THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Fluid arrivals simulation for choice network revenue management

Citation for published version:

Barbier, T, Anjos, M, Cirinei, F & Savard, G 2019, 'Fluid arrivals simulation for choice network revenue management', *Journal of Revenue and Pricing Management*, vol. 18, no. 2, pp. 164-180.
<https://doi.org/10.1057/s41272-018-00172-4>

Digital Object Identifier (DOI):

[10.1057/s41272-018-00172-4](https://doi.org/10.1057/s41272-018-00172-4)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Journal of Revenue and Pricing Management

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Fluid Arrivals Simulation for Choice Network Revenue Management

Authors:

Thibault Barbier is Senior Scientist at ExPretio Technologies. He works on forecasting, simulation and optimization problems mainly for railway and bus companies. Simultaneously, he finishes a PhD from Polytechnique Montréal, Canada, focusing on optimization and simulation for revenue management with choice behavior. He holds a MSc in operations research from Polytechnique Montréal and a MEng in energy from Arts et Métiers Paristech, France.

+14384018688

thibault.barbier@polymtl.ca

Department of Industrial and Mathematical Engineering, Ecole Polytechnique Montreal, Montreal, Canada.

Miguel Anjos is Full Professor of Mathematics and Industrial Engineering of Polytechnique Montreal, where he holds the NSERC-Hydro-Quebec-Schneider Electric Industrial Research Chair on Optimization for the Smart Grid, and an Inria International Chair. His research interests are in mathematical optimization and its applications in engineering. His accolades include a Canada Research Chair, the Mritas Teaching Award, a Humboldt Research Fellowship, the title of EUROPT Fellow, and the Queen Elizabeth II Diamond Jubilee Medal. He is a fellow of the Canadian Academy of Engineering.

+15143404711 Ext. 4568

miguel-f.anjos@polymtl.ca

Department of Industrial and Mathematical Engineering, Ecole Polytechnique Montreal, Montreal, Canada.

Fabien Cirinei is an industrial researcher. He applies optimization, data science and simulation in industrial products. He has been Head of Revenue Management Systems Development in Expretio Technologies. He is currently Lead Data Science in retail industry at SSense. He holds a PhD in Operations research from Polytechnique Montréal, Canada, and Université de Valenciennes, France.

fabien.cirinei@expretio.com

ExPretio, Montreal, Canada.

Gilles Savard is a recognized expert in two-level mathematical programming. He has published over 100 scientific papers in major journals in the field of operations research and supervised nearly 80 students and highly qualified staff members. He set out an innovative methodological approach to revenue management that led ExPretio Technologies, a university spin-off, to markets revenue management tools to airline and railway companies. He was appointed CEO of the Institute for Data Valorisation (IVADO) in October 2016.

gilles.savard@polymtl.ca

Department of Industrial and Mathematical Engineering, Ecole Polytechnique Montreal, Montreal, Canada.

Correspondence:

Thibault Barbier, Department of Industrial and Mathematical Engineering, Ecole Polytechnique Montreal, Montreal, Canada.

Submission:

1

This paper is submitted as a Research paper.

Fluid Arrivals Simulation for Choice Network Revenue Management

Thibault Barbier^a, Miguel Anjos^a, Fabien Cirinei^b, Gilles Savard^a

^a*Department of Industrial and Mathematical Engineering, Ecole Polytechnique Montreal, Montreal, Canada.*

^b*ExPretio Technologies, Montreal, Canada.*

Abstract

Since the beginning of revenue management, simulation has been used to estimate the expected revenue resulting from an availability policy. It has also been used to verify the quality of forecasts by projecting them onto past availability policies. Recently, it has been used in simulation-based optimization approaches to find the best policy. Simulation thus has a central role in revenue management. We focus on the choice network revenue management (CNRM) problem that incorporates multiple resources and customer behavior. The traditional CNRM simulation is based on discrete customer arrivals; we propose a new approach based on fluid arrivals. Our estimator is biased, but we observe that the bias is often insignificant in practice and appears to be asymptotically null. Our approach consistently outperforms the traditional simulation in terms of estimation time and is thus a better choice for large instances. We also prove that it is equivalent to an approximation for the CNRM availability policy optimization problem. This equivalence limits the value of optimization-based simulation methods but allows us to propose heuristics to rapidly support the optimization.

Keywords: revenue management, fluid arrivals simulation, choice behavior, availability control, optimization based simulation

1. Introduction and literature

Revenue management (RM) aims to match offers to demand, given limited and perishable resources, in order to maximize revenue. The resources are sold as products whose availability is controlled over the reservation period. Choice network revenue management (CNRM) considers resources and customer behavior simultaneously. The CNRM variant includes customer buying logic such as *buy-up* and *buy-down* and allows customers to *buy-across* resources. See Talluri and van Ryzin (2004) for a complete review of RM. In this article, we focus on simulation for CNRM.

Simulation has been widely used in CNRM. First, it can measure the quality of the availability policies returned by optimization models. The CNRM availability policy optimization problem can be formulated as a dynamic program (DP; Liu and van Ryzin, 2008). The goal is to manage the product availability over the reservation period in order to generate the highest revenue. However, the DP rapidly becomes intractable, and approximations are used; they must find a balance between simplicity and realism. Simulation can then be used to estimate the expected revenue resulting from an availability policy. This is an indicator of the performance of the approximation model. For example, Bront et al. (2009) benchmarks static and dynamic approximations based on the simulated expected revenue. However, simulation often requires many evaluations, even for small instances. Meissner et al. (2013) obtain a precision of approximately 6% of the expected revenue after 2000 evaluations for an instance with only six products. If greater accuracy is required, the simulation will be slow or intractable.

Second, simulation is used within some CNRM approximation models. Talluri (2010), Kunnumkal and Topaloglu (2011), and Talluri (2014) solve a randomized approximation several times, simulating the customer arrivals to tighten their solutions. The simulation adds a stochastic component to deterministic approximations.

Third, simulation is used to forecast demand and predict behavior. It can apply forecasts to historical availability policies to determine what products are booked. The resulting bookings are then compared to the actual bookings made to evaluate the forecast accuracy, and the forecast can be modified by integrating the insights from the simulation. Cleophas et al. (2009) develop a simulation framework with an artificial demand generator to compare the performance

of forecasting methods. Fiig et al. (2014) propose a forecast accuracy measure for behavioral demand based on historical observations. They minimize the corresponding error to optimize the forecast parameters. In practice, simulation must be able to process large historical data sets in a reasonable time. Simulation is also used for the creation of training data with which to test methods, as in van Ryzin and Vulcano (2015) and van Ryzin and Vulcano (2017).

Fourth, simulation-based optimization methods have been explored for CNRM availability policy optimization because they accurately model the problem. Bertsimas and de Boer (2005) and van Ryzin and Vulcano (2008b) develop stochastic gradient descent for RM without choice behavior. Van Ryzin and Vulcano (2008a) and Chaneton and Vulcano (2011) generalize the method to CNRM for nonparametric choice behavior. Other researchers propose model-free methods. For example, Bijvank et al. (2011) integrate a stochastic gradient technique while Gosavi et al. (2005) experiment with simulated annealing and simultaneous perturbation (SP) methods. Optimization-based simulation approaches generally achieve only local convergence. Moreover, they require many evaluations of gradients or finite differences. For the largest instances, current simulation models are too slow.

Fifth, simulation can be used as a *what-if* tool to support decisions in CNRM. We can change one or more features (e.g., the availability policy or the resource capacity) and measure the changes in terms of revenue, bookings, load factor, or any relevant output. RM analysts use this to select promotions or group reservation deals. The Passenger Origin-Destination Simulator (PODS) was introduced by Boeing in the 1990s (Belobaba and Hopperstad, 1999) to analyze customers' RM preferences. It has since been further developed (Carrier, 2003; Weatherford, 2013; Carrier and Weatherford, 2015) and now belongs to PODS Research LLC. Eguchi and Belobaba (2004) use PODS to highlight the importance of group bookings for the Japanese airline market. Gorin and Belobaba (2004) use the software to investigate the potential of RM in a low-fare airline. Darot (2001) studies RM for airline alliances with PODS, and Frank et al. (2008) explain how to set up a stochastic simulation model for RM analyses. Frank et al. (2006) use simulation to measure the effects of continuous capacity adjustments for different allocation times. Doreswamy et al. (2015) use the Airline Planning and Operations Simulator (APOS) developed by Sabre to explore the impact of different RM methods. Bijvank et al. (2011) developed a complete Java simulation library for CNRM. When used as a what-if tool, simulation must quickly return an accurate expected revenue.

