



Contents lists available at ScienceDirect

European Journal of Operational Research

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

Decision Support

Demand management in time-slotted last-mile delivery via dynamic routing with forecast orders

Mohammad Abdollahi^{a,*}, Xinan Yang^a, Moncef Ilies Nasri^b, Michael Fairbank^c^a Department of Mathematics, University of Essex, Colchester CO4 3SQ, United Kingdom^b Ocado Technology, Hatfield AL10 9UL, United Kingdom^c Department of Computer Science and Electrical Engineering, University of Essex, Colchester CO4 3SQ, United Kingdom

ARTICLE INFO

Article history:

Received 30 June 2022

Accepted 13 January 2023

Available online xxx

Keywords:

Decision support systems

Dynamic pricing

Forecast orders

Attended home delivery

Demand management

ABSTRACT

In this paper, we propose a partially time-windowed dynamic routing approach with forecast orders to tackle the dynamic pricing problem of attended home delivery, one of the challenging problems in last-mile logistics. The purpose of forecast orders is to find a cost-effective route map for delivery with forecast orders to guide the delivery system to accept real orders in potentially better time slots for servicing. Initially, we build and optimise virtual routes using forecast orders without time windows. Upon the arrival of real orders, we replace the forecast orders with incoming real orders to help the delivery system estimate the opportunity cost of making the delivery in each time slot. Dynamic pricing is then used to influence the customer's choices of servicing slots, and based on what the customer selects, the real order replaces the forecast order, and the route map is re-optimised. This strategy leads to better scheduling for the overall route map and more slot availability for incoming customers. Through computational study, we demonstrate the benefit of our approach compared to static pricing and previous dynamic pricing policies, with or without forecast orders. In particular, by employing forecast orders without time windows, we witness a decrease in delivery costs and an increase in the number of accepted orders, leading to higher profits.

© 2023 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

With the rapid growth in Internet and mobile-communication networks, online retailing has started to make up a higher and higher share of grocers' revenue. The service allows customers to shop for goods online and to deliver them directly to their front door. Attended Home Delivery (AHD) is amongst the most crucial topics that have been considered by e-grocers maintained by Wang et al. (2014) for the reason that most grocery orders may contain perishable and frozen goods that need to be dealt with immediately explained by Agatz et al. (2013). As Fleckenstein et al. (2022) and Waßmuth et al. (2022) recently surveyed, the AHD problem integrates both a demand-management problem and a variant of a vehicle-routing problem that generally holds a common framework of providing AHD services by sending out delivery vans to visit customers within committed time slots to drop the orders, which leads to a logistics challenge to effectively manage

the geographical location of customers and their time-window requirements. As customer satisfaction typically relates to the time they have to wait for the delivery, offering narrower time slots becomes one of the significant considerations of AHD services stated by Agatz et al. (2011). A narrower time slot complicates the task of optimising delivery routes, as the shorter the time window, the fewer options are available for scheduling a delivery.

Nowadays, most online grocers use fixed, one-hour time slots. Naturally, time slots receive different customer preferences. Suppose nothing is done through the booking horizon. In that case, the company will likely have to face unbalanced demand for time slots, which leads to inefficient delivery routes and a potential waste of fleet capacity. To deal with this issue, researchers developed demand-management strategies in the past and recent literature aiming to control/incentivise customers to select specific time slots to ease the routing difficulty, including Agatz et al. (2013); Asdemir et al. (2009); Campbell & Savelsbergh (2006); Cleophas & Ehmke (2014); Koch & Klein (2020); Yang et al. (2016). Which time slot is the best to promote is the key question to answer in this context, which, however, has not yet been fully addressed by the existing literature. Difficulties in answering this question lie in two facts:

* Corresponding author.

E-mail addresses: ma20139@essex.ac.uk (M. Abdollahi), xyangk@essex.ac.uk (X. Yang), moncef.nasri@ocado.com (M.I. Nasri), m.fairbank@essex.ac.uk (M. Fairbank).

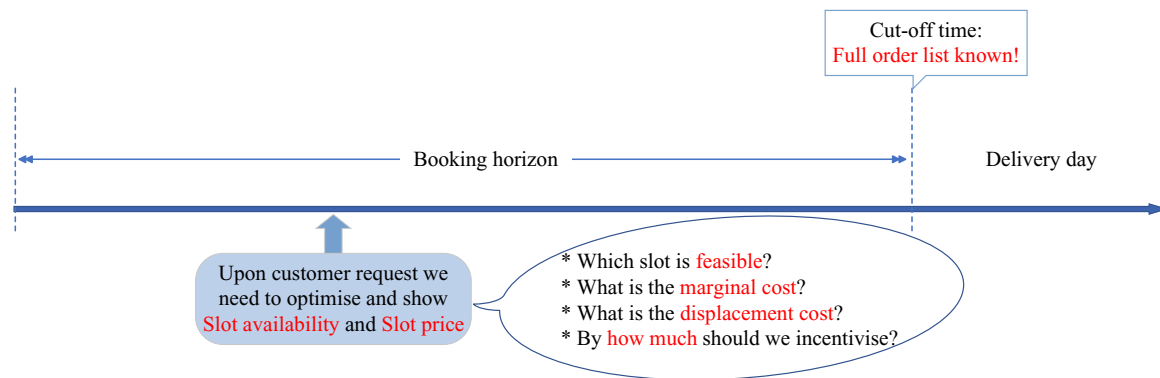


Fig. 1. Time frame.

1. As shown in Fig. 1, the best slot for satisfying an order depends on the insertion cost for placing that order into the final delivery route and the potential revenue loss of displacing another order. However, the full order list is not known during the booking horizon when decisions on which slot to promote have to be made. Therefore, a forecast is needed for the final delivery route so that when the booking is made, the extra cost caused by the likely route deviation to satisfy that new order can be estimated reasonably.
2. The likelihood of a customer accepting an order and its time slot depends on the promotional decisions made during the booking horizon, so the final delivery route is not entirely predictable from historical information where, essentially, a different pricing/order-acceptance approach was used.

There is a mutual dependency between the forecast delivery route and the incentive decision. This work aims to propose a methodology based on dynamic routing to break this loop to allow reasonable demand forecasting while maintaining the dependency of one on the other to the maximum extent.

Past literature in AHD concentrates primarily on how to deal with time windows, both in the demand-management step (via pricing or other incentive means) and in solving the so-called Capacitated Vehicle Routing Problems with Time Windows (CVRPTW) affirmed by Kumar & Panneerselvam (2012). While Yang et al. (2016) and Koch & Klein (2020) suggest route-based approaches by which a virtual route map with time windows is created to predict the final route map to help the CVRPTW find the best time window for servicing an order, an easily overlooked fact is that the best time window can be found by solving the VRP *without* time windows. Inspired by this idea, in this study, we propose a novel dynamic routing and pricing approach based on solving and updating a partially time-windowed CVRP with a combination of real and forecast orders. Specifically, at any time during the booking horizon, we maintain a set of already-committed orders, with fixed known time windows, locations and order sizes, and a set of forecast orders, with given order sizes and locations but without any time windows. We make this distinction because the locations and sizes of forecast orders can be more reliably concluded from historical data (than delivery slots), as they are not meant to be influenced by the incentive policy. As for the delivery-slot choices of forecast orders, however, we do not impose any time windows. Instead, we allow the dynamic routing approach to choose its preferred time slot for each order, which is then offered to customers following the incentive policy. Essentially, this approach assumes that future customers will generally choose the incentivized time slots to receive their deliveries, which aligns with our ultimate aim of developing and deploying an incentive policy to steer customer choices of time slots. In detail, at any point in time, we

solve a partial CVRPTW (p-CVRPTW) made up of accepted orders with time windows and forecast orders without time windows. We feed back the best time window for satisfying the forthcoming customer, based on the solution of the p-CVRPTW. This p-CVRPTW is solved dynamically online and updated whenever new orders are committed.

The major contributions of this article are:

- For the first time, incorporating forecast orders *without time windows* into the vehicle-routing system, to allow the p-CVRPTW to suggest the best time slot to accommodate every forecast order and guide the choice of incoming orders accordingly;
- Proposing a simple-to-implement dynamic opportunity-cost approximation for marginal delivery cost and potential revenue loss, based on the dynamically managed routing system with both actually accepted orders and forecast orders without time windows;
- Presenting an order-replacement and routing re-optimisation framework to capture the influence of new order commitments and facilitate opportunity-cost approximation, which evolves as more information becomes available;
- Presenting an approach that is capable of incorporating the firm's specific routing method, which may include considerations such as clustering, shifting, traffic prediction, etc., to the maximum extent;
- Demonstrating the superiority of the developed approach over four benchmark approaches, on real data-sets taken from four typical geographical and demographic settings;
- Investigating the trade-off between responding time and accuracy of the online decision process.

The article is organised as follows: In Section 2, we explore the previous studies and recent research conducted in the area of the AHD problem. Section 3 explains different aspects of the AHD problem and its dynamic programming model. Section 4, presents our methodology and how to incorporate forecast orders, pricing optimisation and the customer-behaviour model. The experiment settings and results obtained are reported in Section 5. Finally, we conclude in Section 6.

2. Literature review

Following the standard categorisation of demand management in time-slotted deliveries by Agatz et al. (2008), literature can be clustered into four main groups: static slotting (e.g., Agatz et al., 2011), dynamic slotting (e.g., Campbell & Savelsbergh, 2005), static pricing (e.g., Klein et al., 2017) and dynamic pricing (e.g., Campbell & Savelsbergh, 2006). Slotting focuses on time-slot allocation to customer regions. In contrast, pricing assigns delivery prices to bal-

ance demands across time slots and/or steer customer choices towards the best delivery-time options. Amongst static and dynamic approaches, the latter attracts more attention as it reflects the nature of online booking, where customers place their delivery requests dynamically over the booking horizon. Readers are referred to the surveys by Klein et al. (2020) and Snoeck et al. (2020) on the industrial application of revenue management and advances in choice-based models for more information. This article concentrates on dynamic approaches to demand management with pricing. Closely related works in this category are examined in this section.

From the grocer's point of view, delivering an order in a time slot or another typically yields different costs, motivating the company to steer customers' selection of time slots to increase its profit. One can apply many promoting strategies to achieve this aim, such as offering discounted delivery prices/discount vouchers for some slots, highlighting the slot in different colours to reflect their pollution/environmental impacts, etc., Campbell & Savelsbergh (2006) benefit from price discounts to encourage the selection of time slots with cheaper insertion costs, where routing decisions are simplified with an insertion heuristic based on accepted orders. In this work, we consider monetary incentives, which are displayed in terms of reduced delivery prices. In practice, other types of incentives can also be translated into monetary representations, so the same approach applies to more complicated scenarios where some/all customers are not price-sensitive. This study aims to find the optimal incentivising policy in a stochastic and dynamic environment, where orders are collected during the booking horizon. Campbell & Savelsbergh (2006) also use a relatively simple model (linear) of customer behaviour to capture the effect of delivery prices on the probability of a particular time slot being chosen. This work is extended by Asdemir et al. (2009), who use the more advanced, Multinomial Logit (MNL) model to describe customer choices of delivery slots. They formulate a dynamic-pricing problem assuming that delivery costs are fixed and known a priori. This assumption removes the routing part of the problem from profit maximisation, which significantly simplifies the problem. In addition, the state space of the Dynamic Programming (DP) model proposed by Asdemir et al. (2009) grows exponentially in the number of delivery time slots, making it only applicable to impractically small scenarios. Similarly, Bühler et al. (2016) propose several linear mixed-integer programs to approximate delivery costs based on a fixed pool of potential routes. Further down this route, Yang et al. (2016) estimate an MNL choice model from real e-grocer data and numerically demonstrate that using this model for time-slot pricing to influence demand may improve overall profitability. They employ insertion heuristics to update a pool of feasible routes as orders come in over the booking horizon and deploy the marginal delivery cost as estimates of the opportunity cost of accepting an order into a particular time slot. The proposed "foresight" approach, which uses previously planned routes to bring in the effects of forecast orders, is justified as superior to the "hindsight" approach, which only considers accepted orders. While using forecast orders based on previously planned routes can help to build better routes, their method is restricted compared to our approach for the reason that such forecasts are not updated with the acceptance of new orders, and the routing for actual orders is independent of that for forecast orders.

