

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/156191>

Copyright and reuse:


This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it.

Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk



Revenue Management in Product and Service Innovations

by

Yijun Zheng

Thesis

Submitted to the University of Warwick

in partial fulfilment of the requirements

for admission to the degree of

Doctor of Philosophy in Management

Warwick Business School

June 2020

Contents

List of Tables	iv
List of Figures	v
Acknowledgments	vii
Declarations	viii
Abstract	ix
Chapter 1 Introduction and Background	1
1.1 Revenue Management and Applications	2
1.2 Customer Choice Models	5
1.3 Modelling and Solution Approaches	6
1.4 Challenges in Revenue Management	8
1.5 Main Objectives and Contributions of Thesis	9
1.6 Structure of Thesis	13
Chapter 2 Dynamic Pricing of Flexible Time Slots for Attended Home Delivery Management	14
2.1 Home Delivery Services with Time Slots	15
2.2 Literature Review on Demand Management in Attended Home Deliv- ery Services	17
2.3 The Dynamic Pricing Model for Delivery Time Slots	20
2.4 Pricing Policy under MNL Choice Model	24
2.5 A Model-based Opportunity Cost Approximation	28
2.6 Computational Experiments	31
2.6.1 Data and Experimental Design	32
2.6.2 Numerical Results and Analysis	35
2.7 Conclusions	42

Chapter 3	Dynamic Delivery Time Slots Management and Vehicle Routing across Multiple Days	44
3.1	Delivery Service Booking System with Multiple Days	45
3.2	Literature Review on Time Slot Management	48
3.3	The Dynamic Slot Assortment and Vehicle Routing Problem	50
3.4	Construction of Delivery Routes	56
3.5	Demand Model	61
3.6	Online Problems	64
3.6.1	Opportunity Cost	65
3.6.2	Slot Assortment Optimisation Under Choice Models	66
3.7	Numerical Experiments	69
3.7.1	Data and Experiment Design	69
3.7.2	Computational Results and Analysis	73
3.8	Conclusions	80
Chapter 4	Dynamic Capacity Allocation for Airline Upgrade Options with Customer Anticipation	83
4.1	Cabin Upgrade System for Airlines	84
4.2	Literature Review on Service Upgrades	87
4.3	The Dynamic Capacity Allocation under Customer Anticipation	90
4.4	The Assortment Problem using Sequential Choice Model	92
4.5	A Choice-based Linear Approximation Model	96
4.6	Numerical Experiments	99
4.6.1	Design of Experiments and Data	99
4.6.2	Computational Results and Discussion	101
4.7	Conclusions	107
Chapter 5	Summary of Thesis and Future Work	109
5.1	Summary of Research Questions	109
5.2	Contributions to Practice	110
5.3	Contributions to Theory	111
5.4	Limitation of Thesis	112
5.5	Future Work	112
Appendix A	Proof of Propositions	115
A.1	Proposition 2 (Linearization of Slot and Price Assortment Problem)	115
A.2	Proposition 3 (Linearization of Value Function Estimation)	118

Appendix B	Estimated Choice Models	121
B.1	Estimated MNL Model with Substitution Effect	121
B.2	Estimated Nested MNL Model	122
B.3	Estimated MNL Model with No Substitution Effect	123

List of Tables

2.1	Utility parameters of standard slots under the nested MNL model	32
2.2	Specification of flexible slots and their utility parameters	33
2.3	Utility parameters in the estimated MNL model	34
3.1	Parameters in the nested MNL model for slots offered in two delivery days	63
4.1	Utility parameters in the MNL choice model	100
B.1	Estimated parameters of the MNL model involving two delivery days	121
B.2	Estimated parameters of the nested MNL model for two delivery days	122
B.3	Estimated parameters of two independent MNL models	123

List of Figures

2.1	Network structure for flexible slot m and a slot $s \in \mathcal{S}$	25
2.2	Delivery service area provided with flexible slots	34
2.3	Profit increase (%) under P3 relative to never using flexible slots . .	36
2.4	Revenue/order and cost/order increases (%) under P3 relative to never using flexible slots	37
2.5	Vehicle utilisation under different demand patterns under P3	38
2.6	Performance comparison of scenarios (A4 vs P3) relative to using only standard slots	39
2.7	Cost per order and revenue per order increases (%) under scenarios (A4 vs P3) relative to using only standard slots	40
2.8	Final allocation for flexible orders under P3	41
2.9	Final allocation for flexible orders under A4	42
3.1	Time slots from K delivery days presented to customers whose arrival process separated into K booking horizons	46
3.2	An illustration of booking and planning horizons with time slots from multiple delivery days	51
3.3	Choice probabilities for time slots in two delivery days defined under different choice models	64
3.4	Area of attended home delivery services	70
3.5	Total profit increase (%) under MNL-IND, MNL-DAY and TRUE compared to ALL-FEASIBLE	74
3.6	Profit increase (%) per order under MNL-IND, MNL-DAY and TRUE compared to ALL-FEASIBLE	75
3.7	Accepted order increase (%) under MNL-IND, MNL-DAY and TRUE compared to ALL-FEASIBLE	76
3.8	Cost reduction (%) per order under MNL-IND, MNL-DAY and TRUE comparing to ALL-FEASIBLE	76

3.9	CPU time taken to check slot feasibility and decide slot availability for each delivery request under different policies	78
3.10	Total profit increases (%) after offering same-day delivery services .	79
3.11	Accepted order increase (%) after introducing same-day delivery services	79
3.12	Profit per order increase (%) after introducing same-day delivery services	80
3.13	Cost reduction (%) per order after introducing same-day delivery services	80
4.1	Impact of considering upgrade options and customer anticipation under full and partial anticipation effects	103
4.2	Performance comparison of AP and IP in terms of percentage revenue increase (%) generated from selling economy and business tickets and upgrade options	104
4.3	Customer anticipation levels achieved by AP and IP for each individual flight at demand scaling parameter 0.7	105
4.4	Customer anticipation levels maintained during the booking horizon of each flight under AP	106
4.5	Customer anticipation levels maintained by AP during the booking horizon of each flight given different initial anticipation level	107

Acknowledgments

Firstly, I would like to express my sincere gratitude to my supervisors Dr. Nalan Gulpinar and Dr. Arne Strauss for the continuous support of my PhD study and related researches, for their patience, motivation, and immense knowledge. Their guidance helped me in all the time of research and writing of this thesis. I could not have imagined having better supervisor and mentors for my PhD study.

Besides my advisors, I would like to thank all staff members from Operations group for their insightful discussions and encouragement during my study, but also for the hard questions which incited me to widen my research from various perspectives.

I thank my colleagues from WBS for the stimulating discussions and their generous support during my PhD study. Also I thank my friends I have met in Warwick for all the fun we have had during the last four years.

Last but not the least, I would like to thank my family for supporting me spiritually throughout writing this thesis and my life in general.

Declarations

I declare that any material included in this thesis has not been submitted for a degree to any other University. I have made major contributions to all three research topics in revenue management and pricing, including defining research questions, proposing solution methods, conducting numerical experiments and writing research papers. Furthermore, Chapter 2 of this thesis is modified based on the paper titled "Dynamic Pricing of Flexible Time Slots for Attended Home Delivery", which is co-authored with Nalan Gulpinar and Arne K Strauss. It has been accepted by European Journal of Operational Research special issue on OR Approaches for Retail Operations. Chapter 3 is adapted from the working paper titled "Dynamically Managing Delivery Time Slots and Vehicle Routing Across Multiple Days", which is co-authored with Nalan Gulpinar and Arne K Strauss. Finally, the working paper titled "Dynamic Capacity Allocation for Airline Upgrade Options with Customer Anticipation", drawn from Chapter 4 of the thesis, is co-authored with Nalan Gulpinar and Arne K Strauss.

Yijun Zheng

June 2020

Abstract

This thesis is concerned with revenue management problems related with innovative products and services in two merging business sectors: attended home delivery services and airline upgrading system. The attended home delivery service is a fulfilment method offered by e-retailers; and the airline upgrading system is used to resolve demand-capacity mismatch in business and economy cabins. In the first research topic, we define the concept of flexible time slots as any fixed combinations of regular delivery time slots. In exchange for cheap delivery charges, the customer is informed shortly prior to the delivery day in which regular time window the goods arrive. We evaluate the benefit of introducing flexible time slots in terms of increasing total profit and improving delivery efficiency by deriving a dynamic pricing policy. In practice, e-retailers present delivery time slots across multiple days to customers. Different models can be adopted to capture the customer selection behaviour depending on whether the model considers the substitutional effect of delivery days on customer choice decision. We investigate the effect of embedding different choice models within a dynamic slotting policy on the e-retailer's profit in the second research topic. Our third research topic discusses airlines' capacity allocation problem with upgrade options involving multiple flights. Within upgrade options, airlines can postpone upgrade decisions until demand is fully realised and customers only need to pay once their options are executed. However, customers may anticipate the availability of upgrade options based on their past experience and using them to obtain business capacity. Therefore, we address the importance of correctly accounting customer anticipation in managing the capacity in economy and business cabins in the long term.

Those problems are challenging as they all involve real-time decision making under a stochastic customer arrival process and the uncertainty from customer choice behaviour. Solving these problems involves modelling customer choice behaviour, forecasting demand, estimating displacement costs and considering constraints imposed by innovative business models. Moreover, efficient solution methods are required to make decisions in real time. In this thesis, we make contributions to the literature by formulating those problems using DP models and proposing effective choice-based pricing and capacity control policies that can be solved quickly. We also address practical issues and derive managerial insights on those innovative products and services by applying these policies in simulation studies.

Chapter 1

Introduction and Background

This chapter briefly introduces the background on revenue management as well as its recent applications and choice models used in capturing customer purchase behaviour. Since Markov Decision Processes are widely adopted in the literature to address dynamic decisions, we will also cover the basic methodologies used for modelling and solving MDP problems. Then, key challenges involved in making revenue management decisions are highlighted. Moreover, we will describe three research topics discussed in this thesis along with their contributions. Finally, we present the structure of this thesis.

This thesis aims to quantify the benefit of introducing three innovative services/products offered in two business sectors. We firstly investigate the benefit of offering flexible slots along with standard slots in attended home delivery services. By introducing flexible slots, customers have more delivery options to select from and the e-retailer gains flexibility in scheduling final delivery routes. Then, we focus on managing time slots across consecutive delivery days, which grants customers more service options but creates e-retailers complexity in managing their delivery schedules. Finally, we consider airlines' upgrade services offered via options to customers who have repeat purchase behaviour. When upgrade options are offered, the airline can potentially achieve higher revenue to fill its business cabin and customers are able to obtain business cabin capacity through cheap upgrade options. The underlying revenue management problems dealing with management of these products and services are formulated with dynamic programming and solved by linear-programming-based approximation methods. Simulation studies demonstrate their value and contribution to business.

1.1 Revenue Management and Applications

Revenue management (RM) is concerned with methodologies to make demand management decisions such that businesses can improve selling decisions for their products and services and maximise revenue (Huefner, 2011). The idea of RM was first introduced by American Airlines after the flight deregulation in 1978 (Tudor Bodea, 2014). As flights have fixed perishable service capacity and high operational costs, flight ticket prices are strategically set by the airline in order to fill up the flight and further increase revenue.

In general, RM focuses on three types of decisions including structural, price and quantity decisions (Talluri, 2005). Price and quantity decisions are at a tactical level, which are frequently adjusted by the business based on market conditions. Price decisions mainly address the issue of setting differentiated prices to customers over time while quantity related decisions generally deal with the capacity allocation to different products (Ng, 2008). On the other hand, structural decisions are made at the strategical level and related to the mechanism used for selling products and involving commitment to certain price or quantity decisions. For example, the airline is not able to make quantity or price decisions after it offers delayed purchase options, which allow customers to reserve a seat with a fixed fare for a fixed duration before making a final purchase decision (Aydın et al., 2016).

There are two main types of control mechanism used for price and quantity decisions discussed in RM. The first one is static control mechanism which makes the current decision without considering the possibility of updating decisions in the later customer arrival process (Belobaba, 1987). For example, a static policy can set a protection level for each fare class of a flight before accepting customers. Dynamic control is another type of mechanism where decisions are made with respect to each arrival under the consideration of future revenue opportunities (McGill and Van Ryzin, 1999). Under such mechanism, the customer demand is modelled by a stochastic arrival process of individual purchase requests (Gallego and Van Ryzin, 1994). For example, during the booking horizon of a single-leg flight, customers' arrival process is assumed to follow a Poisson process and the presented fare classes are decided for each customer based on displacement costs (Burger and Fuchs, 2005, Feng and Xiao, 2001). Compared to static control mechanism, policies derived based on dynamic control mechanism are able to generate higher revenue (Maglaras and Meissner, 2006, Wright et al., 2010).

In addition to airlines, RM has been applied in various industries, such as hotels, restaurants and retails. Depending on the context of a specific application, RM may focus on price or quantity decisions. As having similar characteristics with airlines, hotel management is another sector that has implemented RM practices

earlier than others (Kimes, 2011). For example, hotels may use a pricing policy to make adjustment over their seasonal reference price to maximise the revenue (Bayoumi et al., 2013). However, restaurants have perishable capacity due to variable service duration and tables with different sizes (Kimes et al., 2007). Therefore, unlike airlines and hotels, RM in restaurants focuses on the turnover rate of tables and the revenue per available seat-hour (Xiao and Yang, 2010). The most profitable combination of tables and the quantity decision of assigning tables to customers are the main concerns of RM practices in restaurants (Guerriero et al., 2014).

RM has been introduced by retailers who bear high variable operation costs for managing their inventories (Lippman, 2003). Deciding which products to display in the limited shelf-space, also known as the assortment problem, is one type of RM problems to maximise revenue (Geismar et al., 2015). In the sales season, dynamically setting the mark-downs on products is another RM problem for retailers that aims to increase the total revenue (Heching et al., 2002).

Globally, business to consumer e-commerce has been growing. The attended home delivery (AHD) services have been introduced by e-commerce as a fulfilment method (Hays et al., 2005). The AHD services involve three processes: sales, purchase and delivery (Agatz et al., 2008). In the sales step, an interface containing all available time windows in one specific delivery day with corresponding prices is displayed to the customer. Following that, the customer selects exactly one time window for the delivery service during the purchase process. Customers can select time windows until a certain cut-off time. At the final delivery stage, routes are generated to accommodate all accepted customers' delivery requests within corresponding time windows.

Having delivery services with specific time windows makes customers' online shopping experience fast and convenient (Ehmke, 2012). However, it increases delivery cost and reduces the total profit for e-retailers because there is little flexibility to make delivery routes more efficient (Joel and Carol, 2006). At the same time, the market report from Mintel (2017) has identified that customers tend to be very sensitive to delivery charges. Therefore, it is not easy to compensate the cost increase by increasing the delivery charge. Moreover, based on the real customer booking data analysed by Yang et al. (2014), customers have preference towards time windows within the delivery day.

Innovative services are also provided by airlines. Airlines are used to offer customers upgrades when approaching the flight departure time in order to balance the mismatched supply and demand among business and economy cabins in a flight (Steinhardt and Gönsch, 2012). Alternatively, Optiontown (2018) has introduced a new approach of offering upgrades via options. Specifically, upgrades are offered as

priced-options to customers who have just purchased economy tickets. The customer needs to immediately decide whether to purchase the option or not. At the end of the booking horizon, the airline executes a number of options based on the availability of the capacity in the business cabin. In this way, an additional revenue can be potentially obtained by the airline from executing upgrade options. However, in the long term, offering upgrade options may lead to cannibalisation. Customers may quickly learn and anticipate the availability of upgrade options such that they would get the business cabin capacity by purchasing the economy cabin capacity and booking the upgrade options (Wu and Chen, 2000). Then, the airline would face the revenue loss from decreasing demand of business class. Therefore, it is essential for airlines to consider customers' anticipation when providing upgrade options such that cannibalisation in business cabin can be avoided in the long term.

This thesis specifically investigates three problems related to novel products and services introduced by e-retailers and airlines. More specifically, we study RM problems in attended home delivery services and airline upgrades. As cost is the main issue in attended home delivery services, our first topic discusses the benefit of introducing flexible time slots in terms of improving profit and reducing cost. Moreover, in practice, customers are presented with time slots across multiple days when they request delivery services. Given such practice, in the second topic, we compare the e-retailer's profit when different choice models are adopted in the dynamic capacity control policy. Finally, the third research addresses the importance of accounting customer anticipation on upgrade options while allocating capacity for a number of flights.

We need to deal with a number of challenges in order to address those problems in two application contexts. For attended home delivery services studied in this thesis, delivery costs are considered by e-retailers in addition to revenues within their objective. As the customer arrival process is stochastic, we do not have full information on all orders at the time when accepting each request. Accordingly, we need to estimate total delivery costs without full order information when evaluating the cost of accepting each request. This brings additional challenges, because these costs are obtained by solving a vehicle routing problem with time windows (VRPTW), which is a NP-hard problem (Golden et al., 2008). It requires large computational efforts to anticipate delivery costs. When a number of delivery days is managed simultaneously to accept customers, it imposes even more computational pressure to the system as the delivery cost is anticipated for individual day. Moreover, when the airline manages the capacity in a number of flights with upgrade options, customers' anticipation is updated based on their previous purchase experience and going to affect customer choice behaviour in their next flight purchase. It requires a sound

identification on factors; this does not only indicate the customer anticipation but also can be linked with customer choice behaviour.

1.2 Customer Choice Models

A revenue management problem in any context requires knowledge on factors influencing customers' price sensitivity and demand. Modelling customers' purchase behaviour using choice models has been well studied in the RM literature. An extensive literature survey on the choice models used in the revenue management problems is provided by Strauss et al. (2018). In an earlier study, the demand of a product is assumed to be independent from the availability of other products (Belobaba and Weatherford, 1996, Talluri and van Ryzin, 2008). Under this independent demand assumption, the customer only considers the purchase of a specific product and leaves the system only when the product is unavailable. On the other hand, dependent choice models, such as the multinomial logit (MNL) choice model, are also widely used to compute customer choice probabilities when availability of other products affects the customer's purchase decision of a specific product.

Suppose a decision maker faces J alternatives. Let U_j denote the utility obtained by decision maker from alternative j . The utility of alternative j is calculated as $U_j = V_j + \epsilon_j$, where V_j is a known parameter and ϵ_j is an unknown part presented by a random variable following a certain distribution. Note that each ϵ_j is assumed to be independently, identically distributed following a Gumble distribution to derive the MNL model. Under the MNL choice model, the probability of decision maker selecting alternative i is defined by

$$p_i = \frac{e^{V_i}}{\sum_{j \in J} e^{V_j}}. \quad (1.1)$$

The pioneer work of Talluri and van Ryzin (2004) analyses the quantity decisions for a single-leg flight RM problem and models the customer purchase behaviour using an MNL choice model. In order to deal with large-scale practical problems, choice-based heuristics have been developed for the network revenue RM problems involving multiple resources and products, for instance, see Gallego and Phillips (2004), Liu and van Ryzin (2007). It is well known that taking customer purchase behaviour into account for solving RM problems brings computational challenges. However, accounting customer choice behaviour in the decision making could create potential revenue increase (Tudor Bodea, 2014). For example, InterContinental Hotel Groups (IHG) has managed to increase the annual revenue by 2.7%

after implementing a choice-based RM system (Koushik et al., 2012).

Apart from the MNL choice model, there is also a Nested Logit (NL) choice model when alternative products can be classified into a number of nests. Moreover, for any two alternatives in the same nest, the ratio of probabilities is independent of the attributes or existence of all other alternatives (Train, 2003). Assume that a set of J alternatives can be partitioned into K non-overlapping subsets (nests) denoted by B_1, B_2, \dots, B_K . The utility obtained by decision maker from alternative j denoted as U_j has the same definition as the one in the MNL model, which also includes an observable parameter and an unobserved part. Let λ_k denote a parameter measuring the degree of independence in un-observed utility among the alternatives in nest k . A higher value of λ_k means greater independence and less correlation. Accordingly, under the nested logit choice model, the probability of decision maker selecting alternative i from nest k is defined by

$$p_i = \frac{e^{\frac{V_i}{\lambda_k}} \left(\sum_{j \in B_k} e^{\frac{V_j}{\lambda_k}} \right)^{\lambda_k - 1}}{\sum_{l=1}^K \left(\sum_{j \in B_l} e^{\frac{V_j}{\lambda_l}} \right)^{\lambda_l}}. \quad (1.2)$$

The NL choice model has also been used in revenue management literature. For example, Anderson and Xie (2012) estimate customers' behaviour for booking hotels through an online platform based on an NL choice model where the nests are defined by the geographic neighbourhood. In this thesis, we adopt both the MNL and NL choice models in solving three different revenue management problems and these models will be explained in details in the corresponding chapter accordingly.

1.3 Modelling and Solution Approaches

In practice, people make sequential decisions under uncertainty where each decision leads to an immediate reward (Bellman, 1954). Markov decision processes (MDPs), also named stochastic dynamic programming, is introduced by Howard (1960) to formulate and solve multi-stage decision-making problems under uncertainty.

MDPs have been widely used in many applications. Sauré et al. (2012) adopt it to formulate the patient appointment scheduling problem. They derive a policy that allocates the treatment capacity to each incoming demand such that the waiting time is reduced. Problems in supply chain management are also formulated as MDPs. For instance, Tempelmeier (2007) focus on the production planning problem for a single product under demand uncertainty. Based on the MDPs formulation, a

solution method is proposed to decide how many units of inventory to order at each time period so that total cost consisting of setup and holding costs is minimised. Furthermore, Kleywegt et al. (2004) formulate an inventory routing problem using MDPs which combines the inventory management and the vehicle routing problem. They find a policy that decides when and how much inventory should be restocked for each customer to minimise the total cost in the long term. In this PhD thesis, we are mainly concerned with formulating AHD and airlines revenue management problems using MDPs and investigating policies to increase the profitability and efficiency of the business operations.

MDPs consider a planning horizon consisting of T discrete time periods, and include a state space \mathcal{S} defining the system. At each time period t , $t = 1, \dots, T$, a decision x_t from the feasible set (so-called action space) \mathcal{X}_t needs to be taken at given state $S_t \in \mathcal{S}$. The decision leads to reward (cost) denoted by $R_t(S_t, x_t)$ and the system is updated into a new state S_{t+1} with probability $P(S_{t+1}|S_t, x_t)$. Note that $P(S_{t+1}|S_t, x_t)$ is a conditional probability which only depends on the current state and action taken at state S_t . The decision maker aims to find a decision policy such that the expected total rewards is maximised. The optimal policy $\mathbf{x}^* = (x_1^*, \dots, x_T^*)$ can be obtained by solving:

$$V_t(S_t) = \max_{x_t \in \mathcal{X}_t} \left(R_t(S_t, x_t) + \sum_{s' \in \mathcal{S}} P(S_{t+1} = s' | S_t, x_t) V_{t+1}(s') \right), \quad (1.3)$$

$$t = 1, \dots, T - 1,$$

with the boundary condition $V_T(S_T) = R_T(S_T)$ calculating the reward obtained at the final time period.

Given a bounded state space \mathcal{S} , backward recursion algorithm can be applied to solve the MDPs in (1.3) to optimality (Bertsekas, 1995). The backward recursion algorithm starts with computing $V_T(S_T)$ for all possible states S_T at the last stage T . Based on these values, as well as the state-dependent action x_{T-1} , we move to stage $T - 1$ and calculate $V_{T-1}(S_{T-1})$ for all possible states S_{T-1} . Such calculation process continues by considering all stages in a backward manner until it reaches the first stage and all calculated values are kept in a table. Then, given initial state $S_1 = \hat{S}$, the optimal value $V_1(\hat{S})$ and its corresponding optimal action x_1 can be easily obtained from the table. The backward recursion algorithm is useful when there is no structural property in the MDPs model to exploit and the state space is relatively small (Powell et al., 2005). However, when the problem is large-scale, the state space \mathcal{S} , the action space \mathcal{X}_t and the outcome space for s' may be too large to evaluate the value function $V_t(S_t)$ for all states within a reasonable time. This is also known as curse of dimensionality.

On the other hand, the MDPs in (1.3) can be rewritten as a linear program (LP). Let σ_{t,S_t} denote the decision variable in the equivalent LP model which represents the value of $V_t(S_t)$ at $S_t \in \mathcal{S}$ for $t = 1, \dots, T$. Given the initial state $S_1 = \hat{S}$, the equivalent LP can be formulated as follows:

$$\begin{aligned}
& \min \sigma_{1,\hat{S}} \\
& s.t. \sigma_{T,S_T} = R_T(S_T), \quad \forall S_T \in \mathcal{S}, \\
& \quad \sigma_{t,S_t} \geq R_t(S_t, x_t) + \sum_{s' \in \mathcal{S}} P(S_{t+1} = s' | S_t, x_t) \sigma_{t+1,s'}, \\
& \quad \forall S_t \in \mathcal{S}, x_t \in \mathcal{X}_t, t = 1, \dots, T-1.
\end{aligned} \tag{1.4}$$

Algorithms, such as Simplex and Interior-Point algorithms, can be applied to obtain optimal solution to the LP problem (1.4) (Murty, 1983). Then, its optimal solution $\sigma_{1,\hat{S}}^*$ is equivalent to the value $V_1(\hat{S})$ defined in MDPs problem (1.3). However, the LP problem (1.4) may involve too many constraints when the underlying MDPs problem in (1.3) is a large-scale problem with too many stages, states and potential actions. The run time of solving the LP problem grows up exponentially in the dimension. Accordingly, the approximate linear programming algorithms have been proposed in the literature to improve the efficiency. Trick and Zin (1993) have developed heuristics that construct an LP by aggregating state. Schuurmans and Patrascu (2002) have proposed constraint generation methods to avoid considering all constraints in the original LP (1.4).

1.4 Challenges in Revenue Management

Business makes pricing/acceptance decisions at every time whenever a request arrives. Since customer requests arrive stochastically during the planning horizon, decisions have to be made under a dynamic environment. Most importantly, a decision made now (at real time) influences future decisions. Thus, those decisions cannot be made independently. As stated by Puterman (2014), making decisions without recognising the relationship between current and future decisions may fail to obtain good overall performance. For example, in the case of selling tickets for one flight, accepting a large number of customers in low-fare class at the beginning may fill up the flight quickly but this may lead to high opportunity costs by rejecting customers from high-fare class (Subramanian et al., 1999). Moreover, because business makes decisions without knowing customers' purchase decisions, uncertainties resulting from customer choice behaviour are involved when the business receives the reward (Morgan et al., 1992). As in the previous example, the customer will not book the flight if the presented fare is higher than its willingness-to-pay (Sierag et al., 2015).

Apart from huge impact of modelling issue related to stochastic multi-stage decision making process, data plays an important role on the performance of dynamic models. In particular, data related to customer booking histories is essential in estimating customer arrival process and the customer choice model, which are further used in the decision policy. However, the dataset may only contain the customers who have made the booking but miss those who have decided to leave without booking. It could create difficulty in defining no-booking customers when we estimate customer behaviour models used in our policies. Moreover, the customer booking data may be censored, also known as constrained demand. In this case, the data doesn't include customers who are rejected due to lack of capacity. Using such data to forecast demand may under-estimate the real demand in solving the RM problems and may lead to revenue decrease (Cooper et al., 2006). Moreover, selecting a suitable customer choice model is important as a sophisticated choice model may not lead to an LP-based decision policy, which can be solved efficiently in polynomial time (Bront et al., 2009).

1.5 Main Objectives and Contributions of Thesis

This thesis consists of three research topics regarding original products and services provided in two areas, more specifically attended home delivery systems and airline upgrades. We evaluate the benefit of introducing each new product and service by constructing dynamic pricing and capacity control policy used in revenue management. We summarise the main objectives in each research topic and its corresponding contributions as follows.

Flexible Time Slots in Attended Home Delivery Services: The market of online grocery in the UK is estimated to grow 13.5% and reach £11.3 billion in 2017 (Mintel, 2018c). Attended home delivery services are commonly used by e-grocers within the market, such as Tesco and Ocado, to fulfil their customers (Hays et al., 2005). Having such home delivery services is one of the top reasons for customers to buy grocery online (Mintel, 2018c). However, the e-grocers' profitability can be affected by the cost of offering delivery services. For example, Amazon failed to achieve its target profit because of the delivery costs increase in the third quarter of 2016 (Dean, 2016). Under attended home delivery services, customers can request a specific delivery time slot from the e-grocer. It has been a trend that e-grocers provide narrower time slots to increase customers' satisfaction (Mintel, 2016). However, it may cause delivery cost increase.

In the first research topic, we introduce flexible time slots to the attended home delivery services. A flexible slot combines several standard slots such that the

e-grocer can decide which standard slot to accommodate the request after the booking horizon. The flexible slots grant the e-grocer flexibility in scheduling deliveries when all delivery requests are realised at the end of the booking horizon. By offering flexible slots, the total delivery costs might be potentially reduced and the total profit could be increased.

The main contributions of studying the dynamic pricing policy of flexible time slots can be summarised as follows:

- A novel linear programming (LP) formulation is proposed to estimate the opportunity cost, which reflects potentially displaced profits from future orders and implications on routing costs.
- An online pricing approach is derived as a computationally tractable LP by exploiting features of the choice model. Accordingly, this approach is able to make pricing decisions within milliseconds.
- The proposed approach is applied in the simulation study under a realistically-sized setting. The simulation study shows that introducing flexible delivery slots can significantly improve e-retailers' profitability by reducing total delivery costs and attracting more customers.

Managing Time Slots Across Multiple Delivery Days: By offering attended home delivery services, e-retailers improve customers' online shopping experience in terms of speed and convenience (Ehmke, 2012). Therefore, the delivery services have been identified as an essential factor that contributes to the future growth of online retailing (Intel, 2018d). In order to provide more delivery options to increase customer satisfaction, customers are offered with time slots from a number of delivery days by the e-retailer, such as Tesco and Ocado. Presented with time slots across different days, customers would like to compare delivery days as well as time slots before they make purchase decisions. Therefore, it is important for an e-retailer to construct an appropriate choice model reflecting customer choice behaviour. Having a 'wrong' choice model, the e-retailer is not able to influence customer preference towards time slots in a manner of optimising its delivery operation. As the demand is managed for a number of delivery days at the same time, the e-retailer may even end up with inefficient delivery operation for all involved delivery days by having a 'wrong' choice model.

In our second research topic, we derive three policies, also known as slotting policies, based on different choice models to decide which time slots to present every time a delivery request arrives. These choice models are different from each other based on how they treat the substitution effect of delivery days on customer slot

decisions. We compare these policies in terms of the profit achieved by the e-retailer of providing the delivery services within one specific day. Moreover, we exploit the opportunity of offering same-day delivery services as an express delivery service to a specific group of customers. These customers' requests can be quickly prepared and easily accommodated by the existing routing plan of the current day without routing plan reconstruction. Therefore, offering such same-day services could potentially increase customer satisfaction, generate more revenue and improve the efficiency of the current routing plans.

The main contributions of studying the dynamic slotting policy of time slots across multiple delivery days can be summarised as follows:

- We propose a dynamic program (DP) formulation to dynamically decide the availability of each slot of multiple delivery days whenever a customer arrives during the booking horizon. Specifically, we neglect the assumption that each delivery day has its corresponding booking horizon but focuses on a calendar-based booking horizon.
- We assume that a nested multinomial logit (NMNL) model is a ground true customer choice model. Two choice models are proposed to estimate this true choice model, which are used in the slotting policy. Both choice models are based on the multinomial logit (MNL) model. However, one model accounts the substitution effect among delivery days whereas the other model treats delivery days independently.
- In the numerical experiments, we find that using the approximated choice model with substitution effect attracts less customers than the case of using the other approximated choice model when demand is relatively low. We also demonstrate that using such choice model leads to inefficient routing plans and increases the delivery cost per order. Moreover, we illustrate that offering same-day delivery services generates extra profit-before-delivery and improves the efficiency of the routing plans on the current day by reducing the delivery cost per order.

Allocating Capacity for Airlines Upgrade Options: Upgrades are offered by airlines to resolve the mismatch between demand and supply in their economy and business cabins. Upgrades are sold as options to customers by an airline. At the end of booking horizon of each flight, the airline executes a number of upgrade options that maximises total revenue of the flight and also refunds payment for unexecuted options. By offering upgrade options, airlines are able to obtain additional revenue from executing upgrade options. Moreover, airlines can postpone the upgrade decisions at the end of the booking horizon after demand is full realised.

On the other hand, there are customers in airline management who regularly purchase the flight and have a network to share their booking experiences. Therefore, by frequent purchase experience and information network, these customers can gradually establish anticipation toward upgrade options and behave strategically to obtain capacity in business class using upgrade options. Moreover, at the end of booking horizon of every flight, information on executed options is released within the network, customers gain new knowledge on upgrade options and modify their purchase decisions for the next purchase. As a result, airlines may have revenue loss due to the introduction of upgrade option in the long term.

In the third research topic, we introduce the definition of customer anticipation level, which influences customer choice behaviour over business and economy cabins. Then, we formulate capacity allocation problem of an airline by accounting the customer anticipation level. We evaluate the benefit of accounting customer anticipation in avoiding cannibalisation from introducing upgrade options in the long term.

The main contributions of studying airlines' upgrade options with customers' anticipation can be summarised as follows:

- We introduce a novel DP model with continuous state space to formulate the capacity allocation problem for consecutive flights offering upgrade options. We define an anticipation level as the main factor evolved during the customer learning process, which affects customer purchase behaviour towards different cabins and options.
- As a solution method, we discretise this continuous state space and construct a choice-based deterministic integer program (CDIP) model to approximate the value function at all discrete states of the DP problem.
- In the numerical experiments, our results show that considering the customer anticipation in allocating capacity with upgrade options can significantly improve the total revenue in the long term. Moreover, based on our solution method, there exists a steady state (customer anticipation level) in solving the DP when the demand remains the same for all considered flights.

Overall, we can conclude that this thesis contributes to the revenue management literature by introducing novel problem formulations and solution methods for three interrelated research topics on innovative products and services introduced in AHD system and airline upgrades. As potential impact of all research topics discussed in the thesis on real-life applications, we construct dynamic control policies to illustrate practical benefits of introducing those products and services in terms of improving business revenue and creating flexibility in operations.

1.6 Structure of Thesis

This thesis consists of five chapters. After summarising the background and highlighting the objectives of this thesis in Chapter 1, we organise its remaining chapters as follows. Chapter 2 focuses on the dynamic pricing policy for flexible time slots introduced in attended home delivery services. Chapter 3 evaluates the benefit of using different choice models in the slotting policy to manage time slots from multiple delivery days. Then, Chapter 4 proposes upgrade options in airlines and discusses the importance of considering the customer learning effect in long-term revenue management. Finally, Chapter 5 highlights the contributions of the thesis and proposes further research directions could be investigated in home delivery services and capacity planning in airlines.

Chapter 2

Dynamic Pricing of Flexible Time Slots for Attended Home Delivery Management

In e-commerce, customers are usually offered a menu of home delivery time windows of which they need to select exactly one, even though at least some customers may be more flexible. To exploit the flexibility of such customers, we propose to introduce flexible delivery time slots, defined as any combination of such regular time windows (not necessarily adjacent). In selecting a flexible time slot (out of a set of windows that form the flexible product), the customer agrees to be informed only shortly prior to the dispatching of the delivery vehicle in which regular time window the goods will arrive. In return for providing this flexibility, the company offers the customer a reduced delivery charge.

In this chapter, we study dynamic pricing of regular and flexible time slots in this context for attended home delivery. The vehicle routing problem (VRP) in the presence of flexible time slots bookings corresponds to a VRP with multiple time windows. The main methodological contribution is the development of a tractable linear programming formulation that links demand management decisions and routing cost implications, whilst accounting for customer choice behaviour. The output of this linear program provides information on the (approximate) opportunity cost associated with specific orders and informs a tractable dynamic pricing policy for regular and flexible slots. Numerical experiments, based on realistically-sized scenarios, indicate that expected profit may increase significantly depending on demand intensity when adding flexible slots rather than using only regular slots.

2.1 Home Delivery Services with Time Slots

Globally, business to consumer e-commerce is growing strongly. According to the Ecommerce-Foundation (2016), global growth in turnover has been around 17.5% in 2016. Online grocery retailing, in particular, is growing at a similar rate: in the United Kingdom (UK), sales exhibit strong annual growth rates of 14.7% in 2016 with similar rates forecast over the next years by Mintel (2017). Attended home delivery is commonly offered in e-commerce when the ordered items are either bulky (like furniture), or in need of refrigeration (groceries), or for other reasons need to be handed over to the customer in person. There is much competitive pressure over the quality of the delivery service. In particular, shorter time windows increase customer satisfaction. Most UK grocery retailers (such as Tesco, Sainsbury's, Morrisons, Ocado, Waitrose) are now offering narrow one-hour delivery time slots, as opposed to the longer time windows offered in the past.

Such narrow delivery time windows lead to high fulfilment costs because there is little flexibility to make vehicle routes more efficient. At the same time, the cost cannot be easily passed on to the customers because they tend to be very sensitive to delivery charges. This is also reflected by market research conducted by Mintel (2017) showing that current online grocery shoppers, lapsed shoppers and non-users would be most encouraged to buy more online if delivery prices were lower. In other words, delivery charges are a major deterrent from online shopping. The combination of high delivery cost and limited capability to recoup the cost via delivery charges (or via increased product prices) indicates that there is a lot of pressure on making delivery services as efficient as possible; but the high sensitivity to delivery prices also means that delivery prices can be used to influence customers' delivery time slot choice behaviour.