Simulation usually handles discrete customer arrivals; we refer to this as discrete arrivals simulation (DAS). The arrival process is stochastic, and the expected revenue is estimated by a Monte-Carlo approach (Gilks et al., 1995). In this approach, each evaluation considers a random discrete arrival sequence. We calculate the revenue by applying the availability policy to the sequence. We then average the revenues obtained to estimate the expected revenue. This estimator is unbiased and approaches the real expected revenue as the number of evaluations increases. The precision relies on the confidence interval (CI), which is proportional to the ratio between the evaluation variance and the root square of the number of evaluations. The variance depends on many complex factors and is thus difficult to calculate a priori. Increasing the number of evaluations will improve the accuracy of the revenue estimate. However, each evaluation must process every arrival in the DAS model, so this estimator is slow for large instances.

In this article, we introduce another way to use simulation to estimate the expected revenue in CNRM: we consider a continuous flow of arrivals. Our approach is called fluid arrivals simulation (FAS). It has been used by Kesidis et al. (1996) for ATM networks and by Figueiredo et al. (2006) for computer networks. To the best of our knowledge it has not been applied to CNRM.

FAS does not use the Monte-Carlo technique because it estimates the expected revenue in a single evaluation by neglecting the order of the arrivals. It takes about the same time as a few DAS evaluations. Consequently, FAS outperforms DAS in terms of estimation time in all our experiments. Because the estimation is direct, FAS is also invariant, but it is biased. This bias is difficult to determine in theory because of the mechanism of the behavior and availability policy. In practice, it is relatively small for large instances and seems to be asymptotically null.

We prove that FAS is equivalent to the choice deterministic linear program (CDLP; Liu and van Ryzin, 2008), which is a widely used approximation for CNRM availability policy optimization. This equivalence limits the value of optimization-based simulation methods for FAS because it is preferable to directly solve the CDLP. We conduct experiments that show the slow convergence of an SP method for FAS and the need for a good initial solution. However, this equivalence allows us to develop new approaches to support the solution of the CNRM problem. They benefit from the speed of FAS and help to reduce the solution time. The two approaches that we propose are: the selection of a good initial CDLP solution by simulation and the estimation of the CDLP policy for a simplified demand. Both methods are simple and greatly accelerate the CDLP in our experiments. Our estimator therefore

potentially has a wide range of applications.

The remainder of this paper is organized as follows. In Section 2, we present the CNRM notation and discuss discrete arrivals. We then present DAS, which is the traditional estimator of expected revenue in CNRM. Section 3 presents our FAS estimator. We describe the discrete changes that occur although the simulation is considered fluid. We then analyze the properties of the bias. In Section 4 we examine the use of our estimator for the CNRM availability policy problem. We start by presenting an SP algorithm for optimization-based simulation with FAS. We then prove that FAS is equivalent to CDLP, and finally we suggest some ways to support optimization with our estimator. Our computational experiments are reported in Section 5, and Section 6 provides concluding remarks.

2. Simulation for the CNRM

In this section, we start by giving the principal CNRM definitions. We then describe the process by which individual customers arrive during the reservation period and eventually buy a product. We finish by presenting the DAS estimator, which is the traditional simulation for CNRM.

2.1. Definitions

CNRM is based on resources $i \in I$, each with a capacity c_i . There are $m = |I|$ resources. These resources are incorporated into products $j \in J$ that are sold at a fare r_j . There are $n = |J|$ products. We denote by I_j the set of resources consumed by each product j . We denote by $S \subseteq J$ a set of products, and we call it an offer. These products are sold during the reservation period, from time $t = 0$ to time $t = T$. The resources perish at the end of the reservation period ($t = T$).

The goal is to control the availability of the products over the reservation period in order to maximize the revenue generated by the sales. Let $x \leq c$ be the vector of remaining capacities and $J(x) \subseteq J$ the set of products with nondepleted resources. We must find the availability policy formed by offers $O(t, x) \subseteq J$ for all $t \leq T$ and $x \leq c$ that maximizes the revenue. The set of products available at any time and for any remaining capacity is $S(t, x) = O(t, x) \cap J(x)$.

A segment $l \in L$ groups customers with the same choice behavior who are interested in the same products $C_l \subseteq J$. This behavior is reflected by the probability $P_l(j|S)$ of buying product j if offer S is proposed. The customers of a segment arrive during the reservation period according to a Poisson process with arrival rate λ_l (a constant). We define $\lambda(S) = \{\lambda_j(S)\}_{j \in J}$ to be the product arrival rate vector if offer S is proposed, where $\lambda_j(S) = \sum_{l \in L} \lambda_l P_l(j|S)$ for each component.

We use the *running instance* (RI) of Figure 1 to explain the following concepts and models. The reservation period

Figure 1: Running instance.

is $T = 1$, and the availability policy offers the first product for 0.3 periods and the second throughout the reservation period.

2.2. Arrivals process

The customers arrive over the reservation period according to a random Poisson process by segment. The segment arrival sequence is therefore a random process; we denote by Ω its set of realizations. Let $\omega \in \Omega$ be a random segment arrival sequence with $\bar{\omega}$ distinct arrivals. It is indexed by $k \in [1, \bar{\omega}]$ to identify each arrival chronologically. The k th arrival occurs at $t_\omega^k \in [0, T]$ for remaining capacities $x_\omega^k \leq c$ and corresponds to a customer of segment $l_\omega^k \in L$. We have $t_\omega^k \leq t_\omega^h$ and $x_\omega^k \geq x_\omega^h$ for all $h \in [k, \bar{\omega}]$.

Let $\psi(l, S)$ be the random purchase function, returning a purchase vector $\{\psi_j(l, S)\}_{j \in J}$. The sole component equal to one corresponds to the product bought. The set of realizations is Ψ , and the purchase function satisfies

$$E_\Psi[\psi(l, S)] = \{P_l(j|S)\}_{j \in J}. \quad (1)$$

We can calculate the random revenue R_ω^k generated by the k th arrival of any arrival sequence ω as follows:

$$R_\omega^k = R_\omega^{k-1} + r^\top \psi(l_\omega^k, S(l_\omega^k, x_\omega^k)), \quad \forall k \in [1, \bar{\omega}], \omega \in \Omega$$

where $R_\omega^0 = 0$. The random revenue R_ω of the entire arrival sequence ω is then:

$$R_\omega = \sum_{k=1}^{\bar{\omega}} r^\top \psi(l_\omega^k, S(l_\omega^k, x_\omega^k)), \quad \forall \omega \in \Omega. \quad (2)$$

We determine the CNRM expected revenue as follows:

$$E[R] = E_{\Omega \times \Psi}[R_\omega]. \quad (E[R])$$

As mentioned in Section 1, this expected revenue is fundamental for CNRM because it reflects the revenue received in practice. However, the combination of two realization sets, the availability policy, and the overlapping segments makes it almost impossible to calculate the expected revenue.

The calculation is possible for our small RI. Segment 1 arrives at least once between 0 and 0.3 with probability $1 - e^{-2 \times 0.3}$, and segment 2 arrives at least once over the reservation period with probability $1 - e^{-3}$. The expected revenue is therefore $(1 - e^{-2 \times 0.3}) \times 100 + (1 - e^{-3}) \times 300 = 330.18$.

2.3. Discrete arrivals simulation (DAS)

The expected revenue is usually estimated, and the traditional RM approach is based on the Monte-Carlo method. Instead of a complete enumeration, this method draws N segment arrival sequences ω^k with $k \in [1, N]$. A revenue $\tilde{R}_{\omega^k}$ for each sequence is obtained from Eq. (2) by choosing a random purchase. The expected revenue is then obtained by averaging these revenues:

$$\mu^D = \frac{1}{N} \sum_{k=1}^N \tilde{R}_{\omega^k} \xrightarrow{N \rightarrow \infty} E[R]. \quad (\text{DAS})$$

This DAS estimator computes each discrete customer arrival. According to the strong law of large numbers, μ^D converges almost surely to $E[R]$ as the number of arrival sequence increases. This estimator is therefore unbiased.