Instead of using marginal delivery costs as estimates of the opportunity costs and concentrating on tentative routes and insertion heuristics, other works exist with a different strategy. They use alternative modelling approaches to simplify the underlying VRPTW solutions to emphasise the potential revenue loss by occupying the slot capacity. Most of these studies investigate the "acceptance scheme" of customer requests, such as Campbell & Savelsbergh (2005); Ehmke & Campbell (2014) and Cleophas & Ehmke

(2014), which aim to maximise the number of requests accepted for delivery. Mackert et al. (2019) proposed a new choice model based on a finite-mixture MNL choice model that can simplify the non-linear optimisation model to a linear one for solving. A model-based, profit-oriented slotting approach is developed to accurately approximate customer choice behaviour.

In dynamic pricing, works also fall under this category, which aims to bring future order acceptance effects by estimating the so-called opportunity cost. The opportunity cost of a delivery slot is interpreted as the potential future-order displacement cost when the current order occupies the limited capacity for this time slot. Given the large size and the stochastic nature of industrial applications in e-fulfilment, the opportunity cost cannot be calculated precisely. Approaches such as Approximate Dynamic Programming (ADP), Linear Approximations and look-ahead heuristics have been deployed to tackle the computational difficulties. In detail, Figliozzi et al. (2007) model the carrier-pricing problem in the dynamic vehicle-routing environment as a stochastic dynamic program, which is solved through a one-step look-ahead heuristic. Since the context is in-freight transportation, the approaches proposed by Figliozzi et al. (2007) are not readily applicable to AHD problems. In the AHD context, Klein et al. (2018) present an approximation approach based on a Mixed-Integer Linear Programming (MILP) reformulation to approximate opportunity costs. However, the MILP suffers from computational challenges; even with further simplifications and parallel computing, their approach has not proven suitable for scenarios with more than 15 vans. To solve the computational difficulty, Yang & Strauss (2017) exploit a continuous approximation of delivery costs and propose an ADP method that estimates the opportunity cost in real-time. Their approach is justified as efficient in industrial-size implementations; however, the routing approach used omits many practical restrictions to achieve this aim. In contrast, in our work here, we propose an approach that can directly deploy the routing package a grocer is currently using, which ensures all practical restrictions of routing are accommodated in the online pricing decisions.

The closest related work to ours is Koch & Klein (2020), which also uses forecast orders and preliminary routes to guide booking decisions. Koch & Klein (2020) propose using two routes, one starting from empty and one (the skeletal route plan) generated by repeatedly deploying a fixed pricing policy (the marginal delivery-cost approximation approach) and removing half of the orders randomly between iterations. New orders are inserted into both routes based on the lowest insertion costs. For the skeletal route plan, artificial orders are replaced by real orders upon acceptance. A major difference between our work and Koch & Klein (2020) lies in the way of generating and using forecast orders.

Unlike all previous works making use of forecast orders, including Koch & Klein (2020) and Yang et al. (2016), in this work, we do not use the "previous route" in our forecast because the "previous route" was constructed with the orders the system received following a fixed-demand scheduling policy (e.g., a fixed pricing policy or a fixed order-acceptance policy). These works consider previously allocated time windows when constructing the forecast route, which implicitly assumes that the previously allocated time window was optimal (or good enough), and assumes that repeating it (or guiding the system to reinforce it) will lead to preferable solutions. In this article, however, we avoid using previous routes and their time-window restrictions imposed on forecast orders. We start with the *optimal* route (or the best route one can find with a heuristic) *without* imposing any time windows to mimic the best possible route. We can do this by performing an "excellent" pricing policy in an ideal world, i.e., where all customers select what turns out to be the optimal time slot (from the route-optimisation software's point of view) to receive their orders. This strategy could be overly optimistic at the very beginning, so we adapt it as time

passes, gradually incorporating more actual information as orders arrive, and re-optimize the routing upon every committed order with a fixed, known time window. With this approach, we do not have to forecast time windows but only the number of orders and their locations. As the approach updates according to actual order arrivals, it is also robust concerning forecast errors. Please refer to Section 5.4 for detailed experiments and results on shifted forecast levels.

On a related note on dynamic routing with forecast orders, but without pricing to influence time-slot choices, our approach shares certain similarities with Bent & Van Hentenryck (2004); Ichoua et al. (2006) and Voccia et al. (2019). However, these works assume known probability distributions of future demands when generating future delivery requests, which, in our case, is unknown because the distribution is influenced by the dynamically-changing pricing policy our system generates. Recently, Soeffker et al. (2022) comprehensively discuss scenario-based approaches in which information models are integrated to tackle stochastic dynamic vehicle routing. Also, Klein & Steinhardt (2022) extend scenario-based approaches to address same-day deliveries and incorporate value function approximation approaches as support to include the dynamism over time to accomplish anticipatory decision-making. Likewise, Ulmer (2020) uses a value function approximation approach to find the best pricing policy in the same-day delivery problem, which proves effective. However, if these approaches rely on a lookup table, they suffer from the curse of dimensionality as the size of the vehicle's fleet grows.

Other articles discussing AHD resolutions under different problem settings exist, such as Dayarian & Savelsbergh (2020) exploring the employment of in-store customers to deliver online orders while they return home, Köhler et al. (2020); Strauss et al. (2020) working on the use of flexible time slots and Agatz et al. (2021) focusing on the effect of green labels. Further, Ojeda Rios et al. (2021) have recently published a survey on Dynamic Vehicle Routing Problems (DVRP) by providing taxonomies of the problems and solution methods. They reported that heuristics and meta-heuristics provided 66% of solutions to DVRP. In this article, we will use DVRP to estimate opportunity costs and guide choices of time slots.

3. Problem specification

For the problem under consideration, the company manages an online booking system that allows customers to book their delivery a couple of days in advance, which we refer to as the booking horizon. The orders committed during the booking horizon have to be delivered to the customer's front door during the agreed time slot by the company, using its fleet. Time slots are predefined by the company which may overlap. A scheme of the slot-booking process is shown in Fig. 2. To purchase goods and book for delivery, a customer has to log in to their account with the grocer, which allows the system to identify their address. We refer to this as a "customer arrival". Next, assuming the customer decides to place an order, the customer chooses their delivery day. This step may happen before or after filling their shopping basket. Note that we do not consider same-day delivery, which means once the truck has been loaded and dispatched, it does not need to come back to the depot to collect more orders until its planned route has finished. After the customer's selection of a delivery day, the company has to identify in real-time all the feasible time slots which could be used to service this order, together with their incentive scales. Based on this information, the customer chooses a time slot, finishing the order commitment. Decisions in this problem must be made in a stochastic dynamic environment, with randomness coming from both customer arrival and customer selection of delivery slots. The requirement for a fast response time, between a

customer's click of delivery day and the display of available slots and prices, adds another layer of difficulty to the problem.

3.1. Dynamic programming model

In this work, we inherit the Markov Decision Process (MDP) model formulated by Yang et al. (2016). We consider a discretized booking horizon with T periods, by which we mean customer arrivals to the website during the booking horizon shown in Fig. 1. Each booking period is sufficiently small such that the probability of having more than one arrival of a booking request is negligible. The final time period T denotes the cut-off time after which no further bookings are taken. The stages of the dynamic program are the time periods $t \in \{1, 2, \dots, T\}$. At time step t within the booking horizon, the system's state can be described by a matrix $X(t)$, with $|A|$ rows and $|S|$ columns. The $[a, s]$ th component of $X(t)$ represents the number of orders accepted up to time t (in the booking horizon), to be delivered to area $a \in A$ in time slot $s \in S$. In what follows, we use \vec{x}_t to denote the matrix $X(t)$ reshaped in column-major order to be stored in a one-dimensional array.

Let $V_t(\vec{x}_t)$ denote the value function at stage t and state \vec{x}_t ; it represents the expected maximum profit obtainable from the sales process from time t until the cut-off time T . The dynamic-programming recursion at stage $t \in \{1, 2, \dots, T\}$ is:

$$\begin{aligned} V_t(\vec{x}_t) &= \max_{\vec{d}} \left[\left(\sum_a \lambda \mu_a \sum_{s \in F_a(\vec{x}_t)} P_{s, F_a(\vec{x}_t)}(\vec{d}_a) \left[r_i + d_{as} + V_{t+1}(\vec{x}_t + \mathbf{1}_{as}) \right] \right) \right. \\ &\quad \left. + \left[1 - \sum_a \lambda \mu_a \sum_{s \in F_a(\vec{x}_t)} P_{s, F_a(\vec{x}_t)}(\vec{d}_a) \right] V_{t+1}(\vec{x}_t) \right] \\ &= \left[\max_{\vec{d}} \sum_a \lambda \mu_a \sum_{s \in F_a(\vec{x}_t)} P_{s, F_a(\vec{x}_t)}(\vec{d}_a) \left[r_i + d_{as} \right. \right. \\ &\quad \left. \left. - (V_{t+1}(\vec{x}_t) - V_{t+1}(\vec{x}_t + \mathbf{1}_{as})) \right] \right] \\ &\quad + V_{t+1}(\vec{x}_t), \quad \forall \vec{x}_t \in \mathcal{X}, \end{aligned} \quad (3.1)$$

where λ indicates the arrival rate of customer requests; μ_a denotes the probability that the arrival comes from area a for a given customer arrival; \vec{d}_a is a vector of length $|S|$ specific to area a , with components d_{as} , where component d_{as} represents the delivery price to area a at time slot s ; \vec{d} is a collection of \vec{d}_a over all areas; $F_a(\vec{x}_t) := \{s : C(\vec{x}_t + \mathbf{1}_{as}) < \infty\}$ denotes all feasible time slots for area a , into which order (a, s) can be feasibly inserted given orders \vec{x}_t that have been accepted, where $\mathbf{1}_{as}$ is the unit vector equal to the flattened single-entry matrix with a 1 in position (a, s) ; $P_{s, F_a(\vec{x}_t)}(\vec{d}_a)$ denotes the probability that a customer chooses slot s when the firm offers the vector of delivery prices \vec{d}_a to feasible slots in $F_a(\vec{x}_t)$; r_i denotes the revenue of the order i that is under consideration. The boundary condition for the MDP model is given by:

$$V_{T+1}(\vec{x}_{T+1}) = -C(\vec{x}_{T+1}) \quad \forall \vec{x}_{T+1} \in \mathcal{X}, \quad (3.2)$$

where $C(\vec{x}_t)$ represents the minimum cost of servicing all accepted orders during their agreed time slots, which is the optimal solution of a Capacitated Vehicle-Routing Problem with Time Windows (CVRPTW); \mathcal{X} denotes the set of all states that allow a feasible delivery schedule. If there is no feasible solution for a given \vec{x}_{T+1} , $C(\vec{x}_{T+1}) = +\infty$.