Despite the competitive pressure over offering narrow delivery time slots, not all customers actually require them to be so narrow. Some may be willing to accept uncertainty over the exact delivery time slot within a given set of potential time slots in return for an incentive. This has recently been exploited by the UK's largest retailer Tesco in that they offer so-called 'Flexi Saver Slots' alongside their regular one-hour slots. Flexi Saver Slots are four hours in length, and the customer is notified on the day of delivery of a one-hour slot (within the booked four-hours window) in which the delivery will be made. The customer pays less in delivery charges in return for giving the retailer more flexibility in their fulfilment operations. This may allow the retailer to accept more orders and/or to fulfil them more efficiently. Accordingly, we generalise this concept further by defining a flexible time slot as any fixed combination of regular delivery time slots, so they do not necessarily need to be adjacent.

We stress that it is not the concept of offering flexible slots that is the new contribution here (this has been used by some retailers for a while already); instead, the challenge lies in quantifying the savings potential of a flexible slot as well as in dynamically pricing these slots under a model that incorporates customer choice behaviour. When making a decision on how to price a flexible slot, we need to take into account how much we may be able to save in the routing due to this flexibility, which is difficult to assess because we do not have full information on all orders at the time of making this pricing decision.

In this research, we study the dynamic pricing problem faced by a firm offering regular and flexible delivery time slots for attended home delivery. We assume that delivery requests for a specific day arrive randomly over a fixed time horizon prior to the delivery day so that the delivery operation takes place after all orders have been received. In other words, we do not consider the same-day delivery where requests may arrive after some delivery vehicles have already been dispatched. The decision problem of the firm arises every time a customer (from a known location and with known order size) requests delivery, and consists of i) evaluating which time windows can feasibly be offered, and ii) deciding which delivery prices to display for all feasible slots. Order size refers to both the number of delivery totes and order profit before delivery cost, and may be estimated from the previous purchasing cases where delivery slots can be booked before completing the shopping session. We assume that all feasible slots are offered so as to increase customer satisfaction, but this assumption can easily be relaxed in our model without structural changes (another dummy price point is needed that drives demand to zero). Delivery charges are assumed to be chosen from a finite set of price points, in line with common industry practice. In response to the firm's decision, the customer chooses a slot or decides to leave without purchase according to a discrete choice model that reflects the set of available options and prices.

After all orders have been received, the firm needs to solve a capacitated vehicle routing problem with multiple time windows for its fleet of delivery vehicles. Note that 'multiple time windows' refers to having some customers with multiple time windows within which the delivery may take place (namely those customers with flexible slots). The objective is to maximise profit after delivery cost by dynamic delivery slot pricing and routing of the delivery vehicles to serve the final set of orders.

Our main contribution is a new approach of how to estimate the opportunity cost associated with accepting a given order in the different delivery options. This opportunity cost reflects the implications on routing costs and potentially displaced profits from future orders in case constraints on the van capacity and/or driving

time are binding. The estimation of opportunity cost is very difficult because the calculation of the final delivery cost is challenging even if we already know the final set of orders (which, however, we do not). To tackle this challenge, we propose a novel linear programming (LP) formulation that accounts heuristically for both delivery costs and future expected order revenue. This LP is solved offline and should be re-optimised throughout the booking horizon so as to provide updated estimates of the opportunity cost. Furthermore, pricing decisions need to be made in a very short time interval. To that end, we propose an online pricing approach that exploits the features of the choice model and the constraint structure and, thereby, equivalently reduces the nonlinear optimisation problem to a small LP that can be solved very quickly. We evaluate the proposed approach in a realistically-sized simulation study and our results show that the concept of flexible delivery slots can significantly improve profitability.

The chapter is organised as follows: in Section 2.2, we review the literature on demand management in the context of attended home deliveries. In Section 2.3, we define the problem as a Markov decision process and present an intractable dynamic programming formulation that is useful to motivate approximate solution methods. In Section 2.4, we formulate the pricing policy with flexible slots. In Section 2.5, we develop our LP approach to opportunity cost approximation. Section 2.6 contains the computational results and we draw conclusions and implications for managers in Section 2.7.

2.2 Literature Review on Demand Management in Attended Home Delivery Services

We focus on the growing literature on combining demand management with delivery slot booking for attended home deliveries. Demand management in our context is to be understood as optimisation of actions that have a direct influence on demand, specifically pricing, deciding on incentives or on the availability of certain offerings. For the reader who is interested in a broader context of e-fulfilment, we refer to the review of Agatz et al. (2008) covering the e-fulfilment literature from an operational research perspective, and Hübner et al. (2016) who review the recent qualitative fulfilment and distribution literature.

The first paper to consider aspects of both demand management and delivery operations is Campbell and Savelsbergh (2005) who investigate a dynamic routing and scheduling problem of a grocery vendor who needs to decide which deliveries to accept or reject, and in which time slot to deliver the accepted orders. All customers have a certain time slot profile; this contains all slots that they are willing to accept.

If the grocer accepts the order, the company assigns one of these slots to the order.

In this first paper, Campbell and Savelsbergh (2005) represent demand as an arrival process that is not affected by the firm’s decisions. In their subsequent work Campbell and Savelsbergh (2006) use a relatively simple customer behaviour model to include the effect of incentives (such as delivery charges) on the probability that any particular time slot is being chosen. The objective is to influence delivery time slot choices to minimise delivery costs, whereas in our work we focus on maximising expected profit.

A more realistic customer choice model was employed by Asdemir et al. (2009), namely the multinomial logit (MNL). They consider a dynamic time slot pricing approach similar to our approach, but propose dynamic programming (DP) as a solution method with fixed delivery costs rather than our LP-based approach. The DP is formulated at the level of a delivery region (such as a postcode sector) under the assumption that the delivery capacity in each time slot for this region is fixed and known a priori. Practical application of this approach may be challenging when there are many delivery time slots because the DP’s state space grows exponentially with the number of slots. In our approach, we also make use of the MNL choice model, but propose a new way of including dynamic delivery cost estimates into a LP model which allows us to solve it for realistically-scaled problem instances.

In contrast to this work on dynamic pricing in attended home delivery, Agatz et al. (2010) focus on the problem of which delivery time slots to offer in which geographic delivery area so as to reduce delivery costs whilst meeting service requirements. They do not consider customer choice behaviour, whereas our focus is on pricing to influence customers’ time slot choices. However, there are some common elements in that they also use the work of Daganzo (1987) to obtain a continuous delivery cost approximation.

Another work that stresses the routing and scheduling aspects (as opposed to demand management) in the attended home delivery context is Ehmke and Campbell (2014). Their objective is to maximise the number of requests accepted for delivery, subject to retaining feasible tours. The company makes decisions on accepting or rejecting delivery slot bookings, and the customers’ slot choices are assumed to be independent of these controls. In our work, we aim to maximise total expected profit by deciding on prices for regular and flexible delivery slots which influence customers’ slot choices. Cleophas and Ehmke (2014) likewise consider accept/reject decisions along with capacity reservations for certain delivery areas and time windows where particularly valuable demand is being expected.

Yang et al. (2014) is more closely related to our work in that the authors consider a dynamic pricing problem for delivery slots under the MNL choice model.

Using real data, they estimate the choice model and find that demand is very sensitive to delivery prices and slot availability. In their pricing policy, they only rely on opportunity cost estimates based on marginal routing costs (derived by insertion heuristics). They do not consider the effect of future lost revenues due to displaced orders in the opportunity cost estimate. This is addressed by Yang and Strauss (2017) who use an approximate dynamic programming approach to incorporate both future revenue and routing cost effects in the opportunity cost. Likewise, Koch et al. (2017) employ approximate dynamic programming to approximate opportunity cost including revenue and cost effects. However, their paper is centred around the idea of quantifying the free delivery time within each time slot for a given route plan. These so-called time budgets are then used to construct value function approximations.

The estimation of routing cost is a major challenge in demand management for attended home delivery. Bühler et al. (2016) discuss various linear mixed-integer programs that approximate the delivery costs for a fixed pool of route candidates. Klein et al. (2018) combine such a linear mixed-integer program (MIP) with the dynamic pricing model of Yang et al. (2014) so as to anticipate future demand. However, the MIP involves a very large number of decision variables for real-life scaled problems, thus making it challenging to solve. Also Song et al. (2018) propose an MIP approach to solving the last-mile delivery problem, however, they focus on the additional requirement of customers having to be served by a specific driver (so as to build trust in the company through repeated interaction with the same individual).

The work of Köhler et al. (2019) is conceptually related in as far as they investigate how to dynamically control the offering of long and/or short delivery time windows to customers in an attended home delivery context. However, here the slots are not ‘flexible slots’ in the sense that we propose in this chapter; instead, their term ‘flexible time window management’ refers to deciding dynamically which short and/or long slots shall be made available to a given customer (so some customers may be shown only long time windows, other only short ones, and yet others a mix of both). They assume fixed delivery fees for all deliveries regardless of time window length or time of day whereas we use dynamic pricing that reflects both customer preferences and opportunity costs associated with having a customer book any given slot.

In the remainder of this section, we review the literature on flexible products. A flexible product was first proposed by Gallego and Phillips (2004). They define it as a product that can be provided in one of a small number of modes. Customers are aware of this set of modes at the point of purchase, but only receive confirmation of the actual mode at a pre-defined time after purchase (usually shortly prior to

product consumption). The product is usually a service such as a flight; in this case, potential modes could correspond to different flights between the same origin and destination but at different departure times. Indeed, Gallego and Phillips (2004) study the problem in the airline context and propose a booking limit control policy for the flexible ticket under a static setting with two time periods and two alternative flights. Gallego et al. (2004) extend this concept to networks, and also consider customer choice modelling. They introduce a deterministic linear program that can approximate the optimal objective of the stochastic optimisation problem. Petrick et al. (2010b) likewise propose a deterministic linear program, but focus on independent demand only. They explicitly incorporate the capacity requirements of requests for flexible products that have been previously accepted and thereby allow them to be rearranged.

Among these deterministic linear programming approaches, Petrick et al. (2010a) investigate how they should be used over time to obtain dynamic control mechanisms under independent demand. Gönsch et al. (2014) pursue this further and find that the deterministic linear programming approximation fails to capture the revenue generated from delaying resource allocation by using flexible products. They propose to use the opportunity cost to obtain a dynamic booking limit policy for general flexible products. Koch et al. (2017) take it a step further by developing a dynamic programming approach for the network revenue management problem with flexible products under customer choice behaviour (which naturally leads to dynamic control policies).

Most studies on flexible products assume that the seller defines a set of potential execution modes of a flexible product. In contrast to this, Mang et al. (2012) investigate a flexible product for which customers self-select the level of flexibility.

In summary, we can build on a growing body of literature on demand management and vehicle routing for attended home delivery, as well as on flexible products. These two concepts have not yet been combined, and indeed the results from the flexible products literature do not carry over directly because future expected vehicle routing implications need to be taken into account.

2.3 The Dynamic Pricing Model for Delivery Time Slots

In this section, we first present the dynamic time slot pricing problem for attended home delivery services and then formulate the problem as a dynamic program (involving both standard and flexible time slots) under a model of customer delivery slot choice (namely the nested multinomial logit model).

We consider an e-grocer having a fixed number of homogeneous trucks, each with capacity c in terms of homogeneous transport totes. The e-grocer provides delivery services to customers located in non-overlapping areas $a \in \mathcal{A}$ for one fixed delivery day. The delivery slots can be offered from a set of non-overlapping standard slots \mathcal{S} , each of the same duration (say, one hour). Flexible slots can be offered from a set \mathcal{M} . Each flexible slot m has a certain set of standard slots $\mathcal{S}_m \subseteq \mathcal{S}$ associated with it that the delivery can be assigned by the retailer. No restriction is imposed on constructing flexible slots. For brevity, we define $\mathcal{F} = \mathcal{S} \cup \mathcal{M}$.

The dynamic slot pricing problem is modelled by a discrete dynamic program. The problem has T stages denoted by $t = 1, \dots, T$ corresponding to the time periods in the booking horizon. The final period T denotes the cut-off time after which no more bookings are accepted. We assume that the time periods chosen are sufficiently small such that the probability of more than one request arrival per period is negligible. Customers are classified into segments $n \in \mathbf{N}$ based on their slot choice behaviour. We assume that an order from a segment- n customer is (on average) worth r_n in profit before delivery costs, and that each order consumes one unit of truck capacity. Let λ represent the probability of a customer arrival in any given time period (the arrival probabilities are assumed to be independent of time only to simplify notation; notice that we can always reduce time-heterogeneous arrival rates to a uniform rate by manipulating the underpinning discrete time grid). Given an arrival, μ_a is the likelihood that the requested delivery address is in area a , and η_{an} is the probability that the request is from customer segment n conditional on there being a request from area a .

In stage t , the state of the system is defined by a matrix of accepted orders $\mathbf{x} \in \mathbb{N}^{|\mathcal{A}| \times |\mathcal{F}|}$, and its component x_{as} indicates the number of orders that have been accepted for delivery in time slot s for area a until time t . At every stage t , given state \mathbf{x} , we need to make pricing decisions for all feasible delivery time slots when delivery services can be provided. In line with common business practice, we assume prices are chosen from a finite set of potential price points $\mathcal{D} = \{d_\kappa : \kappa \in \mathcal{K} = \{0, 1, \dots, K\}\}$, where d_0 denotes the null price that drives demand to zero. We need d_0 to model unavailability of a slot. We also assume that a flexible slot is never higher priced than any feasible standard slot.

Given accepted orders \mathbf{x} , the set of slots in which we can feasibly schedule a delivery in area a is denoted by $\mathcal{F}_a(\mathbf{x})$ consisting of feasible standard slots $\mathcal{S}_a(\mathbf{x})$ and feasible flexible slots $\mathcal{M}_a(\mathbf{x})$. We assume that a flexible slot is feasible as long as it involves at least one feasible standard slot. In practice, this is not an unrealistic assumption; e.g. Tesco is indeed offering their flexible slots even when some of the one-hour regular slots are marked as unavailable. We make this assumption

because we are using a very conservative way of checking feasibility. All feasible slots $s \in \mathcal{F}_a(\mathbf{x})$ will be offered at any stage and state (so we do not consider strategically making certain feasible slots unavailable) because we assume that the retailer wants to maximise the number of available options to improve customer satisfaction.

Let us introduce $\mathbf{g} \in \{0, 1\}^{|\mathcal{A}| \times |\mathcal{F}| \times |\mathcal{K}|}$ where $g_{as\kappa} = 1$ represents assignment of price point d_κ to a feasible slot s for any order received from area a , and $g_{as0} = 1$ indicates the assignment of the null price d_0 to a slot s in area a (which only happens when s is infeasible due to our assumption that all feasible slots are always to be offered to increase customer satisfaction). The action space at state \mathbf{x} is defined as

$$\mathbf{G}(\mathbf{x}) := \{\mathbf{g} \mid \mathbf{d}_{am}^T \mathbf{g}_{am} \leq \mathbf{d}_{as}^T \mathbf{g}_{as}, \forall m \in \mathcal{M}_a(\mathbf{x}), s \in \mathcal{S}_a(\mathbf{x}), a \in \mathcal{A}; \\ \sum_{\kappa \in \mathcal{K} \setminus \{0\}} g_{as\kappa} = 1, \forall s \in \mathcal{F}_a(\mathbf{x}), a \in \mathcal{A}; g_{as0} = 1 \forall s \notin \mathcal{F}_a(\mathbf{x}), a \in \mathcal{A}\}.$$

Let $C(\mathbf{x})$ denote the minimum cost of delivering orders \mathbf{x} ; this minimum cost is the outcome of solving a capacitated vehicle routing problem with multiple time windows. We set $C(\mathbf{x}) = \infty$ when there is no feasible solution for the set of orders \mathbf{x} .

The transition probability to a new state in the next time period is defined by the probability of a customer arrival combined with the customer's slot selection probability. Customers are faced with more delivery time uncertainty with flexible slots than standard slots when booking their deliveries. Therefore, customers may naturally partition presented time slots into standard ones and flexible ones when selecting slots. Accordingly, we use the nested multinomial logit (nested MNL) model to express the customers' slot selection probability where the nests indexed by (\hat{s}, \hat{m}) are defined in terms of flexible and standard slots, respectively.

Let $U_{s\kappa}^n$ denote the utility of booking slot s at price d_κ for a customer from segment n . We compute this utility as $U_{s\kappa}^n = u_{s\kappa}^n + \epsilon_s$ where $u_{s\kappa}^n$ represents a fixed linear predictor function and ϵ_s is a random variable generated from a Gumbel distribution with zero mean. Not booking any time slot is associated with utility u_0^n for a segment- n customer. Let $\omega_{\hat{s}}$ and $\omega_{\hat{m}}$ denote the dissimilarity parameters of standard slot nest \hat{s} and flexible slot nest \hat{m} , respectively. The dissimilarity parameter of no-purchase behaviour is set to 1. We assume that customers from the same segment have the same price sensitivity towards standard and flexible slots. Note that dissimilarity parameters represent the degree of independence in unobserved utility among flexible slots and among standard slots, respectively.

Given the prices identified by $\mathbf{g} \in \mathbf{G}(\mathbf{x})$ in area a , the selection probability of

standard delivery slot s by a segment- n customer is computed as

$$p_{ans}(\mathbf{g}) = \frac{\mathbf{v}_{ns}^T \mathbf{g}_{as} (\sum_{i \in \mathcal{S}_a(\mathbf{x})} \mathbf{v}_{ni}^T \mathbf{g}_{ai})^{\omega_{\hat{s}} - 1}}{(\sum_{i \in \mathcal{S}_a(\mathbf{x})} \mathbf{v}_{ni}^T \mathbf{g}_{ai})^{\omega_{\hat{s}}} + (\sum_{i \in \mathcal{M}_a(\mathbf{x})} \mathbf{v}_{ni}^T \mathbf{g}_{ai})^{\omega_{\hat{m}}} + v_{n0}}, \quad (2.1)$$

where $\mathbf{v}_{ns} = \{v_{ns\kappa} = \exp(u_{ns\kappa}/\omega_{\hat{s}}) \mid \forall \kappa \in \mathcal{K}\}$ for standard slot $s \in \mathcal{S}_a(\mathbf{x})$, $\mathbf{v}_{nm} = \{v_{nm\kappa} = \exp(u_{nm\kappa}/\omega_{\hat{m}}) \mid \forall \kappa \in \mathcal{K}\}$ for flexible slot $m \in \mathcal{M}_a(\mathbf{x})$ and $\mathbf{g}_{as} = \{g_{as\kappa} \mid \forall \kappa \in \mathcal{K}\}$. Note that $v_{ns\kappa}$ is interpreted as the preference weight of slot s priced at d_κ for a segment- n customer. If a customer from area a books slot s , the state \mathbf{x} transitions to state $\mathbf{x} + \mathbf{1}_{as}$. The reader is referred to Train (2003) for further information on the characteristics of nested MNL.

Let λ define probability of customer arrival and μ_a represent probability of that customer coming from area a . Given λ and μ_a , we can define η_{an} as conditional probability of the customer's order from segment n and area a . We can now introduce the value function $V_t(\mathbf{x})$ at state \mathbf{x} in terms of future value functions $V_{t+1}(\mathbf{x})$ as a maximisation problem for action \mathbf{g} ;

$$V_t(\mathbf{x}) = \max_{\mathbf{g} \in \mathbf{G}(\mathbf{x})} (1 - \lambda)V_{t+1}(\mathbf{x}) + \sum_{a \in \mathbf{A}, n \in \mathbf{N}} \lambda \mu_a \eta_{an} \left(\sum_{s \in \mathcal{F}_a(\mathbf{x})} p_{ans}(\mathbf{g}) (r_n + \mathbf{d}^T \mathbf{g}_{as} + V_{t+1}(\mathbf{x} + \mathbf{1}_{as})) + p_{an0}(\mathbf{g}) V_{t+1}(\mathbf{x}) \right). \quad (2.2)$$

By substituting $p_{an0}(\mathbf{g}) = 1 - \sum_{s \in \mathcal{F}_a(\mathbf{x})} p_{ans}(\mathbf{g})$ in (2.2) for a customer's order received from segment n in area a , we can rewrite the value function at state \mathbf{x} as follows:

$$V_t(\mathbf{x}) = \max_{\mathbf{g} \in \mathbf{G}(\mathbf{x})} \sum_{a \in \mathbf{A}, n \in \mathbf{N}} \lambda \mu_a \eta_{an} \left(\sum_{s \in \mathcal{F}_a(\mathbf{x})} p_{ans}(\mathbf{g}) [r_n + \mathbf{d}^T \mathbf{g}_{as} - (V_{t+1}(\mathbf{x}) - V_{t+1}(\mathbf{x} + \mathbf{1}_{as}))] \right) + V_{t+1}(\mathbf{x}). \quad (2.3)$$

Once the booking horizon is finished (i.e., after the cut-off time T), the delivery of accepted orders (\mathbf{x}) takes place. Since the company is concerned with the net profit after delivery cost, the boundary condition at stage $T + 1$ is given by

$$V_{T+1}(\mathbf{x}) = -C(\mathbf{x}). \quad (2.4)$$

The dynamic program (2.3)-(2.4) is intractable because of its large state space. Moreover, computing $C(\mathbf{x})$ in the model is NP-hard since it involves solving a capacitated vehicle routing problem with time windows (Savelsbergh, 1985). Whilst

we cannot solve it directly, it is still useful as it motivates the shape of a pricing policy. If we had at least an approximation of the opportunity cost for an order in time slot s in area a as $\Delta_{as}^t(\mathbf{x}) \approx V_{t+1}(\mathbf{x}) - V_{t+1}(\mathbf{x} + \mathbf{1}_{as})$, we should obtain price \mathbf{g} by solving:

$$\arg \max_{\mathbf{g}_{as} \in \mathbf{G}(\mathbf{x})} \sum_{s \in \mathcal{F}_a(\mathbf{x})} p_{ans}(\mathbf{g}) [r_n + \mathbf{d}^T \mathbf{g}_{as} - \Delta_{as}^t(\mathbf{x})]. \quad (2.5)$$

This problem represents the so-called online decision problem: given state \mathbf{x} at time t , set of feasible delivery slots $\mathcal{F}_a(\mathbf{x})$ and opportunity costs $\Delta_{as}^t(\mathbf{x})$, we need to obtain the price points for all feasible delivery slots within a very short time period (within a few hundred milliseconds as advised by an industry representative). Thus, an efficient solution of (2.5) is crucial and depends to a great extent on the structure underpinning the choice model. Under the nested MNL, this pricing problem is difficult to be solved; however, general attraction models (including MNL as a special case) have strong structural properties that can be exploited in obtaining tractable optimisation routines. Accordingly, we propose to fit an MNL model to the data even though a nested MNL model is a better representative of the actual choice behaviour. As our numerical experiments demonstrate, this can lead to good results even though the simulated customer decisions follow a nested MNL model. We discuss this further in Section 2.4 with the estimated MNL choice model. In Section 2.5, we introduce an approximation to the value function at each state that we use to calculate the opportunity costs (again exploiting the structure of the estimated MNL model to obtain efficient formulations).

2.4 Pricing Policy under MNL Choice Model

The online pricing policy (2.5) should ideally be solved every time a new booking request arrives. Bookings typically arrive in large volume within a relatively short amount of time, and therefore a swift solution method is required to ensure acceptable runtime performance of the online booking system. This section presents an approach of using an estimated MNL choice model in the online pricing policy such that the policy can be equivalently reformulated as a compact linear program.

Given the historical booking data that is generated under an assumption of customers choosing time slots according to a nested MNL model, we can obtain an estimated MNL choice model which approximates the customer choice behaviour under the nested MNL choice model. The reader is referred to Yang et al. (2014) for further information about estimation of MNL choice model parameters from transaction data. Let $\hat{u}_{s\kappa}^n$ denote the utility of booking time slot s at price d_κ for a segment- n customer in the estimated MNL model. Not booking any time slot

has the utility \hat{u}_0^n which is normalised to 1. The selection probability of delivery time slot s by a segment- n customer under prices identified by $\mathbf{g} \in \mathbf{G}(\mathbf{x})$ in area a is computed as follows

$$\hat{p}_{ans}(\mathbf{g}) = \frac{\hat{\mathbf{v}}_{ns}^T \mathbf{g}_{as}}{\sum_{j \in \mathcal{F}_a(\mathbf{x})} \hat{\mathbf{v}}_{nj}^T \mathbf{g}_{aj} + 1}, \quad (2.6)$$

where $\hat{\mathbf{v}}_{ns} = \{\hat{v}_{ns\kappa} = \exp(\hat{u}_{ns\kappa}) \mid \forall \kappa \in \mathcal{K}\}$.

We consider the slot pricing problem (2.5) for a given state \mathbf{x} at time t , customer from area a , order value r and set of feasible slots $\mathcal{F}(\mathbf{x})$, and opportunity cost estimates Δ_s for all $s \in \mathcal{F}(\mathbf{x})$. Let us drop index a to reduce notational clutter. In order to simplify (2.5), we want to linearise the objective function and reformulate the constraints in a way such that the associated coefficient matrix is totally unimodular. This allows us to solve the combinatorial problem exactly as a linear program, and thus offers great advantages in solution speed.




Figure 2.1: Network structure for flexible slot m and a slot $s \in \mathcal{S}$.

Let us first consider the constraints on admissible prices: prices should be chosen from the discrete set \mathcal{D} , and each flexible slot should always be priced no higher than any standard slots since it is an inferior offering (we call the latter price dominance constraints). To formulate the price dominance constraints in a tractable fashion, we draw on a modelling approach of Davis et al. (2013): they model such price dominance constraints as a unit flow problem on a network because this results in a totally unimodular constraint structure which allows us to relax the integer requirements. To do this, we define a network flow problem for each

combination (m, s) of a feasible flexible slot m and one of the feasible standard slots $s \in \mathcal{S}(\mathbf{x})$. There is one source node with unit supply, and one sink with unit demand. Furthermore, we have a node for each combination of m with a price point κ , and likewise for s and each price point κ . The nodes are connected by arcs, as illustrated in Figure 2.1. Recall that the price points are ordered in increasing value in κ . The flow on some arcs corresponds to pricing variables $g_{m\kappa}$ and $g_{s\kappa}$, and on others we have new variables $z_{j\kappa}$ where $z_{j\kappa} = 1$ if time slot j is priced at d_κ or higher; and 0 otherwise. Enforcing the balance constraints at each node of this network for binary variables g and z ensures that the price for flexible slot m must be less than or equal to the price of slot s . By defining such a network for all (m, s) , $m \in \mathcal{M}(\mathbf{x})$, $s \in \mathcal{S}(\mathbf{x})$, we obtain the required constraints to satisfy price dominance with a totally unimodular constraint matrix. The resulting non-linear formulation R_{NLP}^a for a given area a can be stated as follows:

$$\begin{aligned}
R_{\text{NLP}}^a : \max_{\mathbf{g}, \mathbf{z}} \sum_{n \in \mathbf{N}} \mu_n & \frac{\sum_{s \in \mathcal{F}(\mathbf{x})} \sum_{\kappa \in \mathcal{K}} (r_n - \Delta_s^t + d_\kappa) \hat{v}_{ns\kappa} g_{s\kappa}}{1 + \sum_{s \in \mathcal{F}(\mathbf{x})} \hat{v}_{ns}^T g_s} \\
\text{s.t. } g_{m1} + z_{m1} &= 1, \quad \forall m \in \mathcal{M}(\mathbf{x}), \\
g_{m\kappa} + z_{m\kappa} &= z_{m, \kappa-1}, \quad \forall m \in \mathcal{M}(\mathbf{x}), \kappa \in \mathcal{K} \setminus \{0, 1, K\}, \\
g_{mK} &= z_{m, K-1}, \quad \forall m \in \mathcal{M}(\mathbf{x}), \\
g_{m1} &= g_{s1} + z_{s1}, \quad \forall m \in \mathcal{M}(\mathbf{x}), s \in \mathcal{S}(\mathbf{x}), \\
g_{m\kappa} + z_{s, \kappa-1} &= g_{s\kappa} + z_{s\kappa}, \\
&\quad \forall m \in \mathcal{M}(\mathbf{x}), s \in \mathcal{S}(\mathbf{x}), \kappa \in \mathcal{K} \setminus \{0, 1, K\}, \\
g_{mK} + z_{s, K-1} &= g_{sK}, \quad \forall m \in \mathcal{M}(\mathbf{x}), s \in \mathcal{S}(\mathbf{x}), \\
g_{j\kappa} &\in [0, 1], \quad z_{j\kappa} \in [0, 1], \quad \forall j, \kappa \in \mathcal{K} \setminus \{0\}.
\end{aligned} \tag{2.7}$$

The first three groups of constraints state the flow balance at nodes (m, κ) and the next three those of (s, κ) . Note that we can drop the binary restrictions since the constraint coefficient matrix of R_{NLP}^a is totally unimodular. This latter property follows from the fact that the balance constraints of a unit network flow problem with single source and sink satisfy the totally unimodular condition (Nemhauser and Wolsey, 1988).

Proposition 1 (Reformulation of Slot and Price Assortment Problem) R_{NLP}^a is equivalent to the online pricing problem (2.5).

Proof The proof of this proposition is analogous to the argument provided in Davis et al. (2013). \square

Next, we linearise the optimisation problem R_{NLP}^a by introducing the following

decision variables for $n \in \mathbf{N}$, $m \in \mathcal{M}(\mathbf{x})$, $s \in \mathcal{F}(\mathbf{x})$ and $\kappa \in \mathcal{K}$

$$\hat{g}_{ns\kappa} = \frac{\hat{v}_{ns\kappa} g_{s\kappa}}{1 + \sum_{j \in \mathcal{F}(\mathbf{x})} \hat{v}_{nj}^T g_j} \quad \text{and} \quad \hat{z}_{nm\kappa} = \frac{z_{m\kappa}}{1 + \sum_{j \in \mathcal{F}(\mathbf{x})} \hat{v}_{nj}^T g_j}.$$

The linear optimisation model R_{LP}^a for area a can be formulated as follows:

$$\begin{aligned} R_{LP}^a : \max_{\hat{\mathbf{g}}, \hat{\mathbf{z}}} & \sum_{n \in \mathbf{N}} \eta_n \sum_{s \in \mathcal{F}(\mathbf{x})} \sum_{\kappa \in \mathcal{K}} (r_n - \Delta_s^t + d_\kappa) \hat{g}_{ns\kappa} \\ \text{s.t.} & \sum_{s \in \mathcal{F}(\mathbf{x}), \kappa \in \mathcal{K}} \hat{g}_{ns\kappa} + \hat{g}_{n0} = 1, \quad \forall n \in \mathbf{N}, \\ & \frac{\hat{g}_{nm1}}{\hat{v}_{nm1}} + \hat{z}_{nm1} = \hat{g}_{n0}, \quad \forall n \in \mathbf{N}, m \in \mathcal{M}(\mathbf{x}), \\ & \frac{\hat{g}_{nm\kappa}}{\hat{v}_{nm\kappa}} + \hat{z}_{nm\kappa} = \hat{z}_{nm, \kappa-1}, \quad \forall n \in \mathbf{N}, m \in \mathcal{M}(\mathbf{x}), \kappa = 2, \dots, K-1, \\ & \frac{\hat{g}_{nmK}}{\hat{v}_{nmK}} = \hat{z}_{nm, K-1}, \quad \forall n \in \mathbf{N}, m \in \mathcal{M}(\mathbf{x}), \\ & \frac{\hat{g}_{nm1}}{\hat{v}_{nm1}} = \frac{\hat{g}_{ns1}}{\hat{v}_{ns1}} + \hat{z}_{ns1}, \quad \forall n \in \mathbf{N}, m \in \mathcal{M}(\mathbf{x}), s \in \mathcal{S}(\mathbf{x}) \\ & \frac{\hat{g}_{nm\kappa}}{\hat{v}_{nm\kappa}} + \hat{z}_{ns, \kappa-1} = \frac{\hat{g}_{ns\kappa}}{\hat{v}_{ns\kappa}} + \hat{z}_{ns\kappa}, \\ & \quad \forall n \in \mathbf{N}, m \in \mathcal{M}(\mathbf{x}), s \in \mathcal{S}(\mathbf{x}), \kappa = 2, \dots, K-1, \\ & \frac{\hat{g}_{nmK}}{\hat{v}_{nmK}} + \hat{z}_{ns, K-1} = \frac{\hat{g}_{nsK}}{\hat{v}_{nsK}}, \quad \forall n \in \mathbf{N}, m \in \mathcal{M}(\mathbf{x}), s \in \mathcal{S}(\mathbf{x}), \\ & 0 \leq \hat{\mathbf{g}}, \hat{\mathbf{z}} \leq 1. \end{aligned} \tag{2.8}$$

The price of slot s that is indicated by the optimal solution $\hat{g}_{nas\kappa}^*$ can be obtained by solving R_{LP}^a . Specifically, it will be priced at d_κ only if $\hat{g}_{nas\kappa}^*$ is non-zero. Note that only one $\hat{g}_{nas\kappa}$ is non-zero among all $\kappa \in \mathcal{K}$ for slot $s \in \mathcal{F}(\mathbf{x})$.

Proposition 2 (Linearisation of Slot and Price Assortment Problem) *Both R_{NLP}^a and R_{LP}^a problems are equivalent and possess the same optimal value.*

Proof The proof is provided in Appendix A.1. \square

Notice that the opportunity cost $\Delta_{as}^t = V_{t+1}(\mathbf{x}) - V_{t+1}(\mathbf{x} + \mathbf{1}_{as})$ for each customer's order coming from area a needs to be estimated for all available slots s at time t . Then, it becomes an input to R_{LP}^a to determine prices of available time slots. As mentioned earlier, it is crucial that the approximation approach must be computationally efficient to cope with the large-scale problems. The next section introduces an approach by adopting a choice-based linear program to estimate related value functions under the estimated MNL choice model.

2.5 A Model-based Opportunity Cost Approximation

For simplicity, consider value function $V_t(\mathbf{x})$ in (2.2) for a given state \mathbf{x} at time t . We first need to compute total delivery cost to be used for the approximated value function $\hat{V}_t(\mathbf{x})$. Let us introduce binary decision variables $\mathbf{w} = \{w_{ams} \mid \forall a \in \mathcal{A}, m \in \mathcal{M}, s \in \mathcal{S}_m\}$ representing whether the accepted orders for flexible slot m are assigned to their corresponding standard slots s or not. Moreover, let \mathcal{M}_s define a set of flexible slots covering standard slot s . The number of orders x'_{as} from area a to be delivered during time slot s is calculated as $x'_{as} = x_{as} + \sum_{m \in \mathcal{M}_s} w_{ams}$. We apply for the continuous half-width routing method introduced by Daganzo (1987) in order to estimate the total delivery cost. Note that the customer slot selection behaviour is captured by the estimated MNL choice model such that we can obtain a tractable linear model for the opportunity cost approximation. A brief description of this method follows.

We assume that only one vehicle is sent to each area; we emphasise that this assumption does not extend to the solution of the full vehicle routing problem at the end of a booking horizon. Instead, this assumption is only used in the opportunity cost approximation where we anyway do not yet have full information on all orders, but where we need to include forecasted orders so as to arrive at reasonable opportunity cost estimates for all slots. Each vehicle has capacity c and delivers customers' orders within a pre-defined rectangular area (with length α_a and width β_a) during the time window of the standard slot. These rectangles can be of different sizes reflecting different densities of customer locations; this approach has been proposed by Yang and Strauss (2017). If the average mile of any vehicle per hour is ν and the average service time per order takes $\bar{\tau}$, then the maximum number of feasible orders B_a to be delivered within an area a for any standard slot s can be calculated as

$$B_a = \frac{t_0 - \frac{2\alpha_a}{\nu}}{\bar{\tau} + \frac{\beta_a}{6\nu}},$$

where t_0 denotes the time window for each standard slot. Given a transportation cost of δ per mile, the delivery cost of orders x'_{as} from area $a \in \mathcal{A}$ and standard slot $s \in \mathcal{S}$ is computed as

$$C_{as}(x'_{as}) = \delta \left(2\alpha_a L(x'_{as}) + \frac{\beta_a}{6} x'_{as} \right),$$

where the function $L(x'_{as})$ is defined as

$$L(x'_{as}) = \begin{cases} 0 & \text{if } x'_{as} = 0, \\ 1 & \text{if } 0 < x'_{as} \leq B_a, \\ \infty & \text{if } x'_{as} > B_a. \end{cases}$$

Note that the cost $\delta\rho_a$ of travelling from the depot to area a (with ρ_a miles distance) is independent from the customer orders, but still contributes to the total delivery cost. We should also mention that these delivery-cost estimations are not used for constructing the final delivery routes.

We are now ready to present the approximated linear programming model for the dynamic pricing model. Let $\mathbf{g}' = \{g'_{ias\kappa} \mid \forall i \in \{t, \dots, T\}, a \in \mathcal{A}, s \in \mathcal{F}, \kappa \in \mathcal{K}\}$ represent pricing decisions made from time t until the end of planning horizon T . The choice probability of a segment- n customer for slot s with price d_κ is computed as

$$p'_{ians\kappa}(\mathbf{g}') = \frac{\hat{v}_{ns\kappa} g'_{ias\kappa}}{\sum_{j \in \mathcal{F}} \hat{\mathbf{v}}_{nj}^T \mathbf{g}'_{iaj} + 1}, \quad (2.9)$$

where $p'_{ian0}(\mathbf{g}')$ denotes the likelihood of not booking any time slot. Due to notational simplifications, let us define $P_{ians\kappa}(\mathbf{g}') = \lambda\mu_a\eta_{an}p'_{ians\kappa}(\mathbf{g}')$ to denote the probability of a segment- n customer from area a selecting slot s with price d_κ at time i , and accordingly $P_{ian0}(\mathbf{g}') = \lambda\mu_a\eta_{an}p'_{ian0}(\mathbf{g}')$ represents the probability of not booking from a segment- n customer in area a . Then, the expected number of orders to be allocated in standard slot s from area a becomes $\left(x_{as} + \sum_{i=t}^T \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} P_{ians\kappa}(\mathbf{g}')\right)$.

