



HAL
open science

The Monte Carlo first-come-first-served heuristic for network revenue management

Nicolas Houy, François Le Grand

► **To cite this version:**

Nicolas Houy, François Le Grand. The Monte Carlo first-come-first-served heuristic for network revenue management. 2015. halshs-01155698

HAL Id: halshs-01155698

<https://shs.hal.science/halshs-01155698v1>

Preprint submitted on 27 May 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

WP 1514 – May 2015

The Monte Carlo first-come-first-served heuristic for network revenue management

Nicolas Houy, François Le Grand

Abstract:

We introduce the Monte-Carlo based heuristic with first-come-first-served approximation for future optimal strategy (MC-FCFS) in order to maximize profit in a network revenue management problem. Like the randomized linear programming (RLP) model, one purpose of the MC-FCFS heuristic is to have information about displacement costs, considering the full probability distribution of future demands instead of a simplified degenerate distribution as in the deterministic linear programming (DLP) model. However, this information is conveyed by applying the FCFS heuristic as a future strategy rather than using the optimal ex-post profits as in the RLP heuristic. We show that MC-FCFS performs approximately as well as the RLP heuristic at a much lower computational cost and much better than the DLP heuristic at maximizing profit in a multi-night hotel booking setting with or without planned upgrades.

Keywords:

Network revenue management, Monte-Carlo simulations, randomized linear programming

JEL codes:

C44, C63

The Monte Carlo first-come-first-served heuristic for network revenue management.

Nicolas Houy* François Le Grand†

May 22, 2015

Abstract

We introduce the Monte-Carlo based heuristic with first-come-first-served approximation for future optimal strategy (MC-FCFS) in order to maximize profit in a network revenue management problem. Like the randomized linear programming (RLP) model, one purpose of the MC-FCFS heuristic is to have information about displacement costs, considering the full probability distribution of future demands instead of a simplified degenerate distribution as in the deterministic linear programming (DLP) model. However, this information is conveyed by applying the FCFS heuristic as a future strategy rather than using the optimal ex-post profits as in the RLP heuristic. We show that MC-FCFS performs approximately as well as the RLP heuristic at a much lower computational cost and much better than the DLP heuristic at maximizing profit in a multi-night hotel booking setting with or without planned upgrades.

Key-Words: Network revenue management, Monte-Carlo simulations, randomized linear programming.

JEL classification: C44, C63.

*Université de Lyon, Lyon, F-69007, France; CNRS, GATE Lyon Saint-Etienne, Écully, F-69130, France. houy@gate.cnrs.fr.

†EM Lyon Business School, Écully, F-69130, France. legrand@em-lyon.com.

1 Introduction

In revenue management, capacity control aims at using quantity based instruments to maximize revenues. In general, due to the uncertainty in the arrivals of demands and in the arrival pattern of different customer types, deriving an optimal dynamic control strategy is a difficult task. This task is even more difficult –and even in practice infeasible– when considering the revenue management of industries or services offering multiple resources. In this case, the sale of a product to a customer implies the sale of a bundle of resources. As it is standard, we will refer to this situation as network revenue management. Such problems may arise in a number of industries, such as airlines (with multiple-leg flights), hotels and hospitality (with multiple night stays) or car rental (with multiple day rentals).

In addition to demand uncertainty, revenue management is also made difficult by the possibility of upgrades. Upgrade possibilities appear when the seller offers several qualities of the same product and when these different qualities are partly substitutable. In particular, when offered at no extra cost, it is assumed that a customer will accept a better quality product than the one that was initially requested. Upgrading a customer may enable the seller to better match the demand and to avoid denying a request even though this denial would also mean leaving better quality products unused.

In this paper, we study the question of optimal product allocation in an industry facing dynamic stochastic demand arrivals. Our results though valid for any network industries featuring upgrade possibilities, will be applied to the hotel industry in the numerical exercise. Various product qualities can be booked in advance and involve several resources (stays of several consecutive nights). When a demand arrives, several options are available. First, the demand can be accepted at the quality level demanded and then it generates a sure profit. Second, the demand can be upgraded: instead of obtaining the required quality, the customer enjoys for the same price a better product. The seller may therefore

earn a profit instead of rejecting a demand and leaving a product unused (such as an empty room for a hotel). Third, the demand can be rejected. Even though the refusal may be due to a physical constraint –there is no vacant product left– the refusal may also occur due to capacity control: accepting the current demand may harm future profits.

We model the product allocation process as a dynamic optimization problem. Due to the several dimensions of demands –quality, and in the case of hotels, the start date and length of stay– the optimization problem is subject to the Bellman’s curse of dimensionality and the optimal solution is not computationally tractable. This is the case in general for network revenue management. Revenue optimization needs to rely on approximations of various types. These approximations have to reach a good compromise between the performance in terms of generated revenues and of computational efficiency.

Most of capacity control models base the acceptance decision of a demand on the comparison of its displacement cost to its revenue. The displacement cost (or opportunity cost) associated to a demand is the expected future loss in revenue due to the acceptance of this demand and to the usage of the capacity compared to the refusal of the demand and the option to still have an available capacity in the future. The demand will be accepted if its revenue offsets its displacement cost. In the case of several qualities, displacement costs for acceptances in each accessible quality are computed and the demand is allocated in the quality with the smallest displacement cost, and rejected if the demand revenue is smaller than all displacement costs.

We propose a new heuristic to compute the displacement cost. When a demand arrives, our heuristic enables us to compute the expected revenues associated to each of the various options: acceptance in one of the possible qualities or refusal. For each of the possible options, the expected revenues are computed by Monte-Carlo simulations of future demands. The revenue associated to every demand path is computed using a first-come-first-served rule. Revenues are then averaged over all possible paths. The displacement cost, as well

as the acceptance decision, is then straightforward to deduce.

Our Monte-Carlo heuristic with first-come-first-served rule (henceforth, MC-FCFS) presents several appealing features. First, it performs well with respect to other heuristics in terms of generated revenues. Second, it is fast and simple to compute. This heuristic makes it possible to take into account the uncertainty of future demand arrivals, as well as to offer a decision tool that is time-varying and that evolves with accepted demands.

For assessing the performance of our heuristic, we compare it to several standard benchmarks. Besides considering the standard first-come-first-served heuristic (henceforth, MCFS), we also consider two heuristics based on the computation of the displacement cost using deterministic linear programming (DLP) –and randomized linear programming (RLP)– as methods that are respectively denoted DDLP and DRLP (where the first D stands for dynamic). For the sake of homogeneity, we use these benchmark heuristics in conditions that are similar to those of our MC-FCFS heuristic: The displacement cost is computed for every demand arrival as the difference in expected profits. For the DDLP heuristic, we compute the displacement cost using a DLP heuristic. For the DRLP heuristic, we compute the displacement cost using the RLP model: Running Monte-Carlo simulations for future streams of demands, we compute for each simulation the optimal ex-post profit. From a practical perspective, these two options are obviously barely feasible (especially for DRLP), but this enables us to compare the three heuristics –MCFS, DDLP and DRLP– on very similar grounds. We finally compare the average realized profits generated by the three heuristics using the same simulation paths.

