Computers & Operations Research 173 (2025) 106858

Contents lists available at ScienceDirect

Computers and Operations Research

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

An effective mitigation strategy to hedge against absenteeism of occasional drivers

Simona Mancini [a,b,*], Margaretha Gansterer [b], Chefi Triki [c]

[a] University of Palermo, Department of Engineering, Palermo, Italy
[b] University of Klagenfurt, Department of Operations, Energy, and Environmental Management, Universitätsstraße 65-67, 9020 Klagenfurt, Austria
[c] University of Kent, Kent Business School, Canterbury, UK

ARTICLE INFO

Keywords:
Last-mile delivery
Occasional drivers
Drivers absenteeism
Mitigation policy

ABSTRACT

Companies can use occasional drivers to increase efficiency on last-mile deliveries. However, as occasional drivers are freelancers without contracts, they can decide at short notice whether they perform delivery requests. If they do not perform their tasks, this is known as driver absenteeism, which obviously disrupts the operations of companies. This paper tackles this problem by developing an auction-based system, including a mitigation strategy to hedge against the absenteeism of occasional drivers. According to this strategy, a driver can bid not only for serving bundles but also to act as a reserved driver. Reserved drivers receive a fee to ensure their presence but are not guaranteed to be assigned to a specific bundle. The problem is modeled as a two-stage stochastic problem with recourse activation. To solve this problem, this paper develops a self-learning matheuristic (SLM) and an iterated local search (ILS) that exploits SLM as a local search operator. Through an extensive computational study, this paper shows the clear dominance of the newly proposed approach in terms of solution quality, run times, and customers' perceived quality of service compared against three different deterministic approaches. The Value of the Stochastic Solution, a well-known stochastic parameter, is also analyzed. Finally, the identikit of the perfect reserved driver, based on data observed in optimal solutions, is discussed.

1. Introduction and motivation

E-commerce is revolutionizing the way businesses operate, providing platforms to sell their products and services online all the year aroound. This transformation has enabled sellers to reach global customers, breaking down geographical barriers and opening up new markets. E-commerce has also made shopping more convenient for consumers, allowing them to purchase and receive their products at home. With the rapid growth of technology and online shopping, e-commerce is becoming increasingly essential for businesses to remain competitive in today's marketplace. Hence, last-mile delivery of parcels is also becoming a huge business involving an ever increasing number of delivery and courier companies. For example, the USA currently has almost half a million delivery service companies expected to generate 157 billion dollars in 2023 with an annualized growth rate of 6.5% during the period 2018–2023 (Ibisworld, 2023). On a global scale, the percentage of online retail shopping currently amounts to 21% and is expected to increase to 24% in 2026 (Statista, 2023). This growth is expected to be generated by an unspecified number of online stores, which can reach 24 million Shepherd, 2023. Likewise, the global last-mile delivery market size is expected to be 455 billion dollars in 2023 and is poised to grow by 166 billion in 2027 (Ibisworld, 2023).

While e-commerce giants (such as Amazon, ebay, or AliExpress) have their own established distribution network, fulfilling such a high number of deliveries generates several challenging issues for small- and medium-sized companies. Running an owned delivery fleet could be extremely costly for local delivery companies due to the high capital and operations costs. Furthermore, in last mile delivery, demand is far from being constant or regular, but often characterized by frequent peaks of demand concentrated during special seasons (e.g., Christmas). So, maintaining a fixed fleet for the whole year, which is also able to handle demand peaks, would be too costly for the company. On the other hand, an undersized fleet would force the company to reject several requests during high-demand peaks, with a consequent huge loss of revenue.

One of the attractive solutions adopted to reduce delivery cost is to employ occasional drivers (ODs), who are freelancers without a

* Corresponding author at: University of Klagenfurt, Department of Operations, Energy, and Environmental Management, Universitätsstraße 65-67, 9020 Klagenfurt, Austria.

E-mail address: simona.mancini@aau.at (S. Mancini).

https://doi.org/10.1016/j.cor.2024.106858
Received 1 February 2024; Received in revised form 22 August 2024; Accepted 26 September 2024
Available online 10 October 2024

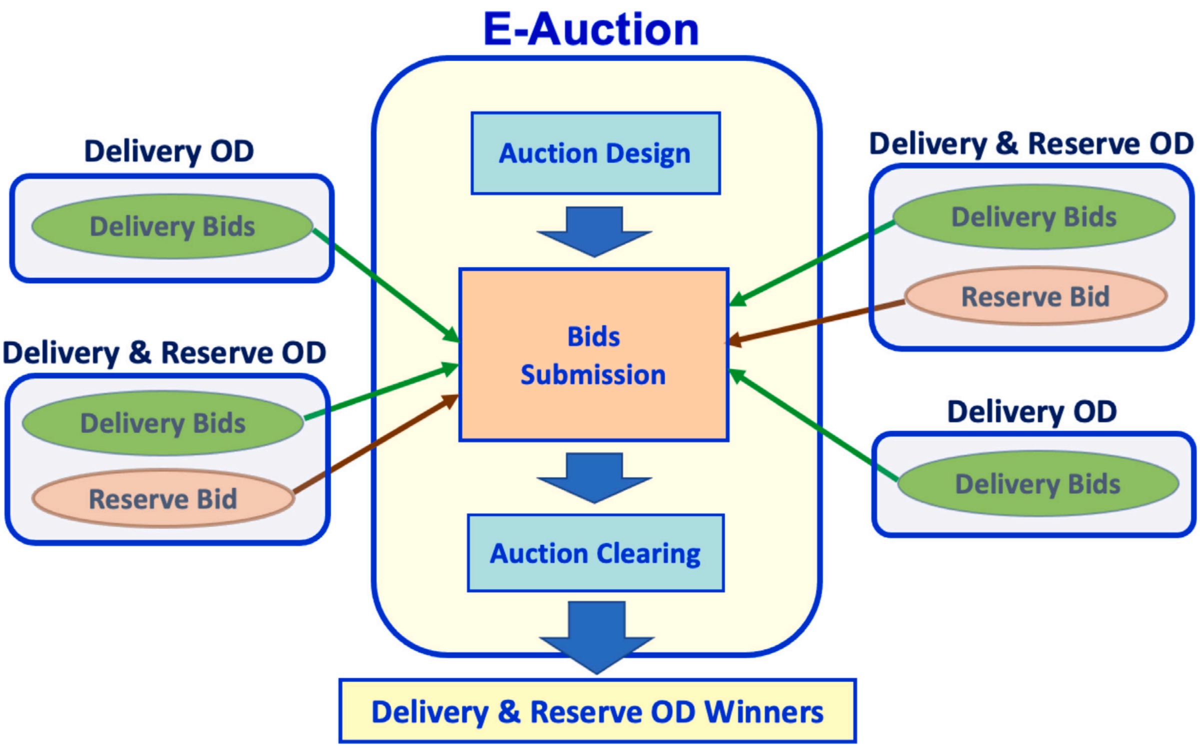

Fig. 1. Auction framework with ODs submitting delivery and/or reservation bids.

fixed contract (Archetti et al., 2016; Triki, 2021; Bortolini et al., 2022; Mancini and Gansterer, 2022; Wang et al., 2023). ODs are usually willing to perform one or more deliveries, according to their availability, for a little compensation. Thus, ODs are cheaper than corporate drivers because they do not imply any fixed costs. Traditionally, ODs are paid a flat price for every delivery. Alternatively, they are compensated according to a distance-based catalogue, which often does not reflect the willingness of the driver to serve any of the requests and neither takes into account the synergies among the assigned deliveries (Le and Ukkusuri, 2019; Cheng et al., 2023). To avoid this kind of distortion, this paper proposes a mechanism based on the employment of Combinatorial Auctions (CAs). In such a delivery system, the company establishes an online platform, on which it advertises a set of potentially overlapping bundles of deliveries to be fulfilled (Rechavi and Toch, 2022). The available ODs join the platform, select the bundles they are willing to serve and submit their bids to the auction with the associated prices (Triki et al., 2023). Every OD can bid on an unlimited number of bundles, but to avoid the over-exposure effect, only one bid per OD can be accepted. After receiving all the bids, the company clears the auction to identify the winning bids and communicates the request assignments to the drivers through the platform (see Fig. 1). The specific design of the CA to be implemented for the delivery-OD assignments is crucial for the success of the platform operations. We consider here a single-round, sealed-bid, first-price auction. This means that the ODs submit secret bids whose prices cannot be seen by the other ODs and the winners will be paid their suggested prices. With this design, collusion among ODs can be avoided and the ODs are encouraged to bid on their real costs to increase their chances of resulting successful in the auction.

The company incurs a penalty for every unfulfilled request, whose value is usually request-dependent. The goal of the company is to minimize the overall cost, given by the sum of the cost of assigning the bundles to the ODs plus the penalty costs.

The drawback of this system is the reliability of the ODs. Indeed, in practice, some ODs bid on different platforms and, if they win in more than one, they choose the most convenient deliveries, withdrawing all the other assignments. Some others might submit bids but then do not follow up the outcomes of the auction because of other obligations. This situation is because ODs do not have regular contracts with the company. Hence, they also do not need to promptly communicate their absenteeism and neither can they be forced to pay no-showing penalties.

The absenteeism of ODs could potentially induce a considerable loss of profits for the company because of the missed service of some deliveries and the consequent refund of unserved customers. Furthermore, the decrement of the company's reliability reduces the quality of service perceived by the customers, who may decide to rely on a competing company having a higher reliability. Hence, the damage from not fulfilling some requests not only affects the current profits but can also negatively impact future demands (Figliozzi and Zhang, 2009).

The absenteeism of ODs falls under the broad field of service disruption for which several approaches such as robustness, flexibility, and redundancy have been proposed in the literature to reduce its effect (Albertzeth and Pujawan, 2018; Pahwa and Jaller, 2023). However, the specific aspect of the absenteeism of drivers has itself received very limited attention in the context of short-term delivery planning (Boysen et al., 2021). In this study, a novel mitigation strategy is proposed, based on the concept of the reservation of ODs to proactively hedge against their possible absenteeism. The company allows the ODs to submit not only delivery bids but also reserve bids (backups) to guarantee their availability as a recourse resource during the delivery period. Even though compensated for their availability, reserved drivers are not assigned any delivery task during the planning phase but may be requested to fulfill one of the bundles left unserved because of the absenteeism of other ODs. In this case, the correspondent delivery bidding price must be paid in addition to the reservation cost. Otherwise, i.e. if no disruption occurs and the service of the reserved OD is not needed, a compensation corresponding to her reserve bidding price must be paid anyway for ensuring the availability. Every OD can choose to submit only delivery bundles or also a bid to act as a reserved driver (see Fig. 1).

Reserve ODs do not need to remain idle while the decision maker is waiting to observe the random scenario that will materialize. They only need to give priority to serve any assigned delivery in case of absence of regular ODs. In some cases, no recourse deliveries are needed and then the reserve ODs will ensure some extra earning without being asked to perform any delivery task, a fact that makes our approach very attractive for them.