Given the average profit-before-delivery (\bar{r}_{an}) received from the accepted order of segment- n customer from area a , we can formulate a non-linear program

(NLP) as follows:

$$\begin{aligned}
(NLP) : \quad \hat{V}_t(\mathbf{x}) = \max_{\mathbf{g}', \mathbf{w}} \quad & \sum_{i=t}^T \sum_{a \in \mathcal{A}, s \in \mathcal{S}} \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} P_{ians\kappa}(\mathbf{g}') (\bar{r}_{an} + d_\kappa) - \sum_{a \in \mathcal{A}, s \in \mathcal{S}} C_{as}(x'_{as}) \\
\text{s.t.} \quad & \sum_{s \in \mathcal{F}} \left[x_{as} + \sum_{i=t}^T \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} P_{ians\kappa}(\mathbf{g}') \right] \leq c, \quad \forall a \in \mathcal{A}, \\
& \sum_{s \in \mathcal{S}_m} w_{ams} = x_{am} + \sum_{i=t}^T \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} P_{ianm\kappa}(\mathbf{g}'), \quad \forall a \in \mathcal{A}, m \in \mathcal{M}, \\
& x_{as} + \sum_{i=t}^T \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} P_{ians\kappa}(\mathbf{g}') + \sum_{m \in \mathcal{M}_s} w_{ams} \leq B_a, \quad \forall s \in \mathcal{S}, a \in \mathcal{A}, \\
& \sum_{i=t}^T \sum_{\kappa \in \mathcal{K}} \frac{P_{ians\kappa}(\mathbf{g}')}{v_{ns\kappa}} = \sum_{i=t}^T P_{ian0}(\mathbf{g}'), \quad \forall a \in \mathcal{A}, s \in \mathcal{F}, n \in \mathbf{N}. \\
& \mathbf{g}' \in [0, 1], \quad \mathbf{w} \geq 0.
\end{aligned} \tag{2.10}$$

The first set of constraints in (2.10) ensures that the capacity of vehicles serving in each area is not exceeded. The second group of constraints expresses the balance equations for allocating orders in flexible slots to standard slots while the third set of constraints are time-window constraints for all standard slots after allocating orders from flexible slots. The final set of constraints enforces to have a single price for each time slot requested by a segment- n customer in any area at each time period.

(NLP) is a difficult optimisation problem that involves the nonlinear choice probability terms. We can decompose the problem in terms of areas via our routing cost approximation using independent delivery areas. We build on the ideas of Gallego et al. (2014) to reformulate (NLP) as a compact linear program. Let $\mathbf{y} = \{y_{asn\kappa} \mid \forall a \in \mathcal{A}, s \in \mathcal{F}, \kappa \in \mathcal{K}\}$ denote new decision variables where $y_{asn\kappa}$ represents the expected number of segment- n customers in area a to select time slot s with price d_κ . Using

$$y_{asn\kappa} = \sum_{i=t}^T P_{ians\kappa}(\mathbf{g}') \text{ and } y_{an0\kappa} = \sum_{i=t}^T P_{ian0}(\mathbf{g}'),$$

we obtain the model (LP) below. Note that the meaning of variables x'_{as} and w_{ams} remains the same. The constraints of the (NLP) model can be easily transformed into the time-aggregated form as presented in (LP).

$$\begin{aligned}
(LP) : \quad R_t(\mathbf{x}) = \max_{\mathbf{y}, \mathbf{w}} \quad & \sum_{a \in \mathcal{A}, s \in \mathcal{F}} \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} (\bar{r}_{an} + d_\kappa) y_{ans\kappa} - \sum_{a \in \mathcal{A}, s \in \mathcal{S}} C_{as}(x'_{as}), \\
\text{s.t.} \quad & \sum_{s \in \mathcal{F}} \left[x_{as} + \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} y_{ans\kappa} \right] \leq c, \quad \forall a \in \mathcal{A}, \\
& \sum_{s \in \mathcal{S}_m} w_{ams} = x_{am} + \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} y_{anm\kappa}, \quad \forall m \in \mathcal{M}, a \in \mathcal{A}, \\
& x_{as} + \sum_{m \in \mathcal{M}_s} w_{ams} + \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} y_{ans\kappa} \leq B_a, \quad \forall s \in \mathcal{S}, a \in \mathcal{A}, \\
& \sum_{\kappa \in \mathcal{K}} \frac{y_{ans\kappa}}{v_{ns\kappa}} \leq y_{an0}, \quad \forall a \in \mathcal{A}, s \in \mathcal{F}, n \in \mathbf{N}, \\
& \sum_{a \in \mathcal{A}, s \in \mathcal{F}} \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} y_{ans\kappa} + y_{an0} = \lambda(T - t + 1) \\
& \sum_{\kappa \in \mathcal{K}} d_\kappa \frac{y_{anm\kappa}}{v_{nm\kappa}} \leq \sum_{\kappa \in \mathcal{K}} d_\kappa \frac{y_{ans\kappa}}{v_{ns\kappa}}, \quad \forall a \in \mathcal{A}, m \in \mathcal{M}, s \in \mathcal{S}, n \in \mathbf{N}. \\
& \mathbf{y} \geq 0, \mathbf{w} \geq 0.
\end{aligned} \tag{2.11}$$

Additionally, we impose a condition to ensure that total number of customers' bookings over all time slots including no-bookings during the remaining time periods must be equal to the expected number of arrivals. Finally, we have to make sure that any flexible slot is not assigned with a higher price than any standard slots.

Proposition 3 (Linearisation of Value Function Estimation) *If the MNL model is considered to describe customer choice behaviour, then (LP) is equivalent to (NLP). Thus, $R_t(\mathbf{x}) = \hat{V}_t(\mathbf{x})$.*

Proof The proof of this proposition is provided in Appendix A.2. \square

2.6 Computational Experiments

The central question that we seek to answer with the numerical studies in this section is to what extent, and under what conditions may flexible slots be able to improve profitability? Furthermore, we are interested in quantifying where potential improvements are coming from. Are we saving on routing costs, or attracting more revenue? How are the results affected by varying ratios of demand to capacity? We begin by describing and justifying the scenarios to be analysed, then report our results and discuss insights and limitations.

2.6.1 Data and Experimental Design

In our experiments, the delivery day has 14 one-hour non-overlapping standard slots. We focus on a single customer segment and define the utility of booking slot s with price d_κ as $u_{s\kappa} = u_s + \gamma d_\kappa$ where u_s is the utility of the slot and γ indicates the price sensitivity. Table 2.1 presents those standard and flexible slots along with their utility parameters defined under a nested MNL model. We construct 7 flexible slots that can be offered to customers as presented in Table 2.2 along with their utility parameters under the nested MNL model. Note that the utility parameter of each flexible slot is set as the average utility of its covered standard slots.

We define two scenarios (abbreviated as P3 and A4) for the design of flexible slots as shown in Table 2.2 in order to test their impact on various performance measures. Scenario P3 features 3 flexible slots (hence ‘3’) covering non-adjacent standard slots according to their popularity (hence the ‘P’). More specifically, each flexible slot covers one popular standard slot and two less popular slots. Scenario A4 provides four flexible slots (hence the ‘4’) each covering adjacent (hence the ‘A’) standard slots. This is similar to current industry practice by Tesco in the UK as they exclusively offer flexible slots consisting of adjacent standard slots. Recall that we are interested in the effect of introducing flexible slots versus not having flexible slots with a nested MNL choice model as the underlying ground truth choice model.

As benchmark decision policy, we use dynamic pricing of all 14 standard slots without the ability to offer flexible slots. We report on the performance of being able to use P3 or A4 relative to this benchmark. Note that we also tested a pricing policy based on the nested MNL model. However, it has not significantly improved on the performance measures compared to using the pricing policy with MNL model (and in some cases even performed worse) and consequently we do not report the corresponding results. It may seem counter-intuitive that using the correct choice model in the policy actually may be worse than using an approximated one; we believe that this effect arises from the fact that our opportunity cost estimation is biased due to the use of MNL in its calculation.

Table 2.1: Utility parameters of standard slots under the nested MNL model

Slot	8–9	9–10	10–11	11–12	12–13	13–14	14–15
u_s	3.2	3.1	3.3	3.2	3.0	2.5	2.7
Slot	15–16	16–17	17–18	18–19	19–20	20–21	21–22
u_s	3.5	3.8	3.9	3.6	4.7	4.2	3.2

$u_0 = 3.5$; $\gamma = -0.45$. Dissimilarity parameters: $\omega_{\hat{s}} = 0.8$ and $\omega_0 = 1$.

Table 2.2: Specification of flexible slots and their utility parameters

Time	P3			A4			
Slot	m_1	m_2	m_3	m_4	m_5	m_6	m_7
8–9							
9–10							
10–11							
11–12							
12–13							
13–14							
14–15							
15–16							
16–17							
17–18							
18–19							
19–20							
20–21							
21–22							
u_m	3.6	3.7	3.3	3.5	2.7	3.7	4.2

Dissimilarity parameters: $\omega_{\hat{m}} = 0.5$

To estimate the MNL-based choice models to be used to construct our decision policies in the simulation studies, we generate booking histories involving 320,000 booking requests. We randomly generate for each request a set of standard and flexible slots to represent the historic set of offered alternatives. Each standard time slot has the probability of 70% to be included in the offer set and each flexible slot (either of slots in P3 or A4) is offered when at least one of its covered standards slots is offered. Half of the 320,000 synthetic offer sets were constructed involving flexible slots sampled from P3 and A4, respectively. The price of each offered slot is randomly selected from the set $\{\mathcal{L}4, \mathcal{L}5, \mathcal{L}6, \mathcal{L}7, \mathcal{L}8\}$ and flexible slots have prices no higher than standard slots. Specifically, we firstly randomly select prices for standard slots from the set. Then, we randomly pick prices for flexible slots from a subset consisting of only price points that are lower or equal to the lowest prices of standard slots.

We simulate each customer slot selection decision based on offered slots and their prices following the nested MNL model in Tables 2.1 and 2.2 (but this model is not known to our decision policy). Based on our generated booking histories, we use the `asclogit` package provided in Stata/SE 15 to estimate the parameters of the MNL choice model in Table 2.3 which are used in the opportunity cost estimation.

We focus on a delivery area with the size $15\text{km} \times 15\text{km}$ with a depot located

Table 2.3: Utility parameters in the estimated MNL model

Slot	8–9	9–10	10–11	11–12	12–13	13–14	14–15
\hat{u}_s	-0.3766	-0.1027	0.2605	-0.2616	-0.73287	-1.3731	-1.1300
Slot	15–16	16–17	17–18	18–19	19–20	20–21	21–22
\hat{u}_s	-0.0983	0.2494	0.1509	0.6240	0.7288	0.4977	-0.0984
Flexible	m_1	m_2	m_3	m_4	m_5	m_6	m_7
\hat{u}_m	1.0681	0.1282	-1.8204	-0.2369	-2.3910	-0.4260	0.5592

Note: $\hat{u}_0 = 0$; $\hat{\gamma} = -0.5507$.

outside the area at (7km, 16km). The area can be equally divided into 25 sub-areas and customers are evenly distributed within each sub-area. We create a pool of customer locations where 75% of customers are located in the 15 shaded sub-areas and 25% of customers located in the 10 white areas as illustrated in Figure 2.2. This simple design mimics the situation of a grocery retailer dispatching from a single depot in the outskirts of a city. Customer orders are always of unit size, and their order profit r before delivery cost (excluding delivery charges) is drawn from a normal distribution (truncated at zero) with mean 25 and standard deviation of 10.




Figure 2.2: Delivery service area provided with flexible slots

We use the number of time periods covered in one sales horizon to reflect the demand level. We assume that exactly one booking request appears at every time period. Based on the estimated capacity level considering the time window constraints and vehicle capacity, we choose the base demand level with 1800 time periods where the ratio of expected demand to capacity is 1. We apply scaling parameters from the set $\{0.6, \dots, 1.7\}$ to the demand level to evaluate the performance of our policies under different demand levels. For example, we consider 1980 time periods when the scaling parameter is 1.1.

Each simulation of the sales horizon iterates over all time periods sampling customer slot booking decisions. Slot feasibility checks during the simulation run

are performed based on the continuous delivery cost approximation (we stress that this feasibility check is rather conservative). A standard slot is feasible if the vehicle has available capacity and if the number of accepted orders does not exceed B_a . A flexible slot is feasible if and only if at least one of its covered standard slots is deemed feasible. During the booking horizon, opportunity costs are estimated using the approach described in Section 2.5 and updated after every 100 customer acceptances. When the scaling parameter is 1.0, opportunity costs are re-optimised 12 times during one booking horizon. Slot price points \mathcal{D} range from £2 to £8 in incremental steps of £1.

At the end of the booking horizon, we calculate the delivery costs $C(x)$ by solving a vehicle routing problem with multiple time windows (‘multiple’ because a flexible slot can be composed by multiple feasible standard slots for a single customer). The company has 25 delivery vans, each with capacity of $c = 100$ units. Travel distance between any two adjacent orders’ locations is measured by the Euclidean distance metric. We multiply the total travel distance with a fuel cost of £10 per kilometre to obtain the total delivery costs (we ignore fixed costs). A van travels with a fixed speed of 25km per hour. The service time for each order is 10 minutes.

We apply the simulated annealing approach of Belhaiza et al. (2014) to minimise the total delivery costs. Starting from an initial delivery route, we make iterative improvement steps by randomly reassigning orders within a route. Note that we modify the cost function and the approach proposed by Belhaiza et al. (2014) to evaluate each route by only considering the total delivery costs and penalties from violating the vehicle capacity and time window constraints. Since we use an existing method for constructing routes (and subsequently evaluating costs), we refrain from re-producing the exact algorithm.

The simulation of each scenario runs 100 times and returns average performance measures on the number of accepted orders, revenue, total delivery costs and total profit. Revenue consists of order revenue and delivery charges. Note that the total profit is computed by subtracting total delivery costs from the total order and delivery charge revenue.

2.6.2 Numerical Results and Analysis

In our experiments, we aim to measure the value of introducing flexible slots, and to derive insights on what drives this value. The scenarios A4 and P3 only differ in the definition of flexible products. We report the computational results in terms of performance measure determined as percentage change relative to the base scenario of having only standard slots (no flexible products). All other parameters remain the same across all scenarios.

First of all, we are concerned with determining the total profitability impact of P3 versus having no flexible slots in dependence of the demand scaling parameter. Total increase of profit corresponds to the product of relative increase in order volume and relative increase in mean profit, both depicted in Figure 2.3. We plot them on top of each other since their sum approximately corresponds to the total profit increase, thus giving an impression of what drives profit. In all scenarios, adding the three flexible slots increases total profit by at least 3%. Moreover, we can see where these profitability increases are coming from: for low demand scenarios, it is through both attracting more orders and higher profit per order; for high demand scenarios, it is mainly from higher profit per order. Intuitively, what happens under low demand is that more orders are attracted by offering flexible slots to fill up available capacity. Delivery capacity stays constant throughout all experiments, meaning that the larger the scaling parameter, the more congested the delivery routes become. Therefore, we cannot attract many more orders under high demand. However, we may be able to attract more high-value orders.




Figure 2.3: Profit increase (%) under P3 relative to never using flexible slots

Let us drill down further on the increase of profit per order with the intention of unearthing insights on what causes these profitability improvements. When we consider the profit per order, we observe that flexible slots significantly increase efficiency regardless of demand levels. We break the profit per order increases further down into percentage changes in revenue per order and percentage changes in cost per order as shown in Figure 2.4. Apparently, the main drivers of profit per order

increases are routing cost savings at low demand, which demonstrates the value of added flexibility in route planning. Under low demand, routes have relatively few orders to serve, and the ability to move some customers can reduce the length of routes considerably. Accordingly, significant delivery cost savings can be achieved. Based on Figure 2.5, fleet utilisation improves for low demand scenarios. Note that the utilisation is defined as total number of served orders of the vehicle. Under high demand, efficiency of deliveries cannot be improved much because the routes are too congested. We conclude that introducing flexible slots tends to be particularly cost efficient when delivery capacity is large relative to demand.

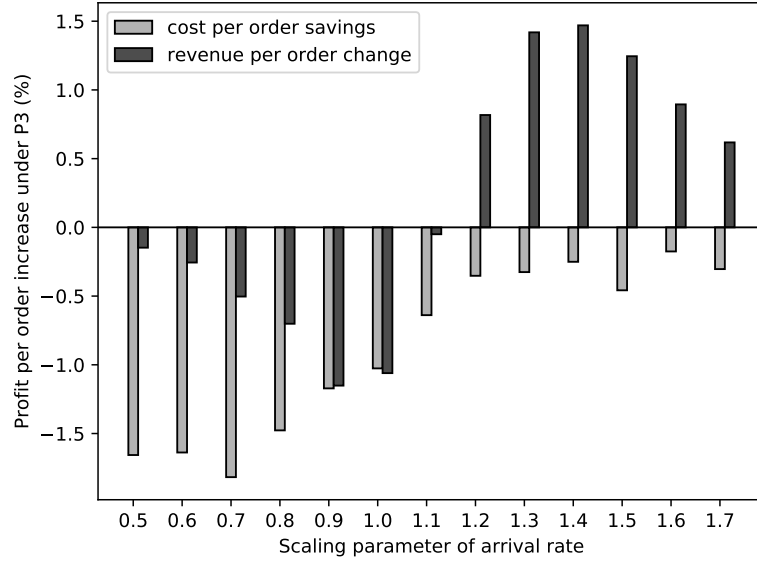


Figure 2.4: Revenue/order and cost/order increases (%) under P3 relative to never using flexible slots




Figure 2.5: Vehicle utilisation under different demand patterns under P3

On the other hand, the revenue per order decreases in low demand scenarios when we introduce flexible slots. Note that the revenue per order consists of the revenue from the order itself and the delivery service revenue. When demand is low, the available capacity needs to be filled up by attracting more customers with low priced delivery services. Since flexible slots cannot be priced higher than standard slots, the delivery service revenue per order decreases in the low demand scenarios, which results in an overall decrease of revenue per order. When demand is high, the policy focuses on attracting high-value orders and both standard and flexible slots are higher priced resulting in increased revenue per order.

Another interesting question is how to design flexible slots to gain more benefit in reducing delivery costs, i.e., should we group adjacent standard slots together to essentially plan for wider time windows within which the customer is ultimately assigned to a specific standard slot as in scenario A4, or to combine some popular with less popular slots as in scenario P3? We compare these scenarios, A4 versus P3, to obtain some insights to that end. As shown in Figure 2.6, P3 performs significantly better on profitability in almost all scenarios. When the scaling parameter is small, the increase is mostly driven by additional order intake.




Figure 2.6: Performance comparison of scenarios (A4 vs P3) relative to using only standard slots

Figure 2.7 compares the changes in cost per order and revenue per order in A4 and P3, respectively, relative to the scenario where no flexible slots are provided. Revenue per order is also highly influenced by the underlying demand level, so we concentrate on the change in cost per order. Both A4 and P3 can reduce cost per order. At low demand levels, P3 results in more cost per order reduction than A4. It makes sense since with P3 we can move customers from popular to less popular slots so as to accommodate more orders on the delivery routes. Therefore, we can conclude that providing wider time slots as flexible slots is not as efficient as combining popular slots with less popular slots.

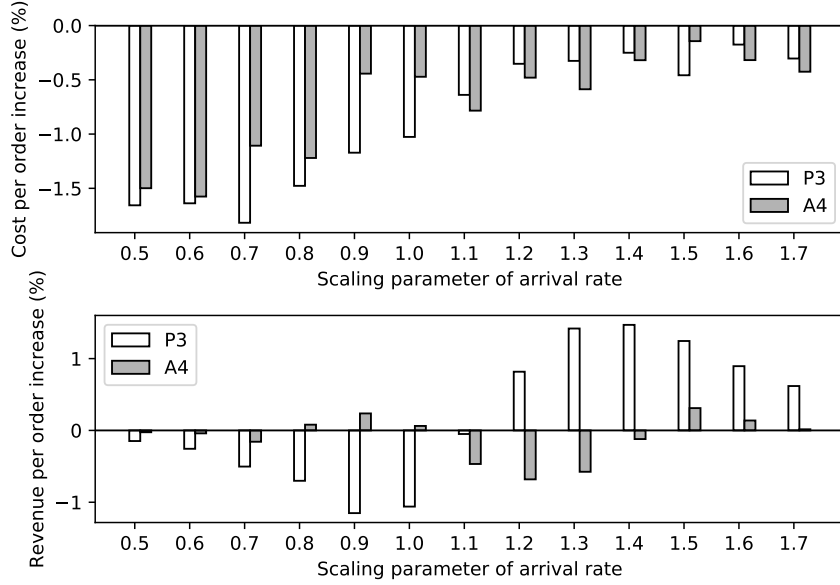



Figure 2.7: Cost per order and revenue per order increases (%) under scenarios (A4 vs P3) relative to using only standard slots

Finally, another interesting practical question is whether customers would typically be assigned to the same mode of a flexible slot (which means that customers' valuations would be affected due to learning effects). To that end, let us classify each standard slot covered by flexible slots as either 'least popular', 'median popular' or 'popular', based on its relative popularity compared to other standard slots covered in the same flexible slots. Figure 2.8 presents final slot allocation decisions for orders in all flexible slots offered in P3. These time slot assignments were made by the routing algorithm at the end of each simulation run; thus, they are not driven by our assumptions on customers' utility parameters of flexible slots. Since popular slots are more congested than other slots, orders in flexible slots are mostly allocated to 'least popular' and 'median popular' slots. As demand is scaled up, the proportion of orders allocated to popular slots increases. When demand is low, our pricing policy tries to retain orders by offering slots at low prices such that a substantial number of customers book into popular standard slots straight away. It results in less delivery (routing) capacity left for accommodating orders in flexible slots. When demand is high, higher prices are charged for slots by our pricing policy, especially for these popular standard slots. It reduces the number of customers who book directly into standard slots but pick to select cheaper flexible slots. Accordingly, more delivery capacity may be found within those popular standard slots such that increases the likelihood of allocating orders in flexible slots to popular standard slots.

Figure 2.9 demonstrates final slot allocation decisions for orders in all flexible slots offered under A4. We can observe similar patterns as under P3 that the likelihood of allocating flexible orders to popular standard slots increases as demand level increases, apart from for flexible slot m_5 . Flexible slot m_5 includes three less popular standard slots (relative to others on the delivery day). When the demand is low, only a small number of customers books directly into those standard slots and we tend to have delivery capacity left even in popular slots. Therefore, most orders are allocated to the popular standard slot such that the cost per order can be reduced. Moreover, we also observe that orders in flexible slots are more evenly allocated to standard slots under A4 than under P3. Hence, we conclude that a customer would find it harder to anticipate in which slot their order will be executed if flexible slots are simply constructed by combining adjacent standard slots as compared to P3.



Note: popular slots are 17-18, 19-20, 20-21; median popular slots are 9-10, 16-17; least popular slots are 12-13, 13-14, 14-15

Figure 2.8: Final allocation for flexible orders under P3



Note: popular slots are 10-11, 12-13, 17-18, 19-20; median popular slots are 11-12, 14-15, 16-17, 20-21; least popular slots are 9-10, 13-14, 15-16, 18-19

Figure 2.9: Final allocation for flexible orders under A4

2.7 Conclusions

In this chapter, we propose a dynamic pricing approach for standard and flexible time slots for attended home delivery. Flexible slots have recently been introduced by a major retailer in the UK in the form of time windows that encompass four hours; customers who choose such a slot are guaranteed to receive delivery in a one-hour slot within this wider time window. Which slot exactly is communicated only shortly prior to the delivery day. Our method can dynamically price such constructs alongside regular narrow time slots under consideration of the customer choice. The approach is based on tractable linear programming formulations and, as such, is scalable to real-life applications.

Several managerial insights have been obtained via a simulation study. First, flexible slots have significant potential in reducing delivery costs, especially when demand is low relative to available delivery capacity. Moreover, we find that retailers would be better off to construct flexible slots as combinations of some more and some less popular slots, as opposed to the current industry practice of using adjacent slots only. Especially, if demand is high relative to available delivery capacity such

flexible slots have the advantage of being able to spread customers more equally across the delivery time slots.

Because of a lack of real data, it is worthy of noting that all results and analysis are derived by assuming that the utility of a flexible time slot is the average utility of the standard slots that it covers. Modifying such assumption only affects the customer choice behaviour. In other words, we may not be able to attract additional customers or obtain more profit-before-delivery by changing the underlying assumption. However, having flexible slots provide the e-retailer with flexibility in scheduling delivery routes. Therefore, we can still observe significant delivery cost savings from offering flexible slots under any assumptions. On the other hand, customers could have low preference towards flexible slots because of the delivery time uncertainty. Accordingly, the e-retailer would have to reduce prices of those slots such that the lost revenue could not be offset by the cost benefit from delivery flexibility. In other words, the overall profit of the e-retailer might not be improved after introducing flexible slots.

A limitation of the proposed approach is that the solution approach makes use of a rather crude approximation of the capacitated vehicle routing problem with multiple time windows, which forms the boundary condition for the dynamic pricing problem. Nevertheless, as the simulation study shows, flexible products still bring significant routing cost savings (the latter being estimated using an established heuristic taken from the existing literature). More refined approximations may improve results, but at the risk of losing scalability. An interesting future research question is how should these flexible slots be best designed, i.e., which regular slots should be combined to form a flexible slot? Furthermore, how could we price flexible slots when we allow customers to design the flexible slot themselves, i.e., if they can freely combine regular slots to form a custom-made flexible slot?

Chapter 3

Dynamic Delivery Time Slots Management and Vehicle Routing across Multiple Days

An e-retailer may offer customers more delivery options from a number of consecutive delivery days. For example, leading UK e-grocers, such as Tesco, Morrisons and Sainsbury's, allow customers to select delivery time slots from 7 consecutive calendar days. This may cause the e-retailer an operational burden from decision making process and vehicle routing construction while it increases customer satisfaction. In this case, customers compare available time slots and their delivery days before they choose one of those time slots for delivery. The e-retailer adopts a certain control policy to influence a customer's choice on time slots such that the total profit of each delivery day is maximised. As a result, it is essential for the e-retailer to establish an appropriate choice model capturing the underlying customer's choice behaviour.

In this chapter, we focus on the dynamic time slot management for multiple consecutive delivery days. We propose three customer choice models to capture the underlying customer choice behaviour and derive three slotting policies accordingly. Specifically, we introduce two choices models to take into account the substitution effect of delivery days by defining customer choice behaviour under a multinomial logit choice model and a nested multinomial logit choice model. We also consider an alternative model that ignores such effect. In this case, we apply an independent multinomial logit choice model to each individual day. Our numerical results counter-intuitively indicate that using the choice model excluding substitution effect of delivery days in the slotting policy may generate higher profit than the other two choice models, especially when the demand is low relative to the delivery capacity. However, integrating such a choice model in the slotting policy does not improve

the delivery efficiency by reducing delivery cost per order. Moreover, our results show that offering same-day delivery services could benefit the e-retailer in terms of attracting more orders and reducing delivery cost per order.

3.1 Delivery Service Booking System with Multiple Days

The UK’s online retail market has grown by 16% and reached £60billion during 2017 (Mintel, 2018d). The home delivery services offered by e-retailers as the main fulfilment method aim to make customers’ online shopping experience become fast and convenient (Ehmke, 2012). Moreover, the delivery services play an important role on the future growth of online retailing. In order to boost sales, Mintel (2018d) also highlights the importance for an e-retailer to offer a wider range of delivery options due to more convenience to customers. In particular, offering time slots from different delivery days expands further delivery options for customers.

Moreover, allowing customers to select delivery days has been empirically found as an important attribute to increase consumer satisfaction in online retailing (Nguyen et al., 2019). When a customer requests a delivery service, it is beneficial for e-retailers to present a number of time slots from several consecutive delivery days at the same time because of flexibility in managing their operations. It has been indeed widely used by leading retailers in UK; for example, Tesco presents time slots in 7 days including the current day for every customer request. Apart from increasing customer satisfaction, e-retailers may also benefit from offering time slots across multiple delivery days to customers. As more time slots are managed at the same time, it creates more opportunities for the e-retailer to exploit the potential low-cost time slots to accommodate each delivery request. As a result, a routing plan with low delivery costs can be obtained by the e-retailer for each delivery day.

However, joint management of time slots across multiple delivery days jointly is a challenging task as the complexity from both customer choice behaviour and vehicle routing planning increases in comparison to single day delivery operations. When the customer is presented with time slots from multiple delivery days, his/her slot choice behaviour is affected not only by the presented time slots but also by the delivery days. Therefore, a choice model including substitution effect of multiple delivery days is required to capture the underlying customer behaviour in selecting a specific day and a time slot. Moreover, such a choice model may not lead to a linear program to decide slot availability, which can be solved efficiently in polynomial time. In addition to the choice model, routing plans on presented delivery days are also anticipated at the same time when considering the costs of accommodating one request. Note that the cost of delivery for each order arriving during a specific day is

obtained by solving a vehicle routing problem with time windows. Since the vehicle routing problem is NP hard, a large computational effort is required to anticipate routing plans for all presented delivery days at the same time.

Those challenging and practical issues in managing time slots from multiple days have been ignored in majority of studies in the literature; they rather consider that time slots from each delivery day are managed independently as shown in Figure 3.1. Independently managing slots from different delivery days largely reduces the complexity of the problem, especially on the vehicle routing side. It is worthwhile to address that if a customer does not select any slot from this delivery day, he/she leaves the system since there is no option offered in other delivery days. This is a strong assumption in those studies that eliminates a situation that the customer may prefer other days. In this case, the substitution effect of delivery days is ignored and consequently the customer choice is reflected inaccurately in the model.




Figure 3.1: Time slots from K delivery days presented to customers whose arrival process separated into K booking horizons

In this chapter, we focus on a rolling-horizon based dynamic slotting problem of an e-retailer presenting time slots from a number of consecutive delivery days. We assume that a stream of customers arrive within a fixed calendar day requesting for delivery services (i.e., they can choose any delivery day). The decision problem of the e-retailer occurs whenever a customer coming from a known location requests a delivery service. Specifically, the e-retailer needs to determine which time slots can feasibly accommodate the request and which time slots need to be displayed to the customer. In response, the customer either selects exactly one slot by comparing all presented time slots and their corresponding days or decides to leave the system. When all orders are received for the delivery day, a capacitated vehicle routing problem with time windows is solved to construct the delivery routes. By dynamically deciding the presented time slots, the objective of the e-retailer is to maximise the total profit of all managed delivery days.

We also consider offering same-day delivery services as express delivery services to a specific group of customers. These customers' requests can be quickly prepared and easily accommodated by the existing routing plan of the current day

without routing plan reconstruction. Offering same-day services in such a way meets customers' demand for express delivery services and also increases their satisfaction of shopping from the e-retailer. Additionally, extra profit-before-delivery can be generated for the current delivery day. Most importantly, the efficiency of the routing plans on the current day can be improved by serving more orders without causing significant changes on the existing routing plan. Average delivery costs of serving one request on the current day could be further reduced because of the improved delivery efficiency.

Our contributions on managing time slots across multiple days can be summarised as follows. Firstly, we propose a new perspective to manage time slots across delivery days by dropping the assumption that each delivery day has its corresponding booking horizon. We deal with the slot management for multiple delivery days as a rolling-horizon problem where one calendar-based booking horizon for all involved delivery days is considered. We propose a DP formulation to dynamically decide the availability of each slot from all delivery days whenever a customer request arrives. Additionally, we explore the opportunity of offering same-day delivery services along with accepting customers for future delivery days. Same-day delivery services are provided only when there is a capacity from the existing routing plan for the current delivery day.

As the customer selects one slot by comparing among offered time slots and delivery days, we propose different choice models to capture the underlying customer choice behaviour based on the nested multinomial logit (NMNL) model and the multinomial logit (MNL) model. They both include the substitution effect of delivery days. Alternatively, we use the MNL models for all delivery days independently by explicitly excluding the substitution effect. The slotting policies are derived based on the proposed choice models using LP-based approximation models. In the numerical experiments, we compare the profit achieved by the e-retailer under these policies. Our results show that using a choice model without considering the substitution effect attracts the highest number of customers if demand is low relative to the delivery capacity. Results also demonstrate that using such choice models (without substitution effect) leads to inefficient routing plans that increase the delivery cost per order. On the other hand, offering same-day delivery services generates extra profit-before-delivery and improves the efficiency of the routing plans on the current day by reducing the delivery cost per order.

The chapter is organised as follows. A literature review is presented in Section 3.2. The problem is described and formulated in Section 3.3. We present the vehicle routing problem in Section 3.4. Then, we describe the demand model and the slotting policy in Section 3.5 and Section 3.6, respectively. Our computational results are

demonstrated in Section 3.7. Conclusions are reported in Section 3.8.

3.2 Literature Review on Time Slot Management

The attended home delivery services can be classified into two groups (namely, the next-day and same-day delivery services) according to specific days when the actual delivery operations take place at the next-day or same-day, respectively. All studies in the literature related to attended home delivery systems deal with management of time slots offered in a specific delivery day. To the best of our knowledge, no study in the literature is concerned with management of multi-days delivery systems.

In this section, we will review studies related to the same-day and next-day delivery services within an attended home delivery system and also summarise different choice models used in the literature to capture customers' behaviour in selecting time slots. We also highlight differences between single-day and multi-days time slot management practices. For a review of the attended home delivery systems and their challenging revenue management practices, the reader is referred to Agatz et al., 2008 and Agatz et al., 2013.

Next-day Delivery Services: Delivery operations of the next-day services begin after the booking process is completed (e.g., Agatz, 2007). Campbell and Savelsbergh (2005) introduce a dynamic decision model that accepts or rejects delivery requests and determines the time slot to accommodate each accepted order. They assume delivery service demand to follow an independent process, not to be affected by the e-retailer's decisions. In their following work, Campbell and Savelsbergh (2006) explore an opportunity of using incentives (such as delivery charges) to influence customer demand for specific time slot such a way that the total delivery cost is minimised. They consider a linear customer choice model to reflect impact of incentives on the probability of a particular time slot being chosen.

The other stream of studies on next-day delivery services focuses on the vehicle routing and scheduling types of decision-making problems with an objective of minimising delivery cost whilst meeting service requirements (Agatz et al., 2010). Another objective is to maximise the number of accepted requests for delivery services by retaining feasible tours; for instance, see Cleophas and Ehmke (2014), Ehmke and Campbell (2014). All these studies under this stream mainly focus on accepting or rejecting delivery time-slot bookings for a single specific time slot. No choice model is considered to reflect customer preference towards other slots.

In the most recent research paper Mackert (2019), the customer's slot choices are assumed to be affected by control decisions on the offered time slots. Specifically, a linear program is developed for making dynamic time slot offering decisions based

on the underlying general attraction model. Moreover, the probability of any specific presented slot is determined from all available and unavailable slots.

The multinomial logit model (abbreviated as MNL) has been widely used in various applications as a customer choice model in managing time slots for one specific delivery day. For instance, Asdemir et al. (2009) consider a dynamic time-slot pricing model using fixed delivery capacity for each time slot. They explore structural properties of the dynamic programming model and solve it by the backward recursion approach using fixed delivery costs. Due to the curse of dimensionality in the state space as the number of time slots and delivery capacities increase, it is difficult to use recursion approach to solve the model for a realistic size of the practical applications. Yang et al. (2014) estimate the MNL choice model using a real data for the dynamic time slot pricing problem for a next-day delivery services. Although the real data involves multiple delivery days, they define the no-purchase event for a customer who is not picking any time slots from a specific delivery day while estimating the MNL model. Moreover, they estimate opportunity costs only based on marginal routing costs. In their following work, Yang and Strauss (2017) use an approximate dynamic programming (ADP) approach to take into account effects of both future revenue and routing cost when estimating the opportunity cost. Moreover, they adopt a continuous approximation model proposed by Daganzo (1987) in order to estimate total delivery cost. In order to improve the conservative delivery costs approximation obtained by the continuous approximation model, Lang et al. (2019) recently propose a simulation-based approach by generating samples of arrival processes to obtain anticipatory delivery routes. As an alternative approach to MNL, Klein et al. (2017) propose a general non-parametric rank-based choice model that takes into account impact of prices on utilities of time slot. Prices of time slots are determined by solving a mixed-integer linear program anticipating future demand and estimate delivery costs. As the underlying optimisation model includes large number of integer decisions, an application of this method in practice remains challenging.

Same-day Delivery Services: The simultaneity of booking process and delivery operation is the key feature of same-day delivery services. A majority of the literature under same-day delivery context focuses on the vehicle dispatch decisions under stochastic customer demand. Under the vehicle dispatch problems, customer requests need to be delivered within specified time windows. Klapp et al. (2016) focus on the dispatch decision made after every fixed time with a single vehicle, where customers and the depot are located on a line. As an extension, Klapp et al. (2018) further discuss the same dispatch problem by relaxing the assumption on delivery locations. Voccia et al. (2017) propose an analytic approach that decides dispatch operation for multiple vehicles by taking into account future requests and

delivery deadlines.

As managing same-day delivery service requests, Ulmer (2017) focuses on the dynamic pricing policy of delivery deadlines offered by the services. It adopts a binary customer choice model where a deadline with the highest utility is picked by the customer. The approximated dynamic programming method is used to estimate opportunity costs used in the pricing policy. Their numerical results show that the dynamic pricing policy provides a higher revenue than a fixed pricing strategy and a geographical pricing policy. For a real application, one needs to consider computational issues such as the curse dimensionality for the ADP approach and the accurate estimation of customer’s slot choice behaviour.

