. As uncertainty increases, both percentages

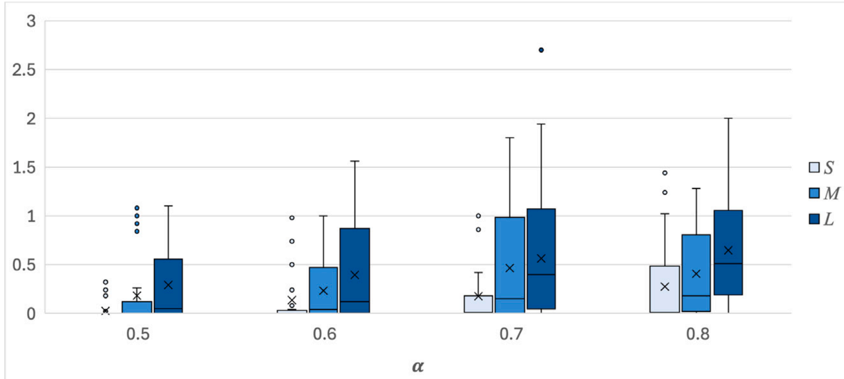

Fig. 8. Average number of detours for RP solutions.

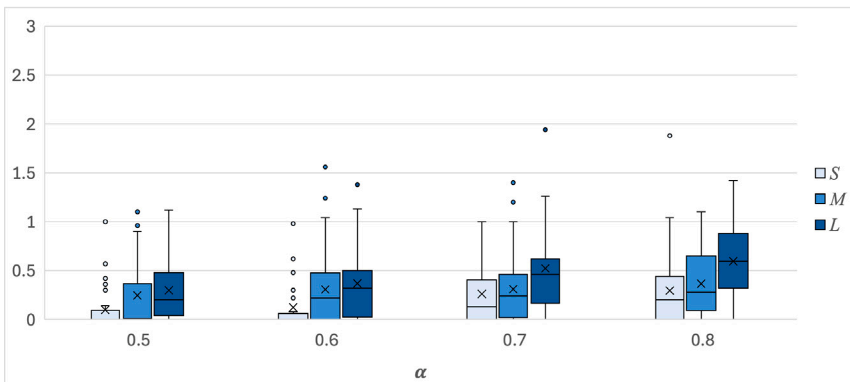

Fig. 9. Average number of detours for EV solutions.

Table 3
Percentages of S instances showing specific recourse features in the RP and EV solutions.

α	$\lvert\mathcal{M}_{RP} \cup \mathcal{D}_{RP}\rvert$	$\lvert\mathcal{M}_{EV} \cup \mathcal{D}_{EV}\rvert$	$\lvert\mathcal{M}_{EV} \setminus \mathcal{M}_{RP}\rvert$	$\lvert\mathcal{M}_{RP} \setminus \mathcal{M}_{EV}\rvert$	$\lvert\mathcal{D}_{EV} \setminus \mathcal{D}_{RP}\rvert$	$\lvert\mathcal{D}_{RP} \setminus \mathcal{D}_{EV}\rvert$
0.5	20.00	50.00	12.50	0.00	22.50	0.00
0.6	30.00	60.00	32.50	0.00	15.00	2.50
0.7	57.50	77.50	37.50	0.00	20.00	2.50
0.8	67.50	92.50	35.00	0.00	22.50	0.00
avg	43.75	70.00	29.38	0.00	20.00	1.25

Table 4
Percentages of M instances showing specific recourse features in the RP and EV solutions.

α	$\lvert\mathcal{M}_{RP} \cup \mathcal{D}_{RP}\rvert$	$\lvert\mathcal{M}_{EV} \cup \mathcal{D}_{EV}\rvert$	$\lvert\mathcal{M}_{EV} \setminus \mathcal{M}_{RP}\rvert$	$\lvert\mathcal{M}_{RP} \setminus \mathcal{M}_{EV}\rvert$	$\lvert\mathcal{D}_{EV} \setminus \mathcal{D}_{RP}\rvert$	$\lvert\mathcal{D}_{RP} \setminus \mathcal{D}_{EV}\rvert$
0.5	37.50	67.50	27.50	0.00	17.50	0.00
0.6	62.50	87.50	32.50	0.00	15.00	2.50
0.7	75.00	92.50	32.50	0.00	12.50	0.00
0.8	87.50	97.50	45.00	0.00	7.50	0.00
avg	65.63	86.25	34.38	0.00	13.13	0.63

grow, albeit to different extents. In fact, $|\mathcal{M}_{RP} \cup \mathcal{D}_{RP}|$ ranges from 20% when $\alpha = 0.5$ for S instances, to 95% when $\alpha = 0.8$ for L ones. The minimum percentage of this range indicates that, with reduced uncertainty and problem complexity, the vast majority of solutions do not need to resort to any recourse action. The maximum percentage of this range reflects the fact that, when the problem becomes more complicated, almost all solutions should involve a recourse action. The $|\mathcal{M}_{EV} \cup \mathcal{D}_{EV}|$ indicator, on the other hand, ranges between 50% and 100%. In particular, Table 3 shows that, even when uncertainty is low, most solutions need to utilize recourse actions if the uncertainty is ignored. This happens in almost all instances, including those with $\alpha = 0.8$. In Tables 4 and 5, we observe that $|\mathcal{M}_{EV} \cup \mathcal{D}_{EV}|$ values over 65% are achieved even with α equal to 0.5, reaching 97.50% and 100% when α is 0.8, respectively. These results align with what we have seen in Figs. 6 and 7. In columns $|\mathcal{M}_{EV} \setminus \mathcal{M}_{RP}|$ and $|\mathcal{M}_{RP} \setminus \mathcal{M}_{EV}|$, we report the number of instances for which only the EV solution resorts to the spot-market and the number of instances for which only the RP solution relies on the spot-market. Similarly, we calculate the values in $|\mathcal{D}_{RP} \setminus \mathcal{D}_{EV}|$ and $|\mathcal{D}_{RP} \setminus \mathcal{D}_{EV}|$, concerning the detour action. A cross-sectional analysis of column $|\mathcal{M}_{EV} \setminus \mathcal{M}_{RP}|$ shows that, especially for M and L instances, in a significant number of cases, the stochastic model tends to avoid spot-market recourse actions, thus limiting the costs. Regarding column $|\mathcal{M}_{RP} \setminus \mathcal{M}_{EV}|$, we observe that there are no cases where the spot-market services are used only in RP, indicating its lack of advantage. Column $|\mathcal{D}_{EV} \setminus \mathcal{D}_{RP}|$, on the other hand, demonstrates that as uncertainty increases, the number of instances for which only