It is worthwhile mentioning that the idea of involving dedicated reserve ODs in the last-mile delivery is not a completely new approach. Indeed, our idea was inspired by the power systems field that, for decades, has successfully relied on implementing the auction mechanism to ensure the exchange of electricity between producers and consumers (see for example Beraldi et al., 2004 and Musmanno et al., 2010). Electricity market operators typically run reserve auctions in which the bidding producers accept to remain available to satisfy peak demand in case of sudden increase in the network's consumption. While our study is clearly focused on the operational implications of the proposed system, legal aspects are not covered here. These include potentially regional restrictions in regards of data privacy, market regulations, or labor laws. To the best of our knowledge, from legal point of view, the newly proposed system does not differ from established platform-based crowd delivery systems.

Clearly, during the phase of CA clearing and bundle assignment, the extent of absenteeism among the ODs is not known in advance. Consequently, such an absenteeism level will be represented as a discrete random variable and the problem can be, thus, formulated as a two-stage stochastic model. The first-stage decisions consist in assigning the delivery bundles to the ODs and in selecting the backup drivers. In the second-stage, once the ODs presence is revealed, the unfulfilled delivery bundles are assigned to the reserved drivers, and the subset of customers that still remain unserved is identified. The goal is to minimize the overall costs related to both the first-stage decisions and the expected value of the second-stage decisions over all the possible scenarios.

To the best of our knowledge, this is the first work explicitly dealing with the problem of the absenteeism of ODs in the context of last-mile delivery, with which the related literature is reviewed. Sections 3 and 4 present the deterministic and the stochastic problem, respectively. The proposed solution approach is discussed in Section 5 followed by the computational study presented in Section 6. Finally, concluding remarks and possible future avenues of research are reported in Section 7.

2. Literature review

The research topic finds its roots in the field of (i) ODs in last-mile delivery, (ii) staff absenteeism and service disruptions, and (iii) mitigation strategies to increase resilience in distribution operations. The related literature in these topics is discussed in the following sections.

2.1. ODs in last-mile delivery

ODs are non-regular drivers used by delivery companies to cover parts of the end-to-end deliveries to reduce operational costs (Wang et al., 2023). ODs can be either private freelancers using their own cars for delivery services, or in-store customers with free capacity in their cars that they use to ensure deliveries once leaving the store.

Inspired by real-life practices adopted by Walmart and Amazon, Archetti et al. (2016) were the first to develop a vehicle routing problem-based model that involves ODs to cover a single delivery each. The authors develop two known metaheuristics (namely, Tabu search and variable neighborhood search) and embed them within a multi-start heuristic solution approach.

Since then, several new variants and extensions of the problem have been proposed. For example, Macrina et al. (2017) allow every OD to serve even multiple deliveries and impose time windows on each delivery. The authors develop a hybrid algorithm combining a genetic algorithm with a local search heuristic. Dahle et al. (2019) considered a pick-up and delivery version of the problem. Di Puglia Pugliese et al. (2022) focus on developing advanced and efficient approaches to solve the same problem such as a machine-learning technique embedded within a variable neighborhood search approach. Likewise, Ahamed et al. (2021) propose a centralized deep reinforcement learning-based method for assigning deliveries to the ODs.

Further extensions proposed by Macrina and Guerriero (2018), Macrina et al. (2020) consist, respectively, in allowing the transshipment of shipments (see also Voigt and Kuhn, 2022) and in considering the fleet fuel consumption minimization. This last aspect has also been addressed by Al Hla et al. (2019) who suggest controlling the drivers' behavior to reduce the fuel consumption and the resulting emissions. Lalang et al. (2019) consider a last-mile delivery system based on the employment of ODs only and extend the problem's settings to consider a multi-depot network. Dos Santos et al. (2022) propose an integrated 2-echelon system in which smart lockers are used, besides being used as collection points, are exploited as transshipment points where customers, willing to act as occasional drivers, collect parcels to be delivered on their way back home. A similar framework is proposed in Yu et al. (2021), with the difference that customers may opt for a specific delivery option (home or locker) or let the company decide for them. This problem setting has been extended in Yu et al. (2022), where compatibility among lockers and customers is considered.

Few scholars, such as Arslan et al. (2019), address the dynamic variant of the problem. The authors consider the pickup operations also, besides the delivery. The ODs have the option to express their preferences in terms of detour length threshold or maximum number of pickups/deliveries to be served (see also Behrend and Meisel, 2018). They claim that their exact method can achieve remarkable savings, which can reach 37% compared to the case in which only regular drivers are employed. Yu et al. (2023) deal with a similar problem by allowing simultaneous pickups and deliveries and considering the fact that the ODs can have different skills. The dynamic variant of the problem is also solved by Dayarian and Savelsbergh (2020), who develop a rolling horizon framework to incorporate real-time details of the in-store ODs, such as their arrival time and their available capacity. Archetti et al. (2021) study a dynamic variant of the problem, in which drivers availability is known, while requests appear dynamically.

The stochastic version of the problem attracted the attention of a few investigators who incorporate into the delivery planning (i) the uncertainty related to the willingness of the ODs to serve the assigned deliveries (Gdowska et al., 2018), (ii) the possible movement of the ODs (Cheng et al., 2017), and (iii) also their availability (see Section 2.2).

For the sake of completeness, it is worth noting that besides the above contributions, several other studies address the planning of deliveries from a managerial point of view. These particularly focus on analyzing the potential benefits of employing the ODs. A non-exhaustive list of studies in this direction include (Horner et al., 2021; Torres et al., 2022a) and Ausseil et al. (2022).

2.2. Service disruption and ODs absenteeism

The delivery industry, as any other transportation activity, is inevitably vulnerable and is subject to various disruptions from natural, political, or human factors (Albertzeth et al., 2020; Pahwa and Jaller, 2023). While several aspects of disruption have been extensively studied in the context of last-mile delivery (see the review by Rivera-Royero et al., 2022), the specific issue of absenteeism still did not attract enough attention. Indeed, Boysen et al. (2021) claim, "Short-term adaptations to account for absenteeism or demand peaks, e.g., during end-of-season sales, are important problems that lack scientific decision support."

Unlike the last-mile delivery, the curse of absenteeism and unavailability had a paramount importance in several staff planning problems such as that related to healthcare (Andrade-Michel et al., 2021; Hosseini et al., 2023) or to education (King et al., 2015; Bakrania et al., 2018). This is due not only to the importance these two sectors have in the daily life of any population but also to the remarkably high level of absenteeism by which they are often characterized. For example, Anderson (2022) reports an alarming rate of healthcare absenteeism in Canada, which amounts to 14.7 days per worker per year, which is 32% higher than the national rate across all industries.

To the best of our knowledge, Haughton, 2009 is the only study focusing on examining the problem of driver absenteeism in the transportation sector. The author, indeed, acknowledges an evident scarcity of dealing with driver absenteeism in the scientific literature, which is due to the fact that "Predominantly, the existing staffing and scheduling models implicitly assume perfect job attendance records by workers." However, the author stresses the fact that this is not true and provides several statistics from Canada, USA, New Zealand, and the UK to show how the absenteeism of drivers can substantially disrupt the transportation systems. He also argues that traditional approaches such as hiring replacement drivers or adjusting schedules can be costly and inefficient.

In the context of last-mile delivery with ODs, a few works incorporate the availability of crowd-shippers in the optimization models. The first study (Dahle et al., 2017) develops a vehicle routing-based mathematical model that involves dynamic ODs. They consider the time and distance constraints of the deliveries, as well as the skills and availability of the ODs. More specifically, the authors assume that the availability of ODs is a function of time and use historical data to estimate the probability with which an OD is available at any given time. The developed model results in a two-stage stochastic approach that first generates a set of feasible routes using a fixed set of regular drivers and then optimizes the incorporation of the ODs on these routes, based on their expected availability. The optimization uses a mixed-integer programming formulation that balances the trade-off between the travel time and cost of using ODs versus regular drivers.

The second study is Torres et al. (2022a), in which the authors take into account both the stochastic nature of customers' demand and the uncertain availability of ODs. The resulting two-stage stochastic model determines, in the first stage, the assignment of customers to the ODs and then adopts a recourse action consisting in defining the routing of the ODs, based on the observed demands and ODs' availability. The objective is to minimize the expected cost of ensuring the deliveries within their given time windows. The model also takes into account the capacity of each OD and the maximum distance they can travel. The authors propose a decomposition-based solution approach combined with a scenario generation technique and then solve the stochastic program for each scenario.

More recently, Silva et al., 2023 developed a data-driven method in which both the uncertainty related to the orders of the customers and the availability of ODs is considered. The resulting model is a two-stage stochastic program that identifies, in its first stage, the fleet routes with respect to the average values of the uncertain data. The recourse action consists of adjusting the vehicles routes while skipping any customer who has no delivery order or who has been outsourced to ODs. The authors also incorporate probabilistic constraints to limit the infeasibility due to the capacity restrictions.

A couple of additional works also proposed 2-stage stochastic models with uncertain availability of ODs while focusing on specific aspects of the problem such as the location of mobile depots (Mousavi et al., 2022) and the design of a personalized compensation scheme for each OD (Hou et al., 2022).

2.3. Mitigation strategies

Like any other supply chain, the last-mile delivery industry is not immune to the dramatic effects of unavoidable disruptions (Muñoz-Villamizar et al., 2021). In a 2011 survey, almost 600 practitioners claimed that delivery chain disruptions were more observable in their organizations than most other risks were (Katsaliaki et al., 2021). Moreover, this study also reported that even though 80% of the companies expressed awareness of the effect of risk disruptions, only 40% of them have adopted mitigation strategies to hedge against disruption.

The topic of mitigation management to ensure supply chain resilience became very popular in recent years (Gurtu and Johny, 2021; Hägele et al., 2023). The efforts have focused on defining both proactive or reactive strategies based mainly on ensuring redundancy to reduce the effect of disruptions and to ensure system reliability (see the reviews by Ivanov et al., 2016; Katsaliaki et al., 2021). Yet, a couple of studies that attempt to explicitly mitigate the effect of the absenteeism of drivers – but not of ODs – do exist. The first attempt in this direction is presented in Haughton, 2009, who proposed a traditional strategy based on providing a reserve team of part-time drivers. These are kept available to cover the unexpected absence of drivers. Similarly, Diab et al. (2014) quantify the number of reserve back-up drivers, known as extraboards, needed to substitute regular drivers in case of absenteeism in the bus transit network of Ottawa, Canada. Finally, Wang and Ozbay (2023) develop a chance-constrained optimization model as a mitigation strategy to take into account the absenteeism of drivers during hurricane evacuations.