Our Approach: As highlighted above, time-slot management across multiple delivery days has not been yet studied in the literature. Indeed customers can select time slots from a number of consecutive days to place their requests in real-life. Comparing to those studies in the literature, our research is concerned with the delivery time-slot management and vehicle routing problem of a retailer providing attended home delivery services. In particular, we consider a choice-based dynamic slotting policy (with a fix delivery price) involving available time slots across multiple future days for both the next-day and same-day delivery services. As time slots allocate across several consecutive days, the influence of presented time slots from other delivery days is explicitly taken into account by the customer choice model. We introduce the nested multinomial logit (abbreviated as NMNL) model as the underlying true customer choice model where the nest is defined by the day. We also use the MNL choice model as an approximation to the underlying true customer choice model with the same definition of a no-purchase event for each delivery day independently and marginal routing costs to estimate the opportunity cost as in Yang et al. (2014).

3.3 The Dynamic Slot Assortment and Vehicle Routing Problem

In this section, we first describe the slot assortment optimisation and vehicle routing problem with multiple delivery days and then present its dynamic programming formulation. We consider an e-retailer having one depot and $|\mathcal{B}|$ number of homogeneous vehicles. Each vehicle is assumed to have a fixed capacity Q_{max} . Delivery operations start from the warehouse and each delivery operation needs to be completed within a certain duration D_{max} (i.e., a vehicle visits all customers in a tour and returns to the depot during the delivery duration). The e-retailer provides delivery services to customers located within a number of non-overlapping areas $a \in \mathcal{A}$. Customers

are classified into segments $n \in \mathcal{N}$ and customers in the same segment have the same slot selection behaviour. Note that customers can be segmented based on area, occupation and size of the family. Accordingly, we further assume that an order from a segment- n customer worth r_n in profit-before-delivery on average and that each order consumes one unit capacity in the vehicle.




Figure 3.2: An illustration of booking and planning horizons with time slots from multiple delivery days

We assume that the retailer receives customer orders during a booking horizon consisting of multiple delivery days. Let us consider K consecutive delivery days of whom all available time slots need to be managed at the same time. As depicted in Figure 3.2, $k = 1$ specifically represents the current delivery day while each of the remaining $K - 1$ days is referred to the future delivery day. Time slots during available delivery days are all presented to customers at the same time and the customer can pick any day based on his/her preference. Customer requests for the same-day delivery services, scheduled on the current day ($k = 1$), need to be all fulfilled while new requests for the current day as well as the future $K - 1$ delivery days (for $k = 2, \dots, K$) are accepted as long as the delivery capacity permits.

We assume that each delivery day possesses the same number of non-overlapping time slots. Let \mathcal{S}_k denote a set of feasible time slots available on delivery day k . We can now define a set $\mathcal{H} = \{s \mid s \in \mathcal{S}_k, k = 1, \dots, K\}$ which includes all time slots from K delivery days.

The dynamic delivery slot assortment and vehicle dispatching problem is a calendar based problem of the e-retailer; therefore, a discrete dynamic program is used to model the problem during one specific delivery day that is also referred to as the planning horizon. Each delivery day has its own planning horizon that is discretised into T decision stages, denoted by time periods $t = 1, \dots, T_s, T_{s+1}, \dots, T_h$ where T_s represents the cut-off time for accepting requests for same-day delivery services and T_h is the last time period of the delivery day to accept a request for other future delivery days. These time periods are assumed to be sufficiently small so that the probability of having more than one arrival during each time period can be neglected. Next, we will present the dynamic programming formulation of the

underlying problem for the current delivery day, $k = 1$. For other delivery days, we can proceed with the same methodology (in terms of the dynamic programming model) in a rolling horizon manner. This will be further explained in computational experiments section.

Before proceeding to the dynamic programming model of the problem, we introduce the necessary notations for state and action sets as well as system dynamics of the system that will be used throughout this study.

States: A state of the system at stage t consists of two components. The first component includes all yet-to-be-delivered orders during K delivery days and we denote it by a vector of $\mathbf{x} \in \mathbb{N}^{|\mathcal{A}| \times |\mathcal{H}|}$. The element x_{as} of \mathbf{x} indicates the number of accepted orders yet-to-be-delivered during time slot $s \in \mathcal{H}$ to area $a \in \mathcal{A}$. The second component involves information related to each vehicle's earliest available departure time period and is denoted by a vector $\mathbf{q} \in \mathbb{N}^{|\mathcal{B}|}$ that only considers availability of vehicles at the current delivery day. More specifically, q_b represents the future time period when vehicle b is available at the warehouse (depot).

Actions: The action space consists of delivery dispatching and slot assortment decisions. Given state (\mathbf{x}, \mathbf{q}) at time $t \leq T_s$ (i.e., before the cut-off time for accepting same-day delivery orders), the retailer needs to make dispatching decisions for all available vehicles and also determine slot availability for all feasible delivery time slots during K delivery days. Let $\mathbf{w} \in \mathbb{N}^{|\mathcal{A}| \times |\mathcal{H}|}$ denote dispatching decisions where its element w_{as} indicates the number of orders to be dispatched during time slot s to area a . Since orders accepted for the current delivery day ($k = 1$) need to be delivered to customers in area $a \in \mathcal{A}$ during time slot $s \in \mathcal{S}_1$, $w_{as} \geq 0$. On the other hand, for time slots $s \in \mathcal{S}_k$ during other delivery days $k = 2, \dots, K$, we can set $w_{as} = 0$.

Let $\mathcal{H}'_{at}(\mathbf{x}, \mathbf{q}) = \{s \mid s \in \mathcal{S}'_k, k = 1, \dots, K\} \subset \mathcal{H}$ denote a set of feasible time slots, which needs to be determined at any state (\mathbf{x}, \mathbf{q}) at stage t to accommodate customer's request coming from area $a \in \mathcal{A}$ before the e-retailer makes a slotting decision. The feasibility of each time slot is checked whether a new coming order violates the specific constraints related to time windows, vehicle capacity and duration of a tour. Time slots for same-day delivery services are only checked before the cut-off time $t \leq T_s$. We assume that no more delivery trips will be added into the existing routing plan. On the other hand, the feasibility of time slots for the future day delivery services needs to be checked at each time period $0 \leq t \leq T_h$ and a new tour could be created if no tour exists for the customer order. The process of feasibility checking for time slots during the current and future delivery days will be further discussed in Section 3.4. Let us introduce a binary decision variable g_{as} for each feasible time slot s and area a . If the feasible slot s is offered to the customer,

then $g_{as} = 1$. Otherwise, $g_{as} = 0$. We can then define the slotting decisions of the retailer as $\mathbf{g} \in \{0, 1\}^{|\mathcal{A}| \times |\mathcal{H}'|}$.

We can now define the action space, denoted as $\Psi_t(\mathbf{x}, \mathbf{q})$ at state (\mathbf{x}, \mathbf{q}) , in terms of slotting and dispatching decisions taken at $t \leq T_s$, before the cut-off time for same-day delivery services, as follows;

$$\begin{aligned} \Psi_t(\mathbf{x}, \mathbf{q}) := \{(\mathbf{g}, \mathbf{w}) \mid & \mathbf{g} \in \{0, 1\}^{|\mathcal{A}| \times |\mathcal{H}'|}; w_{as} \leq x_{as}, s \in \mathcal{S}'_1, a \in \mathcal{A}; \\ & w_{as} = 0, s \in \mathcal{S}'_k, k = 2, \dots, K, a \in \mathcal{A}; \mathbf{w} \in \mathcal{W}\} \end{aligned}$$

where \mathcal{W} involves a set of constraints of the vehicle routing problem. Similarly, for $T_s < t \leq T_h$, the action space at state (\mathbf{x}, \mathbf{q}) includes only the slotting decisions; therefore, we can write the action space as

$$\Psi_t(\mathbf{x}, \mathbf{q}) := \{\mathbf{g} \mid \mathbf{g} \in \{0, 1\}^{|\mathcal{A}| \times |\mathcal{H}'|}\}.$$

State Transitions: We define transition probability from state (\mathbf{x}, \mathbf{q}) at stage t to another state $(\mathbf{x}', \mathbf{q}')$ at next stage $t + 1$ in terms of the probabilities of customer arrivals and the customer's slot selection as follows. Let us consider feasible time slots $\mathbf{g} \in \Psi_t(\mathbf{x}, \mathbf{q})$ at time t to be presented to the customer. Let $p_{ans}(\mathbf{g})$ denote the probability of time slot s being selected by a segment- n customer from area a . Similarly, $p_{an0}(\mathbf{g})$ is the likelihood of not booking any time slot at time t . Since the customer requests arriving by T_s need to be dispatched due to same-day delivery services, state \mathbf{x} is transformed to new state as $\mathbf{x}' = \mathbf{x} - \mathbf{w} + \mathbf{1}_{as}$, if the customer from area a books time slot s . On the other hand, if the customer from area a books time slot s at time $t > T_s$, then the new state is obtained as $\mathbf{x}' = \mathbf{x} + \mathbf{1}_{as}$.

Next, we will explain how the earliest availability time of all vehicles is updated (i.e., state transformation from \mathbf{q} to \mathbf{q}') when the dispatching decisions are made. Let ι denote time period when the delivery operation starts at the current delivery day $k = 1$ such that $1 \leq \iota \leq T_s$. We also introduce τ_b to indicate duration of a tour to be completed by vehicle b on the basis of the existing routing plan to deliver orders. Note that if vehicle b is not included in the routing plan, this means that either the vehicle is not being dispatched yet or it is still in use, then $\tau_b = 0$. If the earliest available time of vehicle b is in the future ($q_b > t$) or the current time period ($q_b = t$), then vehicle b is either under use or at the depot, respectively. Depending on the arrival of requests, beginning of delivery operation and the vehicle's presence in the routing plan, each element q'_b in \mathbf{q}' is updated at stage t as follows.

For a customer arriving before the delivery operation starts ($t < \iota$), the earliest available time of vehicle b is updated as $q'_b = \iota$. On the other hand, if the customer arrives after the delivery operation starts ($t \geq \iota$), one of the following two

situations may arise.

- Case 1: When vehicle b is included in the routing plan ($\tau_b > 0$), we can update it as $q'_b = q_b + \tau_b$.
- Case 2: Suppose that vehicle b is not included in the routing plan ($\tau_b = 0$). If vehicle b is at the depot ($q_b = t$), then $q'_b = q_b + 1$. However, if vehicle b is under use ($q_b > t$), then $q'_b = q_b$.

Note that state \mathbf{q} is not updated after the cut-off time T_s since dispatching decision is considered until the cut-off time of accepting same-day delivery orders.

Dynamic Programming Model: Before presenting the value function at time t , we first introduce notation for probabilities of customer arrival, request coming from a customer segmentation and a time slot selection. Let λ denote the probability of a customer arrival within any time period; this is assumed to be homogeneous during the booking horizon. Given a customer arrival, we define ν_a as a likelihood of the request is from area a and also μ_{an} as the probability of the request is from segment n .

Given state (\mathbf{x}, \mathbf{q}) , the value function, denoted by $V_t(\mathbf{x}, \mathbf{q})$, at stage t for $t = 1, \dots, T_s, T_s + 1, \dots, T_h$ can be formulated as a profit maximisation problem for actions \mathbf{g} and \mathbf{w}

$$V_t(\mathbf{x}, \mathbf{q}) = \begin{cases} \max_{\mathbf{w}, \mathbf{g} \in \Psi_t(\mathbf{x}, \mathbf{q})} \lambda \sum_{a \in \mathcal{A}, n \in \mathcal{N}} \nu_a \mu_{an} \sum_{s \in \mathcal{H}'_{at}(\mathbf{x}, \mathbf{q})} p_{ans}(\mathbf{g}) (r_n + V_{t+1}(\mathbf{x} - \mathbf{w} + \mathbf{1}_{as}, \mathbf{q}')) \\ \quad + p_{an0}(\mathbf{g}) V_{t+1}(\mathbf{x} - \mathbf{w}, \mathbf{q}') + (1 - \lambda) V_{t+1}(\mathbf{x} - \mathbf{w}, \mathbf{q}') \\ \quad - C_1(\mathbf{w}, \mathbf{q}), \quad 1 \leq t \leq T_s \\ \\ \max_{\mathbf{g} \in \Psi_t(\mathbf{x}, \mathbf{q})} \lambda \sum_{a \in \mathcal{A}, n \in \mathcal{N}} \nu_a \mu_{an} \sum_{s \in \mathcal{H}'_{at}(\mathbf{x}, \mathbf{q})} p_{ans}(\mathbf{g}) (r_n + V_{t+1}(\mathbf{x} + \mathbf{1}_{as}, \mathbf{q})) \\ \quad + p_{an0}(\mathbf{g}) V_{t+1}(\mathbf{x}, \mathbf{q}) + (1 - \lambda) V_{t+1}(\mathbf{x}, \mathbf{q}), \quad T_s < t \leq T_h \end{cases} \quad (3.1)$$

Notice that at time t for $T_s < t \leq T_h$, we continue accepting orders for future $K - 1$ delivery days, but dispatching decisions are no longer considered until the end of planning horizon. After the cut-off time for same-day delivery services, the cost $C_1(\mathbf{w}, \mathbf{q})$ of delivery of orders \mathbf{w} to be dispatched on day $k = 1$ can be calculated in the cost-minimised manner by solving a capacitated vehicle routing problem with time windows. When no feasible routing solution for the set of orders \mathbf{w} is found, the delivery cost is set to be infinity; that is $C_1(\mathbf{w}, \mathbf{q}) = \infty$. At the end of planning horizon, we calculate the delivery costs of dispatching remaining orders on the current day. Let $C_k(\mathbf{x}, \mathbf{q})$ denote the minimum cost of delivering orders in day

$k = 1, 2, \dots, K$. Therefore, the boundary condition at stage $T_h + 1$ is

$$V_{T_h+1}(\mathbf{x}, \mathbf{q}) = - \sum_{k=1}^K C_k(\mathbf{x}, \mathbf{q}). \quad (3.2)$$

When no feasible routing solution for set of orders \mathbf{x} in delivery day k exists, we fix $C_k(\mathbf{x}, \mathbf{q}) = \infty$.

By substituting $p_{an0}(\mathbf{g}) = 1 - \sum_{s \in \mathcal{H}'_{at}(\mathbf{x}, \mathbf{q})} p_{ans}(\mathbf{g})$ in (3.1), we can rewrite the value function for $1 \leq t \leq T_h$ at given state (\mathbf{x}, \mathbf{q}) as

$$V_t(\mathbf{x}, \mathbf{q}) = \begin{cases} \max_{\mathbf{w}, \mathbf{g} \in \Psi_t(\mathbf{x}, \mathbf{q})} \lambda \sum_{a \in \mathcal{A}, n \in \mathcal{N}} \nu_a \mu_{an} \sum_{s \in \mathcal{H}'_{at}(\mathbf{x}, \mathbf{q})} p_{ans}(\mathbf{g}) (r_n - \Delta_{tas}^1(\mathbf{x}, \mathbf{q})) \\ \quad + V_{t+1}(\mathbf{x} - \mathbf{w}, \mathbf{q}') - C_1(\mathbf{w}, \mathbf{q}), & 1 \leq t \leq T_s \\ \max_{\mathbf{g} \in \Psi_t(\mathbf{x}, \mathbf{q})} \lambda \sum_{a \in \mathcal{A}, n \in \mathcal{N}} \nu_a \mu_{an} \sum_{s \in \mathcal{H}'_{at}(\mathbf{x}, \mathbf{q})} p_{ans}(\mathbf{g}) (r_n - \Delta_{tas}^2(\mathbf{x}, \mathbf{q})) \\ \quad + V_{t+1}(\mathbf{x}, \mathbf{q}), & T_s < t \leq T_h \end{cases} \quad (3.3)$$

with opportunity costs defined as $\Delta_{tas}^1(\mathbf{x}, \mathbf{q}) = V_{t+1}(\mathbf{x} - \mathbf{w}, \mathbf{q}') - V_{t+1}(\mathbf{x} - \mathbf{w} + \mathbf{1}_{as}, \mathbf{q}')$ and $\Delta_{tas}^2(\mathbf{x}, \mathbf{q}) = V_{t+1}(\mathbf{x}, \mathbf{q}) - V_{t+1}(\mathbf{x} + \mathbf{1}_{as}, \mathbf{q})$.

Due to the large state space, the dynamic program (3.3) is intractable. Additionally, calculating delivery costs ($C_1(\mathbf{w}, \mathbf{q})$ and $C_k(\mathbf{x}, \mathbf{q})$) at each state in the dynamic model is NP-hard because it involves solving a capacitated vehicle routing problem with time windows. Although the problem is hard to be solved directly, it provides some useful insights which motivate the slotting policy and the vehicle dispatching policy. At time t for $1 \leq t \leq T_s$, if we could make the dispatching decision \mathbf{w} independently before the slotting decision, and also estimate the opportunity cost $\hat{\Delta}_{tas}^1(\mathbf{x} - \mathbf{w}, \mathbf{q}') \approx V_{t+1}(\mathbf{x} - \mathbf{w}, \mathbf{q}') - V_{t+1}(\mathbf{x} - \mathbf{w} + \mathbf{1}_{as}, \mathbf{q}')$ of a customer from area a , then the slotting policy is obtained by simply solving the following model

$$\mathbf{g} = \arg \max_{\mathbf{g} \in \Psi_t(\mathbf{x}, \mathbf{q})} \lambda \sum_{a \in \mathcal{A}, n \in \mathcal{N}} \nu_a \mu_{an} \sum_{s \in \mathcal{H}'_{at}(\mathbf{x} - \mathbf{w}, \mathbf{q}')} p_{ans}(\mathbf{g}) [r_n - \hat{\Delta}_{tas}^1(\mathbf{x} - \mathbf{w}, \mathbf{q}')], \quad (3.4)$$

Notice that the action space is now simplified as $\Psi_t(\mathbf{x}, \mathbf{q}) := \{\mathbf{g} \mid \mathbf{g} \in \{0, 1\}^{|\mathcal{A}| \times |\mathcal{H}'|}\}$. On the other hand, if the customer order arrives after cut-off time of the same-day delivery services (that is $T_s < t \leq T_h$) and if we could just estimate the opportunity cost $\hat{\Delta}_{tas}^2(\mathbf{x}, \mathbf{q}) \approx V_{t+1}(\mathbf{x}, \mathbf{q}) - V_{t+1}(\mathbf{x} + \mathbf{1}_{as}, \mathbf{q})$, we should obtain the slotting policy

by solving

$$\mathbf{g} = \arg \max_{\mathbf{g} \in \Psi_t(\mathbf{x}, \mathbf{q})} \lambda \sum_{a \in \mathcal{A}, n \in \mathcal{N}} \nu_a \mu_{an} \sum_{s \in \mathcal{H}'_{at}(\mathbf{x}, \mathbf{q})} p_{ans}(\mathbf{g}) [r_n - \hat{\Delta}_{tas}^2(\mathbf{x}, \mathbf{q})]. \quad (3.5)$$

Thus, if we could separate the vehicle dispatching decisions from the time slotting policy at each time period, we could simplify the problem and obtain what time slots to offer based on the dispatching decision as presented in (3.4) and (3.5). Next, a brief summary of the decision making process follows. For any state (\mathbf{x}, \mathbf{q}) at stage $t \leq T_s$, we first decide orders \mathbf{w} to be delivered and compute its delivery costs associated with the decision. Then, for any request from area a , a set of feasible delivery slots $\mathcal{H}'_{at}(\mathbf{x} - \mathbf{w}, \mathbf{q}')$ needs to be identified on the basis of updated state and the approximate opportunity cost $\hat{\Delta}_{tas}^1(\mathbf{x} - \mathbf{w}, \mathbf{q}')$ is computed in order to decide slot availability by solving (3.4). On the other hand, for stage $t > T_s$, we only need to find a set of feasible delivery slots $\mathcal{H}'_{at}(\mathbf{x}, \mathbf{q})$ and approximate opportunity cost $\hat{\Delta}_{tas}^2(\mathbf{x}, \mathbf{q})$ to decide slot availability by solving (3.5). An efficient method to find \mathbf{w} is crucial for obtaining feasible slots and maximising total after-delivery-cost profit. We will further discuss this in Section 3.4. The customer choice probabilities will be presented in Section 3.5. Then, the slotting policy under a specified customer choice model along with the method to approximate opportunity costs will be covered in Section 3.6.

3.4 Construction of Delivery Routes

As mentioned earlier, delivery operations for each consecutive delivery day are managed independently, and the orders \mathbf{x}_k yet-to-be-delivered during a specific delivery day k can be determined at given state (\mathbf{x}, \mathbf{q}) . We assume that the availability of each vehicle and the lead time of same-day orders need to be taken into account if a delivery request is fulfilled on the current delivery day. Moreover, fulfilling a same-day delivery request does not increase the number of trips in the existing delivery plan. On the other hand, vehicle delivery routes are modified to accommodate same-day delivery services at the current day $k = 1$ and future delivery services taking place during $k = 2, \dots, K$. Recall that routes of current delivery day $k = 1$ need to be updated based on state $(\mathbf{x}_1, \mathbf{q})$ since dispatching decisions are also made for same-day delivery services whereas the routes of future delivery days $k = 2, \dots, K$ are modified based on only state \mathbf{x}_k since only slotting decisions are made for orders received for future delivery days.

In this section, we first introduce necessary notation for the vehicle routing problem with time windows (VRPTW) over multiple delivery days and then present

our approach to solve the problem. In order to construct and update delivery routes for the VRPTW, we adopt the algorithm introduced in Azi et al. (2012). The optimal routes are determined by using the simulated annealing method. Let us start describing key elements of the VRPTW along with notations.

Locations: We consider J number of customers who are located at different locations (represented by indices i and j) within an area. In particular, the depot where delivery operations begin and end is stated at location 0. Thus, we have a set $\mathcal{J} = \{0, 1, \dots, J\}$ consisting of all locations at any state. Let c_{ij} and δ_{ij} denote the cost and duration of travelling from location i to j , respectively.

Vehicles: As stated before, there are $|\mathcal{B}|$ number of homogeneous vehicles each of whom starts and finishes its delivery operation at times E_s and E_e , respectively, and involves the same amount of loading time $\hat{\delta}$.

Delivery Routes: Each vehicle can make several trips by following different routes during a delivery day. Let us assume that vehicle b can conduct M trips and σ_{bm} represent the m -th route (for $m = 1, \dots, M$) to be conducted by vehicle b . For notational convenience, we introduce 0^- and 0^+ to distinguish the beginning and finishing points of routes although they geographically coincide with the same location of the depot. For instance, if the vehicle follows the route m defined as $\sigma_{bm} = \langle 0^-, 1, \dots, I, 0^+ \rangle$, then it departs from the depot located at 0^- and visits all customers at locations j for $j = 1, \dots, I$ and then returns to the depot 0^+ . We use i^- and i^+ to specify predecessor and successor locations of location i to be visited in a delivery route, respectively.

Routing Plan: We define a routing plan at time t as a set of planned routes for all vehicles and is denoted by $\Omega_t = \{\Omega_{k,t}(b) : b \in \mathcal{B}, k = 1, \dots, K\}$. Notice that Ω_t consists of only to-be-delivered orders and needs to be updated at each time t . The total travel cost $C_k(\Omega_t)$ of routing plan Ω_t is calculated as

$$C_k(\Omega_t) = \sum_{b \in \mathcal{B}} \sum_{i \in \Omega_{k,t}(b)} c_{ii^+}, \text{ for } k = 1, \dots, K.$$

Pool of Routing Plans: Let Ω_t denote a pool of feasible routing plans at time t for all K delivery days and Ω_t^* be the cheapest routing plan which is actually used by the e-retailer. The routing plans in the pool Ω_t are maintained for the time slot feasibility check since it is possible that a new request cannot be feasibly inserted into Ω_t^* . The details of feasibility check is presented in the next section.

Model Constraints: We now describe model constraints to be included into the vehicle routing problem with time windows. Note that the following constraints are also considered in the feasibility check while accepting new customers.

- *Route Capacity:* We assume that each customer's request from location i

consumes η_i units of capacity and requires a service time $\bar{\delta}_i$. We introduce $Q(\sigma_{bm}) = \sum_{i \in \sigma_{bm}} \eta_i$ to denote the total capacity used by customers served on route σ_{bm} . We have to ensure that total capacity required for any route does not exceed the total available capacity, Q_{max} . Thus, a linear capacity constraint that is formulated as $Q(\sigma_{bm}) \leq Q_{max}$ is imposed for each route σ_{bm} conducted by vehicle b .

- *Time Windows:* Both customer orders and the depot have their own time windows. Let an interval $[e_j, l_j]$ denote a time window with the earliest and latest delivery times of customer from location j . Similarly, a time window of the depot is given by an interval $[E_s, E_e]$ with the earliest start time E_s and latest completion times E_e of all delivery operations during day k . Given the routing plan $\Omega_t(b)$ of vehicle b , a forward sweep is applied to obtain the earliest arrival time to the customer locations. Let z_i denote the earliest arrival time of vehicle to location i . We initialise the sweep by setting the earliest arrival time to start delivery operation from the depot as the earliest available time of vehicle b . Accordingly, the earliest arrival time for location i can be determined recursively as

$$z_i = \begin{cases} z_{i-} & \text{if } i^- = 0^+ \text{ and } i = 0^- \\ \max\{z_{i-} + \hat{\delta} + \delta_{i-i}, e_i\} & \text{if } i^- = 0^- \\ \max\{z_{i-} + \bar{\delta}_{i-} + \delta_{i-i}, e_i\} & \text{otherwise.} \end{cases} \quad (3.6)$$

This procedure is completed by finding the earliest completion time of all M routes completed by vehicle b , that is denoted by $z_{M,0+}$. Similarly, we need to determine the latest arrival time of vehicles to all customers. Let y_i denote the latest arrival time for location i . We can initialise it as the earliest completion time of all operations conducted by vehicle b , that is $z_{M,0+}$. Then, the latest arrival time for all customers $i = 0, 1, \dots, J$ is computed recursively by moving backwards as follows;

$$y_i = \begin{cases} y_{i+} & \text{if } i^+ = 0^- \text{ and } i = 0^+ \\ \min\{y_{i+} - \delta_{ii+} - \hat{\delta}, l_i\} & \text{if } i = 0^- \\ \min\{y_{i+} - \delta_{ii+} - \bar{\delta}_i, l_i\} & \text{otherwise.} \end{cases} \quad (3.7)$$

When reached to the beginning of route $\Omega_t(b)$, the latest starting time for vehicle from the depot is obtained. As a result, a time window restriction, $z_i \leq y_i$, for each customer request coming from location i needs to imposed.

Note that we record the latest arrival times $y_{m,0^-}$ for position 0^- in all routes $m = 1, \dots, M$ since it will be used for calculating route duration.

- *Route Duration:* The shortest time that is required by a vehicle to finish a route is called as route duration. Given route σ_{bm} from routing plan $\Omega_t(b)$, let ω_i denote the real arrival time based on the route plan for customer i . The dispatching time for route σ_{bm} is set by the latest arrival time as $\omega_{m,0^-} = y_{m,0^-}$. Then, the arrival time for customer i is computed recursively as

$$\omega_i = \begin{cases} \max\{\omega_{i^-} + \hat{\delta} + \delta_{i-i}, e_i\} & \text{if } i^- = 0^- \\ \max\{\omega_{i^-} + \bar{\delta}_{i^-} + \delta_{i-i}, e_i\} & \text{otherwise.} \end{cases} \quad (3.8)$$

Let $\tau(\sigma_{bm})$ denote the duration of the m -th route σ_{bm} conducted by vehicle b . The duration is calculated as the difference between arrival time to 0^+ and departure time from 0^- . We have to ensure that duration of the route cannot exceed the maximum delivery completion time. This can be formulated as a set of linear constraints: $\tau(\sigma_{bm}) \leq D_{max}$ for all routes v of vehicle b .

- *Order Availability:* This constraint is only applied for same-day delivery requests. The same-day delivery orders scheduled on one route must be available before the vehicle is dispatched. Let us introduce θ_i to represent the time when a same day-delivery for location i is available to be dispatched. This also includes the lead time of the order. We have to ensure that the latest departure time of the route involving this order is no less than θ_i .

VRPTW Across Multiple Delivery Days: Given a pool of feasible plans Ω_{kt} at time t , the VRPTW for multiple delivery days aims to determine a routing plan Ω_{kt}^* for the accepted orders to be delivered during a specific delivery day k such that the total delivery cost is minimised by satisfying the model constraints such as capacity, duration, time windows restrictions. The order availability also needs to be satisfied for the same-day delivery services. In order to find the optimal delivery plan, one can apply an enumeration method where total delivery costs of all feasible plans Ω_{kt} are computed and the routing plan Ω_{kt}^* with the minimum cost is determined. Note that this pool of routing plans for each delivery day is also used for the feasibility check. One slot becomes infeasible only when we cannot find feasible insertion position from all routing plans maintained in the pool.

A Solution Approach for VRPTW Across Multiple Delivery Days: When a new customer arrives at time t , we first apply an insertion heuristic to check whether the request can be inserted with the cheapest insertion cost into the current routing plan on day k . A new delivery route is added to the current routing plan if

the request cannot be feasibly inserted into the current plan. Otherwise, all routing plans $\Omega_{k,t}$ in the pool are also updated with the new request; this leads to $\Omega_{k,t+1}$. Note that if any routing plan in the pool cannot feasibly incorporate the request, then the request will be dropped out of the pool.

Given the updated routing plan $\Omega_{k,t+1}$, the structure of simulated annealing is then adopted to re-optimize the routing plan such that the total delivery cost $C_k(\Omega_{t+1})$ is minimised. In order to improve the initial routing plan $\Omega_{k,t+1}$ towards the cost minimisation, we apply *destruction and reconstruction procedures* until the stopping criteria is satisfied. A brief description of these procedures for VRPTW across multiple delivery days follows. At each iteration, we randomly select one level of destruction operation among workday, route and customer levels. All delivery locations in routing plan $\Omega_{k,t+1}(b)$ for randomly-selected vehicle b are removed from $\Omega_{k,t+1}$ in the workday-level destruction whereas locations within one randomly-selected route from $\Omega_{k,t+1}$ is removed in the route-level destruction. In the customer-level destruction, only one randomly-selected location is withdrawn from $\Omega_{k,t+1}$. Then, all these removed delivery locations are inserted back to the routing plan to obtain a new routing plan Ω'_{kt} with the new total delivery costs $C_k(\Omega'_{t+1})$. By the end of this improvement process, a routing plan $\Omega_{k,t+1}^*$ is obtained with minimal total delivery costs. The pseudocode of obtaining the optimal delivery routes is presented in Algorithm 1.

Algorithm 1 Simulated Annealing for optimal delivery plan

- 1: Initialise the starting routing plan Ω , the starting temperature $T = T_{max}$ and set the cooling rate r .
 - 2: Generate a new routing plan Ω' following *destruction and reconstruction procedures*.
 - 3: Compute change in delivery cost $\Delta = C(\Omega') - C(\Omega)$.
 - 4: **if** $\Delta \leq 0$ **then**
 - 5: Update the current routing plan as Ω' .
 - 6: **else**
 - 7: Accept Ω' as the current routing plan with a probability $e^{\frac{-\Delta}{T}}$.
 - 8: **end if**
 - 9: Set $T = rT$ and repeat from Step 2 until a stop criterion is met.
 - 10: **return** Current routing plan Ω .
-

It is important to note that, since the number of routes cannot be changed for the same-day delivery services, we can only apply customer-level and route-level destruction to improve the initial routing plan $\Omega_{1,t+1}$ towards the cost minimisation. As the feature of same-day delivery services, the routing plan and the pool of feasible plans are updated whenever a dispatch decision is made. More specifically, given a vehicle departs from the depot to deliver orders at time t , the corresponding route is

removed from the current routing plan Ω_{1t}^* to obtain $\Omega_{1,t+1}^*$. All requests involved in this dispatched route are removed from every feasible routing plan in the pool that is updated as $\Omega_{1,t+1}$.

3.5 Demand Model

In this section, we firstly discuss the assumption we make to conduct the study and present the underlying nested MNL model that customers follow when making the slot selection. Then we illustrate three estimation methodologies of the choice model used in our slotting policy.

Customer requests for the delivery services arrive over time according to a time-dependent Poisson process. Customers are able to see the time slot availability for K delivery days, including the current delivery day. It is also common in practice that customers may be able to view the slot availability/pricing of several days. We assume that the customer choice depends on alternative delivery days and the availability of slots in these days. Admitting the interaction between delivery days would move our study closer to the practice. Therefore, we assume that the underlying customer slot selection behaviour follows a nested MNL model.

Let us consider K consecutive delivery days which are defined as nests. Each delivery day has a set of time slots denoted by \mathcal{S}_1, \dots , and \mathcal{S}_K accordingly. We define the utility that a segment- n customer obtains from selecting slot s from day k as $u_{nks} + \epsilon_{nks}$ where u_{nks} describes the observable utility of slot s in delivery day k and ϵ_{nks} is a random variable generated from a Gumbel distribution with zero mean. Not picking any slot belongs to another nest indexed by $K + 1$ and has the observable utility u_0 . In addition, each delivery day k has a dissimilarity parameter γ_k presenting the degree of dissimilarity of slots that can be offered within this day. Note that the dissimilarity of no-purchase set is set as 1. For a segment- n customer, we set the preference weight of slot s in delivery day k as $v_{nks} = \exp(u_{nks}/\gamma_k)$. Further details on the nested MNL model can be found in Train (2003). Given slots $(\mathcal{S}'_1, \dots, \mathcal{S}'_K)$ we offer across K delivery days, the probability that a segment- n customer picks slot s in day k is computed as follows:

$$p_{nks}(\mathcal{S}'_1, \dots, \mathcal{S}'_K) = \frac{v_{nks}(\sum_{i \in \mathcal{S}'_k} v_{nik})^{\gamma_k - 1}}{v_0 + \sum_{j=1}^K (\sum_{i \in \mathcal{S}'_j} v_{nij})^{\gamma_j}}.$$

In order to determine the slotting policy by (3.4) and (3.5), we can use the nested MNL model as the choice model, however, we are not able to reach an equivalent slotting policy as a linear program. Therefore, we adopt the MNL model as an alternative choice model that approximates customer choice behaviour captured

by the nested MNL model. Such MNL model considers the day effect on the utility of time slots by assuming that all time slots across all delivery days are different from each other. Let \hat{u}_{ns} denote an estimated utility of the segment- n customer obtains from choosing slot s . Note that the utility \hat{u}_{ns} has captured the delivery day dependence in time slot s . Not booking any slot, denoted as \hat{u}_{n0} , is assumed to be 0 for all segments. The preference weight associated with slot s is defined as $\hat{v}_{ns} = \exp(\hat{u}_{ns})$. Accordingly, the probability of selecting slot s by a segment- n customer under slotting policy \mathbf{g} can be computed as

$$p_{ns}(\mathbf{g}) = \frac{\hat{v}_{ns}g_s}{\sum_{j \in \mathcal{H}'(\mathbf{x}-\mathbf{w}, \mathbf{q}')} \hat{v}_{nj}g_j + 1}. \quad (3.9)$$

The probability of not booking any slot is determined as follows

$$p_{n0}(\mathbf{g}) = \frac{1}{\sum_{j \in \mathcal{H}'(\mathbf{x}-\mathbf{w}, \mathbf{q}')} \hat{v}_{nj}g_j + 1}. \quad (3.10)$$

Next, we will use an example to explain how to obtain the MNL choice models that approximate the nested MNL model when time slots from two consecutive delivery days are offered to customers.

Example 3.5.1 Consider a case where a retailer offers customers with time slots from two consecutive delivery days and there are 12 non-overlapping one-hour time slots in each delivery day between 8am and 20pm. We assume that there is one customer segment.

In order to compute the MNL model approximating the choice behaviour under the nested MNL model, we first start obtaining the booking history by randomly generating a set of available time slots to be presented to each customer. Note that the retailer has recorded time slots that were offered to each delivery request in practice, despite of whether or not the customer chose a time slot. When generating a set of displayed time slots for one customer, we assume that each slot in Day 1 is presented with the probability of 0.2 and 0.8 otherwise. Each slot in Day 2 is displayed with a probability of 0.7 and 0.3 otherwise. The customer's decision is drawn from the probability distribution defined by the nested MNL model presented in Table 3.1. We repeat this process for 160,000 times (customers) to generate a stream of booking histories.

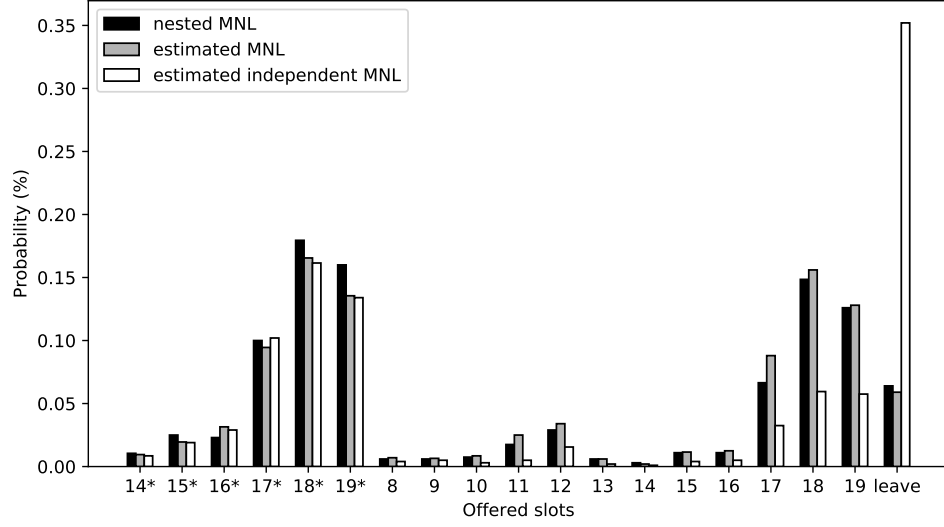
Slot(s)	8	9	10	11	12	13	14	15	16	17	18	19	no-purchase
u_s	2.2	2.4	2.6	3.0	3.4	2.3	1.9	2.7	2.9	3.9	4.3	4.2	3.0

Note: each one-hour time slot is labelled by its starting time. Dissimilarity parameters are 0.8, 0.6 and 1 for Day 1, Day 2 and no-purchase event.