Our numerical results unambiguously show that MC-FCFS performs better than DDLP and obviously better than FCFS. The gains from using MC-FCFS over DDLP are approximately as much as 100% to 200% more than those from using DDLP over FCFS. They are therefore very sizable and all the more so as DDLP is recomputed for every demand. Compared to DRLP, MC-FCFS performs in much the same way but DRLP does not perform

significantly better than MC-FCFS, although it is far more costly from a computational point of view. To conclude, besides performing well, MC-FCFS is scalable and can be used in a realistic operational environment.

Our paper belongs to several strands of the literature. First, it belongs to the very large body of literature on revenue management. Weatherford and Bodily (1992), McGill and Ryzin (1999), Chiang et al. (2007) and Çetiner (2013) provide detailed surveys on the different approaches to revenue management in general. Excellent reviews can also be found in the textbooks by Phillips (2005) and Talluri and Van Ryzin (2005).

Second, we also contribute to the literature on capacity control in network revenue management. In this framework, determining an optimal strategy is not feasible due to the Bellman's curse of dimensionality. The problem needs to be approximated. Talluri and Van Ryzin (2005) distinguish two types of methods: (i) approximating the value function or (ii) decomposing the network problem into single-leg problems that can be solved independently.¹ Method (i) consists in imposing some restrictions to limit the dimensionality of the problem and to make it tractable using mathematical programming models. The DLP model (de Boer et al., 2002; Glover et al., 1982; Williamson, 1992) assumes that only the average demand affects the optimization. The optimization becomes therefore deterministic and can be achieved using standard linear programming techniques. The probabilistic non-linear programming (PNLP) model (e.g., Talluri and Van Ryzin, 1998) considers expected sales instead of expected demands, which yields a non-linear optimization problem. Even if these models can be used to implement a strategy based on booking limits, they have mainly been used to implement bid-price strategies. Bid-price is a first-order approximation of the displacement cost using Taylor expansion and the gradient of the value function. The bid-price associated to a demand for a product involving several resources is thus equal to the sum of the individual bid-prices of these multiple resources. For example in the DLP model, these individual bid-prices can be computed from Lagrange

multipliers and thus use the dual program. For computational efficiency reasons, bid-prices are not computed for every demand but once for a predetermined period. Williamson (1992) and de Boer et al. (2002) show that the DLP model generates bid-prices yielding higher expected revenues than the PNL model. However, weaknesses of the DLP based bid-prices have been pointed out by Talluri and Van Ryzin (1999), who have proposed the RLP model to address some of these issues. Instead of computing bid-prices using demand averages as in DLP, RLP computes average bid-prices using different realizations of the demand. RLP is shown to generate better performances than DLP.

Bertsimas and Popescu (2003) discuss the drawbacks of approximating the displacement cost by additive bid-prices, that may not even be well-defined in the absence of a unique dual solution. They show that directly using the displacement cost (as the discrete difference of expected future profits) provides better performances than using additive bid-prices. In our paper, we follow their approach and directly compute displacement costs.

Another drawback of the bid-price approach –or even of the one proposed by Bertsimas and Popescu (2003)– is that bid-prices are only computed once in a while and exhibit a very weak dependence on the dynamics of demand arrivals and on the current availability of products. Bid-prices are constant between two computations and the acceptance decision is therefore independent of the time in the day and of the seasonality of demand arrivals. Adelman (2007) and Topaloglu (2009) propose two different methods –that rely on a linear approximation of the value function– to generate bid-prices that are time-dependent, and even capacity-dependent in Topaloglu (2009)’s case. We circumvent this difficulty in our approach by proposing a heuristic that is fast enough to be computed at each arrival –that is therefore time-and-capacity-dependent by construction. We directly compare our heuristic to DLP and RLP methods that are computed at each demand arrival to avoid discussing the frequency of reoptimizations. Even in this case, our heuristic delivers performances better than DLP and similar to RLP.

Finally our paper also belongs to the literature on planned upgrades, that we take into account in our model. Among the few papers integrating planned upgrades and capacity control, we can cite Shumsky and Zhang (2009), Gallego and Stefanescu (2009) and Steinhardt and Gönsch (2012) for instance.

The remainder of the paper is organized as follows. In Section 2, we describe our setup. In Section 3, we explain the different heuristics we use and in particular the original MC-FCFS. Section 4 presents results of the quantitative comparison. Section 5 concludes.

2 Setup

As a preliminary, we clarify a couple of notations. We denote by $\mathbb{N} = \{0, 1, \dots\}$ (resp. \mathbb{N}_*) the set of natural numbers including 0 (resp. not including 0). We denote by $\llbracket p, n \rrbracket$ the set of natural numbers from $p \in \mathbb{N}$ to $p \leq n \in \mathbb{N}$. The set of (strictly) positive real numbers is \mathbb{R}^+ (\mathbb{R}_*^+) and a segment on \mathbb{R}^+ is denoted $[x, y]$ as it is standard.

We consider a hotel with rooms of different qualities. Let $I = \llbracket 1, I \rrbracket$ be the set of qualities.² Qualities are ordered: rooms of quality 1 offer the best quality, while rooms of quality I offer the lowest quality. The number of rooms of quality $i \in I$ is equal to N_i .

A booking for the hotel is an element $(B_{i,n})_{(i,n) \in I \times \mathbb{N}}$ of $\mathbb{N}^{I \times \mathbb{N}}$ such that $\forall (i, n) \in I \times \mathbb{N}, 0 \leq B_{i,n} \leq N_i$. The set of bookings is denoted \mathcal{B} . The quantity $B_{i,n}$ specifies the number of rooms of quality $i \in I$ for night $n \in \mathbb{N}$ that have already been attributed. When $B_{i,n} = 0$, all rooms of quality i for night n are empty; they are all booked when $B_{i,n} = N_i$. Time is continuous. At the start our study, at $t = 0$, the hotel is empty.

The hotel faces dynamic demands which can be described as follows. A demand d is a triplet (i, h, l) where (i) $i \in I$ is the requested quality, (ii) $h \in \mathbb{N}$ the requested first night of the stay and (iii) $l \in \mathbb{N}$ the requested length of the stay (the length of a stay is 0-indexed so that requested nights are $\llbracket h, h + l \rrbracket$). We denote by \mathcal{D} the set of demands. Demands arrive following inhomogeneous and independent stochastic Poisson processes in

continuous time. The instantaneous arrival rate of demand $d \in \mathcal{D}$ at time t is $\lambda_t(d)$. This abstract modeling enables us to consider demand seasonality (such as differences between workweek and weekend nights) and to discard the bookings of past nights. Hence, we will assume that the requested first night h is greater than the demand arrival time t : $h \geq \lfloor t \rfloor$.

For any booking $B \in \mathcal{B}$, any demand $d = (i, h, l) \in \mathcal{D}$ and any quality $i' \in I$, we define the element $(B \oplus_{i'} d)$ of $\mathbb{N}^{I \times \mathbb{N}}$ by

$$\forall i'' \in I, \forall n'' \in \mathbb{N}, (B \oplus_{i'} d)_{i'', n''} = B_{i'', n''} + \mathbb{1}_{i''=i'} \mathbb{1}_{n'' \in \llbracket h, h+l \rrbracket},$$

where $\mathbb{1}_A$ is a characteristic function. Then, $(B \oplus_{i'} d)$ is the new state of the booking if demand d has been allocated to a room with quality i' . We say that d is acceptable for B at quality i' if $(B \oplus_{i'} d)$ is a booking, i.e. $(B \oplus_{i'} d) \in \mathcal{B}$. For the sake simplicity, we also define, for any booking $B \in \mathcal{B}$ and any demand $d = (i, h, l) \in \mathcal{D}$, $(B \oplus_0 d) = B$.