The central limit theorem gives an α CI CI_α^D for this estimator:

$$\text{CI}_\alpha^D = \frac{\gamma_\alpha \sqrt{\text{Var}[R]}}{\sqrt{N}}, \quad \forall \alpha \in [0, 1] \quad (\text{CI})$$

where γ_α is the $\frac{1-\alpha}{2}$ percentile of the normal distribution and $\text{Var}[R]$ is the variance, which measures the volatility of the revenue. The CI establishes that $\alpha\%$ of the values are in $\left[\mu^D - \frac{\text{CI}_\alpha^D}{2}, \mu^D + \frac{\text{CI}_\alpha^D}{2}\right]$. The precision of the DAS is thus inversely proportional to the square of the number of evaluations.

The variance is almost impossible to calculate and is thus estimated by the sample variance:

$$\sigma^D = \frac{1}{N} \sum_{k=1}^N (\tilde{R}_{\omega^k} - \mu^D)^2 \xrightarrow{N \rightarrow \infty} \text{Var}[R].$$

The speed of the DAS convergence depends on this variance. If the variance is high, many evaluations are necessary to increase the precision, as we can see in Eq. (CI).

Each arrival requires the central reservation system (CRS) to process the available products. The complexity of each evaluation is therefore proportional to the average number of CRS calls per evaluation, denoted CRS^D . Each

Figure 2: DAS convergence for the RI.

evaluation covers a sequence of $\bar{\omega}$ arrivals, so $\text{CRS}^D \approx E_\Omega[\bar{\omega}]$. The overall DAS complexity therefore depends on the revenue variance and the number of CRS calls per evaluation. Unfortunately, the arrival stochasticity and the availability policy logic often lead to a high variance. Moreover, for large instances the number of arrivals, and thus of CRS calls, can be considerable. DAS may be unable to estimate the expected revenue accurately in the time allowed.

We plot in Figure 2 the DAS estimation convergence of the RI expected revenue with a 95% CI. We observe that DAS is unbiased, as expected.

3. Fluid arrivals simulation (FAS)

In this section, we present our new approach to estimate the expected revenue. It is inspired by work on fluid simulation for queues and computer networks. We also detail here its mechanisms and properties.

3.1. Model formulation

For DAS, the order of the arrivals leads to stochasticity. The main idea of our approach is to consider the arrivals of each segment as a fluid rather than individuals. For example, five customers of a segment arriving over a reservation period of two intervals are considered as a segment arriving with a rate of 5/2. The order of the arrivals thus becomes unimportant, and the expected revenue can be calculated as follows:

$$\mu^F = E_\Psi \left[\int_{t=0}^T \sum_{l \in L} \lambda_l r^\top \psi(l, S(t, x)) \delta t \right] = \int_{t=0}^T \sum_{l \in L} \lambda_l r^\top E_\Psi [\psi(l, S(t, x))] \delta t.$$

In this continuous case, the expected value of ψ is simplified as in Eq. (1) and the FAS is

$$\mu^F = \int_{t=0}^T r^\top \lambda(S(t, x)) \delta t. \quad (\text{FAS})$$

The FAS estimator is therefore continuous and deterministic because it is calculated in a single evaluation ($N^F = 1$). It is also invariant ($\sigma^F = 0$).

3.2. Discrete changes

FAS is continuous but the function $S(t, x)$ is a set of products with discrete additions and removals. We assume that the number of changes is finite, which is the case in practice. We index by $k \in K$ these changes over time where K is the set of changes of size $n_K = |K| - 1$. Each change occurs at time t_k and for the remaining capacities x_k . The first change, $k = 0$, corresponds to the start of the reservation period: $t_0 = 0$ and $x_0 = c$. The final change, $k = n_K$, corresponds to the end of the reservation period: $t_{n_K} = T$ and $x_{n_K} = 0$. Between two changes, the set $S(t, x)$ is constant and denoted by $S_k = S_k(t_k, x_k)$. We can now rewrite the FAS calculation as follows:

$$\mu^F = \sum_{k=0}^{n_K} \int_{t=t_k}^{t_{k+1}} r^\top \lambda(S_k) \delta t = \sum_{k=0}^{n_K} r^\top \lambda(S_k) d_k \quad (3)$$

where $d_k = t_{k+1} - t_k$ is the time between two consecutive changes. There are three possible changes:

- A resource depletion that changes $J(x)$;
- A change in availability policy $O(x, t)$;
- The end of the reservation period, $S(t, x) = \emptyset$.

When a change occurs, we can easily determine the next change by calculating the minimum time step to the next resource depletion, the next change in the availability policy, or the end of the reservation period.

For the RI, the first change corresponds to the beginning of the reservation period. The offer is $S_0 = \{1, 2\}$ and the product arrival rate is $\lambda_0 = (2 \ 3)$. We then calculate the above time steps; they are indicated by an $\times$ in Figure 3. The

Figure 3: Determination of second change for RI.

minimum time step is $d_0 = 0.3$, corresponding to the policy change at $t = 0.3$. We have sold $d_0\lambda_0$ products between these two changes, and we have $t_1 = 0.3$, $S_1 = \{2\}$ and $\lambda_1 = (0 \ 3)$. We illustrate the possibilities for the third change in Figure 4. The minimum time step corresponds to the depletion of resource 2 at $t = 1/3$. The final change is the end

Figure 4: Determination of third change for RI.

of the reservation period. We have sold $d_0\lambda_0 + d_1\lambda_1 + d_2\lambda_2 = 0.3(3 \ 2) + (1/3 - 0.3)(0 \ 3) + (1 - 1/3)(0 \ 0) = (0.6 \ 1)$ products and $\mu^F = 360$.

The complexity of the FAS depends on the number of changes. At each change the CRS must compute the minimum time step to the next change. This CRS call, denoted CRS^F , may be more complex than that for DAS. It is almost impossible to determine CRS^F a priori, but in some cases we can find a bound on the number of changes. If there is no reopening of product sales over the reservation period, the only possible changes are resource depletion, product removal, and the end of the reservation period. Each of these changes occurs only once, so $CRS^F \leq n + m + 1$.

3.3. Bias and properties

With the RI, the estimate $\mu^F = 360$ is not equal to the theoretical expected revenue, $E[R] = 330.18$. This shows that FAS is a biased estimator of the CNRM expected revenue. We denote the FAS bias by $\theta^F = E[R] - \mu^F$ and approximate it by $\tilde{\theta}^F$ as follows:

$$\tilde{\theta}^F = \mu^D - \mu^F \xrightarrow{N \rightarrow \infty} \theta^F. \quad (\text{FAS Bias})$$

The relative estimated bias is $\Delta\tilde{\theta}^F = \frac{\mu^D - \mu^F}{\mu^D}$. The bias is explained by the situations where a discrete resource capacity is sold in fractional quantities to multiple customers. In contrast, in the DAS model a resource cannot be partially sold. For the FAS of the RI, the first resource is sold to 0.6 customers of the first segment.

The number of fractional situations depends on the instance. It is difficult to predict because it depends on the number of resources, the policy, and the demand.

We now show that FAS can underestimate ($\theta^F \geq 0$) as well as overestimate ($\theta^F \leq 0$).

Proposition 1. *The FAS bias can be positive or negative.*

Proof of Proposition 1. Consider two resources with $c_1 = 2$ and $c_2 = 1$, and two products with prices r_1 and r_2 such that $I_1 = \{1\}$ and $I_2 = \{1, 2\}$. Suppose there are two segments, both arriving at the rate 2 during a reservation period $T = 1$. We can easily show that $E[R] \approx \frac{7}{6}r_1 + \frac{5}{6}r_2$ and $\mu^F = r_1 + r_2$. Hence, $\theta^F = \frac{7}{6}(r_1 - r_2)$. By adjusting the values of r_1 and r_2 , we obtain either positive or negative bias. $\square$

Proposition 2. *The FAS bias can be arbitrarily large.*

Proof of Proposition 2. By adjusting r_1 and r_2 in the proof of Proposition 1, we obtain an arbitrarily large bias. $\square$

However, in practice the size of the bias is reasonable. Moreover, the relative bias is $\Delta\theta^F = \frac{\theta^F}{\mu^F} = \frac{r_2 - r_1}{7r_1 + 6r_2}$, which is in the range $-\frac{1}{7} \leq \Delta\theta^F \leq \frac{1}{6}$ and is thus relatively insignificant. It is difficult to theoretically determine the bias because it depends on the policy, the arrival stochasticity, and the resource capacity. Furthermore, it is mainly caused by phenomena occurring when one or more resources have a capacity close to one.

4. FAS and CNRM optimization

In this section, we focus on how FAS can solve or support the CNRM availability policy problem. We start by presenting an SP algorithm for FAS. We then prove the equivalence between FAS and one of the principal CNRM approximations. We finally propose two simple methods that use FAS to support the solution of this problem.