Table 5
Percentages of L instances showing specific recourse features in the RP and EV solutions.

α	$\lvert\mathcal{M}_{RP} \cup \mathcal{D}_{RP}\rvert$	$\lvert\mathcal{M}_{EV} \cup \mathcal{D}_{EV}\rvert$	$\lvert\mathcal{M}_{EV} \setminus \mathcal{M}_{RP}\rvert$	$\lvert\mathcal{M}_{RP} \setminus \mathcal{M}_{EV}\rvert$	$\lvert\mathcal{D}_{EV} \setminus \mathcal{D}_{RP}\rvert$	$\lvert\mathcal{D}_{RP} \setminus \mathcal{D}_{EV}\rvert$
0.5	65.00	92.50	40.0	0.00	20.00	0.00
0.6	77.50	90.00	32.50	0.00	10.00	7.50
0.7	90.00	95.00	20.00	0.00	2.50	0.00
0.8	95.00	100.00	45.00	0.00	0.00	0.00
avg	81.88	94.38	34.38	0.00	8.13	1.88

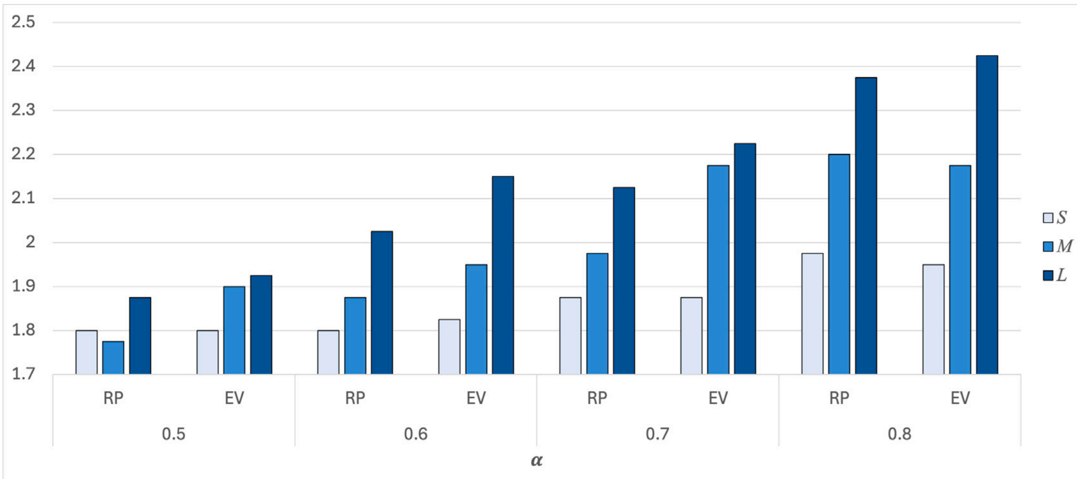

Fig. 10. Average number of vehicles in RP and EV solution for S, M, and L instances.

the EV solution resorts to the detour decreases. The growing complexity of the problem contributes to situations where relying on the detour is essential even when considering uncertainty. In this case, as seen in the previous graphs, RP solutions do so more judiciously, containing costs more effectively than EV. Finally, the last column shows that there exist a few cases where the detour action is used solely by the RP solution: in the vast majority of cases, recourse to the detour is applied in both RP and EV solutions.

6.1.3. Fleet dimension and divisibility

In this last section, we focus on some by-product aspects of the problem optimization, namely, the number of vehicles to use for the service and the possibility of splitting pickup and delivery services at different moments (the divisibility property).

In Fig. 10, we show the average number of vehicles used as α varies. As uncertainty and problem complexity increase, there is a natural and progressive rise in the number of vehicles. Particularly in the RP solutions, the growth is more moderate compared to the EV solutions. This implies that, without considering uncertainty, there is a tendency to use vehicles inefficiently, contributing to cost escalation.