Prices for renting rooms are considered exogenous. The price charged for a room with demanded quality $i \in I$ and for the night $n \in \mathbb{N}$ is denoted $p(i, n) > 0$. To avoid irrelevant situations, we assume that prices must be strictly positive. For any demand $d = (i, h, l) \in \mathcal{D}$, we define the price of demand d as $p(d) = \sum_{k \in \llbracket h, h+l \rrbracket} p(i, k)$.³ Although individual night prices depend on the night and may include seasonal effects, demand prices are a linear mapping of individual night prices. Hence, our pricing function does not include any discount or any rebate related to demand bundles. Notice that for any demand $d \in \mathcal{D}$, $p(d)$ depends on the requested quality. As we will see in the remainder, we will only allow for upgrades (allocation to a better quality at no extra cost) but not for upsells (allocation to a better quality subject to an extra payment).

Once a demand arrives at the hotel, the hotel manager needs to decide upon the acceptance (or the refusal) of the demand and in the case of acceptance which quality to award. Then, an allocation strategy is a function $X : \mathcal{B} \times \mathcal{D} \times \mathbb{R}^+ \rightarrow I^0 = I \cup \{0\}$ such that for any booking $B \in \mathcal{B}$, any demand $d \in \mathcal{D}$ arrived at time $t \in \mathbb{R}^+$, $X(B, d, t)$ is the quality of

the room awarded to demand d . If $X(B, d, t) = 0$, demand d is rejected. For the sake of simplicity, we denote $I_i = \{1, \dots, i\}$ as the set of qualities better than i and $I_i^0 = \{0\} \cup I_i$ as the set of qualities better than i or rejection.

We do not allow overbooking so that the hotel manager needs to make sure that the following room availability condition is satisfied.

$$\forall B \in \mathcal{B}, \forall d \in \mathcal{D}, \forall t \in \mathbb{R}^+, (B \oplus_{X(B,d,t)} d) \in \mathcal{B}. \quad (1)$$

Moreover, we will impose that there can be no downgrade of demands.

$$\forall B \in \mathcal{B}, \forall d = (i, h, l) \in \mathcal{D}, \forall t \in \mathbb{R}^+, X(B, d, t) \in I_i^0. \quad (2)$$

Let \mathcal{X} be the set of allocation strategies satisfying conditions (1) and (2).

In order to write the profit maximization dynamic problem faced by the hotel manager, we define the function $V : \mathcal{B} \times \mathbb{R}^+ \rightarrow \mathbb{R}$ denoting the maximum expected revenue for any given booking and at any time. Function V must satisfy the following Bellman equation:

$$\forall (B, t) \in \mathcal{B} \times [0, T], V(B, t) = \mathbb{E} \left[\max_{X \in \mathcal{X}} \left\{ p(\tilde{d}) \mathbb{1}_{X(B, \tilde{d}, \tilde{\tau}) \in I} + V(B \oplus_{X(B, \tilde{d}, \tilde{\tau})} \tilde{d}, \tilde{\tau}) \right\} \right], \quad (3)$$

$$\forall B \in \mathcal{B}, \forall t \geq T, V(B, t) = 0, \quad (4)$$

where the expectation $\mathbb{E}[\cdot]$ is taken with respect to the first arrival demand time $\tilde{\tau} \geq t$ and the demand type \tilde{d} . Equation (4) includes the boundary condition that holds for some $T \in \mathbb{R}^+$. The date T is the program horizon, possibly arbitrarily large. We can rewrite Equation (3) to introduce the displacement cost:

$$V(B, t) = \mathbb{E} \left[\max_{X \in \mathcal{X}} \left\{ p(\tilde{d}) \mathbb{1}_{X(B, \tilde{d}, \tilde{\tau}) \in I} - DC(B, \tilde{d}, X(B, \tilde{d}, \tilde{\tau}), \tilde{\tau}) \right\} \right] + \mathbb{E} [V(B, \tilde{\tau})], \quad (5)$$

where $DC(B, d, j, \tau)$ for $B \in \mathcal{B}, d \in \mathcal{D}, j \in I^0, \tau \in \mathbb{R}^+$ is the displacement cost of allocating the demand d in quality j (possibly refusing if $j = 0$) of booking B at date τ . Note that if $j = 0$, the displacement cost is zero. A demand will be accepted if the benefit associated

to this demand –here the demand price– is not smaller than all displacement costs and it is allocated in the quality with the smallest displacement cost.

As a straightforward consequence, the optimal allocation function $X^* \in \mathcal{X}$ must satisfy the following condition: $\forall B \in \mathcal{B}, \forall d = (i, h, l) \in \mathcal{D}, \forall t \in \mathbb{R}^+$,

$$X^*(B, d, t) = \arg \max_{j \in I_i^0, B \oplus_j d \in \mathcal{B}} \{p(d)\mathbb{1}_{j \in I} - DC(B, d, j, t)\}. \quad (6)$$

Due to the Bellman’s curse of dimensionality, function V , and hence the displacement cost DC and the optimal allocation function X^* cannot be computed in general. Instead, we will try to test different heuristics in \mathcal{X} in order to obtain as high a profit as possible.

Not solving our problem but still giving a good benchmark, we introduce the ex-post maximum profit. Once stochastic demands have realized their values with a number of occurrences n_d for each demand $d \in \mathcal{D}$, we can easily compute the ex-post optimal solution by solving the following linear maximization problem: $\forall B \in \mathcal{B}, \forall (n_d)_{d \in \mathcal{D}} \in \mathbb{R}^{\mathcal{D}}$,

$$V^{ExPost}(B, (n_d)_{d \in \mathcal{D}}) = \max_{(x_{d,j})_{(d,j) \in \mathcal{D} \times I}} \sum_{d=(i,h,l) \in \mathcal{D}, j \in I_i} p(d)x_{d,j}, \quad (7)$$

subject to:

$$\sum_{d=(i,h,l) \in \mathcal{D}} x_{d,j} \mathbb{1}_{n \in [h, h+l]} \leq N_j - B_{j,n}, \quad \forall j \in I, \forall n \in \mathbb{N}, \quad (8)$$

$$\sum_{j \in I} x_{d,j} \leq n_d, \quad \forall d \in \mathcal{D}, \quad (9)$$

$$x_{d,j} \geq 0, \quad \forall d \in \mathcal{D}, \forall j \in I. \quad (10)$$

The quantity $x_{d,j}$ corresponds to the number of demands of type d allocated in quality j weakly better than i . The ex-post program (7)–(10) maximizes the realized profit (equation (8)) subject to three constraints. Equation (8) states that the number of allocated nights, for any night and any quality, cannot exceed the hotel capacity. Conditions (9) and (10) impose the allocated night for a given quality to be positive and not larger than the number of occurrences of this demand. Notice that the ex-post program depends on the initial booking B and on the vector of demand arrivals $(n_d)_{d \in \mathcal{D}}$ but not explicitly on time.