2.4. Main contributions

This study proposes several novel features that can be summarized as follows:

- As claimed by Pahwa and Jaller (2023), there is an evident scarcity of studies considering ODs absenteeism in the context of last-mile deliveries. The reason is that researchers and practitioners consider the employment of ODs as a mitigation strategy itself to face demand peaks, shortage of regular drivers, etc. Thus, there is a clear gap in defining mitigation strategies to deal with the absenteeism of ODs and our study represents an attempt to fill in this gap by anticipating and swiftly responding to such random disruptions.
- We design a CA that allows the ODs to decide not only their compensation but also the set of deliveries that strongly synergize with their own trajectories. All previous studies adopt simplistic schemes based on a fixed fee per delivery and/or per deviation miles. In that case, the ODs do not have any capability to negotiate the delivery price or to express their willingness to serve through choosing their tariffs (Cheng et al., 2023).
- We introduce for the first time the concept of reserve drivers who are selected by the same mechanism of CAs and receive a fee to ensure their presence for delivery whenever needed to hedge against the absenteeism of regular ODs.

Some of the above features have been addressed separately in the different references cited within this section, but none of them has integrated all the features together within the same decisional framework. This study has been inspired from one side by the works that implement CAs in last-mile deliveries (such as Triki, 2021 and Mancini and Gansterer, 2024), and from the other side, by those incorporating the randomness in customers orders and drivers availability (Dahle et al., 2017, Torres et al., 2022b, and Silva et al., 2023) and can be considered as a more comprehensive investigation in both directions. However, none of the approaches reviewed above can be straightforwardly extended in order to take into account the possible absenteeism of the ODs and, thus, to suggest mitigation strategies to avoid the delivery disruptions.

Table 1
Notation.

Sets	I	Set of customers to be served
	K	Set of bundles of customers
	J	Set of ODs
	K^{ij}	Set of bundles submitted by OD j containing customer i
	K^j	Set of bundles submitted by OD j
Parameter	b_{kj}	Price submitted by OD j to serve bundle k
Variables	Y_{kj}	Binary variable that assumes value 1 if bundle bid k from OD j is accepted, 0 otherwise

3. The deterministic problem

A set I of customers is considered to be served, starting from a single depot. The company offers, through an online platform, a set K of bundles of customers, in which each bundle contains a subset of the customers that can be handled by a single driver. Such bundles are constructed by means of the *corridors approach* presented in Mancini and Gansterer (2022), which has been shown to be able to provide more attractive and profitable bundles with respect to classical clustering approaches. The method identifies circular sectors, starting at a depot and assigns customers within a sector to a bundle. Being α the angle corresponding to this sector, the method splits α in s identical angles, generating α_s identical sectors. If the demand within a potential bundle exceeds the capacity of the associated vehicle, the bundle is split into feasible smaller bundles, by means of a clustering algorithm. The bundles are non-overlapping and cover all available customers. The algorithm is run for different values of s obtaining different sets of bundles, which are then considered within the auction.

A set J of ODs joins the auction on the platform and submits their bids for the bundles they are interested in. b_{kj} indicates the price offered by driver j to serve bundle k. The set of bundles, containing customers i, for which driver j has bid is defined as K^{ij}. The number of bundles in K^{ij} can potentially be quite high because it is in the interest of the ODs to submit as many bids as possible, as long as they are confident they can serve. However, as observed by Sheffi (2004), generating competitive bids may be a complex task and, consequently, many bidders reduce their operation in the auction by submitting a limited number of trivial bundles. Similarly, the set of bundles for which driver j has placed bids is defined as K^j (i.e., $K^j = \bigcup_{i\in I} K^{ij}$). The auction system selects the subset of bids to accept, such that auction clearing costs are minimized, and every customer is assigned to exactly one driver. To avoid the problem of overexposure, i.e. the case in which the same driver wins several bundles but is unable to fulfill all of them, additional constraints are imposed stating that at most, one bid per driver can be accepted. This is in line with the literature on auction-based mechanisms in logistics (see Gansterer and Hartl, 2018, Gansterer et al., 2020). However, in case of infeasibility, the decision maker may decide to relax such a constraint by increasing the number of bids to be accepted per driver to 2 or more.

The resulting decision problem is similar to the winner determination problem described in Gansterer et al. (2018). The decision variables involved, named Y_{kj}, represent acceptance/rejection decisions of bids. Variable Y_{kj} is binary and assumes a value of 1 if the bid from driver j, to serve bundle k is accepted, and 0 otherwise. Note that Y_{kj} is defined only if driver j has submitted a bid for bundle k. This way, the number of variables involved is limited, avoiding the generation of unnecessary variables. Our notation is summarized in Table 1.

The mathematical deterministic model can be expressed as follows:

$$\min \sum_{j\in J} \sum_{k\in K^j} b_{kj} Y_{kj} \tag{1}$$

$$\sum_{j\in J} \sum_{k\in K^{ij}} Y_{kj} = 1 \quad \forall i \in I \tag{2}$$

$$\sum_{k\in K^j} Y_{kj} \le 1 \quad \forall j \in J \tag{3}$$

$$Y_{kj} \in \{0,1\} \quad \forall j \in J, k \in Kj \tag{4}$$

The objective function minimizes the auction clearing cost, i.e. the sum of the values of the accepted bids. Constraints (2) ensure that each customer belongs to exactly one accepted bid, and Constraints (3) imply that at most, one bid per driver can be accepted. Finally, Constraints (4) define the variables' domain.

4. The stochastic problem

The overexposure phenomenon within the same auction system can be avoided by adding Constraints (3). However, drivers cannot be prevented from simultaneously participating in several auctions organized by different companies. In this case, the overexposure may still occur, and a driver who wins several auctions may be forced to drop some of the assigned tasks. This situation could generate a high absenteeism rate, which would strongly affect the delivery planning of the company and risks some of the customers remaining unserved. ODs, in contrast to regular drivers, do not have fixed contracts with companies. Therefore, they do not incur a penalty if they withdraw on short notice or do not show up. However, their absence negatively influences their trustability level, τ_j, which is tracked by the company. This index could be used by the company to determine the probability of a specific driver fulfilling the planned service. This information is important to build robust delivery plans. To limit the effect of the absenteeism of ODs, a mitigation-strategy-based reserving drivers is proposed. The company considers the possibility of reserving some drivers by compensating them to guarantee their availability during the delivery period. Reserved (or backup) drivers are not assigned a specific bundle in the planning phase but are requested to fulfill a bundle of customers for whom delivery failed due to the absenteeism of the designated driver. If the service of reserved drivers is not needed, compensation for ensuring availability must be paid anyway, since they accepted to not participate in other auctions with other companies during the same time period. Each driver can choose to bid only on bundles or also to become a reserved driver.

The overall auction system is then composed of two phases: (i) in the first phase, drivers submit bids for the bundles they are willing to serve (*regular auction*) and eventually for the reserved driver role (*reserve auction*); (ii) the company decides which bids to accept and whom to reserve (*recourse auction*). After driver absences are revealed, the company decides which bundle to assign to which reserved driver. All non-served customers have to be rejected. Note that reserved drivers can be assigned only to bundles for which they have submitted a bid in the first-phase of the auction. The bid offered by driver j to act as a reserved driver is denoted as r_j. If drivers are reserved but unused, their compensation will be equal to r_j. Conversely, if designated to serve a bundle k, their compensation will equal $r_j + b_{jk}$. The overexposure avoidance rule holds also in the recourse auction, i.e. a reserved driver can also be assigned at most one bundle. It is not necessary that each reserved driver fulfills a bundle selected in the first phase. In fact, a different set of bundles can be selected if more beneficial for the company. An illustrative example is reported in Fig. 2, in which the goal is to minimize the overall costs composed of costs related to

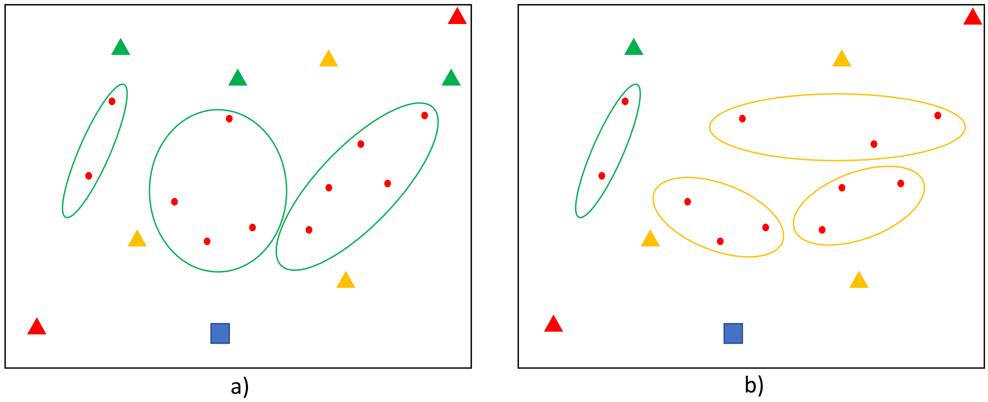

Fig. 2. An illustrative example (depot: blue; customers: red; selected ODs and their bundles: green; reserved drivers: orange; drivers not selected: red). Subfigure (a) shows the result of the first phase. Subfigure (b) depicts the second phase, in which two selected drivers are absent and replaced by three reserve ODs serving different bundles.

accepted bids (in the first and in the second phase), reservation costs, and rejection costs.

The above described decision problem can be formulated as a two-stage stochastic model. In the first stage, bundles are assigned to ODs and ODs are selected to be reserved. In the second one, when the presence of ODs are revealed, unfulfilled customers are assigned to reserved drivers, and a penalty is paid for customers remaining unserved. The goal is to minimize the overall costs related to both the first and the second stage decisions. A set of available scenarios, S, having cardinality $|S|$, is considered. In each scenario, there is information about drivers' presence. This information is tracked by a parameter, δ_{js}, which is equal to 1 if driver j is present in scenario s and 0 otherwise. Scenarios are based on the estimated trustability levels of drivers, τ_j, which is a parameter value between 0 and 1. This parameter is computed based on historical data. For drivers accessing the system for the first time, for whom historical data are unavailable, τ_j is initialized at a default value equal to the average trustability level of the drivers in the system. Since the probability of being absent is explicitly considered in the scenario generation, all scenarios have the same probability of being realized. A penalty p_i must be paid for each scenario in which customer i is rejected, i.e. neither served by a driver selected at the first stage nor by a reserved driver in the second stage. The penalty may vary among customers due to their premium status. In fact, premium customers are guaranteed that their order will be delivered in time (i.e. one day after their ordering), while orders from standard customers must be fulfilled within five working days. Customer-dependent rejection costs are also considered. However, if the obligation-to-serve scheme should be ensured, then it is enough to impose a high penalty cost p_i for all customers and the model will select, in this case, even more expensive bids in order to guarantee the full service coverage.