We finally discuss the possibility of splitting the services for pickup and delivery (divisibility) in contrast to imposing both services to be performed by the same vehicle during the same visit to the hub (simultaneity). In the deterministic setting, the convenience of divisible services with respect to simultaneous ones seems to be strictly related to specific features of the instance solved, as exploited in [12] and other later works. More precisely, in their experiments, the service split appears more frequently for hubs placed very near the depot or aggregated in clusters, and for those having large pickup or delivery demand in proportion to the capacity of the vehicles. In our experiments, even if all the above-mentioned features are considered in the instances generated and tested, we obtained instead a negligible number of solutions that involve split hubs and that are strictly better than their simultaneous counterparts. Note that, in our two-stage perspective, a split must be planned in the first stage to appear also in the second-stage solution. However, to show that in our SVRPDDP$_{op}$ there exist strictly convenient split services, in Fig. 11, we compare the SVRPDDP$_{op}$ solution in a specific scenario $\bar{s}$ and its simultaneous counterpart (that we name SVRPSPD$_{op}$),[5] respectively, on a sample instance taken from our benchmark set. The figure shows the operations of the only vehicle involved in the splitting in the SVRPDDP$_{op}$ solution, including 6 hubs (the black dots) and a depot (the green dot). For each hub j, we indicate in round brackets the values corresponding to the tuple $(q_{j+|H|}, q_j, o_j + \tilde{o}_j^{\bar{s}})$. The first-stage route is drawn in solid lines, the detour arcs are drawn in dashed lines, while a red cross marks the arcs no longer traveled in the scenario solution. Next to each arc (i, j) eventually traveled in the scenario solution, we indicate in square brackets the values corresponding to the tuple $[D_{ij}, P_{ij} + \tilde{P}_{ij}]$. We can see that, in the SVRPDDP$_{op}$, the possibility of splitting the service on the node with quantities $(49, 4, 41)$ yields a much shorter route for the considered scenario compared to the case in which splitting is not allowed. In fact, when the service must be simultaneous, the solution operates a detour on the node with quantities $(96, 0, 62)$ (which is on average far from the depot), thus increasing the total cost. More precisely, the cost of the SVRPSPD$_{op}$ solution is approximately 1.55 times greater than the cost of the SVRPDDP$_{op}$ one (namely, 48667 vs 31334). We, therefore, conclude that the specific data uncertainty generated and the recourse actions proposed, together with the relative recourse costs, reduce the convenience of splitting the two services in most cases, but the saving of allowing divisibility can still be significant.

6.2. Results on case-study instances

In this section, we summarize and discuss similar indicators for the solutions of instances directly inspired by the case study (see Section 5.1).

[5] A mathematical formulation for the SVRPSPD$_{op}$ can be easily obtained by adding to formulation (3a)–(3v) a constraint $x_{j,j+|H|} + x_{j+|H|,j} = 1$ for each node $j \in N_P$ such that $q_{j+|H|} > 0$.

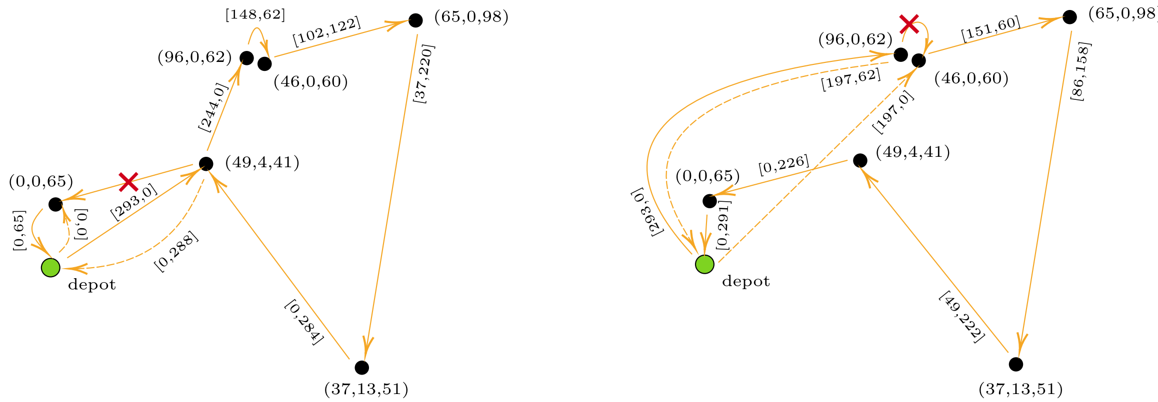

(a) SVRPDDP$_{op}$ (b) SVRPSPD$_{op}$

Fig. 11. A sample instance showing a strictly convenient vehicle route in the SVRPDDP$_{op}$ solution with respect to the corresponding one in the SVRPSPD$_{op}$ solution. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

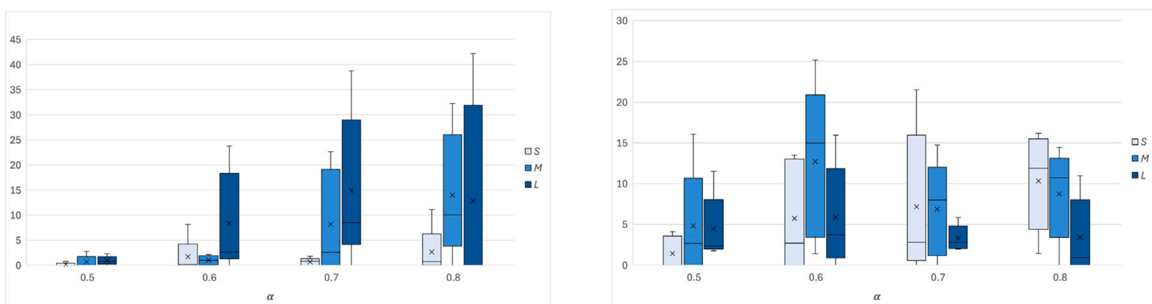

(a) $EVPI(\%)$ (b) $VSS(\%)$

Fig. 12. Case-study (T_B) instances: box-plots of SP indicators for different distributions and α values.

In Fig. 12(a), we show the $EVPI(\%)$ index as α increases and for each level of uncertainty. The trend is similar to that shown in Fig. 3, although the increase is much more pronounced already when α is 0.6 and 0.7. The variability of the index also increases with the minimum amount of optional pickup demand to be collected. The width of the box ranges from 1%–2% when α is 0.5, to a maximum width of 33% when α is 0.8 in the L instances. The median is generally positioned very close to the first quartile of the boxes, indicating that there are still several cases in which the impact of uncertainty is limited. This once again highlights how the topology of the studied network has a strong impact on the problem. Despite the presence of uncertainty, the first-stage route often does not need to be modified unless there are critical changes in demand at specific nodes. In Fig. 12(b), we show the $VSS(\%)$ indicator as α varies and for each level of uncertainty. Unlike what was observed for the T_A instances (see Fig. 4), there is no clear trend in the index. In the S instances, we observe an increase in the average value, while in the M and L instances, the trend is generally oscillatory. Specifically, in the L instances, the average value consistently hovers around 5%, and the width of the box, as well as the third quartile, increases up to α equal to 0.6, then drops sharply at α equal to 0.7, and rises again at α equal to 0.8. In this particular topology, it is evident that the accuracy of the uncertainty estimation through its mean value is not closely correlated with the value of α rather with the degree of uncertainty.