The dynamic program is intractable due to the large state space and the fact that the optimal solution of large-scale CVRPTW alone is intractable. Nevertheless, the formula (3.1) shows that the time-slot pricing decision is a trade-off between the immediate income, $(r_i + d_{as})$, and the expected opportunity cost $(V_{t+1}(\vec{x}_t) - V_{t+1}(\vec{x}_t + \mathbf{1}_{as}))$ arising from reserving the delivery capacity in (a, s) at time t

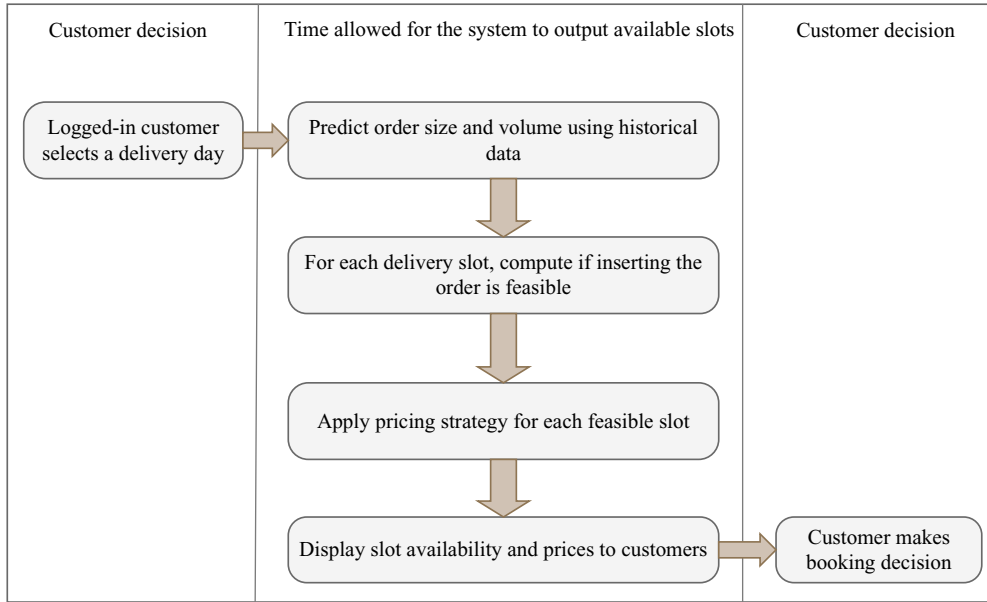


Fig. 2. Slot booking process.

for a future order. Suppose the opportunity cost can be estimated, then the problem can be divided into single-stage decision problems and becomes tractable. There are two major components of the opportunity cost:

1. the marginal delivery cost of servicing one more order in (a, s) , and
2. the potential revenue loss from filling fleet capacity at t with an order in (a, s) .

Both of these two terms depend on the final delivery routes.

This study aims to estimate the opportunity cost via dynamic routing with forecast orders. Both marginal delivery costs and potential revenue loss will be estimated by solving a CVRPTW dynamically over the booking horizon, with a set of already-accepted orders with fixed time windows and a set of forecast orders with relaxed time windows. More details about the approximation are discussed in Section 4.2.

4. Methodology

This section presents the solution methodology for the stochastic dynamic-pricing problem (3.1), which is intractable via backward induction. In detail, we will address how forecast orders are generated, integrated and updated in the dynamic setting and used to estimate opportunity cost in the following sub-sections.

4.1. Forecast orders

As explained, we aim to incorporate forecast orders into the dynamic routing process, to enable making better incentive decisions. Full information on orders in AHD consists of the order's arrival time, customer address, order size and delivery time window. In this work, however, we only forecast the total number of orders over a day, their addresses and order volumes, but not the delivery time window of each order. The reason is that, while we are optimising the incentive decision, we are aiming to steer customers' choices of time windows. Any forecasting model ignoring the impact of the incentive decision will not do a good job of predicting how many orders would select each time slot in the end. On the other hand, the incentive decision is optimised dynamically over

the entire booking horizon, changing over time and highly dependent on previously placed orders.

Therefore, we propose a simplification by assuming that all customers will select the time window most beneficial for the route planning to receive their orders, which we also assume is consistent with what the company aims to achieve by providing incentives. How to calculate the best time window for an upcoming order will be discussed in more detail in Section 4.2. Here we only need an approach to predict the total number of forecast orders and their locations/sizes and assume that all forecast orders are granted a 24-hour time window.

For the total number of orders on a specific delivery day, n , we use a Simple Moving-Average (SMA) model:

$$n = \frac{1}{k} \sum_{i=1}^k \hat{n}_i \quad (4.1)$$

where \hat{n}_i denotes the number of orders we received i weeks prior on the same weekday; k indicates the number of samples we consider for the prediction. We argue that this number n is not influenced significantly by the slot-price/incentives we offer, as the number of customers in a fixed area and their intentions to purchase from the e-retailer are mainly concerned with the demography of the area and the loyalty of customers. We also note that the moving average might be not the best possible approach that one can choose to predict the number of orders. More complicated machine-learning methods could be used to forecast the number of orders based on historical data. However, for this study, we only aim to demonstrate that incorporating forecast orders without any time windows into the routing process helps improve delivery efficiency and increases the total profit, even if a simplified model generates the forecast orders. For every single forecast order, its address, order size and order revenue are randomly simulated from historical data. Specifically, to generate one forecast order for a particular day of the week, we randomly choose (with uniform probability distribution) one order from the previous k weeks on that weekday and note its address, order size and revenue.

4.2. Opportunity cost approximation

As noted above, model (3.1) is not tractable for real implementations. In this section, we present an efficient approximation of it

using dynamic routing. Many firms use dynamic routing to build their fulfilment plan while orders are still collected.

This work calculates approximations of opportunity cost by creatively using the given dynamic-routing package to make the system (3.1) solvable in a practical dynamic setting. One key idea here is to distribute the delivery cost in boundary condition (3.2) into stages and calculate the incremental delivery cost of accepting one more order in a specific area and time window. Similar ideas have been shown to be effective by works such as [Campbell & Savelsbergh \(2006\)](#) and [Yang et al. \(2016\)](#); however, the approach used in this article is more advanced, since it considers the potential revenue loss as well as incremental delivery cost, by incorporating a forecast of future accepted orders.

The number of forecast orders and their locations/order sizes can be generated using the methodology presented in [Section 4.1](#). We indicate the list of forecast orders by a vector \vec{f} , with $|\vec{f}| = n$, ($|\cdot|$ indicates the cardinality of a set), where n is given by (4.1). These virtual orders are put into the problem to help predict the final routes. The time window is the most significant difference between committed and forecast orders. Committed orders have their own time windows, as selected by customers, which are not changeable. However, forecast orders can be placed in whichever time window is most suitable because they have yet to be agreed with customers. This procedure allows the optimisation algorithm to choose an optimised delivery slot for the forecast orders, based on the location and agreed slot of all actual orders collected so far. The optimised delivery time window for these forecast orders is then used as the time window to promote when an actual arrival is seen in the same area as that of the forecast order. In summary, forecast orders serve as dummy orders without time windows that guide the booking process.

4.2.1. Insertion cost

Let $DC_t(\vec{x}_t, \vec{f}_t)$ denote the total delivery cost obtained from the dynamic routing system at booking-horizon time t , of a list of already-accepted orders, \vec{x}_t , and a set of forecast orders, \vec{f}_t that remains in the system until time t . When it is infeasible to fulfil all orders (\vec{x}_t, \vec{f}_t) with the given capacity and committed time slots for \vec{x}_t , we define $DC_t(\vec{x}_t, \vec{f}_t) = \infty$. In the interim periods, we build a series of incremental delivery-cost approximations via dynamic routing with forecast orders. In more detail, the insertion cost of having one more delivery in area a and slot s , $IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, is approximated by

$$IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) = \begin{cases} \min_{j_s \in \vec{f}_{rad}} [DC_{t+1}(\vec{x}_t + \mathbf{1}_{as}, \vec{f}_t - j_s) - DC_{t+1}(\vec{x}_t, \vec{f}_t)] & \text{if } \vec{f}_t \neq \emptyset \\ DC_{t+1}(\vec{x}_t + \mathbf{1}_{as}, \emptyset) - DC_{t+1}(\vec{x}_t, \emptyset) & \text{otherwise} \end{cases} \quad (4.2)$$

where $\vec{f}_{rad} \subseteq \vec{f}_t$ denotes all forecast orders that are not exceeding a radius equals to rad miles away from the new order, $j_s \in \vec{f}_{rad}$ indicates the forecast order to be removed while inserting the new order (a, s) , and $j_s^* \in \vec{f}_{rad}$ denotes the best forecast order which is identified to be removed.

Note that the location of the removed order might be different from that of the new order due to the forecast error and the rule of replacement (explained in detail in [Section 4.3](#)).

Note further that we did not need to re-run the VRP when computing [Eq. \(4.2\)](#). Instead, to calculate the extra driving time/cost in reaching order (a, s) and omitting the order j_s^* , we just estimated the extra driving distance and time which would be required, as a deviation from the existing route found by the VRP, to accommodate these two changes.

4.2.2. Displacement cost (revenue loss)

The insertion cost (4.2) forms one part of the opportunity-cost estimation, whereas the other part comes from the expected revenue loss by accepting order (a, s) , denoted by $RL_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$. Provided that the current route (consisting of both actual and forecast orders) is always treated as the optimal route at the end of the booking horizon, we can construct our displacement cost/revenue loss estimation using it. According to the state-of-the-art approach used in attended home delivery literature, such as [Koch & Klein \(2020\)](#), we interpret the potential revenue loss as the “additive monetary value of the time window consumption” due to the acceptance of a new order. Let $w_{s'}(\vec{x}_t, \vec{f}_t)$ denote the idle time of the current route in time slot s' ; then, after performing the replacement of forecast order j_s^* by the new order $\mathbf{1}_{as}$, the idle time is represented by $w_{s'}(\vec{x}_t + \mathbf{1}_{as}, \vec{f}_t - j_s^*)$. The revenue loss, therefore, is formulated as:

$$RL_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) = \sum_{s' \in S} \theta_{t,s'} (w_{s'}(\vec{x}_t, \vec{f}_t) - w_{s'}(\vec{x}_t + \mathbf{1}_{as}, \vec{f}_t - j_s^*)) \quad (4.3)$$

where $\theta_{t,s'} \in \mathbb{R}$ denotes the expected future revenue income per unit-time in slot s' , that is evaluated at booking horizon t . While unlike [Koch & Klein \(2020\)](#) who learn the $\theta_{t,s'}$ values through sample simulation, in this work we estimate the value of $\theta_{t,s'}$ using the current best route from the dynamic vehicle-routing solutions, as:

$$\theta_{t,s'} = \frac{\sum_i \{r_i | i \in \vec{f}_t, u_{s'-1} \leq \tau_i \leq u_{s'}\}}{u_{s'} - u_{s'-1}} \quad (4.4)$$

where τ_i indicates the delivery time of order i in the current best route, and $u_{s'}$ denotes the finishing time of slot s' . The numerator of (4.4) represents the total revenue (i.e., $\sum r_i$) for all forecast orders i which are scheduled to be delivered in slot s' ; and the denominator represents the duration of time slot s' .