Note that this model differs and introduces some novelty with respect to classical two-stage stochastic models. In fact, in classical models, decisions are made at the first stage without perfect information about some parameters, and, once perfect information is revealed, recourse actions to improve the solution are realized. In this problem, however, to gain the right to apply a recourse action at the second stage, additional costs have to be paid at the first stage, and the set of available recourse actions depends on the reservation choices in the first stage. In fact, if no drivers are reserved at the first stage, then no recourse actions will be available at the second. Reserved drivers can be used to fulfill only bundles which they are willing to serve, i.e. for which they have submitted a bid in the first stage. Consequently, they are able to cover only customers belonging to at most one bundle they have bid for. If customers are unserviceable by any of the reserved drivers, it means that they cannot be covered with a recourse action.

Since there is a substantial difference between classical two-stage models and this model, we define it as a *Two-Stage Stochastic Model with Recourse Activation*. Not only is this new modeling framework valid for this specific problem but it can also be seen as a general approach suitable to model a broader class of problems in which, to be able to access the recourse phase, one or more activation variables at the first stage need to be selected after paying a cost.

4.1. The two-stage stochastic model with recourse activation

The first-stage (deterministic) decision variables related to the formulation are as follows:

- Y_{kj}: binary variable taking value 1 if a bid for bundle k submitted by OD j is successful in the regular auction, and 0 otherwise
- R_j: binary variable taking value 1 if OD j is successful for a mitigation delivery service in the reserve auction, and 0 otherwise

whereas the second-stage (recourse scenario-based) decision variables are:

- W_{kjs}: binary variable taking value 1 if reserved OD j will be asked to serve bundle k whenever a recourse action is needed for mitigation under scenario s, and 0 otherwise
- Z_{is}: binary variable taking value 1 if customer i remains unserved in scenario s, and 0 otherwise

The two-stage stochastic IP program can be formulated as follows.

$$\min \sum_{j\in J}\sum_{k\in K^j} b_{kj}Y_{kj} + \sum_{j\in J} r_j R_j + \sum_{s\in S}\frac{1}{|S|}\Big(\sum_{j\in J}\sum_{k\in K^j} b_{kj}W_{kjs} + \sum_{i\in I} p_i Z_{is}\Big) \quad (5)$$

$$\sum_{k\in K^j} W_{kjs} \le R_j \quad \forall j\in J,\ \ \forall s\in S \quad (6)$$

$$\sum_{j\in J}\sum_{k\in K^{ij}} \delta_{js}Y_{kj} + \sum_{j\in J}\sum_{k\in K^{ij}} W_{kjs} + Z_{is} = 1 \quad \forall i\in I,\ \ \forall s\in S \quad (7)$$

$$\sum_{k\in K^j} Y_{kj} + R_j \le 1 \quad \forall j\in J \quad (8)$$

$$Y_{kj}\in\{0,1\} \quad \forall j\in J, k\in Kj \quad (9)$$

$$R_j\in\{0,1\} \quad \forall j\in J \quad (10)$$

$$W_{kjs}\in\{0,1\} \quad \forall j\in J, k\in K_j, s\in S \quad (11)$$

$$Z_{is}\in\{0,1\} \quad \forall i\in I, s\in S \quad (12)$$

The objective function, reported in (5) minimizes the overall cost for the company, computed as the sum of (i) the cost of the accepted bids for ODs in the regular auction, (ii) the cost of reserving ODs, (iii) the cost of accepted bids in the recourse auction, and (iv) the penalty cost related to the unserved customers. Constraints (6) imply that the

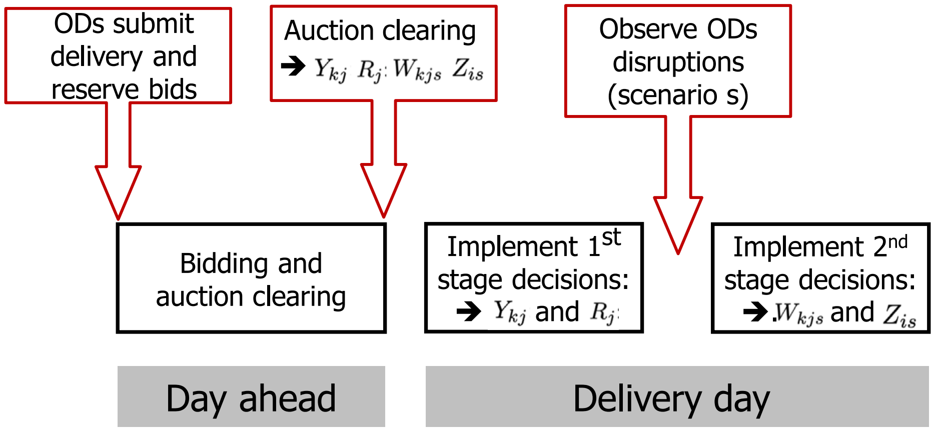

Fig. 3. Graphical representation of the overall auction process.

bid of an OD can be accepted in the recourse phase only if the OD has been reserved in the first phase. Constraints (7) impose that each customer who has not been successfully served either by a regular OD selected in the first stage or by a reserved OD in the second stage remains unserved. Constraints (8) ensure that if drivers are selected for reservation, no bundles are assigned to them in the first stage. Finally, Constraints (9)–(12) state that the decision variables should all be binary. An explanatory draw illustrating the mode of operation of the overall auction process is depicted in Fig. 3.

5. Solution approach

A broad class of algorithms use *learning* to improve their performances. This class can be split into three categories. The first includes algorithms that use different search operators and modify the probability of selecting a specific one, based on the performances the operator obtains during the search process. Thus, the algorithm learns from past experience in the search process. The main and most known member of this category is the Adaptive Large Neighborhood Search (ALNS) proposed by Ropke and Pisinger (2006). Many extensions of the ALNS have been proposed in the literature. For a complete survey on ALNS algorithms and applications, interested readers are referred to Windras Mara et al. (2022). A concept similar to that used in ALNS is exploited in Hyper Heuristics (Burke et al., 2013), in which data collected during the search process is used to guide the algorithm in selecting which heuristic to use, instead of which search operator is within the same heuristic. For a complete survey of Hyper Heuristics, readers are referred to Drake et al. (2020).

The second category includes Machine-Learning-based heuristics. Algorithms belonging to this category may differ greatly from each other, but all follow a common philosophy. The main idea is to analyze optimal solutions of several problem instances to derive knowledge on the characteristics commonly seen in optimal solutions. This information then guides the search process. A detailed classification of Machine-Learning-based heuristics is reported in Karimi-Mamaghan et al. (2022).

The third category, to which our proposed algorithm belong to, is composed by self-learning heuristics. Differently from Machine-Learning-based heuristics, which need a training set of optimal solutions to learn from, self-learning ones learn by themselves during the search process on a single instance. Differently from ALNS and Hyper Heuristics, that learn which operator/heuristic to use, self-learning heuristics iteratively solve a simplified version of the model involving only a subset of the variables, study how the obtained values of these variables would impact the other variables, and modify the simplified version of the problem, taking into account the information derived. Due to their nature, these heuristics are suitable for addressing two-stage stochastic problems. In fact, the simplified version of the problem could include only first-stage decisions, whose impact on second-stage decisions can be evaluated and the derived information used to guide the first-stage problem by artificially modifying the contribution of the variables to the objective function, or to a priori force some first-stage variables to assume a specific value. An example of self-learning heuristics is the Progressive Hedging method introduced by Rockafellar and Wets (2017). This method is based on a decomposition approach according to which each scenario is solved separately, giving lieu to different first-stage choices. First-level variable coefficients in the objective function are then artificially modified to achieve a consensus among scenarios, guiding all the scenarios to converge towards a common set of first-stage decisions. Another example is the consensus-based matheuristic, first proposed by Mancini et al. (2022) and then generalized by Mancini et al. (2023). In this method, each scenario is solved separately, and first-stage variables are ranked according to the number of scenarios in which they have been selected. Then, only the most selected variables are open at the first stage and the problem is iteratively solved, changing the variables selected at the first stage, trying to achieve consensus among scenarios by inserting variables preferred by the scenarios which are penalized most by the current first-stage solution. The proposed self-learning matheuristic, denoted *SLM*, does not require decomposition of the problem and solves each scenario separately. In fact, the idea on which the method is based is to iteratively solve the first stage problem, analyze the impact of first-stage decisions on the second-stage costs, and modify the first-stage variable costs accordingly to take into account this impact.

5.1. A self-learning matheuristic

A two-phases self-learning matheuristic is proposed to address the two-stage stochastic problem with recourse activation. The phases are sequentially executed for a fixed number of iterations, or until a stopping criterion is reached. In the first phase, the deterministic problem (1)–(3) is solved. Hence, all ODs are assumed to be present with a 100% probability. The costs associated with bids (b_{jk}) are substituted by the perceived costs, $\tilde{b}_{jk}$, which are given by the actual bidding costs, b_{jk}, plus an estimation of the recourse/mitigation cost incurred if the OD would be absent (γ_{jk}). For simplicity, these costs are referred to as *proxy costs*, and are all fixed to zero at the first iteration, since there is no a-priori information about them, but they are updated during the search process, based on a self-learning procedure.

In the second phase, the bundles selected in the first stage are fixed (i.e. the values of the Y_{kj} variables), according to the solution obtained in the first phase, and the stochastic model are allowed to optimize reservation-related and recourse-related variables. The set of the selected Y_{kj} bids at a given iteration t is referred to as $\bar{Y}_t$. The mitigation cost found in the second phase, defined as the sum of

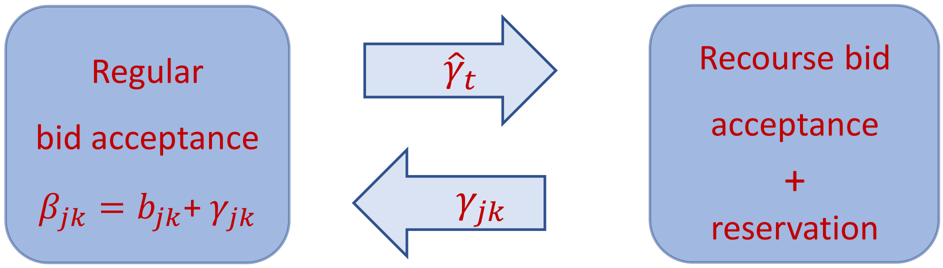

Fig. 4. Procedure of the self-learning matheuristic.

reservation costs (RC^t) plus the second-stage bidding costs (SSC^t), is then attributed to the failed bids according to their contribution to these costs. For each accepted bid in the first stage, the number of scenarios in which the corresponding OD, j, results to be absent, m_j is calculated. Being $M = \sum_{j \in J} \sum_{k \in K} m_j Y_{jk}$ the total number of failures, the reservation cost associated with the selected bid Y_{kj} is defined as follows.

$$rc_{jk} = \frac{m_j}{M} RC^t \tag{13}$$

β_j^t is defined as the set of customers assigned in the regular auction to driver j in iteration t and $\tilde{K}_i$ as the set of bundles containing customer i and n_k as the number of customers belonging to bundle k.