Table 3.1: Parameters in the nested MNL model for slots offered in two delivery days

Based on sampled booking histories, we estimate \hat{u}_s in the MNL model using `asclogit` method in Stata/SE 15. The no-purchase option is selected as the reference variable that $\hat{u}_0 = 0$ in the estimation process. When the no-purchase event is defined as the customer doesn't select any slot from any delivery day, the estimated value of parameters on the MNL model, the robust standard error (RSE) and the p-value are reported in Appendix B.1. Alternatively, we can estimate two independent MNL models for two delivery days respectively by defining no-purchase event with respect to delivery days. Under this case, for example, if a customer picks a slot from Day 2, the customer will be considered as no-purchase for Day 1. The corresponding estimated parameters are presented in Appendix B.3. Note that we can retrieve parameters for the nested MNL choice model from the same booking history. The corresponding estimated parameters are presented in Appendix B.2.

We generate a set of offered time slots to test the ability of those two estimated MNL choice models in approximating the true customer choice behaviour. The set includes 6 time slots from Day 1 and 12 time slots from Day 2. Given this set of time slots, we simulate the slot selection process for 2,000 times (customers) under the nested MNL model defined in Table 3.1, our estimated MNL model and the estimated independent MNL model. Figure 3.3 presents the proportion of customers in selecting each slot under these three choice models. One can easily observe that the estimated independent MNL model overestimates the no-purchase (abbreviated as 'leave') probability and under-estimates the probability of purchasing any slot.



Note: '*' indicates slots in Day 1

Figure 3.3: Choice probabilities for time slots in two delivery days defined under different choice models

We conduct Kolmogorov–Smirnov (KS) test to evaluate whether probability distributions defined our two estimated choice models are equivalent to the one defined by the nested MNL model. By comparing the customer selections generated from the estimated MNL and the true choice models, the p -value in KS test is calculated as 0.0586 so the estimated MNL model is statistically the same as the nested MNL choice model at 95% confidence level. On the other hand, from the comparison of selections generated from the estimated independent MNL and the true choice models, the p -value in KS test is found as less than 0.0001. This indicates that the probability distribution of the estimated independent MNL choice model is not the same as the one defined by the true nested MNL model. Overall, we can conclude that the estimated MNL can effectively approximate the underlying customer choice behaviour where time slots across multiple delivery days are offered. However, the estimated independent MNL models fails to model the customer choice behaviour and it overestimates the no-purchase event.

3.6 Online Problems

We refer to online problems as issues that need to be solved upon arrival of a customer request for delivery. A request consists at this stage only of information on the customer's location and order volume (in terms of delivery totes) and order value.

In practice, we may not yet have information on the order's volume nor value since some retailers allow customers to book their slot before choosing the goods to be delivered; in that case, we assume that an estimate of this volume and value is available (e.g. based on historic observations). The objective is to solve the slot assortment problem, (3.4) and (3.5), for a customer request from area a and segment n ; we re-state this problem here for convenience:

$$\mathbf{g}_a = \begin{cases} \arg \max_{\mathbf{g}_a \in \Psi(\mathbf{x}, \mathbf{q})} \sum_{s \in \mathcal{H}'_{at}(\mathbf{x} - \mathbf{w}, \mathbf{q}')} p_{ans}(\mathbf{g})[r_n - \hat{\Delta}_{tas}^1(\mathbf{x} - \mathbf{w}, \mathbf{q}')], & 1 \leq t \leq T_s. \\ \arg \max_{\mathbf{g}_a \in \Psi(\mathbf{x}, \mathbf{q})} \sum_{s \in \mathcal{H}'_{at}(\mathbf{x}, \mathbf{q})} p_{ans}(\mathbf{g})[r_n - \hat{\Delta}_{tas}^2(\mathbf{x}, \mathbf{q})], & T_s < t \leq T_h. \end{cases} \quad (3.11)$$

To solve this, we first need to assess for which time slots in each potential delivery day the insertion of this request is feasible; in other words, the set of feasible slots $\mathcal{H}'_{at}(\mathbf{x} - \mathbf{w}, \mathbf{q}')$ or $\mathcal{H}'_{at}(\mathbf{x}, \mathbf{q})$ is to be constructed. Accordingly, we have to estimate the opportunity cost $\hat{\Delta}_{tas}^1(\mathbf{x} - \mathbf{w}, \mathbf{q}')$ or $\hat{\Delta}_{tas}^2(\mathbf{x}, \mathbf{q})$. Finally, we use these inputs to decide which of these feasible slots to display as being available by solving the slot assortment problem above. We have discussed the feasibility check process in Section 3.4. In this section, we focus on the opportunity cost estimation and solution method to the slot assortment problem.

3.6.1 Opportunity Cost

The opportunity cost $\hat{\Delta}_{tas}^1(\mathbf{x} - \mathbf{w}, \mathbf{q}')$ or $\hat{\Delta}_{tas}^2(\mathbf{x}, \mathbf{q})$ in (3.4) and (3.5) can be interpreted as the cost arising when accepting an order in slot s for a corresponding delivery day due to the marginal fulfilment cost incurred and due to potential revenue displacement effects. This indicates that by accepting this order, we are not able to satisfy some other orders in the future. In this work, we assume that the latter effect is negligible and focus exclusively on the fulfilment costs. This has the advantage of not having to estimate the expected revenue displacement, which is a hard optimisation problem. The reader is referred to Yang and Strauss (2017) for an approach that takes this into account. In this research, we focus on multi-day choice and investigate how slotting policies can affect profitability by improving route efficiency.

We adopt an approach introduced by Yang et al. (2014) in order to estimate marginal fulfilment costs dynamically. A brief description of this method follows. The main idea to quantify insertion costs for a specific request and specific time slot is to use a weighted average between an insertion cost estimate based only on orders accepted to date and another one based on historic final routing plans. Initially, all weights are on the insertion cost estimate, which is derived as an average over all insertion costs based on historic final routing plans. This estimate is likely

much more accurate at this point in time than the initial high insertion cost to add orders to an (almost) empty tour. As we approach the end of the booking horizon, increasingly more weight is attributed to the insertion cost derived on only accepted orders.

More formally, let us define the insertion cost estimate based on the current orders on the books. For every routing plan Ω_t in the overall collection of routing plans $\mathbf{\Omega}_t$ at time t in the booking horizon, we check whether the order in a given slot s from day k can be feasibly inserted. If the request with time slot s in day k can be feasibly inserted into the routing plan Ω_{kt} with an incremental delivery cost $\bar{c}_{s,\Omega_{kt}}$, then the corresponding insertion cost is computed as $c_{\Omega_{kt}}^s = \bar{c}_{s,\Omega_{kt}} + C_k(\Omega_t) - C_k(\Omega_t^*)$, where $C_k(\Omega_t)$ represents the cost of routing plan Ω_{kt} in day k and $C_k(\Omega_t^*)$ is the cost of the cheapest routing plan in day k in our collection $\mathbf{\Omega}_t$. This reflects that insertion cost may be low, but the cost of the overall plan may be very high relative to the cheapest available alternative plan. If the request cannot be feasibly inserted into routing plan Ω_{kt} , we will set $c_{\Omega_{kt}}^s$ to a large value (we chose $c_{\Omega_{kt}}^s = 50$ in our simulation study) so as to force the slot availability policy to make this slot unavailable. The insertion cost for a service request in slot s at delivery day k for a given location is computed by $c_{\mathbf{\Omega}_{kt}}^s = \min\{c_{\Omega_{kt}}^s \mid \Omega_{kt} \in \mathbf{\Omega}_{kt}\}$.

Similarly, we estimate insertion costs based on a set Φ of historical final routing plans. For each historic plan $\Phi \in \Phi$, we set $c_{\Phi}^s = 0$ if a request in this location already exists in this slot s . Otherwise, c_{Φ}^s will be the insertion cost if the insertion is feasible. In case of having infeasibility, we set the insertion cost to a high value (we use $c_{\Phi}^s = 20$ in our simulation study). The insertion cost for the request across all historic plans is calculated as the average $c_{\Phi}^s = \frac{1}{|\Phi|} \sum_{\Phi \in \Phi} c_{\Phi}^s$.

Finally, the two estimates $c_{\mathbf{\Omega}_{kt}}^s$ and c_{Φ}^s are combined using a weighted average that depends on the remaining time periods t_k in the booking horizon T_k of day k as follows:

$$\hat{\Delta}_{ts} = (1 - \frac{t_k}{T_k})c_{\mathbf{\Omega}_{kt}}^s + \frac{t_k}{T_k}c_{\Phi}^s. \quad (3.12)$$

3.6.2 Slot Assortment Optimisation Under Choice Models

According to the analysis in Section 3.5, we can use a MNL choice model to approximate a nested MNL model (which we will assume to represent actual customer choice behaviour in our numerical experiments). Using the estimated opportunity costs $\hat{\Delta}_{tas}^1(\mathbf{x} - \mathbf{w}, \mathbf{q}')$ or $\hat{\Delta}_{tas}^2(\mathbf{x}, \mathbf{q})$ in the manner discussed above, the online assortment problem (3.11) under the standard MNL model for a given customer request arriving at $1 \leq t \leq T_s$ from segment n and area a can be reformulated as a linear program (this result is due to unimodularity of the constraint matrix in the optimisation

(Davis et al., 2013)):

$$\begin{aligned}
& \max_{\hat{\mathbf{g}}_{\mathbf{a}}} \sum_{s \in \mathcal{H}'_a(\mathbf{x}-\mathbf{w}, \mathbf{q}')} \hat{g}_{ans}(r_n - \hat{\Delta}_{tas}^1(\mathbf{x} - \mathbf{w}, \mathbf{q}')) \\
& \text{s.t.} \quad \sum_{s \in \mathcal{H}'_a(\mathbf{x}-\mathbf{w}, \mathbf{q}')} \hat{g}_{ans} + \hat{g}_{an0} = 1, \\
& \quad \frac{\hat{g}_{ans}}{\hat{v}_{ns}} \leq \hat{g}_{an0}, \quad \forall s \in \mathcal{H}'_a(\mathbf{x} - \mathbf{w}, \mathbf{q}'), \\
& \quad 0 \leq \hat{\mathbf{g}} \leq 1,
\end{aligned} \tag{3.13}$$

where

$$\hat{g}_{ans} = (\hat{v}_{ns}g_{as}) / \left(\sum_{j \in \mathcal{H}'_a(\mathbf{x}-\mathbf{w}, \mathbf{q}')} \hat{v}_{nj}g_{aj} + 1 \right) \text{ and } \hat{g}_{an0} = 1 / \left(\sum_{j \in \mathcal{H}'_a(\mathbf{x}-\mathbf{w}, \mathbf{q}')} \hat{v}_{nj}g_{aj} + 1 \right).$$

Slot s will be offered if \hat{g}_{ns}^* in the optimal solution to (3.13) is non-zero. A similar optimisation problem as in (3.13) is constructed to decide the set of slots to offer to the customer from segment n and area a arriving at t , for $T_s < t \leq T_h$. In this case, we consider the set of feasible slots $\mathcal{H}'_a(\mathbf{x}, \mathbf{q})$ and use the opportunity cost $\hat{\Delta}_{tas}^2(\mathbf{x}, \mathbf{q})$ in the objective function.

It is also possible in principle to optimise the slot assortment using a nested MNL model; assuming that the true customer behaviour is indeed represented by a nested MNL model, one would expect that the resulting policy will perform better than a policy using MNL. Given opportunity cost approximations, we adapt the method proposed in Gallego and Topaloglu (2014) to solve the slot assortment problem under the nested MNL for each request. This method re-formulates the non-linear assortment optimization problem (3.11) into an equivalent linear program. As we will present in our computational experiments in Section 3.7, this method is practically implementable in an online environment and can be solved efficiently. Next, we briefly describe the method for customer requests arriving at t , for $1 \leq t \leq T_s$ and then highlight methodological differences for requests arriving at t , for $T_s < t \leq T_h$.

The method starts with obtaining a collection of slotting decision candidates for each delivery day independently by applying the same method. We adopt $\mathcal{H}'_{ak}(\mathbf{x} - \mathbf{w}, \mathbf{q}')$ to indicate the set of feasible time slots in day k for area a . We introduce a parametric optimisation with parameter $\rho \in \mathbb{R}^+$ to exploit all possible \mathbf{g} that gives an α -approximate solution to problem (3.11) where $1 \leq t \leq T_s$. For any given $\rho \in \mathbb{R}^+$, a slotting decision candidate \mathbf{g} for delivery day k can be calculated by

solving the following slot assortment optimisation model:

$$\max \left\{ \sum_{s \in \mathcal{H}'_{ak}(\mathbf{x}-\mathbf{w}, \mathbf{q}')} \hat{v}_{ns}(r_n - \hat{\Delta}_{tas}^1(\mathbf{x}-\mathbf{w}, \mathbf{q}') - \rho) g_s : g_s \in \{0, 1\}, \forall s \in \mathcal{H}'_{ak}(\mathbf{x}-\mathbf{w}, \mathbf{q}') \right\}. \quad (3.14)$$

Due to binary decisions in (3.14), it can be easily solved as an unconstrained knapsack problem. Specially, we evaluate slots with respect to their values defined as $\hat{v}_{ns}(r_n - \hat{\Delta}_{tas}^1(\mathbf{x}-\mathbf{w}, \mathbf{q}') - \rho)$ and include the slot into the candidate set as long as its value exceeds zero. A collection of candidates can be obtained by enumerating all possible $\rho \in \mathbb{R}^+$ and solving (3.14) for every ρ . Based on the example presented in Gallego and Topaloglu (2014), we can graphically obtain a set of intervals such that the value $\hat{v}_{ns}(r_n - \hat{\Delta}_{tas}^1(\mathbf{x}-\mathbf{w}, \mathbf{q}') - \rho)$ for slot s remains positive or negative if ρ belongs to one of those intervals.

Having obtained a collection of slotting decisions candidates Θ_k for delivery day k in this manner, the optimal slotting decision as a solution to (3.11) for $1 \leq t \leq T_s$ is found by combining candidates from each collection. To simplify notation, let us introduce $\hat{U}_k(\mathbf{g})$ to denote the expected total preference weight of slots in delivery day k under slotting decisions $\mathbf{g} \in \Theta_k$, which is computed as $\hat{U}_k(\mathbf{g}) = \hat{\mathbf{v}}_n^T \mathbf{g}$. Given that slots offered on day k are denoted by $\mathbf{g} \in \Theta_k$ and the customer has decided to book a slot from day k , then the expected profit obtained from the customer request is computed as

$$\hat{R}_k(\mathbf{g}) = \frac{\sum_{s \in \mathcal{H}'_{ak}(\mathbf{x}-\mathbf{w}, \mathbf{q}')} (r_n - \hat{v}_{ns}(r_n - \hat{\Delta}_{tas}^1(\mathbf{x}-\mathbf{w}, \mathbf{q}') - \rho)) \hat{v}_{ns} g_s}{\hat{U}_k(\mathbf{g})}. \quad (3.15)$$

According to the Lemma 1 in Gallego and Topaloglu (2014), if we could somehow determine the value z^* that satisfies:

$$\hat{v}_{n0} z^* = \sum_{k=1}^K \max_{\mathbf{g} \in \Theta_k} \{ \hat{U}_k(\mathbf{g})^{\hat{k}} (\hat{R}_k(\mathbf{g}) - z^*) \}, \quad (3.16)$$

then we can compute the best slotting decision for delivery day k based on

$$\max_{\mathbf{g} \in \Theta_k} \{ \hat{U}_k(\mathbf{g})^{\hat{k}} (\hat{R}_k(\mathbf{g}) - z^*) \}. \quad (3.17)$$

In order to compute z^* , we can linearise problem (3.16) by introducing decision variables y_k for $k = 1, \dots, K$. Accordingly, we solve the corresponding linear program

$$\min \{ z \mid \hat{v}_{n0} z \geq \sum_{k=1}^K y_k; y_k \geq \hat{U}_k(\mathbf{g})^{\hat{k}} (\hat{R}_k(\mathbf{g}) - z), \forall \mathbf{g} \in \Theta_k, k = 1, \dots, K \}. \quad (3.18)$$

Then, we can have the slotting decision for each delivery day by solving problem (3.17) based on the optimal solution z^* from (3.18).

Similarly, at time t for $T_s < t \leq T_h$, the method considers a set of feasible slots $\mathcal{H}'_a(\mathbf{x}, \mathbf{q})$ and uses the opportunity cost $\hat{\Delta}_{tas}^2(\mathbf{x}, \mathbf{q})$ to compute the expected profit as follows:

$$\hat{R}_k(\mathbf{g}) = \frac{\sum_{s \in \mathcal{H}'_{ak}(\mathbf{x}, \mathbf{q})} (r_n - \hat{v}_{ns}(r_n - \hat{\Delta}_{tas}^2(\mathbf{x}, \mathbf{q}), \mathbf{q}')) \hat{v}_{ns} g_s}{\hat{U}_k(\mathbf{g})}. \quad (3.19)$$

3.7 Numerical Experiments

The objective of our numerical study is to obtain insights into the performance of different slotting policies to enable us to answer various research questions including (but not limited to):

- To what extent and under which circumstances does a policy that incorporates a more accurate model of customer choice across days improve performance?
- Are there constellations under which a more accurate representation of choice is undesirable for some reason?
- Are the run times realistic for online decision making?
- What is the added value of allowing same day deliveries in addition to next day ones?

3.7.1 Data and Experiment Design

We give a full description of the data and the setup in the experiments where time slots from two future consecutive days are offered to customers ($K = 3$ without same-day delivery services). In addition to next-day delivery, we will test the benefit of offering same-day delivery services ($K = 2$).

We work with a simplified spatial distribution of customers that loosely reflects the situation of a single depot at the outskirts of an urban area. Specifically, we assume that customers are located within a $12 \times 12 \text{ km}^2$ area. The depot is located outside of the delivery area. The delivery area is divided into two large sub-areas, with a third of the total customer population being uniformly distributed in the lightly shaded area and the rest being uniformly distributed over the remaining white part as shown in Figure 3.4.

On the capacity side, we have 16 homogeneous vehicles located at the single depot operating from 7am until 9pm. Each vehicle has 70 units of capacity and travels at a speed of 20 km per hour. Preparing an order requires a 4-hours lead time and serving one customer requires a service time of 15 minutes. The maximum duration




Figure 3.4: Area of attended home delivery services

of each route for any vehicle is 4 hours and each vehicle can have 4 routes maximum during one operation day. It takes 30 minutes for the vehicle to load customer orders before departing from the depot. The total delivery costs are estimated based on the travel distance multiplied by the fuel cost £1.5/km. We assume that the capacity for each delivery day is more restricted (by the one-hour time window constraint) than the vehicle capacity. By setting each vehicle having four trips during one specific delivery day, we estimate the delivery capacity only with respect to time window constraints based on the continuous approximation method from Daganzo (1987). This results to a capacity of serving 30 customers for each vehicle. Note that it is a really conservative capacity estimation.

Customers are modelled to choose a time slot according to the nested MNL defined in Table 3.1. Each day, there are 12 non-overlapping one-hour time slots starting from 8am until 8pm as defined in Example 3.5.1. For the sake of simplicity, we use a single customer segment. The profit before delivery costs that is associated with a given order is sampled from a normal distribution with mean of £20 and standard deviation of £5; if a sample is negative, we replace it with zero. We emphasise that this sampled quantity represents profit (before delivery costs) instead of revenue; this is important in as far as using revenue figures distorts the optimisation since revenue is typically much larger than routing costs. As discussed in Section 3.5, we generate a history of booking records using the ground truth choice model so as to estimate parameters for the choice models to be used in the different policies. Note that the same booking history is used to estimate the choice model for same-day slots as we adopt the same ground truth choice model.

All experiments are conducted based on a computer with processor Intel(R) Core(TM) i5-700 CPU 3.40GHz. The underlying optimisation models are implemented in Java and solved by using CPLEX. We evaluate the performance of each simulation scenario based on the average performance obtained by repeating the corresponding experiments for 100 times.

Next-days Delivery Services: Let us assume that deliveries can be accepted today for any of the next two days; however, deliveries cannot be fulfilled today. In the construction of our scenarios, we focus on a rolling three-days time horizon (today and the next two following days); this is because in e-grocery retailing most demand typically arrives within three days prior to dispatch. This is a conceptual difference to most other related work where there is a fixed booking horizon that leads to a particular delivery day that is being optimised for.

In all scenarios, we already start with a fixed set of orders that we assume to have been accumulated over previous days (note again that we do not work with a fixed but a rolling time horizon). It does not matter at what time an order comes in today since we do not anticipate potential displacement costs associated with displacing future orders due to receiving an order now for a given time slot. Therefore, we do not need to explicitly model a time grid. In the base scenario, we assume that exactly 450 requests arrive each day, chosen such that the expected demand equals total expected delivery capacity. We scale the number of requests up and down to assess the behaviour of slotting policies relative to varying degrees of capacity tightness. The base scenario corresponds to a scaling parameter of 1 (ratio of total expected demand equal to total expected delivery capacity having taken time constraints and the spatial distribution of customers into account). Note that the time constraints stemming from the delivery time windows are in e-grocery retailing typically much more restrictive than vehicle capacity. In order to obtain the pre-accepted orders in Day 1, we simulate one booking process for 450 customers and each customer is only presented with all feasible time slots in Day 1. This leads to 155 pre-accepted orders.

For our simulation experiments, we test four different slotting policies and a brief description of these policies follows.

- **Slotting policy with true choice model (TRUE):** A nested MNL model is used in this policy and its parameters are estimated from the booking histories. The resulting parameters are presented in Appendix B.2.
- **Slotting policy with merged choice model (MNL-DAY):** We use the MNL model with parameters estimated based on the booking histories where only slots in Day 1 and Day 2 are offered. The resulting parameters are

presented in Appendix B.1.

- **Slotting policy with independent choice models (MNL-IND):** This slotting policy ignores interaction between delivery days as assumed in most studies in the literature. To represent this approach, we estimate MNL models for the two delivery days independently. The no-purchase event is defined as in Yang et al. (2014) for each day so that the customer decides not to select any slot of the day. In other words, no-purchase event is accounted for each individual day. If the customer does not select any slot in one delivery day, we classify such case as no-purchase for this delivery day. Parameters of MNL models are presented in Appendix B.3.
- **Offer all feasible slot policy (ALL-FEASIBLE):** As a simple benchmark against which we can assess the benefit of making slotting decisions, the ALL-FEASIBLE policy consists of always offering all feasible slots across all delivery days.

Same-day and Next-day Delivery Services: We then extend the previous experiments by combining same-day and next-days delivery services. The models developed in the literature consider these two problems separately. Our results show potential benefit of combining these two, even when same-day orders are incorporated in a fairly simplistic way.

The setting of this experiment is different from the one described above since we now have additionally the opportunity to offer same-day slots to customers as well as slots on the next days. Specifically, we restrict the experiments to include only one day into the future. At the beginning of today we will have already a number of orders on the books for today, we always start from the same set of previously accepted orders including 372 customers. We obtain this set of orders by presenting 500 customers with all feasible time slots in Day 2 and simulating their slot selection decisions accordingly.

With the opportunity to offer same-day delivery, the lead time required to assemble an order becomes important to determine which slots could potentially be offered given a request at a particular point in time. We assume that the customer arrival process follows the same non-homogeneous Poisson process during each day with rates $\lambda(t)$ defined as

$$\lambda(t) = \begin{cases} 2.3, & \text{before 8am,} \\ 2.8, & \text{between 8am and 5pm,} \\ 2.5, & \text{after 5pm.} \end{cases} \quad (3.20)$$

We apply a scaling parameter to these arrival rates during one booking horizon in order to generate different demand levels (ranging from 0.6 to 1.6 with step size 0.1, where level 1 corresponds to 620 requests).

In addition to policies MNL-IND (applied to slot options today and on the next day), we specifically introduce another policy to test with our setting:

- **Slotting policy with next-day only (NEXT-ONLY):** We never offer same-day slots and use the MNL choice with utility parameters as in Appendix B.3 for Day 2.

3.7.2 Computational Results and Analysis

We discuss first the effects of modelling choice between delivery days as independent (between days) versus modelling them using the (true) nested MNL model over multiple (here two) days. Then, we quantify the effect of incorporating same-day delivery with next-day delivery.

On Choice Modelling for Delivery Days Independently: Most literature in the domain of attended home delivery services assumes that deliveries are being planned for a single delivery day. Figure 3.5 shows the percentage profit increases of the various slotting policies relative to using the benchmark ALL-FEASIBLE of always offering all feasible slots. A low scaling parameter (< 1) represents low expected demand relative to available capacity, a large parameter ($\gg 1$) represents tight capacity. We observe a somewhat counter-intuitive result: the independent policy MNL-IND delivers profits at least as good as the benchmark across all scenarios (discounting statistically insignificant differences), whereas the TRUE policy (using the true choice model) is performing very poorly under low demand and only outperforms all others under high degrees of capacity tightness.




Figure 3.5: Total profit increase (%) under MNL-IND, MNL-DAY and TRUE compared to ALL-FEASIBLE

To explain this, let us drill a bit deeper. Profit increase may be driven by increased the profit per order (Figure 3.6) and/or by increasing the number of orders (Figure 3.7). As one would expect, using the true choice model (TRUE) is superior to the other policies in its ability to attract highly profitable orders; this is because TRUE can better steer customers towards slots with the lowest opportunity costs (i.e. expected marginal routing costs). To see this, we plot the percentage cost reduction per order relative to ALL-FEASIBLE in Figure 3.8: TRUE dominates all other policies in cost reduction per order across all demand scenarios. As demand increases, routes tend to become more congested and thus the potential for routing cost savings diminishes, but for TRUE at a slower pace than for the other policies. However, these advantages of TRUE are more than outweighed under scenarios of low or medium demand by receiving up to 20% fewer orders as compared to the benchmark ALL-FEASIBLE although we are using the true nested MNL choice model. We emphasize that all policies are using the same opportunity cost approximation method (which only anticipates delivery costs but not the expected value of future orders and related displacement costs); the only difference is in the choice model used in the dynamic slot availability decisions. For scenarios of low demand, displacement costs associated with lost future revenue would typically be zero anyway, and since TRUE uses the true choice model, the reason for its poor performance must be in poor marginal routing cost estimates. This makes sense in as far as marginal routing costs will be particularly difficult to predict when routes are sparsely populated, and small changes in demand may lead to large changes

in the marginal routing cost. TRUE is sensitive to these imperfect cost estimates and tends to restrict the range of available slots so as to nudge customers towards anticipated cheap ones, thereby however incurring a higher risk of non-purchase for the perceived (but often overestimated) benefit of more efficient routes.




Figure 3.6: Profit increase (%) per order under MNL-IND, MNL-DAY and TRUE compared to ALL-FEASIBLE

MNL-IND on the other hand has the exact same difficulties with routing cost estimation but happens to compensate them by the overestimation of the non-purchase probability, leading overall to a higher volume of orders. In other words, given a set of feasible slots with the same corresponding opportunity costs, MNL-IND tends to make more slots available than TRUE or MNL-DAY. In fact, MNL-IND attracts about as many orders as the benchmark.

Finally, we note that using an MNL model per day that includes a feature for the attractiveness of the corresponding day (MNL-DAY) does perform poorly across all scenarios. It models days independently (hence it does better than TRUE under low demand) but it is not able to realise the cost improvements of TRUE under high demand - therefore, it is overall not recommended.




Figure 3.7: Accepted order increase (%) under MNL-IND, MNL-DAY and TRUE compared to ALL-FEASIBLE




Figure 3.8: Cost reduction (%) per order under MNL-IND, MNL-DAY and TRUE comparing to ALL-FEASIBLE

In conclusion, our main finding is that modelling choice independent of cross-delivery-day effects may actually be a better approach than using a model that does take this dependency into account. The independent day policy MNL-IND provided robust results across all demand scenarios, whereas using the true choice model only performs well when capacity is very tight. If the business situation is such that demand substantially varies (meaning that the capacity tightness varies significantly), we may be better off by deploying a slotting policy that treats slotting decisions for different delivery days as independent. Only if capacity is typically

tight, or if we have a means of estimating accurately marginal routing costs, then we should use a choice mode that includes cross-day effects.

Runtime: We consider two tasks to be necessarily online computations given a specific incoming customer request for delivery: checking slot feasibility across all slots on all days, and solving the slot availability optimisation problem. Calculation of marginal routing costs is considered to be an offline task since we could maintain this information in a look-up table for every combination of location and time slot. In principle, the same holds for feasibility checks, but for these are more important to get right since we must not end up in a situation where we commit to an infeasible routing plan, whereas using poor marginal routing cost may only affect overall profitability

Using the compact formulations as presented in this chapter, the slot optimisation turns out to be almost instantaneous for all policies. The feasibility check, however, is more time-consuming and represents the lion’s share of the times reported in Figure 3.9. The benchmark ALL-FEASIBLE does not involve any optimisation yet still takes the most time, whereas TRUE has the most sophisticated optimisation problem but is (on average) the fastest.

To explain, it may be useful to revisit the workings of the feasibility check: for a given order request from a given location, checking feasibility of a specific slot consists of attempting to insert the request into the current delivery plan and the plans maintained in the pool. We stop the search as soon as a feasible insertion location has been found. Therefore, routing plans that are better constructed (and therefore have more slack for additional orders) or that have fewer orders (such that there are fewer potential insertion positions) will be quicker to check. This can be seen in the runtime results: although MNL-IND and TRUE have similar volumes, TRUE is slightly faster.

Overall, MNL-IND and TRUE seem reasonably quick for practical deployment, and it is worth highlighting the somewhat counter-intuitive finding that adding the slot availability optimisation step actually tends to decrease runtime significantly rather than increase it (relative to simply offering all feasible slots). We feel that this is an important argument to present to practitioners who may be reluctant in adopting slotting policies due to concerns over website loading speed implications.

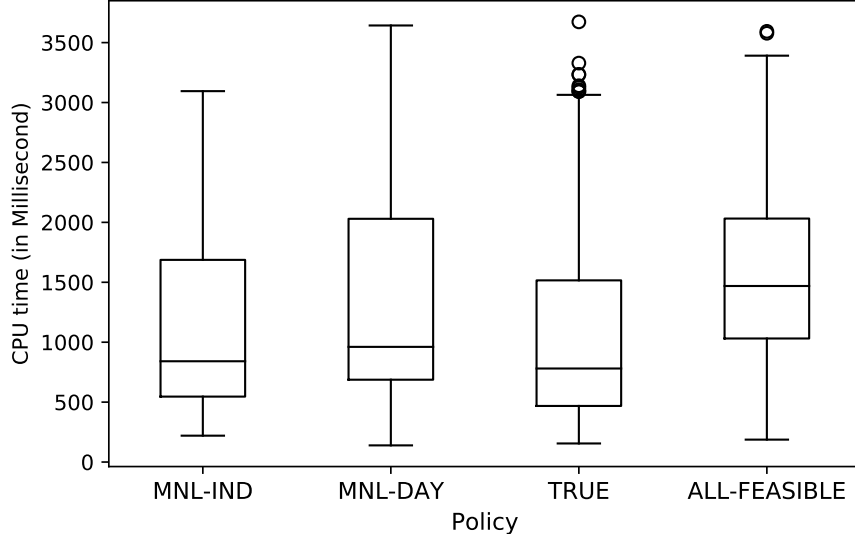


Figure 3.9: CPU time taken to check slot feasibility and decide slot availability for each delivery request under different policies

Integrating Same-Day Delivery: In the literature, typically either next-day or same-day services have been investigated. In this section, we consider to combine same-day and next-day delivery services within multi-day delivery context. To the extent, this has been already being implemented in practice by some retailers: for example, Sainsbury’s (one of the leading retailers in the UK) offers both same-day and next-day deliveries of groceries.

Now, we start to evaluate the benefit of offering same-day delivery services along with next-day delivery services (still managing time slots in two consecutive days). We are concerned with the total profit obtained on one specific day which is computed by the total revenue-before-delivery from all pre-accepted orders and same-day orders excluding the total delivery costs. Figure 3.10 presents the total profit increase (%) in one delivery day after the introducing the same-day delivery services. We see that offering same-day delivery services can increase total profit by around 5% at all demand levels.

We further show how the total profit increase is broken into the number of orders and the profit per order in Figure 3.11 and Figure 3.12, respectively. Accordingly, we observe that the number of accepted orders increases, as well as the profit per order, after offering same-day delivery services at varying demand levels. As the opportunity cost estimation only involves delivery cost, the profit per order increase (%) is mainly driven by obtaining a more efficient delivery plan. Figure 3.13 presents the cost reduction (%) per order after introducing same-day delivery services.

Recall that same-day delivery requests are accepted without dramatically changing the existing routing plan. Therefore, hypothetically, more requests are accepted without causing total delivery costs increase in our settings such that the cost per order decreases at varying demand levels.

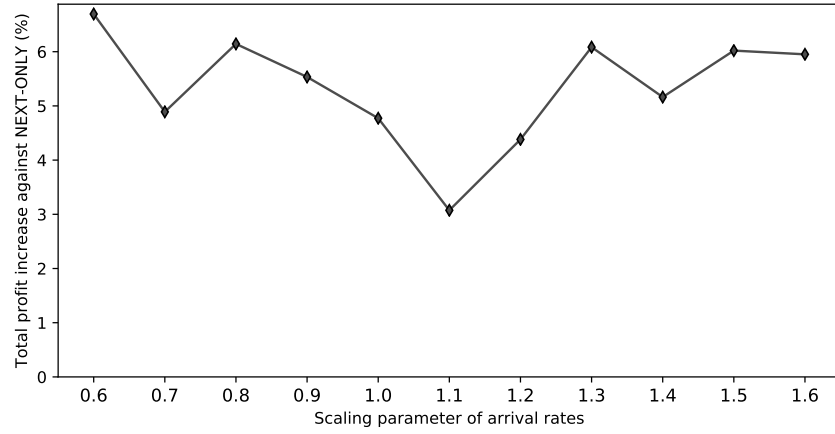


Figure 3.10: Total profit increases (%) after offering same-day delivery services

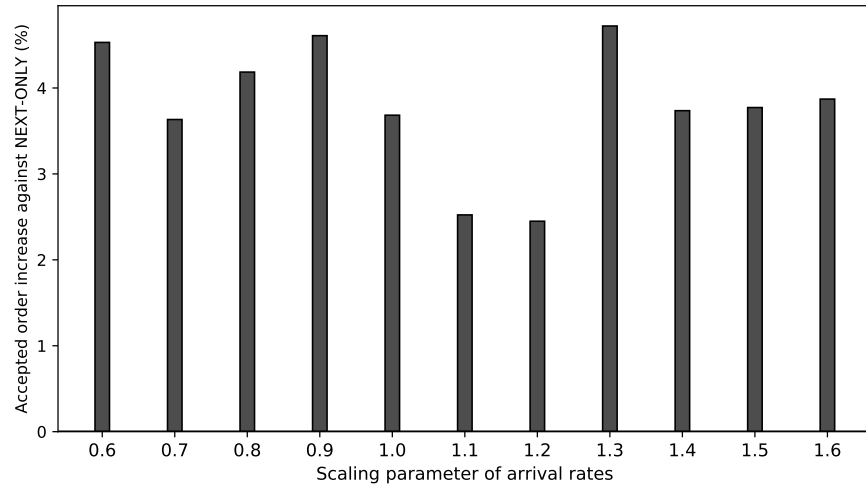


Figure 3.11: Accepted order increase (%) after introducing same-day delivery services

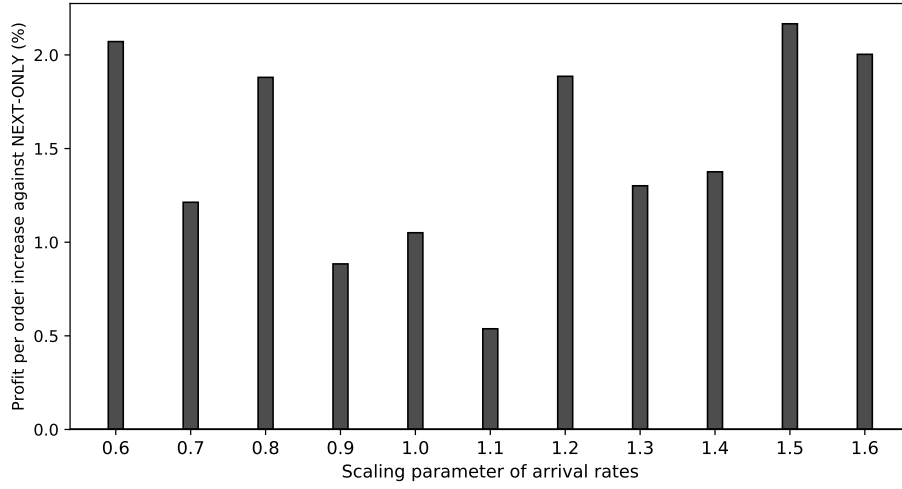


Figure 3.12: Profit per order increase (%) after introducing same-day delivery services

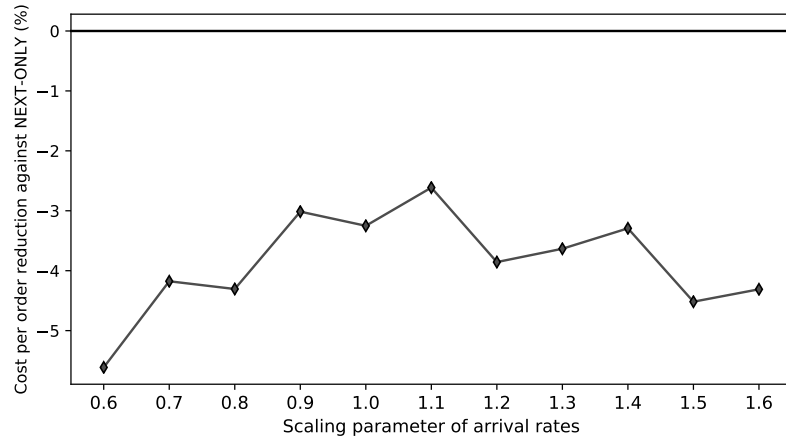


Figure 3.13: Cost reduction (%) per order after introducing same-day delivery services

3.8 Conclusions

In this chapter, we discussed a dynamic slotting problem of managing time slots from multiple delivery days for attended home delivery services. When customers request delivery services, they are presented with time slots from a number of consecutive days and they need to select exactly one slot or leave without purchase. The underlying customer choice behaviour is influenced by the presented delivery days and the delivery time slots available within them. Depending on how the

influence of delivery day on customer choice behaviour is captured, we investigate three approaches to model customer choice and to dynamically decide the availability of time slots from multiple delivery days accordingly.