4.1. Optimization-based simulation

One of the most widely used methods in optimization-based simulation is the SP algorithm; see Spall (1998) for more details. For this algorithm, we use a product closing (PC) availability policy. It specifies a time $0 \leq t_j \leq T$ to close the sale of each product such that:

$$S^{PC}(t, x) = \{j \mid j \in J, t_j \geq t\}.$$

SP is a gradient descent method. We denote by t^k the vector of products closing at iteration k . This technique is based on the following approximation:

$$\frac{\partial \mu(t)}{\partial t_j} \Big|_{t=t_k} \approx \frac{\mu(t+h) - \mu(t-h)}{2h_j}$$

where $h = \{h_j\}_{j \in J}$ and $h_j = \frac{B}{k^\beta}$ with B a Bernoulli random variable and $\beta \in [0, 1]$ a tuning parameter. The next PC policy is thus obtained as follows:

$$t_j^{k+1} = \Pi_{0 \leq t \leq T} \left[t_j^k + \alpha^k \frac{\partial \mu(t)}{\partial t_j} \Big|_{t=t_k} \right], \quad \forall j \in J$$

where $\alpha^k = \frac{\alpha}{k}$ is the step size of the descent.

We did not study the properties of the function $\mu(t)$; see Spall (1998) or Gosavi (2015) for the convergence properties of SP.

4.2. Equivalence to CNRM optimization

Static approximations of the CNRM availability policy optimization problem avoid the discrete customer arrival complexity of the DP by considering a continuous and deterministic flow of customers. They all have the same structure:

$$\begin{aligned} R &= \max_q r^\top q & (\text{STATIC}) \\ \text{s.t.} \quad & Aq \leq c, \\ & q \geq 0. \end{aligned}$$

The CDLP is a static approximation based on an availability policy that controls the time $d_S \geq 0$ for which each offer must be proposed:

$$\begin{aligned} q &= \sum_{S \subseteq J} \lambda(S) d_S \\ \text{s.t. } \sum_{S \subseteq J} d_S &\leq T, \\ d_S &\geq 0, \quad \forall S \subseteq J. \end{aligned} \tag{CDLP}$$

The relationship between the CDLP and the DP is quite similar to the relationship between FAS and the exact $E[R]$. In both cases, the individual arrivals are replaced by their expected arrival rate. We can prove that the CDLP revenue is equivalent to the FAS estimation for any offer duration.

Proposition 3. $\mu^D = R^{CDLP}$.

Proof. The CDLP set of products with non-null duration $\tilde{S}_u$ is arbitrarily ordered over $u \in [0, m-1]$ because the CDLP has at most m sets with non-null durations. Each set has a duration $\tilde{d}_u$, giving $\tilde{t}_{u+1} = \tilde{t}_u + \tilde{d}_u$ with $\tilde{t}_{m+1} = T$ and $\tilde{S}_{m+1} = \emptyset$. We apply Eq. (3) to this policy. The first change occurs at $\tilde{t}_0 = 0$ and corresponds to the initial CDLP set $\tilde{S}_0$. The next change does not occur until $\tilde{t}_1$ because the STATIC capacity constraint ensures that no resources are depleted during this period, and the end of the reservation period is not reached because of the second CDLP constraint. We therefore prove by recurrence that we have $k = u$ until $k = m$ and hence

$$\mu^F = \sum_{k=0}^{n_K} r^\top \lambda(S_k) d_k = r^\top \left(\sum_{k=0}^m \lambda(\tilde{S}_k) \tilde{d}_k + \sum_{k=m+1}^{n_K} \lambda(\tilde{S}_k) \tilde{d}_k \right) = r^\top \sum_{k=0}^m \lambda(\tilde{S}_k) \tilde{d}_k = R^{CDLP}.$$

For $k > m$, we necessarily have $\tilde{S}_k = \emptyset$ to ensure the STATIC capacity constraint, so $\lambda(\tilde{S}_k) = 0$. $\square$

This proof was developed for the CDLP, but it is similar for many static approximations, so the equivalence is easily extended. The equivalence and the fact that the simulator is based on fluid arrivals show that FAS cannot improve the robustness of the static approximation solution by taking into account the arrival order stochasticity; this is in contrast to DAS.

4.3. Optimization support

Our approach can support the optimization although it does not improve the solution quality because of the arrival order stochasticity. Its rapidity and the equivalence result allow several applications. We propose two simple ideas:

- We can solve the CDLP for only the most significant part of the demand to reduce the solution time and then simulate the corresponding optimal policy by FAS to evaluate it for the full demand.
- We can use FAS with a metaheuristic to provide a good initial solution for the CDLP.

Both approaches are tested in the experiments.

5. Computational results

In this section, we perform numerical experiments on the following two estimators:

DAS is the traditional estimator described in Section 2.3. Recall that μ^D is the estimate of the expected value $E[R]$, and σ^D is the estimate of the variance $\text{VAR}[R]$. CPU^D is the running time, and the number of evaluations is N^D . CRS^D is the number of CRS calls, which for this estimator is equal to the number of arrivals. To stop the convergence, we use a 95% relative CI width $\Delta[IC_{95\%}^D] = \frac{[IC_{95\%}^D]}{\mu^D}$ or a maximum number of evaluations.

FAS is our new estimator introduced in Section 3. Recall that μ^F is the estimate of the expected value $E[R]$. There is only one evaluation and thus no variance ($N^F = 1$, $\sigma^F = 0$). The running time is denoted CPU^F . CRS^F is the number of CRS calls, corresponding for this estimator to the number of changes.

Figure 5: Markets for bus-line instance.

Bus-line is an instance of eight buses leaving every two hours from 07:00 to 21:00 from city A to cities B, C, D, E, and F. A total of 15 markets are served, as illustrated in Figure 5. Each bus has a capacity of 30 and there are $5 \times 8 = 40$ resources. Two fares (low, high) are offered for each trip, giving a total of $15 \times 8 \times 2 = 240$ products. In the bus industry, tickets are usually available at least two months in advance, so we set $T = 60$ days. In total there are $15 \times 8 = 120$ segments with nonparametric choice behavior and on average 5.3 products.

Airline is an instance based on the Delta Air Lines network limited to the five major US airports. The network has 20 markets, as illustrated in Figure 6. It has 115 resources that correspond to the flights. There are 1591 products, and the reservation period is $T = 360$ days. There are 438 segments with nonparametric choice behavior and on average 7.9 products.

Figure 6: Markets for airline instance.

The load factor (LF) is defined as $LF = \sum_{l \in L} \lambda_l / \sum_{i \in I} c_i$. We use PC times for the availability policy; these set the times when the sales of each product are closed.

5.1. Convergence and bias

In this section, we compare the convergence of the two estimators and analyze the bias of FAS. The term convergence is imprecise for FAS since it calculates the expected revenue in a single evaluation. However, this is a way to illustrate the differences between these two estimators. Figure 7 illustrates the convergence of the two estimators for the optimal availability policy returned by the CDLP. This approximation is explained in Section 4.2. We stop the DAS simulation after 1000 evaluations.

The most relevant observation is that FAS always overestimates the real expected revenue, by 6.9% on average for the bus-line and 2.8% for the airline. This overestimation arises because the availability policy is optimal. At

Figure 7: Expected revenue estimates μ^F and μ^D with respect to the number of evaluations N for the optimal CDLP availability policy. The FAS relative bias $\Delta\bar{\theta}^F$ and the DAS 95% confidence interval $IC_{95\%}^D$ are also reported.

optimality, the fluid aspect is emphasized because the optimizer takes it into account to maximize the revenue. The difference with DAS is thus at its peak because FAS returns exactly the optimizer revenue, as proved in Section 4.2.

We observe that the bias increases as the LF increases from 0.8 to 1.0 and is then approximately constant. This is verified by further experiments for the bus-line. With the notation $(LF, \Delta\bar{\theta}^F)$, we report $(0.2, 3.63\%)$, $(0.4, 4.24\%)$, $(0.6, 4.55\%)$, $(1.4, 7.39\%)$, $(1.6, 7.63\%)$, and $(1.8, 7.41\%)$. Therefore, the bias becomes constant once a certain LF is reached. This could be because the optimizer does not improve the revenue with an additional fluid aspect. Another explanation is that the fluid aspect situations seen in Section 3.3 are all captured from a certain LF.