Finally, in Table 6 we report the percentage weights of each cost in the RP and EV solutions, similarly to what was done in Table 2 for the T_A instances. The values and their trends are generally similar to those already shown for an average artificial instance, both in the RP and EV solutions. However, if we consider only the RP, we observe that the detour cost has a significantly greater weight for this topology, especially as uncertainty increases. More precisely, we observe an increase in its average value of about 0.2, 4, and 3.5 percentage points for S, M, and L instances, respectively. Instead, the average weight of the spot-market cost is about 0.5 percentage points lower, becoming null in all cases except two. This is probably due to the specific location of the depot in the case-study topology, which is far from the cluster of hubs located in the city center, resulting in higher necessary detour costs and making the spot-market option almost always inconvenient.

7. Conclusions

Reverse logistics covers issues including the return, repair, recall, defect, damage during transportation, newly introduced products, and product replacement [58]. It is gaining more and more attention due to the environmental and economic implications it imposes. E-commerce returns management represents a specific and important activity within reverse logistics. Many studies in the scientific literature concerning this topic show that vehicle routing used in freight transportation becomes more effective and balanced in a forward and reverse logistics network. Consequently, a successful option for the e-retailers is to design an integrated system where the collection of returns is ensured along with the traditional distribution of goods to hubs. In practice, a vehicle routing problem with deliveries and pickups must be solved.

In this article, we focused on a vehicle routing problem with divisible deliveries and pickups in uncertain settings. More specifically, we analyzed a combinatorial optimization problem where a demand-aggregating hub requiring both a delivery and a pickup service could be serviced, if beneficial, in two separate visits. In order to carry out the service, a fleet of homogeneous vehicles was used. Delivery demands were assumed to be mandatory and deterministic as well as a part of

Table 6
Case-study (T_B) instances: percentage of the three cost components in the RP and EV solutions.

	α	S			M			L		
		routing	*spot*	*detour*	*routing*	*spot*	*detour*	*routing*	*spot*	*detour*
RP solutions	0.5	100.00	0.00	0.00	95.01	0.00	4.99	99.84	0.00	0.16
	0.6	94.52	0.00	5.48	99.83	0.00	0.17	87.91	0.00	12.09
	0.7	100.00	0.00	0.00	87.46	0.00	12.54	90.09	0.00	9.91
	0.8	97.76	0.00	2.24	86.21	0.44	13.35	84.91	2.13	12.96
	avg	98.07	0.00	1.93	92.13	0.11	7.76	90.69	0.53	8.78
EV solutions	0.5	−1.59	0.98	0.61	−4.82	1.60	3.22	−5.17	0.85	4.32
	0.6	−5.60	2.70	2.91	−13.39	5.45	7.94	−6.32	2.69	3.63
	0.7	−7.65	1.15	6.51	−6.74	5.62	1.12	−6.15	1.76	4.39
	0.8	−12.38	5.90	6.49	−10.31	5.20	5.11	−4.47	0.29	4.18
	avg	−6.81	2.68	4.13	−8.81	4.47	4.35	−5.53	1.40	4.13

pickup demands. Additionally, we assumed the presence of optional and stochastic pickup. To model this problem, we employed a two-stage SP model with recourse. In particular, our recourse strategies involved the possibility of redirecting vehicles to the depot to unload already picked-up items (and then continue to fulfill the remaining requests) as well as the possibility of activating spot-market pickup services. To solve the problem, we developed an exact algorithm exploiting specific structural properties, namely, an Integer L-shaped method enhanced with valid inequalities. Extensive computational experiments showed the effectiveness of the modeling paradigm adopted for e-commerce returns management under uncertainty and provided several managerial insights.

Future developments may be carried out in several directions. In particular, to address larger instances in terms of the number of hubs of our stochastic problem, we plan to further enhance the Integer L-shaped method with stronger cuts (e.g., logic-based Benders ones) or to develop different procedures. Of particular interest are those methods that somehow exploit the possibility of efficiently solving the same problem in mono-scenario settings. A good candidate algorithm is the well-known *progressive hedging* [59], which first decomposes the problem per scenario and then tries iteratively to find aggregated first-stage policies. In the context of non-convex problems, as in our case, the progressive hedging can be efficiently used as a meta-heuristic convergence framework (see, e.g., [60,61]).

CRediT authorship contribution statement

Alessandro Gobbi: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Daniele Manerba:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Methodology, Funding acquisition, Conceptualization. **Francesca Vocaturo:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

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

Acknowledgments

The work has been supported by "ULTRA OPTYMAL - Urban Logistics and sustainable TRAnsportation: OPtimization under uncertainTY and MAchine Learning", a PRIN 2020 project funded by the Italian University and Research Ministry, Italy (number 20207C8T9M, website: https://ultraoptymal.unibg.it). The authors also acknowledge the Associate Editor and three anonymous Reviewers for having contributed to improving the quality of this paper.

Appendix A. Stability analysis

To find a sample $\bar{S}$ of scenarios that approximates well enough the involved stochastic variables, we performed several preliminary analyses on a representative subset of the artificial benchmark set, namely, on 40 instances spanning all the characteristics in terms of hubs, stochastic demand distribution, and minimum threshold α.