Several managerial insights have been obtained via a simulation study using the common industry practice of offering all feasible slots as a benchmark:

- First, actively managing slot availability using choice-model-based optimisation can significantly reduce fulfilment costs. According to Figure 3.8, the cost per order can be reduced by around 7% when the NMNL choice model is embedded in optimisation.
- Second, this optimisation does not come at the expense of increased runtime for online computations presented in Figure 3.9; in fact, it tends to reduce runtime due to better routes that help to speed up the feasibility checks.
- Third, if the estimation of expected marginal cost-to-serve is difficult (e.g. when demand is low relative to available capacity), then we are better off using an MNL choice model that ignores the effect of other delivery days because its overestimation of the non-purchase probability leads to higher slot availability and thus higher order volume as presented in Figure 3.5. Only if we have a means of obtaining high quality estimates of expected cost-to-serve (which tends to be easier when demand is high relative to available capacity), then using more sophisticated choice models that incorporate cross-day substitution effects is beneficial as shown in Figure 3.8. As a consequence, if there is a high degree of uncertainty over the level of demand, then it is better to use MNL without accounting for cross-day substitution effects because the resulting policy is more robust to demand level shocks and consistently performs at least at the level of the benchmark or better.
- Fourth, offering same-day delivery services along with next-day delivery services improves delivery efficiency and increases profit by around 5% at all demand levels.

In this Chapter, we assume that the NMNL model is the ground true model and customers naturally partition time slots with respect to delivery days. Under this assumption, our discussions and conclusions are obtained because the probability of no-purchase event in the specific delivery day is over-estimated when substitution effect of delivery days is not considered in the policy. However, time slots may be classified/nested based on different factors, such as morning, afternoon and evening. Our conclusions may not be valid if the nest in the NMNL model is defined differently and the probability of no-purchase event in the specific delivery day is not

over-estimated. Moreover, our numerical experiments are conducted with a single customer segment and each order has the same average profit-before-cost, which may not be the case in practice. However, having one segment may have minor impact on our conclusions because our choice model without substitution effect overestimates the no-purchase probability regardless of segmentation.

A limitation of the proposed approach is that the opportunity cost approximation in our solution approaches only considers the future delivery costs. Including displacement costs is beneficial when demand is high relative to available capacity, but should not affect results in the opposite scenarios as displacement costs in that case would typically be zero. Also, in our experiments we limited ourselves to choice between slots over two delivery days. In practice, often 7 days or more are available to choose from; however, in e-grocery the vast majority of orders arrives within three days of delivery, so we feel that this limitation is minor. Naturally, it would be most interesting to see whether these results could be replicated using real data to calibrate multi-day choice models.

Chapter 4

Dynamic Capacity Allocation for Airline Upgrade Options with Customer Anticipation

Upgrades are usually offered by airlines to balance the demand and capacity for economy and business cabins as well as to increase customer satisfaction. To gain more flexibility in upgrading customers, the airline can sell upgrade options that allow reallocation decisions to be postponed until demand is fully realised. Moreover, there exists some customers who regularly purchase the flight tickets and have a network to share their booking experiences. After upgrade options are introduced to the market, these customers are able to develop anticipation on the execution of upgrade options and take advantage of obtaining a business cabin seat much cheaper than the normal price. It may lead to potential revenue loss for the airline in the long run.

In this chapter, we study a dynamic capacity allocation problem of an airline issuing upgrade options for a sequence of identical single-leg flights. We propose a novel dynamic program (DP) model to decide capacity allocation decisions for economy and business cabins and upgrade options of each flight. The states of the DP model are defined by customer anticipation level, which directly affects customer purchase behaviour. By discretising the continuous state space of the DP model by a finite number of specific anticipation levels, backward recursion algorithm is adapted to approximate the value function at each discrete state of the DP model. A choice-based deterministic integer program (CDIP) model is proposed to estimate the revenue from one flight given the anticipation level. Moreover, the optimal solution of the CDIP model is used for optimal capacity allocation decisions for each flight. We design a series of computational experiments to evaluate the performance of the

solution approach using simulation. We find that the airline could potentially obtain additional revenue by offering upgrade options when the demand of business cabin is low relative to its capacity. The results also show that considering the customer anticipation in allocating capacity with upgrade options can significantly improve the total revenue of the airline in the long term.

4.1 Cabin Upgrade System for Airlines

The UK airline market has been growing steadily by nearly 7% and the number of passengers in domestic and international routes has reached to 264 million during 2017 (Mintel, 2018b). In particular, business travellers are the main contribution to airlines and 38% business trips are through air traffic (Mintel, 2018a). In practice, airlines offer upgrades to increase customer satisfaction and resolve supply-demand mismatch between economy and business cabins (Yu et al., 2015). Upgrades can be offered free of charge or at a certain price. In 'free of charge' case, offering upgrades increases customer satisfaction obviously (Mattila et al., 2013). However, this may create unfairness perception among passengers who do not receive upgrades (Park and Jang, 2015). Airlines more often provide cabin upgrades for passengers who purchase economy tickets by charging a certain fee (Çakanyıldırım et al., 2017). Since it is up to customers to accept or reject upgrade offers, the unfairness imposed to customers by offering free upgrades is somehow overcome. However, implementing such strategy requires a sound stopping rule for airlines to limit upgrade offers because they are sold at low price. Offering too many upgrades may result in insufficient capacity for demand of business cabin, which could reduce total revenue (Steinhardt and Gönsch, 2012).

Upgrades can also be alternatively offered through options to customers who purchase economy tickets as introduced by Optiontown (2018). Presented with the upgrade option, the customer needs to immediately decide whether to purchase the option or not. If the airline does not execute upgrade options at the end of the booking horizon, prices associated with options are refunded to customers. Because there is a possibility of upgrading to business cabin at low cost, it could steer customer demand for economy cabin but reduce demand of business cabin. Additionally, customer satisfaction can be increased as options are only charged when they are executed. For airlines, additional revenue can be earned by executing upgrade options. Moreover, since option execution decisions are postponed until demand is fully realised, it provides flexibility for airlines to allocate capacities of economy and business cabins during the booking horizon.

On the other hand, offering upgrade options may lead to revenue loss for

airlines in the long term. Especially, in short haul markets, a large proportion of capacity in one flight are purchased by business travellers (Mason and Gray, 1995). These business travellers are also regular customers commuting by one specific flight. Based on the concept of customer learning proposed by Wu and Chen (2000), these regular customers form attitudes and acquire beliefs on upgrade options, such as the frequency of options being executed, from past purchase experience. Their subsequent purchase behaviour is influenced by these attitudes and beliefs such that they can behave strategically by purchasing economy tickets and waiting for upgrade options being offered. As upgrade options are priced much lower than fares of business cabin, airlines may lose revenue resulted from such customer strategic behaviour. Therefore, it is important for airlines to consider customers' strategic behaviour when providing upgrade options such that cannibalisation in business cabin can be avoided in the long term.

In this chapter, we focus on a capacity allocation problem with upgrade options for a sequence of identical single-leg flights. A brief problem description is as follows. We assume that customers regularly purchase flight tickets and have a network to share their booking experience on upgrade options, such as whether the option is provided and executed. Specifically, the capacity of the business cabin is the same for all flights and the capacity of the economy cabin is the same for all flights. Fixed fares are set for both cabins as well as upgrade options accordingly. During the booking horizon of each flight, we assume that the airline presents available cabins to every request and decides whether to offer upgrade options. Based on the presented cabins and knowledge gained on upgrade options, the customer needs to purchase a capacity or to leave the system. If the customer purchases the economy cabin and the upgrade option is also offered, then the customer further decides whether to purchase the option. At the end of booking horizon of each flight, the airline executes a number of upgrade options that maximises total revenue of the flight and also refunds payment of unexecuted options. As a result, the customer's experience adds new knowledge on upgrade options and this affects their purchase decisions for the next flight. In the long term, the airline wishes to maximise total revenue received over all flights under the consideration of customers' strategic purchase behaviour with upgrade options.

The main contributions of this chapter are summarised as follows. We consider a customer learning process in a long-term capacity allocation problem for airlines. This learning process enables customers i) to gradually establish belief towards upgrade options, ii) to change their purchase behaviour in the future, and iii) to cause possible cannibalisation in business class for airlines. However, formulating customer learning in the capacity allocation problem of consecutive flights is difficult

because learning is a dynamic and noisy process based on individual's past experience. In order to overcome this challenge, we propose a novel dynamic program (DP) model focusing on the capacity allocation decisions for each flight where customer learning process is reflected by the state of the DP model. Note that we do not consider a dynamic capacity allocation problem at customer level in this chapter as a number of consecutive flights is involved in our problem and customer belief is static during the booking horizon of each flight. Specifically, the DP model involves the capacity allocation decisions for each flight such that the discounted total revenue over all flights is maximised.

Furthermore, the capacity allocation in view of customer learning about upgrade options boils down to identifying key factors influencing customer purchase behaviour. We assume that customers' belief affects their purchase behaviour. Accordingly, we introduce the concept of anticipation level defined as the ratio of the number of executed upgrade options to the capacity of business cabin. This anticipation level is then used to represent customer belief in our problem formulation. Such customer anticipation level shares similar characteristics with the reference price as they both consider the impact of customers' purchase experience on the future demand. However, reference prices address customers' perceptions on historical charged prices (Boer, 2015). In our problem, the customer anticipation levels focus more on customers' experience with upgrade options in terms of the number of executed options. During the booking horizon of each flight, the anticipation level is assumed to affect customer purchase behaviour, which is captured by an multinomial logit (MNL) choice model in this chapter. The anticipation level is also used as the system state in the DP formulation of this capacity planning problem.

Due to our approximation method to anticipation level, the state of dynamic system is continuous. As a solution method, we first discretise this continuous state space into a number of specific anticipation levels. We then construct a choice-based deterministic integer program (CDIP) model to approximate the value function at each discrete state in the DP model. Given the anticipation ratio, CDIP model is also used to make capacity allocation decisions at the beginning of the booking horizon of each flight. Based on realised demand, the airline executes upgrade options to maximise the revenue of the current flight. We evaluate the proposed solution method in a simulation study. Our results show that considering the customer anticipation in allocating capacity with upgrade options can significantly improve the airline's total revenue in the long term.

The structure of this chapter is as follows. Literature review is presented in Section 4.2 and the problem is formally described and formulated in Section 4.3. We present the assortment optimisation model and the demand model at customer

level in Section 4.4. Then, an approximation method is introduced in Section 4.5. Our computational experiments and analysis of results are presented in Section 4.6. Conclusions are reported in Section 4.7.

4.2 Literature Review on Service Upgrades

In this section, we review the most related studies in the literature on upgrade and their use in different applications. Those studies are based on the dynamic service upgrade problem. Moreover, we discuss studies concerned with strategic customer behaviour under repeat purchase. Finally, contributions of this research in terms of distinguishing features from the literature are highlighted.

The dynamic service upgrade problem deals with an optimisation of capacity allocation among vertically differentiated products. The main objective is to develop strategies to match demand and supply. Shumsky and Zhang (2009) focus on a dynamic upgrade problem with customer requests arriving during a specific booking horizon. In their problem, the initial capacity of each product needs to be decided before accepting any customer request. When a customer request arrives, they also decide whether to upgrade the customer. Specifically, they only consider single-step upgrades and assume that customers always accept upgrades. A dynamic program has been used to formulate the problem. In their solution method, the same-class demand is always satisfied first before considering upgrades. A protection limit is calculated for each product as the minimum capacity reserved for same-class demand. Although we also consider single-step upgrade in our problem, our upgrade decisions are made when demand is fully realised. Additionally, Shumsky and Zhang (2009) extend their DP model to include multiple booking horizons by allowing replenishment for each product. They assume that the remaining capacity left from the previous booking horizon can also be used in the current booking horizon. Under such a replenishment assumption, their solution method can be only applied for general non-perishable products. Likewise, Yu et al. (2015) focus on dynamically deciding multiple-step upgrade, where a number of protection limits are calculated for each product with respect to the other products. As upgrades are generally used to resolve the mismatch between demand and capacity for each product, they are considered as free of charge in problems studied by Shumsky and Zhang (2009) and Yu et al. (2015).

Recently, Ceryan et al. (2018) study a dynamic pricing and availability control problem for upgrades with non-perishable products. This paper has the closest setting to our problem. They discretise the booking horizon into time periods where the pricing and replenishment decisions for all products and upgrades are made. They

assume that a number of customers may appear during each time period and prices of all products and upgrades are fixed for these customers. Moreover, Ceryan et al. (2018) focus on upgrades as commitments rather than options; therefore, all customers purchasing the upgrades are upgraded at the end of each time period. Similar to our problem, Ceryan et al. (2018) offer upgrades after customers purchase standard products. However, they don't account for customer anticipation on upgrades since they assume that customers don't consider the possibility of being offered upgrades while making their purchase decision.

The concept of free upgrading has also been considered for perishable products or services, such as in airlines and car rental. Gönsch et al. (2013) formulate a DP model to make upgrade decisions for a single-leg flight. In order to obtain an efficient solution method to solve a real-size problem, they introduce a linear programming (LP) approximation for the DP problem and calculate the protection level for each product, which can be integrated into expected marginal seat revenue (EMSR) rules used by airlines. As an alternative approach to the static method in Gönsch et al. (2013), McCaffrey and Walczak (2016) propose a solution method for solving the exact DP to decide upgrades under a single-leg flight.

When the upgrade decisions are made for customers requesting a multi-leg flight or multi-day car rental, the problem becomes a network revenue management problem (Oliveira et al., 2017). Steinhardt and Gönsch (2012) formulate a DP model that accepts/rejects an upgrade request from a customer renting cars for multiple days. As a solution method, Steinhardt and Gönsch (2012) decompose the network DP problem with respect to resources and solve the small size DP model for each resource.

The concept of upselling upgrades is introduced by Gallego and Stefanescu (2011). They develop a DP model for dynamically pricing upgrades in network revenue management using the MNL choice model. They derive the dynamic pricing policy based on a deterministic linear program (DLP) approximation. Similar to our work, Gallego and Stefanescu (2011) also offer upgrades through products rather than commitments. However, these upgradable products are presented simultaneously with other products and customers do not receive any refund in case of no upgrade. Moreover, they only consider the problem during the booking horizon of one flight, and they don't consider customer strategic behaviour in using upgrades to obtain premium products.

Upgrades offered as options have been only studied under hotel revenue management context. Yilmaz et al. (2017) are concerned with finding an optimal pricing policy for premium rooms and standby upgrades in the hotel sector. With standby upgrades, the customer is only charged if the upgrade is available when the

customer arrives at the hotel. Similar to our problem, they also assume that the strategic customer behaviour is influenced by customer's belief on the frequency of getting upgraded through standby upgrades. However, Yilmaz et al. (2017) assume that strategic customers always purchase standard products first in order to check the price of standby upgrades and then purchase upgrade directly for premium products. Moreover, their model does not consider the learning process behind customer belief and only focuses on maximising the total revenue for a specific check-in date under the given customer strategic behaviour.

In their following work, based on the hotel booking data, Yilmaz et al. (2019) empirically show the existence of strategic customers when standby upgrades are offered. They derive a pricing policy that differentiates strategic and myopic customers. They capture customer purchase behaviour with a sequential logit model. However, they only focus on the problem for one specific date without modelling the customer learning process.

The customer learning process, as a result of repeat purchase, has been studied by several researches in the marketing field. They address the learning process for the purpose of improving the sales prediction. The concept of learning is first introduced by Wu and Chen (2000) under the customer repurchase process. They capture customer's learning as the relationship among the repurchase probabilities displayed at different times of purchase. According to an empirical study on consumer purchase data for tea, they find that learning is an important factor influencing customer purchase behaviour. The importance of learning in modelling customer purchase behaviour has also been discussed by Fader et al. (2004) and Meade and Islam (2010) under the context of launching new products.

Product pricing history has also argued to affect customer purchase decision. Popescu and Wu (2007) have discussed the optimal pricing problem under the consideration of reference price which is updated once a new price has been decided and released to customers. By assuming a bounded and continuous reference-dependent demand model, they have proved that the optimal pricing path is monotonic in the long run and the optimal price reaches a steady state. In addition to loss-neutral customers, they have extended their problem for loss-averse ones. However, due to the strong assumption on the demand model, conclusions in Popescu and Wu (2007) may not hold for the case where a number of substitutional products are offered and a discrete choice model is required to capture customer behaviour.

The customer learning process has also been considered in order to determine a dynamic pricing policy. For instance, Zhang et al. (2014) introduce the concept of trust to explicitly reflect the knowledge gained during the learning process. A hidden markov model is developed by using historical purchase data to indicate

transition probabilities among trust levels. Under the uncertainty of trust levels, Zhang et al. (2014) propose the pricing policy by assuming that customers have the same purchase behaviour at the same trust level. Similar to our work, Zhang et al. (2014) consider the evolution of trust influenced by the pricing policy. However, they adopt a simulation-based optimisation as a solution method, which enumerates over all feasible price paths to obtain the optimal pricing decisions. This can be a computationally expensive solution method if a longer time horizon is considered or the number of trust levels increases.

In summary, we can contribute to the literature on the capacity planning problem for perishable products with upgrades by considering the customer learning through repeat purchase. To the best of our knowledge, this research is the first study combining these two main concepts. Therefore, the capacity planning approaches studied in the literature cannot be adopted to solve the underlying DP problem as the customer choice behaviour evolves through learning under repeat purchase. Unlike the other studies in the literature, we introduce a dynamic program formulation to the problem whose states reflect the customer learning process. As the solution method, we introduce an approximation method based on the integer program model.

4.3 The Dynamic Capacity Allocation under Customer Anticipation

Consider an airline providing single-leg identical flights (labelled by $n = 1, \dots, N$) between the same origin and destination points. All flights provide seats from both economy and business cabins. Specifically, the economy cabin in all flights has the same capacity and so does their business cabin. Attached to the economy cabin, the airline sells upgrade options. We denote economy and business cabins by indices of e and b , respectively, and the upgrade option is represented by u .

The products are defined as combinations of seats from economy and business cabins with upgrade options. A set of products in terms of flight tickets in both business cabin, economy cabin and option that can be offered to customers are listed in $\mathcal{S} = \{\emptyset, \{e\}, \{b\}, \{e, b\}, \{e, u\}, \{e, b, u\}\}$. Note that \emptyset refers to the case of rejecting a customer request (i.e., no product is offered). When a customer request arrives, exactly one element s of set \mathcal{S} is selected to display to the customer. We name s for $s \in \mathcal{S}$ as offered products; for instance, $\{e, b, u\}$ are offered products. Note that when $\{e, b, u\}$ are offered products, the upgrade option is not presented to the customer at the same time with economy and business cabins. We assume that the seat capacity of both economy and business cabins has fixed prices, represented by r_e and r_b , respectively. The upgrade option also has a fixed price r_u .

Each flight has its own booking horizon and these booking horizons are independent and non-overlapping. During the booking horizon of one flight, customers select one product or leave with no purchasing based on offered products. Note that even if the customer buys an economy seat, then the upgrade option may not be presented. Accordingly, the uncertainty resulting from each customer's choice behaviour is influenced by presented cabins and customer anticipation. At the end of the booking horizon of the current flight, the airline executes upgrade options given the remaining capacity of business cabin, which updates customers' experience with upgrade options. Note that customers' anticipation remains constant within the booking horizon of each flight. However, since our problem considers the long-term revenue management with a number of flights, there is gradual change with respect to flights. The airline needs to decide which products from set \mathcal{S} to offer to each customer request arriving during the booking horizon of each flight so that the discounted total revenue to be gained over N flights is maximised.

We consider a Markov Decision Process (MDP) to formulate the airline's capacity allocation problem for economy and business cabins with upgrade options. A decision stage is defined by flight n , which includes the entire booking horizon of flight n for $n = 1, \dots, N$. Let T_n denote the number of time periods obtained from discretising booking horizon n and at most one customer may request capacity at each time period. A state at stage n , denoted by α_n , represents customers' anticipation on getting upgrades by using the option in flight n . At each state α_n , the airline decides a set of products from \mathcal{S} to display to every request arriving during the booking horizon of flight n . Note that we introduce a dynamic program formulation for these decisions made at flight (macro) level and construct the link between consecutive flights based on customer anticipation. Then in Section 4.4, we formulate the dynamic assortment planning problem at customer (micro) level during the booking horizon of each flight.

Let $\mathbf{g} \in \mathcal{G}$ denote airline's decisions during the booking horizon of flight n at state α_n , where \mathcal{G} is the search space. Specifically, \mathbf{g} includes all decisions on which set of products from \mathcal{S} should be offered to all requests arriving during the booking horizon of flight n . In other words, the airline decides the number of time periods in the booking horizon where they offer each set of products to customers. We assume that the choice behaviour for each customer is stochastic and follows a known distribution given presented cabins and upgrade option. Let $\mathbf{d} \in \mathcal{D}(\mathbf{g})$ denote one possible aggregated demand realisation for flight n under decisions \mathbf{g} at state α_n , where $\mathcal{D}(\mathbf{g})$ is a set of aggregated demand realisations. Let $P(\mathbf{d}|\mathbf{g})$ represent the probability of realising aggregated demand \mathbf{d} under policy \mathbf{g} . Note that realised demand \mathbf{d} contains information on accepted customers from both economy and

business cabins and executed options. Based on demand \mathbf{d} under decision \mathbf{g} at state α_n , customer anticipation for the next flight α'_{n+1} is updated by a function $\phi(\mathbf{d}, \alpha_n)$, that will be explicitly defined in Section 4.4. In other words, $\alpha'_{n+1} = \phi(\mathbf{d}, \alpha_n)$. We assume that the customer anticipation during the booking horizon of the first flight α_1 is initialised as a constant value. Let $R(\mathbf{d}|\mathbf{g}, \alpha_n)$ denote the total revenue obtained from demand \mathbf{d} for flight n under policy \mathbf{g} . Note that $R(\mathbf{d}|\mathbf{g}, \alpha_n)$ includes the revenue collected from flight tickets as well as the executed upgrade options. Accordingly, the value function $V_n(\alpha_n)$ at state α_n for flight n can be formulated recursively as follows:

$$V_n(\alpha_n) = \max_{\mathbf{g} \in \mathcal{G}} \left\{ \sum_{\mathbf{d} \in \mathcal{D}(\mathbf{g})} P(\mathbf{d}|\mathbf{g}) \left[R(\mathbf{d}|\mathbf{g}, \alpha_n) + \gamma V_{n+1}(\alpha'_{n+1}) \right] \right\}, \quad (4.1)$$

where γ is the discount factor on future revenue (pre-specified as $\gamma < 1$). The boundary condition of the dynamic program at flight level is defined as $V_{N+1}(\cdot) = 0$.

The capacity planning problem involves N flights with individual booking horizon. Accordingly, the dynamic program (4.1) suffers from the curse of dimensionality because of its large state space. Moreover, it is crucial to understand the distribution of demand realisation \mathbf{d} for each flight, because it influences customer anticipation for the following flight. However, identifying such a distribution is complicated as it is affected by both the assortment planning policy and the choice behaviour over all customers. Moreover, theoretically, the assortment problem should be solved for every request for flight n at each state α_n in order to compute $V_n(\alpha_n)$ in (4.1). In the next section, we will formulate the assortment planning problem at customer level as a DP by using a sequential multinomial logit choice model.

4.4 The Assortment Problem using Sequential Choice Model

In this section, we first present a DP formulation of the assortment optimisation problem and then introduce the definition of the underlying anticipation levels as well as a sequential multinomial logit (MNL) choice model. Note that we omit the stage index in state α_n in the following formulation for simplification.

Let us consider a flight n and the current customers' anticipation level at α for receiving upgrade options. The flight has fixed capacity in economy and business cabins denoted as C_e and C_b , respectively. We assume that customers are classified into a set of segments \mathcal{L} such that customers from the same segment have the same choice behaviour towards products (flight tickets and upgrade options). The arrival of customers is uncertain and it follows a homogeneous Poisson process during the

booking horizon of flight n . According to the arrival rate, the booking horizon is discretised into time points $t \in \{1_n, \dots, T_n\}$ where the probability of having more than one arrival during each time period can be neglected. After the cut-off time T_n (that may coincide with the departure time), no more request can be admitted in flight n .

A discrete-time, discrete-state MDP is used to formulate the dynamic assortment planning problem for flight n . We consider T_n stages defined by arrival of a customer request for a seat in flight n , where $t_n = 1_n, \dots, T_n$. A state of the system at stage t_n is denoted by (\mathbf{x}, y) , where $\mathbf{x} \in \mathbb{Z}_+^{[2]}$ represents the remaining capacity in economy and business cabins and $y \in \mathbb{Z}$ represents the number of purchased upgrade options.

Let λ denote the probability of one customer arriving at stage t_n and μ_l indicate the probability of the arriving customer from segment- l . As mentioned before, the offered products for customers are selected from set \mathcal{S} for customers. Let $s \in \mathcal{S}$ denote the offered products at stage t_n . Given the offered products ($s \in \mathcal{S}$) being presented at stage t_n , a segment- l customer may select a business cabin (or an economy cabin) seat with probability p_{lsb}^α (or p_{lse}^α). The state (\mathbf{x}, y) is updated as $(\mathbf{x} - \mathbf{1}_b, y)$ at stage $t_n + 1$ if a business capacity is sold. Accordingly, if the customer has purchased an economy ticket, then he/she purchase an upgrade option, when offered, with probability p_{lsu}^α . The state then becomes $(\mathbf{x} - \mathbf{1}_e, y + 1)$ at stage $t_n + 1$. If the upgrade option is not offered or the customer prefers not to purchase even though it is offered, the state is updated as $(\mathbf{x} - \mathbf{1}_e, y)$ at stage $t_n + 1$. The customer may also leave the system without purchasing any ticket with probability p_{ls0}^α . In this case, the state remains as (\mathbf{x}, y) at stage $t_n + 1$. Therefore, the value function $\Gamma_{t_n}^\alpha(\mathbf{x}, y)$ at state (\mathbf{x}, y) given anticipation α can be formulated recursively as follows:

$$\begin{aligned} \Gamma_{t_n}^\alpha(\mathbf{x}, y) = \max_{s \in \mathcal{S}} \left\{ \lambda \sum_{l \in \mathcal{L}} \mu_l \left[p_{lse}^\alpha [r_e + p_{lsu}^\alpha \Gamma_{t_n+1}^\alpha(\mathbf{x} - \mathbf{1}_e, y + 1) + (1 - p_{lsu}^\alpha) \Gamma_{t_n+1}^\alpha(\mathbf{x} - \mathbf{1}_e, y)] \right. \right. \\ \left. \left. + p_{lsb}^\alpha [r_b + \Gamma_{t_n+1}^\alpha(\mathbf{x} - \mathbf{1}_b, y)] + p_{ls0}^\alpha \Gamma_{t_n+1}^\alpha(\mathbf{x}, y) \right] + (1 - \lambda) \Gamma_{t_n+1}^\alpha(\mathbf{x}, y) \right\}. \end{aligned} \quad (4.2)$$

We define the customer anticipation α' during the booking horizon of the next flight $n + 1$ as the ratio of the number of executed upgrade options for the current flight n to the capacity of business class. This basically indicates that the variable space of customer anticipation level is a continuous interval $[0, 1]$. We can now explicitly state the function $\alpha' = \phi(\mathbf{d}, \alpha)$ in the DP model (4.1) as $\alpha' = \frac{\min\{y, x_b\}}{C_b}$.

Accordingly, the boundary condition $\Gamma_{T_n+1}^\alpha(\mathbf{x}, y)$ of the DP problem (4.2) is formulated as $\Gamma_{T_n+1}^\alpha(\mathbf{x}, y) = r_u \cdot \min\{y, x_b\} + \gamma V_{n+1}(\alpha')$ that includes the revenue received by executing upgrade options as well as the discounted total revenue to

be obtained from the remaining flights. Then, the value function at the first stage $\Gamma_{1_n}^\alpha(\mathbf{x}, y)$ consists of total revenue from selling the current flight n and the discounted total revenue from the remaining flights. When the flight-level DP model (4.1) defined in Section 4.3 is broken down to the customer-level, we can obtain $V_n(\alpha) = \Gamma_{1_n}^\alpha(\mathbf{x}, y)$. Therefore, we can propose an approach to approximate $V_n(\alpha)$ by exploiting the structure of the customer level DP model (4.2).

The Choice Model: Given the definition of α , we first describe the sequential customer choice model defining probabilities in (4.2). Specifically, we discuss two alternative approaches of defining the utilities used in the choice model. The complexity of the DP model (4.2) will be revisited after we introduce the choice model.

On the basis of the customer's anticipation about receiving upgrade options driven by his/her past experience, we assume that the customer could buy (or not) an upgrade option. Therefore, the increase in anticipation ratio will lead to the increase in utility of economy capacity and the increase in utility of upgrade options (full anticipation effect). Let us define utilities gained by the customer request arriving from segment l for each product type (including upgrade options). Let β_{lb} , β_{le} and β_{lu} denote utilities gained by a segment- l customer for economy, business cabins and upgrade options, respectively. Let σ_l denote a sensitivity parameter with respect to the given anticipation level α for a segment- l customer. Given the anticipation level α , the utility functions of economy cabin, business cabin and upgrade option for a segment- l customer are defined as follows:

$$\begin{aligned} v_{lb}(\alpha) &= \beta_{lb} - \sigma_l \alpha + \epsilon, \\ v_{le}(\alpha) &= \beta_{le} + \sigma_l \alpha + \epsilon, \text{ and} \\ v_{lu}(\alpha) &= \beta_{lu} + \sigma_l \alpha + \epsilon, \end{aligned} \tag{4.3}$$

where ϵ represents the error term following a Gumbel distribution with mean 0. Notice that when α approaches zero ($\alpha \rightarrow 0$), the anticipation level has no influence on utilities of any product. On the other hand, when α approaches one ($\alpha \rightarrow 1$), the anticipation increases the utility of economy cabin as well as the upgrade options, but decreases the utility of business cabin accordingly. When the option execution rate for the previous flight increases, customers' anticipation for getting upgrade via options increases so that customers prefer to purchase the economy class and also book the upgrade option if it is offered.

Alternatively, as upgrade options can be refunded at the end of the booking horizon, customer purchase behaviour might not be affected by their experience. The value of the business class ticket should also be independent of the anticipation.

Therefore, we assume that increase in anticipation ratio will only lead to the increase in utility of economy capacity (partial anticipation effect). Moreover, we assume that the utility value of an option should depend on the utility difference between business class and economy class. Accordingly, given the utility parameters introduced previously, the utility functions of economy cabin, business cabin and upgrade option for a segment- l customer at anticipation level α are defined as follows:

$$\begin{aligned} v_{lb}(\alpha) &= \beta_{lb} + \epsilon, \\ v_{le}(\alpha) &= \beta_{le} + \sigma_l \alpha + \epsilon, \text{ and} \\ v_{lu}(\alpha) &= \beta_{lb} - \beta_{le} - \beta_{lu} + \epsilon, \end{aligned} \tag{4.4}$$

where β_{lu} indicate the utility change from the cost of options. Notice that when α approaches zero ($\alpha \rightarrow 0$), the anticipation level has no influence on utility of any product. On the other hand, when α approaches one ($\alpha \rightarrow 1$), the anticipation only increases the utility of economy cabin.

As mentioned before, the upgrade option cannot be observed by the customer until an economy flight ticket is purchased. Thus, we employ a sequential choice model to reflect an unobservable feature of the upgrade options, using a sequential multinomial logit (MNL) choice model. We assume that no-purchase has 0 utility for all customers. Given the offered products $s \in \mathcal{S}$, let \mathcal{S}_s^1 and \mathcal{S}_s^2 define two sets of products to be offered at the first and second stages, respectively. For instance, if $s = \{e, b, u\}$, then $\mathcal{S}_s^1 = \{e, b\}$ and $\mathcal{S}_s^2 = \{u\}$. Note that $\mathcal{S}_s^2 = \emptyset$ if the upgrade option is not included in the offered products s . Given α , the probability of segment- l customers selecting product i from the offered products $s \in \mathcal{S}$ is computed as

$$p_{lsi}^\alpha = \begin{cases} \frac{\exp(v_{li}(\alpha))}{\sum_{j \in \mathcal{S}_s^1} \exp(v_{lj}(\alpha)) + 1}, & \text{if } i \in \mathcal{S}_s^1, \\ \frac{\exp(v_{le}(\alpha))}{\sum_{j \in \mathcal{S}_s^1} \exp(v_{lj}(\alpha)) + 1} \times \frac{\exp(v_{li}(\alpha))}{\exp(v_{li}(\alpha)) + 1}, & \text{if } i \in \mathcal{S}_s^2. \end{cases}$$

Based on the definition of those choice probabilities, as $p_{lse}^\alpha + p_{lsb}^\alpha + p_{ls0}^\alpha = 1$,

we can rewrite the customer-level DP problem in (4.2) as follows:

$$\begin{aligned} \Gamma_{t_n}^\alpha(\mathbf{x}, y) = \max_{s \in \mathcal{S}} \Bigg\{ & (1 - \lambda) \Gamma_{t_n+1}^\alpha(\mathbf{x}, y) + \lambda \sum_{l \in \mathcal{L}} \mu_l [p_{lse}^\alpha [r_e + p_{lsu}^\alpha (\Gamma_{t_n+1}^\alpha(\mathbf{x} - \mathbf{1}_e, y + 1) \\ & - \Gamma_{t_n+1}^\alpha(\mathbf{x} - \mathbf{1}_e, y)) + \Gamma_{t_n+1}^\alpha(\mathbf{x} - \mathbf{1}_e, y) - \Gamma_{t_n+1}^\alpha(\mathbf{x}, y)] \\ & + p_{lsb}^\alpha [r_b + \Gamma_{t_n+1}^\alpha(\mathbf{x} - \mathbf{1}_b, y) - \Gamma_{t_n+1}^\alpha(\mathbf{x}, y)] + \Gamma_{t_n+1}^\alpha(\mathbf{x}, y) \Bigg\}. \end{aligned} \quad (4.5)$$

The DP model (4.5) is also intractable because of its large state space. However, this formulation indicates a possible approach to derive a time slot assortment policy. If we had the marginal revenue for accepting a customer into economy cabin (as $\Gamma_{t_n+1}^\alpha(\mathbf{x} - \mathbf{1}_e, y) - \Gamma_{t_n+1}^\alpha(\mathbf{x}, y)$), business cabin (as $\Gamma_{t_n+1}^\alpha(\mathbf{x} - \mathbf{1}_b, y) - \Gamma_{t_n+1}^\alpha(\mathbf{x}, y)$) and upgrade option (as $\Gamma_{t_n+1}^\alpha(\mathbf{x} - \mathbf{1}_e, y + 1) - \Gamma_{t_n+1}^\alpha(\mathbf{x} - \mathbf{1}_e, y)$), we should be able to find an assortment policy by exactly computing $V_n(\alpha)$ defined in (4.1). Nonetheless, since we want to maximise the discounted total revenue from a number of flights in the problem, the approximated marginal revenue for each individual customer would be too small that can be neglected. Therefore, it would be reasonable to consider the capacity allocation problem at flight level rather than focus on the assortment planning problem at customer level in order to approximate $V_n(\alpha)$. Furthermore, due to our definition on α , the DP problem (4.1) has a continuous state space. To overcome this issue, in Section 4.5, we first discretize the state space in (4.1), and then introduce a choice-based method to approximate the total revenue of flight n . This approximation method provides the expected number of executed upgrade options so that we can calculate the expected customer anticipation for flight $n + 1$.

4.5 A Choice-based Linear Approximation Model

In this section, we focus on a solution method for the DP model (4.1) by approximating the revenue gained from each flight under a given customer anticipation level. First of all, the state space of (4.1) is continuous, $\alpha \in [0, 1]$, due to the definition of customer anticipation in Section 4.4. Accordingly, we can discretize the state space into K discrete states and denote the set of discrete anticipation levels as $\mathcal{A} = \{\hat{\alpha}_1, \dots, \hat{\alpha}_K\}$ for all stages. Note that states in \mathcal{A} are maintained in ascending order. Since the state space only involves a finite number of discrete points due to discretization, we can use the backward recursion algorithm to estimate the value function for flight n at each state $\hat{\alpha}_k$, $k = 1, \dots, K$.