The opportunity-cost estimation is then:

$$OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) = IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) + RL_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) \quad (4.5)$$

which can be used to replace the $(V_{t+1}(\vec{x}_t) - V_{t+1}(\vec{x}_t + \mathbf{1}_{as}))$ in [Eq. \(3.1\)](#), so that the DP program can be reformulated as:

$$\tilde{V}_t(\vec{x}_t) \approx \left[\max_{\vec{d}} \sum_a \lambda \mu_a \sum_{s \in F_a(\vec{x}_t, \vec{f}_t)} P_{s, F_a(\vec{x}_t, \vec{f}_t)}(\vec{d}_a) [r + d_{as} - OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})] \right] + V_{t+1}(\vec{x}_t), \forall \vec{x}_t \in \mathcal{X}, \quad (4.6)$$

with all elements known (except for $V_{t+1}(\vec{x}_t)$; but this term is not relevant in pricing optimisation) for every new order arriving at the system. This approximation decomposes the MDP into single-stage decision problems, provided that the opportunity cost $OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$ is evaluated dynamically over time. Note that the difference in (4.3) can be negative, indicating that replacing a forecast order with an actual one can help to travel less. We can interpret the negative value for RL as accepting the incoming order will not endure revenue loss and lead to higher slot availability for future customers. Therefore, the opportunity cost regarding RL , in this case, favours accepting the incoming order.

In this work, we inherit the state-of-the-art MNL model to describe customer-selection behaviour introduced by [McFadden et al. \(1973\)](#). In this scheme, the selection probability of time slot s under price d_s , $P_s(\vec{d})$, is given by:

$$P_s(\vec{d}) = \frac{\exp(\beta_0 + \beta_s + \beta_d d_s)}{\sum_{k \in F_s(\vec{x}_t, \vec{f}_t)} \exp(\beta_0 + \beta_k + \beta_d d_k) + 1}, \quad (4.7)$$

where β_0 is the base utility on all choices, β_s is the utility of slot s itself, and β_d holds the utility sensitivity to delivery charge d_s . These β values are found by numerical optimisation for the historical data of purchases made and reflect the popularity of different

times of day, and the inferred price elasticity of demand. Similarly, the probability of no-booking under the delivery charge \vec{d} , $P_0(\vec{d})$, is given by:

$$P_0(\vec{d}) = \frac{1}{\sum_{k \in F_0(\vec{x}_t, \vec{f}_t)} \exp(\beta_0 + \beta_k + \beta_d d_k) + 1} \quad (4.8)$$

As Dong et al. (2009) show, under this choice model, given $OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, the optimal solution d_s^* to the online pricing problem can be achieved for $s \in F_d(\vec{x}_t, \vec{f}_t)$ with:

$$d_s^* = OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) - r - \frac{h}{\beta_d}, \quad (4.9)$$

where h is the unique solution to:

$$(h-1)\exp(h) = \sum_{s \in F_d(\vec{x}_t, \vec{f}_t)} \exp(\beta_0 + \beta_s + \beta_d(OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) - r)) \quad (4.10)$$

In Section 4.3 we give more details about how the order to remove, j^* , can be identified, and how the order replacement is carried out in an online setting.

4.3. Insertion-cost evaluation and order replacement

Upon a new arrival in area a , the insertion feasibility and insertion cost have to be evaluated for fulfilling this order in time slot s , based on what orders have been accepted so far. As mentioned earlier, in this study, we maintain a delivery plan of all accepted orders and forecast orders dynamically and assume that at any interim stage, a “current best route” is available from the company's CVRPTW solver. As time goes by, we aim to replace forecast orders with actual incoming orders, one by one. For every potential replacement in a particular time slot s , we identify the best forecast order to remove, i.e., j_s^* , and calculate the incremental delivery cost involved in the replacement, and make it an estimate of the insertion cost for this time slot, i.e., $IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$. The methodology is summarised in Algorithm 1.

Algorithm 1 Opportunity-cost estimation for a potential new order i in area a .

- 1: Compute set \vec{f}_{rad} (set of candidate forecast orders, within radius rad , to be replaced by the new order) for new order i in area a .
- 2: **for** slot $s \in S$ **do**
- 3: Compute insertion cost, $IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, according to (4.2), denote the best order to remove as j_s^* .
- 4: **if** $IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) \neq \infty$, i.e., the insertion into slot s is feasible **then**
- 5: Record the best forecast order to remove, j_s^* , for later use in Alg. 2
- 6: Calculate the potential revenue loss of accepting one more order for slot s , $RL_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, according to (4.3)
- 7: Calculate opportunity cost for slot s , $OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, according to (4.5)
- 8: **else**
- 9: Set slot s as unavailable
- 10: **end if**
- 11: **end for**
- 12: Solve (4.6) with the opportunity costs $OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, and display available slots with the optimised pricing 4.9 to customers.

Algorithm 1 is called when we need to calculate the opportunity cost approximation to display available slots with the optimised pricing to customers. Let us consider a customer request i from area a to book a delivery slot. For this request, we can easily identify a subset of neighbouring forecast orders to replace by

measuring their distance to the customer's location. To cover different implementation scenarios, we can apply different rules to calculate the subset, such as radius, road distance and/or post-code sectors. Let us denote the subset of forecast orders to remove by \vec{f}_{rad} . In this study, to find \vec{f}_{rad} , the algorithm will check available forecast orders within radius rad . If there are no forecast orders within radius rad , we increase the radius gradually with pre-set radius bands until forecast orders are found. The algorithm's performance is affected by the largest possible radius considered in implementation. If it is too small, the availability of slots might be restricted; if it is too large, the delivery cost prediction will require significant processing time to explore all forecast orders in the range. The trade-off between processing time and forecast accuracy is presented in Section 5.5 with numerical results.

When all forecast orders are removed already, i.e., $\vec{f}_{rad} = \emptyset$, the algorithm will check the feasibility of inserting the new order without replacement. This is the same as the typical insertion heuristics carried out by, for example, Campbell & Savelsbergh (2006). This helps mitigate forecasting errors. Indeed, forecast orders are particularly important at the start of the booking horizon in guiding orders to suitable time slots. They become less important towards the end when the accepted orders with time windows nearly fix the routing plan. When the number of forecast orders is higher, we remove all remaining forecast orders in the end. Section 5.4 focuses on the forecast order levels and how the number of forecast orders may affect the performance.

After the customer's selection of their preferred slot, Algorithm 2 will perform the actual replacement/insertion of

Algorithm 2 Order replacement/insertion upon customer selection.

- 1: Denote the slot which the customer has selected by \hat{s} ; the new order being made in area a by i ; and the forecast order to be replaced by $j_{\hat{s}}^*$
- 2: **if** $j_{\hat{s}}^* \neq \emptyset$ **then**
- 3: Remove $j_{\hat{s}}^*$ from the route
- 4: Update $\vec{f}_t \leftarrow \vec{f}_t - j_{\hat{s}}^*$
- 5: **end if**
- 6: Insert order i into slot \hat{s} of the route, Update $\vec{x}_t \leftarrow \vec{x}_t + \mathbf{1}_{a\hat{s}}$
- 7: Re-optimize the route for order set (\vec{x}_t, \vec{f}_t)

the new order i in the selected slot \hat{s} . Note that the resulting approach is more robust if traffic conditions, schedule tightness, potential lateness, etc., are considered in detail in the dynamic routing system. Upon the acceptance of a new order, the CVRPTW is re-optimised/updated for a better delivery schedule until the next arrival comes into the system. This means that the routing optimisation could adjust the delivery time of the remaining forecast orders in the system to re-optimize the route.

5. Numerical results

In order to investigate how our approach performs in practice, we test the methodology above on four typical delivery areas, each with different customer densities and spread patterns. Real customer locations and historical booking data are deployed in the tests, with essential manipulations to protect commercial information and customer privacy.

5.1. Routing package

As explained in Section 4.2, the designed approach can be implemented with any routing packages for CVRPTW. This includes dynamic-routing packages with automatic updating schemes based

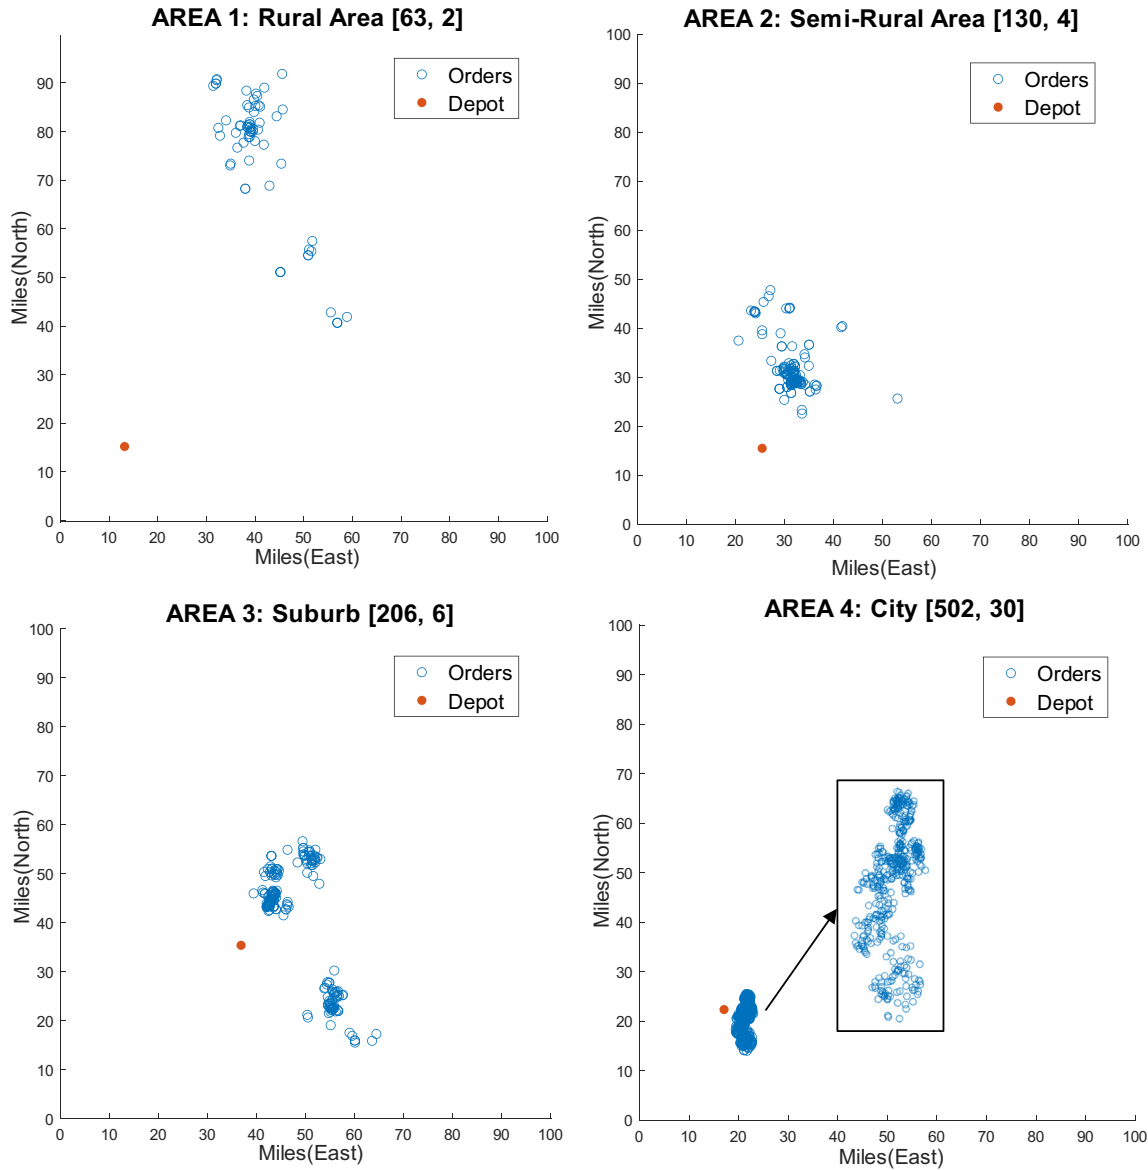


Fig. 3. Different areas with different densities and spread patterns of orders. A detailed number of forecast orders and vehicles for each area are given as [#forecast orders, #vehicles] in the subtitles.

on the acceptance of new orders and static routing packages that allow warm-starting from known solutions that are obtained from replacing/inserting the new orders in the current best routes. Such routing packages are understood to be standard tools currently used by delivery companies. In our experimental tests, however, we do not rely on any company's specialised routing package but use a generic meta-heuristic, i.e., Simulated Annealing (SA), to solve the CVRPTW. The application of SA to the VRPTW was introduced and proven effective in terms of accuracy and execution time by [Chiang et al. \(1996\)](#).