The second-stage cost associated with a bid of driver j on bundle k, ssc_{jk}, is defined as

$$ssc_{jk} = \sum_{s \in S} \sum_{i \in \beta_j^t} \sum_{k' \in \tilde{K}_i} \sum_{j' \in J} \frac{1}{n_{k'}} w_{j'k's}. \tag{14}$$

The proxy cost γ_{jk} is then defined as:

$$\gamma_{jk} = rc_{jk} + ssc_{jk}. \tag{15}$$

For all the bids which do not belong to $\bar{Y}_t$, the proxy cost γ_{jk} is not modified and takes the value of the previous iteration.

The method stops after a maximum number of iterations T or when a steady state is reached, meaning that modifying the proxy costs does not imply any change in the selected first-stage bids (i.e. $\bar{Y}_t = \bar{Y}_{t-1}$). Note that a steady state might differ from a local minimum as there is no guarantee that the current solution at a given iteration t is better than the solution at iteration $t-1$, as happens in the case of the Progressive Hedging method (Rockafellar and Wets, 2017).

The core idea of the approach is that the algorithm is self-adapting its decisions, based on the information achieved by computing the effect of a set of first-stage decisions on the recursion and mitigation costs. The novelty of the approach is twofold. First, the method is able to identify how much a single first stage decision is responsible for the second stage and for the mitigation cost. If a first stage solution leads to very high recourse and mitigation costs, this does not necessarily imply that all the variables selected on the first stage should not be selected, since the high second stage costs might be due to only a subset of them. The identification of the single "bad" decisions made at the first stage, allows us to exclude these from future solutions. The second main element of novelty relates to the ability to identify the impact of first-stage decisions not only on the recourse costs but also on the mitigation cost needed to achieve them. The method terminates after a maximum number of iterations (*Itermax*) or after a maximum number of not improving iterations (*NImax*). In Fig. 4, the procedure of SLM is illustrated.

A pseudocode of the algorithm is reported in Algorithm 1.

Algorithm 1 Self Learning Matheuristic (SLM)

$\hat{Y}_t \leftarrow \{\emptyset\}$
$\hat{\gamma} \leftarrow 0$
$t \leftarrow 0$
while $t < T$ **do**
  $\hat{\beta} \leftarrow \hat{b} + \hat{\gamma}$
  $\hat{Y}_t \leftarrow regular\_bids\_auction(\hat{\beta})$
  if $\hat{Y}_t \neq \hat{Y_{t-1}}$ **then**
    $\hat{\gamma} \leftarrow recourse\_reservation_bids\_auction(\hat{\beta}, \hat{Y}_t)$
    $t \leftarrow t + 1$
  else
    $t \leftarrow T$
  end if
end while

The method can be generalized and easily adapted to a broad class of two-stage stochastic problems with recourse activation, in which a direct correlation between each first-stage decision and its impact on the second stage can be established. Moreover, the algorithm can be easily embedded in more complex metaheuristics, such as Iterated Local Search (ILS), in which diversification operators are exploited to escape from steady states.

5.2. Iterated local search based on the self-learning matheuristic

To avoid a premature convergence towards a local minimum, our proposal is to apply a diversification procedure, which is designed similar to an ILS (Lourenço et al., 2018). Each time, SLM reaches a steady state (i.e. $\bar{Y}_t = \bar{Y}_{t-1}$), a random perturbation value μ_{jk} is generated, taking values $[0; \mu^{MAX}]$ for all the bids belonging to the current solution while fixing $\mu_{jk} = 0$ for all those not selected in the current solution. The proxy costs are then updated as $\gamma_{jk} = \gamma_{jk} + \mu_{jk}$. This way, an attempt is made to artificially make not-selected bids more attractive to push the model to select at least one of them and get out of the steady state. If, despite the perturbations, no changes occur in the set of selected bids, the perturbation is iteratively applied until a change is reached. The procedure stops after a maximum number of perturbations (MAX_{PERT}).

A pseudocode of the algorithm is reported in Algorithm 2.

Algorithm 2 Iterated Local Search (ILS)

$t \leftarrow 0$
while $t < T$ **do**
  $\hat{\gamma} \leftarrow MLS(\hat{\gamma})$
  $\hat{\gamma} \leftarrow \hat{\gamma} + \hat{\mu}$
  $t \leftarrow t + 1$
end while

These two newly proposed matheuristics are compared with a more classical Variable Neighborhood Search (VNS), properly adapted to handle our problem, which will be described in the next section.

6. Computational study

For the computational study, 4 sets of 10 instances each are used. The first two have 20 customers and 200 and 400 scenarios, respectively, while the others have 40 customers and 200 and 400 scenarios, respectively. All of them include 20 ODs. The final destinations and customers' locations of both ODs are generated randomly. Instances are publicly available at Mancini et al. (2024).

Bundles of customers, which are given as input, are generated following the corridors-based approach introduced by Mancini and Gansterer (2022), which has been shown to be effective in generating attractive bundles. Although our system allows drivers to autonomously enter their bids based on their own preferences, in order to simulate realistic drivers' bidding behavior, we use the automatic bidding system provided by Mancini and Gansterer (2022). It considers the two following parameters: (i) *flexibility*, which represents the maximum acceptable detour for a driver ω from the path from the depot to her final destination, and (ii) *willingness to work* (ϕ_ω), where $\phi_\omega = 1$ describes a neutral behavior (truthful bidding) where the ODs' bids reflect exactly the actual detour implied. In case of a lower willingness (i.e., $\phi_\omega > 1$), bid prices are increased since the ODs agree to perform a delivery only for a very high compensation. Values smaller than 1 ($\phi_\omega < 1$) indicate that the driver reduces the bid price to have a greater chance of winning the order. The value of a bid kj is calculated as the detour length, δ_{kj}, needed by driver j to serve bundle k, multiplied by a unitary distance cost c^u, plus a fixed cost c^f for each customer belonging to the bundle. This value is further multiplied by the willingness-to-work parameter (ϕ_j) associated with the OD who submitted the bid:

$$b_{kj} = (c^u \delta_{kj} + c^f |\tau_k|)\phi_j. \tag{16}$$

The flexibility level, instead, only indirectly impacts the bidding process. In fact, an OD places a bid for a bundle only if the related detour δ_{kj}, is lower than the maximum value allowed (δ_j^{MAX}) which corresponds to the OD's flexibility level.

For what concerns reservation bidding, we generate bids as an integer number randomly selected in the interval [1,10]. We considered these values to be uncorrelated to the willingness to work. In fact, for example, a driver could have a very low willingness to work (i.e., accepts only very high compensation) but at the same time could be available to guarantee her presence for a relatively small compensation. It could also happen that a driver with a very high willingness to work makes a very competitive bids for the regular auction but asks a high compensation for acting as a reserved driver, since she does not want to risk to ensure her availability bearing the risk to not be assigned to any order and receive just a very low fee for the reservation.

For each set, a basic version is used with a medium trustability level, which is uniformly distributed in [0.6,1] and medium rejection costs, which are uniformly distributed in the set {5,10,15,20}. There are four additional versions for each set. In the first two, rejection costs are kept constant on the respective level of the basic set, whereas low [0.6, 0.8] and high [0.8, 1] trustability levels are considered, respectively. Similarly, in the third and fourth version, a medium trustability level (taken from the basic set) is assumed, while rejection costs are low {5, 10} and high {15, 20}, respectively. Hence, in total, this study is based on 200 instances.

In SLM, the maximum number of iterations $Itermax$ is fixed to 100, while $NImax$ is fixed to 10. In ILS, the maximum size of perturbation μ^{MAX} is fixed to 1, while the maximum number of perturbations, i.e. the stopping criterion, is fixed to 10.

Four analyses have been performed. The first one aims at comparing the performance of both SLM and ILS (which is based on SLM as explained in Section 5.2), with the exact model solved by a commercial solver. The second analysis is devoted to analyzing the advantage of applying the proposed mitigation strategy. In the third one, the optimal solution of the stochastic problem (SP) is compared with those obtained by applying three different deterministic approaches. This analysis is used to determine a well-known and broadly used stochastic indicator: the Value of the Stochastic Solution (VSS), which measures the benefit achievable by solving the stochastic problem instead of its deterministic counterpart (Birge and Louveaux, 2011).

Both the second and the third analysis are performed for different average trustability levels of ODs as well as different average rejection costs. This procedure analyzes the impact of these parameters. The last analysis aims at finding an identikit of the perfect reserved driver by analyzing the optimal solutions.

All experiments use a machine equipped with a 11th Gen Intel Core i7-1185G7 with 32 GB of RAM. The mathematical model is run under Xpress 8.13 with standard settings and a time limit of 3600 s. The optimality gap tolerance was set to 10^{-5}, which is the common default value.

6.1. Performance comparison

In this section, we compare the optimal solution obtained by solving the mathematical model (MODEL) by means of a commercial solver, with the best solution obtained by SLM, by ILS in which SLM is used as local search, and by a more classical VNS.

Here we briefly describe the structure of the VNS. It is based on the same framework of SLM in which, at each iteration, a deterministic problem is solved considering that the perceived cost of each bid is composed by actual bidding cost plus an estimation of the recourse/mitigation cost (proxy cost). But differently from SLM, where proxy costs are learned during the search process, here they are considered as further variables in the solution space. At the first iteration, as in SLM, proxy costs are all fixed equal to 0 and the corresponding deterministic problem is solved. Then a random perturbation in the range [-0.5; 0.5]*ρ is applied to all the bids, and then the deterministic model is solved again. The parameter ρ is initially set equal to 1 and it is multiplied by 2 every time we fail obtaining an improvement. As soon as an improvement is found, the value of ρ is restored to 1. The procedure terminates once the maximum value of $\rho = \rho^{MAX} = 16$ is reached, or after a maximum number of iterations, T.

Results are reported in Table 2, which is organized as follows. For MODEL, for each basic set of instances, the value of the optimal solution and the total time elapsed (in seconds) are reported, whereas for SLM, ILS, and VNS the percentage gap to the optimal solution and the time elapsed (in seconds) are reported.

Both approaches (SLM and ILS) show a good performance with more than 90% reduction of computational times. As expected, ILS guarantees lower gaps, but it is more than twice slower than SLM. For both methods, the gap to MODEL decreases when increasing the instances size. Both of them also scale well and are suitable for addressing larger instances that MODEL cannot handle. VNS is always outperformed by SLM and ILS, both in terms of solution quality and computational times, showing the benefit achievable using the self learning mechanism. Such benefit becomes much larger for instances with a larger number of customers and/or of scenarios, showing also a better scalability of SLM and ILS with respect to VNS.

6.2. Analysis of the proposed mitigation strategy

This analysis aims to quantify the cost reduction achievable by applying the proposed mitigation strategy based on the reservation of drivers. To that end, the solution of SP is compared with three deterministic versions of the problem, in which we do not allow reservation, and therefore, we clear the regular auction and then, once possible driver absences are revealed, we simply reject those customers belonging to bundles that remained unserved.

The first one, named DP, simply consists of solving the deterministic problem described in Section 3. Hence, the acceptance of a subset of bids that minimizes the regular auction clearing costs is optimized,

Table 2
Comparison of MODEL, SLM, ILS, and VNS. Instances are denoted as XXcYYYs, with XX being the number of customers and YYY the number of scenarios. OF indicates the objective function value.