We adopt a choice-based deterministic model to approximate the expected revenue obtained from one flight. Given the anticipation level α for flight n , we can compute the expected revenue when a segment- l customer is presented with offered

products s as

$$r_{ls}^\alpha = \sum_{i \in \mathcal{S}_s^1} r_i \times p_{lsi}^\alpha.$$

Note that \mathcal{S}_s^1 does not involve upgrade options, so the expected revenue only computes revenue from selling the capacities of economy and business cabins, but not the revenue from executing upgrade options. Let $\mathbf{h} = \{h_s \mid \forall s \in \mathcal{S}\}$ denote a vector of decision variables representing the number of time periods that customers are presented with offered products $s \in \mathcal{S}$ during the booking horizon of flight n . Recall that the airline can also reject a customer's request by not offering any product, which is indicated by \emptyset in \mathcal{S} . Let w denote a decision variable representing the number of upgrade options executed at the end of booking horizon of flight n . Given w , we can exactly compute the updated anticipation for flight $n+1$ as $\hat{\alpha} = \frac{w}{C_b}$ which is continuous value arising between 0 and 1.

We assume that $V_{n+1}(\frac{w}{C_b})$ can be approximated by a piece-wise linear function defined by discrete states in \mathcal{A} and the corresponding value functions are denoted by $V_{n+1}(\hat{\alpha}_k)$. We use a special order set (SOS) constraint to model the continuous piece-wise linear function by introducing non-negative continuous variables z_k for $k = 1, \dots, K$ and binary variables y_k for $k = 1, \dots, K$. Note that y_k takes 1 if z_k is non-zero, and 0 otherwise. Let Ω_i indicate one unique pair of non-adjacent z_k , where $i = 1, \dots, {}^2C_K - K + 1$ for given K discrete anticipation levels. Note that 2C_K means the number of 2-combinations from K elements. For example, $\Omega_i = \{1, 3\}$ represents two non-adjacent z_1 and z_3 . Accordingly, value function $V_n(\alpha)$ can be approximated by the following integer optimization problem developed for flight n

at anticipation level α :

$$\begin{aligned}
CDIP_n(\alpha) : \quad & \max_{\mathbf{h}, w, y_k, z_k} \sum_{l \in \mathcal{L}} \mu_l \sum_{s \in \mathcal{S}} \bar{r}_{ls} h_s + r_u w + \sum_{k=1}^K z_k V_{n+1}(\hat{\alpha}_k) \\
\text{s.t.} \quad & \sum_{s \in \mathcal{S}} h_s = \lambda T_n, \\
& \sum_{l \in \mathcal{L}} \mu_l \sum_{s \in \mathcal{S}} p_{lse}^\alpha h_s \leq C_e, \\
& \sum_{l \in \mathcal{L}} \mu_l \sum_{s \in \mathcal{S}} p_{lsb}^\alpha h_s + w \leq C_b, \\
& w \leq \sum_{l \in \mathcal{L}} \mu_l \sum_{s \in \mathcal{S}} p_{lsu}^\alpha h_s, \\
& \sum_{k=1}^K z_k \hat{\alpha}_k = \frac{w}{C_b}, \\
& \sum_{k=1}^K z_k = 1, \\
& z_k \leq y_k, \quad k = 1, \dots, K, \\
& \sum_{k=1}^K y_k = 2, \\
& \sum_{j \in \Omega_i} y_j \leq 1, \quad i = 1, \dots, {}^2C_K - K + 1, \\
& \mathbf{h} \geq 0, \quad w \in \mathbb{Z}^+, \\
& z_k \geq 0, \quad y_k \in \{0, 1\}, \quad k = 1, \dots, K.
\end{aligned} \tag{4.6}$$

The first constraint in $CDIP_n(\alpha)$ describes the capacity allocation with respect to the total demand. The second and third constraints, respectively, state that the seat allocation for the economy and business cabins cannot exceed the capacity. The fourth constraint limits the demand on the number of executed upgrade options. The next five sets of constraints impose SOS-2 restrictions: namely, i) at most two of z_k for $k = 1, \dots, K$ can be nonzero, and ii) these two nonzero z_k must be adjacent, which leads to ${}^2C_K + 4$ integer constraints.

The complexity of solving $CDIP_n(\alpha)$ is influenced by the number of integer constraints, which depends on the number of discrete anticipation levels. As anticipation level is assumed to range between 0 and 1, we maintain a reasonable number of discrete anticipation levels, such as $K = 10$ in our experiments. Based on backward recursion algorithm, we can maintain a table keeping the $CDIP_n(\alpha)$ -based estimations to the value function at each discrete state for all flights. According to the maintained table, given anticipation level α for flight n , the assortment planning

policy \mathbf{g} defined in DP (4.1) can be expressed by the optimal solution \mathbf{h}^* obtained by solving $CDIP_n(\alpha)$.

4.6 Numerical Experiments

In this section, we first describe design of experiments and data issue. Then, we present our computational results and analysis of these results. We design a set of numerical experiments to illustrate performance of the capacity planning models and impact of upgrade options on airline revenue management. In particular, we aim to show the importance of accounting customers anticipation under the use of upgrade options over a long-term planning horizon. Finally, we numerically investigate a steady level of customer anticipation that the airline wishes to achieve to maximise the total revenue at the given demand.

4.6.1 Design of Experiments and Data

We consider 200 single-leg flights of an airline that provides upgrade options. These flights are identical in terms of capacity of economy and business cabins. On the basis of the seating plan of an aircraft A380 operated by British Airways, we assume that all flights have 97 and 313 seats in business and economy cabins, respectively. The prices of flight tickets in economy and business cabins, respectively, are set to be £50 and £150. An upgrade option is sold at price of £70 if offered. The booking horizon of each flight consists of a number of discrete time periods. For simplicity, we assume that there is exactly one customer requesting the flight capacity at each time period.

Policies: The following three policies that the airline can adopt to allocate flight capacity are obtained for the policy comparison purposes:

- **Benchmark Policy (BP):** We consider a benchmark policy to highlight the value of upgrade option in airline RM. Under this policy, we assume that the airline does not offer upgrade options to customers at all. Moreover, the capacity allocation decisions are made for each flight to maximise the expected revenue of the single flight based on the customer choice model without having upgrade options.
- **Independent Policy (IP):** Customers' anticipation on upgrade option is not taken into account when computing the policy; thus, the third component in the objective function of $CDIP_n(\alpha)$ (4.6) is set to zero. The capacity allocation decisions are made independently for each flight to maximise the total revenue

of the flight, but customer anticipation is implicitly considered in the customer choice behaviour during the simulation process.

- **Anticipation-integrated Policy (AP):** The policy is obtained by the capacity planning problem at flight level as proposed in this chapter. We involve customers' anticipation on upgrade option into the policy while managing the capacity of a sequence of identical flights. The capacity allocation decisions for each flight are made to maximise the discounted long-term total revenue as formulated in (4.1). The discount factor is set to be 0.9 in the numerical experiments.

Note that overbooking is not considered in all policies described above.

The Customer Choice Model: Without the support of real data, we conduct simulations under two utility definitions in (4.3) and (4.4), respectively. Under each assumption, we consider one customer segment for our numerical experiments. We assume that customers from this segment prefer the economy cabin to business cabin, and this is captured by the MNL choice model. The utilities of no-purchase, the economy and business seats, as well as upgrade option, are presented in Table 4.1. According to the market research from Mintel (2019), around 80% of passengers in UK fly with economy class. Therefore, utility parameters used in our experiments are set in such a way that the conditional probability of a customer booking business cabin given the customer has booked a ticket is around 80% when both cabins are presented.

Economy	Business	Upgrade option	No-purchase
-1.5	-2.5	0.3	0.0

Table 4.1: Utility parameters in the MNL choice model

Customers' anticipation on upgrade option is scaled by sensitivity parameter (σ_l) to reflect its impact in the choice models. Note that utilities of both normal flight tickets and upgrade option have the same sensitivity parameter on customers' anticipation, but their impact is differentiated by the (+ or $-$) sign in utilities. As we wish to compare the importance of customer anticipation under different demand scenarios, we focus on a fixed sensitivity parameter $\sigma_l = 1.0$.

Simulation Experiments: We conduct simulation experiments to evaluate the benefit of customer anticipation in managing the cabin capacity with upgrade options in the long term. In this study, we assume that each flight has the same arrival process and the booking horizon of individual flight involves the same number of time periods. As a base demand scenario, we use 1800 time periods so that

the ratio of expected unconstrained demand to the total capacity (410 this case) is around 1. Since upgrade options can be only beneficial when demand is lower than the capacity, we vary the scaling parameters between 0.5 and 1.1 with the incremental step 0.1 in the base demand scenario and evaluate our policies under different demand levels.

The main steps of our simulation experiments are briefly explained below. Each simulation consists of simulating customer arrivals during the booking horizon of 200 flights, where the customer anticipation is always set to zero for the first flight. Given the current level of customers' anticipation, we determine capacity planning decisions for each flight in terms of the number of time periods of offering each set of products by solving the $CDIP_n(\alpha)$ problem (4.6). During the booking horizon, one set of products is offered to each customer in the sequence of $\{e, b, u\}$, $\{e, u\}$, $\{e, b\}$, $\{e\}$, and $\{b\}$. We start to reject customers after there is no remaining time period of offering $\{b\}$. Based on the offered set of products, the customer booking decision is randomly generated based on the choice model.

At the end of the booking horizon of each flight, we make adjustments for the accepted customer requests for both economy and business cabins. We first deal with the overbooked seat for business cabin. Those customers overbooked the business cabin are rejected. If the business cabin has more remaining capacity, then we start executing options. In other words, customers who have already purchased options are to be upgraded as long as the business cabin's capacity permits; thus, more capacity in the economy cabin is created. At the final stage, the economy cabin customers are rejected if the economy cabin is overbooked, after executing the upgrade options. The associated revenue from flight tickets and upgrade options is deducted from the total revenue if any customer request is rejected. After the adjustment process, we can compute the customer anticipation level during the booking horizon of the next flight based on the number of executed upgrade options. This procedure is carried out for 200 flights in the same manner by simulating their corresponding booking horizons. In order to obtain robust results, we repeat the simulation process for 100 runs and use the average results to evaluate the performance of those policies under each scaling parameter.

4.6.2 Computational Results and Discussion

We are firstly concerned with investigating the financial benefit of upgrades as one of products for the airline's revenue management. Then we aim to illustrate the potential impact of taking the customer anticipation into account when the airline makes the capacity allocation decisions over a long planning horizon. Figure 4.1 presents the average revenue increase in one flight under two policies AP and IP with

respect to BP at different demand scaling parameters. Under both utility definitions (4.3) and (4.4), we observe that offering upgrade options as additional products can generate more revenue under AP and IP when the demand scaling parameters are smaller than 1.0. The result seems counter-intuitive as using options is cheap/inferior way of the airline to sell business capacity and it should cause revenue loss. However, with the underlying MNL model, when anticipation is included in either of the utility definition, introducing upgrade option could also attract more people to purchase the capacity. In other words, the probability of no purchase may decrease when anticipation ratio increases under our choice model. For example, when anticipation has full effect on utilities defined in (4.3) and anticipation level is 0.5 under our simulation setting, the no-purchase probability at the first stage increases by 5% compared with the case of anticipation level of 0. When anticipation has partial effect on utilities defined in (4.4) and anticipation level is 0.5, the no-purchase probability at the first stage would increase by 8% compared with the case of anticipation level of 0. Therefore, under both utility definitions, we can observe additional revenue is obtained through offering upgrade options under AP and IP. Note that more utility is added to the case where customers will purchase under the partial anticipation effect. However, under full anticipation effect, some utility are removed from business and some are added to economy. As utility of no-purchase stays the same, no-purchase probability reduces more if the utility of overall purchase behaviour increases more largely. Therefore, the benefit of considering anticipation has larger impact on no-purchase probability in the partial anticipation effect than in the full effect.

As the demand scaling parameter increases (from 1.0 onwards) in Figure 4.1, there is no revenue benefit of offering upgrades in both AP and IP under both assumptions. These results confirm that upgrade options should be offered only when low demand realises. It intuitively makes sense that upgrade options should be no longer offered when there is enough demand to fill up the business cabin at full price. Figure 4.1 also shows that AP is able to create more revenue increase than IP when the demand is less than the capacity. For instance, the gap between revenue increase achieved by two policy is around 9% under when anticipation has full effect on utilities and the demand scaling parameter is set to 0.6. Note that all findings obtained under partial anticipation effect are consistent with results under full anticipation effect, so we focus on our discussion on the results under definition (4.3) with full anticipation effect in the remaining section.

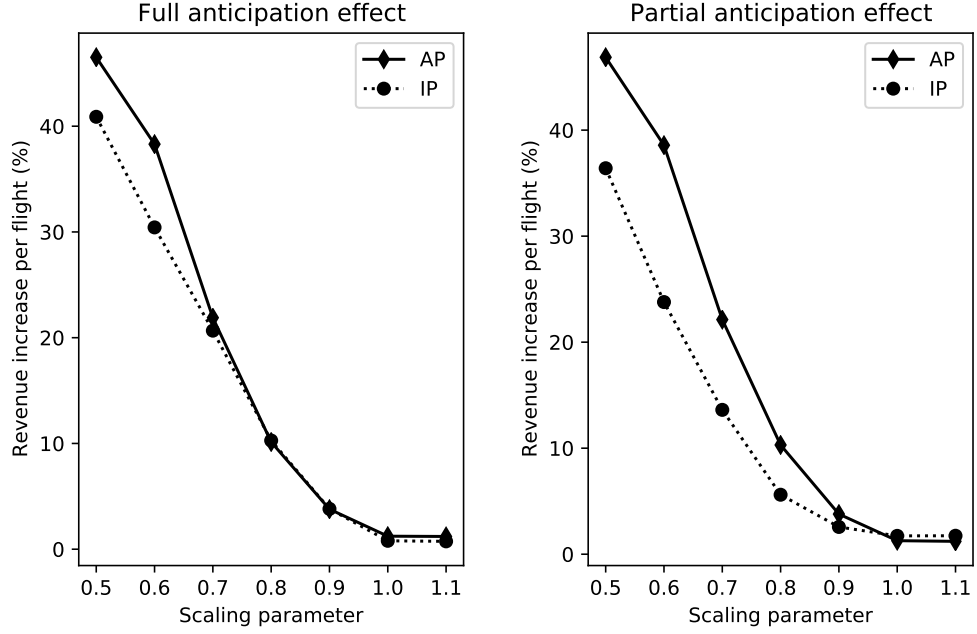


Figure 4.1: Impact of considering upgrade options and customer anticipation under full and partial anticipation effects

We now illustrate what percentage revenue increase is received from selling capacity in business and economy cabins and upgrade options by comparing two policies AP and IP in Figure 4.2. When the demand scaling parameter increases from 0.9 onwards, AP obtains more revenue from selling business class tickets and less revenue from executing upgrade option compared to IP. Moreover, it is worthwhile mentioning that AP does not sell business capacity at full price when the demand scaling parameters are set to 0.5 and 0.6. Under our setting, AP generally sells less business capacity at full prices, but achieves higher revenue than IP by selling economy capacity and upgrade options. Especially, when demand is low (such as 0.5 and 0.6), AP produces higher profit as it sells more business capacity via upgrade options rather than business cabin tickets. Therefore, in comparison with IP, AP achieves a better trade-off between promoting upgrade options by manipulating customer anticipation and reserving the full-priced business capacity.

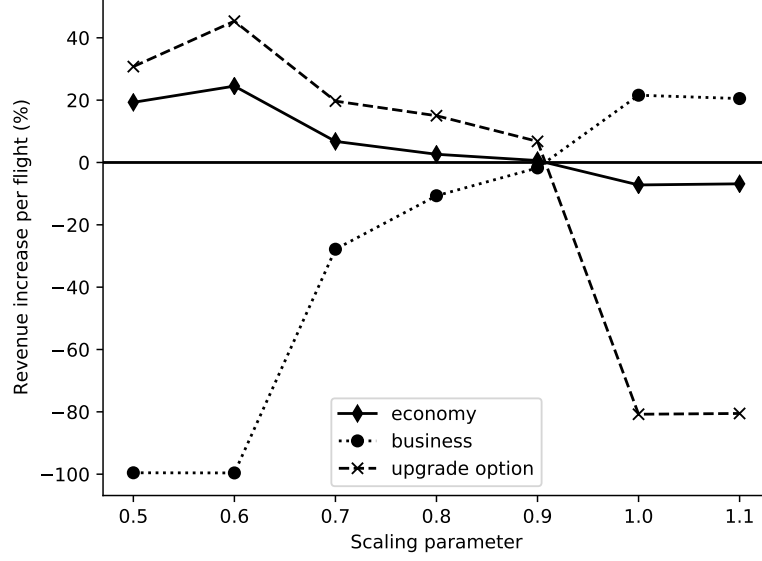


Figure 4.2: Performance comparison of AP and IP in terms of percentage revenue increase (%) generated from selling economy and business tickets and upgrade options

As the airline considers the long-term revenue over 200 flights, we intend to numerically investigate how the customer anticipation evolves at each flight. Figure 4.3 illustrates the evolution of customer anticipation under both policies, AP and IP, for each individual flight when the demand scaling parameter is set to 0.7. Given the same initial customer anticipation level at zero for the first flight, customer anticipation level is stabilised at different levels under two policies. Recall that the demand level remains the same for all flights in our computational experiments. As IP always makes the same capacity allocation decision based on customer anticipation as zero for all flights, it allows customers gradually learn about the execution rate of upgrade option. This leads to a stable customer anticipation level as 0.6 in this case. On the other hand, AP also achieves a stable anticipation level at around 0.7. As a result, we can conclude that there exists a steady state (customer anticipation level) in the DP problem (4.1) based on the $CDIP_n(\alpha)$ approximation (4.6) for a given the demand level.

Furthermore, we investigate the steady state conditions obtained by the DP problem (4.1) at varying demand levels. Figure 4.4 presents how customer anticipation evolves for each flight under AP at fixed demand scaling parameters as 0.7, 0.8 and 0.9. We can observe that the steady state exists at different demand levels. Moreover, we notice that the customer anticipation level decreases as the demand scaling parameter increases. For instance, the customer anticipation level is

around 0.2 when the demand scaling parameter is set at 0.9. It increases to 0.7 when demand scaling parameter is 0.7. Based on our assumption, one can easily state that the higher anticipation is, the more likelihood that a customer books economy cabin and purchase the upgrade option. Therefore, it intuitively makes sense that AP manipulates customer anticipation at a higher level at low demand case (when the scaling parameter is 0.7) in order to fill up the flight and obtain higher revenue than at high demand case (when the scaling parameter is 0.9).

Theoretically, because anticipation can divert customers from business to economy, anticipation level of 0 should always lead to a higher revenue than the case where anticipation level is above 0. However, due to the underlying MNL choice model, including anticipation in the utilities not only moves customers from business to economy but also could attract more people to purchase the capacity. Therefore, our simulation study shows that the steady anticipation level depends on the demand level as shown in Figure 4.4. When customer arrival rate is relatively low to the total capacity, the company may wish to maintain the anticipation rate at certain level such that more customers are attracted and customers may ‘purchase’ upgrade options.




Figure 4.3: Customer anticipation levels achieved by AP and IP for each individual flight at demand scaling parameter 0.7

Finally, we would like to report a numerical evidence that the existence of a steady state (customer anticipation) achieved by the DP model (4.1) based on the $CDIP_n(\alpha)$ approximation (4.6) does not depend on the choice of the initial

anticipation level. For this experiment, we fix the demand scaling parameter as 0.9 and the initial anticipation level for the first flight is chosen as 0.0, 0.5 and 1.0. As it can be easily seen from Figure 4.5, the same customer anticipation level (around 0.2) is obtained regardless of the choice of initial anticipation level set for the first flight. Moreover, when the initial customer anticipation level is selected as smaller than the steady anticipation level, the DP model (4.1) achieves the steady anticipation immediately after the first flight. However, when the initial anticipation is higher than the steady state level, it takes a number of flight to reduce the customer anticipation to the steady level. Such observation indicates that it would be easier to increase the customer anticipation rather than decrease it within a short term.




Figure 4.4: Customer anticipation levels maintained during the booking horizon of each flight under AP




Figure 4.5: Customer anticipation levels maintained by AP during the booking horizon of each flight given different initial anticipation level

4.7 Conclusions

In this chapter, we have studied the airline’s capacity allocation problem involving both economy and business cabins, where upgrade options are offered to address the demand-capacity mismatch in both cabins. We specifically focus on a long term capacity planning problem over a number of flights, where customers may take advantage of upgrade options to obtain capacity in business cabin based on their anticipation. We introduce a DP formulation for the capacity allocation problem for all flights. As a solution method, we propose a value function approximation method by discretizing state space on customer anticipation and apply the backward recursion algorithm to compute the value function at each state. Given the anticipation level, we adopt a choice-based approach to estimate the revenue of the current flight and calculate the anticipation for the next flight based on expected number of executed upgrade options. The capacity allocation decisions for a flight are also decided by solving the CDIP with current anticipation level based on the estimated value functions at each state from next stage (or flight).

We design a series of computational experiments to demonstrate performance of the proposed policies and present the benefit of introducing upgrade options. Then, we address the importance of including customer anticipation in the capacity allocation policy, since it is able to achieve a better trade-off between reserving

capacity in business cabin and promoting upgrade options given the demand. Finally, we numerically investigate the existence of a steady level of customer anticipation. Our results indicate that this steady state is only influenced by the underlying demand, but not the initial anticipation level.

Due to the lack of real data, conclusions derived from the numerical study are based on two different utility definitions and a specific parameter setting. Two definitions mainly distinguish with each other on whether anticipation has influence on the utility of upgrade option. However, under both definitions, including anticipation in the MNL choice model can not only divert customer from business cabin to economy cabin but also increase customer demand. Therefore, similar conclusions are reached as more revenue are obtained from offering upgrade options. It requires further study on the choice model to evaluate the revenue advantage of using upgrade option where anticipation cannot increase customer demand.

Since we introduce a naive definition on customer anticipation as the percentage of business capacity used by options, there are also issues in whether it truly reflects customer behaviour in the future purchase or not. In particular, when all upgrade options are executed and there is still available capacity in the business cabin, the anticipation level is calculated less than 1 by our approach. However, customer anticipation level should be the highest, 1 in our case, as all upgrade options are executed. Therefore, a more appropriate numerically approximating customer anticipation, such as using the probability of a customer obtaining an upgrade option, needs to be further explored. Note that using such definition of anticipation level may lead to a case that our proposed solution method is no longer applicable. It might be necessary to propose a solution approach for the new approximation method of anticipation.

On the other hand, the assumption that booking horizons of flights are non-overlapping is a very critical assumption made to simplify our problem such that a LP-based approximation approach can be proposed. However, if we relax such assumption, our conclusion on the benefit of including customer anticipation in our capacity planning for consecutive flights with upgrade options will be still valid. Since flight booking horizons are overlapping, there would be less time for each flight to adjust its capacity allocation decision with respect to the latest anticipation level. Accordingly, the revenue for each flight might decrease but the benefit of recognising customer anticipation should remain significant.

As future study, in addition to theoretically proving the existence of a steady state, we aim to establish under which condition a steady state exists by relaxing the underlying assumptions. Moreover, one can investigate capacity allocation policy based on the structural property of the steady state.

Chapter 5

Summary of Thesis and Future Work

In this chapter, we briefly summarise our concluding remarks and the main findings obtained from research topics studied in this PhD thesis. We then highlight limitations of our studies. Finally, some potential research directions will be presented as continuation and extension of this PhD study.

5.1 Summary of Research Questions

This thesis covers three research topics on innovative products and services provided in attended home delivery services and airline upgrades. Specifically, we propose dynamic control policies to manage these innovative products and services in order to evaluate the benefit of introducing each of these products and services in terms of improving business revenue and operations.

In Chapter 2, we introduce the concept of flexible time slots in attended home delivery services. When purchasing a flexible slot, the customer is notified just before the delivery day in which regular time window the delivery arrives in exchange for cheap delivery charges. We aim to evaluate the profit improvement after introducing flexible slots with an appropriate pricing policy. Chapter 3 focuses on a dynamic slotting problem under a practical scenario in offering attended home delivery services, where customers are presented with delivery time slots across multiple days. We assume that customers compare time slots as well as delivery days when selecting their slots. Different choice models can be adopted in the slotting policy and we want to identify the choice model which can help the policy generating the highest total profit. Chapter 4 discusses a long-term airlines' capacity allocation problem with a sequence of flights, where upgrade services are offered via options

and upgrade charges are only taken when options are executed for one flight. This research focuses on a specific group of customers who have repeat purchase behaviour on a flight and they can learn from their past purchase experiences with upgrade options and change their choice behaviour accordingly. We intend to verify whether more revenue could be generated by the airline after introducing upgrade options under the consideration of customer anticipation.

5.2 Contributions to Practice

When new products and services are introduced to the market, companies need to decide the pricing policy and/or availability policy for these products and services. Based on these policies, companies would then evaluate the impact of those innovative products on the existing products such that policies can be modified to improve the profitability of the business.

In Chapter 2, we numerically present the benefit of introducing flexible slots in attended home delivery services with respect to increasing e-retailers' profit and reducing delivery costs. Specifically, we construct flexible slots from standard slots by popularity (P3) and total profit increases by at least round 4% after introducing flexible slots. When demand is low relative to the delivery capacity, using flexible slots has significant potential in reducing delivery costs, e.g. delivery costs per order decreases by around 1.5% after having flexible slots when demand to capacity ratio is lower than 0.8. Moreover, we also introduce flexible slots by simply merging adjacent slots (A4). Offering these flexible slots also increases the e-retailer's profit by at least 2%. Compared to P3, using these slots has the advantage of being able to spread customers more evenly across the delivery time slots, especially, if demand is high relative to delivery capacity.

We compare the performance of using different choice models in the slotting policy when offering time slots across multiple days in Chapter 3. Firstly, we find that the choice model excluding the substitution effect of days over-estimated the customers' no-purchase probability by around 30%. However, we reach a conclusion in the numerical experiments that using such a choice model could potentially increase profit by almost 5% compared to the case of offering all feasible slots, when the demand is low relative to the delivery capacity. However, using this choice model might lead to inefficient delivery schedules. Delivery costs per order could have been reduced further by around 5% if the true choice model is used, when demand to capacity ratio is low. Therefore, we recommend that e-retailer should use a choice model closely capturing the underlying customer choice behaviour in order to achieve an efficient delivery plan with low delivery cost per order.

In Chapter 4, we numerically address the benefit of introducing upgrade options in generating potential revenue to the airline. When the customer demand is relatively low to its capacity, revenue of each flight is increased by around 40%. However, when customer demand approaches capacity, there is merely no additional revenue created for each flight as business cabin capacity can be sold at full price. Moreover, our results suggest that the airline could potentially generate higher revenue (around 10% revenue increase) once its capacity allocation policy includes customer anticipation. Accounting for customer anticipation in the capacity allocation policy is able to achieve a better trade-off between reserving capacity in business cabin and promoting upgrade options given the demand.

5.3 Contributions to Theory

RM problems discussed in the thesis are challenging as they all deal with the uncertainty arising from the stochastic customer arrival process and customer choice behaviour. Solving these problems involves estimating customer choice models and anticipating future revenue. Moreover, it requires sophisticated methodologies that can solve problems in real time.

In order to evaluate the benefit of using flexible time slots in attended home delivery services, we propose a dynamic choice-based pricing policy in Chapter 2. Specifically, we introduce a novel LP-based opportunity cost estimation approach, which includes both delivery costs and displacement cost. Given estimated opportunity costs, we also derive an LP-based pricing policy such that slot prices can be decided in milliseconds.

In Chapter 3, we firstly formulate the problem of offering time slots across multiple consecutive delivery days as a dynamic program. Then, we assume that customers compare time slots as well as delivery days when selecting their slots. Depending on whether the substitution effect of delivery days is accounted, we propose three different choice models to capture customer choice behaviour. Finally, we construct three slotting policies with three choice models accordingly and each policy decides slot availability in real time.

In Chapter 4, we introduce the concept customer anticipation to model such learning process between flights as the result of introducing upgrade options to customers. Then, under the consideration of customer anticipation, we propose a novel dynamic policy to decide capacity allocations for a sequence of single-leg flights. Based on our approximation method of customer anticipation, we finally develop a capacity allocation policy for each flight by discretising the state space of DP and estimating the value function using backward aggregation.

5.4 Limitation of Thesis

There are limitations of policies and methodologies introduced in this thesis. Firstly, we mainly use LP approximation methods to approximate value functions in our DP models as solution methods. Using more refined approximation methods, such as an approximated dynamic programming, may improve the performance of our proposed policies. For research topics in attended home delivery services, more sophisticated routing heuristics could be adopted in order to improve computational efficiency in checking time slot feasibility and constructing delivery plans.

Moreover, the performance of all those policies determined in this thesis is evaluated under specific numerical settings. For example, when customers have pessimistic attitudes towards flexible slots, the e-retailer would have to reduce price to attract customers. Then, the e-retailer might face more the revenue loss than the cost reduction from delivery flexibility, which leads to a total profit decrease. There are also issues in defining the true choice model when customers are selecting slots from a number of delivery days. Customers might cluster slots with respect to morning and afternoon. When upgrade options are offered by the airline, the relationship between anticipation and utilities of products is debatable. The utility of business cabin and the probability of purchasing upgrade option could be independent from customer purchase experience (anticipation). The definition of anticipation would be more appropriate to reflect a probability of an upgrade option being offered and executed. Due to a lack of data, we are not able to test and evaluate our solution methods with real cases. As continuation to the PhD study, we would like to focus on applying our proposed policies with real data as case studies. By analysing real data, we firstly would be able to derive some insights regarding to customer arrival process and customer choice behaviour. Then, we could modify our solution approaches with respect to the information revealed from the data and evaluate the performance of our approaches in terms of improving business revenue and operation.

5.5 Future Work

This thesis specifically discusses research topics on revenue management in last-mile logistics and airlines. Potential research directions under these two areas are listed and discussed as follows:

- **Innovative delivery technologies:** Non-conventional vehicles have been introduced for last-mile delivery in urban areas, such as electric cargo bikes, parcel copters and self-driving parcels (Slabinac et al., 2015). It could lead to a mixture of delivery fleet for the e-retailer and might bring additional

complexity regarding to manage demand for each type of fleet. Once the mixture of delivery vehicle types is determined, we could investigate a dynamic pricing policy to influence customers' choice on delivery fleet such that the total profit of the delivery operation can be maximised.

- **Capacity management in crowd-shipping:** Crowd-shipping has been introduced to last-mile delivery as a platform, which connects the customer wanting to send a parcel with the driver willing to fulfil the delivery. As a platform, given the delivery demand, it is important to understand incentives of drivers to participate in a crowd-shipping system, such as drivers' willingness to work (Yildiz and Savelsbergh, 2019). In order to maximise the total profit of the crowd-shipping platform, we can focus on a dynamic matching policy that decides which to-be-delivered requests are shown every time a driver is asking for delivery opportunity.
- **Strategic customer behaviour:** In airlines, capacities of one flight are frequently updated over the booking horizon (Büsing et al., 2019). When a choice-based revenue management model is adopted in this scenario, capacities in low fare class are released to customers approaching the end of the booking horizon. In the long term, customers are able to anticipate the time of capacity update and strategically postpone their purchase in exchange for potential low fare tickets. If the customer has high anticipation, the likelihood of him/her leave the system with no purchase increases. This anticipation is also time dependent as customer anticipation may increases approaching the end of the booking horizon of a flight. Further research could consider customer anticipation within a choice model to effectively address strategic customers when setting the capacities for airline fare classes.

Remark: The main issue in last-mile logistics is the high delivery cost, so revenue management concepts have been adopted to influence customers' preference towards delivery options in order to improve the business operation and reduce delivery cost. With the development of new delivery technologies (systems), such as crowd-shipping platforms and automated electronic vehicles, the complexity from the transportation side brings challenges in anticipating delivery cost while applying pricing and capacity control mechanisms to manage the logistic service demand. For instance, a customer may require both pick-up and delivery services, or automated vehicles need to go to a power station before their batteries run out.

On the other hand, customers' choice behaviour is dynamic as customers can learn from their past purchase experience. In most of the cases, apart from price, we cannot identify factors, which change customers' choice behaviour and can

be influenced by a RM control policy. For instance, we may not be able to find such a factor that affects customers' preference towards different brands of cookies. Therefore, the customer choice model is periodically reviewed and updated in the RM policy to improve short-term profit. However, when we can explicitly identify those factors, revenue management can be applied to improve business profit in the long run.

Appendix A

Proof of Propositions

A.1 Proposition 2 (Linearization of Slot and Price Assortment Problem)

Consider the nonlinear optimization problem R_{NLP}^a for area a and its optimal solution $(\mathbf{g}^*, \mathbf{z}^*)$. For notational simplicity, we introduce parameter $c_{ns} = r_n - \Delta_s^t$. Let $\mathbf{h} = \{h_{s\kappa} \mid s \in \mathcal{F}(\mathbf{x}), \kappa \in \mathcal{K}\}$ and $\mathbf{f} = \{f_{s\kappa} \mid s \in \mathcal{F}(\mathbf{x}), \kappa \in \mathcal{K} \setminus \{K\}\}$ represent decision variables (corresponding to decisions of the nonlinear optimization model). Given the optimal solution $(\mathbf{g}^*, \mathbf{z}^*)$, we define $V_n^* = \frac{\sum_{s,\kappa} (c_{ns} + d_\kappa) \hat{v}_{ns\kappa} g_{s\kappa}^*}{1 + \sum_s \hat{v}_{ns}^T g_s^*}$. The nonlinear problem (2.7) can then be rewritten as the following linear optimization model:

$$\begin{aligned}
& \max_{\mathbf{h}, \mathbf{f}} \quad \sum_n \eta_{an} \sum_{s,\kappa} (c_{ns} + d_\kappa - V_n^*) \hat{v}_{ns\kappa} h_{s\kappa} \\
& \text{s.t.} \quad h_{m1} + f_{m1} = 1, \quad \forall m \in \mathcal{M}(\mathbf{x}), \\
& \quad h_{m\kappa} + f_{m\kappa} = f_{m,\kappa-1}, \quad \forall m \in \mathcal{M}(\mathbf{x}), \kappa \in \mathcal{K} \setminus \{0, 1, K\}, \\
& \quad h_{mK} = f_{m,K-1}, \quad \forall m \in \mathcal{M}(\mathbf{x}), \\
& \quad h_{m1} = h_{s1} + f_{s1}, \quad \forall m \in \mathcal{M}(\mathbf{x}), s \in \mathcal{S}(\mathbf{x}), \\
& \quad h_{m\kappa} + f_{s,\kappa-1} = h_{s\kappa} + f_{s\kappa}, \quad \forall m \in \mathcal{M}(\mathbf{x}), s \in \mathcal{S}(\mathbf{x}), \kappa \in \mathcal{K} \setminus \{0, 1, K\}, \\
& \quad h_{mK} + f_{s,K-1} = h_{sK}, \quad \forall m \in \mathcal{M}(\mathbf{x}), s \in \mathcal{S}(\mathbf{x}) \\
& \quad \mathbf{h}, \mathbf{f} \in [0, 1].
\end{aligned} \tag{A.1}$$

We first prove that R_{NLP}^a is equivalent to (A.1), and thus $\sum_n \eta_{an} V_n^* = \sum_n \eta_{an} W_n^*$. The optimal solution of (A.1) is $(\mathbf{h}^*, \mathbf{f}^*)$ and the objective value is $\sum_n \eta_{an} W_n^*$ where $W_n^* = \sum_{s,\kappa} (c_{ns} + d_\kappa - V_n^*) \hat{v}_{ns\kappa} h_{s\kappa}^*$. Notice that the feasible sets of

both problems (2.7) and (A.1) consist of the same set of constraints.

- Given the optimal solution $(\mathbf{g}^*, \mathbf{z}^*)$ of R_{NLP}^a , one can easily write the following inequality

$$\sum_n \eta_{an} W_n^* \geq \sum_n \eta_{an} \sum_{s,\kappa} (c_{ns} + d_\kappa - V_n^*) \hat{v}_{ns\kappa} g_{s\kappa}^* \quad (\text{A.2})$$

since the optimal solution is also feasible for (A.1). Furthermore, by substituting

$$V_n^* (1 + \sum_s \hat{v}_{ns}^T g_s^*) = \sum_{s,\kappa} (c_{ns} + d_\kappa) \hat{v}_{ns\kappa} g_{s\kappa}^*$$

in (A.2) we obtain

$$\sum_n \eta_{an} W_n^* \geq \sum_n \eta_{an} V_n^*.$$

This basically implies that the optimal objective value of (A.1) is at least as large as the optimal objective value of R_{NLP}^a .

- Next, let us consider the optimal solution $(\mathbf{h}^*, \mathbf{f}^*)$ obtained from (A.1). This is a feasible solution for the problem R_{NLP}^a because both problems have the same search space. Then, we can write the following valid inequality

$$V_n^* \geq \frac{\sum_{s,\kappa} (c_{ns} + d_\kappa) \hat{v}_{ns\kappa} h_{s\kappa}^*}{\sum_s \hat{v}_{ns}^T h_s^* + 1}$$

that leads to

$$V_n^* (\sum_s \hat{v}_{ns}^T h_s^* + 1) \geq \sum_{s,\kappa} (c_{ns} + d_\kappa) \hat{v}_{ns\kappa} h_{s\kappa}^* \Rightarrow V_n^* \geq W_n^*. \quad (\text{A.3})$$

From this, one can obtain $\sum_n \eta_{an} V_n^* \geq \sum_n \eta_{an} W_n^*$. This shows that the objective value of R_{NLP}^a is at least as large as the optimal objective value of (A.1).