Implementing the SA in dynamic routing can be characterised by an offline phase and an online phase. The offline phase is a Capacitated Vehicle Routing Problem (CVRP) without time windows that occur entirely before the booking horizon, which consists of only forecast orders. The best possible route for the forecast orders (and their best time slots) are found using SA, preliminary to the order acceptance process. The online phase covers the entire booking horizon when actual orders are collected and forecast orders are replaced as time goes by, according to [Algorithms 1 and 2](#). SA

is called to re-optimize the current route after each replacement is done, until the arrival of the following order.

5.2. Experiment settings

We test our model on four typical area settings to investigate different scenarios regarding the spread of orders. These areas are a Rural area, which reflects the countryside and small villages; a Semi-rural area, which represents towns; a Suburb area, which represents the outskirts of big cities where people live in disjoint but not far-way satellite communities; and a City area, where people live with the highest density, e.g., in apartments. Each area has a depot, whose location/distance to the primary service area is set to capture the business settings. See [Fig. 3](#).

As what has been used by [Yang et al. \(2016\)](#), we fit a non-homogeneous Poisson distribution to historical data and use it to generate customer arrival times in simulation. Order details, including customer addresses and order sizes, are randomly simulated from real orders in the past. We also borrow the time slot

Table 1

Results of SP, DP-IC, DP-FR-F, DP-DR-TWF and DP-DR-F methods.

		SP	DP-IC	DP-FR-F	DP-DR-TWF	DP-DR-F
Rural	Total mileage	578.93(27.61)	575.65(22.84)	577.45(23.27)	562.46(22.10)	543.08(24.49)
	Mileage/Order	9.48(1.08)	7.70(0.89)	7.64(0.68)	7.84(0.84)	6.75(0.49)
	Number of orders	60.60(5.25)	74.40(5.97)	75.00(4.69)	71.37(6.41)	79.66(3.71)
	Total order size	213.38(21.20)	273.07(23.23)	273.35(19.90)	267.38(23.38)	296.60(15.98)
	Accepted price	3.00	-4.13(0.65)	-3.86(0.54)	-1.47(0.34)	-3.15(0.63)
	Total Profit	2109.36(206.98)	2168.66(189.12)	2187.86(157.03)	2226.23(208.64)	2433.46(133.50)
	Improvement on SP		2.81%	3.72%	5.54%	15.36%
Semi-rural	Total mileage	459.41(18.35)	446.16(17.39)	450.14(17.61)	419.76(21.10)	415.79(16.38)
	Mileage/Order	3.66(0.36)	3.09(0.26)	3.09(0.22)	2.94(0.26)	2.67(0.14)
	Number of orders	125.43(8.83)	144.40(9.84)	145.03(7.04)	142.33(7.89)	152.00(6.22)
	Total order size	448.47(30.41)	539.69(34.84)	545.63(28.63)	545.02(27.27)	583.50(20.29)
	Accepted price	3.00	-3.79(0.36)	-3.41(0.39)	-0.90(0.13)	-1.11(0.55)
	Total Profit	4431.28(300.94)	4341.81(270.34)	4447.90(231.75)	4685.57(239.68)	5108.86(171.81)
	Improvement on SP		-2.02%	0.38%	5.74%	15.29%
Suburb	Total mileage	640.77(24.33)	622.74(20.18)	622.67(23.20)	613.99(20.11)	577.59(25.82)
	Mileage/Order	3.18(0.35)	2.65(0.17)	2.65(0.18)	2.88(0.18)	2.35(0.11)
	Number of orders	202.03(15.65)	234.73(9.05)	234.20(8.90)	212.97(9.16)	245.27(5.03)
	Total order size	736.01(57.25)	891.57(32.52)	893.51(32.36)	817.10(36.21)	927.40(20.93)
	Accepted price	3.00	-3.80(0.26)	-3.50(0.30)	-0.00(0.34)	-0.59(0.50)
	Total Profit	7263.83(564.16)	7182.95(291.06)	7273.36(280.74)	7396.60(358.20)	8249.52(230.96)
	Improvement on SP		-1.11%	0.13%	1.83%	13.57%
City	Total mileage	453.07(7.63)	449.66(7.74)	449.38(8.63)	432.16(11.60)	407.63(10.15)
	Mileage/Order	0.99(0.04)	0.89(0.04)	0.89(0.03)	0.95(0.04)	0.76(0.03)
	Number of orders	458.43(14.20)	505.07(14.65)	504.43(13.94)	455.33(13.53)	531.27(13.82)
	Total order size	1702.06(57.96)	1966.86(57.56)	1966.51(58.29)	1794.86(72.67)	2060.32(58.49)
	Accepted price	3.00	-3.86(0.21)	-3.45(0.20)	1.53(0.16)	3.25(0.60)
	Total Profit	16777.67(562.93)	15858.71(435.56)	16066.44(463.20)	17676.47(667.65)	20372.63(355.69)
	Improvement on SP		-5.48%	-4.24%	5.36%	21.43%

and MNL parameters from Yang et al. (2016), i.e., 27 slots per day, one hour each; some overlap with others. Please refer to (4.7) and (4.8) for the detailed MNL customer-choice model parameters. Continuous slot prices, in the range £[-10, 10], are considered in the simulation (i.e., these are the slot prices offered to customers), with negative prices indicating discounts offered to customers as an incentive to purchase. This pricing scheme has been approved by the commercial partner we are collaborating with for this project. Detailed revenue/profit and order size information cannot be published due to commercial concerns, but for a meaningful interpretation of the results, the ratio between order revenue and variable delivery cost is set to 40.5 to reflect the real situation.

5.3. Experiment results

We aim to justify the effectiveness of the proposed dynamic-pricing approach by maintaining a set of forecast orders without time windows in the routing system. The comparisons are therefore carried out between a system using forecast orders without time windows, i.e., Dynamic Pricing with Dynamic Routing of Forecast orders without time windows (DP-DR-F), a system with time-windowed forecast orders obtained from historical routes, i.e., Dynamic Pricing with Dynamic Routing of Time-Windowed Forecast orders (DP-DR-TWF), and a system without any forecast orders, i.e., Dynamic Pricing Insertion Cost (DP-IC). All other decisive factors are the same for the tests, including:

- The same routing package is used for each paired test, i.e., the Simulated Annealing (SA) for CVRPTW as described in Section 5.1 with the same tunable parameters;
- All implement dynamic routing where updating (re-optimisation) of the current best route is performed after the acceptance of every new order until the next order arrives;
- All deploy the same MNL customer-choice model estimated from real data for customer selections;

- All deploy the same approach for insertion cost and revenue loss (where forecast orders exist) estimation according to Sections 4.2 and 4.3. The only difference is the forecast routes (or whether there is a forecast route) used.

On top of this, we also benchmark our approach with the commonly-used Static Pricing (SP), i.e., where every feasible slot for the entire booking horizon has a fixed price of £3. Also with a fixed forecast-route approach, i.e., Dynamic Pricing Fixed Routing Forecast (DP-FR-F), which runs dynamic pricing with time-windowed forecast orders/routes obtained from historical routes (the "Foresight" policy of Yang et al., 2016) but not having the forecast route updated as new order information comes in. To have a fair test on the performance of the pricing policy alone, the feasibility of placing an order in a slot under these two benchmark policies is also informed by the dynamic-routing package.

Experiments are carried out using MATLAB on an Intel Core i9-7940X 3.1 gigahertz machine. Since we deployed a meta-heuristic approach to solve the CVRPTW, we conducted 30 independent runs. We reported the average and the standard deviation (mean(s.d.)) to minimise the influence of the randomness involved in the solution approach. Performance on crucial indicators for profit and efficiency are presented in Table 1, with the best one in every row highlighted in bold. The results show that the DP-DR-F outperforms the other approaches in all the measurements, which confirms the effectiveness of the proposed approach.

One important observation here is that dynamic pricing (DP-IC) is not necessarily better than static pricing (SP), mainly when a poor estimation of the opportunity cost is used. Furthermore, the short-sighted incremental delivery cost based on accepted orders alone is insufficient for opportunity-cost estimation. This insight is in line with the conclusions of Yang & Strauss (2017), which emphasise the importance of "incorporating the impact of future profit opportunities from orders".

As this work proposes, incorporating forecast orders (without time windows) provides an easy method for future-profit estima-

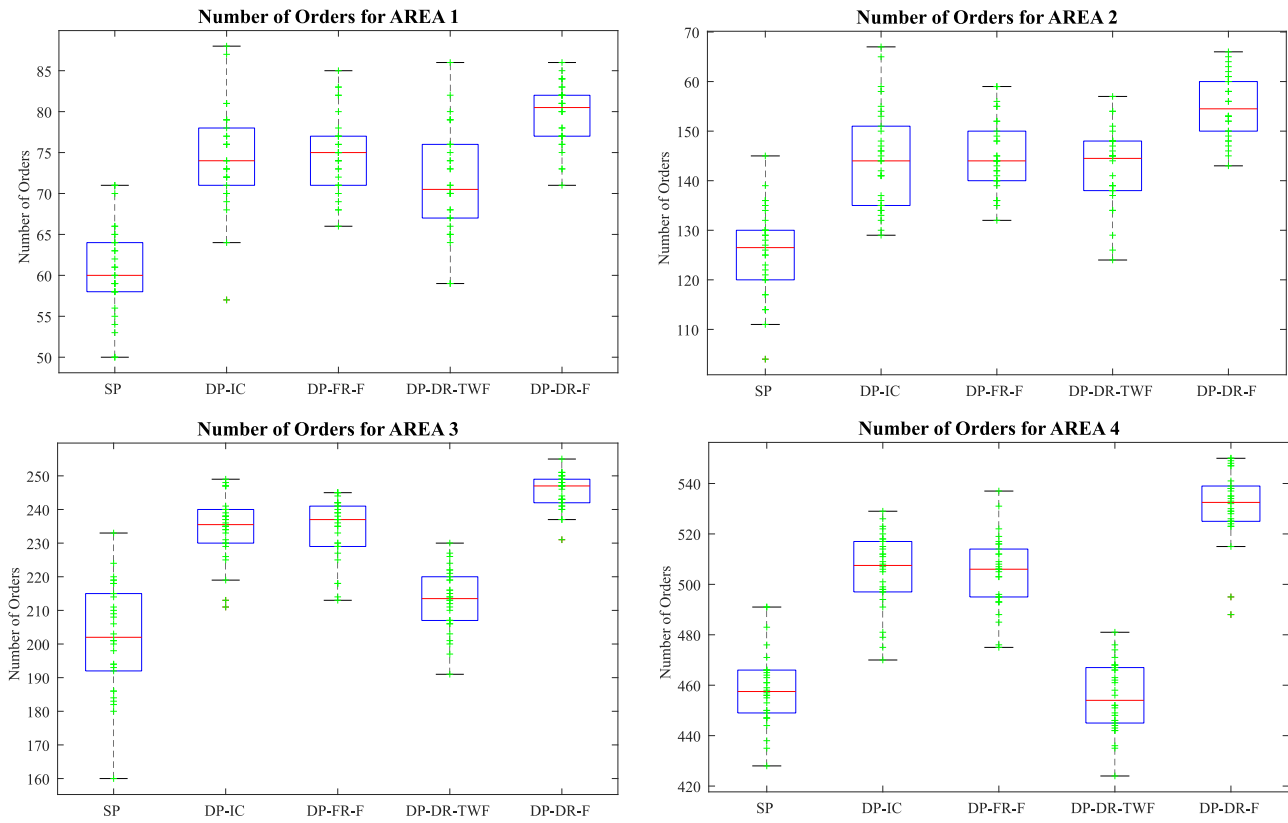


Fig. 4. Total number of orders accepted by the methods.

tion. Together with the better marginal delivery-cost estimation, obtained by maintaining a hybrid route with both forecast and actual orders, the approach achieves a 13.57–21.43% profit increase over the static-pricing method, which is better than that of 2.2–2.5% in Yang & Strauss (2017) and that of 2.6–6.2% in Yang et al. (2016). We also performed paired sample t-tests on the total profits for all the approaches in the studied areas. The p-values produced were all less than 0.05, indicating that the differences were statistically significant at a 5% significance level.