	MODEL		SLM		ILS		VNS	
	OF	TIME	%GAP	TIME	%GAP	TIME	%GAP	TIME
20c200s	60.41	30.21	5.65%	0.55	4.26%	3.72	5.99%	26.77
20c400s	61.05	84.9	5.69%	1.36	5.08%	13.35	6.68%	67.17
40c200s	122.56	64.39	4.05%	2.86	2.79%	9.69	4.18%	27.69
40c400s	123.03	276.53	3.53%	7.44	2.84%	20.12	7.23%	77.12

Table 3
Average percentage gap achieved by solving DP, TDP, and ATDP compared against SP for different average drivers trustability (TRUST) levels.

TRUST	SP	DP	TDP	ATDP
LOW	71.89	39.75%	35.50%	35.62%
MEDIUM	60.41	32.04%	22.18%	22.74%
HIGH	54.57	19.52%	21.88%	21.88%

Table 4
Average percentage gap achieved by solving DP, TDP, and ATDP compared against SP for different average drivers rejection costs (RC).

RC	SP	DP	TDP	ATDP
LOW	58.54	17.34%	12.30%	13.44%
MEDIUM	60.41	32.04%	22.18%	22.74%
HIGH	63.22	42.47%	30.68%	30.81%

neglecting the information about the trustability of ODs. The set of selected bids is then given as input to the availability scenario of each OD, which is then analyzed separately. Customers who were included in a bid from a driver who is absent, are then automatically considered unserved, since no recourse actions are allowed.

The second deterministic problem, TDP, minimizes a modified function of the auction clearing costs, in which the value of each bid is adapted by the trustability of the bidding driver. This can be obtained substituting the objective function reported in (1) by

$$\min \sum_{j\in J} \sum_{k\in K^j} b_{kj}(1-\tau_j)Y_{kj}. \tag{17}$$

By this, the trustability level of its bidder positively influences the attractiveness of a bid. As in the previous case, the selected bids are used to evaluate the actual costs of each scenario, assuming that all the customers, included in a bid assigned to an absent driver, remain unserved.

Finally, for the third deterministic problem, named ATDP, an additional constraint is added, which imposes that the average trustability level of selected drivers cannot be lower than the average level of all the ODs, named $\bar{\tau}$. This allows accepting potentially risky drivers (due to their low trustability), as long as the trustability of all selected drivers reaches a threshold, which is given by the average trustability of all participating drivers. This restriction can be imposed by adding the following constraint.

$$\sum_{j\in J} \sum_{k\in K^j} \tau_j Y_{kj} \geq \bar{\tau} \sum_{j\in J} \sum_{k\in K^j} Y_{kj} \tag{18}$$

All the models (SP,DP, TDP, and ATDP) are solved to optimality by a commercial solver. Therefore, in this analysis, we always compare optimal solutions.

Tables 3 and 4 report the average value of SP and the average gap achieved by each deterministic approximation, for the variants of the trustability level of the average drivers and of the average rejection costs, respectively.

As can be observed in Table 3, the advantage of using the proposed mitigation strategy for drivers absenteeism is very strong. This applies even if the average trustability level of drivers is high. This is further supported by the results reported in Table 4, which show that the approach clearly dominates even if customers' rejection costs are low. Summarizing, it can be stated that the mitigation strategy outperforms not only the naive deterministic strategy (which neglects the trustability of drivers), but also deterministic strategies in which the trustability level is taken into account. TDP and ATDP have been further observed to show very similar performance. Both seem to yield more robust dispatching plans than DP does in case of the average trustability level of low and medium ODs. However, if the average trustability is high, the DP strategy, which neglects trustability and focuses only on the minimization of the clearing costs of bids, performs slightly better than the other two policies that do take trustability into account. Hence, in these cases the absenteeism phenomenon has less impact than that the bidding values have.

Another important observation relates to the percentage of rejected customers, which can be a key indicator of the quality of service. Moreover, a high rejection rate could potentially lead to a loss of customers, who, unsatisfied by the service, might choose to purchase from other companies. After SP is solved for the basic scenario, a rejection rate of only 1.69% is achieved. Note that by applying DP, TDP, and ATDP, 19.10%, 11.24%, and ATDP 11.24%, respectively, are obtained.

6.3. Comparison with deterministic problems with recourse policy

An indicator commonly used to determine the importance of considering the uncertainty in the problem by solving its stochastic variant, is the Value of the Stochastic Solution (VSS). The higher the value of VSS, the higher the need to resort to stochastic optimization. Following the definition given in Maggioni and Wallace (2012), the VSS can be seen as a measure of the advantage achievable by solving the stochastic problem instead its deterministic counterpart, where all random input variables are replaced by their means. Let us define as the EEV the expected value of the solution of the deterministic problem applied to all the second-stage scenarios, and as RP the expected value of the solution of the two stage recourse stochastic problem. Thus, according to Maggioni and Wallace (2012), the VSS can be defined as EEV-RP. However, if replacing random input variables by their means makes perfectly sense when those variables are continuous, this deterministic approximation is meaningless when dealing with binary variables, as in our problem. In fact, a driver can be either present or absent (0 or 1), and no presence implication can be associated with any fractional value in this context. In such cases, the EEV can be obtained considering a *reference problem* (see Mancini et al., 2023) consisting into an associated deterministic problem which is more meaningful respect to considering the expected values of the random variables. The above defined DP, TDP, and ATDP all belong to the category of reference problems in our context. In this analysis we use the above mentioned deterministic approaches (DP, TDP, and ATDP) to determine the outcome of the regular auction, then we fix the value of these decision variables, and let the SP model optimize the reservation and the recourse auctions. These approaches are based on the integration of deterministic and stochastic solution methods and have the advantage to handle uncertainty better than pure deterministic approaches, but, at the same time, require a smaller computational effort with respect to a pure stochastic approach and, therefore, are more suitable to address large-sized instances. We then calculate the VSS with respect to all the three

Table 5
PVSS analysis. Average percentage gaps achieved by solving DP, TDP, and ATDP with recourse actions compared against SP for different drivers' trustability levels.

TRUST	SP	DP	TDP	ATDP
LOW	71.89	9.43%	7.03%	8.79%
MEDIUM	60.41	10.05%	10.12%	10.29%
HIGH	54.57	6.72%	19.54%	6.72%

Table 6
PVSS analysis. Average percentage gaps achieved by solving DP, TDP, and ATDP with recourse actions compared against SP for different rejection costs (RC).

RC	SP	DP	TDP	ATDP
LOW	58.54	9.34%	10.50%	10.08%
MEDIUM	60.41	10.16%	10.12%	10.29%
HIGH	63.22	10.33%	9.14%	9.48%

Table 7
Comparison of SP against the deterministic problems on out-of-sample scenarios.

INSTANCE	SP	DP	TDP	ATDP
1	51.65	50.31	62.22	62.22
2	43.77	44.80	46.03	46.03
3	66.38	84.76	84.76	76.01
4	82.52	82.52	82.52	82.52
5	49.28	49.28	49.86	49.86
6	54.62	54.62	54.62	54.62
7	60.19	62.49	62.49	62.49
8	62.27	69.78	65.41	65.41
9	66.12	69.31	66.12	66.12
10	42.53	42.53	42.53	42.53
AVG	57.93	61.04	61.65	60.78
PVSS		5.37%	6.43%	4.92%

combined deterministic–stochastic approaches. We report the average optimal objective function achieved by the stochastic problem, SP, and the average gap to all three deterministic strategies (i.e. VSS/RP), since we believe that, in this context, this is more significant than the absolute value of VSS. We call this value *Percentage VSS (PVSS)*. Low gaps show that the value of the stochastic problem is not so relevant and that the problem can easily be approximated with the deterministic counterpart. Conversely, high PVSS values show that the stochastic component cannot be neglected. Differently from the previous analyses in which the deterministic strategies do not allow any recourse action, here the deterministic problems (DP, TDP, and ATDP) are solved to derive the set of bids to select in the regular auction and let the SP model make the optimal decisions in the reserve and the recourse auction. Note that the set of available recourse actions depends on the decisions made in the reserve auction.

Results are reported in Tables 5 and 6. The results of the analysis show relatively high values of PVSS (near to 10% on average), independently of the average level of trustability of the drivers and of the average rejection cost. This proves the importance of solving the stochastic problem instead of its deterministic counterpart. Differently from the previous analyses, in which a clear dominance of TDP and ATDP with respect to DP was observed, enabling recourse actions, the three deterministic strategies to select regular auction bids show a very similar performance. This observation means that recourse actions are useful to recover from poor decisions made at the first stage. Moreover, all the three deterministic strategies, if enabling recourse actions, can achieve small rejection rates (for the basic scenario) of 2.32%, 1.77%, and 2.09% for DP, TDP, and ATDP, respectively, which are comparable with the ones obtained by SP (1.69%). These rates prove that recourse actions can provide very competitive solutions from the customers satisfaction point of view but are much more expensive for the companies. In fact, concerning the total cost for the company, although enabling recourse actions reduces the gap to SP from 27% to 10%, the resulting gap is still considerable and fully justifies the implementation of a stochastic model.

6.4. Out-of-sample scenario analysis

The drawback of the VSS is that it is computed analyzing the same set of scenarios which were used in the stochastic problem. In other words, the solution obtained by SP is perfectly *tuned* on those scenarios, while the deterministic problems are solved without exploiting any information about the set of scenarios on which they will be evaluated and this could yield an advantage to the SP. Therefore, to have a fairer comparison of the approaches we perform an analysis of out-of-sample scenarios. For this, we consider a new set of scenarios (out-of-sample scenarios), generated from the same distribution as the in-sample ones. The number of these scenarios is much smaller than the ones of the in-sample set, since the scope of our analysis is to verify if the optimal solution of SP performs well also on scenarios which do not exactly correspond to the large set of in-sample scenarios on which it has been computed. In Table 7, we report, for each instance, with medium flexibility and medium rejection costs, the value of the optimal solution obtained by the different approaches. The last two rows report the average results and the gap between the SP solution and each one of the deterministic problems. It should be noted that, when comparing on out-of-sample scenarios, the relation $SP \leq EEV$, which is always true for in-sample scenarios, does not necessarily hold anymore. Our results show a value of PVSS around 5%–6% which means that considering the stochastic problem is worth the computational effort. SP obtains the best performance in 9 over 10 instances, and it is only slightly outperformed by DP (but clearly outperforms TDP and ATDP) on the remaining one. Conversely, all the three deterministic approaches perform very well in some instances but very bad on other ones (e.g., the third one). However, given the increasing computational efforts needed by SP on very large instances, the three deterministic approaches, and in particular ATDP, could be a suitable tool to address very large instances which cannot be handled by SP, providing good quality results. Comparing the three deterministic approaches, it is important to point out that ATDP, which explicitly takes into account drivers' trustability, by ensuring a minimum average level of trustability among selected ODs, is performing better than DP that totally neglects the trustability aspect. What is surprising is that TDP, which weights the offers of the drivers by their trustability, is performing worse than DP. This means that the weight given by TDP to trustability in the decision process, is too large. In fact, this weighting system strongly penalizes good offers provided by low trustable drivers, neglecting the fact that, thanks to the recourse policy allowed by the reserving options, accepting those offers could be convenient. This means that it could happen that reserved drivers can cover the customers who remained unserved at a reasonably low cost, in case of OD's absence. In fact, the reservation option effectively mitigates the impact of drivers absenteeism, making drivers' trustability a less crucial element in the decision process.