This ex-post optimal profit corresponds to the optimal profit in absence of uncertainty, or alternatively in presence of perfect foresight. It should be clear that this ex-post optimal profit would be greater than the optimal profit (solution of program (3)) –if the optimization would have been possible. Since all our heuristics are not expected to perform as well as an actual optimization, the profit of any heuristic will of course be (strictly) smaller than the ex-post optimal profit. The difference between the profits of our heuristics and the ex-post optimal profit can be split into two parts: (i) part of the difference stems from the fact that the heuristic is not an exact optimization and (ii) the rest of the difference comes from the uncertainty which generates a wedge between the optimal profit and the ex-post optimal profit. Obviously, we are mainly concerned about the first source of difference. In consequence, we will refer to this ex-post optimal profit as a benchmark and it will mainly be used to compare heuristics to each other.

3 Heuristics

The purpose of our paper is to compare the performances of different allocation strategies in terms of profit maximization.

3.1 First-come-first-served

We start with the first-come-first-served heuristic (FCFS, henceforth). This heuristic is straightforward and satisfies all demands as they arrive with the only constraints being room availability and no downgrades. Should several qualities be available to the customer, the latter is served with the lowest possible quality (without downgrade). Formally,

$$\forall B \in \mathcal{B}, \forall d = (i, h, l) \in \mathcal{D}, X^{FCFS}(B, d, t) = \max\{j \in I_i^0, (B \oplus_j d) \in \mathcal{B}\}. \quad (11)$$

Notice that X^{FCFS} is independent of time and we will write $X^{FCFS}(B, d)$ in the remainder.

This allocation strategy does not perform any actual optimization along the acceptance dimension of our problem nor along its update dimension. The acceptance is automatic,

no matter how likely the arrival of a customer offering a better deal (longer stay or better quality) is, *i.e.* whatever the second term in Equation (6). Not surprisingly, as we will see, the FCFS heuristic will not perform very well in optimizing profits. Still, we will use it as a benchmark as well as a building block for defining other heuristics.

3.2 Dynamic deterministic linear programming

The dynamic deterministic linear programming heuristic (DDLP, henceforth) is close to the bid-price one, which is the most widespread revenue management technique implemented in the industry. The core idea of DDLP is to decrease the complexity of the optimal allocation problem by substituting the expected future value of demands in Equation (3) by the value of expected demands.

First, we compute the expected demand $\overline{e_t(d)}$ of any type $d \in \mathcal{D}$ at time t :

$$\overline{e_t(d)} = \int_t^\infty \lambda_s(d) ds. \quad (12)$$

Then, we solve the profit maximization deterministic linear program with the simplification that the future realization of any demand $d \in \mathcal{D}$ will be exactly its expectation, $\overline{e_t(d)}$. This program, also depending on the current booking B , is similar to Equation (7):

$$\forall B \in \mathcal{B}, \forall t \in \mathbb{R}^+, V^{DLP}(B, t) = V^{ExPost}(B, (\overline{e_t(d)})_{d \in \mathcal{D}}). \quad (13)$$

Finally, the idea of the DDLP heuristic is to allocate demand d to quality j such that: (i) the no-downgrade condition (2) holds, (ii) the room availability constraint (1) is satisfied and (iii) the allocated quality minimizes the displacement cost. The DDLP strategy X^{DDLP} can thus be formalized as in equation (6), where the corresponding displacement cost DC^{DDLP} can be expressed as follows: $\forall B \in \mathcal{B}, \forall d = (i, h, l) \in \mathcal{D}, \forall j \in I_i, \forall t \in \mathbb{R}^+$,

$$DC^{DDLP}(B, d, j, t) = V^{DLP}(B, t) - V^{DLP}(B \oplus_j d, t). \quad (14)$$

Notice that this heuristic is very similar to the heuristic proposed by Bertsimas and Popescu (2003), where they consider, as we do, the explicit displacement cost as a discrete difference

of value functions instead of the bid-price. The only difference in our approach is that we dynamically compute the displacement cost for *each* demand. We do so to be able to compare our heuristic (MC-FCFS in Section 3.4) –that can easily be computed for every demand– to other heuristics on similar grounds.

3.3 Dynamic randomized linear programming

As we have just seen, the DDLP heuristic decreases the complexity of the problem faced by the hotel manager by considering the value of expected demands as expected future values. This simplification does not take into account the full distribution of future demands. The dynamic randomized linear programming heuristic (DRLP, henceforth) does not make this strong assumption but rather computes future values with Monte-Carlo simulations. Still, it simplifies the future allocation policies by considering the ex-post optimal profit for each random draw of future demands (as described in Equation (7)).

More precisely, the idea of the DRLP heuristic is to allocate demand d to quality j , such that: (i) constraints (1) and (2) hold and (ii) the chosen quality minimizes the displacement cost computed using average ex-post profits from Monte-Carlo simulations. The DRLP heuristic is parametrized by the number $\bar{k} \in \mathbb{N}_*$ of Monte-Carlo simulations. For each simulation $0 \leq k \leq \bar{k}$ at date t , a vector of future demands $(n_{t,d}^k)_{d \in \mathcal{D}}$ is drawn and the ex-post value function $V^{ExPost}(B \oplus_j d, (n_{t,d}^k)_{d \in \mathcal{D}})$ –defined in equation (7)– is computed for all qualities $j \in I_i^0$ in which the demand d can be allocated (rejection included).

As for DDLP, we can formalize the DRLP strategy $X_{\bar{k}}^{DRLP}$ as in equation (6). The corresponding displacement cost $DC_{\bar{k}}^{DRLP}$ can be expressed as follows: $\forall B \in \mathcal{B}, \forall d = (i, h, l) \in \mathcal{D}, \forall j \in I_i, \forall t \in \mathbb{R}^+$,

$$DC_{\bar{k}}^{DRLP}(B, d, j, t) = \frac{1}{\bar{k}} \sum_{k=1}^{\bar{k}} (V^{ExPost}(B, (n_{t,d}^k)_{d \in \mathcal{D}}) - V^{ExPost}(B \oplus_j d, (n_{t,d}^k)_{d \in \mathcal{D}})). \quad (15)$$

Because it relies on Monte-Carlo simulations, the DRLP heuristic is stochastic.⁴ This

heuristic is very demanding computationally since it requires up to $(I + 1) \times \bar{k}$ linear optimization programs for each incoming demand (and each optimization is quite demanding).

The DRLP method can be seen as an extension of the RLP-based bid-prices introduced by Talluri and Van Ryzin (1999) following the lines of Bertsimas and Popescu (2003). Instead of computing the bid-price, we consider the displacement cost as the discrete difference of (average) ex-post profits.

3.4 Monte-Carlo FCFS

We now introduce the original Monte-Carlo FCFS (MC-FCFS, henceforth) heuristic. This heuristic relies on the same idea as DRLP, taking the full distribution of future demands into account but computationally simplifying the problem at each demand arrival. The displacement cost is computed using Monte-Carlo simulations but for each simulation, the profit is computed using the FCFS heuristic instead of the ex-post optimization.