We also note that there is a difference of magnitude in the bias for the two instances. The bus-line bias is around 2.5 times higher than that for the airline. This suggests the asymptotic nullity of the bias: the average capacity per resource is 30 for the bus-line and 180 for the airline. The fluid aspect situations are more absorbed in the airline, clearly because of the higher capacity. This also highlights the difficult of predicting the bias.

To further investigate these poor results of FAS, we generate random availability policies for different scenarios. These scenarios vary the percentage γ_{close} of PC times fixed to zero; the other PC times are randomly chosen according to a uniform law. When $\gamma_{close} = 0$, every product is offered, and when $\gamma_{close} = 1$ no product is offered. We report in Figure 8 the relative FAS bias for different relative widths of the DAS CI and with respect to the scenarios for the PC times and the LF. We selected 100 and 25 availability policies per scenario respectively for the bus-line and the airline. The full results are reported in the Appendix: see Table A.2 for the bus-line and Table A.3 for the airline.

We first observe that the relative difference in the bias is not as high as before. It is 7.7 times lower (6.9% to 0.9%) and 14 times lower (2.8% to 0.2%) respectively for the bus-line and airline instances. This confirms that the optimal policy emphasizes the fluid aspect to maximize the revenue. Therefore, the optimal availability policy is certainly the one with the highest bias.

We observe that the bias evolution does not seem to follow any specific trend and might be unrelated to γ_{close} . This confirms that the fluid aspect does not depend on any one factor but is a consequence of a more complex interaction between the demand, the policy, and the structure of the instance.

We note that the bias for the 5% relative CI has considerable variability, whereas the 1% and 0.5% biases are smoother and similar. This shows that the DAS convergence is not rapid. For precise estimation, a 1% relative CI seems appropriate.

We now investigate the relationship between the FAS bias and the instance structure. We vary the capacity of both instances by a factor $\gamma_{capacity}$, i.e., $c \rightarrow c \times \gamma_{capacity}$. For each capacity scenario, we maintain the LF by scaling up the

Figure 8: Relative bias $\Delta\tilde{\theta}^F$ with respect to the percentage γ_{close} of closed products.

Figure 9: Relative bias $\Delta\tilde{\theta}^F$ for the optimal and random policies when capacity is scaled and demand adjusted proportionally.

segment arrival ratio. In Figure 9, we report the evolution of the FAS bias for averaged random availability policies and for the optimal CDLP availability policy. The LFs are 0.8, 1.0, and 1.2. The capacity factor varies from 0.1 to 100. We first observe the same features as in the previous experiments. The optimal policy is always the one with the highest bias. It is on average between 3 and 30 times higher than the random policy bias. This figure also shows that the fluid aspect is more prominent in the bus-line instance (0 to 20%) than in the airline instance (0 to 8%). The main observation is that the bias decreases as the capacity and demand are scaled up. We have not proved that the bias is asymptotically null, but the results suggest this. However, an asymptotic result does not in practice determine the bias of an instance.

In conclusion, these experiments show that the bias is caused by a fluid aspect that is difficult to predict. It depends on a combination of mechanisms between the instance structure, the demand, and the policy. It is stronger when the

availability policy is optimized. However, the bias is low for random policies, is reduced for instances with a higher capacity, and appears to be asymptotically null.

5.2. Estimation time

We now compare the estimators in terms of estimation time. Estimations must be returned quickly to allow users to rapidly test several options before making a decision. Also, these estimators are often used in optimization-based simulation methods that require many estimations.

In Figure 10, we report the estimation times with respect to γ_{close} for the experiments of Section 5.1. The estimators compared are FAS and DAS for three relative CI widths: 5%, 1%, and 0.5%.

Figure 10: Estimation time for FAS and DAS ($[IC_{95\%}^D]$ is 5%, 1%, and 0.5%) with respect to γ_{close} .

The most important observation is that FAS is always faster, whatever the value of γ_{close} and the relative CI width. As expected, the difference increases with the relative CI width because DAS needs more evaluations to reach the necessary precision, as explained by Eq. (CI). Overall, FAS is faster than the 5%, 1%, and 0.5% DAS: respectively 2.7, 67.9, and 275.6 times faster for the bus-line and 2.4, 6.5, and 27.5 for the airline.

We note that the superiority of FAS is even more pronounced as γ_{close} increases. It is respectively 33.7, 87.8, and 1673.1 times faster for the bus-line and 3.4, 8.6, and 101.7 times faster for the airline when γ_{close} is 20%, 60%, or 90% in comparison with the 1% DAS. As γ_{close} increases, the number of changes decreases because more products are closed, and the relative DAS variance increases because the revenue depends on what products were closed. This is shown by the measures of variance and the CRS calls reported in Tables A.2 and A.3.

We observe that the estimation times are slightly lower for both estimators as the LF increases. For FAS and the LFs 0.8, 1.0, and 1.2, the average estimation times are respectively 0.05, 0.045, and 0.043 for the bus-line and 3.3, 2.9, and 2.6 for the airline. This is due to the higher demand that tends to consume resources faster and thus decreases the number of changes, as shown in Tables A.2 and A.3. For the 1% DAS and the LFs 0.8, 1.0, and 1.2, the average estimation times are respectively 3.3, 3.2, and 2.9 for the bus-line and 92.6, 79.2, and 73.7 for the airline.

In Table A.1, we report the measures of the FAS and DAS estimations for the optimal availability policy. We note that FAS is 18.5, 34.1, and 53.3 times faster for the bus-line and 5.6, 9.6, and 6.3 times faster for the airline than the 1% DAS for the 0.8, 1.0, and 1.2 LFs. This supports our observation on the relationship between the estimation time and the LF.

As expected, the estimation of the bus-line expected revenue is 7.0 times faster than that for the airline. This is mainly because the airline instance is larger in terms of resources, products, and segments.

In Figure 11, we report the time to compute a CRS call for the two estimators with respect to γ_{close} . For FAS, each CRS call corresponds to a change, and we must compute the next change by calculating the time step. For DAS, it corresponds to a customer arrival, and we must compute the available products. The data used is from Tables A.2 and A.3.

Figure 11: Evaluation time CPU/N for the two estimators, averaged over the load factors 0.8, 1.0, and 1.2.

We note that the CRS call time decreases as γ_{close} increases for both estimators. This is because both calculations are easier when fewer products are offered and thus for a higher γ_{close} .

We observe that it takes longer to compute a FAS change than a DAS arrival. On average, it is 56.3 times longer for the bus-line and 141.1 for the airline. This is because determining the next change is more complicated than simply computing the available products. We must determine for each remaining resource the time step to depletion and for each available product the time step to unavailability. The minimum time step corresponds to the next change. On the other hand, computing the available products involves simply checking if a product is available for the policy and if its resources have remaining capacity.

Note that the results for estimation times depend on how the estimators are coded. We tried to find the best implementation of each approach to give a fair comparison.

In conclusion, the experiments show that FAS outperforms DAS in terms of estimation time. As expected, the time to calculate each change (CRS^F) is greater than the time to compute each arrival (CRS^D). However, the number of arrivals may be large, depending on the desired precision and the number of arrivals per evaluation.

5.3. Optimization

We now compare the two estimators in terms of solving the CNRM problem. The goal is to use simulation to converge to the availability policy returning the best expected revenue.

We start by implementing the SP algorithm (Section 4.1) for the FAS estimator. We use the parameters $\alpha = 0.01$ and $\alpha = 0.001$ respectively for the random and the optimal starting point. We also set $\beta = 0.5$. In Figure 12, we report the convergence of this method (μ^F) for a random and an optimal starting point. The availability policies found during the convergence are simulated by the DAS estimator (μ^D).

The main message from Figure 12 is that the SP needs a good starting point to find a near-optimal solution. We stop the process after 1000 iterations, but the additional improvement is insignificant because of the step size h_j as explained in Section 4.1. We adjusted the parameters α and β , but the convergence was worse.

For the optimal starting point, the SP starts by worsening the solution and then converges to near-optimality. This is because the SP leaves the optimal area corresponding to a specific local maximum.

Figure 12: SP technique applied to the FAS estimator (μ^F). The solutions returned by FAS are simulated by DAS (μ^D).

The method performs 1000 iterations in approximately 12 min for the bus-line and between 50 and 120 min for the airline. It is much slower than solving the CDLP, and the solution is at best equivalent.