6.5. Perfect identikit of reserved drivers

The last analysis covers the identification of a perfect identikit of reserved drivers. For this, the characteristics of the reserved drivers are analyzed in the basic version of the first set of instances with 20 customers, 200 scenarios, medium flexibility, and medium rejection costs. The average results are reported in Table 8. The characteristics recorded are as follows.

- Flexibility (FLEX), which is given by the maximum detour (in Km) the driver is willing to accept.
- Willingness to work (WILL), which represents how much drivers are willing to make a discount with respect to their true costs; a willingness level of 1 indicates no discount (i.e., the bid is

Table 8
Average values (AVG) and ranges (RANGE) of parameters describing reserved drivers characteristics in obtained optimal solutions.

	FLEX	WILL	τ_j	r_j
AVG	6.89	0.77	0.74	2.11
RANGE	[1;10]	[0.6;1]	[0.6;1]	[1;10]

truthful=. A willingness less than 1, however, means that the driver is willing to offer a discount to have a higher probability to be selected.

- Trustability level (τ_j), computed as the number of times the driver being present divided by the number of times this driver being selected. If a drivers are new in the system and has no historical data, they are assigned a default trustability level equal to the average among all the drivers registered in the system. This default value is kept until they win a minimum number of bids, which are then used to calculate their individual values.
- The value of the reservation bids (r_j), computed as the price the driver is offering to act as a reserved driver.

In Table 8, the average and the range of values of these four parameters are reported for all the drivers selected as reserved drivers.

What can be observed is that the main characteristic of a reserved driver is cheapness. In fact, the most relevant parameter is the reservation cost. A reserved driver shows an average reservation cost of 2.11, while the average value across all the drivers' population is 5.5. Willingness to work is not very relevant since the perfect driver has a value slightly lower than the average. The flexibility seems to be important as the average value is 6.89, while the average among all drivers is 5, but not as relevant as reservation cost. Finally, trustability negatively influences the probability of a driver being selected as a reserved driver. This is not unexpected. In fact, highly reliable drivers do not need to be reserved, as the probability they will be present is very high. The reservation is more effective if applied to drivers who are not reliable, to eliminate a potentially higher risk of absenteeism.

While our computational study focuses on algorithmic aspects, several managerial insights of practical relevance are derived (e.g., optimal identikits of drivers). These insights together with the strengths of the proposed auction system could be used by last-mile platform providers to offer high value services to their customers. While the underlying decision problems are extremely challenging, our study shows that the proposed methods are applicable to real world problems.

7. Conclusions and future research

This paper studied the assignment of bundles of customer requests to ODs through an auction-based system under the presence of uncertain drivers. An absenteeism mitigation strategy is proposed for ODs based on reservation of the driver. According to this strategy, drivers can bid not only for serving bundles which they find attractive but also to act as a reserved driver. Reserved drivers receive a fee to guarantee their presence but are not guaranteed to be assigned to a specific bundle. The problem has been modeled as a two-stage stochastic problem with recourse activation. The model involves three types of auctions: (i) the regular auction in which bids are selected, (ii) the reserve auction in which reserve drivers are selected, and (iii) the recourse auction in which unserved customers are assigned to reserved drivers, due to the absence of other ODs. The first two auctions take place at the first stage, while the third one at the second stage, once absences are revealed.

This model differs from classical two-stage stochastic models because, to be allowed to implement recourse actions, an additional activation cost (here represented by reservation of drivers) has to be paid. Moreover, the set of available recourse actions depends on the decision made within the reservation auction, and therefore on a different stage. In fact, the set of bundles of customers a driver is willing to serve affects the possible recourse actions available. If a customer is not inserted in the set of bundles serviceable by reserved customers, it cannot be covered by a recourse action, which means that, in the case of the absence of the regular driver to which it has been assigned, it must be rejected by the company.

To solve this problem, a self-learning matheuristic (SLM) and an ILS (which exploits SLM as a local search operator) were developed. Both methods show a very good performance within short computational times. A rich computational analysis is also presented, in which the effectiveness of applying the reservation strategy was first analyzed, with respect to solving a deterministic problem without the possibility of recourse actions. Furthermore, the stochastic approach was compared with three different deterministic approximations, two of which take drivers' trustability levels into account. Results showed that the advantage of applying the newly proposed reservation strategy yields huge benefits both in terms of costs for the company and in terms of percentage of customers rejected. The latter decreases from 19.10% to only 1.69%.

The second analysis relates to a well-known stochastic parameter (VSS). The solution of the stochastic model was compared with three strategies based on solving the regular auction as a deterministic problem and then forcing the stochastic model to select the bids selected by the deterministic one, letting it free to optimize the reservation and recourse decisions. Even in this case the stochastic solution significantly outperforms the other approaches, showing the importance of explicitly addressing the stochastic nature of the problem. The stochastic solution also performs well, with respect to deterministic approximations, when applied on out-of-samples solutions.

The last analysis aimed at providing the identikit of the perfect reserved driver, based on data observed in optimal solutions.

Future developments from a methodological point of view could address the generalization of the two-stage stochastic problem with recourse activation to a broad class of problems. This approach would mainly apply to situations, in which the effects of disruptions can be reduced by applying mitigation strategies. From an application point of view, analyzing dynamic problems would be interesting, in which multiple auctions take place in different points in time during the planning horizon. Finally, empirical studies (e.g., real-word observations of driver behavior) would be worthwhile integrated in order to gain deeper insights on the impact of the human factor.

CRediT authorship contribution statement

Simona Mancini: Writing – original draft, Validation, Supervision, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Margaretha Gansterer:** Writing – original draft, Validation, Supervision, Software, Methodology, Investigation, Data curation, Conceptualization. **Chefi Triki:** Writing – original draft, Validation, Supervision, Software, Methodology, Investigation, Data curation, Conceptualization.

Acknowledgment

The work of Simona Mancini has been supported by the funding D26_PREMIO_SINGOLI_RIC_2023_MANCINI_SIMONA -Singoli Ricercatori 2023 from the University of Palermo.

Data availability

Data will be made available on request.

References

Ahamed, T., Zou, B., Farazi, N.P., Tulabandhula, T., 2021. Deep reinforcement learning for crowdsourced urban delivery. Transp. Res. B 152, 227–257.

Al Hla, Y.A., Othman, M., Saleh, Y., 2019. Optimising an eco-friendly vehicle routing problem model using regular and occasional drivers integrated with driver behaviour control. J. Cleaner Prod 234, 984–1001.

Albertzeth, G., Pujawan, I., 2018. Cost-effective strategy to mitigate transportation disruptions in supply chain. In: IOP Conference Series: Materials Science and Engineering. 337, (1), IOP Publishing, 012042.

Albertzeth, G., Pujawan, I.N., Hilletofth, P., Tjahjono, B., 2020. Mitigating transportation disruptions in a supply chain: a cost-effective strategy. Int. J. Logist. Res. Appl. 23 (2), 139–158.

Anderson, M.J., 2022. Integer and Stochastic Programming Approaches to Scheduling Nursing Staff at Long-Term Care Homes (Ph.D. thesis). University of Toronto (Canada).

Andrade-Michel, A., Ríos-Solís, Y.A., Boyer, V., 2021. Vehicle and reliable driver scheduling for public bus transportation systems. Transp. Res. B 145, 290–301.

Archetti, C., Guerriero, F., Macrina, G., 2021. The online vehicle routing problem with occasional drivers. Comput. Oper. Res. 127, 105144.

Archetti, C., Savelsbergh, M., Speranza, M.G., 2016. The vehicle routing problem with occasional drivers. European J. Oper. Res. 254 (2), 472–480.

Arslan, A., Agatz, N., Kroon, A., Zuidwijk, R., 2019. Crowdsourced delivery: A dynamic pickup and delivery problem with Ad-Hoc drivers. Transp. Sci. 53 (1), 222–235.

Ausseil, R., Pazour, J.A., Ulmer, M.W., 2022. Supplier menus for dynamic matching in peer-to-peer transportation platforms. Transp. Sci. 56 (5), 1304–1326.

Bakrania, S., Whittington, C., Anderson, T., Nates, R., 2018. Understanding student absenteeism in undergraduate engineering programmes. In: 29th Australasian Association for Engineering Education Conference 2018 (AAEE 2018). Engineers Australia Hamilton, New Zealand, pp. 37–42.

Behrend, M., Meisel, F., 2018. The integration of item-sharing and crowdshipping: Can collaborative consumption be pushed by delivering through the crowd? Transp. Res. B 111, 227–243.

Beraldi, P., Conforti, D., Triki, C., Violi, A., 2004. Constrained auction clearing in the Italian electricity market. 4OR, A Quart. J. Operat. Res 2, 35–51.

Birge, J.R., Louveaux, F., 2011. Introduction to Stochastic Programming. Springer.

Bortolini, M., Calabrese, F., Galizia, F.G., 2022. Crowd logistics: A survey of successful applications and implementation potential in Northern Italy. Sustainability 14 (24), 16881.

Boysen, N., Fedtke, S., Schwerdfeger, S., 2021. Last-mile delivery concepts: a survey from an operational research perspective. OR Spectrum 43, 1–58.

Burke, E., Gendreau, M., Hyde, G., Ochoa, E., Qu, R., 2013. Hyper-heuristics: a survey of the state of the art. J. Oper. Res. Soc. 64, 1695–1724.

Cheng, S.-F., Chen, C., Kandappu, T., Lau, H.C., Misra, A., Jaiman, N., Tandriansyah, R., Koh, D., 2017. Scalable urban mobile crowdsourcing: Handling uncertainty in worker movement. ACM Trans. Intell. Syst. Technol. 9 (3), 1–24.

Cheng, R., Jiang, Y., Nielsen, O., 2023. Integrated people-and-goods transportation systems: from a literature review to a general framework for future research. Transport Rev (ISSN: 0144-1647).

Dahle, L., Andersson, H., Christiansen, M., 2017. The vehicle routing problem with dynamic occasional drivers. In: Computational Logistics: 8th International Conference, ICCL 2017, Southampton, UK, October 18-20, 2017, Proceedings 8. Springer, pp. 49–63.

Dahle, L., Andersson, K., Christiansen, M., Speranza, M.G., 2019. The pickup and delivery problem with time windows and occasional drivers. Comput. Oper. Res. 109, 122–133.