More precisely, the MC-FCFS heuristic allocates the demand d to quality j , such that: (i) constraints (1) and (2) hold and (ii) the chosen quality minimizes the displacement cost. As for DRLP, the MC-FCFS heuristic is parametrized by the number $\bar{k} \in \mathbb{N}_*$ of Monte-Carlo simulations. The difference is that for each simulation $0 \leq k \leq \bar{k}$ at date t , the FCFS (instead of the ex-post) profit $V^{FCFS}(B \oplus_j d, (d_{t,0}, \dots, d_{t,N_k}))$ is computed for all qualities $j \in I_i^0$ in which the demand d can be allocated (rejection included). For the simulation k at date t , the profit V^{FCFS} depends on the initial booking and also on the random draw of future demands $(d_{t,0}^k, \dots, d_{t,N_k}^k)$ –since the arrival time of demands is random, the number of demands N_k is not constant. This profit is straightforward to compute using the strategy X^{FCFS} defined in equation (11). For $B \in \mathcal{B}$ and $(d_1, d_2, \dots, d_N) \in (\mathcal{D})^N$,

$$V^{FCFS}(B, (d_1, d_2, \dots, d_N)) = p(d_1) \mathbb{1}_{X^{FCFS}(B, d_1) \in I} + V^{FCFS}(B \oplus_{X^{FCFS}(B, d_1)} d_1, (d_2, \dots, d_N)), \quad (16)$$

$$V^{FCFS}(B, ()) = 0, \quad (17)$$

where equation (17) is the initialization: in absence of future demands, the value function is zero. Notice that V^{FCFS} depends on the order of demands in the random draw.

The strategy $X_{\bar{k}}^{MC-FCFS}$ of the MC-FCFS heuristic can be computed as in equation (6) using the associated displacement cost $DC_{\bar{k}}^{MC-FCFS}$, which can be formalized as, $\forall B \in \mathcal{B}, \forall d = (i, h, l) \in \mathcal{D}, \forall j \in I_i, \forall t \in \mathbb{R}^+, \forall ((d_{t,1}^k, \dots, d_{t,N_k}^k) \in (\mathcal{D})^{N_k})_{1 \leq k \leq \bar{k}}$,

$$DC_{\bar{k}}^{MC-FCFS}(B, d, j, t) = \frac{1}{\bar{k}} \sum_{k=1}^{\bar{k}} (V^{FCFS}(B, (d_{t,1}^k, \dots, d_{t,N_k}^k)) - V^{FCFS}(B \oplus_j d, (d_{t,1}^k, \dots, d_{t,N_k}^k))). \quad (18)$$

Like the DRLP heuristic, the MC-FCFS heuristic is stochastic (and we also drop the dependence in the random draw of demands). However, it is much faster to compute than the DRLP heuristic. Indeed, at each demand arrival, up to $(I + 1) \times \bar{k}$ implementations of MCFS are required and this does not involve any optimization program solving. As the simulations in Section 4 will show, this heuristic is very quick to compute.

3.5 Analytical results

We now provide a number of analytical results about MC-FCFS. These results show that according to some dimensions, MC-FCFS outputs are comparable to optimal decisions. Even if the scope of these results is admittedly limited, they provide some boundaries on the behavior of the heuristic. Simulations in Section 4 will complete the picture.

We start by providing an example in which MC-FCFS provides the optimal allocation, while DDLP and DRLP heuristics fail to do so. This example aims at illustrating that although the profits generated by the FCFS heuristic are smaller than the ex-post profits, this does not imply that MC-FCFS profits generated are smaller than DRLP ones.

3.5.1 An example showing that MC-FCFS may be optimal

As in Talluri and Van Ryzin (1998), we consider a simple hotel with two nights and one room. In order to simplify computations, demands for these two nights can arrive at

four dates, within the same day, that we denote $t = 0 \leq t_0 < t_1 < t_2 < t_3 < 1$. No demand arrives once the first night (night 0) has started.

We keep our notation for demands but drop the index for quality. For instance, demand $(0, 0)$ starts in night 0 and is of length 1 night. At time t_0 , a demand $(0, 0)$ arrives with probability 1. At t_1 , a demand $(0, 1)$ arrives with probability 0.4 (while no demand arrives with the complementary probability 0.6). At t_2 , a demand $(1, 0)$ arrives with probability 0.6, while at t_3 , a demand $(0, 1)$ arrives with probability 0.6. No other demand arrives otherwise. The price for one night is \$250.

We compare the DDLP, DRLP and MC-FCFS heuristics to the optimal strategy. To do so, we compute the displacement cost in the different cases. We start with the expected profits at date t_0 if the demand $(0, 0)$ is refused. For the other dates, the problem is rather straightforward for all heuristics and does not deserve a thorough investigation.

- DDLP: the strategy (see equation (13)) is to accept one demand $(0, 1)$, which generates an expected profit of \$500.
- DRLP: the strategy (see equation (15)) is to accept one $(0, 1)$ demand with probability $1 - (0.6 \times 0.4) = 0.76$ and one $(1, 0)$ demand with probability $0.6 \times 0.4 \times 0.6 = 0.144$. This generates an expected profit of \$416.
- MC-FCFS: the strategy (see equation (18)) is to accept one $(0, 1)$ demand with probability $0.4 + (0.6 \times 0.4 \times 0.6) = 0.544$ and one $(1, 0)$ demand with probability $0.6 \times 0.6 = 0.36$. This corresponds to an expected profit of \$362.
- Optimum: simple backward computation shows that refusing demand $(0, 0)$ generates an expected profit of \$380.

Then, accepting the demand $(0, 0)$ at date t_0 yields an expected future profit that is constant for all heuristics and equal to \$150. Indeed, once the demand has been accepted,

there is only room left for the demand $(1, 0)$, which may arrive at t_2 with probability 0.6. All heuristics will accept this demand, if it occurs.

Table 1 gathers the displacement costs for the different heuristics as well as the resulting acceptance decision.

Heuristic	Displacement cost	Room price	Decision
DDL	\$ 350	\$ 250	Reject
DRL	\$ 266	\$ 250	Reject
MC-FCFS	\$ 212	\$ 250	Accept
Optimum	\$ 230	\$ 250	Accept

Table 1: Decision made at t_0 .

The MC-FCFS heuristic is the only one to provide the optimal decision and to accept the demand of t_0 . The DDL and DRL heuristics overestimate the possible gains resulting from demands arriving in the future, which explains why they prefer to refuse the incoming demand. This example makes it clear that the MC-FCFS heuristic is not systematically biased and may provide optimal decisions, even when other heuristics do not.

3.5.2 Analytical results

We provide two analytical results regarding the behavior of the MC-FCFS heuristic.

Monotonicity in quality allocation. We discuss the question of the upgrade in two dimensions. First, we show that when all demands concern stays of one night length, the MC-FCFS strategy always allocates a demand in the smallest available quality or rejects it. In other words, when stays only last for one night, the comparison of the displacement cost to the demand payoff only matters for one quality –the lowest admissible one. This result is similar to Steinhardt and Gönsch (2012, Proposition 1), which proved such a result for the optimal policy. Along this dimension, the MC-FCFS heuristic has an identical behavior to the optimum.

PROPOSITION 1 (MONOTONICITY IN QUALITY)

Let $\bar{k} \in \mathbb{N}_*$, $B \in \mathcal{B}$, $d \in \mathcal{D}$, $i, j \in I$ such that $B \oplus_i d \in \mathcal{B}$, $B \oplus_j d \in \mathcal{B}$ and $t \in \mathbb{R}^+$.