From these two cases, we find $\sum_n \eta_{an} V_n^* = \sum_n \eta_{an} W_n^*$ that basically states that the optimization problems R_{NLP}^a and (A.1) are to be equivalent.

Similarly, we can show that R_{LP}^a and (A.1) are equivalent and they produce the same objective function value. Since V_n^* is a parameter in (A.1), we can remove it from the objective function. This provides R_{LP}^a by exploiting the structural property of MNL choice model. Let $(\hat{\mathbf{g}}^*, \mathbf{z}^*)$ denote the optimal solution of R_{LP}^a . For

$K_n^* = \sum_{s,\kappa} (c_{ns} + d_\kappa) \hat{g}_{ns\kappa}^*$, let us consider the following two cases.

- The optimal solution $(\mathbf{h}^*, \mathbf{f}^*)$ of (A.1) constructs a feasible solution for R_{LP}^a as

$$\hat{g}_{ns\kappa} = \frac{\hat{v}_{ns\kappa} h_{s\kappa}^*}{(1 + \sum_j \hat{v}_{nj}^T h_j^*)}, \quad \hat{g}_{n0} = \frac{1}{(1 + \sum_j \hat{v}_{nj}^T h_j^*)}, \quad \text{and} \quad \hat{z}_\kappa = \frac{f_{s\kappa}^*}{(1 + \sum_j \hat{v}_{nj}^T h_j^*)}.$$

Thus, we can then state the following relationship

$$\begin{aligned} \sum_n \eta_{an} K_n^* &\geq \sum_n \eta_{an} \sum_{s,\kappa} (c_{ns} + d_\kappa) \hat{g}_{ns\kappa} \\ &= \sum_n \eta_{an} \frac{\sum_{s,\kappa} (c_{ns} + d_\kappa) \hat{v}_{ns\kappa} h_{s\kappa}^*}{1 + \sum_j \hat{v}_{nj}^T h_j^*} \\ &= \sum_n \eta_{an} W_n^*. \end{aligned} \tag{A.4}$$

This indicates that the optimal value of R_{LP}^a is greater or equal to the optimal value of (A.1).

- In the same way, one can show that $(\hat{\mathbf{g}}^*, \hat{\mathbf{z}}^*)$ constructs a feasible solution of the problem (A.1). In other words, $h_{ns\kappa} = \frac{\hat{g}_{ns\kappa}^*}{(\hat{g}_{n0}^* \hat{v}_{ns\kappa})}$ and $f_{ns\kappa} = \frac{\hat{z}_{ns\kappa}^*}{\hat{g}_{n0}^*}$ is a feasible solution and satisfy the following inequality;

$$\begin{aligned} \sum_n \eta_{an} W_n^* &\geq \sum_n \eta_{an} \sum_{s,\kappa} (c_{ns} + d_\kappa - V_n^*) \hat{v}_{ns\kappa} h_{ns\kappa} \\ &= \sum_n \eta_{an} \sum_{s,\kappa} (c_{ns} + d_\kappa - V_n^*) \frac{\hat{g}_{ns\kappa}^*}{\hat{g}_{n0}^*}. \end{aligned} \tag{A.5}$$

Using the relations (A.4) and $\sum_n \eta_{an} W_n^* = \sum_n \eta_{an} V_n^*$ (as already proven above) in (A.5), we obtain

$$\begin{aligned} \sum_n \eta_{an} W_n^* &\geq \sum_n \eta_{an} \sum_{s,\kappa} (c_{ns} + d_\kappa) \frac{\hat{g}_{ns\kappa}^*}{\hat{g}_{n0}^*} - \sum_n \eta_{an} \sum_{s,\kappa} K_n^* \frac{\hat{g}_{ns\kappa}^*}{\hat{g}_{n0}^*} \\ &\geq \sum_n \eta_{an} \frac{K_n^*}{\hat{g}_{n0}^*} - \sum_n \eta_{an} \sum_{s,\kappa} K_n^* \frac{\hat{g}_{ns\kappa}^*}{\hat{g}_{n0}^*} = \sum_n \eta_{an} K_n^* \left(\frac{1}{\hat{g}_{n0}^*} - \sum_{s,\kappa} \frac{\hat{g}_{ns\kappa}^*}{\hat{g}_{n0}^*} \right). \end{aligned}$$

Using $\hat{g}_{n0}^* = 1 - \sum_{s,\kappa} \hat{g}_{ns\kappa}^*$, we obtain the following inequality

$$\sum_n \eta_{an} W_n^* \geq \sum_n \eta_{an} K_n^*. \tag{A.6}$$

This indicates that the optimal value of (A.1) is not less than the optimal value of R_{LP}^a .

From (A.4) and (A.6), we achieve $\sum_n \eta_{an} K_n^* = \sum_n \eta_{an} W_n^*$, and thus R_{LP}^a and (A.1) are equivalent. In a summary, we can conclude that R_{NLP}^a and R_{LP}^a are equivalent and they possess the same objective value.

A.2 Proposition 3 (Linearization of Value Function Estimation)

Suppose that the MNL model is used to describe the customer choice behaviour. We consider the NLP and LP problems given state \mathbf{x} at time t . Let $(\mathbf{g}^*, \mathbf{w}_1^*)$ denote the optimal solution of NLP with the optimal value \hat{V}^* . Meanwhile, the optimal solution of LP is denoted by $(\mathbf{y}^*, \mathbf{w}_2^*)$ and the optimal value is R^* . In order to show that these models are equivalent and produce the same optimal value (i.e., $\hat{V}^* = R^*$) under the MNL choice model, we follow steps in two cases:

Case 1: We first prove that $(\mathbf{g}^*, \mathbf{w}_1^*)$ constructs a feasible solution for LP so that $R(\mathbf{g}^*, \mathbf{w}_1^*) \leq R^*$. Given $(\mathbf{g}^*, \mathbf{w}_1^*)$, let's define the following decision variables

$$\mathbf{w}' = \mathbf{w}_1^*, y'_{asn\kappa} = \sum_{i=t}^T P_{ians\kappa}(\mathbf{g}^*) \text{ and } y'_{an0\kappa} = \sum_{i=t}^T P_{ian0}(\mathbf{g}^*), \forall \kappa \in \mathcal{K}, a \in \mathcal{A}, n \in \mathbf{N}, s \in \mathcal{F}.$$

It can be easily shown that $(\mathbf{y}', \mathbf{w}')$ satisfies the first four sets of constraints in LP . For the last set of constraints, using $\sum_{a,n} \mu_a \eta_{an} = 1$ and $\sum_{s,\kappa} p'_{ians\kappa}(\mathbf{g}^*) + p'_{ian0}(\mathbf{g}^*) = 1$ in the left-side of the equality, we find

$$\begin{aligned} \sum_{i=t}^T \sum_{a,n} \left(\sum_{s,\kappa} P_{ians\kappa}(\mathbf{g}^*) + P_{ian0}(\mathbf{g}^*) \right) &= \lambda \sum_{i=t}^T \sum_{a,n} \mu_a \eta_{an} \left(\sum_{s,\kappa} p'_{ians\kappa}(\mathbf{g}^*) + p'_{ian0}(\mathbf{g}^*) \right) \\ &= \lambda(T - t + 1). \end{aligned}$$

This implies that $(\mathbf{g}^*, \mathbf{w}_1^*)$ satisfies all constraints of the LP model. Therefore, it is a feasible solution for LP and we can state that $\hat{V}^* \leq R^*$.

Case 2: Let's first reformulate NLP as a choice-based deterministic linear model (CDLP). Then, using the duality theory, we show that the optimal solution for the dual of $CDLP$ is feasible for the dual problem of LP . Thus, $R^* \leq \hat{V}^*$ holds.

Let set \mathcal{G} consist of all possible pricing decisions \mathbf{g} for all slots. We define new decision variable $v_{ia}(\mathbf{g})$ to represent the probability of offering price vector \mathbf{g} at time i in area a . For notational convenience, we introduce $P'_{ians\kappa} = \sum_{\mathbf{g} \in \mathcal{G}} P_{ians\kappa}(\mathbf{g}) v_{ia}(\mathbf{g})$ as the decision variable in CDLP which indicates the probability of slot s having price d_κ at time i for a segment- n customer from area a . Accordingly, CDLP can be

formulated as follows:

$$\begin{aligned}
CDLP : \quad \hat{V}_t(\mathbf{x}) = & \max_{\mathbf{P}', \mathbf{w}} \sum_{i=t}^T \sum_{a \in \mathcal{A}, s \in \mathcal{S}} \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} P'_{ians\kappa} (\bar{r}_{an} + d_\kappa) - \sum_{a \in \mathcal{A}, s \in \mathcal{S}} C_{as}(x'_{as}) \\
\text{s.t.} \quad & \sum_{s \in \mathcal{F}} \left[x_{as} + \sum_{i=t}^T \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} P'_{ians\kappa} \right] \leq c, \quad \forall a \in \mathcal{A}, \\
& \sum_{s \in \mathcal{S}_m} w_{ams} = x_{am} + \sum_{i=t}^T \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} P'_{ianm\kappa}, \quad \forall a \in \mathcal{A}, m \in \mathcal{M}, \\
& x_{as} + \sum_{i=t}^T \sum_{n \in \mathbf{N}, \kappa \in \mathcal{K}} P'_{ians\kappa} + \sum_{m \in \mathcal{M}_s} w_{ams} \leq B_a, \quad \forall s \in \mathcal{S}, a \in \mathcal{A}, \\
& \sum_{i=t}^T \sum_{\kappa \in \mathcal{K}} \frac{P'_{ians\kappa}}{v_{ns\kappa}} = \sum_{i=t}^T P'_{ian0}, \quad \forall a \in \mathcal{A}, s \in \mathcal{F}, n \in \mathbf{N}. \\
& \mathbf{P}' \in [0, 1], \quad \mathbf{w} \geq 0.
\end{aligned} \tag{A.7}$$

Let $\sigma_1 = \{\sigma_{1a} \mid \forall a \in \mathcal{A}\}$, $\sigma_2 = \{\sigma_{2am} \mid \forall a \in \mathcal{A}, m \in \mathcal{M}\}$, $\sigma_3 = \{\sigma_{3as} \mid \forall a \in \mathcal{A}, s \in \mathcal{S}\}$ and $\sigma_4 = \{\sigma_{4ans} \mid \forall a \in \mathcal{A}, n \in \mathbf{N}, s \in \mathcal{F}\}$ denote dual decision variables corresponding to constraints of the (primal) *CDLP*. We also introduce $\boldsymbol{\sigma} = \{\sigma_1, \sigma_2, \sigma_3, \sigma_4\}$ for notational simplicity. The constraints of the dual problem of *CDLP* are

$$\begin{aligned}
\sigma_{1a} + \sigma_{3as} + \frac{\sigma_{4ans}}{\hat{v}_{ans\kappa}} & \geq \bar{r}'_{an\kappa}, \quad \forall a \in \mathcal{A}, s \in \mathcal{S}, n \in \mathbf{N}, \kappa \in \mathcal{K}, \\
\sigma_{1a} + \sigma_{2am} + \frac{\sigma_{4ams}}{\hat{v}_{anm\kappa}} & \geq \bar{r}'_{an\kappa}, \quad \forall a \in \mathcal{A}, m \in \mathcal{M}, n \in \mathbf{N}, \kappa \in \mathcal{K}, \\
\sigma_{2am} + \sigma_{3as} & \geq 0, \quad \forall a \in \mathcal{A}, m \in \mathcal{M}, s \in \mathcal{S}_m, \\
-\sum_{s \in \mathcal{F}} \sigma_{4ans} & \geq 0, \quad \forall a \in \mathcal{A}, n \in \mathbf{N},
\end{aligned} \tag{A.8}$$

where $\bar{r}'_{an\kappa} = \bar{r}_{an} + d_\kappa - \frac{\delta\beta_a}{6}$ represents the marginal profit-after-delivery in area a with slot price d_κ . The optimal value V_D^* of the dual *CDLP* problem is achieved at $\boldsymbol{\sigma}^*$.

Similarly, we define dual decision variables $\phi_1 = \{\phi_{1a} \mid \forall a \in \mathcal{A}\}$, $\phi_2 = \{\phi_{2am} \mid \forall a \in \mathcal{A}, m \in \mathcal{M}\}$, $\phi_3 = \{\phi_{3as} \mid \forall a \in \mathcal{A}, s \in \mathcal{S}\}$, $\phi_4 = \{\phi_{4ans} \mid \forall a \in \mathcal{A}, n \in \mathbf{N}, s \in \mathcal{F}\}$ and ϕ_5 corresponding to constraints of the primal *LP* problem and denote $\boldsymbol{\phi} = \{\phi_1, \phi_2, \phi_3, \phi_4, \phi_5\}$. The dual problem of *LP* involves the following

constraints

$$\begin{aligned}
\phi_{1a} + \phi_{3as} + \frac{\phi_{4ans}}{\hat{v}_{ans\kappa}} + \phi_5 &\geq \bar{r}'_{an\kappa}, & \forall a \in \mathcal{A}, s \in \mathcal{S}, n \in \mathbf{N}, \kappa \in \mathcal{K}, \\
\phi_{1a} + \phi_{2am} + \frac{\phi_{4ams}}{\hat{v}_{anm\kappa}} + \phi_5 &\geq \bar{r}'_{an\kappa}, & \forall a \in \mathcal{A}, m \in \mathcal{M}, n \in \mathbf{N}, \kappa \in \mathcal{K}, \\
\phi_{2am} + \phi_{3as} &\geq 0, & \forall a \in \mathcal{A}, m \in \mathcal{M}, s \in \mathcal{S}_m, \\
\phi_5 - \sum_{s \in \mathcal{F}} \phi_{4ans} &\geq 0, & \forall a \in \mathcal{A}, n \in \mathbf{N}.
\end{aligned} \tag{A.9}$$

R_D^* is obtained by the optimal solution ϕ^* of the dual problem of LP . Next, from the first two sets of constraints in (A.8), we define $A_{ans\kappa}^* = \sigma_{1a}^* + \sigma_{3as}^* + \frac{\phi_{4ans}^*}{\hat{v}_{ans\kappa}} - \bar{r}'_{an\kappa}$, and $B_{anm\kappa}^* = \sigma_{1a}^* + \sigma_{2am}^* + \frac{\phi_{4ams}^*}{\hat{v}_{anm\kappa}} - \bar{r}'_{an\kappa}$. Then, the following relationship holds

$$\sum_{s, \kappa} A_{ans\kappa}^* + \sum_{m, \kappa} B_{anm\kappa}^* \geq 0, \quad \forall n \in \mathbf{N}, a \in \mathcal{A}. \tag{A.10}$$

Notice that ϕ_5 can take any value satisfying (A.10). Since $\phi_5 \geq 0$ and $\sum_{s \in \mathcal{F}} \sigma_{4ans}^* \leq 0$ (that is obtained from (A.8)), one can easily observe that σ^* and ϕ_5 satisfy constraints in (A.9). Therefore, $R_D^* \leq V_D^*$ holds such that $R^* \leq V^*$.

Appendix B

Estimated Choice Models

B.1 Estimated MNL Model with Substitution Effect

Day 1				Day 2			
Slot (s)	\hat{u}_s	RSE	p-value	Slot (s)	\hat{u}_s	RSE	p-value
1	-1.5052	0.0258	0.00	1	-2.4908	0.0403	0
2	-1.2374	0.0230	0.00	2	-2.1623	0.0346	0
3	-1.0109	0.0210	0.00	3	-1.7942	0.0293	0
4	-0.5032	0.0174	0.00	4	-1.1224	0.0220	0
5	-0.0131	0.0149	0.38	5	-0.4647	0.0172	0
6	-1.3774	0.0244	0.00	6	-2.2600	0.0361	0
7	-1.8610	0.0302	0.00	7	-2.9664	0.0505	0
8	-0.8739	0.0199	0.00	8	-1.5948	0.0268	0
9	-0.6193	0.0182	0.00	9	-1.2752	0.0234	0
10	0.6092	0.0128	0.00	10	0.3532	0.0136	0
11	1.1044	0.0117	0.00	11	0.9841	0.0120	0
12	0.9722	0.0120	0.00	12	0.8314	0.0123	0

Note: utility of no-purchase is 0.

Table B.1: Estimated parameters of the MNL model involving two delivery days

B.2 Estimated Nested MNL Model

Day 1				Day 2			
Slot (s)	\hat{u}_s	RSE	p-value	Slot (s)	\hat{u}_s	RSE	p-value
1	-0.7370	0.0606	0.00	1	-0.7586	0.0643	0.00
2	-0.5277	0.0558	0.00	2	-0.5663	0.0589	0.00
3	-0.3504	0.0519	0.00	3	-0.3514	0.0530	0.00
4	0.0460	0.0433	0.29	4	0.0439	0.0429	0.31
5	0.4302	0.0354	0.00	5	0.4319	0.0335	0.00
6	-0.6372	0.0583	0.00	6	-0.6231	0.0605	0.00
7	-1.0143	0.0671	0.00	7	-1.0365	0.0728	0.00
8	-0.2436	0.0495	0.00	8	-0.2343	0.0500	0.00
9	-0.0445	0.0453	0.33	9	-0.0452	0.0452	0.32
10	0.9209	0.0258	0.00	10	0.9214	0.0225	0.00
11	1.3133	0.0188	0.00	11	1.3113	0.0151	0.00
12	1.2081	0.0206	0.00	12	1.2153	0.0167	0.00

No-purchase utility is 0 with dissimilarity parameter 1.

Dissimilarity parameters for Days 1 and 2 are 0.7799 and 0.5840.

Table B.2: Estimated parameters of the nested MNL model for two delivery days

B.3 Estimated MNL Model with No Substitution Effect

Day 1				Day 2			
Slot (s)	\hat{u}_s	RSE	p-value	Slot (s)	\hat{u}_s	RSE	p-value
1	-3.3918	0.0243	0.00	1	-4.6398	0.0393	0.00
2	-3.1243	0.0214	0.00	2	-4.3109	0.0334	0.00
3	-2.8978	0.0192	0.00	3	-3.9432	0.0279	0.00
4	-2.3904	0.0151	0.00	4	-3.2714	0.0201	0.00
5	-1.9003	0.0122	0.00	5	-2.6140	0.0147	0.00
6	-3.2644	0.0228	0.00	6	-4.4089	0.0350	0.00
7	-3.7484	0.0289	0.00	7	-5.1155	0.0498	0.00
8	-2.7611	0.0180	0.00	8	-3.7441	0.0253	0.00
9	-2.5064	0.0160	0.00	9	-3.4243	0.0216	0.00
10	-1.2790	0.0094	0.00	10	-1.7968	0.0103	0.00
11	-0.7846	0.0079	0.00	11	-1.1666	0.0080	0.00
12	-0.9166	0.0083	0.00	12	-1.3197	0.0085	0.00

Note: no-purchase has utility of 0 for both days.

Table B.3: Estimated parameters of two independent MNL models

Bibliography

- N. Agatz. Challenges and Opportunities in Attended Home Delivery. In The vehicle routing problem: Latest advances and new challenges, pages 1–20. Springer, 2007.
- N. Agatz, M. Fleischmann, and J. A. E. E. Van Nunen. E-fulfillment and multi-channel distribution—a review. European Journal of Operational Research, 187(2): 339–356, 2008.
- N. Agatz, A. Campbell, M. Fleischmann, and M. Savelsbergh. Time slot management in attended home delivery. Transportation Science, 45(3):435–449, 2010.
- N. Agatz, A. M. Campbell, M. Fleischmann, J. Van Nunen, and M. Savelsbergh. Revenue management opportunities for internet retailers. Journal of Revenue and Pricing Management, 12(2):128–138, 2013.
- C. K. Anderson and X. Xie. A choice-based dynamic programming approach for setting opaque prices. Production and Operations Management, 21(3):590–605, 2012.
- K. Asdemir, V. S. Jacob, and R. Krishnan. Dynamic pricing of multiple home delivery options. European Journal of Operational Research, 196(1):246–257, 2009.
- N. Aydın, Ş. İ. Birbil, and H. Topaloglu. Delayed purchase options in single-leg revenue management. Transportation Science, 51(4):1031–1045, 2016.
- N. Azi, M. Gendreau, and J. Y. Potvin. A dynamic vehicle routing problem with multiple delivery routes. Annals of Operations Research, 199(1):103–112, 2012.
- A. E. M. Bayoumi, M. Saleh, A. F. Atiya, and H. A. Aziz. Dynamic pricing for hotel revenue management using price multipliers. Journal of Revenue and Pricing Management, 12(3):271–285, 2013.
- S. Belhaiza, P. Hansen, and G. Laporte. A hybrid variable neighborhood tabu search heuristic for the vehicle routing problem with multiple time windows. Computers & Operations Research, 52:269–281, 2014.

- R. Bellman. The theory of dynamic programming. Bulletin of the American Mathematical Society, 60(6):503–515, 1954.
- P. P. Belobaba. Survey paper - airline yield management an overview of seat inventory control. Transportation Science, 21(2):63–73, 1987.
- P. P. Belobaba and L. R. Weatherford. Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. Decision Sciences, 27(2):343–363, 1996.
- D. P. Bertsekas. Dynamic Programming and Optimal Control, volume 1. Athena scientific Belmont, MA, 1995.
- A. V. D. Boer. Dynamic pricing and learning: Historical origins, current research, and new directions. Surveys in Operations Research & Management Science, 20(1):1–18, 2015.
- J. J. M. Bront, I. Méndez-Díaz, and G. Vulcano. A column generation algorithm for choice-based network revenue management. Operations Research, 57(3):769–784, 2009.
- D. Bühler, R. Klein, and M. Neugebauer. Model-based delivery cost approximation in attended home services. Computers & Industrial Engineering, 98:78–90, 2016.
- B. Burger and M. Fuchs. Dynamic pricing—a future airline business model. Journal of Revenue and Pricing Management, 4(1):39–53, 2005.
- C. Büsing, D. Kadatz, and C. Cleophas. Capacity uncertainty in airline revenue management: Models, algorithms, and computations. Transportation Science, 53(2):383–400, 2019.
- M. Çakanyıldırım, Ö. Özer, and X. Zhang. Dynamic pricing and timing of upgrades. Working Paper, 2017.
- A. M. Campbell and M. Savelsbergh. Incentive schemes for attended home delivery services. Transportation Science, 40(3):327–341, 2006.
- A. M. Campbell and M. W. P. Savelsbergh. Decision support for consumer direct grocery initiatives. Transportation Science, 39(3):313–327, 2005.
- O. Ceryan, I. Duenyas, and O. Sahin. Dynamic pricing and replenishment with customer upgrades. Production and Operations Management, 27(4):663–679, 2018.
- C. Cleophas and J. F. Ehmke. When are deliveries profitable? Business & Information Systems Engineering, 6(3):153–163, 2014.

- W. L. Cooper, T. Homem-de Mello, and A. J. Kleywegt. Models of the spiral-down effect in revenue management. Operations Research, 54(5):968–987, 2006.
- C. F. Daganzo. Modeling distribution problems with time windows. Part I. Transportation Science, 21(3):171–179, 1987.
- J. Davis, G. Gallego, and H. Topaloglu. Assortment planning under the multinomial logit model with totally unimodular constraint structures. Working Paper, 2013.
- J. Dean. Amazon fails to deliver as costs rise. Available at <https://www.thetimes.co.uk/article/amazon-fails-to-deliver-as-costs-rise-q6mxqmjmb> (2019/09/15), 2016.
- Ecommerce-Foundation. Global B2C e-commerce report 2016. Available at <https://www.ecommercewiki.org/reports/300/global-b2c-e-commerce-report-2016> (2019/09/15), 2016.
- J. F. Ehmke. Attended Home Delivery, volume 177, chapter 3, pages 23–33. Springer Science & Business Media, 2012.
- J. F. Ehmke and A. M. Campbell. Customer acceptance mechanisms for home deliveries in metropolitan areas. European Journal of Operational Research, 233(1):193–207, 2014.
- P. S. Fader, B. G. S. Hardie, and C.-Y. Huang. A dynamic changepoint model for new product sales forecasting. Marketing Science, 23(1):50–65, 2004.
- Y. Feng and B. Xiao. A dynamic airline seat inventory control model and its optimal policy. Operations Research, 49(6):938–949, 2001.
- G. Gallego and R. Phillips. Revenue management of flexible products. Manufacturing & Service Operations Management, 6(4):321–337, 2004.
- G. Gallego and C. Stefanescu. Upgrades, upsells and pricing in revenue management. Working Paper, 2011.
- G. Gallego and H. Topaloglu. Constrained assortment optimization for the nested logit model. Management Science, 60(10):2583–2601, 2014.
- G. Gallego and G. Van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. Management Science, 40(8):999–1020, 1994.
- G. Gallego, G. Iyengar, R. Phillips, and A. Dubey. Managing flexible products on a network. 2004.

- G. Gallego, R. Ratliff, and S. Shebalov. A general attraction model and sales-based linear program for network revenue management under customer choice. Operations Research, 63(1):212–232, 2014.
- H. N. Geismar, M. Dawande, B. P. Murthi, and C. Sriskandarajah. Maximizing revenue through two-dimensional shelf-space allocation. Production and Operations Management, 24(7):1148–1163, 2015.
- B. L. Golden, S. Raghavan, and E. A. Wasil. The Vehicle Routing Problem: Latest Advances and New Challenges. Springer Science & Business Media, 2008.
- J. Gönsch, S. Koch, and C. Steinhardt. An emsr-based approach for revenue management with integrated upgrade decisions. Computers and Operations Research, 40(10):2532–2542, 2013.
- J. Gönsch, S. Koch, and C. Steinhardt. Revenue management with flexible products: The value of flexibility and its incorporation into DLP-based approaches. International Journal of Production Economics, 153:280–294, 2014.
- F. Guerriero, G. Miglionico, and F. Olivito. Strategic and operational decisions in restaurant revenue management. European Journal of Operational Research, 237(3):1119–1132, 2014.
- T. Hays, P. Keskinocak, and V. M. de López. Strategies and challenges of internet grocery retailing logistics. In Applications of Supply Chain Management and E-Commerce Research, pages 217–252. Springer US, Boston, MA, 2005.
- A. Heching, G. Gallego, and G. van Ryzin. Mark-down pricing: An empirical analysis of policies and revenue potential at one apparel retailer. Journal of Revenue and Pricing Management, 1(2):139–160, 2002.
- R. A. Howard. Dynamic programming and markov processes. John Wiley, 1960.
- A. Hübner, H. Kuhn, and J. Wollenburg. Last mile fulfilment and distribution in omni-channel grocery retailing. International Journal of Retail & Distribution Management, 44(3):228–247, 2016.
- R. Huefner. Revenue Management: A Path to Increased Profits. Business Expert Press, 2011.
- E. C. Joel and C. B. Carol. Measuring service quality in e-retailing. Journal of Service Research, 8(3):260–275, 2006.
- S. E. Kimes. The future of hotel revenue management. Journal of Revenue and Pricing Management, 10(1):62–72, 2011.

- S. E. Kimes, D. I. Barrash, and J. E. Alexander. Developing a restaurant revenue-management strategy. Cornell Hotel and Restaurant Administration Quarterly, 40(5):18–29, 2007.
- M. A. Klapp, A. L. Erera, and A. Toriello. The one-dimensional dynamic dispatch waves problem. Transportation Science, 52(2):402–415, 2016.
- M. A. Klapp, A. L. Erera, and A. Toriello. The dynamic dispatch waves problem for same-day delivery. European Journal of Operational Research, 271(2):519–534, 2018.
- R. Klein, M. Neugebauer, D. Ratkovitch, and C. Steinhardt. Differentiated time slot pricing under routing considerations in attended home delivery. Transportation Science, 53(1):236–255, 2017.
- R. Klein, J. Mackert, M. Neugebauer, and C. Steinhardt. A model-based approximation of opportunity cost for dynamic pricing in attended home delivery. OR Spectrum, 40(4):969–996, 2018.
- A. J. Kleywegt, V. S. Nori, and M. W. P. Savelsbergh. Dynamic programming approximations for a stochastic inventory routing problem. Transportation Science, 38(1):42–70, 2004.
- S. Koch, J. Gönsch, and C. Steinhardt. Dynamic programming decomposition for choice-based revenue management with flexible products. Transportation Science, 51(4):1046–1062, 2017.
- C. Köhler, J. F. Ehmke, and A. M. Campbell. Flexible time window management for attended home deliveries. Omega, 2019. In Press.
- D. Koushik, J. A. Higbie, and C. Eister. Retail price optimization at intercontinental hotels group. Interfaces, 42(1):45–57, 2012.
- M. A. Lang, C. Cleophas, and J. F. Ehmke. Anticipative dynamic slotting for attended home deliveries. Working Paper, 2019.
- B. W. Lippman. Retail revenue management - competitive strategy for grocery retailers. Journal of Revenue and Pricing Management, 2(3):229–233, 2003.
- Q. Liu and G. van Ryzin. On the choice-based linear programming model for network revenue management. Manufacturing & Service Operations Management, 10(2):288–310, 2007.
- J. Mackert. Choice-based dynamic time slot management in attended home delivery. Computers and Industrial Engineering, 129:333–345, 2019.

- C. Maglaras and J. Meissner. Dynamic pricing strategies for multiproduct revenue management problems. Manufacturing & Service Operations Management, 8(2):136–148, 2006.
- S. Mang, D. Post, and M. Spann. Pricing of flexible products. Review of Managerial Science, 6(4):361–374, 2012.
- K. J. Mason and R. Gray. Short haul business travel in the European Union: a segmentation profile. Journal of Air Transport Management, 2(3-4):197–205, 1995.
- A. S. Mattila, L. Hanks, and L. Zhang. Existential guilt and preferential treatment: The case of an airline upgrade. Journal of Travel Research, 52(5):591–599, 2013.
- D. McCaffrey and D. Walczak. Optimal dynamic upgrading in revenue management. Production and Operations Management, 25(11):1855–1865, 2016.
- J. McGill and G. Van Ryzin. Revenue management: Research overview and prospects. Transportation Science, 33(2):233–256, 1999.
- N. Meade and T. Islam. Using copulas to model repeat purchase behaviour - an exploratory analysis via a case study. European Journal of Operational Research, 200(3):908–917, 2010.
- Mintel. Attitudes towards home-delivery and takeaway food. Available at <https://academic.mintel.com/display/747922/> (2019/09/15), 2016.
- Mintel. Online grocery retailing - uk. Available at <https://academic.mintel.com/display/792541/> (2019/09/15), 2017.
- Mintel. Business traveller - uk. Available at <https://academic.mintel.com/display/898836/> (2019/09/15), 2018a.
- Mintel. Airlines - uk. Available at <https://academic.mintel.com/display/859381/> (2019/09/15), 2018b.
- Mintel. Online grocery retailing - uk. Technical report, mar 2018c.
- Mintel. Online retailing - uk. Available at <https://academic.mintel.com/display/859337/> (2019/09/15), 2018d.
- Mintel. Airlines - uk. Available at <https://reports.mintel.com/display/972338/> (2020/02/10), 2019.
- M. G. Morgan, M. Henrion, and M. Small. Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis. Cambridge university press, 1992.

- K. G. Murty. Linear Programming. Springer, 1983.
- G. L. Nemhauser and L. A. Wolsey. Integer and Combinatorial Optimization. Wiley-Interscience, 1988.
- I. C. L. Ng. The pricing and revenue management of services : a strategic approach. Routledge, 2008.
- D. H. Nguyen, S. de Leeuw, W. Dullaert, and B. P. Foubert. What is the right delivery option for you? consumer preferences for delivery attributes in online retailing. Journal of Business Logistics, 2019.
- B. B. Oliveira, M. A. Carravilla, and J. F. Oliveira. Fleet and revenue management in car rental companies: A literature review and an integrated conceptual framework. Omega, 71:11–26, 2017.
- Optiontown. Learn about upgrade travel option. Available at https://www.optiontown.com/home_page.do?processAction=SelectProduct (2019/09/15), 2018.
- J. Y. Park and S. C. S. Jang. You got a free upgrade? what about me? the consequences of unearned preferential treatment. Tourism Management, 50:59–68, 2015.
- A. Petrick, J. Gönsch, C. Steinhardt, and R. Klein. Dynamic control mechanisms for revenue management with flexible products. Computers & Operations Research, 37(11):2027–2039, 2010a.
- A. Petrick, C. Steinhardt, J. Gönsch, and R. Klein. Using flexible products to cope with demand uncertainty in revenue management. OR Spectrum, 34(1):215–242, 2010b.
- I. Popescu and Y. Wu. Dynamic pricing strategies with reference effects. Operations research, 55(3):413–429, 2007.
- W. B. Powell, A. George, B. Bouzaiene-Ayari, and H. P. Simao. Approximate dynamic programming for high-dimensional resource allocation problems. In Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005., volume 5, pages 2989–2994. IEEE, 2005.
- M. L. Puterman. Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley & Sons, 2014.

- A. Sauré, J. Patrick, S. Tyldesley, and M. L. Puterman. Dynamic multi-appointment patient scheduling for radiation therapy. European Journal of Operational Research, 223(2):573–584, 2012.
- M. W. P. Savelsbergh. Local search in routing problems with time windows. Annals of Operations Research, 4(1):285–305, 1985.
- D. Schuurmans and R. Patrascu. Direct value-approximation for factored MDPs. In Advances in Neural Information Processing Systems, pages 1579–1586, 2002.
- R. A. Shumsky and F. Zhang. Dynamic capacity management with substitution. Operations Research, 57(3):671–684, 2009.
- D. D. Sierag, G. M. Koole, R. D. Van Der Mei, J. I. Van Der Rest, and B. Zwart. Revenue management under customer choice behaviour with cancellations and overbooking. European Journal of Operational Research, 246(1):170–185, 2015.
- M. Slabinac et al. Innovative solutions for a last-mile delivery—a european experience. In Proceedings of the 15th International Scientific Conference Business Logistics in Modern Management Osijek, Croatia, pages 111–129, 2015.
- Y. Song, M. Ulmer, B. Thomas, and S. Wallace. Building trust in home services - stochastic team-orienteeing with consistency constraints. Working Paper, 2018.
- C. Steinhardt and J. Gönsch. Integrated revenue management approaches for capacity control with planned upgrades. European Journal of Operational Research, 223(2):380–391, 2012.
- A. K. Strauss, R. Klein, and C. Steinhardt. A review of choice-based revenue management: Theory and methods. European Journal of Operational Research, 271(2):375–387, 2018.
- J. Subramanian, S. Stidham, and C. J. Lautenbacher. Airline yield management with overbooking, cancellations, and no-shows. Transportation Science, 33(2):147–167, 1999.
- K. Talluri and G. van Ryzin. Revenue management under a general discrete choice model of consumer behavior. Management Science, 50(1):15–33, 2004.
- K. Talluri and G. van Ryzin. An analysis of bid-price controls for network revenue management. Management Science, 44(11-part-1):1577–1593, 2008.
- K. T. Talluri. The theory and practice of revenue management. Springer US, 2005.

- H. Tempelmeier. On the stochastic uncapacitated dynamic single-item lotsizing problem with service level constraints. European Journal of Operational Research, 181(1):184–194, 2007.
- K. E. Train. Discrete choice methods with simulation. Cambridge university press, 2003.
- M. A. Trick and S. E. Zin. A linear programming approach to solving stochastic dynamic programs. Working Paper, 1993.
- M. F. Tudor Bodea. Segmentation, Revenue Management and Pricing Analytics. Routledge, 2014.
- M. Ulmer. Delivery deadlines in same-day delivery. Logistics Research, 10(3):1–15, 2017.
- S. A. Voccia, A. M. Campbell, and B. W. Thomas. The same-day delivery problem for online purchases. Transportation Science, 53(1):167–184, 2017.
- C. P. Wright, H. Groenevelt, and R. A. Shumsky. Dynamic revenue management in airline alliances. Transportation Science, 44(1):15–37, 2010.
- C. Wu and H. L. Chen. A consumer purchasing model with learning and departure behaviour. Journal of the Operational Research Society, 51(5):583–591, 2000.
- B. Xiao and W. Yang. A revenue management model for products with two capacity dimensions. European Journal of Operational Research, 205(2):412–421, 2010.
- X. Yang and A. K. Strauss. An approximate dynamic programming approach to attended home delivery management. European Journal of Operational Research, 263(3):935–945, 2017.
- X. Yang, A. K. Strauss, C. S. M. Currie, and R. Eglese. Choice-based demand management and vehicle routing in e-fulfillment. Transportation Science, 50(2):473–488, 2014.
- B. Yildiz and M. Savelsbergh. Service and capacity planning in crowd-sourced delivery. Transportation Research Part C: Emerging Technologies, 100:177–199, 2019.
- Ö. Yılmaz, P. Pekgün, and M. Ferguson. Would you like to upgrade to a premium room? evaluating the benefit of offering standby upgrades. Manufacturing & Service Operations Management, 19(1):1–18, 2017.

- Ö. Yılmaz, M. Ferguson, P. Pekgün, and G. Shang. Empirically estimating strategic behavior for hotel standby upgrade programs. Working Paper, 2019.
- Y. Yu, X. Chen, and F. Zhang. Dynamic capacity management with general upgrading. Operations Research, 63(6):1372–1389, 2015.
- J. Z. Zhang, O. Netzer, and A. Ansari. Dynamic targeted pricing in B2B relationships. Marketing Science, 33(3):317–337, 2014.