We look closely at key-related components to understand where the additional profits come from in the DP-DR-F approach. Figures 4 and 5 show a graphical comparison of the number of order commitments and the average travelling distance per order across all approaches. These elements demonstrate the efficiency of the final routes we end up with using the DP-DR-F approach so as to justify the capability of DP-DR-F in recognising the best time slots to offer and promoting them via dynamic pricing.

Compared to DP-DR-TWF, which uses time-windowed forecast orders, the DP-DR-F approach omits the time windows from the forecast orders, which thus allows the dynamic route planning to adjust the delivery time of a forecast order to the best possible time slot according to its location and fitness to the current best route. This feature provides extra flexibility for dynamic route planning to create more feasible slots over the booking horizon (more information to follow in Fig. 7) and guide the overall booking process to a more compact route in the end. Higher profits are therefore achieved through selling more goods. These are believed as the key reasons for higher order commitments provided by DP-DR-F with the fixed fleet and time window capacity.

Figure 6 shows samples of the final routes obtained by the five approaches in Area 1. Comparing the plots in Fig. 6, we can see that the DP-DR-F approach gives the most efficient routes, with noticeably fewer long links than the others. It performs especially

well in the “remote area”, which directs the orders in this area to adjacent slots, so as to avoid the van coming back to this area multiple times to meet demands at different times of the day. The improved route plan shows DP-DR-F’s ability to promote the correct time slot that complies with the optimal route to reduce unnecessary travel to meet customer needs.

Figure 7 is related to how many slots are available on average as the booking process progresses for all methods. Based on this plot, we can see that all methods begin with a high slot availability, which decreases as time passes. Decreasing rates for all the other four approaches are similar, whereas, for DP-DR-F, the slope magnitude is lower. This outcome justifies that DP-DR-F works well in reserving resources for later usage to provide more stable slot availability over the entire booking horizon than the benchmark approaches. The higher availability leads to a higher selection rate on average and therefore conveys more orders. In addition, by applying DP-DR-F, a significantly higher number of slots are still available when the booking horizon reaches its end. However, this approach has committed a notably higher number of orders than others. This upshot shows further profit growth potentials of this approach suppose a more extensive market can be reached by the firm, which is not possible with any other approaches.

Another critical term influencing the total profit is slot price, which is the fee customers pay for the delivery service. Figure 8 shows the average slot prices offered to the time slots that customers eventually select with every approach. It is not hard to see that the DP-IC approach outperforms the static approach in the number of accepted orders (Fig. 4) and the per-order delivery costs (Fig. 5). However, there is no significant improvement in the overall profit due to the low average price it charges.

Upon the arrival of a new order, there is a trade-off between long-term profit and immediate gain. This balance explains why the prices offered by the DP-IC and DP-FR-F approaches are con-

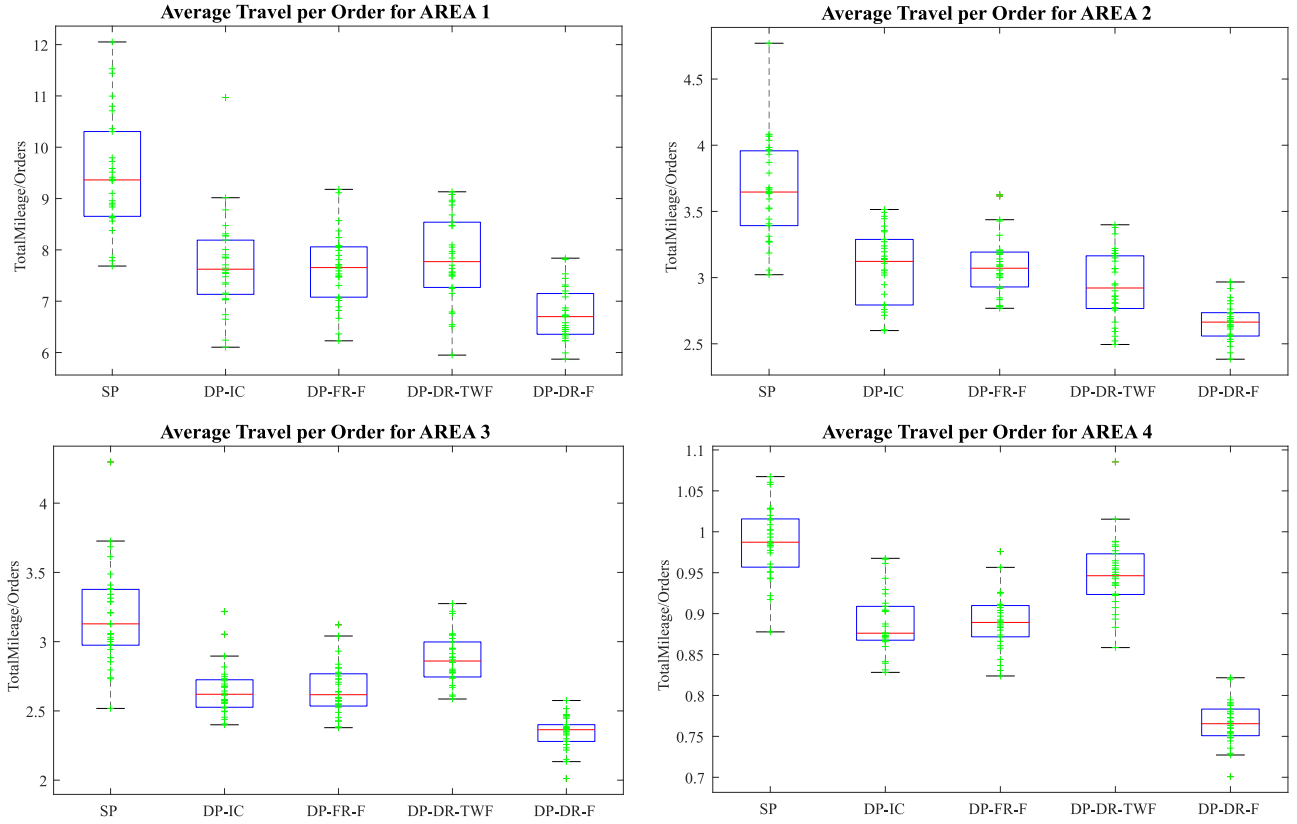


Fig. 5. Average travel distance to fulfil an order.

stantly lower than those of DP-DR-TWF and DP-DR-F. Without considering the expected revenue loss in the opportunity-cost estimation, the DP-IC and DP-FR-F approaches only focus on the immediate gain brought by the order under consideration and try very hard to persuade this customer to buy by lowering the price offered. However, with the DP-DR-TWF and DP-DR-F approach, as they still expect more future orders, the pressure of conveying an order now is lighter, so the price offered is higher on average. Comparing DP-DR-TWF and DP-DR-F, however, the prices charged by DP-DR-F are slightly lower in most cases. This is because, in DP-DR-F, the best time slot is purely identified through the dynamic routing system without time windows, which takes no consideration of the original popularity of time slots. This puts higher pressure on lowering the price to persuade a customer to book an undesirable time slot if their location is deemed the best fit for that time slot.

Note that the average prices offered by dynamic approaches are, in most cases, lower than zero. The finding is consistent with the previous study using the same choice model, i.e., Yang et al. (2016) and Yang & Strauss (2017). As claimed in both previous works, the profit from selling an order is much higher than the profit from making a delivery. The system offers discounts to encourage customers to buy, rather than highly charging them for delivery.

Figure 9 displays the offered prices change over time on a sample run in Area 1. Prices offered by all dynamic pricing schemes are decreasing over time. Such a trend is understandable, as when slot availability decreases, the pricing problem expects a higher no-booking rate, so it tends to lower the price to persuade customers to buy. Also, as explained in Yang et al. (2016), popular slots are filled in earlier than unpopular ones in the booking horizon, so lower prices must be offered further to promote unpopular slots towards the end of booking time.

5.4. Analysis of the number of forecast orders

As explained in Section 4, we use a moving average to estimate the number of forecast orders in each area, which may lead to, sometimes significant, forecast errors. This subsection investigates how sensitive the DP-DR-F approach is to forecasting errors. To this aim, we create scenarios where the number of forecast orders is significantly (20%) higher or lower than the moving average and test the DP-DR-F on them. Table 2 shows the obtained results for 30 independent runs in the format of the average and the standard deviation (mean(s.d.)) for the four studied areas, based on the same set of random examples that have been used in Table 1 to make results comparable across tables.

Concerning the results obtained from the simulation in Table 2, we can infer that if the moving average estimate (4.1) decides the number of forecast orders, we can expect higher performance in total profit. However, if we underestimate or overestimate the number of forecast orders compared to what the moving average method suggests (in our study, 20% lower or higher than the moving average estimate), the performance undergoes a decrease in efficiency. We can notice this point in all studied areas, which affirms the robustness of the proposed method to the number of order estimations. Moreover, the results for a higher number of forecast orders are slightly better than those for a lower number of forecast orders. This fact emphasises the critical role that forecast orders play in the booking process, even towards the end when actual orders with time windows become the majority of the group and the routes are relatively fixed.