Dayarian, I., Savelsbergh, M., 2020. Crowdshipping and same-day delivery: Employing in-store customers to deliver online orders. Prod. Oper. Manage. 29 (9), 2153–2174.

Di Puglia Pugliese, L., Ferone, D., Festa, P., Guerriero, F., Macrina, G., 2022. Solution approaches for the vehicle routing problem with occasional drivers and time windows. Optim. Methods Softw. 37 (4), 1384–1414.

Diab, E.I., Wasfi, R.A., El-Geneidy, A.M., 2014. Extraboard team sizing: An analysis of short unscheduled absences among regular transit drivers. Transp. Res A: Policy Pract 66, 27–38.

Dos Santos, A.G., Viana, A., Pedroso, J.P., 2022. 2-echelon lastmile delivery with lockers and occasional couriers. Transp. Res. E: Logist. Transp Rev. 162, 102714.

Drake, J.H., Kheiri, A., Özcan, E., Burke, E.K., 2020. Recent advances in selection hyper-heuristics. European J. Oper. Res. 285 (2), 405–428.

Figliozzi, M.A., Zhang, Z., 2009. A Study of Transportation Disruption Causes and Costs in Containerized Maritime Transportation. 207815, Transportation Research Forum.

Gansterer, M., Hartl, R.F., 2018. Centralized bundle generation in auction-based collaborative transportation. OR Spectrum 40 (3), 613–635.

Gansterer, M., Hartl, R.F., Sörensen, K., 2020. Pushing frontiers in auction-based transport collaborations. Omega 94, 102042.

Gansterer, M., Hartl, R.F., Vetschera, R., 2018. The cost of incentive compatibility in auction-based mechanisms for carrier collaboration. Networks 74 (4), 490–514.

Gdowska, K., Viana, A., Pedroso, J.P., 2018. Stochastic last-mile delivery with crowdshipping. Transp. Res. Procedia 30, 90–100.

Gurtu, A., Johny, J., 2021. Supply chain risk management: Literature review. Risks 9 (1), 16.

Hägele, S., Grosse, E.H., Ivanov, D., 2023. Supply chain resilience: a tertiary study. Int. J. Integr. Supply Manag 16 (1), 52–81.

Haughton, M.A., 2009. An alternative tactic to deal with the contingency of driver absenteeism. J. Oper. Res. Soc. 60 (9), 1207–1220.

Horner, H., Pazour, J., Mitchell, J., 2021. Optimizing driver menus under stochastic selection behavior for ridesharing and crowdsourced delivery. Transp. Res. E: logistics and transportation review 153, 102419.

Hosseini, S.M.H., Behroozi, F., Sana, S.S., 2023. Multi-objective optimization model for blood supply chain network design considering cost of shortage and substitution in disaster. RAIRO-Operat. Res 57 (1), 59–85.

Hou, S., Gao, J., Wang, C., 2022. Optimization framework for crowd-sourced delivery services with the consideration of shippers' acceptance uncertainties. IEEE Trans. Intell. Transp. Syst..

Ibisworld, 2023. URL https://www.ibisworld.com/industry-statistics/number-of-businesses/couriers-local-delivery-services-united-states/#:~:text=There%20are%20456%2C508%20Couriers%20%26%20Local,over%20the%20past%205%20years%3F, Last accessed on 18 Sept 2023.

Ivanov, D., Mason, S.J., Hartl, R.F., 2016. Supply chain dynamics, control and disruption management. Int. J. Product. Res 54 (1), 1–7.

Karimi-Mamaghan, M., Mohammadi, M., Meyer, P., Karimi-Mamaghan, A.M., El-Ghazali, T., 2022. Machine learning at the service of meta-heuristics for solving combinatorial optimization problems: A state-of-the-art. European J. Oper. Res. 296 (2), 393–422.

Katsaliaki, K., Galetsi, P., Kumar, S., 2021. Supply chain disruptions and resilience: A major review and future research agenda. Ann. Oper. Res. 1–38.

King, N., Dewey, C., Borish, D., 2015. Determinants of primary school non-enrollment and absenteeism: results from a retrospective, convergent mixed methods, cohort study in rural western Kenya. PLoS One 10 (9), e0138362.

Lalang, D., Making, S.R.M., Manapa, I.Y.H., 2019. Modulating the multi-depot vehicle routing problem model with time windows using occasional driver. Int. J. Innov. Creativity Change 5 (3), 39–60.

Le, T.V., Ukkusuri, S.V., 2019. Modeling the willingness to work as crowd-shippers and travel time tolerance in emerging logistics services. Travel Behav. Soc. 15, 123–132.

Lourenço, H., Martin, O., Stützle, T., 2018. In: Gendreau, M., Potvin, J.-Y. (Eds.), Handbook of Metaheuristics. Springer.

Macrina, G., Di Puglia Pugliese, L., Guerriero, F., 2020. Crowd-shipping with time windows and transshipment nodes. Comput. Oper. Res. 113, 104806.

Macrina, G., Di Puglia Pugliese, L., Guerriero, F., Laganà, D., 2017. The vehicle routing problem with occasional drivers and time windows. In: Sforza, A., Sterle, C. (Eds.), Optimization and Decision Science: Methodologies and Applications. Springer International Publishing, Cham, ISBN: 978-3-319-67308-0, pp. 577–587.

Macrina, G., Guerriero, F., 2018. The green vehicle routing problem with occasional drivers. In: Daniele, P., Scrimali, L. (Eds.), New Trends in Emerging Complex Real Life Problems: ODS, Taormina, Italy, September 10–13, 2018. Springer International Publishing, Cham, pp. 357–366.

Maggioni, F., Wallace, S., 2012. Analyzing the quality of the expected value solution in stochastic programming. Ann. Oper. Res. 200, 37–54.

Mancini, S., Gansterer, M., 2022. Bundle generation for last-mile delivery with occasional drivers. Omega 108, 102582.

Mancini, S., Gansterer, M., 2024. Bundle generation for the vehicle routing problem with occasional drivers and time window. Flexib. Serv. Manuf. J http://dx.doi.org/10.1007/s10696-023-09529-3.

Mancini, S., Gansterer, M., Triki, C., 2023. Locker box location planning under uncertainty in demand and capacity availability. Omega 120, 102910.

Mancini, S., Gansterer, M., Triki, C., 2024. Data for an effective mitigation strategy to hedge against absenteeism of occasional drivers. Mendeley Data URL http://dx.doi.org/10.17632/3bnjsssxph.1.

Mancini, S., Triki, C., Piya, S., 2022. Optimal selection of touristic packages based on user preferences during sports mega-events. European J. Oper. Res. 302 (3), 819–830.

Mousavi, K., Bodur, M., Roorda, M.J., 2022. Stochastic last-mile delivery with crowd-shipping and mobile depots. Transp. Sci. 56 (3), 612–630.

Muñoz-Villamizar, A., Solano-Charris, E.L., Reyes-Rubiano, L., Faulin, J., 2021. Measuring disruptions in last-mile delivery operations. Logistics 5 (1), 17.

Musmanno, R., Scordino, N., Triki, C., Violi, A., 2010. A multistage formulation for generation companies in a multi-auction electricity market. IMA J. Manag. Math. 21 (2), 165–181.

Pahwa, A., Jaller, M., 2023. Assessing last-mile distribution resilience under demand disruptions. Transp. Res. E: Logistics and Transportation Review 172, 103066.

Rechavi, A., Toch, E., 2022. Crowd logistics: Understanding auction-based pricing and couriers' strategies in crowdsourcing package delivery. J. Intell. Transp. Syst 26 (2), 129–144.

Rivera-Royero, D., Galindo, G., Jaller, M., Reyes, J.B., 2022. Road network performance: A review on relevant concepts. Comput. Ind. Eng. 107927.

Rockafellar, R., Wets, R.-B., 2017. Scenarios and policy aggregation in optimization under uncertainty. Math. Oper. Res. 16, 119–147.

Ropke, S., Pisinger, D., 2006. An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. Transp. Sci. 40 (4), 455–472.

Sheffi, Y., 2004. Combinatorial auctions in the procurement of transportation services. Interfaces 34 (4), 245–252.

Shepherd, 2023. URL https://thesocialshepherd.com/blog/ecommerce-statistics#:~:text=eCommerce%20Sales%20Are%20Projected%20to%20Reach%20%247.4%20Trillion%20by%202025&text=Over%20the%20next%20four%20years,approximately%20%247.4%20trillion%20by%202025, Last accessed on 18 Sept. 2023.

Silva, M., Pedroso, J.P., Viana, A., 2023. Stochastic crowd shipping last-mile delivery with correlated marginals and probabilistic constraints. European J. Oper. Res. 307 (1), 249–265.

Statista, 2023. URL https://www.statista.com/statistics/534123/e-commerce-share-of-retail-sales-worldwide/, Last accessed on 18 Sept. 2023.

Torres, F., Gendreau, M., Rei, W., 2022a. Crowdshipping: An open VRP variant with stochastic destinations. Transp. Res. C 140, 103667.

Torres, F., Gendreau, M., Rei, W., 2022b. Crowdshipping: An open VRP variant with stochastic destinations. Transp. Res. C 140, 103677.

Triki, C., 2021. Using combinatorial auctions for the procurement of occasional drivers in the freight transportation: A case-study. J. Clean. Prod. 304, 127057.

Triki, C., Akil, J., Al-Asmakh, H., 2023. Optimisation models for the procurement through reverse combinatorial auctions in the logistics and food industries. Int. J. Procurement Manag 16 (4), 530–548.

Voigt, S., Kuhn, H., 2022. Crowdsourced logistics: The pickup and delivery problem with transshipments and occasional drivers. Networks 79 (3), 403–426.

Wang, D., Ozbay, K., 2023. Multi-stage equitable bus-based hurricane evacuation model with a stochastic driver availability component. Transp. Res. Record 2677 (2), 1619–1632.

Wang, L., Xu, M., Qin, H., 2023. Joint optimization of parcel allocation and crowd routing for crowdsourced last-mile delivery. Transp. Res. B 171, 111–135.

Windras Mara, S.T., Norcahyo, R., Jodiawan, P., Lusiantoro, L., Rifai, A.P., 2022. A survey of adaptive large neighborhood search algorithms and applications. Comput. Oper. Res. (ISSN: 0305-0548) 146, 105903.

Yu, V., Aloina, G., Jodiawan, P., Gunawan, A., Huang, T.-C., 2023. The vehicle routing problem with simultaneous pickup and delivery and occasional drivers. Expert Syst. Appl. 214, 119118.

Yu, V., Jodiawan, P., Hou, M.-L., Gunawan, A., 2021. Design of a two-echelon freight distribution system in last-mile logistics considering covering locations and occasional drivers. Transp. Res. E: Logistics and transportation review 154, 102461.

Yu, V., Jodiawan, P., Perwira Redi, A., 2022. Crowd-shipping problem with time windows, transshipment nodes, and delivery options. Transp. Res. E: Logistics and transportation review 157, 102461.