If for all $\tau \in \mathbb{R}^+$ and for all $d' = (i', h', l') \in \mathcal{D}$, $l' > 0 \Rightarrow \lambda_\tau(d') = 0$, then:

- The optimal value function $V_{\bar{k}}^{MC-FCFS}$ is increasing with quality: If $i \geq j$, then $V_{\bar{k}}^{MC-FCFS}(B \oplus_i d, t) \geq V_{\bar{k}}^{MC-FCFS}(B \oplus_j d, t)$.
- The displacement cost $DC_{\bar{k}}^{MC-FCFS}(B, d, i, t)$ is conversely decreasing with quality. If $i \geq j$, then $DC_{\bar{k}}^{MC-FCFS}(B, d, i, t) \leq DC_{\bar{k}}^{MC-FCFS}(B, d, j, t)$.

Proof of Proposition: Similarly to Steinhardt and Gönsch (2012, Proposition 1), the proof can be conducted by backward induction on time. \square

The condition requiring that for any $d' = (i', h', l')$, $l' > 0 \Rightarrow \lambda_\tau(d') = 0$ means that the result only holds when all demands concern stays of one night length. This condition therefore restricts the scope of Proposition 1.

As soon as several night requests are introduced, no result similar to Proposition 1 holds. We provide the following example. As in Section 3.5.1, we consider a hotel with two nights but two qualities 1 (superior) and 2 (standard). We assume that three demands arrive before the first night at dates $t_0 < t_1 < t_2$: (i) demand $(0, 0, 2)$ –demand for a stay of one night starting at the first night and in a room of quality 2– at t_0 ; (2) demand $(0, 1, 2)$ at t_1 ; (3) demand $(1, 0, 1)$ at t_2 . The optimum consists in upgrading the first demand to quality 1 and to accept both remaining demands. The upgrade occurs at t_0 even though a room of a lower quality is available, which shows that the result of Proposition 1 cannot hold in general when relaxing the constraint on the stay length. Finally, it is noteworthy that MC-FCFS also provides the optimal decision in this example. It also recommends an upgrade despite the available room of a lower quality. This does not occur with FCFS.

Asymptotic behavior. Our second result concerns the asymptotic behavior of the MC-FCFS heuristic. MC-FCFS provides the optimal strategies for low demand arrivals or

when the hotel becomes very large. In these rather extreme scenarios, MC-FCFS heuristic is proven to have the same behavior as the optimum. Together with Proposition 1, Proposition 2 shows that along some dimensions, MC-FCFS provides meaningful outcomes.

PROPOSITION 2 (ASYMPTOTIC ANALYSIS)

When the size of the hotel converges for all qualities to infinity, or when the demand arrival intensity converges to zero, the MC-FCFS profits converge to the optimal ones.

Proof of Proposition: For low demand arrivals or when the hotel becomes very large, the constraint (1) never binds almost surely and we can then prove that no demand is (almost surely) rejected. \square

4 Computational results

For our simulations, we consider a hotel with 20 rooms in two quality classes with $N_1 = 2$ and $N_2 = 18$. A thorough formal description of demand is provided in Appendix A. All the heuristics (but FCFS which does not require any consideration for the future) are computed at each demand arrival with a 14 day planning window: We only consider demands (or parts of demands) requesting night stays in the 14 nights following the current date. For robustness checks for this parameter and some others, see Appendix C.

We start with an empty hotel at date 0 and consider the profit for nights between days 21 and 35. We do not consider the first days so as not to take into account the initial ramp up of the hotel filling and we consider a 14-day period to compute profits because all our demand stochastic processes have a weekly periodicity. For computational reasons, the number of streams of future demands generated in order to compute DRLP (\bar{k}) is at most 16. MC-FCFS is computationally much easier to implement and allows us to generate many more streams of future demands. Hence, for MC-FCFS, we use at most 1024 streams of future demands after having checked for robustness of the results for this parameter.

In order to provide a “significant” ranking of the performances of our heuristics, we will test whether the difference between the average profits of the different heuristics is significantly positive. To do so, we compare heuristics using the same realizations of the demand processes and we use the paired t-test for the relative differences of profits of two heuristics. If we denote \bar{D} the average of profit differences over N simulations and σ_D their standard deviation, this test involves computing the statistic $t_p = \frac{\bar{D}}{\sigma_D/\sqrt{N}}$. Since our sample is large (100 simulations in the base case, 50 simulations for robustness checks) and does not contain any extreme outlier, we can assume that this statistic follows a normal distribution *-i.e.*, a Student distribution with an infinite number of freedom degrees. We will report the p-value of this one-sided test, which can be more easily understood than the statistic value itself.

In Table 2, we show the performance of all heuristics and the ex-post optimal profit relatively to the most common DDLP heuristic for 100 simulations. We also show the same results without planned upgrades, *i.e.* for a hotel with 20 rooms of quality 2.

Heuristic	Relative difference to DDLP	
	2 quality classes	1 quality class
Ex-post optimal	3.810%***	1.905%***
FCFS	-0.987%***	-0.976%***
DRLP ₁₆	1.878%***	1.073%***
MC-FCFS ₁₀₂₄	1.730%***	0.975%***

Table 2: Relative profit difference between heuristics and DDLP. 100 simulations. Statistical significance: * for the 5% level; ** for the 1% level and *** for the 0.1% level.

With no surprise, FCFS is the poorest heuristic at maximizing profits. As has already been said, this heuristic does not compute any optimization and only fills the hotel under the constraints of no downgrade and room availability.

MC-FCFS₁₀₂₄ and DRLP₁₆ perform much better than DDLP. They increase the revenue of the hotel substantially and approximately equally. The difference between the

performances of DRLP_{16} and MC-FCFS_{1024} is not significant at the 10% level for both 1 and 2 quality class cases. The gain to adopt MC-FCFS_{1024} or DRLP_{16} in place of DDLP is approximately as large as almost twice the gain to adopt DDLP instead of FCFS in the case of 2 quality classes, about 100% in the case of 1 quality class. Both MC-FCFS_{1024} and DRLP_{16} perform relatively better than DDLP when not only acceptance but also upgrade are embedded in the allocation decision.

In Figure 1, we show the relative profit difference between MC-FCFS or DRLP and DDLP as functions of the number of streams of demands generated, \bar{k} .




Figure 1: Relative difference between MC-FCFS_k (dark line) or DRLP_k (light line) and DDLP as functions of k (log2 scale). Dashed line for the 1% significance level.

With our parameters and one or two quality classes, DRLP performs better than DDLP at the 1% significance level when $\bar{k} \geq 2$. For computational reasons, we could not investigate DRLP for \bar{k} greater than 16. MC-FCFS performs better than DDLP at the 1% significance level when $\bar{k} \geq 8$ in the two quality classes case and when $\bar{k} \geq 16$ in the one quality class case. For $\bar{k} \geq 32$, there is no difference, at the 1% significance level, between the $\text{MC-FCFS}_{\bar{k}}$ heuristics for the different values of \bar{k} , in the one quality class case. The same results hold for $\bar{k} \geq 8$ in the 1 quality class case.