First, to evaluate the *in-sample stability* of our model, we solved, for each one of the considered instances, the DSP formulation by randomly drawing 10 different scenario sets (for the value of the optional pickup demand variation) with cardinality 10, 20, 30, 40, and 50. Fig. A.13 reports box-plots over the instances of the percentage ratios between the standard deviation and the mean of the objective values[6] over the 10 scenario sets for all the cardinalities of $|\mathcal{S}|$. We can see that, at $|\mathcal{S}| = 50$, very small ratio values are achieved, while up to 40 scenarios, the box reaches the 1%, the whisker exceeds the 2%, and many outliers up to the 4.2% appear. Instead, for $|\mathcal{S}| = 50$, the average ratio settles around 0.5%, the median value settles around 0.3%, and the entire box remains below 0.7%. Moreover, no cases exceed the 2%. For these reasons, we decided to use a random sample $\bar{S}$ of 50 scenarios in all experiments, so as to guarantee that the solution values analyzed do not differ too much depending on the scenario set randomly generated.

Then, to understand whether our model with 50 scenarios is *out-of-sample stable*, we solved, for each one of the considered instances, the DSP formulation with 500 scenarios where the first-stage variables are set to the value obtained considering $\bar{S}$. Given the very large number of scenarios considered, this evaluation is expected to return something very close to the *true* optimal solution values. Fig. A.14 reports box-plots over the instances of the percentage difference between the *true* optimal solution values and those obtained by our model considering 10, 20, 30, 40, and 50 scenarios. We can see that our model with 50 scenarios achieves very good stability, meaning that our random generator is not biased and that the analyzed solution values do not differ much from the *true* ones. In particular, for $|\mathcal{S}| = 50$, we observe an average difference of 0.2%, a median value of 0.09%, and a whisker slightly exceeding 0.6%. Looking at the results for fewer scenarios, it can be said that a fairly good out-of-sample stability is also achievable for $|\mathcal{S}| = 40$; however, adding 10 scenarios more basically halves the width of the box-plot.

Finally, we report in Fig. A.15 the box-plots of the average and worst-case percentage optimality gaps obtained through the Sample Average Approximation (SAA) framework, which has strong relations to the previous stability analyses as shown by Seljom and Tomasgard [62]. For each instance, the average and the worst-case percentage optimality gap $\text{SAAGap}_{\text{avg}}$ and $\text{SAAGap}_{\text{worst}}$ are calculated as $100\frac{|UB-LB_{avg}|}{UB}$ and $100\frac{|UB-LB_{worst}|}{UB}$, respectively, where LB_{avg} and LB_{worst} are the average

[6] Standard deviations have been normalized by the mean values in order to give percentage errors with respect to the real magnitude of the objective function value.

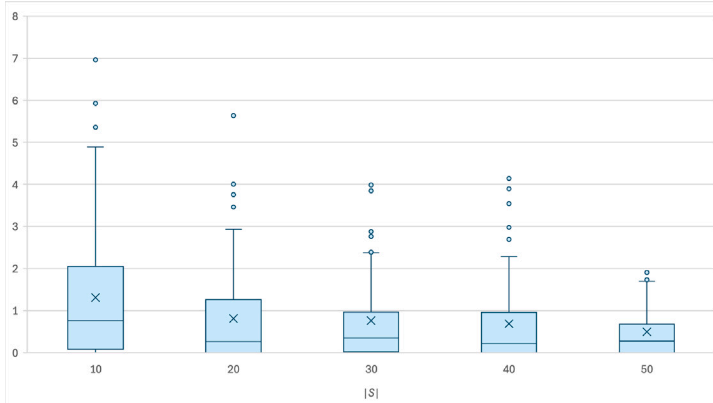

Fig. A.13. In-sample stability: percentage ratio between the standard deviation and the mean of the solution values of 40 instances over 10 different scenario sets of 10, 20, 30, 40, and 50 scenarios.

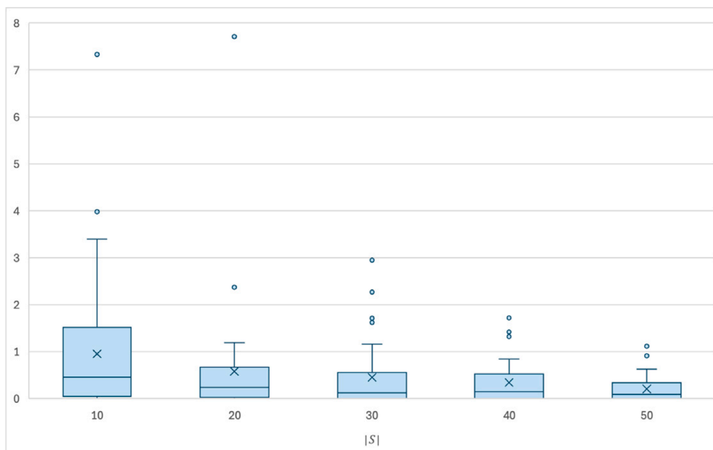

Fig. A.14. Out-of-sample stability: percentage difference between the *true* optimal solution values and those obtained by our model considering 10, 20, 30, 40, and 50 scenarios.

and the minimum over 10 scenario samples (the same introduced for the previous in-sample analysis) of the SP optimal solutions value, while UB is the value obtained against 500 scenarios (the same introduced for the previous out-of-sample analysis) by using the SP optimal solution obtained considering the sample $\bar{S}$. We can see that, apart from a couple of outliers, all the $\text{SAAGap}_{\text{avg}}$ values lie within the 1% and the entire box lies within the 0.5%. Moreover, the average $\text{SAAGap}_{\text{avg}}$ settles around 0.29% while the median is around 0.13%. Regarding $\text{SAAGap}_{\text{worst}}$, apart from a few outliers, all the values lie within the 2%, the average settles around 0.7% while the median is about 0.26%. This further confirms the goodness of our DSP approximation using the scenario sample $\bar{S}$.