Note that both estimators have approximately the same shape over the SP convergence. The estimated DAS revenue is lower than that for FAS because of the bias explained in Section 3.3 and demonstrated in Section 5.1. We note that the bias is approximately constant, and that the optimization on FAS is similarly reflected on DAS.

We conclude that the equivalence with the CDLP makes it difficult for an optimization-based FAS simulation technique to be as efficient as the solution of this mathematical program. Moreover, FAS does not take into account the arrival order stochasticity to improve the robustness of the solutions.

However, the equivalence also allows us to develop methods to support the CNRM optimization. Without going into details, we present two simple examples. The results are reported in Figure 13, and the method is explained below.

Figure 13: Two examples of FAS optimization support in the solution of the CDLP for the airline instance (LF=1).

First, we use FAS to generate a good initial solution for the CDLP. FAS is fast, and we randomly generate as many availability policies as possible in 15 s. The best one is used as the CDLP initial solution. The results are

reported on the left of Figure 13 for the airline instance with $LF = 1$. The best availability policy had a revenue of 1154306.21 (indicated by a circle in the graph), which is 34% lower than the optimal solution. With our approach, we save approximately 700 s, which is worthwhile given the total solution time is 2000 s. However, a tabu search or a genetic algorithm might provide a better initial solution.

Second, we use FAS to determine the CDLP revenue given the full demand for availability policies found by a CDLP solved for a partial demand. We remove the product with the lowest probability from each segment consideration set. We wish to focus on the most important component of the demand. The results are reported on the right of Figure 13 for the airline instance with $LF = 1$. We observe that the final revenue is approximately 7% lower than that for the CDLP with full demand, but we also save approximately 700 s. We could also use the convergence over partial demand to supply a good initial solution for the CDLP with full demand.

In conclusion, the experiments of this section show that FAS is not necessarily a good estimator for an optimization-based simulation because of its equivalence to the CDLP. However, it is fast, and the equivalence allows it to efficiently support the CNRM optimization. The two approaches tested here could potentially be used in other applications.

6. Conclusion

We have proposed a new simulation estimator for CNRM. It estimates the expected revenue of an availability policy by considering fluid arrivals. Requiring only one evaluation, our approach is much faster than the traditional Monte-Carlo simulation based on discrete arrivals. Our estimator is therefore invariant but biased and can underestimate as well as overestimate. The associated bias is almost impossible to measure a priori and can in theory be arbitrarily large. Experiments show that the bias is minimal in practice for a large instance and seems to be asymptotically null. However, it is higher for the optimal availability policy for which the fluid aspect is emphasized. We investigated whether our estimator can solve or support the optimization of the availability policy. Tests on optimization-based simulation methods showed that a good starting point is crucial. We proved that our estimator is equivalent to a widely used approximation for this problem. It is thus preferable to solve the latter rather than use our simulation to converge locally. In particular, our estimator cannot take into account the arrival stochasticity to improve the solution robustness. The equivalence allows us to develop new methods to support the optimization, and we proposed two simple approaches that significantly accelerate the solution of the tested instance. Our estimator is promising because it is fast even for large instances and returns acceptable estimations.

Acknowledgment

The authors are grateful to the Natural Sciences and Engineering Research Council of Canada (NSERC), the Fonds de recherche du Québec en nature et technologies (FRQNT), and ExPretio technologies for funding and supporting this research.

References

- Belobaba, P., Hopperstad, C., 1999. Boeing/mit simulation study : Pods results update. In: 1999 AGIFORS Reservations and Yield Management Study. MIT simulation study: PODS results update. In 1999 AGIFORS Reservations and Yield Management Study Group Symposium, April.
- Bertsimas, D., de Boer, S., 2005. Simulation-based booking limits for airline revenue management. *Operations Research* 53 (1), 90–106.
- Bijvank, M., L'Ecuyer, P., Marcotte, P., 2011. Rmsim: A java library for simulating revenue management systems. In: Proceedings of the Winter Simulation Conference. Winter Simulation Conference, pp. 2703–2714.
- Bront, J. J. M., Méndez-Díaz, I., Vulcano, G., 2009. A column generation algorithm for choice-based network revenue management. *Operations Research* 57 (3), 769–784.
- Carrier, E., 2003. Modeling airline passenger choice: Passenger preference for schedule in the passenger origin-destination simulator (pods). Ph.D. thesis, Massachusetts Institute of Technology.
- Carrier, E., Weatherford, L., 2015. Implementation of mnl choice models in the passenger origin-destination simulator (pods). *Journal of Revenue and Pricing Management* 14 (6), 400–407.
- Chaneton, J. M., Vulcano, G., 2011. Computing bid prices for revenue management under customer choice behavior. *Manufacturing and Service Operations Management* 13 (4), 452–470.
- Cleophas, C., Frank, M., Kliwer, N., 2009. Simulation-based key performance indicators for evaluating the quality of airline demand forecasting. *Journal of Revenue and Pricing Management* 8 (4), 330–342.
- Darot, J. F. J., 2001. Revenue management for airline alliances: Passenger origin-destination simulation analysis. Ph.D. thesis, Massachusetts Institute of Technology.

- Doreswamy, G. R., Kothari, A. S., Tirumalachetty, S., 2015. Simulating the flavors of revenue management for airlines. *Journal of Revenue and Pricing Management* 14 (6), 421–432.
- Eguchi, T., Belobaba, P. P., 2004. Modelling and simulation of impact of revenue management on japan's domestic market. *Journal of Revenue and Pricing Management* 3 (2), 119–142.
- Figueiredo, D. R., Liu, B., Guo, Y., Kurose, J., Towsley, D., 2006. On the efficiency of fluid simulation of networks. *Computer Networks* 50 (12), 1974–1994.
- Fiig, T., Härdling, R., Pölt, S., Hopperstad, C., 2014. Demand forecasting and measuring forecast accuracy in general fare structures. *Journal of Revenue and Pricing Management* 13 (6), 413–439.
- Frank, M., Friedemann, M., Mederer, M., Schroeder, A., 2006. Airline revenue management: A simulation of dynamic capacity management. *Journal of Revenue and Pricing Management* 5 (1), 62–71.
- Frank, M., Friedemann, M., Schrder, A., 2008. Principles for simulations in revenue management. *Journal of Revenue and Pricing Management* 7 (1), 7–16.
- Gilks, W. R., Richardson, S., Spiegelhalter, D., 1995. Markov chain Monte Carlo in practice. CRC press.
- Gorin, T., Belobaba, P. P., 2004. Special issue papers: Revenue management performance in a low-fare airline environment: Insights from the passenger origin-destination simulator. *Journal of Revenue and Pricing Management* 3 (3), 215–236.
- Gosavi, A., 2015. Simulation-based optimization.
- Gosavi, A., Ozkaya, E., Kahraman, A. F., dec 2005. Simulation optimization for revenue management of airlines with cancellations and overbooking. *OR Spectrum* 29 (1), 21–38.
- Kesidis, G., Singh, A., Cheung, D., Kwok, W., 1996. Feasibility of fluid event-driven simulation for atm networks. In: *Global Telecommunications Conference, 1996. Vol. 3. IEEE*, pp. 2013–2017.
- Kunnumkal, S., Topaloglu, H., 2011. A randomized linear program for the network revenue management problem with customer choice behavior. *Journal of Revenue and Pricing management* 10 (5), 455–470.
- Liu, Q., van Ryzin, G., 2008. On the choice-based linear programming model for network revenue management. *Manufacturing and Service Operations Management* 10 (2), 288–310.
- Meissner, J., Strauss, A., Talluri, K., apr 2013. An enhanced concave program relaxation for choice network revenue management. *Production and Operations Management* 22 (1), 71–87.
- Spall, J. C., 1998. An overview of the simultaneous perturbation method for efficient optimization. *Johns Hopkins apl technical digest* 19 (4), 482–492.
- Talluri, K., 2010. A randomized concave programming method for choice network revenue management. Working Papers (Departamento de Economía y Empresa, Universitat Pompeu Fabra).
- Talluri, K., 2014. New formulations for choice network revenue management. *INFORMS Journal on Computing* 26 (2), 401–413.
- Talluri, K., van Ryzin, G., 2004. *The Theory and Practice of Revenue Management*. Vol. 68. Springer Science and Business Media.
- van Ryzin, G., Vulcano, G., 2008a. Computing virtual nesting controls for network revenue management under customer choice behavior. *Manufacturing and Service Operations Management* 10 (3), 448–467.
- van Ryzin, G., Vulcano, G., 2008b. Simulation-based optimization of virtual nesting controls for network revenue management. *Operations Research* 56 (4), 865–880.
- van Ryzin, G., Vulcano, G., 2015. A market discovery algorithm to estimate a general class of nonparametric choice models. *Management Science* 61 (2), 281–300.
- van Ryzin, G., Vulcano, G., apr 2017. Technical note—an expectation-maximization method to estimate a rank-based choice model of demand. *Operations Research* 65 (2), 396–407.
- Weatherford, L. R., 2013. Improved revenues from various unconstraining methods in a passenger origin-destination simulator (pods) environment with semi-restricted fares. *Journal of Revenue and Pricing Management* 12 (1), 60–82.