Then, in terms of performance at maximizing profit, MC-FCFS_{1024} is approximately as good as DRLP_{16} . However, MC-FCFS_{1024} obtains this result with a computation time

about 200 times faster. It is even faster to compute than DDLP for much better results. In Table 3, we show the computation time needed to find the allocation decision at each demand, averaged over all demands that arrived between dates 0 and 35 days. Simulations were run on a 16-cores virtual machine with 16Gbytes RAM and Intel Xeon X5650 2.67GHz CPU. Programs were coded in Java 1.6.0-31 and run in OpenJDK Runtime Environment. Moreover, it would be possible to further improve the computation time of MC-FCFS by trying to move demand draws as fixed costs (and not as variable costs). Indeed, as the computation time of the FCFS strategy is very low (below 10^{-5} seconds), most of the computation time of MC-FCFS is devoted to the draw of future demands.

Heuristic	Avg computation time (s)	
	2 quality classes	1 quality class
DDL _P	0.88	1.59
DRL _P ₁₆	97.47	89.49
MC-FCFS ₁₀₂₄	0.50	0.47
FCFS	8.46×10^{-6}	6.28×10^{-6}

Table 3: Average computation time needed to find the value of the allocation function for one demand depending on the heuristic.

5 Conclusion

The dynamic programming problem faced by the profit maximizing hotel manager boils down to comparing the price of demands to their displacement costs. Computing these costs implies knowing the future optimal strategy and the expected profit generated over the probability distribution of future demands. This future value computation has too many dimensions for the problem to be tractable. One of the solutions (DDL_P) is to decrease the complexity of this problem by considering a degenerate probability distribution of future demands (the average of the original probability distribution). By doing so, the future strategy considered in the computation of displacement costs is exactly considered as the

optimal one but the future demands are simplified. Another solution is to consider the full probability distribution of future demands (or a representative sample of it with Monte-Carlo methods) and to simplify the future optimal strategy. In this line, DRLP considers that the future strategy is the ex-post optimal one. By doing so, DRLP allows the hotel manager to better take into account the distribution of future demands but at a very high computational cost. Instead, we propose that the future strategy be the FCFS heuristic. It is noteworthy that in the literature, DRLP has traditionally been seen as a refinement of DDLP, while in our framework DRLP appears to be related to MC-FCFS.

By choosing FCFS as the future strategy, expected future profits are below the optimal value (while they are above in DRLP when the future strategy is the ex-post optimization). However, the decision of the hotel manager only involves displacement costs, which are discrete differences of expected future profits. The above ranking of FCFS and ex-post profits has no implication for the profits of MC-FCFS and DRLP. We have shown that MC-FCFS performs about as well as DRLP at optimizing profits but does so with a much greater computational efficiency. Moreover, both MC-FCFS and DRLP perform much better than DDLP at optimizing profits. We draw the conclusion that, in our model, the information about the probability distribution of future demands is crucial and is conveyed equally efficiently by MC-FCFS and DRLP.

Notes

¹Liu and Van Ryzin (2008) have pioneered the use of decomposition techniques for network problems. Their work has been influential and extended by many authors (Kunnumkal and Topaloglu, 2010; Meissner and Strauss, 2012; Miranda Bront et al., 2009; Zhang and Adelman, 2009). The focus of our paper is rather on mathematical programming approximations and we will not insist on this strand of literature.

²For the sake of simplicity and with no risk of confusion, we denote by I both the number of qualities and the set of qualities.

³We keep the already defined notation p with a slight lack of rigor but no risk of confusion.

⁴For the sake of conciseness, we drop the dependence in the sequence of random draws in the notation. Still, the reader should keep it in mind. In particular, in case of comparisons (see Proposition 1), it should be understood that the same sequence of random draws is considered.

References

- Daniel Adelman. Dynamic bid prices in revenue management. Operations Research, 55(4):647–661, July 2007.
- Dimitris Bertsimas and Ioana Popescu. Revenue management in a dynamic network environment. Transportation Science, 37:257–277, 2003.
- Demet Çetiner. Selected topics in revenue management. In Fair Revenue Sharing Mechanisms for Strategic Passenger Airline Alliances, volume 668 of Lecture Notes in Economics and Mathematical Systems, pages 3–30. Springer Berlin Heidelberg, 2013.
- Wen-Chyuan Chiang, Jason C.H. Chen, and Xiaojing Xu. An overview of research on revenue management: current issues and future research. International Journal of Revenue Management, 1(1):97–128, 2007.
- Sanne V de Boer, Richard Freling, and Nanda Piersma. Mathematical programming for network revenue management revisited. European Journal of Operational Research, 137(1):72–92, 2002.
- Guillermo Gallego and Catalina Stefanescu. Upgrades, upsells and pricing in revenue management. Working paper, Columbia University, 2009.
- Fred Glover, Randy Glover, Joe Lorenzo, and Claude McMillan. The passenger-mix problem in the scheduled airlines. Interfaces, 12(3):73–80, 1982.
- Sumit Kunnumkal and Huseyin Topaloglu. A new dynamic programming decomposition method for the network revenue management problem with customer choice behavior. Production and Operations Management, 19(5):575–590, 2010.
- Qian Liu and Garrett John Van Ryzin. On the choice-based linear programming model for network revenue management. Manufacturing & Service Operations Management, 10(2):288–310, 2008.
- Jeffrey I. McGill and Garrett J. Van Ryzin. Revenue management: Research overview and prospects. Transportation Science, 33(2):233–256, February 1999.
- Joern Meissner and Arne Strauss. Network revenue management with inventory-sensitive

- bid prices and customer choice. European Journal of Operational Research, 216(2): 459–468, 2012.
- Juan José Miranda Bront, Isabel Méndez-Díaz, and Gustavo Vulcano. A column generation algorithm for choice-based network revenue management. Operations Research, 57(3): 769–784, 2009.
- Robert Phillips. Pricing and Revenue Optimization. Stanford Business Books, 2005.
- Robert A. Shumsky and Fuqiang Zhang. Dynamic capacity management with substitution. Operations Research, 57(3):671–684, May 2009.
- Claudius Steinhardt and Jochen Gönsch. Integrated revenue management approaches for capacity control with planned upgrades. European Journal of Operational Research, 223 (2):380–391, 2012.
- Kalyan T. Talluri and Garrett John Van Ryzin. An analysis of bid-price controls for network revenue management. Management Science, 44(11):1577–1593, November 1998.
- Kalyan T. Talluri and Garrett John Van Ryzin. A randomized linear programming method for computing network bid prices. Transportation Science, 33(2):207–216, 1999.
- Kalyan T. Talluri and Garrett John Van Ryzin. The theory and practice of revenue management. International series in operations research & management science. Springer, New York, 2005.
- Huseyin Topaloglu. Using lagrangian relaxation to compute capacity-dependent bid prices in network revenue management. Operations Research, 57(3):637–649, 2009.
- Lawrence R. Weatherford and Samuel E. Bodily. A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing. Operations Research, 40(5):831–844, September 1992.
- Elizabeth Louise Williamson. Airline Network Seat Control: Methodologies and Revenue Impacts. PhD thesis, Massachusetts Institute of Technology, 1992.
- Dan Zhang and Daniel Adelman. An approximate dynamic programming approach to network revenue management with customer choice. Transportation Science, 43(3):381–394, 2009.