Appendix B. Computational assessment of the exact method

In the following, we deepen some computational aspects related to the exact solution method proposed in Section 4.

In Table B.7, we report the average over all the instances with a certain value of α and type of distribution (S, M, L) of the *time-to-best* (ttb) in seconds obtained, which measures the time needed to find the best solution. Moreover, we report a tailored indicator ($\text{ttb}_{1\%}$) that measures the time-to-best percentage dedicated to finding a solution 1% worse than the best one, and the number of optimality cuts (#optCuts) introduced. First, we can see that the basic ttb tends to increase when the value of α increases (apart from the M instances in which there is no clear trend) and when the uncertainty increases. All the values remain always below 1300 s and show an overall average

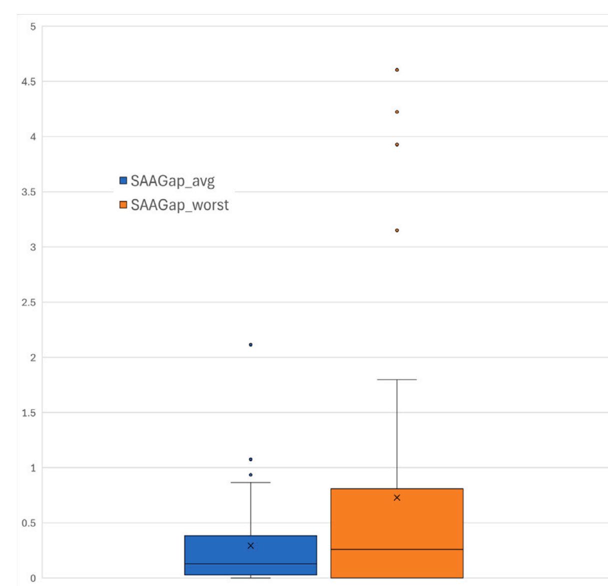

Fig. A.15. SAA-based average and worst-case percentage optimality gap using the sample $\bar{S}$.

Table B.7
Exact method details.

α	S			M			L		
	ttb	$\text{ttb}_{1\%}$	#optCuts	ttb	$\text{ttb}_{1\%}$	#optCuts	ttb	$\text{ttb}_{1\%}$	#optCuts
0.5	56.3	48.68%	196.8	264.7	17.14%	1094.4	307.5	69.75%	2154.7
0.6	111.8	40.02%	834.3	639.3	59.81%	2682.2	1284.6	54.44%	4989.6
0.7	343.6	68.98%	2351.8	1165.3	71.26%	5280.0	1213.0	56.95%	5860.2
0.8	933.1	72.29%	5082.3	369.3	76.95%	6151.2	1056.1	41.08%	7004.6
avg	361.2	57.49%	2116.3	609.6	56.29%	3801.9	965.3	55.56%	5002.3

Table B.8
Performance of the method without IL decomposition.

α	S			M			L		
	ttb_{bc}	$\Delta_{(5)-(6)}$	Δ_{all}	ttb_{bc}	$\Delta_{(5)-(6)}$	Δ_{all}	ttb_{bc}	$\Delta_{(5)-(6)}$	Δ_{all}
0.5	696.3	9.65%	80.62%	1693.5	7.81%	71.44%	2216.1	6.22%	65.88%
0.6	1180.6	7.39%	61.89%	2381.8	5.62%	54.97%	2598.7	4.15%	48.88%
0.7	1438.9	5.76%	46.73%	1988.0	4.10%	40.64%	3072.0	2.80%	35.67%
0.8	2028.5	4.23%	32.72%	2352.5	2.80%	27.56%	3549.8	1.79%	23.86%
avg	1336.1	6.76%	55.49%	2103.9	5.08%	48.65%	2859.2	3.74%	43.57%

of about 650 s. The $\text{ttb}_{1\%}$ indicator is also interesting to analyze. In particular, on average, the $\text{ttb}_{1\%}$ is about 56%, meaning that we could save more than 40% of the time-to-best if we accept a worsening of 1%. This could be useful in the light of a possible heuristic use of a similar decomposition method. As expected, the number of optimality cuts needed grows when the value of α and the uncertainty grow, indicating that harder instances necessitate many more iterations. On average, the number of cuts needed settles between 2 and 5 thousand.

Finally, to assess the contribution of different components of our method, namely, the IL decomposition scheme and the proposed valid inequalities, we study the behavior of the algorithm in the case that no decomposition is applied (i.e., if the complete DSP represents the MP and, therefore, it is directly solved via a branch-and-cut enhanced by the valid inequalities). For this pure branch-and-cut version of the algorithm (*bc*), Table B.8 reports, averaged again over all the instances with a certain value of α and type of distribution, the time-to-best (ttb_{bc}) the percentage improvement of the problem lower bound at the root node of the decision tree when adding only valid inequalities (5) and (6) ($\Delta_{(5)-(6)}$) and when GCCs (7) are also separated (Δ_{all}), respectively. Concerning ttb_{bc}, growing trends similar to those of ttb for the complete method can be observed. However, the average times are considerably larger, with an average CPU effort of about 2100 s. This means that the IL decomposition applied and the corresponding optimality cut generation are worth implementing, with time-to-best savings up to 75%. Instead, the remaining indicators state that the valid inequalities introduced give a net contribution to the value of the LP relaxation of the MP, on average between 43% and 55%, and with some peaks up to 70%–80%. The contribution tends to decrease a bit when the uncertainty increases, and more when the value of α increases. This is in line with the time-to-best results, since worse lower bounds generally correspond to instances harder to solve. As expected, the larger part of the cuts' contribution is given by the GCCs, which directly intervene in the structure of the routes.