5.5. Impact of radius size on performance and run-time

The proposed approach is online, so the time it takes to find feasible time slots and optimise their prices is crucial for

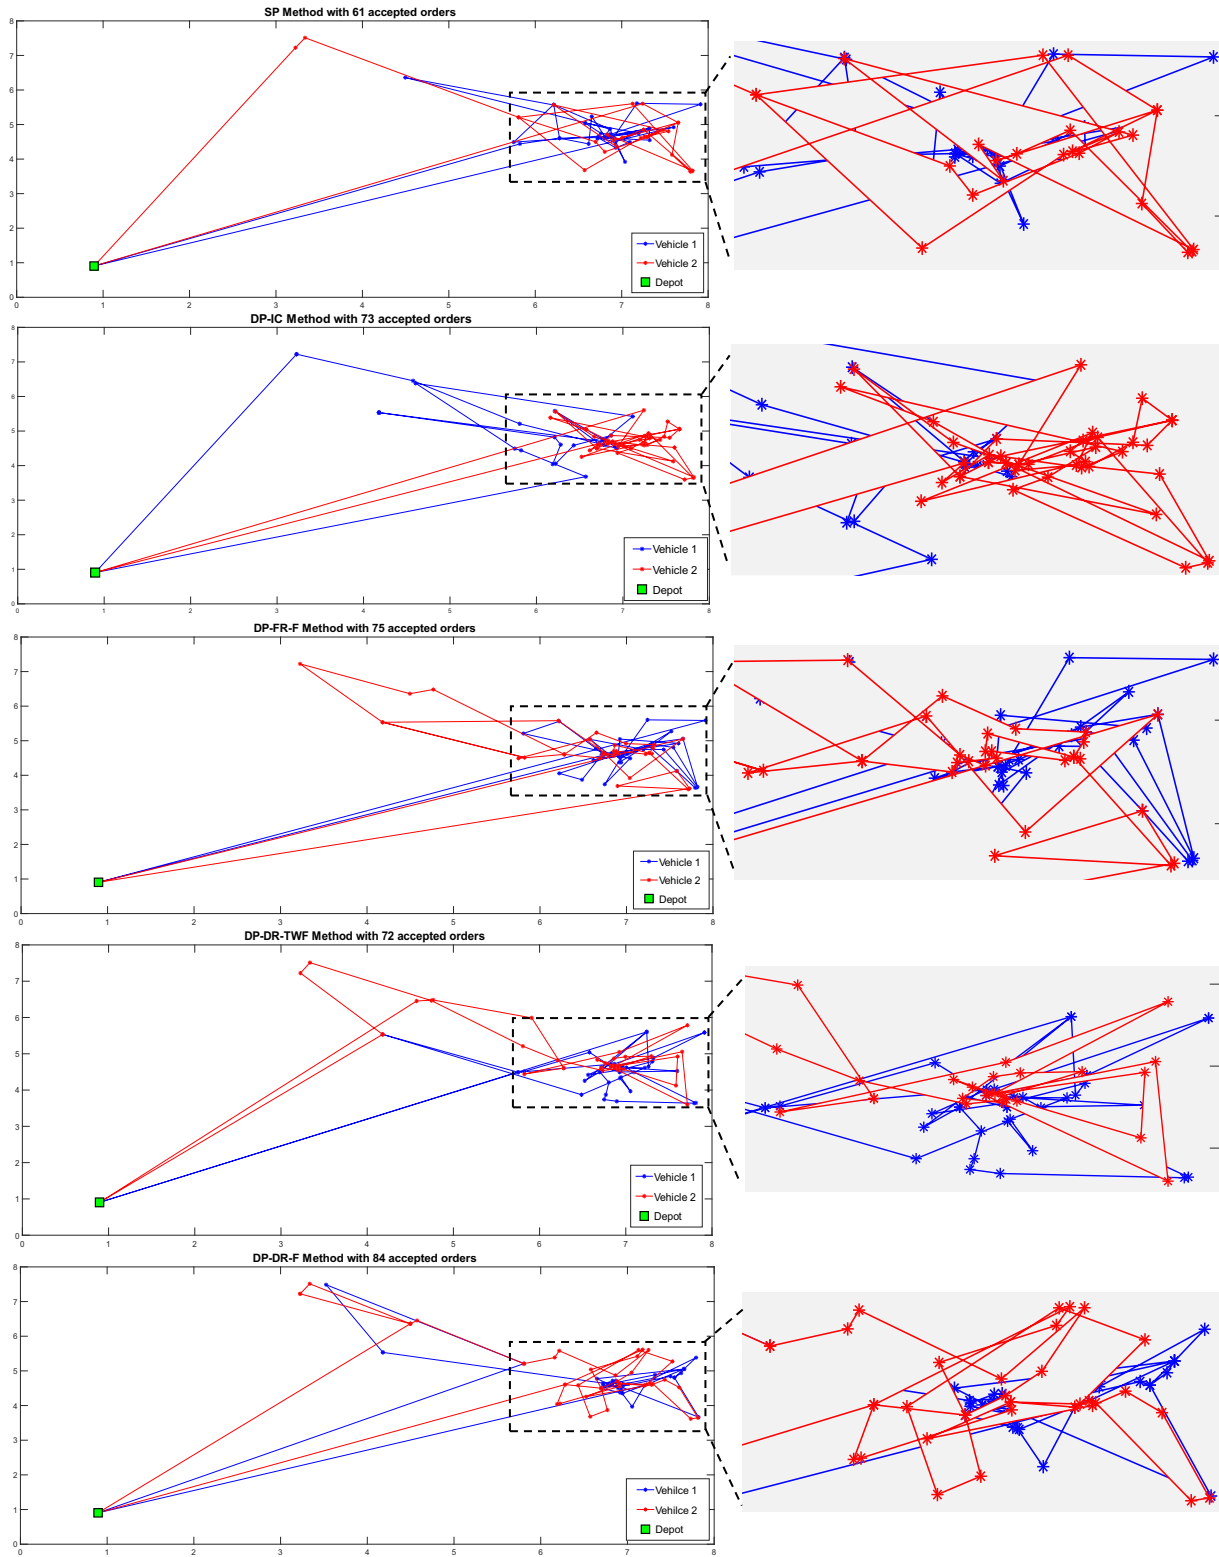


Fig. 6. Final route for Area 1 using SP, DP-IC, DP-FR-F, DP-DR-TWF and DP-DR-F methods.

successful implantation. One way of accelerating the decision process is limiting the number of forecast orders to replace when evaluating the insertion cost. This subsection presents how various forecast order radii can affect the proposed method's functionality and run-time. More precisely, when the DP-DR-F method wants to find the forecast orders for replacement with a new arrival, only

the forecast orders located within a specific radius of the new arrival will be considered. The greater this radius is, the higher the number of neighbouring forecast orders will be, resulting in more computation time.

In Table 3, we define experiments with different radii based on order density, i.e., as a ratio to the average distance between or-




Fig. 7. Comparison of slots availability over time for Area 1.

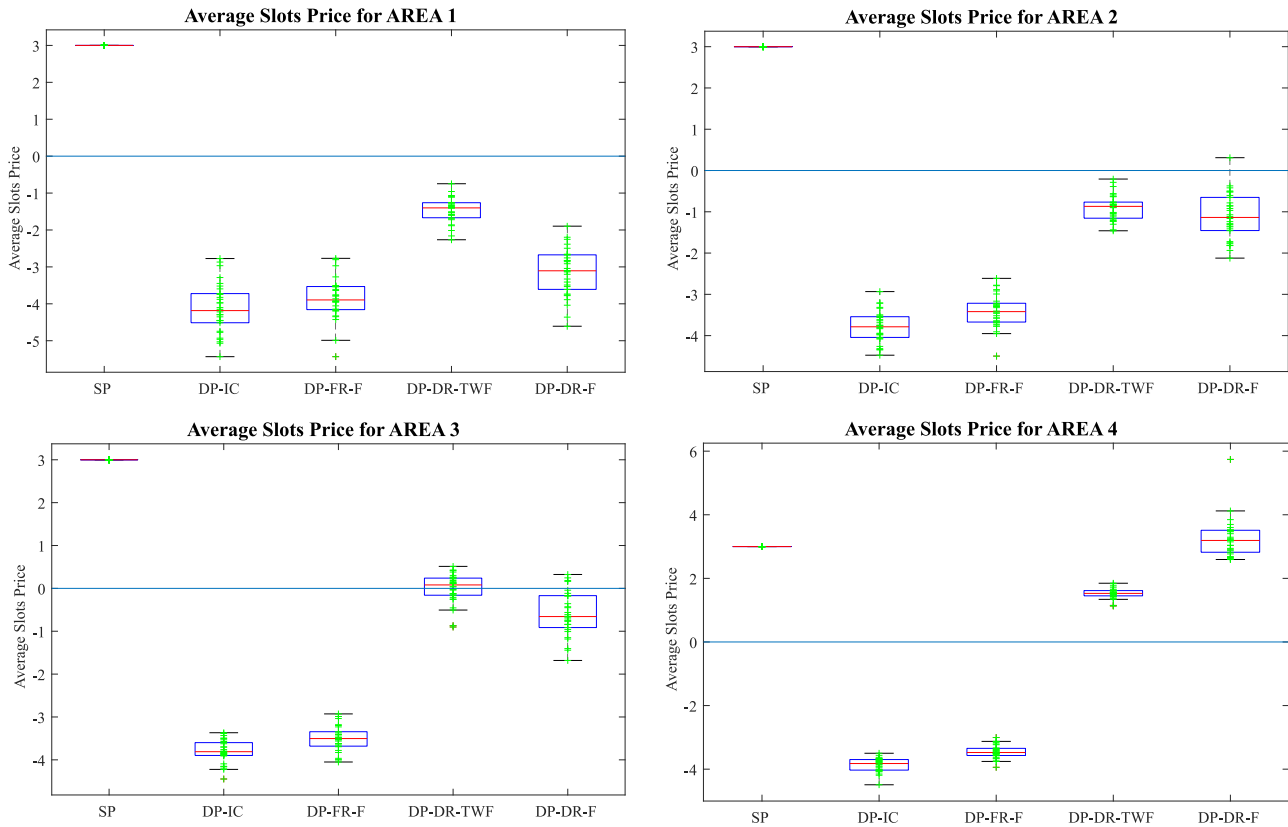


Fig. 8. Average prices for the slots that are booked by customers.

ders (ADO). When the method tries to find the candidate forecast orders, all the forecast orders within a specific radius will be explored, and replacement feasibility will be conducted for each of them. If no forecast orders exist in this radius, the method will double the size of the radius to search for forecast orders. This process will continue until at least one forecast order is found.

The average running time to obtain a list of feasible slots, with their optimised prices, is reported in the last column of Table 3. This measure can be seen as the average online reaction time upon a customer's arrival. According to this table, there is an increasing trend in performance (e.g., profit) when we enlarge the area searched for a replacement. At the same time, the execution time



Fig. 9. Average slots' price offered over time for Area 1.

Table 2

Comparison of different forecast order levels for the 4 studied areas.

		80% Forecast	100% Forecast	120% Forecast
Rural	Total mileage	548.04(24.81)	543.08(24.49)	522.15(21.05)
	Mileage/Order	6.94(0.53)	6.75(0.49)	6.82(0.61)
	Number of orders	78.23(4.34)	79.67(3.71)	76.03(5.90)
	Total order size	288.49(17.93)	296.60(15.98)	282.80(19.53)
	Accepted price	-3.26(0.64)	-3.15(0.63)	-1.70(0.88)
Semi-rural	Total Profit	2356.21(155.07)	2433.47(133.50)	2428.21(151.38)
	Total mileage	417.00(18.27)	415.79(16.38)	375.37(14.96)
	Mileage/Order	2.70(0.17)	2.67(0.14)	2.59(0.23)
	Number of orders	153.77(5.73)	155.00(6.22)	144.67(12.06)
	Total order size	577.47(20.33)	583.50(20.29)	540.01(47.61)
Suburb	Accepted price	-2.22(0.49)	-1.11(0.55)	1.10(1.03)
	Total Profit	4885.20(170.11)	5108.86(171.81)	5040.52(382.49)
	Total mileage	585.34(26.56)	577.59(25.82)	554.94(23.39)
	Mileage/Order	2.40(0.18)	2.35(0.11)	2.45(0.15)
	Number of orders	243.43(9.24)	245.27(5.03)	226.13(9.81)
City	Total order size	921.44(33.65)	927.40(20.93)	847.16(38.49)
	Accepted price	-1.98(0.46)	-0.60(0.50)	2.26(0.74)
	Total Profit	7858.11(276.22)	8249.52(230.96)	8175.12(322.48)
	Total mileage	413.09(10.44)	407.63(10.15)	389.75(8.86)
	Mileage/Order	0.80(0.03)	0.77(0.03)	0.83(0.04)
	Number of orders	516.10(12.05)	531.27(13.82)	469.03(19.45)
	Total order size	2017.98(41.48)	2060.32(58.45)	1793.04(86.05)
	Accepted price	1.85(0.77)	3.35(0.60)	6.58(0.45)
	Total Profit	19229.16(391.95)	20372.63(355.70)	19315.16(799.10)

increases more sharply, which is in line with our expectations. Compared to the largest possible radius, the reduced search range, e.g., to $0.01 \times ADO$, only sacrifices less than 0.5% of the total profit while reducing the average reaction time by 95.38% compared to the full search. The obtained results justify the effectiveness of the proposed simplification. In practice, the online grocer can choose a radius according to the maximum affordable run-time to have the best achievable results within the time limit.