Appendix

Optimal policy

LP	FAS				DAS _{reg}				DAS _{reg}				DAS _{0.5reg}						
	μ	CPU	CRS	μ	CPU	CRS	N	σ	μ	CPU	CRS	N	σ	μ	CPU	CRS	N	σ	
0.8	92600.8	0.124	22.0	85533.37	0.296	944.9	10.0	10209287.0	87450.81	2.294	962.2	126.0	6247656.2	87304.62	9.612	961.9	557.0	6895921.2	
1.0	102771.8	0.048	28.0	95450.90	0.172	1206.6	9.0	11827927.0	95786.65	1.638	1206.2	82.0	4884783.1	95503.05	7.880	1199.6	401.0	5031943.5	
1.2	110140.9	0.031	30.0	102935.03	0.110	1439.3	4.0	5376764.4	102537.98	1.653	1439.9	74.0	5062021.2	102723.51	8.139	1438.7	360.0	6149063.3	
Bus-line	0.8	1577010.7	2.902	76.0	1532008.23	13.603	14788.8	4.0	94893130.2	1532144.54	16.224	14669.8	5.0	66112925.5	1536161.05	102.131	14732.5	31.0	113502139.7
	1.0	170257.4	2.574	87.0	1726482.23	14.530	18477.3	4.0	144123012.0	1718801.66	24.695	18459.0	7.0	121457926.4	1717689.62	74.968	18348.2	21.0	98069964.2
	1.2	1921637.3	2.668	96.0	1857434.60	17.004	21991.3	4.0	237029724.3	1873311.53	16.792	22201.3	4.0	73871673.7	1872608.46	38.686	22072.3	9.0	48479596.2

Table A.1: Optimal policies for both instances.

Random policies for bus-line

LF	Close	FAS			DAS _{5%}			DAS _{1%}			DAS _{0.5%}								
		μ	CPU	CRS	μ	CPU	CRS	N	σ	μ	CPU	CRS	N	σ					
0.8	0	56537.3	0.102	101.3	56094.36	0.107	962.1	7.0	2921638.5	56018.76	2.854	960.1	188.1	3812418.5	56013.19	11.377	960.0	759.1	3858877.2
	10	56104.2	0.090	96.7	55703.74	0.115	962.3	7.2	3007110.8	55620.41	2.937	960.0	192.9	3850863.5	55619.71	11.861	960.0	782.0	3925145.6
	20	55555.9	0.074	91.2	55010.18	0.108	960.5	7.7	32329646.2	55044.11	2.866	959.4	204.2	4004469.6	55039.72	11.667	959.9	830.7	4078813.4
	30	53931.2	0.064	86.1	53458.60	0.110	962.0	8.4	3300121.4	53482.54	2.960	959.5	222.3	4110774.5	53492.86	11.995	959.8	903.6	4194560.8
	40	52847.9	0.055	79.8	52297.73	0.098	960.9	8.0	3118408.1	52380.90	2.701	959.6	219.4	38904023.5	52381.77	11.442	959.9	926.1	4117838.1
	50	50049.1	0.043	70.7	49637.00	0.099	961.0	8.9	3188065.2	49635.22	2.791	959.9	249.9	3981737.4	49634.92	12.502	960.0	1005.1	4018042.3
	60	46558.0	0.032	60.7	46202.32	0.101	959.4	9.4	2902915.1	46172.58	3.136	960.2	287.2	3957792.2	46173.27	12.682	959.9	1152.1	3970582.2
	70	41428.5	0.024	50.0	41103.90	0.118	959.6	12.4	3153471.2	41083.03	3.291	959.9	337.2	3664334.4	41067.55	13.091	960.0	1343.6	3655435.0
	80	34478.5	0.012	37.2	34066.80	0.143	961.1	16.5	2856063.1	34138.44	3.641	960.4	460.7	3937334.5	34114.18	14.088	960.0	1849.8	3407358.0
	90	22142.2	0.004	21.7	21950.56	0.193	959.4	31.3	2151097.9	21968.58	5.725	960.1	913.4	2610901.6	21957.67	23.164	960.0	3693.9	2635124.2
100	0.00	0.000	1.0	0.00	0.017	967.2	4.0	0.00	0.00	0.016	963.5	4.0	0.00	0.00	0.016	957.4	4.0	0.00	
1.0	0	62476.1	0.089	92.1	61726.12	0.105	1198.4	5.8	2754979.4	61843.55	2.333	1199.3	128.3	3168554.3	61836.65	9.426	1200.0	517.8	3217196.1
	10	61699.5	0.080	89.5	61105.22	0.106	1201.6	6.3	3028752.7	61124.78	2.449	1200.4	143.0	3453028.6	61108.19	9.541	1199.9	557.0	3374719.0
	20	61508.9	0.069	85.2	60817.15	0.093	1200.6	5.9	2663804.9	60709.68	2.348	1199.5	146.1	3490234.7	60821.31	9.474	1199.8	588.5	3532049.3
	30	60340.5	0.061	81.4	59611.68	0.091	1200.5	6.0	2785956.9	59707.01	2.432	1199.7	160.0	3684424.6	59703.52	9.711	1200.0	638.2	3687769.0
	40	58718.4	0.050	73.6	58071.25	0.101	1198.4	6.9	3011886.0	58087.20	2.523	1200.3	171.8	3719482.8	58056.05	10.123	1199.8	687.2	3735746.0
	50	55466.6	0.040	67.7	54441.24	0.065	1202.3	4.8	1576191.8	54946.28	2.801	1200.0	205.9	4036720.4	54933.19	11.448	1200.1	839.5	4016927.4
	60	53099.8	0.032	60.9	52783.24	0.094	1199.3	7.5	2564536.5	52648.35	2.176	1200.3	172.6	3074771.0	52687.85	10.644	1200.6	843.4	3783906.9
	70	45590.0	0.022	51.6	45054.71	0.163	1199.4	13.4	4023351.7	45168.85	3.945	1200.0	321.4	4199301.3	45153.91	15.659	1200.6	1275.1	4182286.1
	80	35287.2	0.010	34.8	35017.10	0.185	1198.8	18.3	3178027.4	35133.32	4.452	1201.0	435.0	3380377.7	35058.14	18.080	1200.1	1705.4	3413271.8
	90	24203.1	0.004	21.7	24041.96	0.289	1204.6	31.8	2887619.8	23988.94	7.199	1200.4	784.8	2875117.7	23977.68	28.598	1200.2	3114.5	2865961.3
100	0.00	0.000	1.0	0.00	0.025	1204.9	4.0	0.00	0.00	0.026	1207.0	4.0	0.00	0.00	0.026	1201.3	4.0	0.00	
1.2	0	66206.9	0.084	83.7	65604.91	0.101	1443.9	5.0	2125171.0	65780.12	2.216	1439.3	108.9	3039319.0	65841.30	8.206	1440.8	403.2	2841232.4
	10	65797.6	0.076	82.9	65872.43	0.121	1439.6	6.5	3572049.2	65312.83	2.206	1440.9	115.8	3179389.2	65216.81	8.452	1440.2	442.8	3046957.5
	20	65459.0	0.067	80.9	64778.36	0.106	1445.1	5.7	2751926.2	64708.86	1.895	1442.0	102.8	2798014.8	64791.50	8.503	1439.7	462.0	3148926.7
	30	64675.1	0.057	74.7	64406.83	0.095	1447.8	5.5	2523237.3	64009.93	2.273	1437.9	130.0	3430867.8	64074.90	9.108	1439.9	519.2	3454938.4
	40	63326.3	0.049	71.3	62378.14	0.099	1440.0	5.8	3002225.1	62721.75	2.436	1439.9	144.3	3679479.8	62734.53	10.161	1439.8	598.9	3830511.7
	50	60784.2	0.039	63.4	60334.00	0.122	1446.5	7.8	4304230.0	60188.26	2.132	1441.4	137.5	3229970.5	60151.86	8.903	1440.2	572.1	3376911.4
	60	57095.5	0.030	58.0	56376.70	0.096	1442.6	6.6	2721236.4	56601.59	2.918	1441.1	197.9	4113755.1	56656.19	11.742	1440.0	796.6	4417954.2
	70	48037.9	0.018	45.2	47758.24	0.142	1441.5	10.5	3799538.8	47740.32	3.333	1440.4	248.8	3707782.6	47792.62	13.236	1440.4	988.7	3616478.0
	80	41793.2	0.010	35.0	41548.49	0.146	1438.5	12.0	2934531.4	41525.58	4.054	1441.0	340.9	3611829.0	41505.62	16.075	1439.5	1347.9	3593111.0
	90	28048.0	0.003	20.1	27915.51	0.211	1435.6	21.3	2443094.4	27860.12	5.365	1440.0	527.4	2576515.6	27870.81	21.771	1440.3	2139.3	2625548.8
100	0.00	0.000	1.0	0.00	0.032	1441.6	4.0	0.00	0.00	0.030	1443.7	4.0	0.00	0.00	0.028	1434.4	4.0	0.00	