Data availability

Data will be made available on request.

References

[1] Lone S, Weltevreden JWJ. European E-commerce report. 2022, URL https://ecommerce-europe.eu/wp-content/uploads/2022/06/CMI2022_FullVersion_LIGHT_v2.pdf, Last access: 2024-12-31.

[2] Buldeo Rai H. The net environmental impact of online shopping, beyond the substitution bias. J Transp Geogr 2021;93:103058.

[3] Buldeo Rai H, Touami S, Dablanc L. Autonomous e-commerce delivery in ordinary and exceptional circumstances. The French case. Res Transp Bus Manag 2022;45(Part A):100774.

[4] Allen J, Piecyk M, Piotrowska M, McLeod F, Cherrett T, Ghali K, Nguyen T, Bektaş T, Bates O, Friday A, Wise S, Austwick M. Understanding the impact of e-commerce on last-mile light goods vehicle activity in urban areas: The case of London. Transp Res Part D: Transp Environ 2018;61:325–38.

[5] Zhang M, Pratap S, Zhao Z, Prajapati D, Huang GQ. Forward and reverse logistics vehicle routing problems with time horizons in B2C e-commerce logistics. Int J Prod Res 2021;59(20):6291–310.

[6] National Retail Federation. 2023 consumer returns in the retail industry. 2023, https://cdn.nrf.com/sites/default/files/2024-01/Customer_Returns_Report_2023_Final.pdf, Last access: 2024-12-31.

[7] Bergmann FM, Wagner SM, Winkenbach M. Integrating first-mile pickup and last-mile delivery on shared vehicle routes for efficient urban e-commerce distribution. Transp Res Part B: Methodol 2020;131:26–62.

[8] Jenkins A. A guide to reverse logistics: How it works, types and strategies. 2021, URL https://www.netsuite.com/portal/resource/articles/inventory-management/reverse-logistics.shtml, Last access: 2024-12-31.

[9] Gobbi A, Mansini R, Moreschini L, Ranza F. Optimal locker placement for budget-constrained crowdshipping in public transport networks. IFAC-PapersOnLine 2025;59(10):3041–6.

[10] Wandelt S, Wang S, Sun X. A literature review on hub location-routing models and their solution techniques. Comput Oper Res 2025;173:106861.

[11] Mukundan G. E-commerce returns management: Processes & best practices. 2024, https://hiverhq.com/blog/e-commerce-returns-management, Last access: 2024-12-31.

[12] Nagy G, Wassan NA, Speranza MG, Archetti C. The vehicle routing problem with divisible deliveries and pickups. Transp Sci 2015;49(2):271–94.

[13] Branke J, Mattfeld DC. Anticipation and flexibility in dynamic scheduling. Int J Prod Res 2005;43(15):3103–29.

[14] McCue I. It's raining returns. Here are 7 ways to wring out revenue. 2020, URL https://www.netsuite.com/portal/resource/articles/ecommerce/returns.shtml, Last access: 2024-12-31.

[15] King AJ, Wallace SW. Modeling with stochastic programming. New York, USA: Springer; 2012.

[16] Santos MJ, Jorge D, Ramos T, Póvoa A. Green reverse logistics: Exploring the vehicle routing problem with deliveries and pickups. Omega 2023;118:102864.

[17] Jargalsaikhan B, Romeijnders W, Roodbergen KJ. A compact arc-based ILP formulation for the pickup and delivery problem with divisible pickups and deliveries. Transp Sci 2021;55(2):336–52.

[18] Li J, Cang L, Wu Y, Zhang Z. Two-echelon collaborative many-to-many pickup and delivery problem for agricultural wholesale markets with workload balance. Omega 2025;130:103164.

[19] Wassan NA, Nagy G. Vehicle routing problem with deliveries and pick-ups: Modelling issues and meta-heuristics solution approaches. Int J Transp 2014;2(1):95–110.

[20] Mingozzi A, Giorgi S, Baldacci R. An exact method for the vehicle routing problem with backhauls. Transp Sci 1999;33(3):315–29.

[21] Ropke S, Pisinger D. A unified heuristic for a large class of vehicle routing problems with backhauls. European J Oper Res 2006;171(3):750–75.

[22] Koç Ç, Laporte G. Vehicle routing with backhauls: Review and research perspectives. Comput Oper Res 2018;91:79–91.

[23] Wade AC, Salhi S. An investigation into a new class of vehicle routing problem with backhauls. Omega 2002;30(6):479–87.