6. Conclusion and future work

This paper introduces a novel dynamic pricing method for attended home delivery using forecast orders without time windows. The approach maintains a dynamic route of actual orders (with

time windows) and forecast orders (without time windows). It estimates opportunity costs online using the most up-to-date information in the dynamic route. No extra learning is needed. The approach can be easily integrated with any dynamic-routing package a company is using, which allows the company's specific routing needs and restrictions to be considered as well. The approach is tested on real data with an MNL customer-choice model.

One limitation of this study is that we presume dynamic pricing is powerful enough to balance customer demands so all uncommitted orders can be moved freely across time slots. While in practice demand is also restricted by customer availability. In future research, we will consider bringing the original slot popularity into consideration while planning the forecast route of orders without time windows, so as to maintain the benefit of both approaches.

Table 3

The effect of different radii on performance and run-time of DP-DR-F method in Area 3. ADO stands for average distance between orders.

Radius	Total mileage	Total order size	Number of orders	Total profit	Run-time (second)
0.01 × ADO	580.80	891.90	234.37	8214.05	0.0041
0.10 × ADO	577.59	909.20	240.43	8214.47	0.0104
0.50 × ADO	581.34	914.99	241.50	8160.30	0.0336
1.00 × ADO	573.67	922.15	244.40	8240.20	0.0519
∞	579.47	923.15	244.90	8251.86	0.0887

By conducting experiments on four different areas, the advantages of employing forecast orders without time windows are observed in higher-order commitments, lower delivery costs and higher overall profits compared to all benchmarking approaches. The improvement of 13.57–21.43% profit on Static Pricing is significantly better than the results from former approaches in Yang & Strauss (2017) and Yang et al. (2016) and of an amount likely to be of commercial interest to those managing AHD operations. The robustness of the DP-DR-F approach is also justified through experiments when the number of forecast orders is overestimated or underestimated. Potential accelerations of the approach via restricting the radius of exploration are discussed and tested in the end to improve running efficiency, and make the approach more suitable for online implementations.

References

- Agatz, N., Campbell, A., Fleischmann, M., Savelsbergh, M., et al., (2011). Time slot management in attended home delivery. *Transportation Science*, 45(3), 435–449. <https://doi.org/10.1287/trsc.1100.0346>.
- Agatz, N., Campbell, A. M., Fleischmann, M., Nunen, J. V., Savelsbergh, M., van Nunen, J., Savelsbergh, M., et al., (2013). Revenue management opportunities for internet retailers. *Journal of Revenue and Pricing Management*, 12(2), 1–19. <https://doi.org/10.1057/jrpm.2012.51>.
- Agatz, N., Fan, Y., & Stam, D. (2021). The impact of green labels on time slot choice and operational sustainability. *Production and Operations Management*, 30(7), 2285–2303. <https://doi.org/10.1111/poms.13368>.
- Agatz, N. A. H., Fleischmann, M., van Nunen, J. A. E. E., et al., (2008). E-fulfillment and multi-channel distribution: A review. *European Journal of Operational Research*, 187(2), 339–356. <https://doi.org/10.1016/j.ejor.2007.04.024>.
- Asdemir, K., Jacob, V. S., & Krishnan, R. (2009). Dynamic pricing of multiple home delivery options. *European Journal of Operational Research*, 196(1), 246–257.
- Bent, R. W., & Van Hentenryck, P. (2004). Scenario-based planning for partially dynamic vehicle routing with stochastic customers. *Operations Research*, 52(6), 977–987.
- Bühler, D., Klein, R., Neugebauer, M., et al., (2016). Model-based delivery cost approximation in attended home services. *Computers and Industrial Engineering*, 98, 78–90. <https://doi.org/10.1016/j.cie.2016.05.014>.
- Campbell, A. M., & Savelsbergh, M. (2006). Incentive schemes for attended home delivery services. *Transportation Science*, 40(3), 327–341.
- Campbell, A. M., & Savelsbergh, M. W. P. (2005). Decision support for consumer direct grocery initiatives. *Transportation Science*, 39(3), 313–327.
- Chiang, Russell, W.-C., & Robert, A. (1996). Simulated annealing metaheuristics for the vehicle routing problem with time windows. *Annals of Operations Research*, 63(2), 3–27. <https://doi.org/10.1007/BF02601637>.
- Cleophas, C., & Ehmke, J. F. (2014). When are deliveries profitable? Considering order value and transport capacity in demand fulfillment for last-mile deliveries in metropolitan areas. *Business and Information Systems Engineering*, 6(3), 153–163. <https://doi.org/10.1007/s12599-014-0321-9>.
- Dayarian, I., & Savelsbergh, M. (2020). Crowdsourcing and same-day delivery: Employing in-store customers to deliver online orders. *Production and Operations Management*, 29(9), 2153–2174. <https://doi.org/10.1111/poms.13219>.
- Dong, L., Kouvelis, P., & Tian, Z. (2009). Dynamic pricing and inventory control of substitute products. *Manufacturing and Service Operations Management*, 11(2), 317–339.
- Ehmke, J. F., & Campbell, A. M. (2014). Customer acceptance mechanisms for home deliveries in metropolitan areas. *European Journal of Operational Research*, 233(1), 193–207. <https://doi.org/10.1016/j.ejor.2013.08.028>.
- Figliozzi, M., Mahmassani, H., & Jaillet, P. (2007). Pricing in dynamic vehicle routing problems. *Transportation Science*, 41, 302–318. <https://doi.org/10.1287/trsc.1070.0193>.
- Fleckenstein, D., Klein, R., & Steinhardt, C. (2022). Recent advances in integrating demand management and vehicle routing: A methodological review. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2022.04.032>. URL: <https://www.sciencedirect.com/science/article/pii/S0377217222003460>
- Ichoua, S., Gendreau, M., & Potvin, J.-Y. (2006). Exploiting knowledge about future demands for real-time vehicle dispatching. *Transportation Science*, 40(2), 211–225.
- Klein, R., Koch, S., Steinhardt, C., & Strauss, A. K. (2020). A review of revenue management: Recent generalizations and advances in industry applications. *European Journal of Operational Research*, 284(2), 397–412. <https://doi.org/10.1016/j.ejor.2019.06.034>. URL: <https://www.sciencedirect.com/science/article/pii/S037721719305272>
- Klein, R., Mackert, J., Neugebauer, M., & Steinhardt, C. (2018). A model-based approximation of opportunity cost for dynamic pricing in attended home delivery. *OR Spectrum: Quantitative Approaches in Management*, 40(4), 969–996. URL: https://EconPapers.repec.org/RePEc:spr:orspec:v:40:y:2018:i:4:d:10.1007_s00291-017-0501-3
- Klein, R., Neugebauer, M., Ratkovitch, D., & Steinhardt, C. (2017). Differentiated time slot pricing under routing considerations in attended home delivery. *Transportation Science*, 53. <https://doi.org/10.1287/trsc.2017.0738>.
- Klein, V., & Steinhardt, C. (2022). Dynamic demand management and online tour planning for same-day delivery. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2022.09.011>. URL: <https://www.sciencedirect.com/science/article/pii/S0377217222007238>
- Koch, S., & Klein, R. (2020). Route-based approximate dynamic programming for dynamic pricing in attended home delivery. *European Journal of Operational Research*, 287(2), 633–652. <https://doi.org/10.1016/j.ejor.2020.04.002>. URL: <https://www.sciencedirect.com/science/article/pii/S037721720303209>
- Köhler, C., Ehmke, J. F., & Campbell, A. M. (2020). Flexible time window management for attended home deliveries. *Omega*, 91, 102023.
- Kumar, S. N., & Panneerselvam, R. (2012). A survey on the vehicle routing problem and its variants.
- Mackert, J., Steinhardt, C., & Klein, R. (2019). Integrating customer choice in differentiated slotting for last-mile logistics. *Logistics Research*, 12. https://doi.org/10.23773/2019_5. URL: <https://www.bvl.de/lore/all-volumes--issues/volume-12/issue-1/integrating-customer-choice-in-differentiated-slotting-for-last-mile-logistics>
- McFadden, D. et al. (1973). Conditional logit analysis of qualitative choice behavior.
- Ojeda Rios, B. H., Xavier, E. C., Miyazawa, F. K., Amorim, P., Curcio, E., & Santos, M. J. (2021). Recent dynamic vehicle routing problems: A survey. *Computers and Industrial Engineering*, 160, 107604. <https://doi.org/10.1016/j.cie.2021.107604>. URL: <https://www.sciencedirect.com/science/article/pii/S0360835221005088>
- Snoeck, A., Merchán, D., & Winkenbach, M. (2020). Revenue management in last-mile delivery: State-of-the-art and future research directions. *Transportation Research Procedia*, 46, 109–116. <https://doi.org/10.1016/j.trpro.2020.03.170>. The 11th International Conference on City Logistics, Dubrovnik, Croatia, 12th–14th June 2019. URL: <https://www.sciencedirect.com/science/article/pii/S2352146520303707>
- Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2022). Stochastic dynamic vehicle routing in the light of prescriptive analytics: A review. *European Journal of Operational Research*, 298(3), 801–820. <https://doi.org/10.1016/j.ejor.2021.07.014>. URL: <https://www.sciencedirect.com/science/article/pii/S037721721006093>
- Strauss, A., Gulpinab, N., & Zheng, Y. (2020). Dynamic pricing of flexible time slots for attended home delivery. *European Journal of Operational Research*. <https://doi.org/10.1016/j.ejor.2020.03.007>.
- Ulmer, M. W. (2020). Dynamic pricing and routing for same-day delivery. *Transportation Science*, 54(4), 1016–1033. <https://doi.org/10.1287/trsc.2019.0958>.
- Voccia, S. A., Campbell, A. M., & Thomas, B. W. (2019). The same-day delivery problem for online purchases. *Transportation Science*, 53(1), 167–184.
- Wang, X., Zhan, L., Ruan, J., & Zhang, J. (2014). How to choose “last mile” delivery modes for e-fulfillment. *Mathematical Problems in Engineering*. <https://doi.org/10.1155/2014/417129>.
- Walßmuth, K., Köhler, C., Agatz, N., & Fleischmann, M. (2022). Demand management for attended home delivery—A literature review. *ERIM Report Series Research in Management*. ERS-2022-002-LIS, <https://ideas.repec.org/p/ems/eureri/137109.html>.
- Yang, X., & Strauss, A. (2017). An approximate dynamic programming approach to attended home delivery management. *European Journal of Operational Research*, 263, 935–945. <https://doi.org/10.1016/j.ejor.2017.06.034>.
- Yang, X., Strauss, A. K., Currie, C. S. M., Eglese, R., et al., (2016). Choice-based demand management and vehicle routing in e-fulfillment. *Transportation Science*, 50(2), 473–488. <https://doi.org/10.1287/trsc.2014.0549>.