Random policies for airline

LF	Close	FAS				DAS _{5%}				DAS _{1%}				DAS _{0.5%}						
		μ	CPU	CRS	N	μ	CPU	CRS	N	σ	μ	CPU	CRS	N	σ	μ	CPU	CRS	N	σ
0.8	0	1036884.9	6.158	309.8	1035345.63	7.956	14723.9	4.0	95297381.6	1035097.75	18.037	14734.7	9.1	55766726.6	1035026.42	70.425	14733.8	35.5	59860729.8	
	10	1012809.5	5.528	300.1	1011513.38	7.920	14749.4	4.0	72285052.6	1011869.55	17.791	14746.3	9.1	52324969.9	1011821.45	79.087	14733.7	39.9	65366514.4	
	20	994411.4	5.036	287.4	994189.12	7.553	14725.0	4.0	70366851.1	993855.04	15.366	14747.0	8.1	45966566.6	994389.44	75.610	14742.3	40.1	62598393.6	
	30	977374.5	4.348	267.8	974643.35	7.296	14746.3	4.0	81224022.4	975654.57	18.644	14727.6	10.3	56191646.7	975952.92	76.686	14730.9	42.2	63815236.3	
	40	932862.6	3.759	253.8	931963.23	6.799	14733.2	4.0	70582000.9	931074.37	20.587	14738.1	12.3	63999581.8	931470.17	83.525	14732.9	50.1	69123216.5	
	50	888975.3	2.914	232.5	887217.33	6.127	14737.9	4.0	64117001.0	886392.97	17.140	14726.6	11.2	51859982.3	886970.68	76.340	14730.6	49.8	62434919.6	
	60	842022.9	2.304	212.1	841742.04	6.018	14742.0	4.0	64244840.8	840035.83	17.868	14735.8	11.9	50421832.9	840241.18	72.074	14728.0	48.1	53940928.5	
	70	767356.8	1.650	182.4	765198.97	5.598	14734.1	4.0	42989441.7	764863.24	19.458	14730.5	14.0	49074125.4	765978.79	87.671	14734.7	62.6	58943370.0	
	80	677436.6	0.998	152.0	677084.15	5.655	14719.7	4.0	52735082.5	675961.69	21.888	14735.4	15.5	42905672.2	676054.41	100.622	14730.2	70.7	52082792.6	
	100	479920.2	0.410	108.0	477958.09	5.936	14728.1	4.0	41420370.6	479085.50	47.909	14746.4	32.2	46516668.2	478665.06	204.576	14735.8	137.6	50230505.1	
1.0	0	1151504.4	5.584	285.0	1148955.20	8.739	18403.7	4.0	79943389.2	1149083.74	16.465	18413.1	7.6	55904820.3	1149584.53	69.860	18411.5	32.0	66925193.9	
	10	1134115.7	4.913	270.4	1132963.08	8.728	18413.2	4.0	83204396.3	1131246.66	16.501	18408.2	9.2	53782468.4	1132662.67	60.276	18417.1	28.1	56346434.6	
	20	1109327.5	4.381	253.1	1108683.29	8.053	18408.6	4.0	51590047.9	1108918.89	18.569	18398.4	9.8	65707270.8	1108365.85	59.192	18419.9	29.3	57339959.2	
	30	1081099.5	3.879	244.6	1079690.95	7.705	18444.1	4.0	66725733.3	1079376.48	12.354	18420.4	6.4	39441581.6	1081067.17	54.131	18434.1	28.0	51021333.5	
	40	1032990.5	3.215	222.5	1030358.74	7.253	18409.4	4.0	69253853.8	1031280.26	16.560	18404.3	9.1	55029788.5	1032217.19	71.920	18433.2	39.6	67667746.4	
	50	991126.5	2.668	217.0	991087.37	6.746	18402.5	4.0	74623096.6	987777.64	18.496	18405.3	11.0	62132748.0	989075.00	63.672	18421.1	37.6	58647366.4	
	60	926371.6	1.998	190.2	925409.40	6.716	18405.0	4.0	75171115.0	925650.33	20.050	18423.1	11.8	61010298.7	925020.86	73.316	18408.0	43.4	59392005.2	
	70	860785.5	1.532	173.2	857986.10	6.382	18404.8	4.0	53207193.8	859793.50	18.488	18438.9	11.7	52661987.5	858800.70	85.522	18409.7	53.9	62340181.9	
	80	732914.0	0.925	143.8	731382.05	6.450	18414.3	4.0	59585522.6	730700.18	22.015	18416.3	13.5	42488674.5	731824.68	91.062	18421.7	55.5	47792181.0	
	100	538481.3	0.361	96.7	536110.18	6.801	18422.6	4.0	46290822.8	535763.02	30.602	18412.4	18.0	31735035.1	537460.67	162.564	18420.0	95.9	44038911.1	
1.2	0	1238896.0	4.964	256.4	1236474.39	8.911	22120.9	4.0	120645725.6	1238156.51	13.820	22076.0	6.2	50926248.3	1236286.70	54.377	22101.6	24.4	58061849.3	
	10	1220504.6	4.516	252.7	1220107.06	8.851	22145.0	4.0	45560197.4	1219042.73	18.828	22088.8	8.5	75020619.0	1219110.12	52.301	22105.4	23.6	53667753.7	
	20	1190876.9	3.974	237.2	1188082.03	8.538	22091.8	4.0	522460091.4	1189150.21	11.093	22056.2	5.2	37155266.2	1188819.06	58.742	22092.1	27.5	60523862.6	
	30	1171343.4	3.554	228.0	1170013.16	8.158	22116.7	4.0	76646075.7	1168867.33	14.720	22060.3	7.2	51118588.7	1170273.27	65.820	22111.2	32.2	69501307.5	
	40	1104690.7	2.964	211.1	1102947.04	7.705	22101.9	4.0	62785463.0	1102604.45	16.188	22086.7	8.3	59356321.6	1103385.73	52.615	22112.7	27.4	52884161.6	
	50	1066002.2	2.297	193.5	1065033.98	7.218	22117.6	4.0	90991984.2	1064902.13	15.370	22107.2	8.5	56793062.7	1064111.04	46.662	22090.4	25.7	46625084.1	
	60	998682.2	1.781	176.7	995694.51	7.209	22083.8	4.0	63031155.5	996957.72	14.519	22094.6	8.0	46524614.0	997061.59	76.815	22098.0	42.3	66928650.8	
	70	931859.5	1.378	160.3	930049.79	6.727	22129.7	4.0	55157355.8	930090.75	12.410	22090.6	7.3	32348401.4	930538.88	72.315	22095.7	42.7	58440979.8	
	80	809614.2	0.865	133.7	807634.25	7.069	22064.2	4.0	67003002.3	808300.85	19.695	22093.2	11.1	45431471.6	807825.92	92.551	22096.1	52.0	54188822.0	
	100	595283.1	0.356	94.4	596550.86	7.610	22123.6	4.0	43202173.9	593057.34	36.838	22100.0	19.3	41433174.1	593760.43	164.953	22110.7	89.3	50322242.5	
100	0	0.000	1.0	0.000	5.917	22087.3	4.0	0.0	0.000	5.914	22110.2	4.0	0.0	0.000	5.934	22130.4	4.0	0.0		

Table A.3: Airline with 50 random policies per PC scenario.