Appendix

A Demands

We assume that demands follow a homogeneous Poisson process of intensity Λ . Once a demand arrives at time t , the requested quality is $i \in I$ with probability $P^1(i)$. Independently of the quality, the first night of the stay will start at night h with probability $P_t^2(h)$ (see eq. (20) below). The length of the stay l only depends on h (see eq. (21) below). Formally, the arrival intensity of demand $d = (i, h, l)$ at time t , $\lambda_t(d)$ can be rewritten as:

$$\lambda_t(d) = \Lambda_t \cdot P^1(i) \cdot P_t^2(h) \cdot P^{3|2}(l|h) \quad (19)$$

We consider that customers may book nights that start at most 7 days after t (and cannot start before t). Between the night following t , night $\lfloor t \rfloor$, and night $\lfloor t + 6 \rfloor$, the probability $P_t^2(h)$ of the first night h of the stay is $\forall t \in \mathbb{R}^+, \forall h \in \mathbb{N}$,

$$P_t^2(h) = \frac{\mu \cdot (1 - \mu)^{h - \lfloor t \rfloor}}{\sum_{k \in \llbracket \lfloor t \rfloor, \lfloor t + 6 \rfloor \rrbracket} \mu \cdot (1 - \mu)^{k - \lfloor t \rfloor}} \mathbb{1}_{h \in \llbracket \lfloor t \rfloor, \lfloor t + 6 \rfloor \rrbracket} \quad (20)$$

where μ is the intensity of the first night process truncated to remain within the 7 next nights. Table App.1 contains values of $P_t^2(h)$ for our base case parameters.

First night ($h - \lfloor t \rfloor$)	0	1	2	3	4	5	6
Probability (%)	41.15	24.69	14.81	8.89	5.33	3.20	1.92

Table App.1: Probability distribution of the first night of the stay (relatively to the arrival date t of the demand).

We also consider that lengths of the stays cannot exceed 7 days. We will make the difference between week nights and weekend nights. In particular, the probability of staying two nights in a row will be higher during weekends than during the rest of the week. Keeping our exponential assumption valid for the length of the stay, $\forall h \in \mathbb{N}, \forall l \in \mathbb{N}$,

$$P^{3|2}(l|h) = \frac{\nu_{(h+l)\%7} \cdot \prod_{k \in \llbracket h+1, h+l \rrbracket} (1 - \nu_{k\%7})}{\sum_{j \in \llbracket 0, 6 \rrbracket} \nu_{(h+j)\%7} \cdot \prod_{k \in \llbracket h+1, h+j \rrbracket} (1 - \nu_{k\%7})} \mathbb{1}_{l \in \llbracket 0, 6 \rrbracket} \quad (21)$$

where $\nu_{(h+l)\%7}$ is the intensity of the stay length process truncated to last at most 7 nights ($h < t+7$). The fact that ν is indexed by $(h+l)\%7$ allows us to make the difference between workweeks and weekends. If w is the weekend length, we will set $\forall i \in \llbracket 0, 6-w \rrbracket, \nu_i = \nu_{ww}$ and $\forall i \in \llbracket 6-w, 6 \rrbracket, \nu_i = \nu_{we}$. Table App.2 gathers values of $P^{3|2}(l|h)$ for our base case parameters.

		Stay length						
		0	1	2	3	4	5	6
First day of the stay (%7)	0	80.02	16.00	3.20	0.64	0.13	0.01	0.01
	1	80.02	16.00	3.20	0.64	0.03	0.03	0.08
	2	80.02	16.00	3.20	0.16	0.13	0.41	0.08
	3	80.02	16.00	0.80	0.64	2.05	0.41	0.08
	4	80.02	4.00	3.20	10.24	2.05	0.41	0.08
	5	20.00	16.00	51.21	10.24	2.05	0.41	0.08
	6	20.00	64.01	12.80	2.56	0.51	0.10	0.01

Table App.2: Probability distribution of stay lengths as a function of the first night.

Finally, for normalization purposes, we consider the demand intensity for each quality (instead of Λ and P^1). More precisely, for any quality $i \in I$, the expected weekly demand for rooms of quality i the hotel is facing is⁵ $e_i = \int_{t \in [0,7]} \sum_{l,h \in \mathbb{N}} (l+1) \cdot P_t^2(h) \cdot P^{3|2}(l|h) dt$. We will call the demand intensity for each quality, $D_i = e_i / (7 \cdot N_i)$, the ratio between the expected demand and the available number of rooms for each quality over one week.

As an illustration, in Table App.3, we give, for our base case parameters, the expected demand for rooms of quality 2 and each night of the week.

Day	0	1	2	3	4	5	6
Demand	36.62	21.40	18.35	17.74	17.62	17.60	28.16

Table App.3: Daily expected demand over one week.

B Base case parameters

We summarize our base case parameters' values in Table App.4.

	Parameter	Value
Number of rooms	N_1	2
	N_2	18
Prices (Quality 2)	Sunday (0)	200
	Monday (1)	100.1
	Tuesday (2)	100.01
	Wednesday (3)	100.001
	Thursday (4)	100.0001
	Friday (5)	100.00001
	Saturday (6)	200.000001
Ratio of quality 1 rooms prices to quality 2 rooms prices		2.05
Demand parameters	μ	0.4
	ν_{ww}	0.8
	ν_{we}	0.2
	w	2
Demand intensities	$D_1 = D_2$	1.25

Table App.4: Base case parameters' values.

We roughly impose two price levels for each quality, where the two highest prices correspond to nights with the highest demand –see Table App.3. Due to weekly periodicity, 7 prices describe the full price function p . We set prices with some very small differences in order to avoid, as much as possible, multiple solutions in the maximization problems.

C Robustness checks

In Table App.5a, we show our results for a 21 day planning window and in Table App.5b, for different demand intensities.

Heuristic	Relative difference to DDLP	
	2 quality classes	1 quality class
Ex-post optimal	3.305%***	2.037%***
FCFS	-1.269%***	-0.995%***
DRLP ₁₆	1.637%***	1.152%***
MC-FCFS ₁₀₂₄	1.491%***	1.137%***

(a) Impact of a 21 day planning window.

Heuristic	Relative difference to DDLP	
	$D_1 = 1.0, D_2 = 1.$	
Ex-post optimal	2.854%***	
FCFS	-1.782%***	
DRLP ₁₆	0.654%***	
MC-FCFS ₁₀₂₄	0.703%***	
$D_1 = 1.0, D_2 = 1.5$		
Ex-post optimal	3.349%***	
FCFS	-1.452%***	
DRLP ₁₆	1.738%***	
MC-FCFS ₁₀₂₄	1.323%***	
$D_1 = 1.5, D_2 = 1.$		
Ex-post optimal	3.029%***	
FCFS	-1.393%***	
DRLP ₁₆	0.874%***	
MC-FCFS ₁₀₂₄	0.863%***	
$D_1 = 1.5, D_2 = 1.5$		
Ex-post optimal	2.518%***	
FCFS	-2.392%***	
DRLP ₁₆	1.274%***	
MC-FCFS ₁₀₂₄	0.989%***	

(b) Role of demand intensities

Table App.5: Relative profit difference between the different heuristics or the ex-post optimal and the DDLP heuristic. 2 quality classes, 14 days planning window, 50 simulations.