[24] Tasan AS, Gen M. A genetic algorithm based approach to vehicle routing problem with simultaneous pick-up and deliveries. Comput Ind Eng 2012;62(3):755–61.
[25] Subramanian A, Uchoa E, Pessoa AA, Ochi LS. Branch-cut-and-price for the vehicle routing problem with simultaneous pickup and delivery. Optim Lett 2013;7(7):1569–81.
[26] Koç Ç, Laporte G, Tükenmez İ. A review of vehicle routing with simultaneous pickup and delivery. Comput Oper Res 2020;122:104987.
[27] Gribkovskaia I, Halskau Ø, Laporte G, Vlček M. General solutions to the single vehicle routing problem with pickups and deliveries. European J Oper Res 2007;180(2):568–84.
[28] Hoff A, Gribkovskaia I, Laporte G, Løkketangen A. Lasso solution strategies for the vehicle routing problem with pickups and deliveries. European J Oper Res 2009;192(3):755–66.
[29] Polat O. A parallel variable neighborhood search for the vehicle routing problem with divisible deliveries and pickups. Comput Oper Res 2017;85:71–86.
[30] Bruck BP, Iori M. Non-elementary formulations for single vehicle routing problems with pickups and deliveries. Oper Res 2017;65(6):1597–614.
[31] Gribkovskaia I, Laporte G, Shyshou A. The single vehicle routing problem with deliveries and selective pickups. Comput Oper Res 2008;35(9):2908–24.
[32] Gutiérrez-Jarpa G, Marianov V, Obreque C. A single vehicle routing problem with fixed delivery and optional collections. IIE Trans 2009;41(12):1067–79.
[33] Bruck BP, dos Santos A G. Hybrid approach for the multiple vehicle routing problem with deliveries and selective pickups. In: 12th international conference on hybrid intelligent systems. HIS, 2012, p. 265–70.
[34] Santos MJ, Jorge D, Bonomi V, Ramos T, Póvoa A. Enhancing logistics through a vehicle routing problem with deliveries, pickups, and backhauls. Int Trans Oper Res 2024.
[35] Shapiro A, Dentcheva D, Ruszczyński A. Lectures on stochastic programming: modeling and theory. Philadelphia, USA: MOS-SIAM Series on Optimization; 2009.
[36] Oyola J, Arntzen H, Woodruff DL. The stochastic vehicle routing problem, a literature review, part II: solution methods. EURO J Transp Logist 2017;6(4):349–88.
[37] Oyola J, Arntzen H, Woodruff DL. The stochastic vehicle routing problem, a literature review, part I: models. EURO J Transp Logist 2018;7(3):193–221.
[38] De Maio A, Laganà D, Musmanno R, Vocaturo F. Arc routing under uncertainty: Introduction and literature review. Comput Oper Res 2021;135:105442.
[39] Laporte G, Musmanno R, Vocaturo F. An adaptive large neighbourhood search heuristic for the capacitated arc-routing problem with stochastic demands. Transp Sci 2010;44(1):125–35.
[40] Zhao W, Zhang R, Luo K, Wang W. A two-stage stochastic programming method for a heterogeneous vehicle routing problem with time windows and stochastic demand. Expert Syst Appl 2025;291:128463.
[41] Beraldi P, Ghiani G, Musmanno R, Vocaturo F. Efficient neighborhood search for the probabilistic multi-vehicle pickup and delivery problem. Asia-Pac J Oper Res 2010;27(3):301–14.
[42] Zhang Y, Alshraideh H, Diabat A. A stochastic reverse logistics production routing model with environmental considerations. Ann Oper Res 2018;271:1023–44.
[43] Shuang Y, Diabat A, Liao Y. A stochastic reverse logistics production routing model with emissions control policy selection. Int J Prod Econ 2019;213:201–16.
[44] Wollenberg N. Decision making under uncertainty in routing problems for reverse logistics (Ph.D. thesis), Universität Duisburg-Essen; 2015.
[45] Zhu L, Sheu J-B. Failure-specific cooperative recourse strategy for simultaneous pickup and delivery problem with stochastic demands. European J Oper Res 2018;271:896–912.
[46] Zhang W, Chen Z, Zhang S, Cai Y. Dynamic multi-stage failure-specific cooperative recourse strategy for logistics with simultaneous pickup and delivery. Soft Comput 2021;25:3795–812.
[47] Yang W-H, Mathur K, Ballou RH. Stochastic vehicle routing problem with restocking. Transp Sci 2000;34:99–112.
[48] Birge JR, Louveaux FV. Introduction to stochastic programming. New York: Springer Verlag; 1997.
[49] Wallace SW, Ziemba WT. Applications of stochastic programming. Society for Industrial and Applied Mathematics; 2005.
[50] Laporte G, Louveaux FV. The integer L-shaped method for stochastic integer programs with complete recourse. Oper Res Lett 1993;13(3):133–42.
[51] Van Slyke RM, Wets R. L-shaped linear programs with applications to optimal control and stochastic programming. SIAM J Appl Math 1969;17(4):638–63.
[52] Lysgaard J, Letchford AN, Eglese RW. A new branch-and-cut algorithm for the capacitated vehicle routing problem. Math Program 2004;100(2):423–45.
[53] Gobbi A, Manerba D, Mansini R, Zanotti R. A kernel search for a patient satisfaction-oriented nurse routing problem with time-windows. IFAC-PapersOnLine 2019;52(13):1669–74.
[54] Gobbi A, Manerba D, Mansini R, Zanotti R. Hybridizing adaptive large neighborhood search with kernel search: a new solution approach for the nurse routing problem with incompatible services and minimum demand. Int Trans Oper Res 2023;30(1):8–38.
[55] Boykov Y, Kolmogorov V. An experimental comparison of min-cut/max- flow algorithms for energy minimization in vision. IEEE Trans Pattern Anal Mach Intell 2004;26(9):1124–37.
[56] OpenStreetMap, URL https://www.openstreetmap.org, Last access: 2025-06-01.
[57] LockerMap. URL https://lockermap.com, Last access: 2025-06-01.
[58] Desticioglu B, Hatice Calipinar H, Ozyoruk B, Koc E. Model for reverse logistic problem of recycling under stochastic demand. Sustainability 2022;14:4640.
[59] Rockafellar RT, Wets RJ-B. Scenarios and policy aggregation in optimization under uncertainty. Math Oper Res 1991;16(1):119–47, [Accessed 2024-05-06].
[60] Manerba D, Perboli G. New solution approaches for the capacitated supplier selection problem with total quantity discount and activation costs under demand uncertainty. Comput Oper Res 2019;101:29–42.
[61] Giusti R, Manerba D, Tadei R. Multiperiod transshipment location–allocation problem with flow synchronization under stochastic handling operations. Networks 2021;78(1):88–104.
[62] Seljom P, Tomasgard A. Sample average approximation and stability tests applied to energy system design. Energy Syst 2021;12:107–31.